University of Cergy-Pontoise

Doctoral school "Science and Engineering"

ETIS - CNRS UMR 8051 / ENSEA / University of Cergy-Pontoise

THESIS

Submitted in fulfillment of the requirements for the degree of

Doctor of Philosophy in Robotics and Cognitive Science

INTERACTIONS BETWEEN COGNITIVE AND EMOTIONAL PROCESSES: A STUDY IN NEUROMIMETIC MOBILE AND SOCIAL ROBOTICS

by

Marwen Belkaid

December 2016

Doctoral Committee:

Pr. Yann COELLO	University of Lille	Committee chair
Dr. Frédéric ALEXANDRE	Inria Bordeaux	Reviewer
Pr. Jeffrey L. KRICHMAR	University of California	Reviewer
Dr. Lola CAÑAMERO	University of Hertfordshire	Examiner
Pr. Philippe GAUSSIER	University of Cergy-Pontoise	Thesis supervisor
Dr. Nicolas CUPERLIER	University of Cergy-Pontoise	Thesis co-advisor

Acknowledgments – Remerciements

First, I would like to express my sincere gratitude to Frédéric Alexandre, Jeffery Krichmar, Lola Cañamero and Yann Coello. I am thankful for their meticulous reading and their pertinent comments and advices. It was a great pleasure and an honor for me to defend my thesis in front of this doctoral committee.

The rest of the acknowledgments will be in French because they are directed at colleagues, friends and family members who are French speakers.

Je suis infiniment reconnaissant envers mes encadrants de thèse Philippe Gaussier et Nicolas Cuperlier pour ces trois années de collaboration, pour leur confiance et pour les libertés qu'ils m'ont accordées dans le cadre de mes travaux de recherche. Ils ont également toujours été disponibles et n'ont jamais hésité à mettre la main à la pâte avec le petit thésard que j'étais quand il le fallait. Chacun à sa manière, Philippe et Nicolas ont su me pousser à me donner à fond, et ce jusqu'à la dernière minute. Cela a été un plaisir de travailler avec eux.

Je voudrais ensuite exprimer ma gratitude envers tous les membres du laboratoire ETIS, et en particulier Mathias Quoy, directeur du laboratoire, pour m'avoir accueilli et donné les moyens de mener à bien mon projet de recherche. Je salue aussi particulièrement l'écoute dont Mathias a toujours fait preuve à l'égard des doctorants. Merci également à Astrid et Sokhena, mes drôles de dames, pour leur bonne humeur et leur bienveillance.

Je tiens à remercier les membres de l'équipe Neurocybernétique : Alex, Ghilès, Pierre, Sofiane, Arnaud, Frédéric, Jérôme, Abdelhak, Antoine, Adrien, Artem, Ali, Souheïl, Pierre, Khursheed, Raphaël, Caroline, Ganna, Nils, Alia, Eva et Frédérique. L'ambiance qui règne au sein de cette équipe a été l'une de mes sources de motivation tout au long de mon doctorat. Permanents et doctorants, sans distinction, ont été là pour m'aider, me conseiller, me rassurer et me faire rire. Une pensée particulière pour Adrien, mon maître Jedi, Ali et Souheïl, mes sages du village, Raphaël, mon confident du train, Caroline, mon associée dans le dressage de robots mobiles, Ghilès, mon compagnon de chasse aux ramens et aux gyozas, et Jérôme, mon espion/ambassadeur à l'extérieur. Je remercie aussi Filipe Gama, Bastien Lepesant, Kevin Denamganaï et Jason Henault, que j'ai le plaisir d'encadrer pendant leurs passages plus ou moins au sein de l'équipe.

Pour finir, mes remerciements s'adressent évidemment à ma famille, Kamel, Jouda, Memia et Ahmed, ainsi qu'à tous mes amis. Sans la présence et le soutien de tous mes proches, il m'aurait été impossible de réussir.

A tous ceux qui font de moi une meilleure personne.
A mon étoile.

Abstract

The purpose of my thesis is to study interactions between cognitive and emotional processes through the lens of neuromimetic robotics. The proposed models are implemented on artificial neural networks and embodied in robotic platforms – forming situated systems. In general, the interest is twofold: 1) taking inspiration from biological solutions to design systems that better interact with their physical and social environments, and 2) providing computational models as a means to better understand biological cognition and emotion.

The first part of the dissertation addresses spatial navigation as a framework to study biological and artificial cognition. In Chapter 1, I present a brief overview of the literature on biologically inspired navigation. Then, two issues are more specifically tackled. In Chapter 2, visual place recognition is addressed in the case of outdoor navigation. In that matter, I propose a model based on the notions of *visual context* and *global precedence* combining local and holistic visual information. Then, in Chapter 3, I consider the *interactive learning* of navigation tasks through non-verbal human–robot communication based on low-level visuomotor signals.

The second part of the dissertation addresses the central question of emotion–cognition interactions. In Chapter 4, I give an overview of the research on emotion as a cross-disciplinary enterprise, including psychological theories, neuroscientific findings and computational models. In Chapter 5, I propose a *conceptual model* of emotion–cognition interactions. Then, various instantiations of this model are presented. In Chapter 6, I model the perception of the *peripersonal space* when modulated by emotionally valenced sensory and physiological signals. Last, in Chapter 7, I introduce the concept of *Emotional Metacontrol* as an example of emotion–cognition interaction. It consists in using emotional signals elicited by self-assessment to modulate computational processes – such as attention and action selection – for the purpose of behavior regulation.

In this thesis, a key idea is that, in autonomous systems, emotion and cognition cannot be separated. Indeed, it is becoming well admitted that emotion is closely related to cognition, in particular through the modulation of various computational processes taking place in the brain. This raises the question of whether this modulation occurs at the level of sensory processing or at the level of action selection. In this thesis, I will advocate the idea that artificial emotion must be integrated in robotic architectures through bidirectional influences with sensory, attentional, decisional and motor processes. This work attempts to highlight how this approach to internal emotional processes can foster efficient interactions with the physical and social environment.

Résumé

L'objectif de ma thèse est d'étudier les interactions entre processus cognitifs et émotionnels à travers le prisme de la robotique neuromimétique. Les modèles proposés sont implémentés sur des réseaux de neurones artificiels et encorporés dans des plateformes robotiques – formant des systèmes situés. D'une manière générale, l'intérêt est double : 1) s'inspirer des solutions biologiques pour concevoir des systèmes qui interagissent mieux avec leurs environnements physiques et sociaux, et 2) mettre en place des modèles computationels comme moyen de mieux comprendre la cognition et les émotions biologiques.

La première partie du manuscrit aborde la navigation comme cadre permettant d'étudier la cognition biologique et artificielle. Dans le Chapitre 1, je commence par donner un aperçu de la littérature en navigation bio-inspirée. Ensuite, deux problématiques sont traitées. Dans le Chapitre 2, la reconnaissance visuelle de lieux en environnement extérieur est abordée. Dans ce cadre, je propose un modèle basé sur les notions de *contextes visuels* et de *précédence globale* et qui combine des informations locales et holistiques. Puis, dans le Chapitre 3, je considère l'*apprentissage interactif* de tâches de navigation à travers une communication homme–robot non-verbale basée sur des signaux visuomoteurs de bas niveau.

La deuxième partie du manuscrit s'attaque à la question centrale des interactions entre emotion et cognition. Le Chapitre 4 introduit la recherche sur les émotions comme une entreprise inter-disciplinaire, incluant des théories en psychologie, des découvertes en neurosciences et des modèles computationnels. Dans le Chapitre 5, je propose un *modèle conceptuel* des interactions emotion–cognition dont différentes instantiations sont ensuite présentées. Plus précisément, dans le Chapitre 6, je modélise la perception de l'*espace péripersonnel* quand elle est modulée par des signaux sensoriels et physiologiques ayant une valence émotionnelle. Enfin, dans le Chapitre 7, j'introduis le concept de *métacontrôle émotionnel* comme un exemple d'interaction emotion–cognition. Cela consiste à utiliser des signaux émotionnels élicités par une auto-évaluation pour moduler des processus computationnels – tels que l'attention ou la sélection de l'action – dans le but de réguler le comportement.

Une idée clé dans cette thèse est que, dans les systèmes autonomes, emotion et cognition ne peuvent pas être séparées. En effet, il est maintenant bien admis que les émotions sont très liées à la cognition, en particulier à travers la modulation de différents processus computationnels dans le cerveau. Cela pose la question de savoir si ces modulations se produisent au niveau du traitement sensoriel ou celui de la sélection de l'action. Dans cette thèse, je préconise l'intégration des émotions artificielles dans les architectures robotiques via des influences bidirectionnelles avec les processus sensoriels, attentionnels, moteurs and décisionnels. Ce travail tente de mettre en avant la manière dont cette approche des aspects internes des processus émotionnels peut

favoriser des interactions efficaces avec l'environnement physique et social du système.

Résumé étendu

L'objectif de ma thèse est d'étudier les interactions entre processus cognitifs et émotionnels à travers le prisme de la robotique neurmimétique. Les modèles proposés sont implémentés sur des réseaux de neurones artificiels et encorporés[1] dans des plateformes robotiques – formant des systèmes situés. D'une manière générale, l'intérêt est double : 1) s'inspirer des solutions biologiques pour concevoir des systèmes qui interagissent mieux avec leurs environnements physiques et sociaux, et 2) mettre en place des modèles computationels comme moyen de comprendre la cognition et les émotions biologiques.

La première partie du manuscrit aborde la navigation comme cadre permettant d'étudier la cognition biologique et artificielle. A ce sujet, le Chapitre 1 commence par un récapitulatif d'importants résultats de recherche en neuroscience qui impliquent la régiond hippocampique dans un ensemble de fonctions liées à la navigation. Dans le cadre de ma thèse, un intérêt particulier est porté sur les cellules de lieu : un type particulier de neurones dont les activités sont corrélées avec la localisation de l'agent dans un environnement familier. Un modèle biologiquement plausible de reconnaissance de lieux basé sur ce type ce cellules à été proposé dans de précédents travaux effectués au sein du laboratoire.

Dans le Chapitre 2, je me suis intéressé de plus près à la reconnaissance visuelle de lieux en extérieur. Ces environnements présentent un défi car ils sont beaucoup moins structurés et controlés qu'en intérieur. En particulier, un problème particulier auquel fait face le modèle existant de cellules de lieu est celui de la robustesse de la reconnaissance sur de longues distances. L'approche que j'ai suivi a été d'examiner le type d'entrées visuelles utilisées pour la localisation. Dans la littérature, les méthodes utilisant des informations locales ont généralement de meilleures performances. Cependant, l'information visuelle global est également informative et peut nécessiter un coût computationnel plus faible. Pour confirmer cela, j'ai présenté une expérience comparant sept descripteurs holistiques au descripteur local utilisé dans LPMP. Comme prévu, ce dernier a globalement montré une meilleure performance. Toutefois, les résultats suggèrent aussi que les caractéristiques globales contiennent assez d'information utile pour la navigation. En accord avec ceci, il existe des travaux qui mettent en avant la capacité de l'Homme à capturer l'essentiel d'une scène visuelle à partir de son aspect global et d'informations basse-résolution. De plus, la notion de précédence globale indique que ces contextes visuels facilitent une reconnaissance fine à partir de détails visuels plus locaux. Par conséquent, j'ai proposé un modèle de reconnaissance de lieu basées sur le contexte. Cette solution augmente les performance du modèle existante. Elle montrent aussi des résultats prometteurs en terme de navigation

---

[1]Terme proposé par Jean-Paul Laumond pour traduire le mot *embodied* dans le concept de l'intelligence incarnée dans un corps [Pfeifer and Pitti, 2012].

dans de grands environnements et sur de longues distances. Plus précisément, l'utilisant d'une reconnaissance holistique rapide pour limiter l'espace de recherche lié à la vision locale a non seulement permis daméliorer la capacité de localisation mais aussi de maintenir un coût computationnel quasi-constant.

Il a aussi été montré qu'une implémentation du modèle de cellules de lieux sur des robots (le modèle LPMP) pouvait être utilisée pour leur apprendre à effectuer des trajectoires de navigation basées sur un couplage sensorimoteur. Dans le Chapitre 3, je me suis intéressé à l'apprentissage interactif de telles tâches via une communication homme–robot non-verbale qui repose sur des signaux visuomoteurs de bas niveau. En particulier, j'ai travaillé sur un modèle qui vise à prendre en compte la sensibilité aux interactions rythmiques dont fait preuve l'humain. Pour cela, la détection de synchronie est utilisée comme moyen de sélectionner un partenaire. La reconnaissance et le suivi visuels des partenaires humains ont été utilisés pour permettre au robot de ré-engager et de maintenir les interactions lors des phases d'apprentissage. Les résultats présentés montrent que cette architecture est capable d'apprendre et d'exécuter une tâche de navigation sensorimotrice.

La deuxième partie du manuscrit s'attaque à la question centrale des interactions entre emotion et cognition. Le Chapitre 4 introduit la recherche sur les émotions comme une entreprise inter-disciplinaire. Une vue d'ensemble des principales écoles de pensée ainsi que quelques exemples de théories influentes en psychologie sont donnés. Ensuite sont rapportés quelques découvertes neuroscientifiques liées aux régions cérébrales et aux signaux neurochimiques impliqués dans les émotions. Puis, quelques exemple de modèles computationnels implémentés sur robots ou sur des agents conversationnels virtuels sont présentés. Pour un roboticien, une vision functionnelle des émotion les rend intéressantes à modéliser pour deux raisons : améliorer les interactions robot–robot et humain–robot and accroître l'autonomie et les capacités d'adaptation des robots. D'un autre côté, les modèles computationnels permettre aux chercheurs de tester des hypothèses. Aller de concepts théoriques à des implémentations concrètes révèle aussi les ambiguïtés et pousse vers des définitions plus explicites.

Comme indiqué dans le Chapitre 5, différentes recherches convergent sur le fait que cognition et emotion ne peuvent pas être facilement séparées. Il s'avère en fait qu'elles s'influencent continuellement. Des processus liés à la perception, l'attention, la mémoire et le raisonnement, qui sont associés à la cognition, influencent les expériences émotionnelles. En retour, ils sont aussi modulés par les signaux émotionnels. Biologiquement, ces interactions semblent être représentés par une influence bi-directionnelle entre les activités neuronales et les modulations neurochimiques. Computationnellement, la théorie des systèmes dynamique paraît être un bon cadre pour étudier un tel couplage. Sur un plus haut niveau d'abstraction, j'ai proposé un modèle conceptuel d'interaction emotion–cognition. Ce modèle offre une vision complémentaire à celle d'une description sous forme de réseaux de neurones en représentant des connexions entre processus – de manière similaire à un diagramme d'activité (*workflow*) par exemple. D'abord, cela permet de facilement faire la correspondance avec des descriptions qui sont elles aussi généralement orientées processus. De plus, le modèle fournit ainsi une vue générique du concept indépendamment de l'implémentation et de l'application. Pour résumer, le modèle illustre comment, de manière continue, le système:

(a) évalue des événement se produisant dans le corps ou dans l'environnement extérieur, avec un intérêt particulier pour des stimuli émotionnellement significatifs, et

(b) traite des signaux émotionnellement modulés et les réintègre dans le flux de traitement

d'information dans le but d'effectuer des traitements et des évaluations de plus haut niveau.

Dans le Chapitre 6, je présente la première instantiation de ce modèle abstrait. Je m'intéresse ainsi à l'effet des émotions sur la perception de l'espace péripersonnel; l'espace où l'on peut interagir confortablement et sans danger avec l'environnement physique et social. Pour les humains et les animaux, la capacité à construire une telle représentation est fondamentale, même pour des comportements basiques d'approche et d'évitement. Par ailleurs, il a été montré que les émotions ont une forte influence sur la perception l'espace environnant. Dans le modèle que je propose, des signaux sensoriels et physiologiques ayant une valence émotionnelle modulent le flux de traitement de l'information au niveaux sensoriel et moteur. Je présente différentes expériences montrant que cette modulation émotionnelle l'espace péripersonnel peut servir de paradigm pour un contrôle réactif du comportement du robot. Des résultats supplémentaires suggères que le modèle peut rendre compte de l'impact de la désirabilité des objets dans la perception de leur atteignabilité. En outre, dans un contexte social, les robots exhibent un comportement craintif ou agressif selon les interactions internes entre signaux émotionnels. Ceci illustre comment l'aspect interne des processus émotionnels peut avoir un impact sur des aspects externes, plus liés aux interactions physiques et sociales.

Un autre exemple d'interaction emotion–cognition est donné dans le Chapitre 7. En effet, j'introduis ici le concept de Metacontrôle Emotionnel comme étant l'utilisation de signaux émotionnels élicités par une auto-évaluation pour moduler différents processus computationnel dans le but de réguler son comportement. L'auto-évaluation est la capacité à évaluer ses propres aptitudes et performances. Dans le modèle proposé, elle est effectuée à la volée, en se basant sur l'évaluation de contingences sensorimotrices et la détection de situations nouvelles. Dans la littérature, frustration et ennui caractérisent une incompatibilité entre les compétences et les exigences de la tâche. Par conséquent, dans mon modèle, la frustration est élicitée dans le cas où le robot régresse ou stagne avec une forte nouveauté. A l'opposé, l'ennui est généré par l'absence de nouveauté et de progrès dans la tâche. Dans mon modèle, ces *affects* duaux déclenchent tous les deux des méchanismes de modulation de haut-niveau qui régulent le comportement du robot. Dans la première expérience du chapitre, la frustration est utilisée pour surveiller et pondérer les sous-comportements indépendants d'un robot mobile dans une tâche de navigation. D'abord, je montre que le modèle fournit un moyen efficace de communiquer un échec et le besoin d'une intervention à un partenaire/superviseur humain. De plus, je démontre qu'il permet aussi la sélection de stratégie dans une tâche qui alterne suivi de route et l'exécution d'une trajectoire sensorimotrice apprise. Par la suite, la seconde expérience s'intéresse à la reconnaissance d'objets dans une tâche dont le but est de retrouver des objets cibles dans la scène visuelle. Cette fois, la frustration et l'ennui sont utilisés pour biaiser l'attention. Ils permettent de sortir de situations de bloquage qui sont due à un échec ou à une réussite persistante. Mes résultats montrent que, par rapport à une architecture en boucle ouverte, le métacontrôle émotionnel augmente les performances du système et favorise un comportement exploratoire dans ce types de situations.

Contents

Personal publications

International journals

Belkaid, M., Cuperlier, N., and Gaussier, P. (2016c). Emotional Metacontrol: top-down emotional modulation of sensorimotor processes during active object perception. *Submitted*, pages 1–13.

International peer-reviewed conferences

Belkaid, M., Cuperlier, N., and Gaussier, P. (2016a). Combining local and global visual information in context-based neurorobotic navigation. In *Proceedings of the IEEE International Joint Conference on Neural Networks*, pages 16622:1–16622:8.

Belkaid, M., Lesueur-Grand, C., Mostafaoui, G., Cuperlier, N., and Gaussier, P. (2016c). Learning sensorimotor navigation using synchrony-based partner selection. In *Proceedings of the ACM International Conference Artificial Intelligence and Robotics*,19:1–19:5.

Belkaid, M., Cuperlier, N., and Gaussier, P. (2015b). Emotional modulation of peripersonal space as a way to represent reachable and comfort areas. In *Proceedings of the IEEE International Conference on Intelligent Robots and Systems*, pages 353–359.

Belkaid, M., Cuperlier, N., and Gaussier, P. (2015c). Emotional modulation of peripersonal space impacts the way robots interact. In *Proceedings of the European Conference on Artificial Life*, pages 431–437.

Belkaid, M., Cuperlier, N., and Gaussier, P. (2015a). Emergence of low-level aggressive responses in a reactive goal-directed robotic system. In *Meeting of the International Society of Research on Emotion*.

Belkaid, M., Lesueur-Grand, C., Mostafaoui, G., Cuperlier, N., and Gaussier, P. (2014). When and from who to learn : Synchrony as a way to learn sensorimotor associations. In *Symposium on Biology of Decision Making*.

Jauffret, A., <u>Belkaid, M.</u>, Cuperlier, N., Gaussier, P., and Tarroux, P. (2013b). Frustration: a generic mechanism to improve autonomy in robotics. In *Proceedings of the IEEE Joint International Conference on Development and Learning and Epigenetic Robotics*, pages 1–7.

Jauffret, A., <u>Belkaid, M.</u>, Cuperlier, N., Gaussier, P., and Tarroux, P. (2013d). Frustration: a generic mechanism to improve autonomy in robotics. In *Symposium on Biology of Decision Making*.

French peer-reviewed conferences

<u>Belkaid, M.</u>, Cuperlier, N., and Gaussier, P. (2016b). De la détection de nouveauté au métacontrôle en passant par l'ennui et la frustration. In *Actes du Workshop Affect Compagnon Artificiel et Interaction*.

Introduction

"I propose to consider the question, 'Can machines think?'"

(Turing [1950])

Robots are fascinating. In addition to the technological advances they represent, they serve as a great means to question our knowledge of ourselves (humans) through a process of *understanding by design*. But, far from the *Promethean dream* – the dream of creating a new form of life – robot are progressively leaving the realm of science fiction. Indeed, as scientific and technological progress is made, more and more intelligent machines are introduced in our factories, shops, hospitals, museums and homes. Much like personal computers, Internet and mobile phones, robots are very likely to re-shape our daily life. One must admit that these are very exciting times to be a roboticist.

The word "robot" was first used by Karel Čapek[2] in his play *R.U.R.* (Rossum's Universal Robots) published in 1920. The Slavic word "robota" meaning labor, the play described artificial people made by humans and intended to work for them. Then, the word "robotics" was coined[3] by Isaac Asimov in the early 1940s is his short stories *Liar!* and *Runaround* as the science and technology of robots.

As a scientific field, robotics encompasses aspects of mechanical, electrical and computer science engineering. The latter part constitutes the intersection with another scientific community, that of *artificial intelligence* (AI). This fields emerged when the philosophical and mathematical foundations of formal reasoning could be tested on programmable machines. With the invention of digital computers, it was finally possible to consider the idea of creating an artificial brain. The 1940s were a pivotal period as converging seminal works in that matter appeared. McCulloch and Pitts [1943] proposed the first computational model of neuron. Wiener [1948] introduced cybernetics as the study of control and communication in animals and machines. Shannon [1948] gave rise to the information theory laying the foundations of modern telecommunication. Turing [1950] formulated the Turing test, designed to assess whether a machine exhibits intelligent behavior. A few years later, the field of artificial intelligence was founded as an academic discipline in 1956, at a conference gathering pioneers like Marvin Minsky, John McCarthy, Claude Shannon, Allen Newell and Herbert Simon. With the objective of allowing machines to fulfill complex tasks only accessible to humans so far, the overlap with robotics is obvious. Contemporaneously, Walter [1950] presented his *Machina speculatrix*, better known as the "tortoise" robot. By creating animal-like machines, the purpose of Walter's work was

---

[2]Karel Čapek asserts that his brother Josef Čapek is the actual originator of the word.
[3]Isaac Asimov thought the word already existed when he first used it.

to demonstrate that complex behaviors, like those enabled by the brain, come from the rich connectivity between small number of neurons.

In a broader view, robotics and artificial intelligence are part of the mutli-, cross-disciplinary field of *cognitive science*. Cognitive scientists – also including philosophers, psychologists, sociologists and neuroscientists among others – study the mental processes that are involved in the intelligence and behavior of humans, animals and machines.

Different approaches to cognition

When defining intelligence, three schools of thoughts can be distinguished. The *cognitivist* approach is based on the idea that thoughts are discrete representations (symbols) we manipulate via a set of rules. This belief suffers from various weaknesses. The most famous is the *symbol grounding problem* [Harnad, 1990]. Indeed, being able to manipulate symbols does not mean that one understands their meanings and their relation with the real world. A good illustration of this limit is given by a thought experiment proposed by Searle [1980]. In his argument, he shows that he could easily be thought to understand Chinese without any knowledge of the semantics of the language, provided he has access to a set of rules allowing him to correlate a set of Chinese characters (question) with another (answer). Symbolic artificial intelligence, which was the predominant approach in the early days of the field, has strong links with cognitivism. Therefore, classical AI – or *good old-fashioned AI* – is subject to the issue highlighted above. In more concrete terms, designing an "intelligent" behavior requires human expertise to provide the set of rules that govern the application domain. Additionally, it seems very unlikely that one can define all the potential outcomes of any possible situation in the real world. Therefore, such expert systems are generally limited to very specific use cases; and any new application needs a new knowledge base. So, although this approach can succeed in simulating intelligence, it fails in providing a framework for the study of cognition.

In contrast, the *connectionist* approach describes cognition as emerging from networks of interconnected units. Activities inside these networks depend on internal constraints determining how connections are changed and how nodes are activated by their inputs. The global coherence comes from these local properties. In this view, mental processes are dynamic patterns of activity rather than symbol manipulation. The main computational framework of this movement is provided by artificial neural networks (ANN) – networks of elementary calculating units (neurons) interconnected by weighted links (synapses). But since the first formal model of neuron [McCulloch and Pitts, 1943], this community has had ups and downs. The analogy with biological learning in pioneering works like those of Hebb [1949] and Rosenblatt [1958] was very promising. Then, the "dark ages" of ANN came when the limitation of existing models to linearly separable problems[4] was mathematically proven [Minsky and Papert, 1969]. However, the paradigm has regained popularity since the 1980s thanks to a large variety of models overcoming this limitation [Kohonen, 1982; Hopfield, 1982; Rumelhart et al., 1985; Rumelhart and McClelland, 1986; Carpenter and Grossberg, 1988]. The connectionist approach avoids the symbol grounding problem by allowing the system to self-organize and construct categories (classes) by

---

[4]To simplify, consider elements belonging to two classes are represented in a 2D space. If the two classes can be separated by a single line, ensuring that all elements of a class are on the same side of the line, then it is a linearly separable problem. This type of problem can be solved by a single-layer perceptron [Rosenblatt, 1958]. However, non-linearly separable problems require more layers of neurons. See Appendix A for an illustration

itself. But there still can be a few limitations related to this framework. First, if learning has to be supervised, then again, we need an expert showing the correct output we expect for a given input. Besides, the issue of application-specific and restricted knowledge can also be encountered with this paradigm. For instance, let us consider a system that has to learn what a "chair" is. This problem can be addressed by showing various images of chairs and making the network learn the corresponding category. Yet, designers constantly come up with new ideas, creating objects that are aesthetically different but functionally similar to the chairs we knew before. A system that does not have the capacity to interact with the objects it encounters – and sit on a chair every now and then – will never capture this functional aspect.

This leads us to the *enactivist* approach. Introduced by Maturana and Varela [1980]; Varela et al. [1992], enactivism emphasizes that intelligence needs a body. In this view, cognition emerges from interactions with the environment. Thus, a program alone cannot understand the world if it does not move in it, act on it and alter it. This is in line with the sensorimotor theories of perception in psychology [Piaget, 1951; Gibson, 1979; O'Regan and Noë, 2001]. In this framework, perception is an active process that allows the detection of sensorimotor invariants from a flow of information (sensory input). According to the theory of affordances proposed by Gibson [1979], humans extract meaning from "action possibilities" around them. By definition, affordances depend on the physical capabilities of the agent. In the example I presented earlier, the meaning that we implicitly expect the system to acquire about chairs is the one defined by human morphology, not that of a platypus.

The enactivist approach intends to offer an alternative to the connectionist and cognitivist views of intelligence. Because it suggests that cognition can only be embodied, it implies that artificial intelligence requires robots. However, it does not presuppose a particular computational framework when it comes to the question of how to make machines think. My work lies in the intersection between connectionism and enactivism. Indeed, I will be interested in the phenomena that emerge from interactions between internal processes; as well as from physical and social interactions with the environment.

Robotic approaches and paradigms

Neurorobotics consists in embodying neural systems in robotic platforms. There is also a strong link with neuroscience through the design of brain-like architectures. Thus, these neural systems range from artificial neural networks to biological (cultured) neural tissues wired to electronic devices. Some artificial models attempt to precisely mimic the behavior of biological neurons. For example, spiking neural networks model cells that only fire when the potential – simulated electrical charge accumulated from synaptic inputs – reach a certain threshold [Hodgkin and Huxley, 1952; Maass, 1997]. These models are particularly appropriate to study the temporal aspect of the interactions between neurons. Otherwise, a simpler approach is to consider *mean firing rate* models. As the name suggests, the activity of these artificial neurons illustrate the average firing in a certain period of time rather than actual spikes. In my experiments, I will implement this type of neurons. The inputs of the first layers are provided by the sensors of the robotic platforms. The activities inside the network are updated at each time step, resulting in motor orders controlling the platform actuators.

Bio-inspired robotics encompasses various technologies that synthesize principles originally observed and studied in biological organisms. The idea is to take advantage from natural solutions developed ecologically throughout evolution [Braitenberg, 1984; Brooks, 1986; Gaussier and Zrehen, 1995; Pfeifer and Scheier, 1995]. Defined as such, the domain covers a large spectrum of research and has notably expanded in the last decade [Pfeifer et al., 2007; Lepora et al., 2013]. A significant portion regards the design of hardware solutions, including biosensors, bioactuators and biomaterials. For examples, the relatively recent sub-field of soft robotics relies on non-rigid (compliant) materials to model biomorphological features in order to facilitate robots physical interactions with the environment. More related to artificial intelligence and thus more relevant to my thesis is the sub-field of *biologically inspired cognitive robotics*; i.e. embodied artificial systems that carry out tasks related to attention, learning, decision-making and so on, in a way that mimics aspect of animal and human behaviors [Krichmar, 2012]. Besides, *Animats*, which I will elaborate on below, intend to model animal behaviors and thereby also belong to this family by definition.

Behavior-based robotics is an approach that arose in the 1980s as an alternative to the deliberative paradigm commonly used in what is today called *classical* robotics. Rather than top-down centralized decisions following a sense-plan-act sequence, these architectures rely on sensorimotor coupling and the principle of embodied cognition [Gaussier and Zrehen, 1995; Pfeifer and Scheier, 1995]. This research largely took inspiration from ethology. Converging efforts aiming to reproduce biological behaviors in artificial systems gave rise to the *animat* approach. One of the pioneers of this paradigm is Brooks [1986, 1991]. He proposed the subsumption architecture, organizing the behavior into a hierarchy of layers (sub-behaviors), all running in parallel. The key point lies in the decomposition of the overall behavior instead of decomposing the problem. Another seminal work in that of Braitenberg [1984] showing how rich and seemingly complex behaviors emerge from the interaction of very simple systems with their environment. In this thesis, most of the tasks do not require planning. The sense-act coupling allows fast control and reactivity by only figuring out what to do in the next step. Moreover, the architectures are either implemented on physical robots that interact with the real world or run on simulations that provide realistic physics or the type of feedback that is expected from a real environment in the context of the specific task.

Developmental/Epigenetic robotics aims at understanding the mechanisms that allow lifelong learning and acquisition of increasingly complex cognitive structures through interactions with the environment. In his plenary talk at the *World Congress on Computational Intelligence* in 2016, Leslie Valiant suggested that the world could be viewed as a collection of organisms that learn (and evolve) to adapt to their environment. Whereas evolutionary robotics studies the emergence of new architectures across generations of a population of agents, epigenetic robotics is interested in developmental processes at the scale of individuals. More specifically, the purpose is to model the emergence of higher order cognitive functions through autonomous learning that benefit from physical and social interactions [Asada et al., 2009]. In the context of my thesis, the tasks will not be pre-programmed but learned through actual or simulated interactions. The approach is bottom-up and constructivist: the system is provided with minimalistic prewired reflexes and learning capabilities (inspired by studies on biological cognition) in order to observe potential emerging properties.

Outline of the dissertation

The purpose of my thesis is to study the interactions between cognitive and emotional processes through the prism of robotics. My work will be inspired by converging findings showing that, in biological systems, emotion cannot be separated from cognition. Indeed, it is becoming well admitted that emotion is closely related to cognition, in particular through the modulation of various computational processes. This raises the question of whether this modulation occurs at the level of sensory processing or at the level of action selection. In this thesis, I will advocate the idea that artificial emotion must be integrated in robotic architectures through bidirectional influences with sensory, attentional, decisional and motor processes. Such an approach to internal emotional processes is intended to foster efficient physical and social interactions between the robot and its environment.

The programs that govern the robot behavior will be implemented on artificial neural networks (*neurorobotics*). They will be designed by taking inspiration from biological systems (*bio-inspired robotics*) – i.e. what we know about humans and animals from studies in other cognitive science disciplines. Sensorimotor architectures will be proposed, based on the idea that complex, yet stable, behaviors can emerge from the sensation–action loop in an embodied and situated agent (*behavior-based robotics*). Given the ability to learn and adapt, these robotic system can benefit from a rich physical and social interaction with their environment (*developmental/epigenetic robotics*).

The first part of the dissertation addresses spatial navigation as a framework to study biological and artificial cognition. In Chapter 1, I present a brief overview of the literature on biologically inspired navigation. Then, two issues are more specifically tackled. In Chapter 2, visual place recognition is addressed in the case of outdoor navigation. In that matter, I propose a model based on the notions of *visual context* and *global precedence* combining local and holistic visual information. Then, in Chapter 3, I consider the *interactive learning* of navigation tasks through non-verbal human–robot communication based on low-level visuomotor signals.

The second part of the dissertation addresses the central question of emotion–cognition interactions. In Chapter 4, I give an overview of the research on emotion as a cross-disciplinary enterprise, including psychological theories, neuroscientific findings and computational models. In Chapter 5, I propose a *conceptual model* of emotion–cognition interactions. Then, various instantiations of this model are presented. In Chapter 6, I model the perception of the *peripersonal space* when modulated by emotionally valenced sensory and physiological signals. Last, in Chapter 7, I introduce the concept of *Emotional Metacontrol* as an example of emotion–cognition interaction. It consists in using emotional signals elicited by self-assessment to modulate computational processes – such as attention and action selection – for the purpose of behavior regulation.

Part I

Spatial navigation as a framework for cognition

Neuromimetic approach to robotic navigation

"[T]he only function of the brain from an evolutionary point of view is to control movement and so interaction with the environment. That is why plants do not have brains."

(Wolpert [2002])

Navigation is the act of moving, or finding a way, from one place to another. It is a fundamental cognitive function that is crucial for survival and is common to most species. Navigation provides an ideal framework to study biological and artificial cognition because of the variety of computational processes that are involved [Wystrach and Graham, 2012]. Another interest is the low dimensionality of the behavioral output – compared to object grasping for instance.

In the last decades, extensive research in ethology and neuroscience enabled a tremendous progress in understanding how the brain solves navigation problems. In fact, the discovery of place cells and grid cells – which are specific populations of neurons exhibiting firing activities related to spatial localization – earned John O'Keefe, May-Britt Moser and Edvard Moser the Nobel Prize in Physiology or Medicine in 2014. Nevertheless, there is a lot more we have not learned yet about animal and human navigation.

In robotics, navigation is also a major research topic. Mobile robots, along with robotic arms, are among the oldest systems that roboticists have dealt with. Still, being able to self-localize and consistently move in unknown, dynamic environments raises numerous challenges in the domain. A "classical" approach to these issues is known as the *Simultaneous Localization and Mapping* (SLAM) problem [Smith and Cheeseman, 1986; Leonard and Durrant-Whyte, 1991]. This approach considers that the robot has to build an internal representation of the environment in order to correctly navigate in it. In this view, the main challenge lies in the ability to construct such a map while moving; much like a chicken-and-egg problem.

In this thesis, I will rather be interested in biomimetic approaches to navigation. *Machina speculatrix*, better known as the "tortoise" robot, by Walter [1950] is probably the first of the kind. This wheeled robot was able to avoid obstacles and autonomously find its way to a recharging station when required. By creating animal-like machines, the purpose of Walter's work was to demonstrate that complex behaviors like those enabled by the brain come from the rich connectivity between small number of neurons. Later, the *animat* approach to robotics arose from
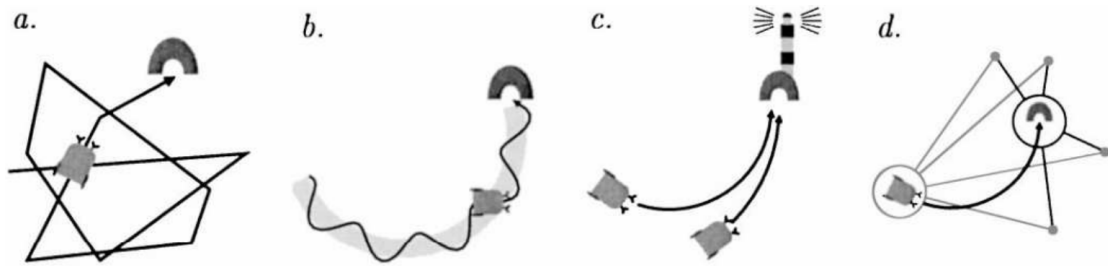
Figure 1.1: Increasingly complex local navigation behavior. From **LEFT** to **RIGHT**: random search, direction-following, aiming, and guidance. The figure is borrowed from Franz and Mallot [2000].

converging efforts aiming to reproduce biological behaviors in artificial systems [Braitenberg, 1984; Brooks, 1991].

Franz and Mallot [2000] proposed a classification of navigation strategies that animals and animats can use. First, a dichotomy is suggested between local and way-finding strategies. In the former, the agent can perform: random search (in the absence of sensory cues), direction-following (in the presence of a trail), aiming (when a salient cues signal the goal location), or guidance (when the spatial relation between cues determines the goal location). The second family of strategies consist in the deduction of a route from an internal representation of the environment; and possibly the ability to adapt to changes and/or deduce new paths or shortcuts. In my thesis, I will rather focus on the first class to study and implement local, unplanned navigation behaviors (i.e. sensorimotor behavior).

1.1 Navigation in the brain

The hippocampus (HP) and neighbor areas are the cornerstone of spatial representations and computations required in navigation. In vertebrates, HP is located in the medial temporal lobe, underneath the cortical surface. The hippocampal system comprises the entorhinal cortex (EC), perirhinal and postrhinal cortices, septum, and the hippocampal formation. The latter, represented in Figure 1.2, constitutes the hippocampus *proper* and includes the dental gyrus (DG), subiculum and Cornu Ammonis divided into CA1, CA2 and CA3.

One of the major findings implicating HP in navigation is the discovery of *place cells* (PC) in rats by O'Keefe and Dostrovsky [1971]. These neurons exhibit activities that are highly localized in space. In other words, the generated firing patterns thus encode different locations, in a sort of 'cognitive map' of a familiar environment. The spatial correlates of the neural firing, called 'place fields', vary in size, depending on the size of the environment and also on their location (e.g. smaller near the arena borders). Place cell activities appear seconds after a rat enters the arena but may take weeks to stabilize [Wilson and Mcnaughton, 1993]. However, once the activities converge, they persists for weeks [Thompson and Best, 1990].

Place cells were initially recorded in hippocampal CA1 and CA3. However, similar neural activities are found in DG [Jung and McNaughton, 1993], EC [Quirk et al., 1992], and even the retrosplenial cortex (RSC) [Smith et al., 2012]. Place fields are smaller in DG while wider and noisier in EC and RSC. In addition, spatially localized firing in DG and RSC is also directionally modulated. These findings suggest many levels of multimodal integration participating in spatial localization.
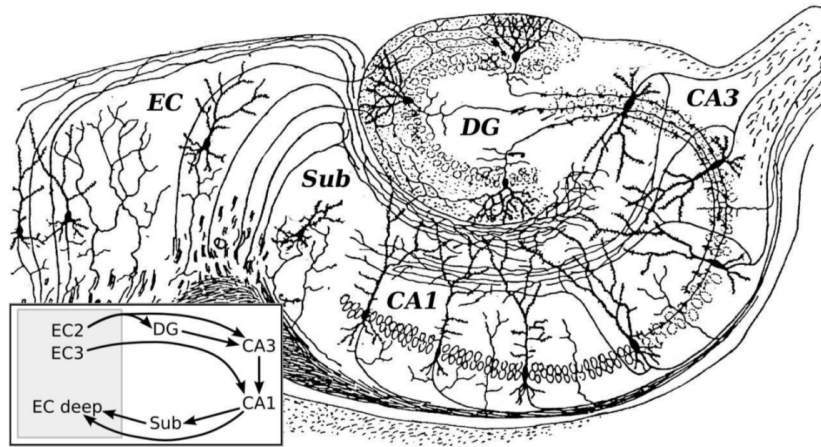
Figure 1.2: Sagittal section of the rat hippocampus. EC: entorhinal cortex. Sub: subiculum. DG: dental gyrus. CA1/CA3: Cornu Ammonis 1 and 3. Connections between these structures are also illustrated. The figure is borrowed from Hanoune [2015], readapted from the original figure by Cajal [1911].

One of the most important sensory modalities in place cell activities is vision [O'keefe and Conway, 1978]. In particular, Cressant et al. [1997] provides evidence of the importance of distant landmarks compared to objects in the arena. Besides, Muller and Kubie [1987] showed that rotating a visual cue in a cylinder-shaped arena induced the same rotation on the place fields. Also scaling up the cylinder made the field scale up likewise. However, the place cell firing remains relatively stable if one of the visual cues is removed [Muller and Kubie, 1987]. Also, darkness does not affect them if it is preceded by a light period [Quirk et al., 1990]. This proves that other modalities intervene in place cell encoding.

In addition to a variety of sensory cues like odors and sounds that may participate in spatial localization, self-motion information is of particular importance. This idiothetic information have diverse sources: vestibular, proprioceptive and, to a lesser extent, visual (optical flow). In the vestibular system located in the inner ear, the semicircular canals and otoliths are respectively sensitive to rotational and linear movements. Various experiments indicate the role of the former in the perception of angular motion in rodents and humans [Mittelstaedt and Mittelstaedt, 1982; Etienne et al., 1988; Metcalfe and Gresty, 1992]. Vestibular information is complemented by proprioceptive feedbacks, as it has been shown that actively walking provides more accuracy to the perception of linear movements [Etienne et al., 1988]. Last, in the absence of the proprioceptive and vestibular information, optical flow can allow for estimating self-motion [Kearns et al., 2002].

The capacities to perceive self-motion is essential for *path integration* (PI): the process of deducing the way back to a starting point by integrating inertial signals. It is also called 'dead reckoning', in reference to the deduced ('ded') reckoning historically used by sailors to navigate across featureless open sea. PI is often observed in *homing* behaviors of a variety of species, from ants and bees to rodents. For instance, desert ants are able to return to their nest after journeys of hundreds of meters without any external cues [Wehner and Srinivasan, 1981]. Figure 1.3 illustrates the difference of PI accuracy across species of vertebrates and mammals, including humans. The vestibular system turns out to be essential in this process, as impairment has been shown in damaged rats [Stackman and Herbert, 2002] and humans Cohen [2000].
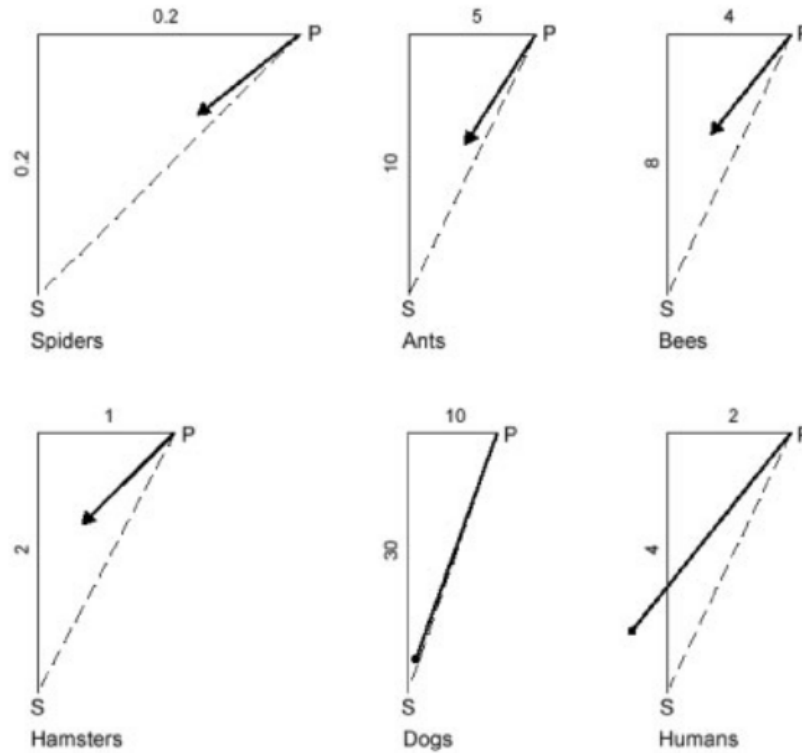
Figure 1.3: Accuracy of path integration in spiders, ants, bees, hamsters, dogs and humans. The figure is borrowed from Etienne and Jeffery [2004].

It is crucial to keep track of one's orientation when navigating only based on self-motion. It is also when using external cues, either to readjust dead reckoning or to perform landmark-based localization. *Head direction cells* merge multimodal information and self-generated movements to provide the neural equivalent for an internal compass [Taube, 2007]. In other words, the firing of this type of neurons increases as the deviation of the current head direction from the preferred direction of the cell decreases. These neurons were initially discovered by Ranck Jr [1984] in the presubiculum and postsubiculum, which are part of the subicular complex along with the hippocampal subiculum. But as stated earlier, directionally selective neurons are also found in other regions like EC and RSC. Studies have implicated the thalamus in the head direction system. For example, neurons in the anterior thalamus fail to fire in their preferred directions if the animal is unable to move [Taube et al., 1996]

Another discovery confirming the central role of the hippocampal area in navigation are *grid cells* in EC by Hafting et al. [2005]. When recorded in sufficiently large environments, these cells present activities forming a regular triangular pattern in the shape of a grid – hence the name. Thus, rather than a specific spatial firing field, a grid cell signals the distance and angle to other firing locations. In rodents, grid cells depend on external landmarks. However, they also persist in darkness with a decrease in the spatial correlation of the rate maps and a dispersal of the grid vertices (Hafting et al., 2005). Moreover, EC hosts another type of neurons called *conjunctive cells* [Sargolini et al., 2006]. These cells exhibit activities combining position and direction information. An interesting property is that only encountered patterns are encoded –

instead of all the possible combinations – thus providing compressed codes that can be used for spatial localization.

1.2 Biologically inspired models for robotic navigation

A large variety of biologically inspired models exists in the literature of robotic navigation [Gaussier and Zrehen, 1995; Burgess et al., 1994; Arleo and Gerstner, 2000; Milford et al., 2004]. One of the first computational models of place cells was proposed by Burgess et al. [1994]. This mathematical model was directly inspired by the finding of O'Keefe and colleagues on biological place cells. Initially tested in a simulated environment, it was later implemented on a real robot [Burgess et al., 1997]. Place cell activities are determined by the perception of visual landmarks in the environment. Then, connections between concurrently active place cells are learned in order to allow the robot to return to a goal location. Contemporaneously, Redish and Touretzky [1997] worked on navigation-related processes in the hippocampus. In their model, place cells are encoded by combining visual (local views) and vestibular (head direction) inputs while path integration merge the latter with self-motion information. These spatial codes formed in HP are provided to basal ganglia for the purpose of action selection. Similarly, Arleo and colleagues proposed a collection of models combining allothetic (visual) and idiothetic (self-motion) information allowing robots to encode places and perform path integration [Arleo and Gerstner, 2000; Arleo et al., 2004; Arleo and Rondi-Reig, 2007].

'RatSLAM', proposed by Milford et al. [2004], is also a model taking inspiration from the rat hippocampus for the purpose of robotic navigation. In comparison to the examples given above, this solution aims less at the biological plausibility than the robustness and scalability from the robotics perspective. This initial work was followed by numerous variants developing certain aspects of the model and evaluating it in different contexts, including indoors and outdoors. For instance, the system was able to build a map of the road in an area of about 0.3 Km by 1.6 Km using only a webcam for both visual recognition and odometry [Milford and Wyeth, 2008]. The upper part of the image was used to detect rotations and estimate the pose while the lower part provided the speed estimation. The visual place recognition took as input a projection of the subsampled image over the horizontal axis to obtain a vector representing the intensity profile. Mounted on a car, this solution combined the pose and view information to passively build the map of the visited area. More recently, features from RatSLAM were incorporated in more plausible modeling of rodents environment mapping, and entorhinal grid cells in particular, namely HiLAM. In this architecture proposed by Erdem et al. [2015], HiLAM builds multiple scales of grid cells by extracting idiothetic information from recorded visual data. The motion estimation is rectified using the loop-closure detection feature of RatSLAM (off-line). This combination provides a scalable solution for goal-directed navigation using a hierarchical grid cell and place cell mapping. However, the learning is performed on recorded datasets and no experiments with an on-line interaction between RatSLAM and HiLAM is presented.

1.3 PerAc and the LPMP model for place cells-based navigation

Part of my thesis research relies on the works of Gaussier and colleagues from the Neurocybernetics team of ETIS laboratory and partner laboratories. Pioneering to this research is the Per-Ac model proposed by Gaussier and Zrehen [1995]. This model advocates the idea that perception,

as the ability to build and use representations based on sensory inputs, requires action. PerAc, illustrated in Figure 1.4, thus proposes a generic solution to build sensorimotor associations from the interactions of an agent with its environment. In the context of navigation, these associations can consist in place–action coupling controlling the robot trajectory: go in direction D when in location L. This model is complemented with a large body of research on biologically plausible models of the hippocampal region and place cells, head direction cells, and grid cells, as well as cognitive maps formation [Gaussier et al., 2002, 2007] implemented in simulations and real robot experiments both indoors and outdoors [Gaussier et al., 1997, 2000; Giovannangeli et al., 2006; Giovannangeli and Gaussier, 2008; Cuperlier et al., 2006, 2007; Jauffret et al., 2015].
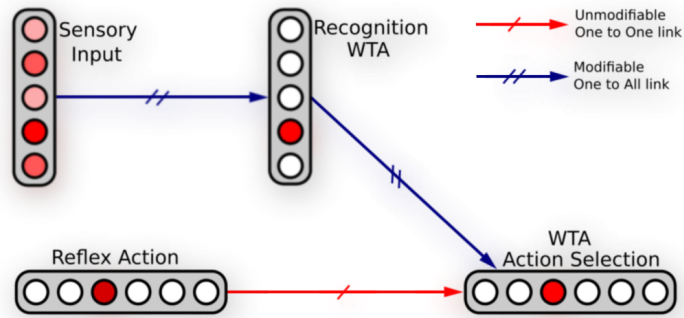
Figure 1.4: PerAc model for sensation–action coupling. WTA: winner-takes-all, i.e. strict competition. The unconditional pathway (bottom, unmodifiable link) triggers reflex behavior. The conditional pathway (top, modifiable link) associates the most active category representing the sensory input with the reflex action. After learning, the recognition of that sensory category with trigger the learned behavior. The figure is borrowed from Jauffret [2014], readapted from the original figure by Gaussier and Zrehen [1995].

In particular, Gaussier and Zrehen [1995]; Gaussier et al. [2002] presented a model of places cells based on visual inputs. In this model, the ventral stream provides the *'what'* information – i.e. the identification of objects in the visual scene – while the dorsal stream provides the *'where'* information – i.e. the location of the objects. The 'what' and 'where' information, transmitted through the perirhinal cortex and parahippocampal cortex respectively, can thus be merged to encode places in EC or DG.

Here, I will detail a particular implementation of this model that I will refer to as the LPMP (Log-Polar Max-Pi) model: encoding local visual descriptors using log-polar transformation and merging what and where information by a max-pi operation (max of tensor product). In this model, a saliency map is obtained by a method of corner detection in order to extract points of interest. To do so, a Deriche [1987] filter is applied on the grey scale image. The so obtained gradient image is convolved with a DoG (Difference of Gaussian) filter to extract corners. Then, local views around these salient points are transformed using log-polar mapping. Such visual descriptors representing regions of interests in the image are then encoded as signatures to provide the 'what' information; whereas their position along the horizontal axis (i.e. azimuths) represents the 'where' information. The rest of the model is illustrated in Figure 1.5.

The activity of each neuron s_i representing visual signatures at time t is given by the follow-
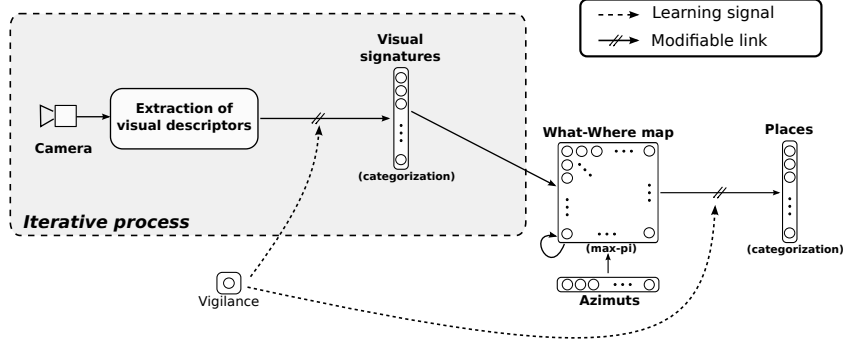
Figure 1.5: The model used for place recognition. The W–W map allows for active vision in the case of local descriptor. It is also used to virtual reconstruct a panorama from various images.

ing equations:

$$s_i(t) = 1 - \frac{1}{N_D} \sum_{j=1}^{N_D} |w_{ij}^D(t) - d_j(t)| \tag{1.1}$$

where d_j is the j^{th} element of the descriptor vector of size N_D and w_{ij}^D is the weight of the synaptic link between s_i and d_j. A binary vigilance signal triggers the recruitment of a new neuron to encode a new category of local view. The learning rule of these neurons is the following:

$$\frac{dw_{ij}^D}{dt} = \delta_j^R |d_j^g(t) - w_{ij}^D(t)| \tag{1.2}$$

where R is the index of the newly recruited neuron and δ_a^b the Kronecker delta, equal to 1 if $a = b$ and 0 otherwise. With this type of neurons, the idea is to save the input pattern in the links weights. Thus, the closer the input is to the learned pattern, the stronger the neurons activations. This method has good generalization properties for robotic place recognition [Giovannangeli et al., 2006].

In addition, the model relies on the 'where' information; that is to say the orientation in which the visual feature is observed. The input is a vector in which each neuron has a preferred direction around the yaw axis. The activities of the azimuth neurons α_i are obtained after a lateral diffusion around the neuron coding for the direction of the current visual input. In our case, it corresponds to a gaussian bell with standard deviation σ_{azim}. It integrates information about the orientation of the body (robot), head (camera) and fovea (center of the visual feature). Using a magnetic compass is the easiest way to obtain the body/head direction but we could also extract it from vision, odometry and other modalities [Giovannangeli and Gaussier, 2007; Delarboulas et al., 2014].

The 'what' and 'where' 2D map (W–W map) is a second-order tensor M in which each neuron codes for a signature-azimuth couple. The serial exploration of the visual scene requires a short-term memory. Therefore, the W–W map stores previous activities while the visual scene exploration is still in progress (i.e. before the end a panorama). In order to reduce the computational cost, the 360^o surrounding field is discretized in $N_A = 5$ orientations before the computation of the tensorial product. Then, the activity of the $N_M = N_S \times N_A$ W-W tensor is:

$$M(t) = \max[(s \otimes a), M(t - dt).(1 - R(t))] \tag{1.3}$$

where N_S is the size of the signatures vector, s and a are the signatures and azimuths vectorial representations, R a binary reset signal triggered at the end of a panorama and $\otimes$ is the tensorial product operator.

Pattern of activities in the W–W map code for the current place. Such patterns can be categorized in a place vector, in which each neuron p_i has the following activity at time t:

$$p_i(t) = 1 - \frac{1}{\rho_M . N_M} \sum_{j=1}^{\rho_M . N_M} \max_{i=1..N_M}{}^{k} |w_{ij}^M(t) - m_j(t)| \qquad (1.4)$$

where m_j is the j^{th} element of the tensor M, w_{ij}^M is the weight of the synaptic link between p_i and m_j, ρ_M is the proportion of W–W couples required for place recognition and $\max_{i \in I}{}^{k}$ is an operator that gives the k^{th} maximum in I. The learning rule is the same as the equation (2.3), replacing $w_{ij}^D(t)$ by $w_{ij}^M(t)$.

Please note that, as compared to the equation (1), the factor ρ_M has been introduced for the purpose of robustness to occlusions. This parameter also compensates the absence of activation thresholds in the previous layers of the neural network. It represents an estimation of the ratio of the W–W pattern used to code for a place that is necessary and sufficient in order to recognize it [Giovannangeli et al., 2006].

For the purpose of fast learning of desired trajectories in robotic tasks, the category recruitment (i.e. leaning signatures and places) is performed in a one-shot way. In other words, the vigilance signal represented in Figure 1.5 is binary. Figure 1.6 illustrates the kind of place cell activities the we obtain with this model. The firing rates are the highest at the center of the place field – that is where learning occurred. This type of activities with important generalization properties could correspond to entorhinal rather than hippocampal place cells.
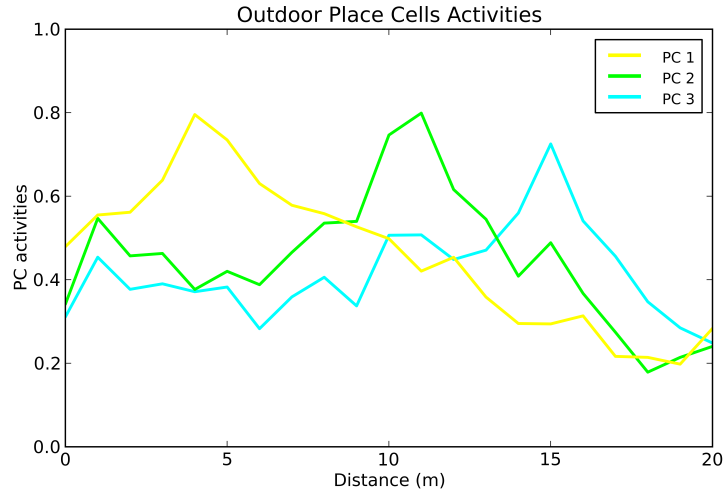
Figure 1.6: Activities of 3 visual place cells recorded on a linear track in a real outdoor environment. The highest activities correspond to the learned positions of the associated cells. The architecture exhibits good generalization properties and neurons encode large place fields.

1.4 General discussion

In this chapter, I presented a brief overview of the literature on biologically inspired navigation. After a review of the main neuroscientific findings influencing robotic and computational models, I focused on place cells modeling and place recognition. Indeed, whereas some basic local strategies like road-following can do without it, self-localization seems essentials whenever a specific, more or less complex trajectory has to be adopted. According to Levitt and Lawton [1990], "where am I?" is the first question that defines navigation – followed by "Where are other places relative to me?" and "How do I get to other places from here?".

In that matter, unlike current robots, mammals are able to recognize visited places and to localize themselves in nearly any sort of environments. Therefore, taking inspiration from biological systems has often been a good way to build robust and adaptive systems that can also be competitive with non-biologically inspired models in terms of performance [Giovannangeli et al., 2006; Milford and Wyeth, 2008; Siagian and Itti, 2009; Jauffret et al., 2015; Erdem et al., 2015]. Here, I detailed the PerAc and the LPMP models developed by Gaussier and colleagues [Gaussier and Zrehen, 1995; Gaussier et al., 2002; Giovannangeli et al., 2006]. The next two chapters will tackle the use and extension of the existing models in the case of outdoor navigation and in the context of interactive learning of sensorimotor navigation behavior respectively.

Outdoor place recognition based on visual contexts

As seen in the previous chapter, place recognition is one of the abilities that mobile robots may need in order to navigate properly. Some approaches rely on range sensors (e.g. sonars and ladars) or GPS, either exclusively or combined with other sensors. However, the former generally shows little robustness to irregularities that are typical in outdoor environments and the latter are useless indoors when satellite signals are shadowed. On the other hand, vision is an essential modality that allows mammals to recognize visited places and to localize themselves in nearly any sort of environments, also with the help of other types of sensory information (e.g. olfactory, tactile, or proprioceptive).

Visual place recognition is still a hot topic in robotics [Lowry et al., 2016]. The most difficult challenges lie in outdoor navigation where the environment is much less structured and controlled. Part of the research on the topic focuses on the condition-invariance problem, studying the robustness to perceptual changes occurring when switching from day to night or from sunny to rainy weather for example [Milford et al., 2014]. This issue will not be tackled in this work. Moreover, there is a distinction between on-road and off-road navigation. In the first case, the robot moves within a determined path or trail. From the visual place recognition perspective, it implies very few or no deviations at all from the initial trajectory; and thereby very few variations in the visual scene. On the other hand, off-road navigation implies less constrained trajectories and the need to recognize places from different points of view; which is known as the pose-invariance problem. In the extreme case, the goal is to classify locations regardless to the robot position and the camera orientation. In my work, I am more interested in obtaining generalization properties that provide robustness with regard to variations in the robot trajectories. Another issue that is inherently (although not exclusively) linked to outdoor navigation is the efficiency over long distances. The implication – in the sense of pattern recognition in general – is the need to store and retrieve a much larger amount of information. In that matter, the goal of this chapter is to introduce a context-based model intended to overcome the computational cost of place recognition in wide environments.

Existing approaches in visual place recognition can be based on global or local descriptors: some try to capture the global aspect of the scene by considering the image as a whole while others extract and encode particular subparts of the image, called regions of interest. While

the former methods are faster[1], they have generally shown less effective than the latter [Van de Sande et al., 2008; Siagian and Itti, 2009]. In the first part of this chapter, I will list some of these descriptors and compare their performance in the specific context of neurorobotic outdoor navigation.

There is also a notable number of works proposing the combination of local and global visual information in an attempt to take advantage of the strengths of both type of methods [Lisin et al., 2005; Siagian and Itti, 2009; Rostami et al., 2014; Milford et al., 2014]. In that matter, the second part of this chapter will present a model for context-based place recognition, inspired by the notions of visual contexts [Bar, 2004; Chun, 2000] in the case of global precedence [Navon, 1977; Oliva and Torralba, 2006].

2.1 Local and global visual descriptors

Scene recognition models can be classified based on whether they use local or global visual descriptors. Global (holistic) vision methods consider the image as a whole and encode it as a single vector. Because they are typically very compact, they allow fast computation.

Various models using this kind of descriptors can be found in the literature. For instance, Milford et al. [2014] proposes to subsample the whole image and use it as a global signature. In former work Milford and Wyeth [2008], the subsampled image is also projected over the horizontal axis to obtain a vector representing the intensity profile. Global descriptors can be encoded in histograms as well [Van de Sande et al., 2008; Rebai et al., 2014]. In addition, some solutions construct signatures based on statistical moments (e.g. mean, variance or possibly higher order moments) [Van de Sande et al., 2008; Siagian and Itti, 2009]. Three types of information are commonly used to compute visual descriptors. First and foremost evident is the luminance channel Milford and Wyeth [2008]; Milford et al. [2014]. Besides, chrominance channels can also be informative [Van de Sande et al., 2008; Rebai et al., 2014]. Moreover, orientations are very useful to describe textures both indoors (e.g. doors and computers screens) and outdoors (e.g. trees, roads and distant buildings) [Riesenhuber and Poggio, 1999; Oliva and Torralba, 2006]. Notably Siagian and Itti [2009] encode these three types of information simultaneously.

On the other hand, local descriptors only carry information relative to certain regions of interest in the image. State-of-the-art methods, like SIFT Lowe [2004] and SURF Bay et al. [2008] typically implement this kind of solution. In addition, the LPMP model I presented in the previous chapter also encodes local information. Place cells are thus obtained from patterns of landmark–azimuth observed by the robot [Gaussier and Zrehen, 1995; Gaussier et al., 2002; Giovannangeli et al., 2006]. As for the visual descriptor, it is a log-polar transformation of the most salient regions of interests. The Table 2.1 summarizes the differences between this technique (labeled *rhotheta*) and SIFT Lowe [2004] and SURF Bay et al. [2008].

Holistic methods have generally shown less effective than local features-based ones [Van de Sande et al., 2008; Siagian and Itti, 2009]. Indeed, the latter use richer information: a certain number of signatures per image ('what' information) and their corresponding positions ('where' information). However, they can be useful in situations where stability of points of interest detection is difficult to ensure (e.g complex textures like tree leaves or condition variation). But

---

[1]While this is generally true, it obviously depends on a variety of factors: size of image, size of regions of interest, number of regions extracted per image, encoding methods, etc.

Table 2.1: Comparison between Local rhotheta, SIFT, SURF and GIST

	rhotheta	SIFT	SURF	GIST
Multiscale input	No	Yes	Yes	No
Saliency map	DoG	DoG	DoB(Box)	-
PoI extraction	Local extrema	Local extrema	Hessian matrix	-
Descriptor type	Log polar mapping	Orientation histogram	Orientation histogram	Orientation histogram

usually, visual descriptors are evaluated in the context of object recognition [Van de Sande et al., 2008; Kihl et al., 2015], face recognition Yan et al. [2007], or video retrieval Van de Sande et al. [2008]. Siagian and Itti [2009], who tackled robotic navigation, only included SIFT and GIST descriptors.

So this first experiment will aim at identifying common properties, strengths and weaknesses of a set of global descriptors as well as situations where they could outperform a local vision method. Thus, I consider seven global descriptors. First, based on the above listed methods, I select a subset of possible combinations between three visual channels (luminance, chrominance and orientation/texture) and three encoding techniques (subsampling, histograms and statistical moments). Subsampling is tested on the intensity and the chromatic channels (referred to as *proflum* and *profcol*). Histograms are used to encode color and orientation information (*histocol* and *histogab*). Statistical moments are calculated for orientations (*momgab*). This choice has been made so that each type of channel and encoding technique is tested at least once. There was also an attempt to avoid as much as possible using descriptors that are too poor or too large (like a histogram of light intensity or a subsampling of the output of a large filter bank for instance). Indeed, global descriptors are meant to be lightweight, as long as they contain enough information. Two additional descriptors from the literature are added to the list: GIST descriptor [Siagian and Itti, 2007] (referred to as *gridgab*) and Milford's subsampling of the whole image [Milford et al., 2014] (labeled *subslum*). More details on the implementation of these visual descriptors will be given below.

These seven global descriptors are compared to the log-polar transformation of local views (referred to as *rhotheta*) used in the LPMP model. This instance of local descriptors is biologically plausible [Schwartz, 1977], has relatively little computational cost, is invariant to small rotations and scale variations, and gives good place recognition results in neurorobotic experiments [Giovannangeli et al., 2006]. It is expected to outperform holistic descriptors.

2.1.1 Implementation details

rhotheta is used for local views encoding. First, a Deriche [1987] filter in applied on the grey scale image. It consists in a derivative and a smoothing filter that allow edge detection. In this method, the parameter α determines the gradient resolution.

Then, the output is convolved with a DoG filter consisting in two gaussians of standard deviations σ_{DoG_1} and σ_{DoG_2}. The result is a saliency map providing points of interest. Local views are extracted around the N_{PoI} most salient points between two disks of radius r_{small} and

r_{big}. To avoid redundancies, two PoI cannot be closer that $r_{big}/2$. Then, local views are encoded using a log-polar transformation. Thereby, a descriptor of size $N_{\rho\theta}$ is obtained.

proflum represents a scanline intensity profile, like in RatSLAM by Milford and Wyeth [2008]. The grey scale image is subsampled at a factor of κ_l. Then, a 1D vector represents the normalized sum of the pixels intensity in each column. Which means that a descriptor of size $(view_W/\kappa_l)$ is obtained.

profcol, likewise, represents a profile for each of the chromatic dimensions of an image represented in the *Lab* color space. In this representation, *a* and *b* are color-opponent dimensions, respectively coding for the Red-Green and the Yellow-Blue axis. The *Lab* color space was designed to approximate human vision. In this experiment, unlike the *RGB* space, it allows for easily removing the lightness component of the image and only encode the chrominance. Also, as compared to the *HSV* space used in similar work [Rebai et al., 2014; Van de Sande et al., 2008], the two remaining color dimensions are homogeneous. Thus, two vectors represent the normalized sum of the pixels intensity in each column of the *a* and *b* images subsampled at a factor of κ_c. The descriptor size is thus $(2.view_W/\kappa_c)$.

histocol associates a histogram of size $N_{H_{col}}$ to each of the chromatic dimensions of an image represented in the *Lab* color space. Thus, a descriptor of size $(2.N_{H_{col}})$ is obtained.

histogab associates a histogram of size $N_{H_{gab}}$ to each of the outputs of a Gabor filter bank. Gabor filters are defined by a sinusoidal wave multiplied by a Gaussian function and allow orientation detection. The complex representation for a 2D filter is the following:

$$g(x,y) = e^{-\frac{x'^2 + \gamma^2 y'^2}{2\sigma_g^2}} . e^{i(2\pi\frac{x'}{\lambda} + \psi)} \tag{2.1}$$

with $x' = x\cos\theta + y\sin\theta$ and $y' = y\cos\theta - x\sin\theta$ where γ is the spatial aspect ratio, σ_g the standard deviation of the gaussian function, λ the sinusoidal factor wavelength, ψ the phase offset and θ the preferred orientation of the filter. Here, the filter bank is used to detect 4 orientations (0, $\pi/4$, $\pi/2$ and $3\pi/4$). In addition, I consider 3 spatial scales (represented by λ_1 λ_2 λ_3) with the same filter size (represented by σ_g). Thus, the descriptor size is $(12.N_{H_{gab}})$

momgab captures the 1st- and 2nd-order statistical moments of the outputs of a Gabor filter bank used to detect 4 orientations on 3 spatial scales (detailed above). Since only the mean and the standard deviation are encoded, the descriptor size is equal to $(24 = 2 \times 12)$.

gridgab is an implementation of GIST descriptor proposed by Siagian and Itti [2007, 2009]. The outputs of a Gabor filter bank used to detect 4 orientations on 3 spatial scales are divided into 4-by-4 grid subregions. The descriptor is a $(192 = (4 \times 4) \times 12$ values vector containing the mean of each subregion.

subslum is an implementation of the whole image descriptor as described by Milford et al. [2014]. The luminance channel is subsampled at a factor of 100. The descriptor size is thus $(view_W * view_H/100)$

noise is a random descriptor of size N_{noise} based on a uniform white noise, intended to serve as a baseline.

All the parameters are listed in Table 2.5.

2.2   Experiment 1: Comparison of local and global descriptors

This study is performed on three datasets[2] recorded in the area around the university of Cergy-Pontoise using the RobuROC platform. The images were captured by a fisheye camera and a magnetic compass was used to acquire the orientation data. The datasets were recorded in a way that is similar to how real robots navigate and localize themselves in our field experiments: 1) In order to learn a new place, the robot camera captures 15 images over a 360 degrees panorama. During this process, the robot stays still in order to avoid distortions in the representation of the place. 2) When exploring the environment (i.e. the rest of the time), the robot only captures 7 images per panorama, for the purpose of faster place recognition.

The datasets consist in trajectories recorded in visually different environments (see Fig. 2.1). Dataset *C* (23.1 meters long) simulates on-road navigation in a quite structured but highly dynamic environment (moving cars and pedestrians). Datasets *A* and *B* (20 meters long each) simulate off-road navigation. These environments are less structured and buildings are more distant. As illustrated in Figure 2.1, each dataset contains a set of exploration databases including trajectory variations to address the pose-invariance problem.

New datasets were recorded instead of using existing datasets because one of the goals of the experiment is to evaluate the generalization properties of the model within distances that require the ground truth of the robot position at the scale of a meter. Also, variations in orientation of the robot were made possible. For the sake of evaluation simplicity, all trajectories are linear on x and y axis, although perturbations were induced by the rugged nature of the field. However, the camera is stabilized using a Kalman filter to limit the pitch and roll as much as possible. The robot speed is constant during exploration. Thereby, we calculate the ground truth by estimating the robot position at every step, knowing that: 1) a place is learned every $d_{learn} = 3.9 \pm 0.1$ meters in average; and 2) an exploration panorama is completed after traveling $d_{explo} = 0.71 \pm 0.01$ meters. Despite the small distance, the interest of these datasets is the variety of trajectories. The precise acquisition of the ground truth on which this evaluation is performed would be much more difficult to do in very large scale and long term experiments.

In order to compare visual descriptor performances on these datasets, we feed them to the neural network described in the LPMP model. Which means that like *rhotheta*, all the other descriptors serve as inputs to the feature categorization layer which provides the 'what' information. The latter is combined with the 'where' information (i.e. orientation of the center of the image in case of a global descriptor) to obtain patterns of signature–azimuth which are categorized to represent places.

This work aims at a rigorous evaluation of the above presented descriptors in order to characterize their performances in the specific context of on-road and off-road robotic navigation. In this context, two criteria are essential: generalization and recall-precision trade-off. Measuring recall and precision aims to verify that the system capacity to return relevant answers does not decrease dramatically when tuned to return as much elements as possible. It is an evaluation from the information retrieval perspective. The generalization criterion consists in the system's ability to return relevant elements in new situations by recognizing common characteristics shared with learned patterns. In navigation, such property is crucial so that the robot can correctly perform in the real world. In particular, topologically close locations should have similar recognition levels. It also allows for using control mechanisms that generate smooth

---

[2]All datasets can be downloaded via this link: https://cloud.ensea.fr/index.php/s/lvJrBOPLTvLAkc8
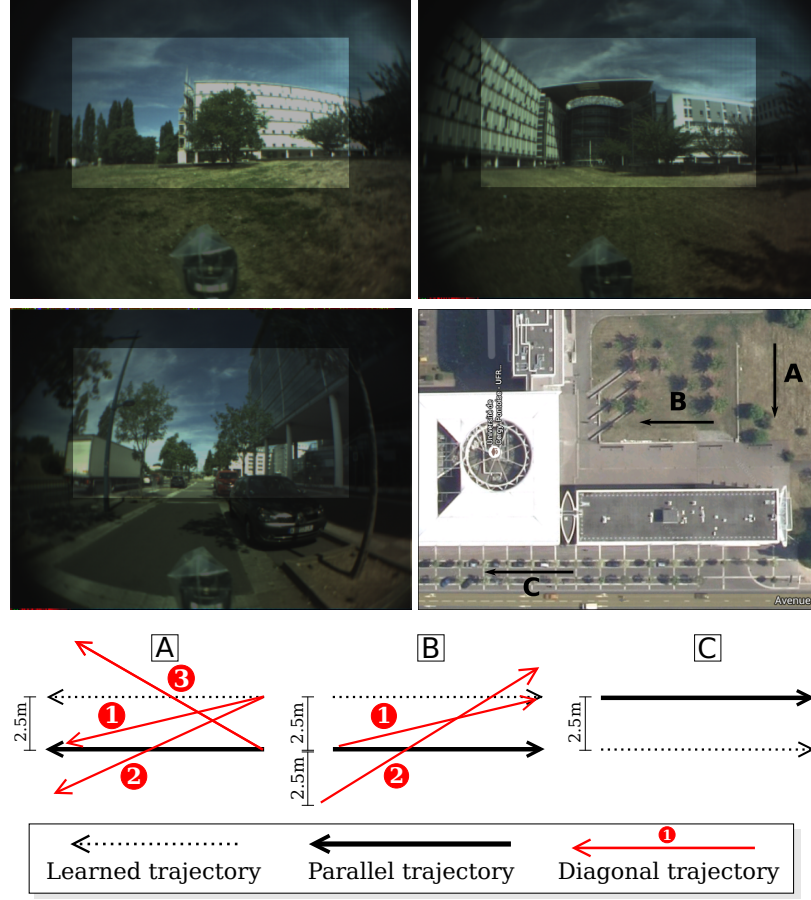
Figure 2.1: TOP: Samples from the 3 dataset (A: top-left, B: top-right, and C: center-left) and a map representing locations where they were recorded. In order to reduce distortions introduced by the fisheye lens and to remove the irrelevant bottom part, the image is cropped to consider only the highlighted 440×240 subimage. BOTTOM: The exploration trajectories recorded on each dataset. In addition to "parallel" and "diagonal" trajectories, another exploration database is recoded on the learning trajectory and is not represented for clarity purpose.

trajectories by averaging their responses when several situations are recognized well enough.

In this experiment, I propose three measures of performance:

3WD (3 winners distances): It measures the average distance (in meters) between the 3 best recognized places and the position where they were originally learned. It characterizes the systems ability to generalize and not only recognize places at the precise location where the corresponding visual features were learned. For the sake of uniformity among trajectories evaluation, all exploration positions are projected on the learning trajectory and normalized as explained in Figure 2.2. Thus, the results are greater or equal to 1; 1 being the best results.

MAP (mean average precision): It is a traditional compact representation of the recall-precision curves. It is the mean of the average precision at every position. Indeed, the neural network generates new place cells activities every time the W-W maps are updated by a new visual panorama. This is analogous to a new query in classical information retrieval systems. So we can calculate the system precision (i.e. well ranked place cells responses) depending how
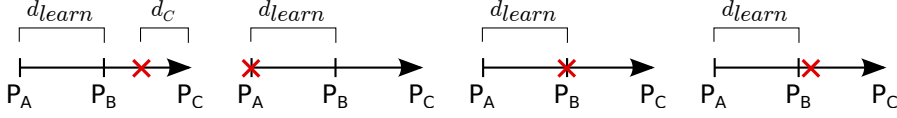
Figure 2.2: Examples of distances from place "centroid". Four positions X (red crosses) are considered on the segment $[P_A, P_C]$ of a trajectory where the closest learned places P_A, P_B and P_C are the winners. In these examples, the sum of distances between X and P_A, P_B and P_C is $2.5 \times d_{learn}$, $3 \times d_{learn}$, $2 \times d_{learn}$, and $2.2 \times d_{learn}$ respectively. This value is bounded by $2 \times d_{learn}$ and $3 \times d_{learn}$ which respectively correspond to a retrieval on a learned spot in the middle of the trajectory and one on the starting or ending points. The average value $2.5 \times d_{learn}$ is used to normalize the 3WD measure.

many place cells activity would be considered in the output. MAP scores are in $[0, 1]$ and 1 corresponds to a perfect precision.

WNR (winners-to-noise ratio): Similarly to a classic signal-to-noise ratio, it compares the level of the desired responses to the level of background noises. In this case, I consider that the most relevant information consists in the average of the 3 winners levels while the noise is the average of the remaining place recognition levels. This measure assesses whether a place cell activity decreases when the robot is far from the initial learning location. Also, higher WNR scores indicate there's less sensitivity to noise and more robustness to small variations. In other words, little risk that the fifth closest place cell be more active and win the competition mistakenly. The WNR ratio is greater or equal to 1 and the greater the better.

3WD and MAP are based on the ground truth estimated from starting and ending positions as well as elementary distances calculated on site in the pretest phase. The three measures evaluate the system from different evaluation perspectives. The results should be taken together, with particular focus on the application context of navigation and place recognition. Indeed, in practice, it does not make much sense to expect high precision for too many responses. Considering the case of a linear navigation trajectory, the place learned at the beginning is very likely to be badly recognized when close to the end of the trajectory. Three responses seem to represent a satisfying threshold beyond which place discrimination is allowed to be noisy and less precise. Put differently, at each point of the line, we can expect the system to i) localize itself at the closest place "centroid", and ii) to correctly determine if it is closer to the previous place or to the next one. This leads to the intuition of considering 3 winners as a satisfying trade-off between a strict recognition (1 winner) and the evaluation of all categories. Hence the consideration of 3 winners in 3WD and WNR. On the other hand, MAP assesses the response of places beyond that threshold to extend the generalization capacity. Although I expect precision to decrease at some point (depending on the trajectory and the visual variation), MAP is calculated for all places of a learned trajectory in order to avoid the selection of another arbitrary threshold.

2.2.1 Results

Off-road datasets

On Datasets A and B, the local descriptor generally obtain the best MAP and WNR scores (See Table 2.2 and 2.3). However, *rhotheta* is outperformed several times by the other descriptors in the 3WD test. Among the global descriptors *gridgab* gets the best overall results. But almost all the others also show better performances than *rhotheta* on distinct trajectories.

Table 2.2: Dataset A

		3WD	MAP	WNR
on-trajectory	rhotheta	1.203	**0.538**	**1.106**
	histocol	**1.084**	0.498	1.097
	histogab	1.359	0.453	1.030
	momgab	1.400	0.416	1.023
	profcol	1.176	0.334	1.005
	proflum	1.228	0.428	1.080
	subslum	1.203	0.497	1.074
	gridgab	1.288	0.476	1.062
	noise	2.357	0.066	–
parallel	rhotheta	1.403	**0.392**	1.073
	histocol	1.662	0.286	**1.076**
	histogab	1.464	0.303	1.026
	momgab	1.393	0.325	1.021
	profcol	1.462	0.269	1.003
	proflum	1.383	0.212	1.058
	subslum	1.422	0.276	1.054
	gridgab	**1.302**	0.298	1.045
	noise	2.417	0.130	–
diagonal 1	rhotheta	**1.126**	**0.555**	**1.088**
	histocol	1.710	0.330	1.070
	histogab	1.512	0.422	1.023
	momgab	1.414	0.353	1.018
	profcol	2.021	0.186	1.003
	proflum	1.638	0.254	1.052
	subslum	1.484	0.295	1.047
	gridgab	1.362	0.277	1.040
	noise	2.266	0.082	–
diagonal 2	rhotheta	**1.439**	0.436	**1.072**
	histocol	1.847	0.218	**1.072**
	histogab	1.545	0.240	1.023
	momgab	1.702	0.347	1.016
	profcol	2.348	0.176	1.003
	proflum	1.848	0.194	1.046
	subslum	1.756	0.213	1.048
	gridgab	1.794	0.237	1.036
	noise	2.312	0.117	–
diagonal 3	rhotheta	1.797	**0.323**	**1.109**
	histocol	**1.359**	0.288	1.096
	histogab	1.921	0.221	1.037
	momgab	1.605	0.239	1.026
	profcol	1.548	0.299	1.006
	proflum	1.503	0.277	1.089
	subslum	1.533	0.294	1.086
	gridgab	1.666	0.284	1.048
	noise	2.193	0.126	–

Notes: 3WD: The closer to 1 the better;
MAP and WNR: The great the better

Table 2.3: Dataset B

		3WD	MAP	WNR
on-trajectory	rhotheta	**1.029**	**0.605**	**1.223**
	histocol	1.657	0.247	1.076
	histogab	1.301	0.379	1.044
	momgab	1.383	0.328	1.029
	profcol	1.537	0.217	1.005
	proflum	1.891	0.136	1.105
	subslum	1.893	0.181	1.103
	gridgab	1.090	0.456	1.083
	noise	2.335	0.076	–
parallel	rhotheta	1.333	0.464	**1.208**
	histocol	1.839	0.215	1.070
	histogab	1.185	0.399	1.050
	momgab	1.370	0.386	1.025
	profcol	1.527	0.306	1.005
	proflum	1.635	0.205	1.098
	subslum	1.651	0.284	1.091
	gridgab	**1.059**	**0.535**	1.079
	noise	2.466	0.127	–
diagonal 1	rhotheta	1.121	**0.613**	**1.224**
	histocol	1.922	0.163	1.073
	histogab	1.321	0.355	1.045
	momgab	1.468	0.332	1.021
	profcol	1.784	0.131	1.004
	proflum	1.848	0.293	1.082
	subslum	1.779	0.283	1.078
	gridgab	**1.109**	0.493	1.072
	noise	1.901	0.144	–
diagonal 2	rhotheta	1.213	0.416	**1.159**
	histocol	1.907	0.240	1.069
	histogab	1.875	0.218	1.034
	momgab	1.719	0.266	1.021
	profcol	1.687	0.218	1.006
	proflum	2.155	0.240	1.101
	subslum	2.092	0.235	1.106
	gridgab	**1.158**	**0.464**	1.061
	noise	2.241	0.100	–

Notes: 3WD: The closer to 1 the better;
MAP and WNR: The great the better

On-road dataset

In general, *rhotheta* outperforms global descriptors as expected but it is second to *gridgab* once (See Table 2.4). Among holistic descriptors, the best 3WD and MAP scores on dataset C are obtained by *gridgab* then *profcol* and *histogab*. However, WNR scores are quite similar.

Table 2.4: Dataset C

		3WD	MAP	WNR
on-trajectory	rhotheta	**1.263**	**0.395**	**1.201**
	histocol	2.237	0.170	1.026
	histogab	2.059	0.201	1.018
	momgab	2.281	0.182	1.014
	profcol	2.158	0.254	1.004
	proflum	2.494	0.201	1.092
	subslum	2.445	0.206	1.087
	gridgab	1.828	0.263	1.048
	noise	2.929	0.115	–
parallel	rhotheta	**1.452**	0.236	**1.159**
	histocol	2.749	0.096	1.024
	histogab	2.583	0.119	1.019
	momgab	2.943	0.093	1.016
	profcol	2.448	0.170	1.003
	proflum	2.679	0.161	1.077
	subslum	2.539	0.197	1.076
	gridgab	1.783	**0.265**	1.031
	noise	2.595	0.094	–

Notes: 3WD: The closer to 1 the better;
MAP and WNR: The great the better

2.2.2 Discussion

Considering the existing literature [Van de Sande et al., 2008; Siagian and Itti, 2009], the local descriptor *rhotheta* was expected to generally outperform global descriptors. All in all, this hypothesis is confirmed by my results. More particularly, it shows the best overall performance on diagonal trajectories which confirms that local descriptors are best suited for the pose invariance problem. Moreover, its mean average precision (MAP) is higher in most of the experiments. So, from the information retrieval perspective, local vision-based systems indeed do better.

Yet, in several situations, holistic descriptors interestingly obtain better 3WD scores. In fact, all of the tested global descriptors perform better than *rhotheta* at least once. In robot navigation, visual place recognition is generally involved in action selection and robotic control. In this context, we are more interested in the activity of the winners than in the overall responses precision. Also, getting better generalization between topologically close places is very important. Thus, the results suggest that holistic descriptors can be efficient for robot navigation in some situations.

It is worth noting that *gridgab* – which is an implementation of GIST Siagian and Itti [2007] – exhibits the best performances among global descriptors. It is based on the division of feature maps into grid subregions. Thus, it encodes localized information although describing the whole image. This might explain why its results are close to those of the local descriptor. Nevertheless, each of the tested holistic descriptors shows effective in a different situation, making it hard to eliminate any of them based on this study. In particular, they perform better on datasets A and B. This can be due to the fact that the environment is more dynamic in C. It might also be due to characteristics on the off-road context. Or a combination of both. Further investigation is required in order to assess this.

2.3 Visual contexts and global precedence

Humans are able to coarsely recognize a scene at a glance. The *gist*[3] of the scene is extracted from its global aspect based on low-resolution visual information [Navon, 1977; Chun, 2000; Bar, 2004; Oliva and Torralba, 2006]. Such visual context identification resolves ambiguities and facilitates objects recognition – which is, in contrast, based on the processing of higher resolution information [Bar, 2004]. For instance, contextual cueing has been shown to drive spatial attention and increase performances in search tasks [Chun, 2000].

The notion of context is very related to spatial representation as it often refers to background cues Smith [2008]. So context identification is essential for navigation. But, while place recognition *per se* is mainly based on geometric visual information (distances and directions of the landmarks in the environment), it is also influenced by additional sensory (e.g. colors, sounds, odors, etc.) and behavioral cues Anderson and Jeffery [2003]. Thus, the term "context" can carry diverse, more or less abstract meanings in the literature Smith [2008].

In pure localization tasks, context detection can refer to the recognition of a broad area in which several locations (places) can be discriminated. For example, Cuperlier et al. [2008] proposed a model in which a coarse place recognition modulates the responses of a more refined recognition level. This initial work provides proof of the interest of contextualizing place recognition in a simulated navigation task. Nevertheless, the two levels were based on the same geometric cues. On the other hand, according to the global precedence concept, the global aspect of the image can be used to drive the recognition of local details Navon [1977].

To verify whether holistic descriptors can successfully discriminate large regions of the environment, I concatenated the three datasets used in the previous experiment. All places of A, B and C are learned then all *on-trajectory* exploration panoramas are run. This time, instead of evaluating recognition place by place, I considered dataset as units. That is to say, larger places, or visual contexts in a sense. Just like *rhotheta*, global descriptors successfully discriminate each dataset. More precisely, places learned as part of a dataset are best recognized within that dataset. Also, some of them generalize between A and B which present some visual similarities. For the sake of conciseness, Figure 2.3 illustrate maximum activities among places learned in each dataset. It is noteworthy that similar activity profiles are obtained with mean activities within dataset. Put differently, places learned as part of a dataset have an average recognition level that is higher that places learned from the other datasets. This suggests that they can be

---

[3]The term "gist" here is used in its general meaning – that is to say, the summary or the essential aspect of the scene. I use "GIST" (capital letters) when referring the particular implementation in the form of a visual descriptor proposed by Siagian and Itti [2007, 2009].

used as a way to encode visual contexts. Due to their low computational cost, it can be interesting to use them in an earlier recognition stage. Thus, they could drive the recognitions of smaller places in order to reduces the search space of the local vision system Oliva and Torralba [2006].
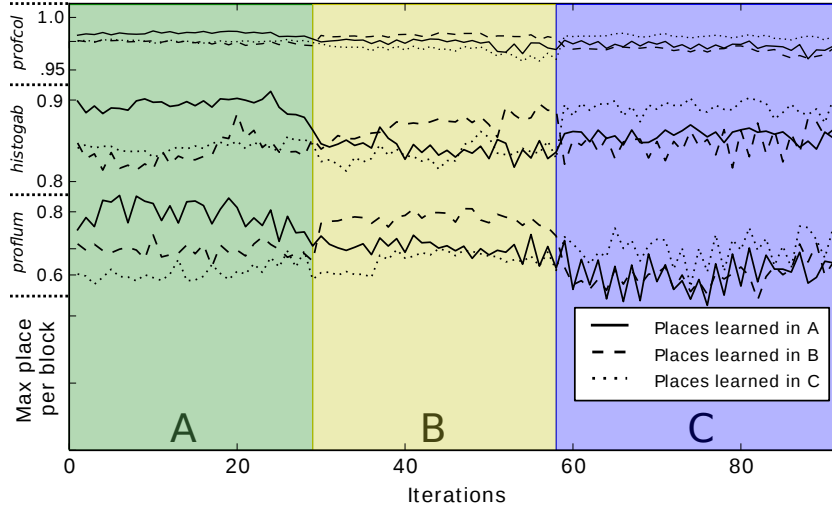
Figure 2.3: Maximum activities among places that were respectively learned in A, B and C are represented. The response of *proflum*, *profcol* and *histogab* descriptors are superimposed. Despite differences in terms of response amplitude and WNR, holistic descriptors correctly discriminate the datasets.

2.4 Proposed model for context-based place recognition

Based on the ideas introduced above, I propose an extension of the LPMP model, illustrated in Figure 2.4. In this model, two identical neural networks are respectively dedicated to global low-resolution and local high-resolution visual processing. The former pathway modulates the activity of the latter. As a result, a coarse localization (*Context*) and a refined one (*Place*) are obtained. The processing chain is based on a biologically plausible model of place cells in the hippocampal system Gaussier et al. [2002] but also integrates modifications that allow for better results in the case of robotic implementations [Giovannangeli et al., 2006].

The place recognition level drives the learning process in a one-shot way. That is to say, whenever none of the learned places activities is greater than a vigilance level v, new categories (signatures, contexts and places) and associations are learned. This way, the system learns independently without human supervision.

On both levels, a position in the environment is encoded as a constellation of neural activities merging 'what' and 'where' information – that is to say signature–azimuth couples. *Gist* and *Landmarks* signatures encode global low-resolution and local high-resolution descriptors respectively.

The activity of each neuron g_i representing a gist signature at time t is given by the following equations:

$$g_i(t) = 1 - \frac{1}{N_{GD}} \sum_{j=1}^{N_{GD}} |w_{ij}^{GD}(t) - d_j^g(t)| \tag{2.2}$$

Figure 2.4: Context-based model for place recognition. Two parallel neural networks extract visual features, merge them with azimuth information and categorize What-Where patterns to obtain two levels of localization respectively based on global and local descriptors. The combination link serves for the contextualization of information in the local vision pathway based on the global vision one. Learning links are only shown for the gist signatures and contexts for the sake of readability; but all the plastic links are modified based on the vigilance signal.

where d_j^g is the j^{th} element of the global descriptor vector of size N_{GD} and w_{ij}^{GD} is the weight of the synaptic link between g_i and d_j^g. The vigilance signal triggers the recruitment of a new neuron, such that the learning rule is the following:

$$\frac{dw_{ij}^{GD}}{dt} = \delta_j^R |d_j^g(t) - w_{ij}^{GD}(t)| \tag{2.3}$$

where R is the index of the newly recruited neuron and δ_a^b the Kronecker delta, equal to 1 if $a = b$ and 0 otherwise. Like in the LPMP model, with this type of neurons, the idea is to save the input pattern in the links weights. Thus, the closer the input is to the learned pattern, the stronger the neurons activations. Although far from modeling single neurons in a realistic way, this method shows good generalization properties for robotic place recognition [Giovannangeli et al., 2006].

Moreover, the recognition of a landmark signature l_i not only depends on the input descriptor but also on the recognition of associated gist signatures. Thus, the activity of each neuron l_i at

time t is given by the following equations:

$$l_i(t) = \frac{1}{N_{cxt}} \sum_{q=1}^{N_{cxt}} \left(\max_{k=1..N_G} {}^q (w_{ik}^G(t).g_k(t)) \times \right.$$
$$\left. \left(1 - \frac{1}{N_{LD}} \sum_{j=1}^{N_{LD}} |w_{ij}^{LD}(t) - d_j^l(t)| \right) \right) \qquad (2.4)$$

where d_j^l is the j^{th} element of the local descriptor vector of size N_{LD}; N_G the size of the gist signatures vector; w_{ij}^{GD} and w_{ik}^G are the weight of the synaptic links between l_i and d_j^g; and g_k respectively and $\max_{i \in I}^q$ is an operator that gives the q^{th} maximum in I.

A crucial feature of this model is that, when implementing this type of neurons, only the N_{cxt} most recognized gist signatures are read. In other words, the activities of landmark neurons that are not associated to them are not computed in order to reduce the computational cost. As I will show in the experimental results, this feature allow the system to maintain a quasi-constant execution time by only processing a limited amount of information at each time step.

The learning rule for w_{ij}^{LD} is the same as w_{ij}^{GD} in (2.3) (replacing g subscripts and superscripts by l). In contrast, the links between gist and landmarks signatures use a hebbian-like learning rule. The purpose is to capture co-activations between contextual information (global aspect of the image) and local visual input (regions of interest):

$$\frac{dw_{ik}^G}{dt} = l_i(t).g_k(t) \qquad (2.5)$$

If the learning process were slow rather than one-shot, the equation could benefit from a decay factor. That was the case in the initial work by Cuperlier et al. [2008]. The learning converged to very few contextual associations performed after very long simulations.

In addition, the model relies on the 'where' information; that is to say the orientation in which the visual feature is observed. The input is a vector in which each neuron has a preferred direction around the yaw axis. The activities of the azimuth neurons α_i are obtained after a lateral diffusion around the neuron coding for the direction of the current visual input. Here, it corresponds to a gaussian bell with standard deviation σ_{azim}. It integrates information about the orientation of the body (robot), head (camera) and fovea (center of the visual feature).

The what–where maps (W–W maps) are second-order tensors MG and ML in which each neuron codes for a signature-azimuth couple. A short term memory stores previous activities while the visual scene exploration is still in progress (i.e. before the end a panorama). In order to reduce the computational cost, the 360^o surrounding field is discretized in N_A orientations before the computation of the tensorial product. Then, the activities of the N_{MG} and N_{ML} W-W tensors are:

$$MG(t) = \max[(g \otimes a), MG(t - dt).(1 - R(t))] \qquad (2.6)$$
$$ML(t) = \max[(l \otimes a), ML(t - dt).(1 - R(t))]$$

where $N_{MG} = N_G \times N_A$ and $N_{ML} = N_L \times N_A$; N_G and N_L are the size of the signatures vector; g, l and a are the signatures and azimuths vectorial representations; R a binary reset signal triggered at the end of a panorama; and $\otimes$ is the tensorial product operator.

Patterns of activities in the W–W maps code for the current location. Such patterns are categorized in context and place vectors, in which c_i and p_i neurons respectively have the following activities at time t:

$$c_i(t) = 1 - \frac{1}{\rho_{MG}.N_{MG}} \sum_{j=1}^{\rho_{MG}.N_{MG}} \max_{i=1..N_{MG}}{}^q |w_{ij}^{MG}(t) - m_j^g(t)| \tag{2.7}$$

$$p_i(t) = \frac{1}{N_{cxt}} \sum_{q=1}^{N_{cxt}} \left(H\left(\max_{k=1..N_C}{}^q (w_{ik}^C(t).c_k(t)) \right) \times \right. \tag{2.8}$$

$$\left. \left(1 - \frac{1}{\rho_{ML}.N_{ML}} \sum_{h=1}^{\rho_{ML}.N_{ML}} \max_{i=1..N_{ML}}{}^h |w_{ij}^{ML}(t) - m_j^l(t)| \right) \right)$$

where m_j^g is the j^{th} element of the tensor MG; m_j^l is the j^{th} element of the tensor ML; w_{ij}^{MG} is the weight of the synaptic link between c_i and m_j^g; w_{ij}^{ML} is the weight between p_i and m_j^l; ρ_M is the proportion of W–W couple required for context and place recognition; $H(x)$ is the heaviside function; and $\max_{i \in I}{}^q$ is an operator that gives the q^{th} maximum in I.

Like the landmark layer, the activities of place neurons that are not associated with the N_{cxt} best recognized contexts are not computed. Also, like in the LPMP model, the factors ρ_{MG} and ρ_{ML} were introduced for the purpose of robustness to occlusions. This parameter also compensates the absence of activation functions in the previous layers of the neural network.

The purpose of binarizing the contextual term through the heaviside function is to ensure that the dynamics of place neurons activities only depends on the input pathway (instead of being reduced by the contextual factor which is < 1). This way, we do not alter the recruitment mechanism and the same vigilance threshold can be used across all the methods considered in the comparison.

2.4.1 Implementation details

In the following experiments, I will use subsampling-based descriptors as inputs for the global vision pathway. Subsampling is a simple and biologically plausible process. Here, the scanline profile technique [Milford and Wyeth, 2008] will be applied on the three visual channels. For luminance and chrominance, *proflum* and *profcol* exhibited satisfying results in Experiment 1. The *profgab* is obtained similarly from the output of a Gabor filter bank used to detect 4 orientations $(0, \pi/4, \pi/2$ and $3\pi/4)$ on a single spatial scale represented by λ.

I refer to the combination of those three global descriptors as *allprof*. The complete holistic descriptor size is $(7.view_W/\kappa)$ where $view_W$ is the views width. It gives a low-resolution representation of the image.

2.5 Experiment 2.A: Different granularity using local and global descriptors

In this experiment, the same datasets as those described earlier are used. However, unlike Experiment 1 and as explained in the previous section, the system automatically chooses when to

trigger learning depending on the vigilance parameter and the place recognition levels. So we have to provide a larger number of learning panoramas.

The three A, B and C are concatenated. The so obtained dataset includes $74 = 23 + 25 + 26$ learning panoramas (i.e. sets of 15 images captured while robot is stopped) and $92 = 29 + 29 + 34$ exploration panoramas (i.e. sets of 7 images captured while moving). The closest distance between two possible places is $d_{learn} = 0.93 \pm 0.03$ meters in average; and as in Experiment 1, an exploration panorama is completed after traveling $d_{explo} = 0.71 \pm 0.01$ meters. We calculate the ground truth by estimating the robot position at every step based on that.

Here, I consider only one pathway at a time in order to test *rhotheta*, *proflum*, *profcol*, *profgab* and *allprof* separately. The objective is to compare the descriptors capacity to discriminate places in terms of granularity. A new place is learned whenever none of recognition level of the known places is higher than a vigilance threshold v. Consequently, the number of recruited place cells is a good metric for our test: the more places are learned, the small the generalization.

I run several tests for $v \in \{0.6, 0.7, 0.8, 0.9, 0.98\}$. There are 74 learning panoramas that we feed to the visual processing chain one by one. One can note that after a place is learned, an intermediate panorama has to be used to test the recognition. Thus, the system can learn 34 place at the most on our dataset.

2.5.1 Results

As shown if Figure 2.5, global descriptors recruit less place cells than the local descriptor *rhotheta*. Since *allprof* is simply a concatenation *proflum*, *profcol* and *profgab* without normalization, its results not only depend of the variation observed on each individual descriptor but also on their sizes. However, for $v = 0.8$, it exhibits an averaged response.

The difference in terms of discrimination granularity can be seen for almost all of the vigilance values we considered. Here we are mainly interested in studying the learning process. Yet, it is worth pointing out that limit v values do not allow for a satisfying recognition during exploration.
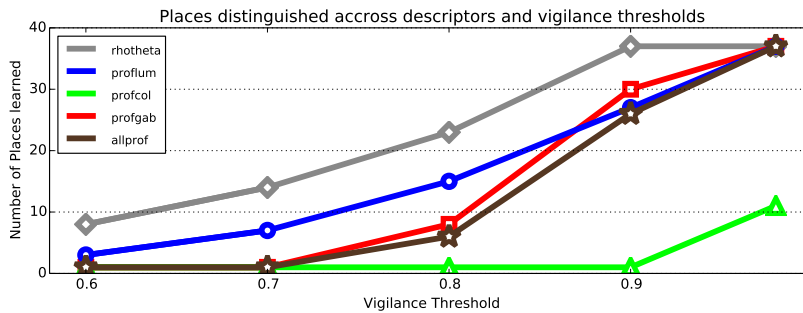
Figure 2.5: Experiment 2.A, Number of discriminated places by local and global descriptors for several values of the vigilance threshold. In general, global descriptors recruit less place cells than *rhotheta*

2.5.2 Discussion

As expected, global descriptors inherently discriminate larger place fields than the local descriptors since they are less sensitive to visual variations. Indeed, more place cells are learned using

rhotheta than all the other descriptors. Holistic descriptors tend to capture more global information; hence the coarse granularity. This also confirms the interpretation of the results presented in Figure 2.3. Therefore, the proposed model for context-based place recognition should benefit from the fast recognition of the spatial context to drive the finer localization.

2.6 Experiment 2.B: Visual context-based place recognition

The issue of combining local and global visual information has been addressed in related work. The existing models can be classified in two families. On the one hand, the two levels of visual descriptors are concatenated and fed simultaneously to the scene classifier [Lisin et al., 2005; Rostami et al., 2014]. On the other hand, the global features are used for a first level of recognition (i.e. bigger regions of the navigation environment), which is then refined by the local information (i.e. more precise location) [Siagian and Itti, 2009; Milford et al., 2014]. Like Lisin et al. [2005], I will refer to the first class as *stacking* methods and the second one as *hierarchical* methods.

The model I propose in this chapter belongs to the second class. The purpose of this experiment is to evaluate its performance and computational. In this model, the global vision pathway modulates the activities in the local vision pathway at two levels: gist signatures serve as contexts for landmarks and coarse localization as contexts for places. In order to assess the role of the contextualization in each of these layers, I also implemented two architectures where it is only done at the landmarks level or at the places level. Those three versions will be labeled **Cxt_LP**, **Cxt_L** and **Cxt_P** respectively. Cxt stands for context, L for landmarks and P for Places.

Additionally, I compare the context-based models to the non-hierarchical, stacking method. First, I test the case where all descriptors are concatenated as an input for the landmarks categorization (labelled **Stack_L**). Second, I consider the case where the gist signatures are processed separately (without modulation of the landmarks activities) but the two W-W maps are put together for place recognition (labelled **Stack_P**). In other words, the *combination links* in Fig. 2.4 are replaced by a simple concatenation operation and the contextualization terms in equations (2.2) and (2.7) are omitted. Last, the **LPMP** model (using the local vision pathway only) is used as a baseline.

Given the results obtained in the previous test, we set the vigilance threshold to $v = 0.8$.

The same measures as in Experiment 1 are used to evaluation the performance of the above mentioned architectures (3WD, MAP and WNR). But the evaluation of the computational cost is also necessary. Indeed, more visual input to process (global+local) could make the architecture too slow to run on a real robot with real time constraints. Therefore, I also measure the framerate, i.e. the number of images processed per second. These experiments are run on a 6-core 12-thread 3.33GHz CPU with 16 GB of RAM.

2.6.1 Results

The performance scores of the six tested models are presented in Figure 2.6. First, all models combining local and global visual information outperform the LPMP model in the 3WD test. The best results are obtained by **Cxt_LP** then **Cxt_P**. Besides, context-based model get higher MAP score than stacking versions. We note that for this test, LPMP performs as well as context-

based models. As for the WNR, we observe a real impact of the contextualization – **Cxt_LP** obtaining the highest score.
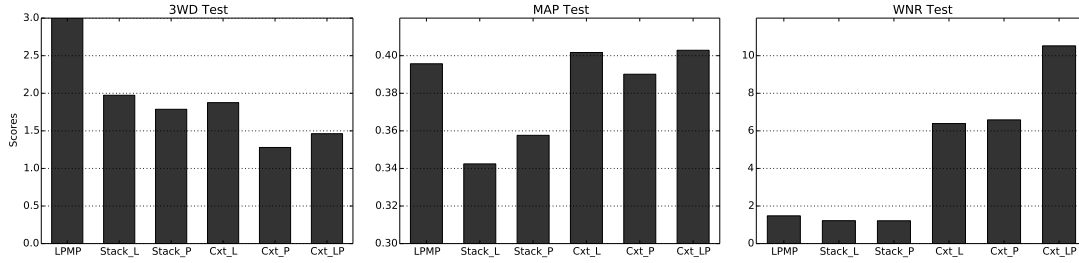
Figure 2.6: Experiment 2.B, Performance scores of all tested models. Our context-based model **Cxt_LP** gets the best results in almost all the tests or is second to one of its variants **Cxt_P**. 3WD: The closer to 1 the better; MAP and WNR: The great the better.

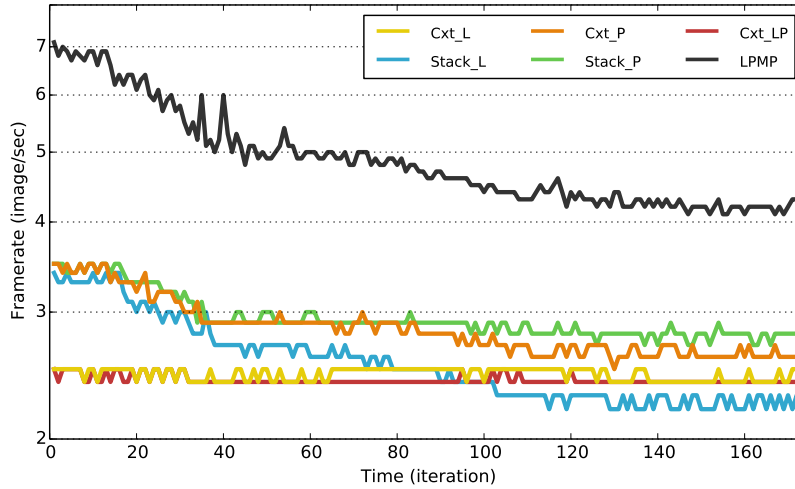On the other hand, adding the global vision pathway induces an additional computational cost, as shown in Figure 2.7. Except for LPMP, the highest framerate is obtained with **Stack_P**. However, beside the absolute value, we note that the framerate stays almost constant in the case of **Cxt_LP** and **Cxt_L** while it decreases over time for all the other models.

Figure 2.7: Experiment 2.B, Computational costs of all tested models, reflected by the framerate. For the sake of readability, the y axis is represented in a logarithmic scale. Beside the absolute value, we note that the framerates decrease over time for most of the models except **Cxt_LP** and **Cxt_L** which maintain a rather constant computational cost.

2.6.2 Discussion

Stack_L and **Stack_P** were implemented in order to compare the context-based model to non-hierarchical methods. Interestingly, as opposed to my results, Lisin et al. [2005] obtained better performances with the stacking method than with the hierarchical one. However, their experiments were based on images of multicellular organisms (planktons) on a uniform background.

In contrast, we use images recorded in outdoor, dynamic environments. This highlights the specificity of robotic navigation, as compared to other application domains of scene recognition and computer vision.

The fact that all models using local and global visual input outperform the LPMP model in the 3WD test confirms the interests of using global descriptors for place recognition in a robotic system. Indeed, in this application domain, we want topologically close locations to have close recognition levels. However, the stacking models get lower MAP scores than the others which means that the place cells following the 3 winners often correspond to wrong recognition. Thus, from the information retrieval point of view, these methods are less efficient than the hierarchical ones. Additionally, as expected, the context-based method proposed in this paper considerably increases the WNR by penalizing landmarks that are recognized independently from the global scene to which they are associated as well as place cells that are active outside of their contexts.

Also expected was the computational cost induced by the additional processing pathway dedicated to the global vision. Although the global descriptors are relatively compact and their computation not costly, their processing across the categorization layers of the neural network reduces the system framerate. Nevertheless, the results demonstrate this is not a critical draw-back. A framerate of 2.5 images/sec means it takes 3 sec to capture a 7-image exploration panorama and update the robot localization. Since the linear speed is generally set lower than 1 m/sec in field experiments, it is considered an acceptable framerate for a moving camera dedicated to place recognition. Additional faster sensors can be used for more critical functions like obstacle avoidance. But more importantly, beside the absolute value, one can note that the framerates decrease over time for most of the models. Indeed, given that there are $N_{PoI} \times 15$ landmarks recruited per place cell, the larger the environment (or dataset) the more information has to be processed at the end of the experiment. However, the framerate stays almost constant in the case of **Cxt_LP** and **Cxt_L**. This is due to an interesting feature of our model: in the local vision pathway, only the activities of neurons associated to a limited set of well recognized ones (N_{cxt}) from the global vision pathway are computed. This is not observed with **Cxt_P** because there are not enough places learned in this experiment so that the sole contextualization of places exhibits this effect. Nevertheless, we except that on longer trajectories, the benefit of the context-based model persists while the framerate of other techniques keeps dropping up.

2.7 On the broader notion of context

While the concept of visual context exploited in my work mainly originates from research on human vision, there is also significant evidence of the influence of contextual information on place recognition in animals. For instance, like in my work, color stimuli modify the firing patterns of place cells in rats [Bostock et al., 1991]. It is also worth noting that visual textures – e.g. solid color vs. wire mesh design [Eacott and Norman, 2004] – have also been used to manipulate context information.

But the notion of context encompasses various types of information including sensory, behavioral and interoceptive cues [Bouton et al., 2006; Smith, 2008]. Indeed, Anderson and Jeffery [2003] showed that place cell activities are modulated by combinations of background colors and odors (multisensory cues). Place cell firing may also vary depending on the starting point [Mizumori et al., 2007] or the path [Wood et al., 2000] of the rat inside the maze. Moreover, it has been suggested that ingested drugs might generate different interoceptive contexts [Bouton et al.,

1990]. Interestingly, Hayman and Jeffery [2008] formulated the Contextual Gating Hypothesis in which contextual information (non-geometric) modulates the input of place recognition similarly. Although this idea was proposed to explain how similar grid cell activation patterns results in distinct place recognition, it is consistent with the context-based model I propose.

Since "context" gathers such a variety of information, it is often associated with episodic memory. Indeed, the memory of autobiographical events relies on both external and internal cues. In that matter, two brain areas are implicated: the hippocampus (HPS) and retrosplenial cortex (RSC) [Mizumori et al., 2007; Smith, 2008; Miller et al., 2014]. Interactions between these two regions are very likely to be the key to the representation of episodic memory. In particular, RSC is thought to support the integration of self-motion and visual information. For instance, Alexander and Nitz [2015] show that it encodes conjunctions of left-right turning, position in a route and position in the larger environment. Mizumori et al. [2007] indicate that one the role of RSC might be to provide HPS with contextual, integrated sensory and movement-related information.

In the processing of visual contexts, RSC activities are correlated with gist-based recognition [Bar, 2004]. Besides, spatially localized firing is observed in this area, representing large and noisy place fields [Smith et al., 2012]. Similar features are exhibited by the proposed model. Yet, this model can still be extended to a broader view of contexts, as conjunctions of sensory, behavior, and even emotional states.

2.8 General discussion

The context-based neuronal architecture appears promising both from the performance and computational cost perspectives. But further evaluation is required in order to validate the interest of the model in outdoor experiments. Indeed, the results I presented were based on off-line tests. This made it possible to rigorously compare various architectures and provide performance measures. On the other hand, there are certain obvious shortcomings: i) the datasets covered rather short distances, ii) all trajectories were linear, iii) there was no real interaction with the environment. Although the datasets were recorded in way that reproduced as much as possible on-line visual exploration, this passive evaluation was only intended as a preliminary test. Experiments on real robot should be performed in order to address the above mentioned weaknesses.

Additionally, as I pointed out, a broader notion of context encompasses a variety of information; including behavioral cues. The model I proposed could be extended by combining visual contexts with movement-related information for example. This might also be a step toward the representation of memory episodes. Nevertheless, such a model can only be meaningful if the system is situated and actually interacts with environment.

For future works, it also seems important to investigate the vigilance mechanisms handling category recruitments in the local and global pathways. For the sake of simplicity, the same vigilance threshold was used in these experiments. Yet, the results of Experiment 2.A raise questions about this choice. Indeed, different descriptors have different variation intervals and might benefit from an adapted recruitment mechanism. Additionally, to fully benefit from the difference in terms of granularity of discrimination and generalization between local and global features, it seems essential to handle recruitments independently between the two pathways.

Personal publications

Belkaid, M., Cuperlier, N., and Gaussier, P. (2016a). Combining local and global visual information in context-based neurorobotic navigation. In *Proceedings of the IEEE International Joint Conference on Neural Networks*, pages 16622:1–16622:8.

Table 2.5: Experiment 1 and 2: Parameter values

	Value	Description
$view_W$	440	Width of the subimages (views)
$view_H$	240	Height of the subimages (views)
σ_{azim}	30	Std. dev. of azimuths diffusion (in degrees)
N_a	5	Nb. of orientations in the W-W Map
$\rho_M G$	0.33	Ratio of required W-W couples (places)
$\rho_M L$	0.33	Ratio of required W-W couples (contexts)
Local descriptor		
α	0.4	Gradient resolution (edge detection)
σ_{DoG_1}	8	Std. dev. of 1^{st} DoG gaussian (in pixels)
σ_{DoG_2}	2	Std. dev. of 2^{nd} DoG gaussian (in pixels)
N_{PoI}	5	Nb. of PoI extracted per image
r_{small}	10	Small disk radius in local views (in pixels)
r_{big}	64	Big disk radius in local views (in pixels)
$N_{\rho\theta}$	54	Size of the rhotheta descriptor
Global descriptor (Exp 1)		
κ_l	4	Subsampling factor (grey scale image)
κ_c	4	Subsampling factor (*a* and *b* images)
$N_{H_{col}}$	4	Nb. of bins in the color histogram
$N_{H_{gab}}$	4	Nb. of bins in the gabor histogram
γ	0.7	Spatial aspect ratio of the gabor
ψ	0	Phase offset of the gabor filter
σ_g	8	Std. dev. of the gabor gaussian (in pixels)
$\lambda_1,\lambda_2,\lambda_3$	2, 8, 32	Wavelengths of gabor filters (in pixels)
N_{noise}	64	White noise baseline vector size
Global descriptor (Exp 2)		
κ	4	Subsampling factor (all global desc.)
γ	0.7	Spatial aspect ratio of the gabor
ψ	0	Phase offset of the gabor filter
σ_g	16	Std. dev. of the gabor gaussian (in pixels)
λ	32	Wavelength of gabor filters (in pixels)
N_{cxt}	7	Nb. of best recognized contexts

Sensorimotor navigation through human–robot interactions

Future robots are supposed to become our partners and share the environments where we live in our daily life. This is confirmed by several emerging applications for human–robot interactions such as companion robots helping elders and impaired people, robots for pedagogical tools, monitoring robots, and guiding robots (museums, hotels). In his plenary talk at the *International Conference on Intelligent Robots and Systems* in 2015, Oussama Khatib organized the following sub-fields in robotics as increasingly challenging: 1) industrial robotics, 2) mobile robotics, 3) social robotics, and 4) cognitive robotics.

As robots become part of our daily life, their acceptability becomes a central issue [Goodrich and Schultz, 2007]. One the major difficulty in this field of research is to make interactions with robots as efficient as pleasant for the human user. Indeed, the design of social robots must take into account the fact that these systems have to interact with "non-experts". More specifically, it has to allow the latter to efficiently and intuitively make robots learn the behavior they expect from them. The two previous chapters introduced an architecture for place recognition which can be used in sensorimotor navigation behaviors. Of particular interest in this chapter is the learning of such behaviors through human–robot interactions. Thus, I will present an architecture aimed at the interactive learning of navigation task. This work is inspired by findings on low-level human interactions and dynamical motor control. More specifically, the focus will be put on rhythmical interactions and interpersonal synchrony as catalysts for learning.

3.1 Interactive learning of navigation behaviors

Based on the PerAc principle, we can couple the codes representing place cells with directions of movements. This place–action strategy provides the basis of sensorimotor navigation behaviors. Indeed, as illustrated in Figure 3.1, three or more place–action associations that are directed toward a goal location build an attraction basin guiding the robot. Likewise, the same principle can be applied for trajectories like a round path.

These sensorimotor associations can be made during an autonomous exploration of the environment and reinforced when they form a path to a goal [Cuperlier et al., 2007; Hanoune et al., 2015]. The PerAc framework also allows for teaching navigation tasks via human–robot inter-
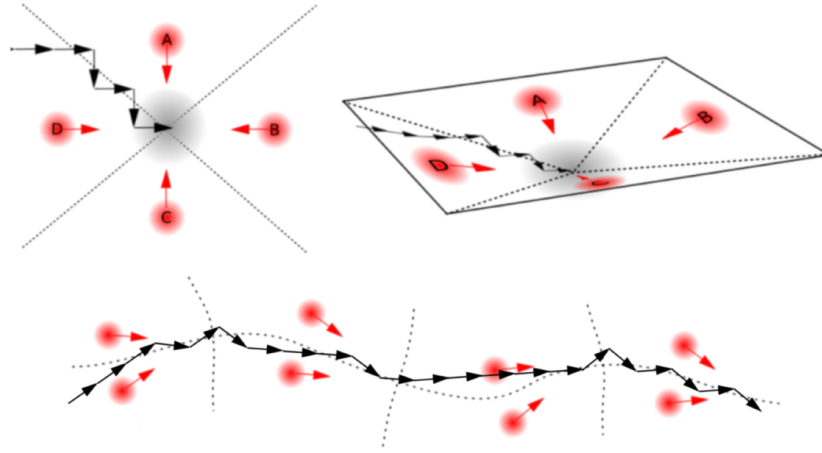
Figure 3.1: Place–action strategy in homing and round path learning. **TOP-LEFT**: Four associations directed toward the goal location guide the sensorimotor trajectory of a homing behavior. **TOP-RIGHT**: A schematic illustration of the attraction basin built with the place–action couples. **BOTTOM**: A portion of a round path trajectory learned by "converging" place–action associations. The figure is slightly readapted from [Jauffret, 2014].

actions. In the literature, it is becoming more and more admitted that interactive learning helps improve the system autonomy, performance and ability to evolve in a complex world [Goodrich and Schultz, 2007]. Accordingly, there are several examples of the use of such techniques in the literature, both in pure machine learning systems (e.g. [Fails and Olsen, 2003]) and in robotics (e.g. [Nicolescu and Mataric, 2006; Saunders et al., 2006; Giovannangeli and Gaussier, 2010; Knox et al., 2013]). Some of these works rely on the notions of teaching and imitating [Nicolescu and Mataric, 2006; Saunders et al., 2006; Giovannangeli and Gaussier, 2010].

Giovannangeli and Gaussier [2010] studied this in the context of an interactive learning of round paths. Of particular interest in this work was the comparison between *prescriptive* and *proscriptive* teaching. In the former, the human teaching attempts to show the perfect behavior by guiding the robot along the desired trajectory. In the latter, he only corrects the robot by indicating the direction toward the desired path when it gets too far from it. The authors advocate a trade-off (or alternation) between these two strategies. By always showing the right path in a prescriptive way, the teacher does not get any feedback on the quality of the learned trajectory. Also, as shown in Figure 3.1, the attraction basin is more efficiently obtained with "converging" place–action associations. On the other hand, the proscriptive teaching enables a non-verbal communication mediated by each one's actions: human guidance and robot behavior as a feedback on performance. Such interactions allow more efficient learning but the obtained attraction basin may be too wide and lead to a significant variability in the robot trajectories. Hence, is it argued that alternating between prescription and proscription provides the best results.

The work by Giovannangeli and Gaussier [2010] also highlights the importance of a situated, embodied interaction in this task learning. In their experiments, the robot is embedded with a joystick representing an artificial neck. Similarly to a stereotypical human/dog training, the teacher pulls a leash attached to the joystick in order to show the desired direction. As a communication medium, the leash has two functions: 1) it draws the robot attention to initiate a learning step, 2) it provides information on the desired behavior. In the rest of this chapter,

I will show another type of human–robot interaction that can be used to teach navigation task only based on visual communication.

3.2   Learning place–action navigation through synchrony-based partner selection and visual tracking

The work I will present here is a collaboration with Caroline Lesueur-Grand, another doctoral student in my laboratory who's research focuses on intuitive human–robot interactions. In this experiment, we take a radical simplification route taking advantages of recent discoveries in low-level human interactions and dynamical motor control. Our objective is thus to propose a bio-inspired model for human–robot interactions that takes into account the human sensitivity to synchronous interactions. We believe that exploiting the natural stability and adaptability properties of rhythmic activities can solve several acceptability problems. Also, according to studies on interpersonal coordination, bi-directionality is a fundamental factor in intuitive interactions. Therefore, the robot must not only have abilities to adapt its dynamics to that of the human partners, but also to "entrain" them unintentionally (i.e. without cognitive load) by giving a feedback (stimuli) representing its own dynamics (head movement, arm movement, auditive stimuli etc.). In the context of navigation task learning, we propose that adult–infant-like interaction allow a mobile robot to 1) initiate an interaction with a human partner, 2) maintain the interaction by learning to recognize the partner and to focus the visual attention on him, and 3) learn to navigate in a "tolerated" area shown by the selected partner (through the interaction).

As robots become part of our daily life, the acceptability of human–robot interactions becomes a central issue [Goodrich and Schultz, 2007]. One of the major difficulty in this field of research is to make these interactions as efficient as pleasant for the human user. Notably, most of the existing approaches have neglected the importance of understanding the dynamics of the interactions. In particular, studies on developmental psychology acknowledged synchrony as a prime requirement for interaction between a mother and her infant. In fact, infants stop interacting with their mother when she stops synchronizing with them Nadel et al. [1999]. Synchrony detection mechanism in young infants plays a pervasive role in learning and cognitive development (word learning, self-awareness and control, learning related to self etc.Gogate and Bahrick [1998]Rochat and Striano [2000]). Moreover, it is worth noting that studies on interpersonal motor coordination point out *unintentional* synchronizations among people. Issartel et al. [2007] studied interpersonal motor-coordination between two participants when they were instructed not to coordinate their movements. The results showed that participants could not avoid unintentional coordination with each other. This reflects that when visual information is shared between two people in an interpersonal situation, they coordinate (unintentionally) with each other.

Synchrony in social interactions has also been studied and used in robotics. Prepin and Gaussier [2010] and Andry et al. [2011] used synchrony as an internal reinforcement signal for learning. Blanchard and Canamero [2005] proposed a velocity detection system to synchronize the movements of two robots to improve the reactivity of agents to changes in their environment. Moreover, Marin et al. [2009] showed that motor resonance between robots (humanoid) and humans could optimize the social competence of human-robot interactions. Additionally, Hasnain et al. [2012] and Grand et al. [2014] proposed a synchrony-based neural network architecture

capable of selecting a partner and locating its focus of attention in order to initiate and then to sustain the interaction.

3.3 Experiment 1: Learning a navigation area

In this experiment, we use a customized mobile robot shown in Figure 3.2 based on the Robulab platform by Robosoft. This platform is embedded with a small added arm controlled by an oscillator representing its interacting dynamic. Two panning cameras are also used in this experiment. One for Landmarks recognition (simulating eye saccades during navigation) and the other one for optical flow detection, partner recognition and visual tracking (simulating head motion during interactions). They respectively have rotation fields of approximately 300 degrees (the whole panorama minus the pole position) and 200 degrees. The robot additionally has a laser sensor for the detection of obstacles and human presence.
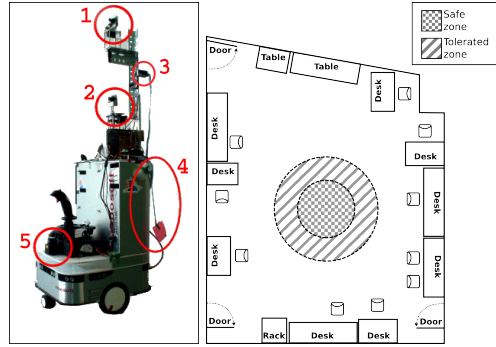

Figure 3.2: Experience 1, setup. **LEFT**: Robotic platform: Two panning cameras – respectively simulating head motion during interactions (1) and eye saccades during navigation (2) –, an oscillating arm (3) and a laser sensor (4). **RIGHT**: Experimental environment: in a 8 m long by 6.5 m wide lab room, the tolerated navigation zone (stripes) is a 3.40 m diameter area around a safe zone (checkerboard) in the center of the room.

The typical scenario of our application, summarized in Figure 3.3, is as follows :

- The robot moves autonomously according to its own dynamics (physically expressed by the rhythmic oscillations of its arm) and visually scans the environment with its first panning camera,

- If a human stands in front of it and starts moving, the induced optical flow influences the robot dynamics and consequently the arm oscillations; the human partner can thus be rhythmically and (possibly) unintentionnally "entrained" while the robot arm tries to synchronize with the perceived rhythm.

- If this mutual (bi-directional) influence converge to common motion dynamics, synchrony is detected, the robot stops scanning the environment and focuses its visual attention on synchronous regions permitting to learn the partner shape,

- When the selected partner starts moving, the robot's tracks and follows the human using the second camera; thereby the partner shows the desired direction for the current robot location (place)

- If the difference between the robot body and head orientations (resp. current and desired orientations) is sufficient, a new place–action learning is triggered

- After the human intervention, whether a new place is learned or not, the robot starts scanning the environment again to perform place–action navigation.
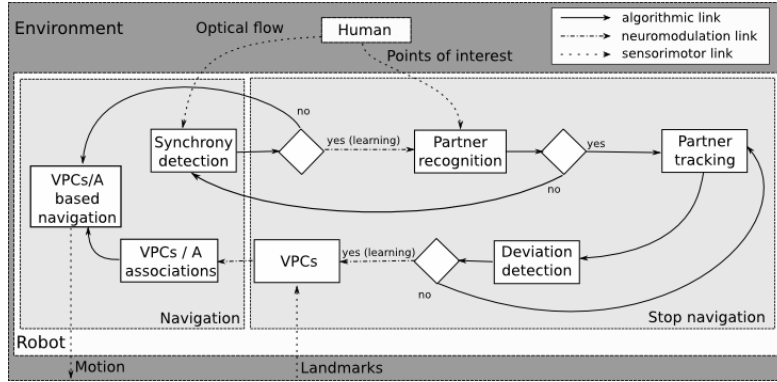
Figure 3.3: Synoptic schema illustrating the human/robot interaction dynamics. In order to learn a navigation task, we establish a bidirectional action-based interaction. When the human imitates the robot arms motion, synchrony detection allows for partner selection. Deviation detection according to the orientation shown by the selected partner triggers the sensorimotor learning of associations between visual places and directions (VPC). The robot interacts with the environment by performing place–action navigation.

We define a tolerated navigation zone with a diameter of 3.40 m in the center of the experimental room shown in Figure 3.2. Taking into account the linear and rotational speed, the safe zone where the robot can freely move is a 1.80 m diameter concentric area. The 0.80 m-wide ring between these two circles determines the area where a human intervention is required in order to correctly build the attraction basin inside the tolerated zone.

Interaction (re-)engagement

Interactions are initiated through synchrony-based partner selection. Initially, the robot arm oscillates at its own frequency and amplitude illustrating the robot internal dynamics. The neuronal model is based on a pair of neurons o_1 and o_2 inducing the oscillations [Revel and Andry, 2009]:

$$o_1(t+1) = o_1(t) - \beta o_2(t) + \alpha_1 \qquad (3.1)$$
$$o_2(t+1) = o_2(t) - \beta o_1(t) + \alpha_2 \qquad (3.2)$$

where α_1, α_2 and β determine the oscillation frequency. When a person interacts with the robot by moving her arm, the motion in the visual field is estimated by an optical flow algorithm. As proposed by Hasnain et al. [2012] and Grand et al. [2014], the so-deduced motion energy is used to entrain and adapt the robot's arm oscillations to the human movement dynamics. Thus, the equation (3.1) is replaced by the following:

$$o_1(t+1) = o_1(t) - \beta o_2(t) + \alpha_1 + of(t) \qquad (3.3)$$

where of is the energy induced by the optical flow.

Synchrony detection then consists in extracting the maxima from the human and the robot oscillators and calculating their correlation so that synchronized interactants can be selected as partners. Additionally, in order to facilitate future interaction initiations, the robot learns to visually recognize those partners. To do so, we use a neural network similar to that of the LPMP model I described earlier. First, points of interest are detected on the image and local views around them are encoded by a log-polar transformation. In the dedicated neural network, a first layer categorizes those local views. Then, a short term memory allows a second layer to captures a set of local views in order to recognize a partner. Here, only the 'what' information (i.e. local view identification) is used, not the 'where' information (i.e. local view azimuths).

Place–action learning

Teaching navigation behaviors using the place–action strategy consists in associating a direction of movement to the current location. Therefore, attention is focused on recognized partners through visual tracking so that they can show the correction orientation. Because points of interest stability cannot be guaranteed when the partner starts moving, visual tracking relies on a low-resolution shape recognition. More precisely, the original 320×240 image is subsampled at a factor of 2 and the central 65×45 subimage is captured while facing the human partner. Then, the tracking method consists in finding the maximum of correlation in a 11 pixels neighborhood around the center of the image. We use a 1D correlation to reduce computational cost Ancona and Poggio [1995].

At the beginning of the experiment, the robot is put inside the safe zone. Before any learning, it navigates randomly due to noise activities in place cells. Whenever it gets close to the limits of the area, the human experimenter intervenes to correct its trajectory. As explained earlier, if the robot does not recognize the human as a partner, synchrony detection is necessary to engage the interaction. In case of a known partner, the robot visually tracks the person in front of it. If a sufficient deviation (20 degrees in this experiment) is detected between its current direction and the one that is suggested by the partner, the robot visually triggers the place–action learning. Place cells are learned using the LPMP model and sensorimotor coupling is made by means of the Widrow and Hoff [1960] rule.

We note that the initial condition is not determining for the resulting behavior. If the robot starts the experiment outside the tolerated zone, only a couple of additional place–action associations is required to guide it toward the safe zone.

3.3.1 Results

For illustration purpose, I shall first describe the human–robot interactions obtained with the architecture I described above. Figure 3.4 shows images from phases of interaction engagement by two partners and results regarding partner selection and visual tracking. Here, the robot is static while three different people try to engage an interaction. Two of them had already synchonized their arm movements with the robot but the third one did not. Consequently, the robot only learned the shapes of two synchronized partners. Therefore, visual tracking is performed only when one of them is recognized (recognition level has to be higher than a threshold).

As for navigation area learning, Figure 3.5 shows the results of an experiment that lasted 14 minutes. Some neural activities are represented more than once for clarity purpose. If so, they
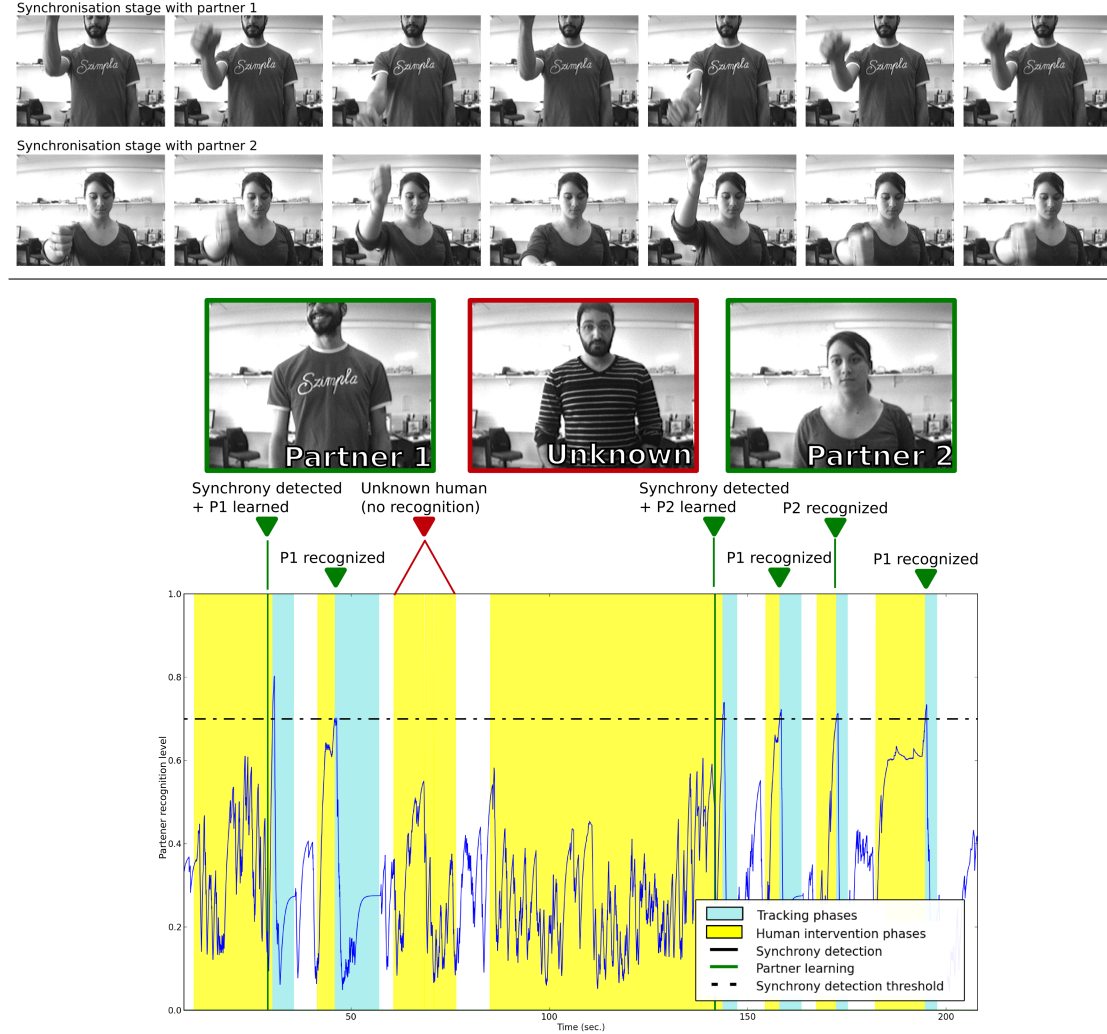
Figure 3.4: TOP: Interaction initiation with robot. Two partners move their arms in front of the camera. BOTTOM: Sequence of human/robot interactions with three partners. Two selected partners are recognized and visual tracking is launched after each of their interventions. However, a third person (unknown) is not able to engage the interaction without the synchronization phase.

have the same color in (A), (B), (C), or (D).

(A) shows the evolution of synchrony detection during human intervention. When synchrony detection signal is higher than a threshold (dotted black curve) and a human (or object) presence is detected (based on laser sensor) the partner learning is launched (green curve). There is only one exception after $t = 500$. Partner learning is not launched because no human (or object) presence is detected.

(B) shows the partner recognition (blue curve) level during the experiment. When it is higher than a recognition threshold, it allows for the partner tracking. The reader can see that when the robot camera is moving (tracking), the recognition level decreases due to the sensitivity of points of interest to small variations. This highlights the interest of a low-resolution shape tracking.

(C) shows the robot head orientation (black curve). In some cases, the partner verifies if the
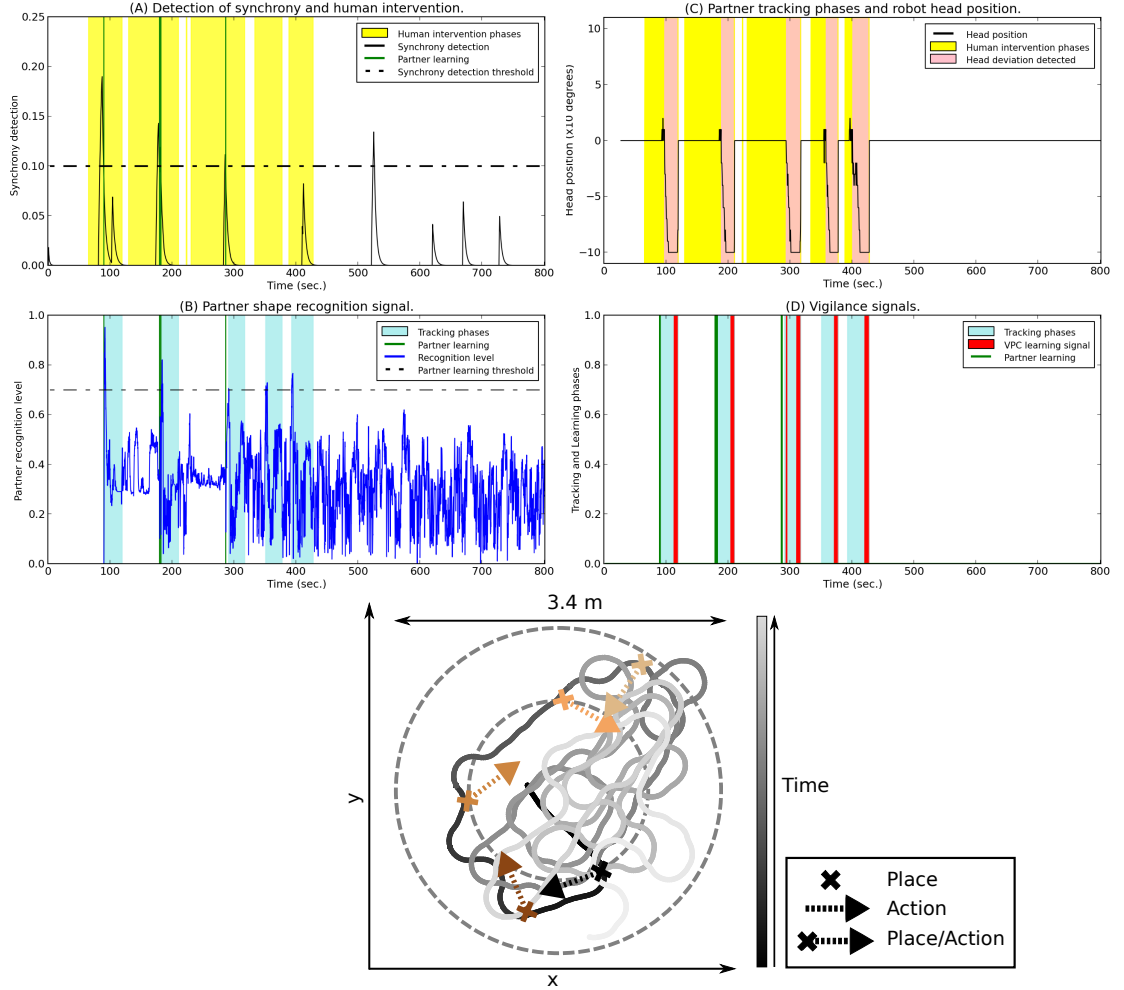
Figure 3.5: Results of the experiment on interactive learning of place–action navigation. **TOP**: Neural activities. The same signals have the same color in (A), (B), (C), and (D). (A) During human intervention, synchrony detection activities higher than threshold phases trigger partner learning. (B) When the parter is recognized the tracking phase is launched. (C) When the robot head orientation changes a deviation is detected. (D) Tracking phase induces a change of orientation which triggers place/action learning. **BOTTOM**: Robot trajectory and spots where place/action learning occurred. The gray gradient illustrates the time course. The robot learned to stay inside the 170 cm radius tolerated zone.

robot is really tracking him. This can be observed when the head is turned in the opposite direction at the beginning of some tracking phases. The deviation is calculated from the difference between the orientation of the robot head and body.

(D) shows the place/action learning signals following each tracking phases that induce a sufficient deviation from robot initial orientation.

Figure 3.5 also illustrates the robot trajectory and the spots where the learning occurred (bottom part). Only a few place–action couples are required for this task.

3.3.2 Discussion

The overall obtained behavior is satisfying: after only a few human interventions, the robot was able to learn the navigation area. However, one can notice that the last place–action association is learned too far from the central zone. This causes the robot to navigate very close and even beyond the limits of the desired area. Such a results highlights the drawbacks of a purely proscriptive teaching strategy. Indeed, the experimenter waited until the robot clearly got away from the navigation area before correcting it. The results is an attraction basin that is wider than the optimal one. A more efficient teaching would have consisted in an earlier, prescriptive intervention showing the way back to the navigation area.

3.4 General discussion

The proposed method advocates the consideration of the dynamics of interaction in the development of companion robots that will share our environment. Indeed, synchrony and rhythmic adaptation can serve as catalyst for sensorimotor learning. More specifically, an example of the use of a previously developed partner selection model [Hasnain et al., 2012; Grand et al., 2014] in the context of a navigation task is given. Combining synchrony-based partner selection and visual tracking allowed us to extend previous works on interactive learning [Giovannangeli and Gaussier, 2010] to purely visuomotor-based human–robot communication. A major advantage of this approach is to allow non-experts to efficiently and intuitively make robots learn the behavior they expect from them. In fact, initiating, re-engaging and maintaining the interaction using a simple low-level visuomotor mutual rhythmic adaptation (optical flow) can be achieved by common and untrained people through different ways as arm movements, hand/head/object shaking facilitating intuitive non-verbal communication without any additional costly or invasive device requiring training.

Several examples of interactive learning architectures can be found in the literature. For instance, Fails and Olsen [2003] and Knox et al. [2013] get interesting results in terms of learning performance. But the interaction situatedness is questionable since the teacher does not directly communicate with the system. In contrast, Saunders et al. [2006] rely on an ecological learning inspired by studies about primates social behavior. Yet, this work separates the learning mode from the execution one. On the other hand, Nicolescu and Mataric [2006] implement action-based interactions between the human and the mobile robot and interesting results are shown regarding the learning of complex high-level tasks. This work also highlights the importance of situated interactions where task learning occurs during the execution. Thereby, the teacher can directly observe the effect of his/her interventions and the robot can more easily communicate its inners states and demonstrate the acquired competences.

The work I presented in this chapter is rather exploratory and is intended to serve as a proof of concept. Learning navigation areas can be useful in service robotics (e.g. cleaning, air purification) or surveillance. But, as seen above, the place–action action strategy enables the learning of specific trajectories like round paths for instance. In fact, part of Caroline's research (with whom I collaborating in this work) is to study dynamical gait synchronization. The robot should be able to detect the human rhythm on his legs motion during locomotion. New learnings would be possible without stopping the robot. For example, we could use this method to learn a round navigation task. The partner could show the trajectory by walking side by side with the robot.

As far as my thesis is concerned, this work also raises important questions regarding the robot's evaluation of its own behavior and skills. For example, in this experiment, a basic threshold prevents from learning place–action associations that do not have a strong impact on the robot trajectory. This rudimentary novelty detector allows the robot to be slightly more "selective" with regards to the human teaching. But it would be far more interesting to use adaptive mechanisms that take the dynamics of the robot behavior into account. Self-assessment could allow to determine situation where the robot competences do not match the task demand and where learning is required. Thereby, the robot could communicate the need for human intervention. These questions will be addressed in Chapter 7.

Personal publications

Belkaid, M., Lesueur-Grand, C., Mostafaoui, G., Cuperlier, N., and Gaussier, P. (2014). When and from who to learn : Synchrony as a way to learn sensorimotor associations. In *Symposium on Biology of Decision Making*.

Belkaid, M., Lesueur-Grand, C., Mostafaoui, G., Cuperlier, N., and Gaussier, P. (2016c). Learning sensorimotor navigation using synchrony-based partner selection. In *Proceedings of the ACM Artificial Intelligence and Robotics*,19:1–19:5.

Part II

Emotion–Cognition interplay

Biological and Artificial Emotions: A cross-disciplinary research

"Over the past three decades, emotion, its neural substrates, activation, regulation, and functions have become hot topics in many areas of psychology and related disciplines. [...] Now we have a cornucopia of emotion books – amazon.com has 347,272 titles, and it is not unusual for a university library to have more than 400 scholarly books on the topic. Today there are at least five scientific journals with "emotion" in their titles and there are many more that publish research on emotion, resulting altogether in 2,732 articles in the past decade. There appears to be more agreement on the significance of emotion and much greater acceptance of its place in science than was evident 25 years ago."

<div align="right">(Izard [2010])</div>

The study of emotional phenomena goes back centuries. But most of the fundamental ideas in modern emotion research can be found in relatively "recent" theories. In the 17th century, Spinoza defines an emotion space in which all emotional experiences can be described according to 3 dimensions (axes): Joy, Desire and Sadness. A contemporary theory is that of Descartes, founded on the idea that emotions are a combination of emotional primitives. Despite the clear distinction in their basic definitions, the major disagreement between the two theories lies in the fact that Descartes makes a clear separation between the mind (cognition, the rational brain) and the body (emotion, instincts) while Spinoza unifies them and sees emotions as the foundation of mind. Later, in the 19th century, Darwin introduces the evolutionary theory. In his view, emotions are innate, universal and part of a genetic heritage. He also links emotions to the nervous system.

Another pioneer of modern emotion psychology is William James. His work is associated to Carl Lange's and referred to as the first "scientific" emotion theory. James and Lange separately introduced the so-called *peripherist* view: emotional experiences are physiological patterns, bodily changes induced by the autonomic nervous system due to certain stimuli or event. It is worth noting that Lange's position was more radical and that James himself did not agree with some of his statements [James, 1890]. Opposed to the "revolutionary" James-Lange approach, Walter Cannon and Philip Bard argue in favor of a *centralist* view: emotions are indeed triggered by the central nervous system processing the causing stimuli. The Cannon-Bard theory thus

goes back to the orginal view that physiological changes are consequences rather that causes of emotional experiences.

Do we cry because we are sad? Or do we feel sad because we cry? Psychologist have debated for decades, failing to agree on a clear definition of what is an emotion [Izard, 2010; Lindquist et al., 2013]. Despite the lack of consensus, significant progress have been made in understanding emotional processes[Izard, 2010]. Some of the concepts provided by Spinoza's, Descartes's and Darwin's seminal works are the bedrocks of more recent theories and models. The James-Lange and Cannon-Bard theories, both based on a physiological approach, had a strong influence on research of emotion (see [Coppin and Sander, 2012] for more details). Today, it is a macrocosm including researchers from as numerous fields as philosophy, psychology, sociology, ethology, neuroimaging, neurophysiology, psychiatry but also artificial intelligence and robotics.

With the emergence of cognitive science as a conglomerate of scientific disciplines investigating cognition, there has been considerable effort to engage in cross-disciplinary research on emotion. The International Society for Research on Emotions (ISRE) was founded in 1984 to provide a bridge between research areas studying emotional phenomena from different perspectives. Given their complex multi-faceted nature, emotions provide a perfect framework for interdisciplinary discussions and collaborations [Cañamero, 2014].

Of particular interest in the context of this thesis are the collaborations between computational approaches and those investigating biological emotions. In other words, between roboticists, cyberneticists and computer scientists on the one hand, and psychologists, neuroscientists and philosophers on the other hand.

Why should a roboticists be interested in the research on emotions? Nature seem to have provided emotion as a collection of problem-solving tools guiding more or less complex behaviors [Damasio, 2003; Rolls, 2005]. Indeed, many researchers highlighted the importance of emotions from the functional perspective [Fellous, 2004; Rolls, 2005; Scherer, 2009; Izard, 2010]. Among other functions, assessing events, focusing attention, enhancing communication, and motivating cognition and action are rather consensual [Izard, 2010]. Thus, for roboticists, the interest of emotional mechanisms is twofold: enhancing robot–robot and human–robot interactions, and increasing robot autonomy and adaptation capabilities [Arbib and Fellous, 2004; Cañamero and Gaussier, 2005]. Modeling both external and internal aspects of emotion in artificial systems led to new paradigms of thinking in robotics and artificial intelligence [Breazeal and Brooks, 2005; Nair et al., 2005; Hasson et al., 2011; Krichmar, 2013].

Why would a robotic/computational approach be beneficial to the research on emotions? One way computational models of emotion can be interesting is by allowing researchers to test hypotheses [Picard, 1997; Cañamero and Gaussier, 2005; Spackman, 2004]. Going from theoretical concepts to concrete implementations reveals ambiguities and forces more explicit definition [Marsella et al., 2010]. Thus, emotion theorists have incorporated computer science vocabulary – e.g. "information processing", "symbol" – in order to formalize their models [Ortony et al., 1988; Scherer, 2009]. But, the use of robots as *models* of biological emotions not only implies the realization and evaluation of existing theories but also the construction of new ones [Cañamero and Gaussier, 2005; Marsella et al., 2010]. Advances in the research on robot emotion should lead to new insights on the functions of emotions and suggest new avenue

for research on their neural bases [Fellous, 2004]. Also, the study of artificial emotions as an emergent property of a dynamical system sheds lights of the processes underlying this complex machinery [Cañamero and Gaussier, 2005; Hasson et al., 2011]. Less relevant to my thesis but still noteworthy is the way artificial systems led to new forms of human subject experimentation [Marsella et al., 2010]. Because they can be programmed to interact in a systematic way, they provide powerful tools for the design of psychological experiments [Boucenna et al., 2014, 2016; Oker et al., 2015].

4.1 Theories of emotions

In psychology, three major classes of emotion theories can be identified [Scherer, 2009; Coppin and Sander, 2012]:

- Basic emotions theories: Based on Darwin's work on the expression of emotions, this approach was further developed by Ekman and Izard among others. In this view, certain events or stimuli trigger certain expressions and physiological responses, which are characteristic of the corresponding emotion.

- Constructivist emotion theories: Rooted in James's idea that bodily changes are the basis of emotions, the notion of *core affects* representing neurophysiological states and leading to emotional experiences have been more recently advocated by Russel (dimensional theories) and Barret among others.

- Appraisal theories of emotions: Founded by Arnold and Lazarus, this school of thought gave rise to many models like Scherer's CPM (Component Process Model) or Ortony, Clore and Collins's OCC. Their aim is to describe the computational processing of information that lead from an external event to a change in the behavior.

An additional class of models, called anatomical, can be distinguished as well [Marsella et al., 2010]. This approach consists in proposing comprehensive theories built on neurobiological findings. In this this section, I will provide examples of such models which provide a higher level view of emotion. Then, the next section will focus more on the neural substrates implicated in emotional phenomena.

Obviously, some theories and models do not exactly fit in this classification, or fall into more than one category. Therefore, throughout the rest of the section, I shall detail some of these influential works regardless to classes and labels. For the sake of conciseness though, I will focus on those that seem relevant to computational implementations on artificial systems.

4.1.1 Ekman's basic emotions

Ekman is probably one of the most important proponents of the discrete and basic emotions approach [Ekman and Friesen, 1971; Ekman et al., 1983; Matsumoto and Ekman, 2004]. He argues that emotions correspond to distinct patterns of activation in the autonomous nervous system. Also, distinct emotions are governed by different neural circuitries that have evolved for the purpose of survival functions. In his theory, Ekman distinguishes basic emotion from more complex, moral and pro-social emotions [Matsumoto and Ekman, 2004]. The former, evolutionarily acquired are universal, in the sense that they have common properties across species. In

Ekman's theory, there seven basic emotions: initially anger, fear, sadness, disgust, surprise and joy [Ekman and Friesen, 1971], to which contempt was more recelty added [Matsumoto and Ekman, 2004]. To develop his theory of discrete emotions, he conducted numerous cross-cultural studies on facial expression. According to him, the fact that individuals are able to recognize the expressions of people from cultures they had little or no contact with (New-guinea with respect to the Western culture) proves that emotions are innate and universal.

4.1.2 Russell and Barrett's core affects

Like Spinoza's emotional space, many theorists based their definition of emotion on the idea that they are expressed by a certain number of dimensions. In their early work, Russell and Mehrabian [1977] proposed the three-dimension model PAD (Pleasure-Arousal-Dominance). Then, Russell [1980] introduced the circumplex model, illustrated in Figure 4.1, in which emotions are represented in a circular two dimensional space. Evidence from affective neuroscience in line with the existence of two independent neurophysiological systems representing the valence and arousal neural circuitries are presented to support this model [Posner et al., 2005]. Russell and Barrett [1999] proposed the concept of core affect to refer to the "cognitive" interpretation of the neurophysiological state; in other words, the conscious state induced by an emotion.
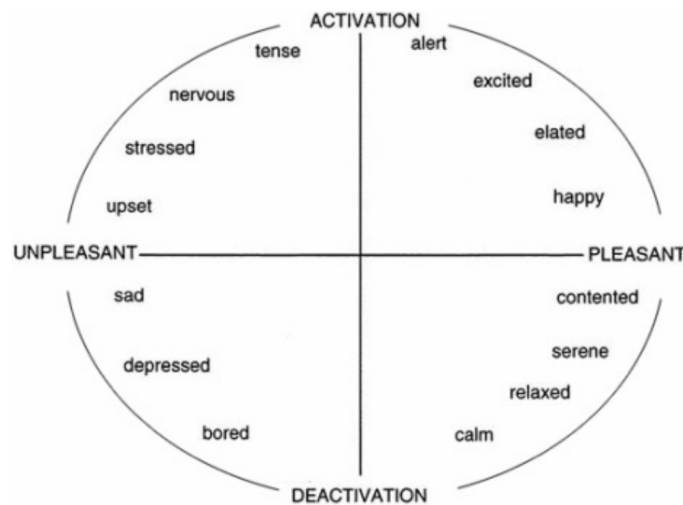
Figure 4.1: Russell's circumplex model. The horizontal and vertical axes respectively represent the valence and arousal dimensions. The figure is borrowed from Posner et al. [2005].

4.1.3 Ortony, Clore and Collins's cognitive structures

According to Ortony et al. [1988], emotions are discrete and elicited from a cognitive appraisal of events. The OCC theory distinguishes twenty-two types of emotions and proposes a semi-formal description of their cognitive structure. For instance, fear and hope result from the reaction to the prospect of negative and positive events. One of the interests of this model is to describe more complex emotions like shame and gratitude. Additional details will be given in Chapter 5.

4.1.4 Scherer's Component Process Model

The Component Process Model (CPM) elaborated by Scherer [1984, 2001] is one of the most influential model in the appraisal theories of emotion. It is based on the idea that the appraisal process is a rapid succession of stimuli-processing stages called "stimulus evaluation checks" (SEC). Scherer [2009] emphasizes the notion of synchronized changes in these subsystems and the need to think about emotions as a dynamic interaction between all the components of emotion. Scherer [2010] also argues that CPM provides a suitable blueprint for a comprehensive computational models of emotion. This model will be described in more details in Chapter 5.

4.1.5 Rolls's reward-based model

Rolls [1999, 2005] proposes to define emotions as states elicited by rewards (anything that is desirable) and punishments (anything that should be avoided). This view is based on the notion that it is more efficient to specify goals than particular behaviors. Thus, as illustrated in Figure 4.2, emotions depend on reinforcement contingencies: the presentation, omission or termination of reinforcers. For example, happiness is produced by the presentation of a reward while sadness, anger and frustration in contrast are caused by its omission (short period) or termination (end of presentation). On the other hand, the absence of a punisher may elicit relief.
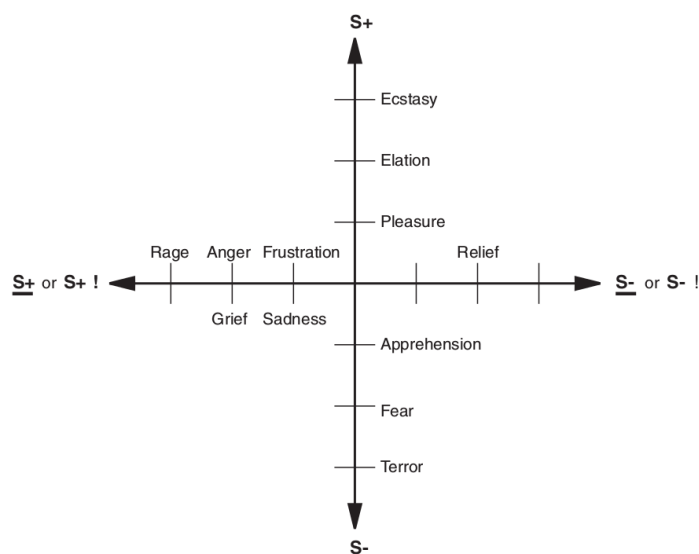
Figure 4.2: Rolls's reward-based model. (S) represents the presentation of a reinforcer. (S̲) and (S!) represent the omission and termination of a reinforcer. The figure is borrowed from Rolls [2005].

4.1.6 Damasio's conceptual tree

Associated to the neo-Jamesian school of thoughts, Damasio [2003] views emotion as emerging from basic physiological signals. This vertical relation is illustrated in the form of a conceptual tree shown in Figure 4.3, rooted in homeostatic mechanisms responsible for survival-related approach and avoidance behaviors. The nesting principle of the tree metaphor accounts for the universality of low level aspect of emotion across species. He also incorporates the notion of

basic emotions (fear, anger, joy, etc.) as non-linear combinations of lower regulatory reactions. On top of this nesting architecture and relying on basic emotions are more evolved capacities, characteristic of human beings like social emotions and conscious emotional experiences (feelings). Moreover, the appraisal of external stimuli and internal bodily signals drives the regulatory mechanisms at all levels to generate more or less complex emotional behaviors.
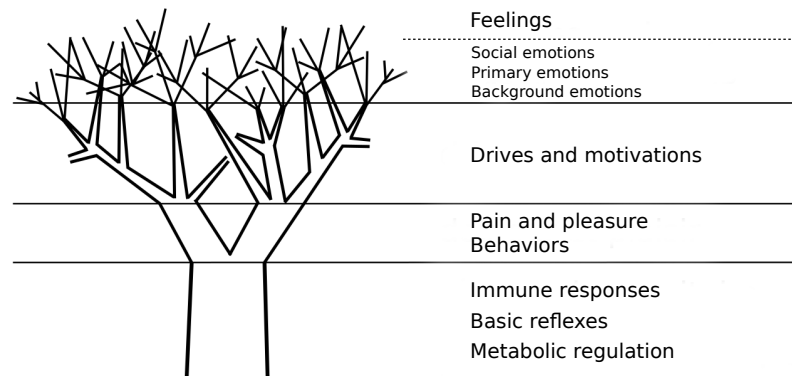
Figure 4.3: Damasio's conceptual tree. Nested homeostatic mechanisms build on simple reflexes to obtain more and more complex emotional mechanisms. On top of this hierarchy, feelings represent conscious experiences of emotion. The figure is a mix of two figures by Damasio [2003].

4.1.7 Discussion

Although the above listed views of emotions seem to have fundamental disagreements, they are not necessarily mutually exclusive. They rather focus on different aspects of emotion. While the basic approach studies what distinguishes an emotion from another (different patterns of neural activations), the dimensional view tend to describe what they have in common (shared circuitries associated to the valence and arousal). In fact, a dimensional model can provide a good readout method and be easily translated in a discrete emotion framework.

On the other hand, appraisal theories focus on the information processing that elicits emotions. Despite the emphasis put on cognition in this school of thought, they do not dismiss the bodily grounding of emotion. The strength of appraisal theories lies in the precise high level description of the combination of processes eliciting emotion. However, less is said about the functioning at the lower level (stimulus level). Taking the bodily (physiological and neurological) states as a point of departure to describe emotion is probably the main advantage of the former two approaches in this regard. Indeed, investigating the neural substrates involved in emotion provides good insights of the various operations that are responsible of the emotion elicitation.

4.2 The emotional brain

One of the most important notions one can encounter when studying the neurobiological aspect of emotion is that of the *limbic system*. It has its origins in the *triune brain* theory developed by MacLean [1949] based on the work of Papez [1937]. In this theory, a three-brain structure

sequentially emerged through the course of evolution: the reptilian, paleomammalian and neo-mammalian complexes. The limbic system represented by the second layer is thus thought to be responsible for motivational and emotional behaviors.

This view gave rise to the very common idea that emotions are primitive, evolutionarily conserved mechanisms that take place in subcortical brain areas. In contrast, more evolved functions like reasoning and planning are operated by the neocortex. This assertion is partly true since structures like the hypothalamus and amygdala play a significant role in emotional processes. However, included in the limbic system, the hippocampus and surrounding areas are today much more implicated in navigation and episodic memory (as seen in Part 1) than in emotion. On the other hand, cortical areas like the orbitofrontal cortex, the anterior cingulate cortex and the prefrontal cortex were shown to also significantly participate in emotion (see Figure 4.4). More importantly, it is today widely admitted that there are no specialized areas, but rather networks assuming functions as a whole. This will be reflected in the next section and more clearly highlighted in the next chapter.
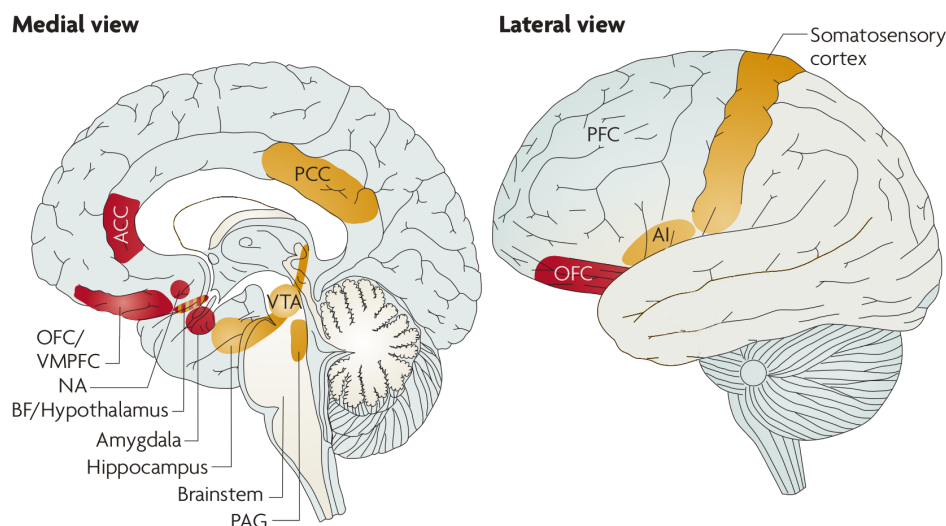
Figure 4.4: Brain areas involved in emotion. OFC: orbitofrontal cortex. VMPFC: ventromedial pre-frontal cortex. NA: nucleus accumbens. ACC: anterior cingulate cortex. VTA: ventral tegmental area. PAG: periaqueductal grey. AI: anterior insula. The figure is slightly readapted from Pessoa [2008].

4.2.1 Brainstem (BS)

BS is the part of the brain which is adjoining and structurally continuous with the spinal cord. Thus, it controls the flow of messages between the brain and the rest of the body (e.g. sensori-motor signals). BS host structures that are central to the chemical signaling systems. Thereby, it assumes important functions related to vital reflexes and regulatory control.

Periaqueductal Grey (PAG) is a primary center for the transmission and modulation of noci-ceptive signals. For example, PAG lesions in rats either attenuate or accentuate freezing behav-iors normally caused by nociceptive electric shock – depending on whether the dorsolateral or the ventral parts are damaged [Fanselow et al., 1995]. PAG also projects to Raphe Nuclei which secrete serotonine. The latter is thought to modulate harm aversion and pain processing [Cools

et al., 2008]. Locus Coeruleus, which another nucleus of the Reticular Formation, is the principal site for brain synthesis of norepinephrine involved with physiological responses to stress and panic. All in all, the Reticular Formation is thought to modulate the arousal level of the central nervous systems [Cardinal et al., 2002]. Moreover, the largest group of dopamine neurons is found in the Ventral Tegmental Area (VTA). This neurotransmitter is involved in reward processing and motivation [Berridge, 2012]. More details on these chemical messages linked to emotion and affect will be given in the next section.

4.2.2 Hypothalamus (HTH)

The importance of HTH in emotion was highlighted by Cannon and Bard "decortication" experiments [Bard, 1934]. For instance, emotional expressions are suppressed due to HTH ablation. In is also present in the Papez [1937] circuit via the mammillary bodies.

HTH is commonly associated with survival-related functions. As the major center of neuroendocrine integration in the brain (i.e. the release of hormones at the reception of neuromodulatory input), HTH links the nervous system to the endocrine system. It intervenes in various bodily functions such as monitoring physiological parameters and regulating hunger and thurst. More details on these chemical messages linked to emotion and affect will be given in the next section. Besides, more recent findings indicate massive bidirectional connections with the cortex in addition to subcortical regions. Subsequently, not only does it have basic *descending* control functions, it is also implicated in higher cognitive functions [Pessoa, 2010].

4.2.3 Amygdala (AM)

AM is one of the most connected areas of the brain and of the most important structures with regard to emotional processes. It is a *connector hub* linking multiple *provincial hubs* with high connectivity within specific functional groups [Pessoa, 2010]. Among others, AM is connected to the thalamus, nucleus accumbens, sensory cortices and prefrontal cortex.

The role of AM in emotion has been extensively investigated, particularly by LeDoux and colleagues [LaBar et al., 1998; LeDoux, 2003, 2012]. Most of this research is based on Pavlovian fear conditioning, which consists in pairing meaningless environmental stimuli (Conditioned Stimuli, CS) with biologically potent ones (Unconditioned Stimuli, US) in order to associate them with emotional defensive responses (Unconditioned responses, UR). For example, LaBar et al. [1998] showed that AM activation is higher during acquisition and extinction phases – that is to say respectively when US and CS are first being associated and when CS starts to repeatedly appear without US afterward. Although AM is better known to mediate aversive learning, it is also active during appetitive conditioning [Balleine and Killcross, 2006; Weymar and Schwabe, 2016].

A commonly admitted function of AM is thus to attribute an affective valence to sensory stimuli (Paton et al. [2006]; Lewis [2005]; Pessoa [2010]). LeDoux [1996] proposed that a direct subcortical route from sensory periphery (e.g. retina) to thalamus and AM allows the rapid emotional evaluation. This pathway is referred to as the *low road*, as opposed to the so-called *high road* indirectly transmitting information to AM through cortex. The former thus relies on coarse representations while the latter is based on more detailed, refined information. However, this strict opposition is challenged by recent findings suggesting additional rapid pathways

through cortical regions such as the orbitofrontal cortex [Barrett and Bar, 2009; Pessoa and Adolphs, 2010].

4.2.4 Basal Forebrain (BF) and Basal Ganglia (BG)

There is a large literature acknowledging the existence of multiple cortico-basal circuits implementing a variety of functions [Alexander et al., 1986; Middleton and L, 2001; Haber, 2008; Jahanshahi et al., 2015]. Among the latter, emotional/limbic loops includes interactions with the orbitofrontal and anterior cingulate cortices, which I will describe in the following paragraphs. This functional description in the shape of loops encompasses neuronal connections (information transmission) and neurochemical projections (modulation).

The Nucleus Accumbens (NAc) is located in the ventral striatum of the basal ganglia. It is part of the pleasure and reward circuitry. It is where the dopaminergic circuitry departing from the brainstem splits into the mesolimbic and mesocortical pathways. NAc has been implicated in drug addictions. For instance, repetitive, non-decremental stimulation of dopamine transmissions by drugs abnormally strenghtens stimulus–drug associations [Di Chiara, 2002]. Also central to the experience of pleasure, NAc is identified as a hedonic hotspot [Berridge and Kringelbach, 2008]. More specifically, Peciña and Berridge [2005] suggest that opioid circuits for stimulating hedonic 'liking' circuits are tightly localized in the rostromedial shell of NAc.

A collection of other structures are also involved in chemical signaling: the substantia nigra in BG is a source of dopamine; the nucleus basalis in BF and the striatum in BG are major sources of acetylcholine. More details on these chemical messages linked to emotion and affect will be given in the next section.

4.2.5 Prefrontal Cortex (PFC)

The prefrontal cortex (PFC) in the cerebral cortex located in the front part of the brain (frontal lobe). Based on connectivities and functional specificities, PFC can be decomposed into three substructures: the dorsolateral (dlPFC), ventromedial (vmPFC) and orbitofrontal (OFC) parts.

The role of the prefrontal cortex in emotion was made famous by the case of Phineas Gage. Due to an explosion, a metal bar went through the patient's skull, face and brain. Surprisingly, Gage recovered a few months later. But, the lesion of the brain caused important changes in his personality and emotional behavior. The case occurred in 1848 but was put back under the spotlight by Damasio and colleagues. Indeed, Damasio et al. [1994] reconstituted the lesion and determined that the most impacted regions were the prefrontal cortices, and the left OFC more specifically.

In addition, Damasio et al. [1996] proposed the *somatic markers* hypothesis. They consist in associations between certain situations and the patterns of elicited physiological and emotional reactions that are established since early developmental stages. According to Damasio, these associations take place in vmPFC. The somatic markers – carrying the emotional value of encountered situations – then influence decision making when a quick response is required or when when logical reasoning fails in driving one's behavior [Damasio et al., 1996; Damasio, 2003; Bechara and Damasio, 2005].

OFC has numerous connections to limbic areas and is also thought to strongly participate in emotional processes. While the orbital part receives extensive sensory information, the medial part projects to HTH, NA and BS. Öngür and Price [2000] suggest that OFC plays a role

in sensory-visceromotor associations and guidance of emotional, reward-related behavior (especially food-related). Rolls [2004] also supports the idea that OFC participates in learning the emotional and motivational values of stimuli. He proposes that OFC and AM jointly establish associations between primary (e.g. taste and touch) and secondary reinforcers. According to Barrett and Bar [2009], OFC represents a key circuitry responsible for the creation of contextually sensitive, multimodal representations of the world and its emotional and hedonic value.

4.2.6 Insula

The insular cortex appears to be a central area in conscious emotional and survival-related experiences [Damasio, 2003]. For example, consciousness of thirst causes activation in this area while satiation deactivates it [Denton et al., 1999]. This is also observed in the cingulate cortex which will be addressed later on. Also, the experiment by Casey [2000] compared a condition where participants' hands were put in ice-cold water to a condition where they were subjects to a non-painful vibratory stimulus. Insula was shown to be particularly activated by pain sensation. Moreover, when thinking of emotional episode to the point of feeling the corresponding emotion, insula, among other somatosensory cortices as well as the cingulate cortex and brainstem tegmentum, exhibits significant activities [Damasio et al., 2000].

4.2.7 Anterior Cingulate Cortex (ACC)

ACC is a complex cortical area with numerous bi-directional projections to limbic regions like AM and OFC, and particularly to autonomic regulatory areas such as HTH and PAG. It is assumed to participate in emotional evaluation and cognitive control [Bush et al., 2000; Shackman et al., 2011; Inzlicht et al., 2015]. Botvinick et al. [2004] and Inzlicht et al. [2015] highlighted the role of ACC in conflict monitoring and the triggering of motivated behavior. ACC is also thought to be involved in the conscious experience of emotion, in conjunction with other regions like somatosensory cortices and insula [Denton et al., 1999; Damasio et al., 2000; Damasio, 2003]. Bush et al. [2000] proposed an emotion/cognition division of ACC, distinguishing the ventral/rostral part and the dorsal part for these processes respectively. But it is noteworthy that these "emotional" and "cognitive" subregions are also inter-connected. Also, these is evidence challenging the emotion/cognition separation in the case of emotionally triggered control [Shackman et al., 2011].

4.3 The neurochemical brain

Another aspect of the emotional brain that needs to be addressed lies in the chemical signaling systems that modulate the neuronal activities throughout the nervous system. They are central in the study of emotion because of their impact on metabolic regulation, bodily responses, pleasure and pain sensations, motivation and motor activation [Fellous, 2004; Kelley, 2005].

There are three kinds of chemicals messengers in the brain: neurotransmitters, neuromodulators and neurohormones. Neurotransmitters are released by the nerve ending of a neuron in order to communicate with adjacent neurons. In contrast, neuromodulators are released by neurons farther away from the target cells which are larger groups of neurons. Yet, some neurotransmitters – such as dopamine, serotonin, acetylcholine and norepinephrine – act like neuromodulators in the sense they have an effect on distant neurons. On the other hand, neurohormones are not

released by neurons but by glands called neuroendocrine cells, which receive neuromodulatory inputs from other cells in order to secrete those hormones. Some hormonal molecules also have a role of neuromodulation. Among them are opioid peptides like endorphins. For convenience, although specified otherwise, I will use the term neuromodulator to refer to all these endogenous substances in the rest of this document.

Neuromodulators are secreted by specialized structures receiving inputs from brain areas involved in all levels of behavior (reflex, motivation, cognition). They influence the neural activities of nerve cells by activating associated receptors in neuronal membranes. The global effect of a neuromodulatory system may vary considerably depending on the combination of receptors and the circuitry within the modulated brain region and the connections of that region within the brain. For example, even two neurons of the same structure may have different mixtures of receptors on their membranes. However, the effect of a neuromodulator on a particular receptor is very specific. For example, it can modulate the excitability of the target neuron. Besides, all neuromodulatory systems have both tonic and phasic activity responses: respectively corresponding to a general elevation of the baseline activity and to short bursts of activity.

It is noteworthy that an important part of what we know about the neuromodulatory systems comes from research on drug addiction and psychiatric disorders.

4.3.1 Dopamine DA

Given its role in motor activation, appetitive motivation, reward processing and cellular plasticity, DA is a key neuromodulator as far as emotion is concerned. There are two major classes of DA receptors: D1 and D2. As seen previously, DA is mainly produced in VTA (but also in other structures like subtantia nigra) and targets many cortical and subcortical area like NA, AM and PFC.

Historically, dopaminergic pathways where often associated with the notion of pleasure. Indeed, rewarding stimuli such as highly palatable food strongly activate DA release [Bassareo and Di Chiara, 1999]. However, in line with other findings, the work of Berridge and colleagues [Robinson and Berridge, 1993; Berridge and Kringelbach, 2008; Berridge, 2012] confers a more motivational function to this neurotransmitter. Rather than being linked to the hedonic impact of the reward, DA is thought to be important for attributing "incentive salience" to neural representations of rewards. In other words, for 'wanting' instead of 'liking'. For example, DA deprived rats are unable to generate the motivation arousal necessary for ingestive behavior and can starve to death although they are able to move and eat [Ungerstedt, 1971]. In addition, there is evidence supporting the hypothesis that DA mediates reward-related learning. Predictive reward signals are often coded by dopaminergic systems [Schultz, 2004; Baldwin et al., 2002]. Interestingly, the dopamine system can also be activated by aversive stimuli [Scott et al., 2006].

4.3.2 Serotonin 5-HT

5-HT has been implicated in various behavioral functions like motor pattern generation, arousal, sleep, mood and social behavior. There are 14 different 5-HT receptors in human nervous system. The main serotonergic source is the raphe nuclei of the reticular formation. Descending projections reach brainstem, and spinal motor and sensory regions. Ascending inputs project to widespread regions in the cortex in addition to limbic areas like HTH and AM.

Importantly, the number and variety of 5-HT receptors and projections suggest that this neuromodulation is very complex. In nonhuman primates and mammals in general, there is converging evidence of the 5-HT modulation of aggression. More precisely, 5-HT depletion is inversely correlated with an increase in aggressive behaviors [Mehlman et al., 1994]. In contrast, high levels of serotonin enhances social behavior and decreases aggression [Raleigh and McGuire, 1991]. In humans, while abnormally low 5-HT levels are also known to correlate with increased aggressive or antisocial behavior, alcoholism, and impaired impulsive control, patients with reduced serotonin function have been shown to have higher rates of major depression and suicide attempts [Mann et al., 1996]. More recent findings link 5-HT to the neural processing of punishment and harm aversion. In that matter, Cools et al. [2008] proposes that serotoninergic modulation might affect aversive motivational processes in the form of punishment prediction error.

4.3.3 Opioids

The term opioids refers to a large variety of substances that act on opioid receptors (mu, delta, kappa receptors), like drugs derived from opium (opiates). Here, we are interested in endogenous opioid peptides, which include endorphins, enkephalins and dynorphins. The latter are involved in emotional regulation, responses to pain and stress, endocrine regulation and food intake.

The most basic role of opiods in emotion is the generation of pleasure in the hedonic circuits. While extensive research investigated their effect in the modulation of food intake, such effect is thought to be mediated by an increase of food reward and palatability [Kelley et al., 2002]. For instance, the activation of opioid receptors in rats highly increases the number of 'liking' reactions to a sucrose taste [Peciña and Berridge, 2005]. Opioids are also known to reduce pain. Indeed, it has been shown that exposure to physical and psychological stress activates opioid system and provokes analgesia [Lester and Fanselow, 1985; Bolles and Fanselow, 1982]. Such mechanism might have evolved to facilitate adaptive escape in case of threat. Yet another remarkable effect is to foster social attachment. Investigating separation distress, Panksepp et al. [1980] proposed that brain opioid systems might provide neurochemical mediation of social-bonding among animals. Thus, such a higher order behavioral process as social attachment would have arisen from primitive neural circuits which subserved compatible functions like hedonic liking.

4.3.4 Acetylcholine (ACh) and Norepinephrine (NE)

In the central nervous system, ACh is secreted by cholinergic nuclei in the brainstem and the basal forebrain. Norepinephrine (NE) is also called noadrenaline. Although the former name is generally preferred, the system that produces NE is referred to as noradrenergic. The prime source of NE in the central nervous system is the locus coeruleus situated in the brainstem.

Both neurotransmitters are involved in arousal, alertness and attention. Angela and Dayan [2005] proposes that ACh and NE respectively signal expected and unexpected uncertainty. The former consists in the known unreliability of predictive cues within a context while the latter refers to strongly unexpected observations. For example, NE locus coeruleus neurons in rats respond to novelty and change in the environment [Vankov et al., 1995]. Besides, NE and ACh have complementary role in sleeping and waking phases [Jones, 2005].

4.3.5 Discussion

The study of the neurobiology of emotion demonstrates that there is no specialized brain area implementing an emotional system. Rather, emotions involve decentralized processes that operate on different behavioral levels. Several brain regions participate, with different functions supported on different levels. Emotional experiences imply complex interactions between these regions. Therefore, it is difficult to even decompose the brain in specialized structures responsible for specific functions.

Yet, some regions seem to be more involved in the appraisal processes that trigger emotion. Among these are AM, OFC and ACC, which allow more or less complex emotional evaluations, from pure stimuli-based level, to contextualized appraisal, to conscious experiences. Thus, they implement an important part of what can be referred to as the neural 'valence circuitry'. On the other hand the execution of emotional responses is at least partially orchestrated by areas like HTH, BS and NA. Through their connection with the neuromodulatory system, they participate in the 'arousal circuitry'. From the neurobiological perspective, the emotion theories have strong overlaps. However, they can be distinguished based on whether they consider there is a unique circuitry appraising events and triggering emotional responses, or distinct circuitries handling different dimensions of emotion, or even more circuitries each one dedicated to certain types of emotions. Another distinction concerns whether the emphasis is put on what nature provides in these circuitries (i.e. what is common across species) or what nurture gives (i.e. what is specific in human emotion).

4.4 From biological to artificial emotions

The literature reviewed so far shows the variety of functions in which emotions are implicated. They also highlight the complexity of the machinery related to emotional phenomena in living beings. In the last decades, substantial research has been interested in reproducing these mechanisms in artificial systems. The resulting computational models can consist in top-down implementations of existing models (e.g. OCC, PAD, CPM) or bottom-up simulations of elementary properties related to emotion (e.g. physiological functions, facial expressions, motivated behavior). They adopt approaches that range from symbolic, rule-based to dynamic, network-based programming. Some of them are implemented on virtual conversational agents while others on robotic platforms. Like with theories of emotions, I will detail some of these computational model regardless to any classification.

4.4.1 Gebhard's ALMA model

Because of the simplicity and implementability of the OCC model [Ortony et al., 1988], various computational models of emotion implemented it in formal descriptions of emotional responses [Adam et al., 2009; Dastani and Lorini, 2012]. The "layered" model proposed by Gebhard [2005], called ALMA, also relies on the set of appraisal variables described by OCC. But, instead of giving a strict implementation of OCC, the appraisal rules are specified in a XML-like language before the runtime phase to determine the agent's personality profile. The emotion engine simulates three interacting kinds of affect that operate at different timescales: emotion, mood and personality. Another specificity is the combination of OCC with the PAD dimensional model [Russell and Mehrabian, 1977]. This consists in viewing moods as medium-term affects

averaging the agent's emotional states in time. Gebhard thus proposes a discretization of the mood states as represented in the three dimensional PAD space into eight distinct mood octants; to which he also maps the OCC discrete emotions. ALMA provides a good example of how implementing emotions in artificial systems can foster reflections on how to link different emotion theories. It also addresses the issues of interactions between affects of different timescales. Thanks to the use of an XML-like language, it is a flexible framework that can allow other virtual agents designers to extend it. Yet, like in other OCC-based models, affect results from a purely cognitive appraisal. It also suffers from the symbol grounding problem and the need to hard-code the rules that guide the emotional behavior.

4.4.2 Morgado and Gaspar's agent flow model

Morgado and Gaspar's agent flow model aims at giving a unified view of emotional and cognitive process [Morgado and Gaspar, 2007, 2004]. The authors advocate the idea that the emotional appraisal depends on two base factors: the agent potential to produce change in environment (achievement potential) and the environment conduciveness or resistance to that change (achievement conductance). These dimensions respectively correspond to the agent skills and the environmental challenges. In their model, the emotional disposition resulting from these variables is expressed in the cognitive space in terms of the distance from a goal/target situation and the velocity of the generated "movement". This model provides an interesting insight on how emotional states can be translated into actions in an artificial agent. Moreover, by arguing for the importance of their two base appraisal dimensions, the authors intend to influence appraisal theories through a computational approach. However, this model is evaluated in a simple 2D environment where the cognitive attributes determining the agent emotional dispositions are easy to compute (Morgado and Gaspar [2007]). There is a possible gap in terms of scaling up to real world experiments. Also, some of the factors ruled out by Morgado and Gaspar [2004] from the fundamental processes of emotion elicitation (like novelty) could turn out to be essential in real conditions.

4.4.3 Courgeon and colleagues' MARC agents

Multimodal Affective & Reactive Characters (MARC) is a framework for 3D interactive expressive virtual agents designed by Courgeon and his colleagues. Interestingly, this system was used to implement and evaluate different theoretical models of emotion in the context of human–agent interactions [Courgeon et al., 2008, 2009; Courgeon and Clavel, 2013]. A discrete emotions architecture mainly considered the six basic emotions listed by Ekman and Friesen [1971]. The dimensional version relied on the PAD model proposed by [Russell and Mehrabian, 1977]. Last, the cognitive approach to emotion was addressed by adapting the CPM model introduced by Scherer [1984, 2001] to which a social appraisal aspect is also added. In the evaluation of these models, it is for example argued that the continuous space provided by a dimensional approach allows smoother transitions between different states and thereby more realistic emotional feedback as compared to discrete emotions. However, in contrast to the appraisal approach, both discrete and dimensional theories are judged purely descriptive and insufficient for an autonomous generation of emotional responses by an interactive artificial agent. Unfortunately, all the implemented architectures are compared from the perspective of virtual agents design and the results are not discussed in terms of their impact on the understanding of human emotions.

4.4.4 Velásquez's Cathexis

Cathexis is a distributed computational model proposed by Velásquez [1997] to model the dynamic nature of emotional phenomena like emotions, moods and temperaments. Inspired by the work of Minsky [1986], "emotion families" are represented by proto-specialists, i.e. distributed modules that monitor external and internal events. The term "emotion families" refers to groups of related affective states sharing certain properties like the physiological states, the eliciting events or the responses; and influencing each other. For instance, fear, fright, terror and panic. In addition, Velásquez [1997] considers four types of emotion elicitators: neural (e.g. physiological signals), sensorimotor (e.g. facial expressions and body postures), motivational (e.g. hunger- and thirst-related drives) and cognitive (e.g. beliefs and desires). Although this classification/labeling may be questionable, mixing cognitive and non-cognitive elicitators is a significant characteristic of Cathexis. Yet, this model was only tested in a simplified, structured human–agent interaction via a 2D user interface, which may seem quite insufficient to evaluate its performance in real situations.

4.4.5 Breazeal's Kismet social robot

Kismet is an expressive face robot designed by Breazeal [2003] exhibiting emotional behaviors during interactions with humans. The architecture proposed is also based on the idea of distributed specialists ([Minsky, 1986]) handling separate perceptual, behavioral or emotional processes [Breazeal and Brooks, 2005]. The emotion system consists in mapping prototypical "precipitating" events to behavioral and facial responses. It includes the six basic emotions listed by Ekman and Friesen [1971] to which others emotional states like boredom and interest are added. For instance, the presence of an undesired stimulus elicits disgust which triggers a withdrawal behavior in order to signal rejection to the human caregiver. The considered discrete emotions are mapped to a dimensional space similar to that proposed by [Russell and Mehrabian, 1977]: valence-arousal-stance. Interpolation between various state provides continuous expressiveness in facial movements. Interestingly, this seems to enable a variety of behaviors although the elementary specialists are hard-wired.

4.4.6 Canamero and colleagues' hormone-based models

Cañamero [1997] introduced a behavior-based architecture simulating synthetic physiological variable. Following a developmental approach, the artificial agents are seen as newborns whose behavior is driven by survival-related parameters modeling bodily needs and basic emotional responses. This seminal work was later extended and tested on real robots. Avila-garcía and Cañamero [2004, 2005] particularly evaluated the hormone-based paradigm in the context of competitive scenarios. Put in an environment containing various types of resources and obstacles, the robots must select their actions from a repertoire of behaviors in order to survive and maintain their well-being. Internal variable included parameters related to food and temperature while behaviors could be consumatory (feeding and warming up), appetitive (searching) or aversive (avoidance). Moreover, Lones et al. [2014] more specifically studied the emergence of aggressive behaviors in certain situations due to the influence of such physiological and well-being related variables. Cañamero and colleagues' work on hormonal modulation of perception as a paradigm for action selection demonstrate how survival-related homeostatic mechanisms

can implement an efficient architecture and how emotional responses can emerge from the dynamics of interaction between the robot and the environment.

4.4.7 Krichmar's neuromodulation-based models

The approach proposed by Krichmar [2008, 2013] is to use neuromodulation as a controller for artificial agents. More specifically, Krichmar [2008] views the tonic and phasic modes of neuromodulatory system as a way to implement a biologically plausible network mechanism for the exploration/exploitation trade-off. The author suggests that tonic responses favor exploratory behaviors – because the system is more prone to random fluctuations – while the transient bursts of activity in the phasic mode strongly bias the output toward preferred responses. More recently, Krichmar [2013] shows how neuromodulatory interactions can control the approach and avoidance behaviors of a real robot. The dopaminergic system modulates the curiosity and reward-seeking responses while the serotonergic system influences harm aversion. Cholinergic and noradrenergic systems provide an attentional filter that regulates the curious vs. anxious behaviors. The computation framework provided by Krichmar helps to shed light on the interactions between neuromodulatory systems that are generally studied separately in neuroscientific research. It gives an insight on the mechanisms by which the nervous system might handle action selection and the exploration/exploitation trade-off.

4.5 General discussion

As pointed out in the beginning of this chapter, the interest for artificial emotions has two complementary but distinguishable objectives: 1) taking inspiration from biological solutions to design systems that better interact with their physical and social environments, and 2) providing computational models as tools to better understand biological emotions. Despite their complementarity, these goals are not easy to conciliate. As a matter of fact, most of the artificial emotion systems are intended to facilitate interactions with humans. Generally, this kind of approach translates into implementations of existing models with no interest in providing more insights on biological emotions. The main objective lies in the application. In contrast, other computational models target basic components of emotional phenomena and may help understand emotion, although the application scope may be more limited.

Theoretical and experimental research on human and animal emotion have largely inspired models of artificial emotion. Earlier in this chapter, I gave an example of how different approaches to emotion could be classified according to certain schools of thoughts. But in terms of how they influence/contribute to the computational modeling of emotions, a particular criterion seems relevant: are emotions *more importantly* considered as cognitive processes or bodily responses. Rather than a strict dichotomy, this results in a continuum between cognitive and embodied views as far as theory is concerned. However, in computational models, the former view is generally more related to rule-based implementations whereas the latter is more likely to be adopted in distributed, network-based models. The parallel with the cognitivist and the connectionist approaches to cognition that I detailed in the introduction of my thesis can easily be drawn. Thus, my position will be to take an embodied view of emotion as a point of departure. A particular interest will be put neurobiological aspects, considering both neural and chemical brain networks.

Conceptualizing emotion–cognition interaction: a process-oriented model

The review of the multidisciplinary study of emotion I presented in the previous chapter has a set of important implications. The first one is that emotions are decentralized processes involving a complex neuronal and neurochemical machinery. There is no central emotional region which activity can be studied in order to solve the mystery of emotional phenomena. The second implication is that emotions operate at many time scales and at many behavioral levels. A core system that seems to be shared accross species provide basic regulatory, survival-related mechanisms. But more evolved species like primates and humans also benefit from more complex emotional capabilities such as social emotions and conscious emotional experiences. Last, because of the complexity and the richness of emotional processes, it seems hard to make a clear separation between emotion and cognition.

The term *cognition* refers to a large panel of mental processes that are necessary to control behavior and attain goals. It encompasses perception, attention, memory, learning and motor control; but also more sophisticated mechanisms that are involved in reasoning, perspective-taking and complex decision-making. In particular, cognitive control consists of a set of executive functions allowing goal-driven behaviors [Diamond, 2013]. Among them are:
1) inhibition, which includes response inhibition (self-control related to impulsive actions) and interference control (selective attention),
2) working memory, i.e. maintaining information for a short period of time,
3) cognitive flexibility, e.g. switching from one concept to another, changing perspective, etc.

Whereas the meaning of cognition is relatively vague, that of *emotion* is even less consensual, as I pointed out in the previous chapter. However, there is quite an agreement that it involves partially dedicated neural circuits, response systems, and feeling state/process that allow for assessing events, focusing attention, enhancing communication, and motivating cognition and action [Izard, 2010]. This description, although it cannot serve as a clear and comprehensive definition of emotion, highlights the close relation between emotion and cognition.

When addressing the central question of how emotion interacts with cognition, one of the following two perspectives is generally taken: how cognition influences emotion, or how emotion influences cognition. It probably depends on the school of thoughts one is attached to, and on whether emotions are considered as (*mostly*) originating from cognitive processes or from bodily responses. However, some research works also try to describe the emotion–cognition

interplay in an integrated view. These models generally insist on the dynamic interaction between the two kinds of processes. They will also allow me to introduce the generic model that I propose and implement in this thesis.

## 5.1 Cognitive influence on emotion

Appraisal theories of emotion provide a very interesting framework for studying the set of cognitive processes involved in the activation of emotions. These models suggest that organisms constantly explore their environment and react to relevant stimuli. Thus, the aim is to describe the computational processing of information that leads from an external event to a change in the behavior. Appraisal theories postulate that different emotions result from different patterns of evaluation (processing). This idea was originally proposed by Arnold [1960]. In her seminal work, she proposed that the appraisal process is distributed over four components: physiological reaction, motor reaction, motivation for action, and subjective feeling.

Another pioneer of this school of thought is Lazarus [1991]. In his model, he identified two levels of appraisal. The primary appraisal evaluates stimuli with regards to one's goals and well-being. Then, the secondary appraisal analyzed the resources that one can use to respond and adapt to the situation. Re-appraisal allows modifications of the primary and/or secondary appraisal while the interactions between the individual and the environment occur.

In this framework, various theories arose, aiming to specify a set of standard criteria that are presumed to underpin the emotional appraisal process. Some of them conceptualize emotion as an information-processing system just like any cognitive mechanism [Oatley and Johnson-Laird, 1987; Ortony et al., 1988]. For instance, event-based emotions in the OCC theory [Ortony et al., 1988] such as joy and fear are triggered by the evaluation of the desirability of event consequences. The notion of desirability here is quite close to that of utility; that is to say whether the event satisfies or interferes with one's goals. In addition to the desirability of events, the central criterion in OCC allows for evaluating the approbation of actions and the attraction of objects. Then, secondary appraisal variables, such as the likelihood, the unexpectedness or the praiseworthiness of a state of affairs, influence the intensity of the triggered emotion. This model of emotion is widely used in computer science because of its implementability. However, it is particularly subjects to general criticisms of appraisal/cognitive views of emotion [Scherer, 2009; Coppin and Sander, 2012]: e.g. relying on symbolic representations and high-level information processing or not accounting for the rapid onset of some emotional reactions.

The CPM model proposed by Scherer [2010] attempts to address these criticisms with a multilevel approach to appraisal. In this model, the appraisal module responsible for the evaluation of objects and events is organized as a rapid succession of stimuli-processing stages called "stimulus evaluation checks" (SEC). As shown in Figure 5.1, four major appraisal objectives are listed: relevance, implication, coping potential, and normative significance. For instance, to evaluate relevance, there is a check for novelty, for intrinsic pleasantness, and for the importance with respect to goals and needs. To evaluate implication, there is a check for cause, for probable outcomes, for discrepancy from expectation, for conduciveness for goals, and for urgency. These appraisal checks require cognitive processes related to attention, memory, motivation, reasoning and self.

To counter the above mentioned criticisms of appraisal theories, Scherer [2009, 2010] claims that all the criteria can be processed in parallel at three hierarchically organized levels. First, the
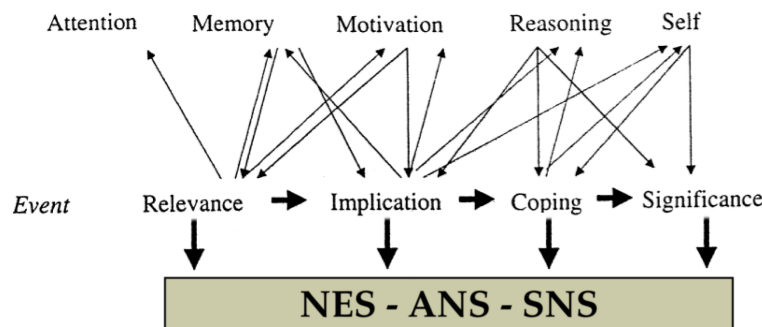
Figure 5.1: Scherer's representation of the connections between cognitive components and emotional appraisal processes. Cognitive components like attention and memory are required during the appraisal processes carried out by the central nervous system (CNS) and influencing the neuroendocrine system (NES), autonomic nervous system (ANS), and somatic nervous system (SNS). The figure is slightly readapted from Scherer [2001].

sensorimotor level, in which the checking mechanisms are mostly genetically determined and based on pattern matching. Second, the schematic level, in which the processing is automatic and unconscious based on social learning and repeated experiences. Third, the conceptual level, involving propositional knowledge and requiring consciousness. However, while the checks take place in parallel, it is assumed that the outcome of the evaluation is sequential: novelty, intrinsic pleasantness, task/goal relevance, goal conduciveness, etc. In particular, Grandjean and Scherer [2008] tested the sequence hypothesis in two experiments manipulating novelty, intrinsic pleasantness, and goal relevance. The measured brain activities support the assumption that early checks (novelty and intrinsic pleasantness) occur in an automatic, unconscious mode of processing, whereas later checks (specifically goal conduciveness) require more extensive, effortful and controlled processing. Moreover, a neuroimaging experiment conducted by Sander et al. [2005] demonstrates that the processing of affectively and socially salient signals is modulated by voluntary attention. The experiment used a dichotic-listening paradigm; meaning that angry and neutral prosody were presented simultaneously to both ears while participants were asked to perform a gender decision task on voices heard in one ear, while they had to ignore voices presented on the other side. The results show that, although the amygdala responded to anger stimuli independently of attention, the orbitofrontal cortex showed greater activation to the same emotional stimuli presented on the to-be-attended side compared to the to-be-ignored side.

Last, although out of the scope of my thesis, it is worth mentioning that another example of cognitive influence on emotion is *emotion regulation* [Gross, 1998; Ochsner and Gross, 2005]. It involves various strategies aiming to alter ongoing emotional responses or initiate new ones. Among them are: selective inattention to emotional stimuli, suppression of expressive behavior, or (re-)interpretation of emotion-eliciting situations in ways that limit emotional responding. The prefrontal cortex and the anterior cingulate cortex appear to be key regions for these regulatory processes [Ochsner and Gross, 2005].

5.2 Emotional influence on cognition

There is also significant evidence of the influence of emotion on cognition. Here, I will focus on the emotional modulation of three types of processes: perception, attention and memory.

5.2.1 Perception

The effect of emotion on perception can be observed from the modulation of visual processing. For example, emotional stimuli induce an increase in the activation of the visual cortex [Padmala and Pessoa, 2008] after affective conditioning. Also, Phelps et al. [2006] showed that emotionally arousing stimuli (fearful faces) lowered the contrast threshold in comparison with neutral stimuli. In other words, participants were more sensitive to visual contrast when they had previously seen emotional faces than when they had seen neutral faces.

More evidence can be found in studies investigating the perception of space and distance from objects and individuals. For example, positively valenced objects tend to be perceived as closer and more reachable than negative ones [Valdés-Conroy et al., 2012; Balcetis and Dunning, 2009]. Also, Coello et al. [2012] showed that a knife seems farther when oriented toward us, i.e. when potentially dangerous. On the other hand, a positive affective state, induced by pleasant music for instance, reduces the area needed to feel comfortable in over-crowded spaces [Tajadura-Jiménez et al., 2011]. The perception of our peripersonal space indeed seems to depend on a emotional evaluation of external stimuli. For example, a study by Kennedy et al. [2009] implicates the amygdala in this process.

5.2.2 Attention

Emotions also influence attentional processes. Öhman et al. [2001] showed that fear-related stimuli (snakes and spiders, but also angry faces in earlier experiments) are detected faster than non-threatening ones.

Another example can be found in the emotional stroop tasks used by psychologists [Williams et al., 1996; McKenna and Sharma, 2004; Kunde and Mauer, 2008; Frings et al., 2010]. Typically, two stimuli are presented simultaneously. The task or action to perform is related to the non-emotional stimulus (i.e. press a button corresponding to the color in which a word is written). A delay is noted when the second, co-occurring stimulus (the meaning of the word in the previous example) carries an emotional valence. This phenomenon is called the emotional stroop effect. Although the name is a reference to the classical stroop tasks, the cognitive interference observed in each case is different. Indeed, a Stroop test creates a conflict between incongruent stimuli; typically the color name and the color in which the name is written (e.g. "red" written in green). In contrast, the (not-so-appropriately-called) emotional stroop tasks capture an attentional bias toward emotionally significant stimuli. The term *emotional intrusion* was claimed by McKenna and Sharma [2004] to better describe this phenomena.

Yet another evidence comes from the attentional blink effect. It consists in the impairment of the detection of a target stimulus (T2) when presented rapidly after a first target stimulus (T1). Anderson [2005] suggests that the effect depends on the arousal but not the valence of the emotional stimuli. The attentional blink has also been shown to be modulated by the emotional significance of both T1 and T2 [Schwabe et al., 2011]. That is to say, the emotional relevance of T1 increases the effect while that T2 reduces it. The functioning of locus coeruleus, central

in the noadrenergic system, is thought to account for the attentional blink effect [Nieuwenhuis et al., 2005].

5.2.3   Memory

Studies on working memory – i.e. the ability to maintain information in mind – provide additional support for the influence of emotion on cognition. For example, Perlstein et al. [2002] found that the valence of the face stimuli that participants were asked to memorize modulated the activity of dlPFC; which is thought to be a critical area for working memory. Besides, Gray and colleagues investigated the effect of emotion induction on the ability to memorize word and face stimuli in 3-back tasks [Gray, 2001; Gray et al., 2002]. Interestingly, the results showed that emotional states consistently exerted opposite effects on working memory for verbal versus nonverbal information. More precisely, the performance on the face task was enhanced by a unpleasant state and impaired by a pleasant one, and inversely in the word task. They also showed that the dlPFC neural activity was greater in the word-unpleasant and face-pleasant conditions, intermediate in the neutral conditions, and lower in the word-pleasant and face-unpleasant conditions.

Besides, McGaugh [2015] provides a review of findings indicating that emotion also consolidates long-term memory. For example, low- and high-intensity stimulation of the amygdala respectively impair and enhance memory [Gold et al., 1975]. Also, epinephrine and corticosterone are known to be released in case of emotional arousal – experimentally induced through footshocks for instance. And it has been shown that these adrenal stress hormones produce dose- and time-dependent memory enhancement. In humans, there is evidence that arousing pictures or words tend to be well remembered. Therefore, McGaugh [2000] proposes that the basolateral AM modulates memory consolidation through noradrenergic (NE) projections, which are influenced by the release of stress hormones due to emotional arousal.

5.3   Emotion–cognition interaction

In the examples presented above, either cognitive processing (appraisal) is at the root of emotional elicitation and continuous updating or non-cognitive, emotional processing influences cognition. These works are not in contradiction. As a matter of fact, several researchers attempt to provide an integrated view of the relation between emotion and cognition. For example, Inzlicht et al. [2015] propose that cognitive control can be understood as an emotional process. The authors note that cognitive errors are associated with physiological changes such as increased skin conductance, cardiac activity and pupil dilation, which are also considered as emotional primitives. In addition, they point out that negative affect increases the saliency of goal conflicts and motivates goal-directed behavior in order to minimize the conflict. Although giving an interesting perspective on the emotion–cognition interaction, this view is restricted to "integral emotions" – that is emotions that are directly related to the current task or goal, as opposed to "incidental emotions", secondary and elicited or manipulated by other factors.

Pessoa [2008, 2010] also makes a case against the segregation between emotion and cognition. Based on anatomical and functional evidence such as those I presented earlier, he highlights the interactive and integrative potential that exists in brain structures like the amygdala and prefrontal cortex. He also argues that complex behaviors have their basis in dynamic coalitions of

networks of brain areas involved in both emotion and cognition. Instead of a one-to-one mapping between areas and functions, he proposed the conceptual model represented in Figure 5.2. Brain areas have many-to-many links with different types of neural computations that generate behavior. The behavior space is described using affective and cognitive axes. Thus, any behavior is by definition both cognitive and affective. Importantly, the axes are not orthogonal, such that any change in one of the behavior dimension affects the other. Additionally, specific brain areas belong to several intersecting networks. Therefore, neural computations have to be seen as implemented by the interaction of multiple areas.
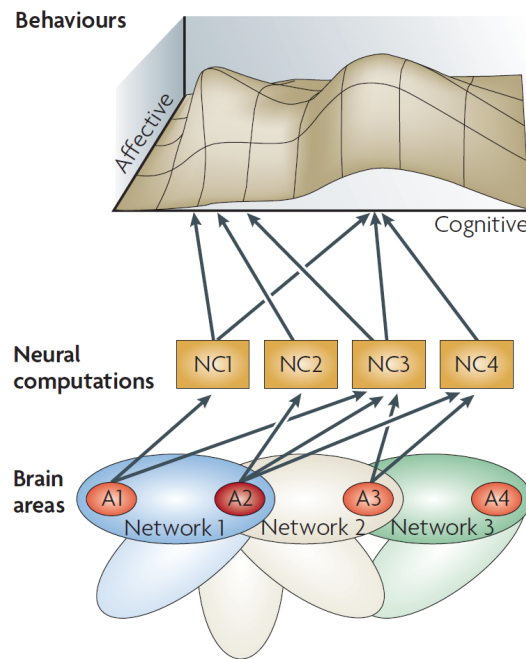

Figure 5.2: Pessoa's model of the mapping between brain areas and emotional–cognitive behaviors. Brain areas are part of networks and have a many-to-many relation with different types of neural computations. These neural computations collectively underlie behaviour. The behavior space is described using non-orthogonal affective and cognitive axes. The figure is borrowed from Pessoa [2008].

In line with this notion of overlapping networks, there is a large literature acknowledging the existence of multiple cortico-basal circuits implementing both motor and non-motor functions [Alexander et al., 1986; Middleton and L, 2001]. The non-motor circuits include emotional/limbic and cognitive/associative loops [Haber, 2008; Jahanshahi et al., 2015]. As shown in 5.3, the former represent interactions with the orbitofrontal and anterior cingulate cortices while the latter concern interactions with the dorsolateral prefrontal cortex. This functional description in the shape of loops encompasses neuronal connections (information transmission) and neurochemical projections (modulation).

Moreover, Fellous [2004] argues that emotions should be understood as dynamical patterns of neuromodulation rather than patterns of neural activity. In fact, two notions are encompassed in this view. First, that emotion and cognition are integrated and implemented by the same structures: the former corresponds to the state of neuromodulation while the latter corresponds to the state of neural computation. This approach radically puts the chemical signaling system at
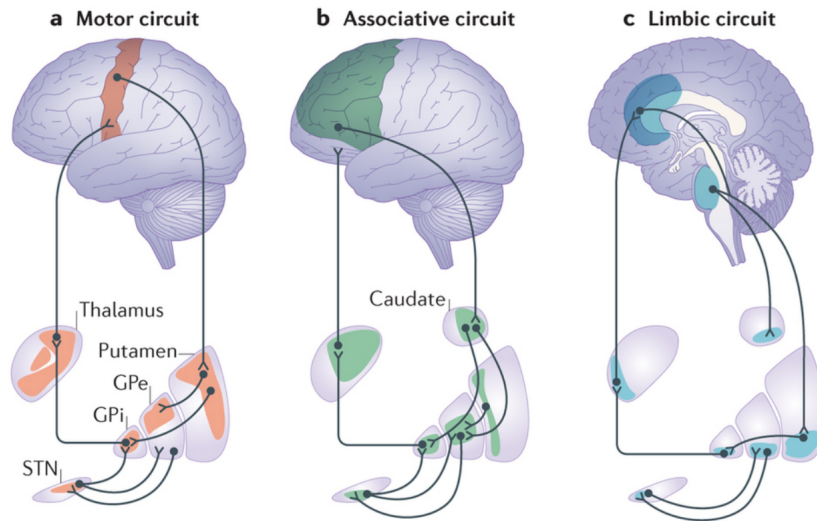
Figure 5.3: Illustration of motor, cognitive/associative and emotional/limbic cortico-basal loops. The figure emphasizes the caudate and putamen of the striatum, which is the largest collection of neurons in the basal ganglia. The cognitive pathway includes projections to the dorsolateral prefrontal cortex. The emotional pathway includes projections to the orbitofrontal and anterior cingulate cortices. All circuits pass through the thalamus. The figure is borrowed from Jahanshahi et al. [2015].

the center of the question of emotion-cognition interaction. The neuromodulatory state is biased by the various levels of computation with different influence potential, as represented in Figure 5.4. In return, the neural activity depends on the aforementioned neuromodulatory state.
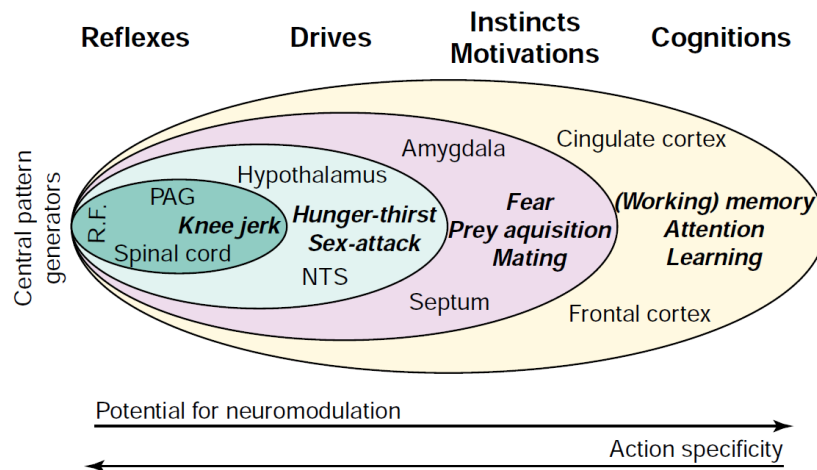
Figure 5.4: Fellous's model of the organization of behavior. Ellipses represent zone of direct influence and neural recruitment during emotional expression and experience. The potential for neuromodulation increases from reflexes to cognition while action specificity decreases. The figure is borrowed from Arbib and Fellous [2004] and merges the organization of behavior with a mapping of the brain structures as originally proposed in two separate figure by Fellous.

The second fundamental notion encompassed in Fellous's view is that of dynamical states. Emotions such as fear and anger that are identified as basic emotions thus correspond to attractors that are temporally and/or spatially stable. This idea is also developed by Lewis [2005]. Indeed, in order to account for the bidirectional relation between cognition (appraisal) and emotion, Lewis proposes to study them through the lens of dynamical systems theory. He builds on a set of principles from this theory to describe internal states as attractors and transitions. In this view, emotional episodes, triggered by perceptual events, physiological events or memories, emerge from appraisal-emotion interactions. The meaningfulness of changes in the world or the body is evaluated by the amygdala, orbitofrontal cortex and nucleus accumbens that initiate new neurochemical patterns. This can be characterized as a phase transition that disrupts the orderliness of the existing state. This results in a self-amplification phase then a self-stabilization phase: positive feedback recruits more neural components to an emerging state and negative feedback couples them in a stabilizing regime

The work by Hasson et al. [2011] in robot emotion is in line with these inspiring models and also illustrates the approach I follow in this thesis. Hasson and colleagues propose to represent an embodied system as two coupled abstract controllers, respectively dedicated to interactions with the physical and social environments. The purpose is not to claim that interactions with the physical and social environments must be handled by separate modules or structures, but rather to put together the processes that are related to the same type of interaction in one abstract entity in order to insist on the interplay between them. Thereby, emotions result from the dynamics
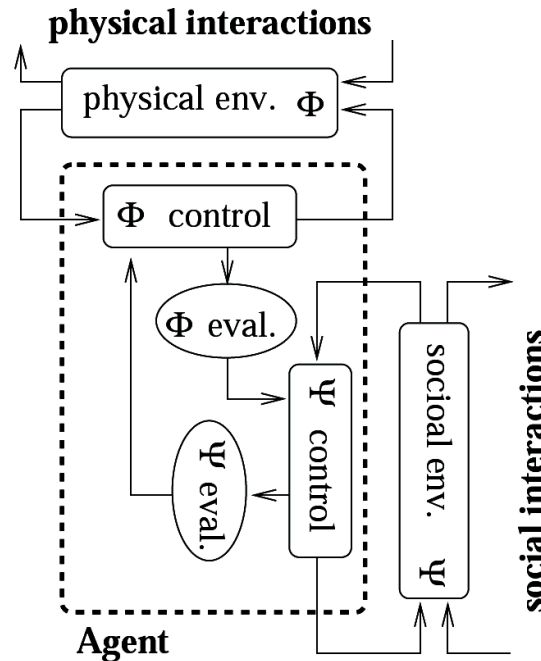
Figure 5.5: Hasson and colleagues' model of the coupling of an embodied agent with its environment. The agent behavior is handled by two coupled abstract controllers, which are respectively dedicated to interactions with the physical and social environments. Emotion emerges from the dynamics of interactions between the two controllers as well as between the agent and the environment. The figure is borrowed from Hasson et al. [2011].

of 1) internal interactions between those two kinds of processes (physical and social) and 2) external interactions with the environment. This idea is represented in Figure 5.5. Emotions are grounded in the whole architecture through the integration with other perceptual, behavioral, attentional and regulatory processes; which are also handled by the two coupled controllers. Thereby, they are not mere responses to external stimuli. They result from the dynamics of the interactions between the two coupled controllers as well as with the physical and social environments. That is to say, for a robot controlled by an artificial neural network like the ones I use in my experiments, emotions should not be merely modeled by the activation of some neurons but rather be read in the network dynamics.

## 5.4   Proposed model of the emotion–cognition interaction

In this section, I will propose an abstract model of the emotion–cognition interaction. Similarly to a workflow diagram, it illustrates interactions between processes rather than structures. The model will then be instantiated in the following chapters of my thesis and neural networks implementing the involved processes will be described in more concrete use cases.

Before I introduce my model, I shall first illustrate what a non-emotional sensorimotor architecture looks like in this kind of representation. As it is shown in Figure 5.6, an embodied system interacts with the environment by receiving sensory input and performing action that influence the forecoming sensations. The information processing flow involves parallel computational processes such as reflexes (pre-wired, evolution-based), memory (temporal integration), conditioning (predictions) and categorization (higher level representations). Representations obtained from the cognitive processing can be reintegrated as input to the information processing flow. For instance, local views from the visual input categorized as landmarks can be reintegrated to encode objects and places. This may also require the integration of additional input – e.g. merging the what and where information. Additionally, among the parallel computational processes are those that allow for shifting to the action space. Thereby, the information processing uses the sensory input (in different levels of representations) to trigger behaviors in the sensorimotor pathway. The actions that are actually performed by the system result from the competition between those sub-behaviors computed in parallel. The competition can be strict or soft, allowing the cooperation between sub-behaviors. To complete the example given so far, this model is instantiated in Figure 5.7 to describe the place–action architecture presented in Chapter 1 and 3.

In line with the literature reviewed above, implementing emotional competence in such an architecture requires (partially) dedicated networks that assess events with regard to survival and well-being. That is to say, capturing physiological and sensory inputs which carry an emotional valence. This is illustrated in Figure 5.8 by the *Valence extraction* block. Like in the works by Scherer [2001, 2010]; Grandjean and Scherer [2008], the term *valence* – in a broad sense – refers to pleasantness, novelty, goal relevance, etc. Such type of processing can be unlearned, evolutionarily acquired in order to regulate bodily functions and handle stimuli with intrinsic affective properties (e.g. auditory or visual stimuli [Koch and Schnitzler, 1997; Öhman et al., 2001]. But the emotional value can also be acquired through pavlovian (stimulus–stimulus) or instrumental (stimulus–reward) learning. The result of this emotional evaluation (appraisal) is the modulation of the computational processes that govern the system behavior. This models the neuromodulatory function of the chemical brain system highlighted earlier [Fellous, 2004] (and see also Chapter 4). Emotion thus influences sensing-related and action-related processes in
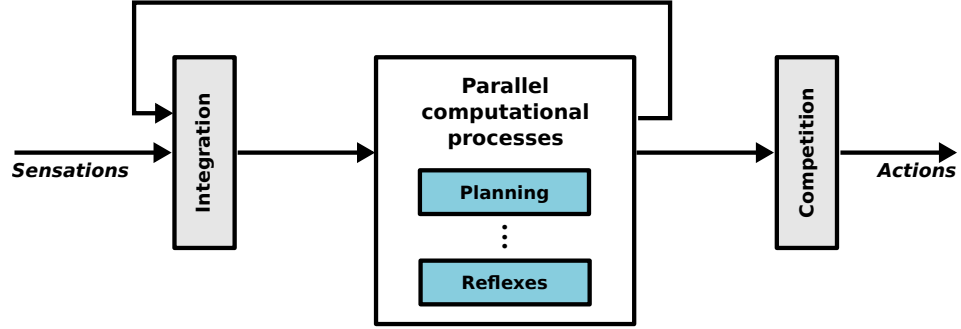
Figure 5.6: Abstract model of sensorimotor architectures. The system, situated in its physical and social environment, constantly processes sensory inputs that are reintegrated in the information processing flow for the purpose of higher order processing. *Integration* refers to the construction of a new representations either by sensory associations or reintegration of processed lower level representations. *Parallel computational processes* include memory, conditioning, categorizations and space shifting (from sensory to motor spaces). Higher order processing can go on to allow for the construction of cognitive maps and goal-driven planning.
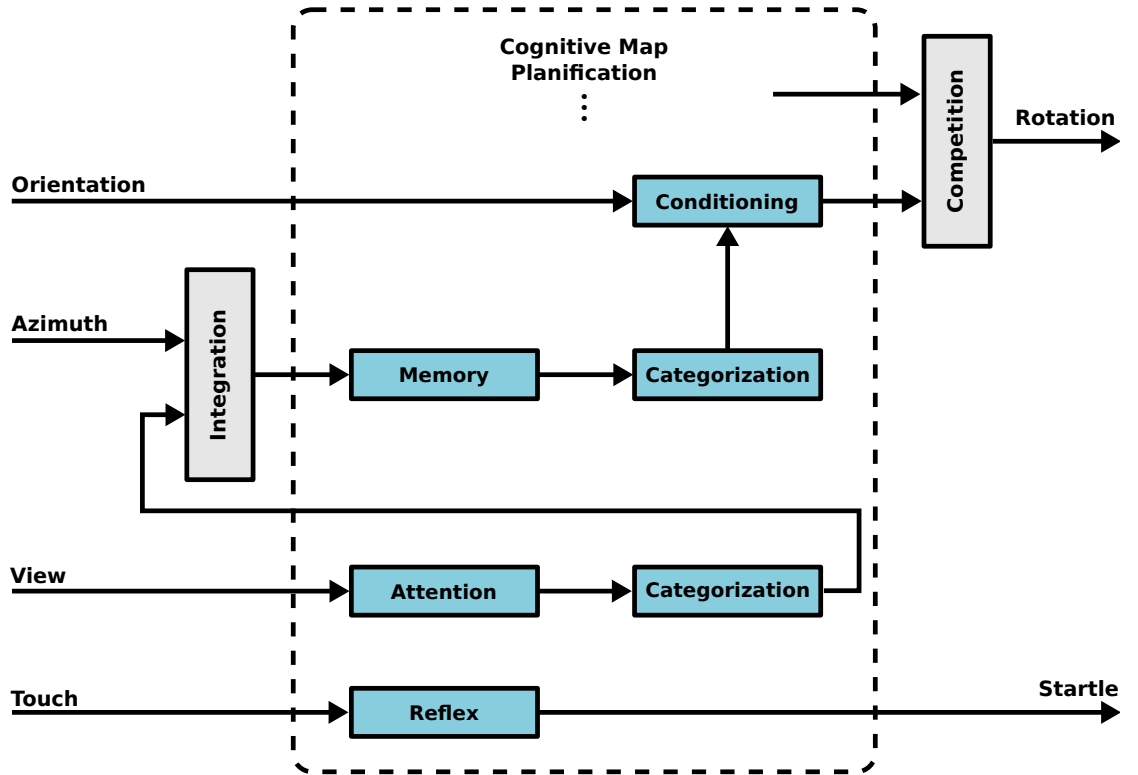
Figure 5.7: Instantiation of the abstract model of sensorimotor architectures. Two behaviors are represented: raw stimulus-driven startle reflex and place/action navigation strategy. The *parallel computational processes* included in this example are reflex, attention, memory, categorization and place–action learning (conditioning). *Integration* here refers to the combination of what-where information: visual landmarks and their azimuths.
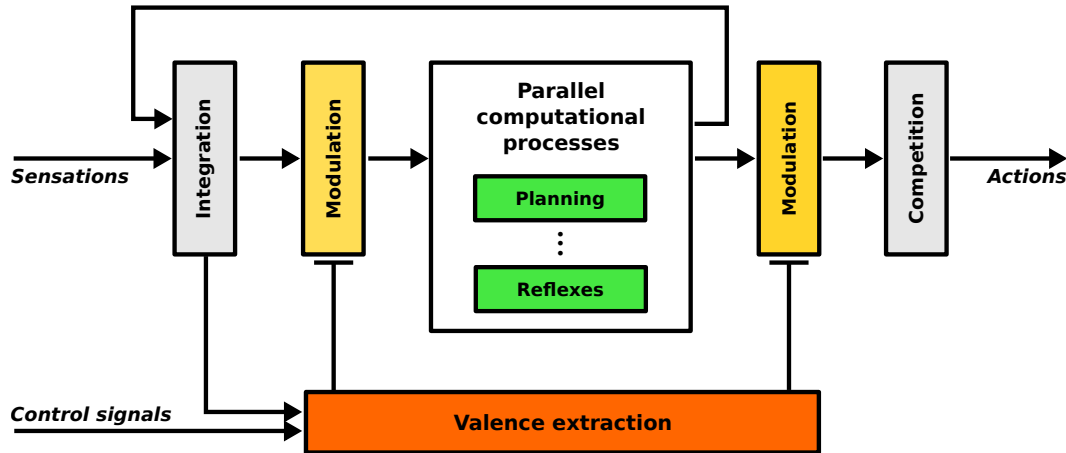
Figure 5.8: Abstract model of emotion–cognition interaction. The system, situated in its physical and social environment, constantly appraises events from the body and the world with a particular interest to emotionally relevant stimuli that affect other processes. It continuously processes emotionally modulated signals and reintegrates them in the information processing flow for the purpose of higher order processing. *Integration* refers to the construction of a new representations either by sensory associations or reintegration of processed lower level representations. *Parallel computational processes* include memory, conditioning, categorizations and space shifting (from sensory to motor spaces). Higher order processing can go on to allow for the construction of cognitive maps and goal-driven planification. *Valence extraction* consists in the evaluation (appraisal) of the emotional values of more or less complex representations.

the information processing flow. For example, when occurring on the sensation space, emotional modulation affects perception and memory. When occurring on the action space, it can modulate attentional responses, action selection, and motor expression. Thereby, the system sensations and actions are no longer neutral and objective, but rather emotionally colored.

5.5 General discussion

In this chapter, I presented evidence of the bi-directional influence between emotion and cognition. These effects are historically studied and theorized separately, taking a particular perspective either departing from cognitive or emotional processes. But more and more researchers advocate an integrated view of this relation. A variety of terms exist in the literature: interplay, interaction, coupling, integration. These terms reflect subtle nuances with regards to the kind and strength of the linkage. However, this again highlights the need for a clearer definition of what belongs to the cognition domain and what belongs to the emotion domain. If most (if not all) of the computational processes are modulated by emotion, what is "pure" cognition ? In this documents, these terms will be used almost interchangeably.

Here, I introduced my approach to emotion–cognition interaction in the form of a conceptual, process-oriented model. In the next two chapters, concrete implementations and experimental results will be presented. First, we will be interested in the perception of the *peripersonal space* when modulated by emotionally valenced sensory and physiological signals. Second, we will address the use of emotional signals elicited by self-assessment to modulate computational processes – such as attention and action selection – for the purpose of behavior regulation.

Emotional modulation of peripersonal space to control robot behavior

Peripersonal space refers to the area around the body that is perceived as secure and reachable. The ability to build such a representation is necessary in both approach and avoidance behaviors. Here, I will argue that building such an action-centered perception of the surrounding space is an efficient way to control the robot behavior in a way that inherently takes emotional information into account.

I previously advocated the integration of artificial emotion with other computational processes. Several studies show that the perception of reachability and comfort distances depends on emotion. Thus, modeling peripersonal space is a good framework to instantiate the conceptual model I proposed in the previous chapter and study interactions between emotional and perceptual processes. This will also allow me to investigate the effect of internal emotional processes on the interaction with the physical (robot–obstacle) and social (robot–robot) environment.

6.1   Peripersonal space and the plastic perception of it

When interacting with the physical and social environment, the perception of the space surrounding the body is of fundamental importance. Two notions are encompassed. First, the area in which an external intrusion can be perceived as possibly threatening, or at least uncomfortable; *the space where one can be*. It has mostly been studied in social psychology, for example to investigate the impact of the familiarity with a person or the impersonality of an environment [Little, 1965]. Numerous designations exist for this concept: 'personal space', 'interpersonal space', 'comfort space' [Little, 1965; Tajadura-Jiménez et al., 2011; Lloyd, 2009]. The other related notion is that of the space within arm distance; *the space where one can act*. It is generally referred to as the 'reachable space' or 'peripersonal space' in the research on sensorimotricity and action-centered perception [Rizzolatti et al., 1997; Lourenco and Longo, 2009]. For instance, studies in neuroscience and cognitive psychology investigate it in the context of body schema remapping in case of tool use [Iriki et al., 1996; Longo and Lourenco, 2006]. The term 'near-space' is also used, in contrast with 'far-space' or 'extrapersonal space' where objects are out of reach.

Although these two notions were historically studied separately and can sometimes be distinguished in the literature, I consider them as two aspects of the same process of integrating information to represent the space around oneself: peripersonal space (PPS) as *the space where one can safely and comfortably interact with the physical and social environment*. In line with this view, a whole body of evidence suggest that comfort and reachable spaces share common properties and rely on the same mental processes. For example, Graziano and Cooke [2006] highlight the defensive behavior inherent to near-space perception during self-motion and object reaching. Also, a study by Lourenco et al. [2011] indicates that people with larger peripersonal space representations report higher rates of (non-clinical) claustrophobia, a state of anxiety in restricted spaces like tunnels, elevators and crowded areas. Additionally, schizophrenic patients, who have deficits in specifying interpersonal distances appropriately, are more inclined to misjudge the boundaries of action space in physical interactions with objects and persons [Delevoye-Turrell et al., 2011]. Iachini et al. [2014, 2016] explicitly tested the relation between interpersonal (social) and reachable (action) spaces. It appears that the perception of comfort and reachability distances is affected by the same factors; in the experiments in question, participant gender, type of virtual stimuli and approach condition.

As an interface between the body and its environment, the representation of the PPS requires multimodal integration. Indeed, interactions between the parietal and premotor cortices seem to be essential in the defensive behavior associated with comfort space. More precisely, it has been shown that neurons in the ventral intraparietal area (VIP) and a polysensory zone (PZ) in the precentral gyrus of the premotor cortex similarly respond to visuotactile stimuli [Rizzolatti et al., 1997; Graziano and Cooke, 2006]. Also, electrical stimulation in both areas leads to similar defensive reactions like blinking, ducking, withdrawing body parts, or protecting the face with the arm. However, Graziano and Cooke [2006] suggest that VIP is rather involved in sensory processing for near-space while PZ is relatively more involved in defensive motor output. Many researchers argue that the perception of PPS does not only rely on sensations but is based on motor-related information as well [Rizzolatti et al., 1997; Coello, 2005; Coello et al., 2008; Lloyd, 2009]. For instance, an experiment based on Transcranial Magnetic Stimulation presented by Coello et al. [2008] showed that the motor cortical area contribute to the conscious judgment of what is reachable.

Several findings demonstrate that the perception of the peripersonal space is plastic. For example, reachability distance increases in case of (active) tool use [Iriki et al., 1996; Longo and Lourenco, 2006] and contracts when additional weights increase the difficulty of actions [Lourenco and Longo, 2009]. Of more interest in the context of my thesis, the representation of PPS is also modulated by emotional factors. Tajadura-Jiménez et al. [2011] showed that a positive affective state induced by pleasant music reduces the comfort distance in over-crowded spaces. Conversely, personal space expands when placed in threatening contexts [Dosey and Meisels, 1969]. On the other hand, positively valenced objects tend to be perceived as closer than negative ones Balcetis and Dunning [2009]; Valdés-Conroy et al. [2012]. Coello et al. [2012] show that there is a significant difference in the perception of distance from dangerous objects (e.g. knife, syringe) as compared to neutral ones only when the former are oriented toward participants. This proves that an on-line evaluation of the harmful consequences of physical interactions alter the perception of peripersonal space, but not the semantic knowledge about the object. Such an emotional evaluation also appears to be required in social interactions, as Kennedy et al. [2009] showed that amygdala lesions impair the estimation of interpersonal space.

6.2 Approach and avoidance: modeling appetitive/aversive pathways

From a constructivist perspective, pleasure and pain are at the basis emotional phenomena [Russell and Barrett, 1999; Damasio, 2003]. Another key component is motivation – in a broad sense, including drives – [Scherer, 2001; Damasio, 2003; Inzlicht et al., 2015]. In Chapter 4, I provided a brief review of the neuronal and chemical substrates that are involved in the processing of these signals. For example, in mammals brain, the brainstem and hypothalamus appear to be crucial. These regions have ascending and descending connexions transmitting information between the body and the brain. They also represent keys centers for the neurochemical system. Dopaminergic, serotonergic, opioidergic, cholinergic and noradrenergic systems interact to compose the rich dynamics of emotional states [Fellous, 2004; Lewis, 2005; Berridge and Kringelbach, 2008].

In this work, I do not aim to provide a detailed model of these complex brain circuitries. However, I am interested in functionally mimicking simple interactions between appetitive and aversive signals to construct dynamic, meaningful internal states. These states can thus be used to drive the robot behavior and modulate its perception of the environment. Indeed, basic motivated behavior in biological organisms can be represented in terms of approach and avoidance. This idea is also predominant in behavior-based robotics where the sub-behaviors that guide the robot actions are generally related to objects, locations or states that have to be reached or avoided.

I propose to define the approach mot_{ap} and avoidance mot_{av} motivation levels at time t as follows:

$$mot_{ap}(t) = [ap(t) + \varepsilon_{mot}.mot_{ap}(t - dt) - \gamma_{mot}.mot_{av}(t - dt)]^{0,1} \qquad (6.1)$$
$$mot_{av}(t) = [av(t) + \varepsilon_{mot}.mot_{av}(t - dt) - \gamma_{mot}.mot_{ap}(t - dt)]^{0,1}$$

where ap and av respectively represent the level of appetitive and aversive motivational drives, ε_{mot} and γ_{mot} respectively represent the integration and inhibition factors of the interaction, and $[x]^{0,1} = 0$ if $x < 0$, 1 if $x > 1$, and x otherwise.

Similarly, a medium-term affective state aff that results from positive and negative states aff_{pn} and aff_{rw} integrating reward and punishment signals rw and pn at time t is obtained as follows:

$$aff(t) = aff_{rw}(t) - aff_{pn}(t) \qquad \text{with} \qquad (6.2)$$
$$aff_{rw}(t) = [rw(t) + \varepsilon_{aff}.aff_{rw}(t - dt) - \gamma_{aff}.aff_{pn}(t - dt)]^{0,1}$$
$$aff_{pn}(t) = [pn(t) + \varepsilon_{aff}.aff_{pn}(t - dt) - \gamma_{aff}.aff_{rw}(t - dt)]^{0,1}$$

where ε_{aff} and γ_{aff} respectively represent the integration and inhibition factors, and $[x]^{0,1} = 0$ if $x < 0$, 1 if $x > 1$, and x otherwise.

The affective and motivational states will be used to modulate the information processing flow in the continuous representation of the robot peripersonal space as described in the next section.

6.3 Proposed model for emotional modulation of peripersonal space

> "The crucial factor modulating the size of personal space, then, could be represented by approach-avoidance actions that reduce or increase this size depending on the social-emotional valence of external stimuli and by motor plans to react to rewarding or threatening objects."
>
> <div align="right">(Iachini et al. [2014])</div>

As seen earlier, peripersonal space is a multimodal sensorimotor interface with the world; a representation that merges action-centered perceptions of the near-space. Influenced by emotion, it is inherently related to both approach and avoidance behaviors. If we consider a mobile robot with two degrees of freedom, we can represent various states of its perception of its peripersonal space like in Figure 6.1. Its comfort zone can contract or dilate according to the pleasantness of the current affective state. Also, appetitive and aversive stimuli respectively induce an extension or a retraction of the reachable space in the corresponding direction.
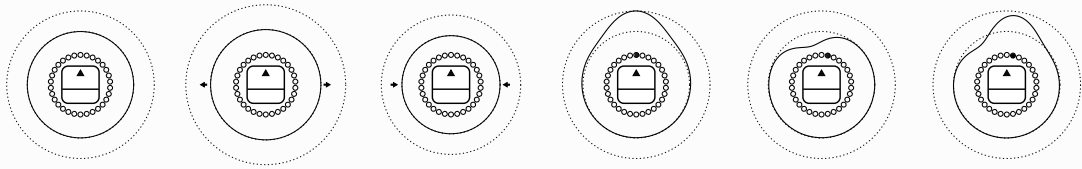
Figure 6.1: Top view of different forms of modulation of robot peripersonal space. FAR-LEFT: No modulation. LEFT and CENTER-LEFT: The comfort zone contracts or dilates according to the pleasantness of the affective state. CENTER-RIGHT, RIGHT, and FAR-RIGHT: Appetitive and aversive stimuli respectively induce an extension or a retraction of the reachable space in the corresponding direction.

Maintaining a representation of the surrounding space relies on memory. For example, while writing these lines, I am perfectly able to reach for the cup of coffee I laid on the table a minute ago while still looking at the screen of my computer. Graziano et al. [1997] reported neural activities in the ventral premotor cortex that keep track of objects located near a monkey's body after the lights are turned off. These neurons continue to fire when, unknown to the monkey, these objects are withdrawn. In this model, I thus propose that peripersonal space perception is based on a working memory that temporally integrates sensorimotor information. For instance, the robot can remember the position of an obstacle it avoided. Also, it can update a path integration vector associated with a goal according to the speed and direction of instantaneous movement. As illustrated in Figure 6.2, information from the working memory is combined in a unique representation that merges action-centered perceptions at any time. Merging comfort (avoidance-related) and reachable (approach-related) spaces, is performed using a dynamic neural field (DNF, described in Appendix and based on [Schöner et al., 1995]). This will be described in more details in the next section.

In line with the conceptual model presented in Chapter 5, these perceptions are emotionally modulated in the sensory and in the action space. The affective and the motivational states respectively modulate the input of the working memory (subjective perception) and the way this information is used to illustrate action tendencies in the representation of the peripersonal space (motivated perception). For example, an appetitive drive makes a desirable object seem more
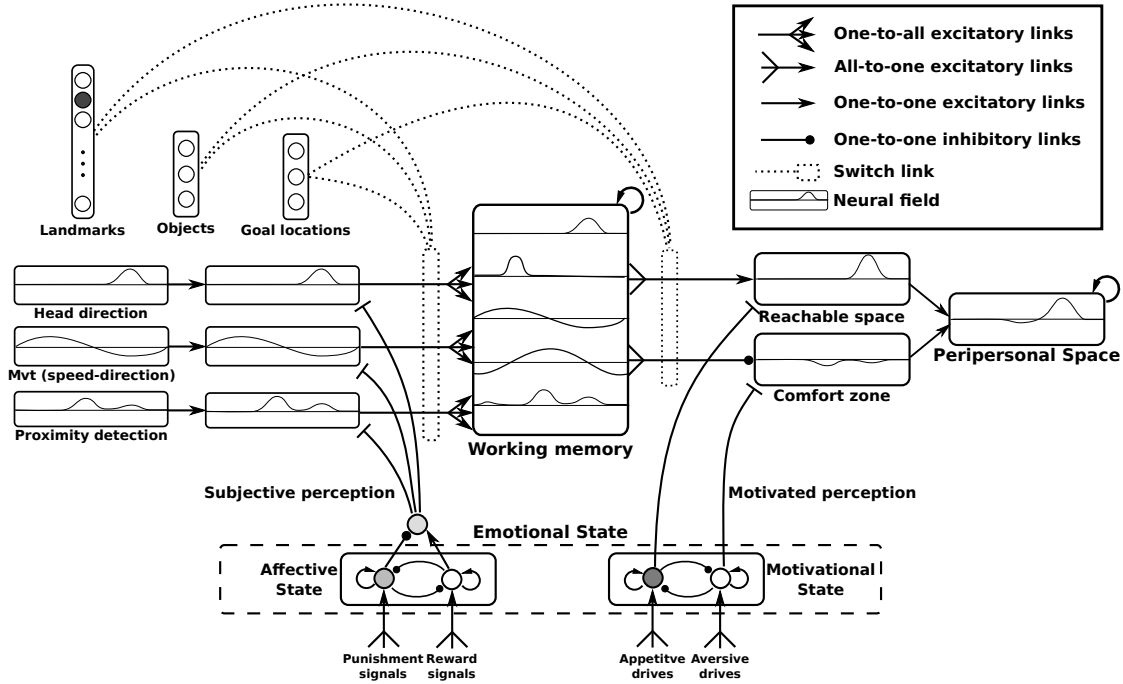
Figure 6.2: Model for building a representation of the robot peripersonal space. It is based on working memory taking input from various sensory modalities. PPS is modulated by the robot emotional states in order to integrate a subjective and motivated perception of its environment.

reachable. Likewise, a defensive motivation highlights aversive stimuli in the comfort space and generates an avoidance behavior.

Hasson and Gaussier [2010] proposed a model of working memory allowing robots to associate several goal locations to the drives they satisfy (e.g. hunger and thirst). Proprioceptive path integration fields are used to return to the resource locations when needed. Although the working memory has limited capacity, the architecture is able to handle multiple goals by replacing the least used memory field when new resources are discovered. Since the implementation of such mechanisms is out of the scope of my thesis, I invite the reader to refer to the principles involved in this architecture for more complex use cases. In this work, the number of sensory input is small enough to bypass these issues.

6.4 Single-resource problem: a survival task

As stated earlier, the representation of emotionally modulated perception of the near-space can be used to control robots approach and avoidance behaviors. In what follows, I present two experiments that illustrate this idea in the context of a survival problem. The robot has two drives: feeding and protecting its own physical integrity. Of particular interest in terms of action selection will be the situations where approach and avoidance are in contradiction. But before I can present these experiments, I shall first describe some implementation details and the experimental method.

6.4.1 Implementation details

To determine the motivational states as described above, let us first consider two survival-related drives: feeding (appetitive) and safety (aversive). In the latter, we calculate the mean activity on the N_s proximity sensors s_i to obtain the level th of threat at time t:

$$th(t) = \frac{\sum_i s_i}{N_s} \tag{6.3}$$

It is possible to draw a parallel with the rodents medial part of the superior colliculus which activation is linked to the detection of predators and other sources of threat (Comoli et al. [2012]). A particular sensory information is thus associated with avoidance without recognition.

The feeding drive is guided by the perceived level of a simulated physiological variable, based on the model of hypothalamus (HTH) proposed in Hasson [2011]. The level fd of the physiological variable associated to the food resource at time t is:

$$fd(t) = \alpha_{fd}.(FD_{max} - fd(t - dt)) \times F(t) - \beta_{fd}.fd(t - dt) \times (1 - F(t)) \tag{6.4}$$

where $FD_{max} = 1$ is the maximal variable level, α_{fd} and β_{fd} respectively indicate the ingestion and the consumption speed factors and F is the food ingestion signal. Using this HTH model gives the robot the ability to anticipate the lack of food in order to trigger the appropriate behavior. For example, let us consider the level of the physiological variable increases and decreases linearly. Considering both functions reach the maxima simultaneously, with the HTH model, the perceived level of the physiological variable drops below the satisfaction level more quickly in the consumption phases. Figure 6.3 shows a comparison between a direct perception of the physiological variable level and the HTH model to illustrate the interest of such a mechanism.
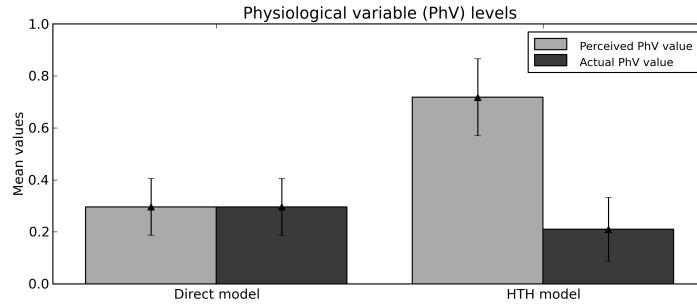
Figure 6.3: Comparison between a direct perception of the level of a physiological variable (PhV) and the HTH model (30 feeding cycles each). We suppose the PhV increases and decreases linearly in our system. The results shows how using the HTH model allows for anticipating the lack of food and prevents from depletion.

When the level of its physiological variable drops below the satiation threshold, the robot has to return to the resource in order to feed. To do so, starting from the resource location, the return vector is calculated by integrating the speed and direction of instantaneous movements. The activity p of each neuron x in the path integration field of size N at time t is given by the following equation:

$$p_x(t) = \Big[p_x(t - dt) + v_{lin}(t)cos\big(d(t) + 2\pi x/N\big) \Big].(1 - Res(t)) \tag{6.5}$$

where v_{lin} the linear speed, d the direction, Res the binary reset signal occurring when the resource is detected. It is worth noting that the mathematical proof provided by Gaussier et al. [2007] requires the application of a cosine-shaped mask but that a Gaussian bell also works in practice.

The feeding drive becomes active whenever the level of the physiological variables drops below a satisfaction threshold f_{thld}. The robot then uses the path integration vector to return to the resource. Similarly, obstacle detection triggers the defensive drive and generates an avoidance behavior. In some case, these two low-level motivations can be contradictory, e.g. if there is an obstacle (object or other robot) on the way to the resource. The dynamic interactions between the appetitive and aversive drives allow them to inhibit each other, which either favors the approach or the avoidance behavior (see equation (6.1)).

To build a combined perception of the comfort (avoidance-related) and reachable (approach-related) spaces, a dynamic neural field (DNF, described in Appendix and based on [Schöner et al., 1995]) is used. Appetitive and aversive stimuli given as input generate attractors and repulsors. Indeed, dynamic neural fields [Amari, 1977] have interesting properties in the sense of dynamical systems (bifurcation, fusion, hysteresis, memory) that allow dynamical filtering (distal competition and proximal cooperation). As a result, the approach and avoidance information is merged into a unique action-centered perception of the peripersonal space. This representation contains two pieces of information that are necessary for the control of robot behavior. The level of the most activated neurons is used to calculate the robot linear speed. Rotational speed is deduced by a readout of the output derived signal with respect to the current orientation. This is explained in more detailed in Appendix A.

6.4.2 Experimental method

To study the impact of the emotional modulation of the peripersonal space on the robot behavior, I compare the full architecture described so far (labeled *Model* version) to two altered versions:

– *NoCompet* version: The robot perception is modulated according to its emotional state. However, there is no lateral inhibition between punishment and reward signals nor between appetitive and aversive drives.

– *NoModul* version: No modulation of approach/avoidance is performed at all. Robot drives only serve for triggering direct objects avoidance and returns to the resource location. This version is the closest to a classical reactive architecture in which approach and avoidance have the same weight.

6.5 Experiment 1: Real robot Vs. static obstacle

In this experiment, I use a customized mobile robot shown in Figure 6.4 based on the Robulab platform by Robosoft. The latter has 9 ultra-sound proximity sensors, of which only a subset –covering a 180 degrees-wide front field – is used. A light sensor is placed underneath to detect red-colored surfaces. As shown in Figure 6.4, a 45 cm x 45 cm red zone in the room simulates a resource. To compare the three versions of the control architecture, a cardboard box is put on the resource location to create a conflict between approach and avoidance. The parameters described above are set to the values given in Table 6.1
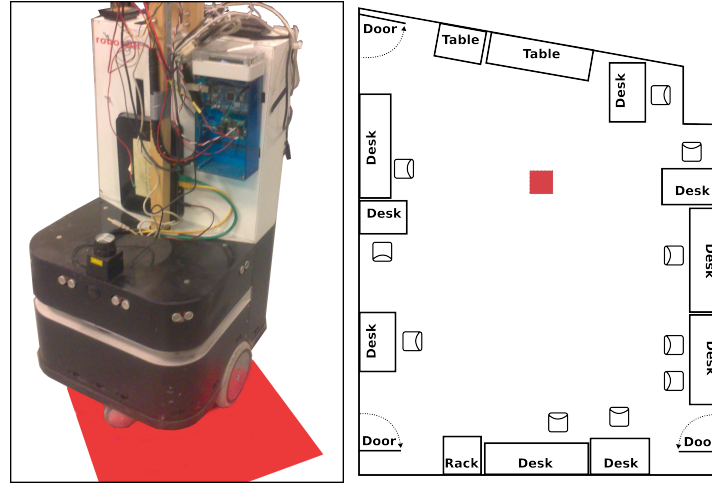
Figure 6.4: Experiment 1 setup and environment. LEFT: The mobile robot platform is 40 cm-wide, 50 cm-long, 140 cm-high, with 2 driving wheels and 2 idler wheels. RIGHT: The experiments are run in one of the rooms of our lab, approximately 8 m long by 6.5 m wide.
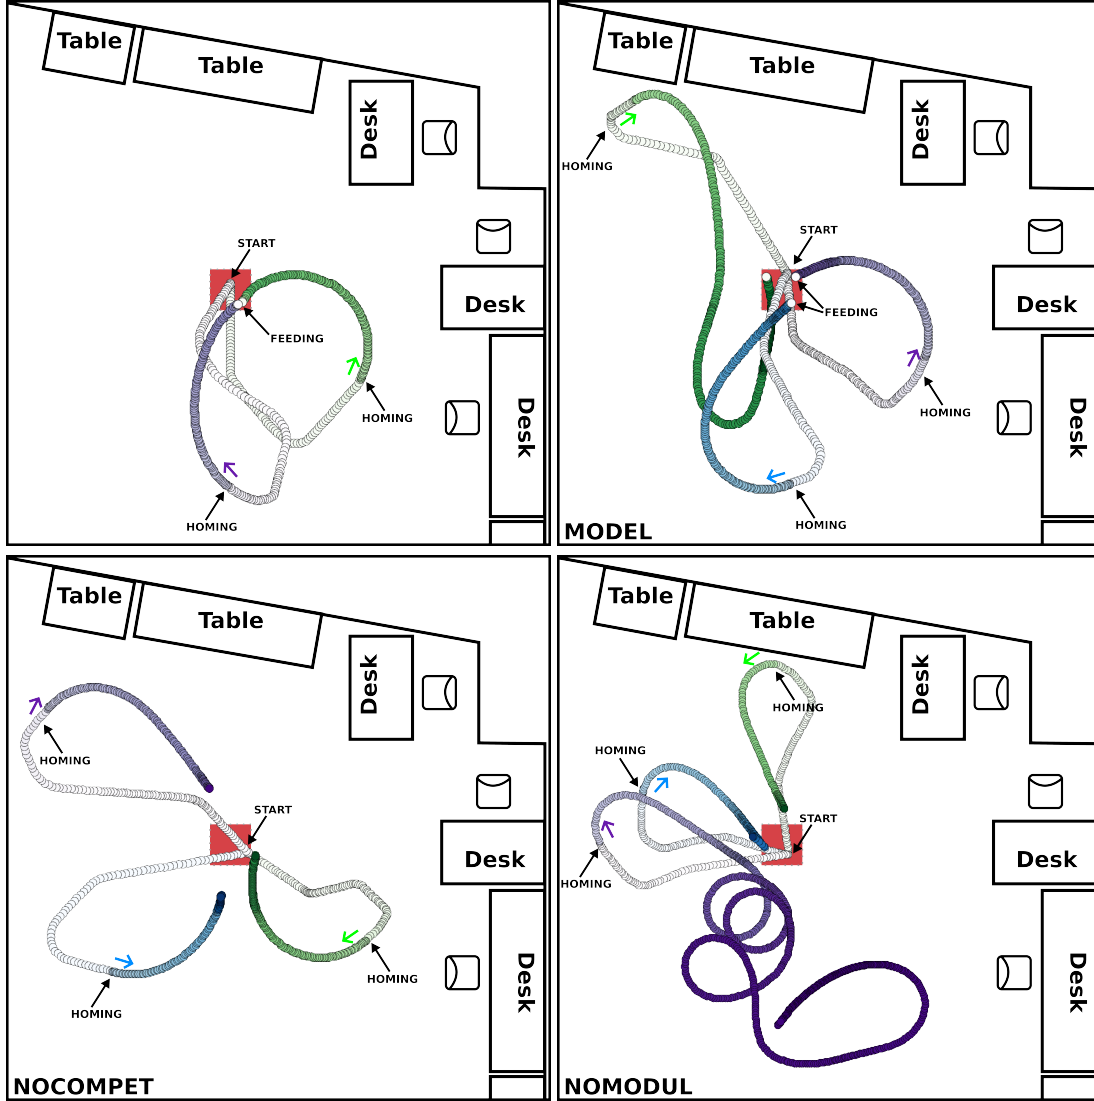
6.5.1 Results

Fig. 6.5 shows multiple trajectories observed in a regular case (no obstacle) as well as when the three architecture versions – *Model*, *NoCompet* and *NoModul* – face a static obstacle put on the resource. In the *Model* case, the robot is able to reach the resource, even after a deviation due to its avoidance behavior. This is allowed by the inhibition of the defensive sub-behavior by the appetitive one when the level of the physiological variable becomes critical. On the other hand, with the two other architectures, the robot fails to feed. Either the defensive drive dominates the robot behavior (*NoCompet* version, and green and blue trajectories with *NoModul* version) or the latter oscillates between approach and avoidance (purple trajectories with *NoModul* version).

6.5.2 Discussion

With the *NoCompet* version, there is an implicit prioritization of the sub-behaviors due to their respective dynamics. Thus, the defensive drive has a bigger impact on robot behavior. The appetitive drive increases more slowly. When, it's high enough it can only cancel the avoidance behavior, making the robot stop in front of the obstacle.

In the *NoModul* case, the robot either stops in front of the obstacle or diverges. Depending on the situation, if the approach and the avoidance fields are perfectly opposed, they cancel each other quite quickly. Otherwise, the system oscillates between the attractor and the distractor respectively created by the appetitive and the aversive sub-behavior. Thus, the robot accumulates errors in the path integration fields and is unable to reach the resource. It is worth noting in the former case, with the *NoCompet* architecture, deadlocks occur farther from the resource in comparison with the *NoModul* architecture. This is due to the robot extending its comfort zone because of negative affective states induced by collision with the box.

On the other hand, with the *Model* version, the approach sub-behavior can inhibit the defensive one when the level of the physiological variable becomes too low. Even if the robot deviates from the resource due to the avoidance mechanism, it is able to access the resource afterwards.

Figure 6.5: Experiment 1 results. **TOP-LEFT**: Robot trajectories in a regular case of returning to the resource. Colored gradients from light to dark represent the decay of the physiological variable. Different colors are used for different trajectories. **TOP-RIGHT**: Trajectories where the robot is able to reach the resource by pushing the obstacle (*Model* version). White dots at the end of trajectories indicate a new food ingestion. **BOTTOM-LEFT**: Trajectories where the robot stops before reaching the resource (obstacle-induced deadlock with *NoCompet* version). **BOTTOM-RIGHT**: Trajectories where the robot stops or diverges before reaching the resource (obstacle-induced deadlock with *NoModul* version).

In conclusion, the model provides an action-centered perception of the peripersonal space that allow efficient control of the robot behavior.

6.6 Experiment 2: Multirobot competition in simulation

This experiment was performed on the Webots simulator in order to avoid damaging real robots. We simulate two robots identical to the one described previously moving in a 17.5 m x 15 m environment. In this experiment, $\alpha_r = 0.05$, meaning the resource ingestion is slower than in the previous experiment in order to observe the competition between the two robots. All other parameters are kept as in experiment 1 (see Table 6.1).

We define a cycle as an interval in which a robot, initially satisfied (non-hungry), consumes the energy obtained from the previous ingestion and returns to the resource in order to feed once again. Each of these cycles is considered as an independent sample of the multirobot competition for the resource. Once its feeding drive satisfied, a robot get away from the resource. It randomly navigates in the environment updating its path integration field to be able to return to the resource when needed. Similarly to the previous experiment, with *NoModul* and *NoCompet* versions, robots are expected to be unable to access the resource until the other is done feeding. However, in the *Model* version, robots should be able to fight for the resource by pushing each other.

To test this hypothesis, three measures are used:
– *min_phyvar*: Lowest level of the physiological variable associated with the feeding drive,
– *nb_own_access*: Number of own accesses to the resource within a full cycle,
– *nb_other_access*: Number of other robot accesses to the resource within own cycle.

The first one is a measure of food deplation, i.e. how close to starvation the robots get. The two latter quantify cycle interruptions. Besides, we consider two variables:
– *version*: Which version is used (each of the 3 versions is associated with a different group),
– *modcomp*: Whether our model is used (the NoCompet and NoModul versions are gathered in one group).

6.6.1 Results

Statistical comparison

I compare the three architectures in 15-minute simulations, which corresponds to a total of N=51 cycles (samples, N_{Model}=16, $N_{NoCompet}$=18, $N_{NoModul}$=17). The means, standard deviation, skewness and kurtosis of the depend variables are the following:
*min_phyvar*: μ=0.68, σ=0.07, skewness = 0.81, kurtosis = 2.54;
nb_own_access: μ=1.92, σ=1.78, skewness = 2.13, kurtosis = 3.82;
nb_other_access: μ=1.47, σ=1.14, skewness = 2.96, kurtosis = 10.89.

Since none of these variables follows a normal distribution so we perform non-parametric tests, the non-parametric Kruskal-Wallis (K-W) test is performed [Corder and Foreman, 2009]. This test shows that there is no effect of *version* nor *modcomp* on *min_phyvar* (resp. $\chi^2 = 0.45, p = 0.80$; and $\chi^2 = 0.00, p = 0.95$). Mean ranks are 25.81 for *Model*, 27.72 for *NoCompet*, and 24.35 for *NoModul*. Regarding *nb_own_access*, there is a strong tendency with *modcomp* ($\chi^2 = 3.81, p = 0.05$) but no significant effect of *version* was found ($\chi^2 = 5.36, p = 0.07$) with a mean rank of 30.94 for *Model*, 21.25 for *NoCompet*, and 26.38 for *NoModul*. However, there is a main effect of both *modcomp* and *version* on *nb_other_access* (resp. $\chi^2 = 6.03, p = 0.01$; and $\chi^2 = 8.19, p = 0.02$). Here, mean ranks are 31.78 for *Model*, 20.61 for *NoCompet*, and 26.26 for *NoModul*.

Let us consider two additional measures *bin(nb_own_access)* and *bin(nb_other_access)* respectively equal to 1 if *nb_own_access* and *nb_other_access* are greater than 1, and 0 otherwise. Indeed, in the case of a perfect alternation of the robots over the resource, each should access it exclusively and only once in every cycle. Any different configuration could correspond to a feeding cycle being interrupted by another robot. In this case, we find a strong tendency on *bin(nb_own_access)* with both *modcomp* and *version* (resp. $\chi^2 = 3.68, p = 0.05$; and $\chi^2 = 6.01, p = 0.05$) as well as a significant effect on *bin(nb_other_access)* (resp. $\chi^2 = 5.19, p = 0.02$; and $\chi^2 = 6.73, p = 0.03$).
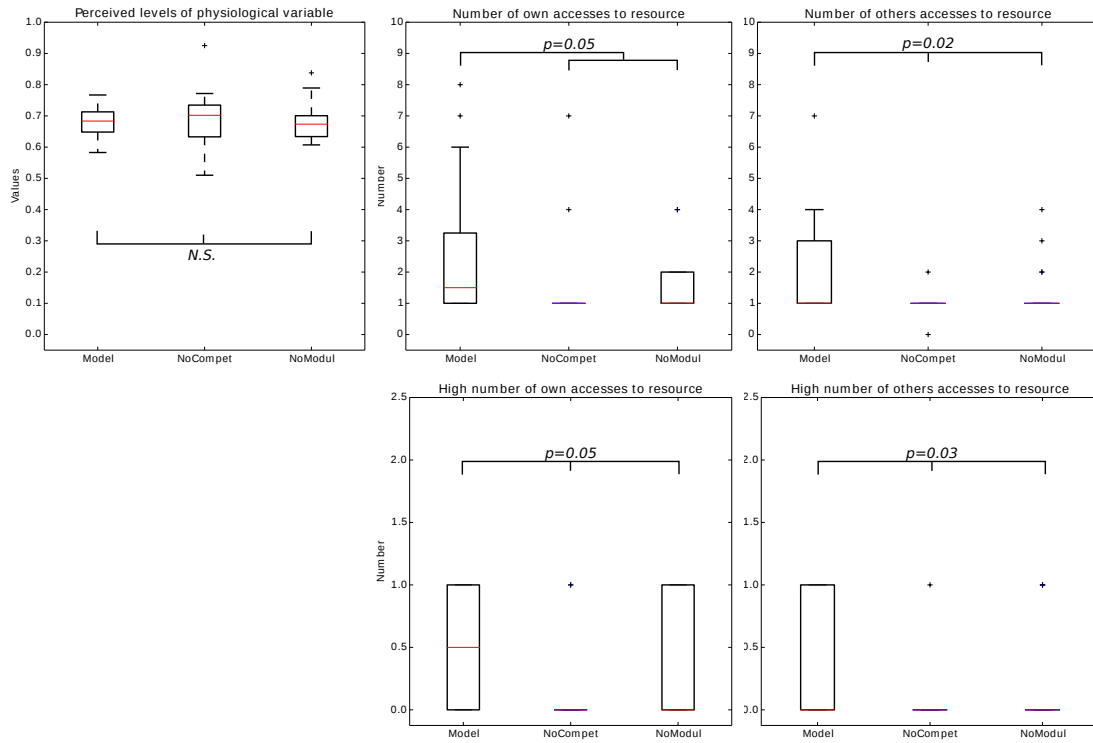
Figure 6.6: Experiment 2, Statistical results. Red lines represent the median and boxes illustrate the two middle quartiles of the data ([25–75%]). Whiskers extend the box ranges by a factor of 1.5. Plus marks represent data that fall out of the whisker range. No effect has been revealed on food depletion (*min_phyvar*). But, a strong tendency is found with *modcomp* variable on *nb_own_access* and with *version* on *bin(nb_own_access)*. There is also a main effect of *version* on *nb_other_access* and *bin(nb_other_access)*.

Behavioral comparison

When the *NoCompet* or *NoModul* architectures are used, the robots tend to be unable to access the resource before it is free – i.e. before the other robot is done feeding. However, with the *Model* version, they can push one another and compete for the resource. This is due to the approach sub-behavior inhibiting the defensive one.

Figure 6.7 shows arousing areas as well as positively and negatively valenced ones in the environment. Arousal and valence are extracted from equations (6.1) and (6.2) as follows:

$valence = aff = aff_{rw} - aff_{pn}$ and $arousal = 1/2(mot_{ap} + mot_{av})$

We notice that in the *NoCompet* case, the area around the resource is one of the most arousing because both drives are simultaneously active. The robots generally need to feed but avoid collisions with the other one currently on the resource. Also, rewards are only obtained right on the center of the resource and punishments around it. However, with the *Model* and *NoModul* architectures, the emotional states are generated exactly in the same way even though, in the latter, they do not modulate robots PPS; and thereby their behavior. We see that most arousing areas are less localized than in the *NoCompet* case. Indeed, the competition between approach and aversion makes them inhibit each other and avoids arousal saturation. But, with the *Model* version, negative valence reaches a lower level than with the two other architectures. This is due to the robot tendency to avoid collisions less than with *NoCompet* and *NoModul*.
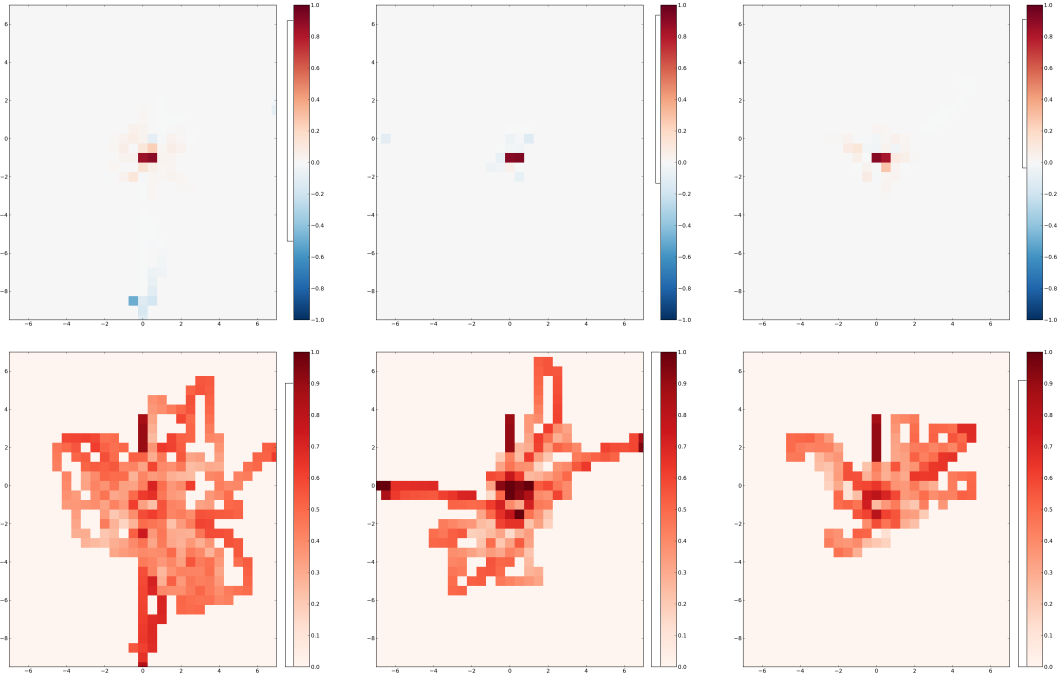
Figure 6.7: Experiment 2, Heatmaps representing arousing (**BOTTOM**) and positively and negatively valenced (**TOP**) areas in the environment. They are averaged for both robots. The brackets on the color-bars show the intervals between min and max values in each case. The columns respectively correspond to *Model*, *NoCompet* and *NoModul*.

6.6.2   Discussion

The results regarding the *min_phyvar* measure show there is no significant difference between the architecture version in terms of food depletion. This is due to the random exploration following feeding phases and to resource consumption being slower that its ingestion. This leaves the possibility for the robot to alternate in resource access. Yet, it is interesting to observe how this alternation occurs, i.e. how robots interact in this survival task. Indeed, there is a social dimension that emerges from the emotional modulation of the robot peripersonal space.

Similarly to Experiment 1 (Sect. 6.5), the robots tend to be unable to access the resource

before it is free when the *NoCompet* or *NoModul* architecture is used. In the *NoModul* case, collisions between the robots occur due the appetitive and aversive behaviors having exactly the same weight. Therefore, some cycle interruptions can happen. Yet, the robot generally deviates from the resource in order to avoid the other robot currently feeding. We argue that interactions between the robots can be seen as emotionally and socially communicative in the *NoCompet* and *Model* versions. With the former, robots seem either patient or fearful. Their modulation of their PPS makes them extend their comfort zone. They are more sensitive to aversive stimuli and defensive sub-behavior tends to take over the appetitive one. On the contrary, using the *Model* version, the robots seem more proactive and determined. When the resource is not available they try to push whatever is on their way. In both architectures, the robots build a representation of their reachable and comfort areas. They also interact in a way that expresses aspects of their internal states. However, when the appetitive and aversive motivations cannot inhibit each another, it can lead to deadlock situations like in Experiment 1 (Sect. 6.5).

6.7 Impact of emotion on vision-based object reaching

The two previous experiments showed the behaviors obtained with the emotional modulation of peripersonal space on a reactive navigation architecture. The environment was dynamic but the sensory information quite basic. Here, I will consider a vision-based object reaching task. A 60 degrees firewire camera is directed toward a platform where the experimenter puts one of the three objects shown in Figure 6.8. One is desirable (toy), one is neutral (cable roll) and one is threatening (snap-off utility knife). The purpose is to compare the objects reachability as it is perceived by the robot depending on their desirability. That is to say, the level of distortion induced by the detection of these objects on the robot reachable space. For the sake of simplicity, there is no actual object reaching using a robotic arm.

These objects are learned by the system in a pre-experiment phase based on a visual system that is similar to the one described in Chapter 1 and 3: i) a saliency map is created using a corner detection algorithm convolving gradient images with a DoG (Difference of Gaussian) filter, ii) additionally, a grayscale-signatures vector classifies the regions of interest encoded by log-polar transformation of region pixels (intensity image), iii) a color-signatures vector classifies the mean color of the regions of interest. The visual signatures are associated to the corresponding signatures through supervised learning (conditioning with the Widrow and Hoff [1960] rule). The system also learns a set of local views that should not be associated to any of the objects in order to be robust to background "noise".

In the experiment, the head direction field is used to integrate visual stimuli – i.e. recognized local views – in the perception of the robot peripersonal space. An approach behavior is simulated by a motivation to reach recognized objects; proportionally to the level of recognition. This will serve as a baseline for the perception of reachability in the following results. Additionally, desirable and threatening objects respectively trigger supplementary appetitive or aversive motivations. These inputs are merged in neural field PPS of size N representing the peripersonal space, like in the previous experiments. Thus, the perception of object reachability can be calculated as follows: *reachability* $= \max\limits_{i=1..N} PPS_i$
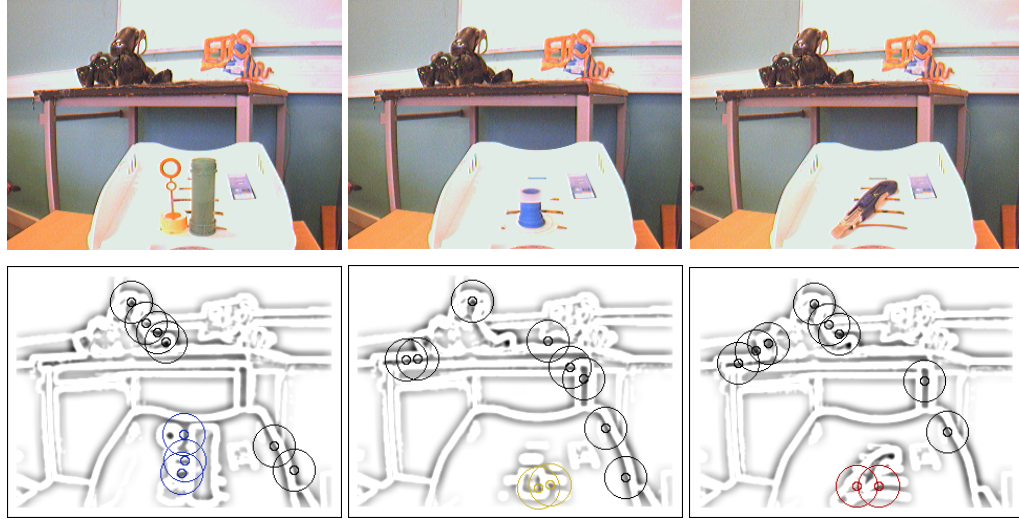
Figure 6.8: Visual object recognition. **TOP**: 3 objects are put in front of the camera, from the left to the right: a toy (positive/desirable), a cable roll (neutral) and a snap-off utility knife (negative/threatening). **BOTTOM**: Gradient image convoluted with DoG used to detect landmark. Local views of the desirable, neutral and threatening objects are recognized. They are represented by blue, yellow and red circles respectively. Several other points of interest are detected on background objects but are not relevant to the task (white circles).

6.7.1 Results

Fig. 6.9 shows how reachability perception varies from one object to another using the three architectures. In our system, the reachability depends on the object recognition level. Therefore, in the following, we will consider reachability perception proportionally to the baseline corresponding to an estimation of reachability only determined by the object recognition level.

Using the *Model* version, the desirable object seems more reachable than the neutral one. Also, the threatening object is perceived as less reachable due to the harm aversion inhibiting the approach behavior. On the contrary, in the *NoCompet* case, the negative one reachability in not altered by its aversive caracteristics since no competition between motivations occurs. However, similarly to *Model*, the positively valenced object is indeed perceived as more reachable than the others. Lastly, the *NoModul* version perceives reachability in a rather binary way because no modulation of PPS is operated.

6.7.2 Discussion

Our model succeeds in accounting for the objects desirability in their reachability perception. Thanks to the emotional modulation and the competition between aversion and approach, positive and negative objects respectively seem more and less reachable than neutral ones. In addition, this experiment shows how visual sensory input can be used in the framework of the proposed model of peripersonal space. For example, in navigation tasks like in Experiments 1 and 2, such additional sensory stimuli can be merged with path integration and proximity detection in order to build robot PPS. Thus, the objects to that the robots needs to reach or to avoid could be recognized visually.
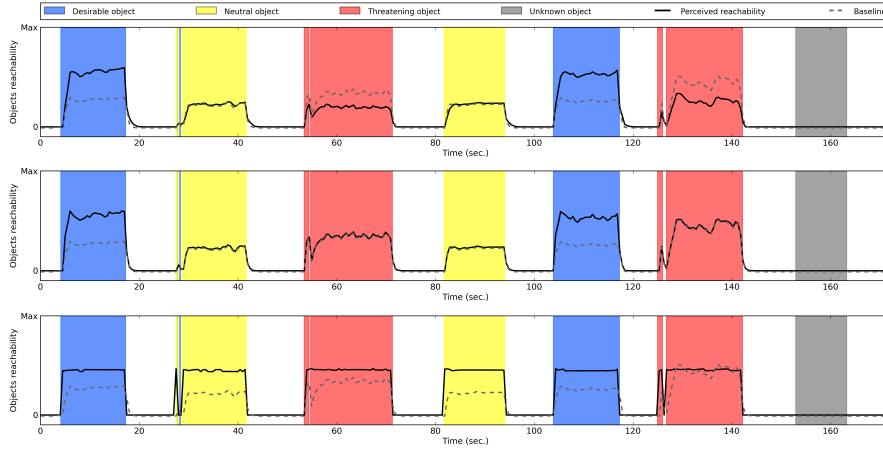
Figure 6.9: Levels of perceived reachability with respect to the three objects. The system perceives reachability depending on the object recognition level. Therefore, the results should be considered in comparison with the baseline (estimation of reachability only determined by the object recognition level, shown in grey dotted line). **TOP**: *Model* version. Positive and negative objects respectively seem more and less reachable than neutral ones. **CENTER**: *NoCompet* version. The desirable object seems more reachable than the other. However, neutral and threatening one reachability are perceived equally. **BOTTOM**: *NoModul* version. Reachability is quite binary, only depending on whether an object is recognized. In all cases, an unknown object is not considered in the reaching task and therefore do not alters the robot peripersonal space.

6.8 General discussion

In this chapter, I instantiated the conceptual model I proposed in Chapter 5 in the case of emotional modulation of the peripersonal space. In line with the related literature [Lourenco et al., 2011; Iachini et al., 2014, 2016], the term of peripersonal space here refers to the process of integrating sensorimotor information in order to represent the space around oneself. Two notions are unified: the personal space needed to feel safe and comfortable, and the reachable space in which it is possible to act on objects. In both cases, there is evidence of the effect of emotions on this representation of the near space [Balcetis and Dunning, 2009; Tajadura-Jiménez et al., 2011; Coello et al., 2012; Valdés-Conroy et al., 2012].

Following a constructivist approach, the robot emotional state emerges from interactions between an appetitive and an aversive pathway using simple pleasure (reward), pain (punishment) and motivation (approach/avoidance) signals. The dynamics of the emotional state influence other computational processes to provide the robot with a subjective and motivated perception of the objects and space around it. This approach was tested in two types of experiments. I demonstrated that the so obtained representation of the peripersonal space is sufficient to control the robot behavior in a survival task. I also showed the emotional modulation of perception has an impact on the way robots interact. With the *NoCompet* and *Model* versions of the architecture, the robots behave in a way that expresses aspects of their internal states. With the former, the robots seem either patient or fearful. Their modulation of their PPS makes them extend their comfort zone. They are more sensitive to aversive stimuli and the defensive sub-behavior tends to take over the appetitive one. On the contrary, using the *Model* version, the robots seem more proactive and determined. When the resource is not available they try to push whatever is on

their way in order to feed. Thereby, they display an aggressive behavior. In a way, these results prove that when emotion is integrated in an interaction loop with cognition, internal dynamics have an effect on external interactions with the physical and social environment. Moreover, in the last experiment, I use the model in a different application context and show that the model accounts for the impact of desirability on the perception of reachability. This is consistent with the literature [Balcetis and Dunning, 2009; Tajadura-Jiménez et al., 2011; Coello et al., 2012; Valdés-Conroy et al., 2012] and demonstrates the genericness of my model, in terms of the type of stimuli and the nature of the tasks.

Neural population coding is a biologically inspired form of information coding [Taube et al., 1990; Bullmore and Sporns, 2009]. I chose to uniformly use it for sensory inputs, working memory and peripersonal space representations. This form of coding could also directly be exploited for action selection and motor control. Moreover, the use of DNF to merge representations of the reachable and comfort space allowed us to take advantage from its fusion and bifurcation properties for the purpose of stable and smooth control

Using the information about a distant "food" resource in the perception of the peripersonal space might seem contradictory with the notion of near-space. Indeed, peripersonal and extrapersonal space representations involve distinct neurological systems [Previc, 1998]. Although the architecture proposed here do not aim to precisely model the brain structures involved in these processes, it would be interesting to more-closely investigate how the far-space information is used to perform actions in the near-space. In fact, in our system, the information about the distant destination is mixed with information about proximal obstacle to deduce actions to perform in the area near the body. And, similarly to the perception of reachable objects, there is also evidence that the desirability of distant locations modulate the perception of their distance [Alter and Balcetis, 2011].

Personal publications

Belkaid, M., Cuperlier, N., and Gaussier, P. (2015b). Emotional modulation of peripersonal space as a way to represent reachable and comfort areas. In *Proceedings of the IEEE International Conference on Intelligent Robots and Systems*, pages 353–359.

Belkaid, M., Cuperlier, N., and Gaussier, P. (2015c). Emotional modulation of peripersonal space impacts the way robots interact. In *Proceedings of the European Conference on Artificial Life*, pages 431–437.

Belkaid, M., Cuperlier, N., and Gaussier, P. (2015a). Emergence of low-level aggressive responses in a reactive goal-directed robotic system. In *Meeting of the International Society of Research on Emotion*.

Table 6.1: Experiments 1 and 2: Parameters values

	Value	Description
ε_{mot}	0.8	Integration factor (motivational state)
γ_{mot}	0.2	Inhibition factor (motivational state)
ε_{aff}	0.2	Integration factor (affective state)
γ_{aff}	0.8	Inhibition factor (affective state)
α_{fd}	0.1	Ingestion speed (Exp 1)
α_{fd}	0.05	Ingestion speed (Exp 2)
β_r	0.01	Consumption speed
f_{thld}	0.5	Satiation threshold

Emotional Metacontrol: top-down regulation based on self-assessment

In this chapter, I will introduce the concept of *Emotional Metacontrol* as an additional example of emotion–cognition interaction. It consists in using emotional signals elicited by self-assessment to modulate other computational processes for the purpose of behavior regulation. By self-assessment, I mean the capacity to evaluate one's own abilities and performance.

In the model I will present, this process is performed on the fly, relying on novelty detection. The so-computed self-assessment may lead to two emotional states – namely frustration or boredom – both signaling the need for top-down regulation of behavior. I will present experiments highlighting the interest of *Emotional Metacontrol* in two different contexts: action selection in a navigation task (Section 7.5) and attention modulation in a visual search task (Section 7.8).

7.1   From novelty to emotion via self-assessment

In humans and animals, emotions are often elicited by extrinsic factors: e.g. fear by threatening objects or agents, joy by pleasant events, disgust by undesirable stimuli. But intrinsic information can also carry an emotional valence. In particular, novelty plays a significant role in emotional processes. It is related to surprise which is one of the basic emotions identified by Ekman and Friesen [1971]. In addition, it is considered by appraisal theorists to be an important factor in the evaluations responsible for emotion elicitation [Scherer, 2009; Grandjean and Peters, 2011].

In artificial intelligent systems, the need to notice changes in the environment was identified long ago. In the pioneering work of Brooks [1986], monitoring changes is one of the levels of competence in the layered control paradigm proposed for behavior-based robotics. There is also a large literature on novelty detection in artificial neural networks and machine learning (see the review by Marsland [2003]). Kohonen and Oja [1976] proposed the first novelty filter based on neural networks. It consist of a pattern recognition model where an input is compared to the learned patterns in order to extract the amount of novelty in it. Besides, self-organized maps provide methods of novelty detection for the purpose of unsupervised learning [Kohonen, 1993]. In the Adaptive Resonance Theory developed by Carpenter and Grossberg [1988], the vigilance threshold determines if an input is novel enough to recruit a new category encoding it.

More recently, Marsland et al. [2005] proposed a novelty filter implemented on mobile robots for the purpose of inspection. A habituation mechanism in the output layer reduces the weights of the most active categories of a clustering layer in order to detect abnormal inputs.

Noticing novelty amounts to identifying differences between current and usual/previous situations. All the above methods filter external novelty using pattern matching. Another way is to assess the ability to predict inputs based on past experience. Researchers in the field of artificial intrinsic motivation have used this idea of measuring prediction errors for the purpose of self-improvement [Schmidhuber, 1991; Oudeyer et al., 2007]. In psychology, this theory distinguishes activities that are undertaken because of their inherent interest and enjoyability from those motivated by a separable outcome [Deci and Ryan, 1975; Csikszentmihalyi, 1991]. Intrinsic motivation is thus implicated in spontaneous exploration, cumulative learning and open-ended development. In the work by Schmidhuber [1991], the system builds a model of the world based on which it predicts the next sensations given current sensations and planned actions. Prediction progress is used as an intrinsic reward for the system. Another artificial curiosity model is the one developed by Oudeyer et al. [2007]. In order to increase the robot learning capacities, a meta-predictor estimates the errors of a module predicting sensorimotor signals. Thereby, the robot engages in actions of increasing complexity, avoiding situation that are either too familiar or too novel.

This notion of "optimal" level of novelty is conceptualized by Csikszentmihalyi [1991] as the flow channel, illustrated in Figure 7.1. In his flow theory, the optimal experience lies in the compatibility between skills and challenges. When people are too competent for a task, they get bored. But when the task is too difficult, people get anxious and frustrated. Moreover, according to Fellous [2004], curiosity and boredom are emotions that decouple stimuli from responses. From a functional perspective, these states foster exploration and allow us to try new things. On the other hand, in the reward-centered model proposed by Rolls [2005], frustration may be produced by the omission of an expected reward, or the termination of a reward . Another way to put is, as Scherer describes it, that it is due to someone or something preventing us from
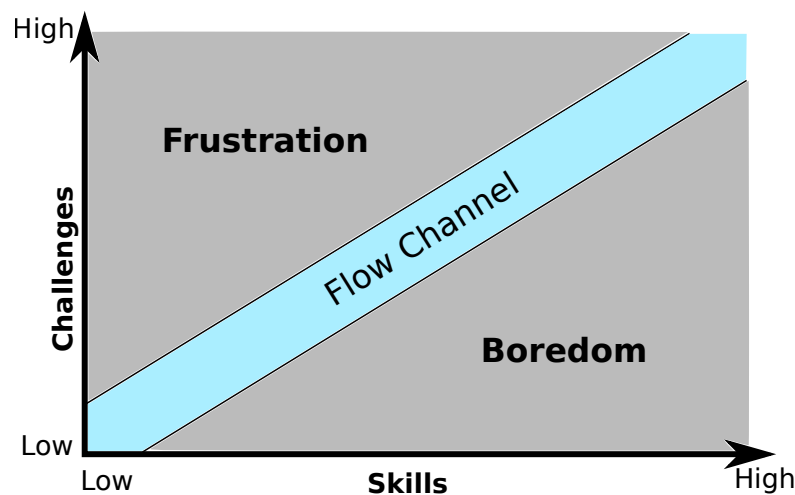
Figure 7.1: Optimal experience. Interesting and enjoyable activities are those maintaining an equilibrium between skills and challenges. Outside the flow channel, people experience either frustration or boredom. The figure is readapted from [Csikszentmihalyi, 1991].

reaching a highly relevant goal [Scherer, 2009].

The purpose of this work is not to give a definition of 'frustration' and 'boredom'; nor to exhaustively enumerate their triggering factors, functions or associated responses. In contrast, I consider the dual, undesirable states characterizing an incompatibility between skills and task demands. These states emerge from a process of self-assessment analyzing the dynamics of novelty detection, that is prediction errors anchored in the agent sensorimotor experience. The terms 'frustration' and 'boredom' turn out to describe these states well [Csikszentmihalyi, 1991]. In addition, as states above, these emotional states may function as signals affecting other computational processes in order to maintain the system in an intermediate state – where it is neither frustrated nor bored.

7.2 Self-assessment in humans

The evaluation of one's own behavior and performance is an important feature of the self-concept; i.e. the internal representation of oneself. People may engage in this activity because of social feedback, out of intellectual curiosity, or due to enforcing events (e.g. failure, illness). Self-evaluation has been studied by psychologists in order to determine which factors motivate these processes and when they operate Sedikides and Strube [1995]; Taylor et al. [1995]; Kwang and Swann [2010]. In this literature, *self-assessment* refers to the desire to have accurate information about oneself; and one's abilities in particular. It is thought to allow individuals to select goals and activities that are in accordance with one's skills. Self-assessment is distinguished from self-enhancement (seeking a more positive self-representation), self-verification (seeking consistency in the self-representation) and self-improvement (seeking the improvement of one's abilities). Taylor et al. [1995] point out that these processes may be elicited by situations of threat/failure or the anticipation of future threat/failure.

Self-evaluation capacities appear quite early in toddlers and progressively develop over the years [Stipek et al., 1992; Kelley et al., 2000]. Three stages are observed. First, children experience success-induced joy and easily switch to other tasks in case of failure. Yet, they do not seem to process evaluation in a self-reflective representation. It is only around the age of two that infants emotionally react to failure (e.g. frowning, calling for help) and express pride or shame. They also start anticipating adult reactions, seeking approbation and positive feedback when succeeding, and avoiding negative reactions (e.g. looking away) when failing. Later, children begin to evaluate and emotionally react to their performance independently of adult reactions. They also show signs of competitiveness as they enjoy winning and completing task as compared to others.

The medial prefrontal cortex plays a key role in social knowledge [Beer et al., 2010; Wagner et al., 2012]. More specifically linked to self-referential cognition, its ventral part (vmPFC) is associated with abstract representations about self. For instance, Moran et al. [2009] asked participants to press a button when they saw green words (oddball items) in a series of sequentially presented words. Also included in the sequence where neutral and biographical items colored in white. The study revealed that vmPFC was activated by self-relevant words (e.g. participants first name or pet name). But self-evaluation also occurs when people evaluate their actions, behaviors, and abilities in the moment. In an experiment conducted by Beer et al. [2006], patients with orbitofrontal (OFC) damage – who are generally known to exhibit objectively inappropriate social behavior – evaluated the appropriateness of behavior after a self-disclosure task

with a stranger, then after watching a video of the interaction. This study suggests an important role of OFC in on-line self-evaluation, associated with the generation of social emotions (e.g. embarrassment) needed for interpersonal behavior. Additionally, Beer et al. [2010] provide supplementary evidence of the role of OFC in on-line self-evaluation, more specifically the suppression of a natural bias toward more positive evaluations (see [Sedikides and Strube, 1995; Kwang and Swann, 2010] on self-enhancement). The results also show the involvement of mPFC in on-line self-evaluation, possibly because participants relied on abstract self-knowledge in the absence of feedback on their task performance.

Less dependent on semantic knowledge, and perhaps more related to my work, is the implication of the anterior cingulate cortex (ACC) in evaluative processes. As mentioned in Chapter 4, ACC is thought to mediate conflict monitoring [Botvinick et al., 2004]. More precisely, conflicts may arise at the level of stimulus and/or response. For example, in the stroop effect, incongruence may be detected with "ineligible" color names (e.g. colors that are not in the set of possible responses) [Milham et al., 2001] or with rectangle-shaped response cues in "eligible" colors [Kim et al., 2011]. But it can also appear when overriding prepotent responses (i.e. responses associated with "eligible" incongruent color names). Put more clearly and making the analogy with my model, conflict – or novelty, in our case – can be detected at a purely sensory level (unexpected stimulus) or at a sensorimotor level (inappropriate response with respect to stimulus, task, or context). Interestingly, ACC activation is observed in both sorts of conflict monitoring [Kim et al., 2011] and in the top-down control of response selection and preparation

7.3 Self-assessment in artificial systems

In the previous section, I mentioned the research works by Schmidhuber [1991] and by Oudeyer et al. [2007] as examples of self-assessment components aiming to foster learning and development in artificial systems. These models pertain to the knowledge-based class of intrinsic motivation models: what is evaluated is the systems ability to build representations and predict the outcome of its actions. A second family of models try to rather focus on competences [Barto et al., 2004; Schembri et al., 2007]. This kind of method is highly influenced by reinforcement learning approaches [Sutton and Barto, 1998]. In the model proposed by Schembri et al. [2007], skills are represented by actor-critic modules, the intrinsic reward being the critic's evaluation error. The reader can refers to [Mirolli and Baldassarre, 2013] for a review of intrinsic motivation models.

Other robotic models take a rather different approach to self-assessment. For example, Wallace [2005] addresses the issue through the lens of the software validation. In the author's view, monitoring the run-time behavior of an intelligent system is a solution to overcome the problem of incomplete specifications. In a symbolic architecture where the agent behavior is determined by a decision-tree organizing, the self-assessment module represents a set of high-level constraints that are used to bias the decision process. Yet another different but related approach is that of Arkin and Ulam [2009]. From the perspective of military applications, the proposed model is aimed at autonomous lethal robot with ethical capabilities. In this context, the self-evaluation module assesses the battle damages. Interestingly, the effect of this monitoring is to trigger moral emotions like guilt and remorse. These internal states are used to limit or cease lethal actions during run-time or to adapt the robot ethical module for an improvement of future behavior.

7.4 Modeling novelty-based self-assessment

The model of self-assessment I describe here is based on novelty detection. For that purpose, the robot has to build an internal representation of the 'normal', 'usual' situations (Grandjean and Peters [2011]). Thereby, it can notice when the current situation is different from what it learned. In order to represent the current sensorimotor experience, I rely on the mathematical definition of the perception proposed by Gaussier et al. [2004]. More specifically, I will consider that perception is the integration of sensations and actions over a sliding time window. Thus, Per is a tensorial product between these two input vectors Sen and Ac with a recurrent link of weight α:

$$Per(t) = \alpha.Per(t - \delta t) + (1 - \alpha)(Sen(t - \delta t) \otimes Ac(t - \delta t)) \tag{7.1}$$

The system uses this neural representation of the current sensorimotor experience to learn how to predict the forthcoming sensations. In other words, the objective is to learn sensation-action contingencies to capture invariants in the sensorimotor behavior. This is done by feeding the Per matrix to a neural layer that implements a classical conditioning by means of the Widrow and Hoff [1960] rule (see Appendix). This gradient descent reduces the squared error between the actual output $\widehat{Sen}$ and the desired output Sen based on the patterns of activities in the Per matrix. Thus, the output estimates the mean of the sensory input $\mu = \mathrm{E}[Sen]$ for a given perceptual state. Thereby, we can compute a prediction error $e^{(k)}$ for this mean ($k = 1$) as well as for higher order moments $\mu^{(k)} = \mathrm{E}[e^{(k-1)}]$ also estimated from Per. For instance, $\mu^{(2)}$ and $\mu^{(3)}$ respectively capture the pseudo-variance and pseudo-skewness using the L1-norm instead of the L2-norm.

$$e^{(k)} = \begin{cases} |Sen - \mu| & \text{if } k = 1 \\ |e^{(k-1)} - \mu^{(k)}| & \text{otherwise} \end{cases} \tag{7.2}$$

The novelty $\mathcal{N}^{(k)}$ at the k^{th} order is:

$$\mathcal{N}^{(k)} = \frac{1}{E} |\sum_{n=1}^{E} e^{(k)}(n)| \tag{7.3}$$

where E is the size of the $e^{(k)}$ vector. Higher orders of novelty estimate higher orders (in the statistical sense) of deviations from the average sensorimotor behavior learned by the model; with an obviously increasing computational cost (See [Jauffret et al., 2013a] for a discussion on this issue). Since, different contexts of application imply different requirements in that matter, Experiments 1 and 2 will use distinct orders of novelty, noted $\mathcal{N}$ in the rest of the equations.

The self-assessment of the system performance must not only rely on the instantaneous novelty level, but also monitor the evolution of such information [Hasson et al., 2011]. Therefore, we compute the progress $\mathcal{P}$ and regress $\mathcal{R}$ as follows:

$$\mathcal{P}(t) = \mathcal{N}(t - \delta t) - \mathcal{N}(t) \tag{7.4}$$
$$\mathcal{R}(t) = \mathcal{N}(t) - \mathcal{N}(t - \delta t) \tag{7.5}$$

These self-assessment signals are derived from the inner dynamics of the system. They also characterize the relation between the system skills/performance and the task demands. I propose

that frustration can be elicited by the integration of these signals in a way that captures regress and stagnation in high level of novelty. On the other hand, boredom can be generated by the absence of novelty and progress. Thus, the levels of frustration $\mathcal{F}(t)$ and boredom $\mathcal{B}(t)$ are:

$$\mathcal{F}(t) = \gamma.\mathcal{F}(t - \delta t) + \beta_{\mathcal{N}}.\mathcal{N}(t) + \beta_{\mathcal{R}}.\mathcal{R}(t) - \beta_{\mathcal{P}}.\mathcal{P}(t) \tag{7.6}$$

$$\mathcal{B}(t) = \gamma.\mathcal{B}(t - \delta t) + (1 - \gamma) - \beta_{\mathcal{N}}.\mathcal{N}(t) - \beta_{\mathcal{P}}.\mathcal{P}(t) \tag{7.7}$$

where γ in an integration factor and β_X is the weight corresponding to the input X.

Like in intrinsic motivation theories Csikszentmihalyi [1991]; Oudeyer et al. [2007], neither of these dual affects represent a positive experience. Instead, the system should seek an equilibrium between its skills and the task demands; a state in which it is neither frustrated nor bored. The concept of *Emotional Metacontrol* that I introduce in this chapter consists in using this cognitive–emotional signals as input to top-down regulatory mechanisms (second order control). More precisely, these feedback loops generate an inhibition potential I that is used to bias the targeted computational processes.

7.5 Monitoring and selecting strategies in a navigation task

The work on strategy monitoring is a collaboration with Adrien Jauffret, a former doctoral student in my laboratory who's research focuses on robot autonomy in navigation. The purpose of this experiment is to use the *Emotional Metacontrol* to monitor and regulate independent sub-behaviors of a mobile robot. Therefore, two complementary navigation strategies are considered. The first one is the place–action strategy described in Chapter 3. It is based on visual recognition of locations (places) to which actions (movement direction) are associated in order to iteratively learn a particular trajectory. In the outdoors, due to the difficulty to find stable and discriminating landmarks, a second type of strategy can be very useful: road following. Likewise, in the indoors, such a strategy can be efficient in long corridors. In any case, proximity detection using infrared sensors ensures obstacle avoidance as a basic sub-behavior that competes/cooperates with the latter strategies.

I will present two use cases for the *Emotional Metacontrol*. The level of frustration associated to a strategy can be used to: A) detect failure in order to call for help, or B) reduce the weight of the corresponding sub-behavior in the competition/cooperation process of action selection.

7.5.1 Implementation details

Various methods for road following can be found in the literature. They can be based on range sensors (e.g. Ma et al. [2000]; Sparbert et al. [2001], respectively radar- and lidar-based), on splines models (e.g. Wang et al. [2004]), or on Gabor filters (e.g. Kong et al. [2009]). These techniques are either inefficient if the road edges are not significant or computationally costly. On the other hand, Jauffret et al. [2013b] proposed a model of road following based on rapid and biologically plausible detection of vanishing points. The method is based on the initial work of Revel [1997]. It consists in a competition between N_{vp} neurons, each encoding the direction of a potential vanishing point on the horizon. Four preferred gradient orientations are defined for each of these N_{vp} points. The activity of a vanishing point neuron is determined

by the total number of pixels where the gradient is aligned with its preferred directions. More implementation details are provided by Jauffret et al. [2013b].

As explained above, the road following and place–action strategies are evaluated by the self-assessment network separately. Consequently, two Per matrices capture the sensorimotor contingencies of each sub-behavior. The Sen vector in equation (7.1) respectively represents the direction of the detected vanishing point and the pattern of activity of place cells. In both cases, Ac is the direction of movement (robot absolute orientation).

Given the variability of the sensorimotor patterns that the model can encounter in this navigation task, the 3rd order novelty was empirically determined as a good trade-off between precision and latency. Thereby, the system should be able to filter unusual situations while being robust to periodical or sporadic events. For the purpose of self-assessment and emotional metacontrol, $\mathcal{N} = \mathcal{N}^{(3)}$ is thus used to elicit a frustration signal $\mathcal{F}$, which generates an inhibition potential I computed as follows:

$$I(t) = \mathcal{F}(t) - \beta_h.h(t) \tag{7.8}$$

where $h(t)$ signals the intervention of the human experimenter. This negative term avoids affect-based inhibition during teaching phases. The parameters described above are set to the values given in Table 7.1

## 7.6 Experiment 1.A: Failure detection and call for help

This experiment is performed on the Webots simulator to avoid the problem of very long robotic experiments. Snapshots of the robot and the environment are shown in Figure 7.2. The simulated platform is similar to the one used in Chapter 6. Here, it is equipped with a pan-tilt camera used for the visual place recognition and a fixed camera used for the vanishing point detection. Also, the full range of proximity sensors is used for obstacle avoidance.

The first part of the experiment consists in learning a round path inside a 17.5 m × 15 m room with the place–action strategy. The robot is trained by a human experimenter in a proscriptive, semi-supervised way: the supervisor only corrects the robot trajectories when the latter escapes from the desired path (Giovannangeli and Gaussier [2010]). The learning of eight place–action couples is sufficient to correctly perform the task. The self-assessment module also learns the sensorimotor contingencies of the place–action strategy as a model of a 'normal' experience/situation; a correct performance of the task. This learning requires three laps in order to be able to predict place recognition with a relatively low and flat level of novelty.

In the second part of the experiment, the experimenter deviates the robot away from the learned path and inside a second room. In this 10 m × 10 m unknown room, there is no prior learning. Therefore, this situation is expected to generate a high level of novelty and be detected as a failure of the place–action strategy.

7.6.1 Results

Figure 7.4 shows the results of the self-assessment of the place–action strategy. After the system learned to correctly perform the desired trajectory, the experimenter deviates the robot toward the second, unknown room. The robot tries to perform the task by using the learned place–action associations. However, the previous learning fails in this novel context. Thus, the incapacity to predict place recognition based on previous experience causes the level of frustration to increase
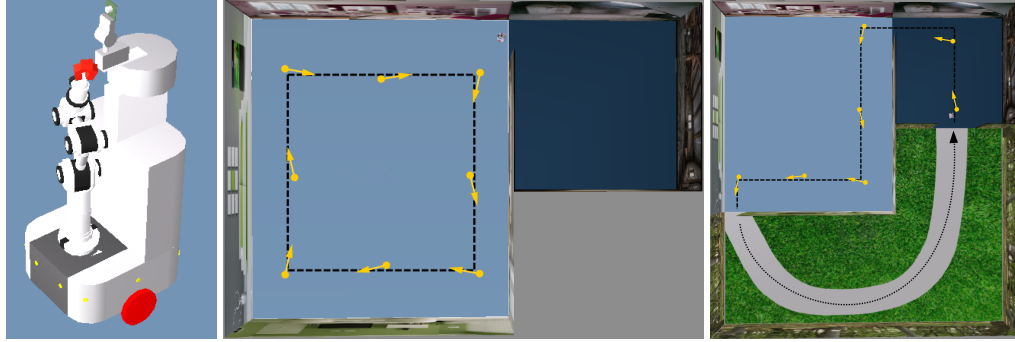
Figure 7.2: Experiment 1, Setup. **LEFT**: The simulated Robulab platform is 40 cm-wide, 40 cm-long, with 2 driving wheels. **CENTER**: The experimental setup in experiment 1.A comprises two rooms. One is used to learn a trajectory using the place–action strategy and the other to test the ability to detect the failure of this strategy in an unknown location. **RIGHT**: The experimental setup in experiment 1.B comprises the same two rooms and an additional path connecting them from the outdoors. In the indoors, place–action association determine the robot trajectory while road following is suitable in the garden.

and reach the threshold. The robot calls for help so that the human experiment can teach it how to escape from the unknown room. This intervention triggers the learning of new place–action couples as well as sensorimotor contingencies for the purpose of novelty detection. Back in the first room, the robot converges again to a 'normal' behavior consisting in the learned trajectory.

Additionally, when put back in the second room, the robot is again able to escape and perform the round path in the first room. Since the escape situation is not encountered enough times, the self-assessment module does not perfectly predict place recognition but the level of frustration stayed below the threshold.

7.6.2 Discussion

In this experiment, the system succeeds in detecting failure and asking for the help of the human experimenter. Once it is back to the initial room, the robot is able to perform the navigation trajectory learned in the beginning of the experiment. One can note that the robot exhibits some sort of "hesitation" when escaping again from the second room due to the door that activates obstacle avoidance (narrow passage). However, it successfully reaches the initial room and resumes its round path. As explained in Chapter 3, to build a robust attraction basin, more than one place–action couple is required. So, additional learning could improve the escape trajectory. But also the self-assessment. Indeed, the robot would be able to better predict the transition between the two rooms and learn this situation as 'normal'. Thus, it will become part of its past experience, on which novelty detection is based.

Besides, this experiment provides another example of how internal processes of emotion can also benefit to the communication aspect. Being able to detect situations of failure is a very useful feature. But it requires other mechanisms that can allow the use of this information for the purpose of problem solving. Thanks to *Emotional Metacontrol*, the emotionally influenced behavior displayed by the robot (i.e. slowing down, stopping, and signaling helplessness) provides social cues to the human experimenter. Such interactions help the robot learn from the latter so that it is able to autonomously solve problem in the future.
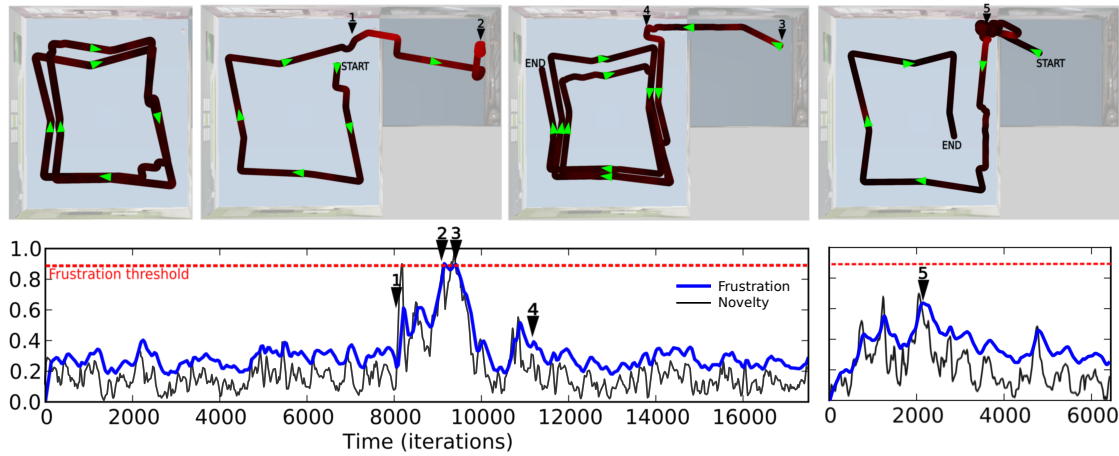
Figure 7.3: Experiment 1.A, Results. **TOP**: Robot trajectories inside the simulated environment in various phases of the experiment. Green arrows show the movement direction. Colored gradients from black to red indicate the increase of frustration. **BOTTOM**: Novelty and frustration levels during these phases. Red-dotted line represents the frustration threshold. **TOP-FAR-LEFT**: Robust performance of the learned round path. **TOP-LEFT**: The experimenter deviates the robot from this trajectory to the second, unknown room (marked 1 on figure), which increases the level of frustration until the threshold is reached. In this situation of failure, the robot calls for help (marked 2). **TOP-RIGHT**: The experimenter teaches the right direction toward the known room (marked 3) and the robot is able to perform the round path again (marked). **TOP-FAR-RIGHT**: The robot is placed in the second room second time and is again able to go back to the learned trajectory (marked 5).

7.7 Experiment 1.B: Competition/cooperation between strategies

As shown in Figure 7.2, in this experiment, the Webots environment is extended with a garden in which a road links the two rooms from the outdoors. Here, the robot has to perform the entire loop using the two strategies presented earlier. Due to the environment specificities, place–action association should perform well in the indoors as seen in the previous experiment, but the road following strategy is more suitable in the outdoor part. This is due to the absence of discriminating visual landmarks in this part.

Like in the previous experiment, the first part consists in learning the indoor half of the trajectory using the place–action strategy. The robot is trained in the same way by a human experimenter and seven place–action couples are necessary to correctly perform this part of the task. However, as stated earlier, no learning is required for road following. But, for the purpose of self-assessment, sensorimotor contingencies are learned for both strategies.

Based on that, in the second part of the experiment, the robot is expected to correctly perform the whole round path, switching from a strategy to another depending on the type of environment. Indeed, place–action and road following strategies should respectively fail in the outdoors (absence of learned landmarks) and the indoors (inconsistency between movements and direction of vanishing points). The inhibition of the failing strategy should allow the other to drive the robot behavior in order to perform the task.

7.7.1 Additional implementation details

Similarly to the previous chapter, a competition/cooperation process for action selection is implemented using a dynamic neural field (DNF, described in Appendix and based on Schöner et al. [1995]). The place–action and the road following sub-behavior independently propose an orientation. Each strategy also has an associated affect-based inhibition signal due to the self-assessment process. As the name suggests, this signal inhibits the corresponding strategy proportionally to the associated frustration level so that it has less influence on the robot global behavior. As in the previous experiment, obstacle avoidance sub-behavior participates in action selection to avoid collisions. The DNF merges the obtained weighted sub-behavior in way that make proximate and distant orientations respectively activate or inhibit each other (respectively, cooperate and compete).

7.7.2 Results

Figure 7.4 shows the results of emotionally modulated competition/cooperation between strategies. First, the separate evaluation of this mechanism on each strategy confirms the expected behavior: due to the environment specificities, place–action and road following perform better in the indoors and outdoors respectively. Then, for the performance of the complete round path, the robot starts in the middle of the outdoor path. The results show that, most of the time, the frustration levels of the two strategies are in phase opposition. Thus, the global behavior is mostly driven by the place–action association in the indoors and by the road following in the outdoors.

7.7.3 Discussion

As it can be seen in the robot trajectory, the sharp turn at the exit the big room makes the transition to the garden quite difficult to achieve. Due to the last place–action couple taught by the experimenter, the most salient vanishing point at that location corresponds to the bottom-left corner of the environment. The resulting oscillation between the place–action and road following sub-behaviors makes the robot go in circles until it converges again to the right path. It is worth noting that despite this imperfection, and the resulting lack of robustness in this particular learned trajectory, the robot should always be able to find the path and perform the complete task. Indeed, place–action make it go down and wall avoidance make it go right. Yet, this gives an example of situations that are challenging for the current system of competition/cooperation between strategies.

Additionally, the experiment illustrates some interesting synergies between the two strategies. For example, the contrast induced by open doors generates vanishing points that are coherent with the orientation learned with place–action associations. Likewise, vanishing points inside the rooms generally direct the robot toward corners, thus increasing the robustness of the trajectory between learned places. This demonstrates the interest of cooperation in the soft competition implementing action selection in the architecture (see Appendix). Yet, it is worth noting that such a property might be counterproductive in other situation and that further investigation in that matter is required.
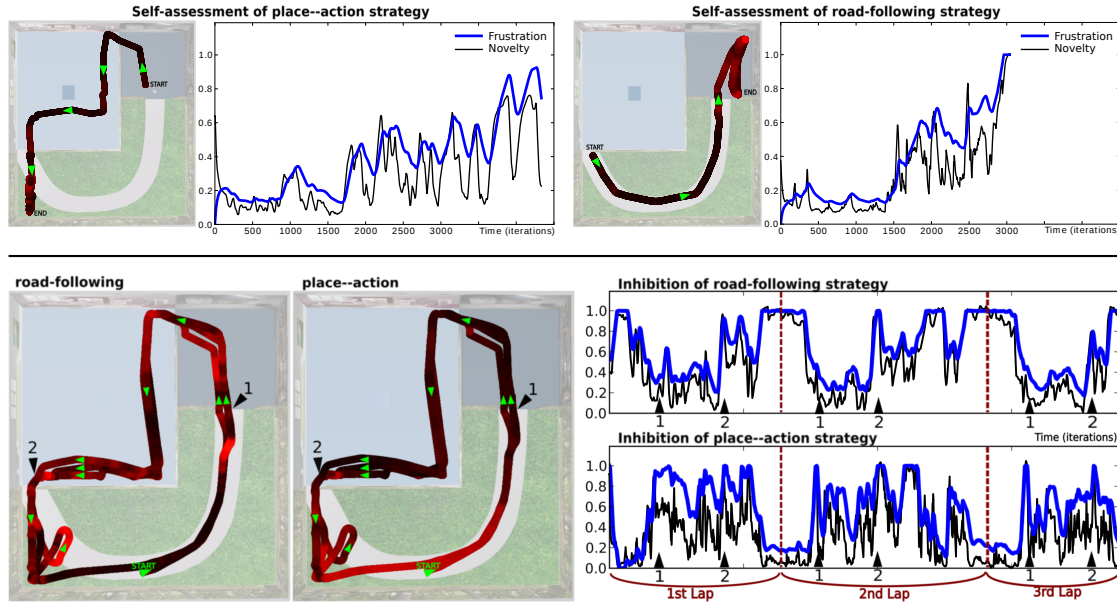
Figure 7.4: Experiment 1.B, Results. **TOP-LEFT**: Separate evaluation of the place–action strategy. Robot trajectory in the indoor half of the environment follows the learned path. Green arrows show the movement direction. When the robot gets out of the big room, the self-evaluation module detects the failure of the strategy. Colored gradients from black to red indicate the increase of frustration. The level of frustration reaches the threshold at the end of this unitary test. **TOP-RIGHT**: Separate evaluation of the road following strategy. Robot trajectory in the outdoor half of the environment follows the road. But when the robot enters the small room, this strategy fails to perform the correct trajectory and the level of frustration in high once again. **BOTTOM**: Strategies competition/cooperation in the performance of the whole round path. The performance of the two strategies is almost in phase opposition. The behavior is mostly driven by the place–action association in the indoors (from mark 1 to 2) and by the road following in the outdoors (from mark 2 to 1).

7.8 Top-down attentional bias in an visual search task

In the previous sections, I introduced the notion of *Emotional Metacontrol* and showed the impact of frustration on the monitoring and selection of strategies in the context of navigation tasks. The rest of this chapter includes an additional (dual) affect: boredom. Also, in order to demonstrate the genericness of the model, the next experiment addresses another context of application: top-down emotional bias of the robot visual attention during active object perception.

More specifically, I consider a visual search task, which is a common experimental paradigm in psychology. This type of perceptual task involves attentional mechanism that allow to scan the visual environment in order to find a target object among distractors. Here, I will use the three objects shown in Figure 7.5: Target1 (Dalek), Target2 (Minnie) and Distractor (Darth Maul). The goal of the task is to search for (recognize) as many target objects as available in the configuration and perform the corresponding (learned) actions to confirm the recognition.

The protocol includes two phases. First, during task learning, the three objects are presented one by one in front of the robot. The latter learns a set of local views of the objects for the purpose of visual recognition. It also associates an action with each of the target objects. For instance:

Target1 $\rightarrow$ Action1 = move arm to the right;

Target2 $\rightarrow$ Action2 = move arm to the left.

Second, the experiment consists of a set of trials in which the objects are presented pairwise; interchangeably on the left or right position. There are three configurations illustrated in Figure 7.5: T1D = (Target1 + Distractor), T2D = (Target2 + Distractor), T1T2 = (Target1 + Target2). The robot performs one of the learned actions to show the experimenter that one of the target objects was recognized.
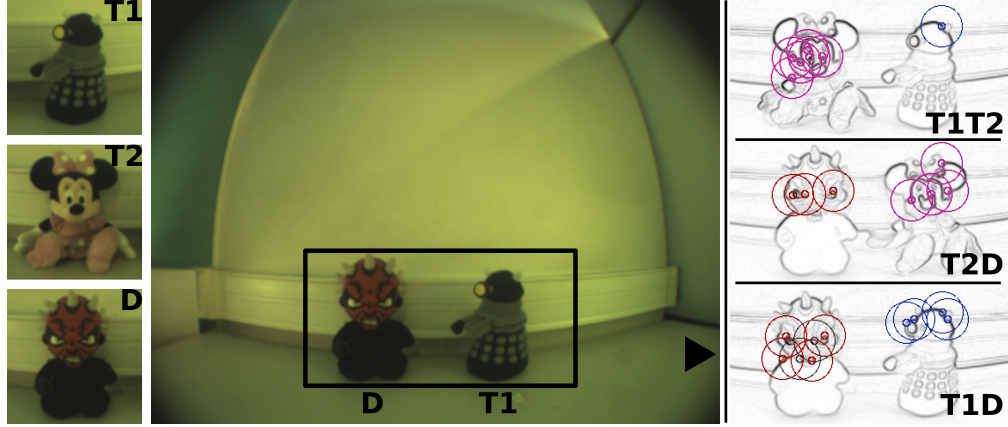

Figure 7.5: Objects and configurations. **LEFT**: RGB images of the objects T1, T2 and D that are used in the experiment. **CENTER**: RGB image of the T1D configuration. Points of interest are detected inside of the rectangular region in the middle. Luminance and chrominance information are useful to discriminate the objects. **RIGHT**: Gradient images in the three configurations (T1D, T2D and T1T2 from the bottom to the top). The more salient the objects, the more points of interest they gather with the bottom-up attentional mechanism. Colored circles show the object to which each local view is associated. Blue is for Target1, purple for Target2 and red for Distractor. Given the visual system used in this experiment, Target2 is more salient than the Distractor, but Target1 is not. Which means that based solely on the bottom-up attentional mechanism, Target2 would be recognized in both the T2D and the T1T2 configurations but the system would not be able to focus its attention on Target1 in any of the T1D or T1T2 configurations.

7.8.1 Implementation details

Figure 7.6 gives an overview of the whole active perception architecture. Local views are extracted according to a saliency map S' combining bottom-up saliency and top-down inhibition induced by the *Emotional Metacontrol*. The bottom-up saliency map S is obtained by means of a corner detection algorithm applied on the grayscale image:

$$S = \frac{G_x^2 <G_y^2> + G_y^2 <G_x^2> - 2G_xG_y <G_xG_y>}{<G_x^2> + <G_y^2>} \tag{7.9}$$

where G_x and G_y are the oriented gradient images given by the convolution of the grayscale input image with the following vectors: $\begin{bmatrix} -1 & 0 & 1 \end{bmatrix}$ and $\begin{bmatrix} -1 & 0 & 1 \end{bmatrix}^T$, and $<G>$ the output
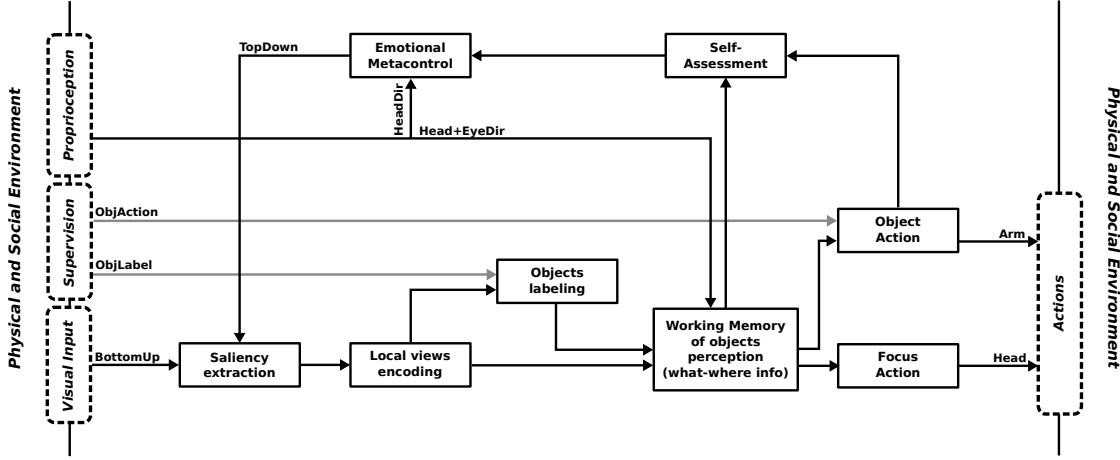
Figure 7.6: Overview of the architecture used for object perception. This figure shows the information flow between all the modules involved. The grey links are only used in the pre-task phases i.e. during objects and task learning. Object recognition is based on a bottom-up attentional mechanism that sequentially extracts local views from visual input. The objects labeling performed during task learning allows for sorting localview–position activities by objects (what-where information) during the continuous exploration of the visual scene in order to store them in the Working Memory. Most salient objects generate head movements toward them as an active focus of attention (FocusAction). When a target object object is recognized, the corresponding action is performed (ObjectAction). The metacontrol mechanism serves as a top-down attentional bias.

of the convolution of an image G with the following mask $\begin{bmatrix} 1 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 1 \end{bmatrix}$.

This method for bottom-up saliency is more adapted to this experiment than the one used is the LPMP model (see Chapter 2 – obtained with a Deriche [1987] filter and a convolution with a DoG. First, the sharp gradients capture textures on the objects. Second, the obtained points of interest are more often located on the objects, which allows local view to generally contain more information on the object than the background.

On the other hand, the top-down attentional bias consists in applying an inhibition of magnitude I on the bottom-up map S. Thereby, we obtain the final saliency map S':

$$S' = S - M.I \tag{7.10}$$

where I is the inhibition potential obtained with equation (7.6) and M an image containing the shape and position of the inhibition:

$$M(x,y) = \begin{cases} 1 & \text{if } x \in [p_x - 1/2.m_{size}, p_x + 1/2.m_{size}] \\ & \text{and } y \in [p_y - 1/2.m_{size}, p_y + 1/2.m_{size}] \\ 0 & \text{otherwise} \end{cases} \tag{7.11}$$

where m_{size} is the side of the square-shaped mask centered on the pixel (p_x, p_y). p_x is obtained from the proprioception merging the head and eye direction while p_y is a constant depending on the elevation of the objects in front of the camera.

The result is a set of N_{PoI} points of interest on the most salient parts of the visual scene. Similarly to the LPMP model I described in chapters 1 and 2, local views are extracted around these salient points between two disks of radius r_{small} and r_{big} (see Figure 7.5 for an illustration). The so defined regions of interest are encoded using a log-polar transformation on the luminance and chrominance channels of the *Lab* color space. Thereby, we obtain three intermediate descriptors of size $W_{LP} \times H_{LP}$. For the sake of conciseness, let us define the final descriptor d of size $N_D = 3 \times W_{LP} \times H_{LP}$ as the concatenation of these three intermediate descriptors. The interest of the log-polar transformation lies in its relatively small computational cost and its robustness to small rotations and scale variations. It is also a biologically plausible operation [Schwartz, 1977]. Besides, the *Lab* color space intends to approximate human vision by mimicking the color opponency processing of early visual stage (bipolar and ganglion cells).

The last element of this part of the architecture is a layer of neurons that categorize the patterns observed in the visual descriptors in order to learn the local views. This stage gives the "what" information by creating signatures identifying the local views that the system observes. Here, the same type of neurons as in the LPMP model were used (see Section 1.3 of Chapter 1 for description and equations). During learning phase (see below) whenever none of the activities of learned visual signature is greater than the vigilance level, a new category is recruited. Two different vigilance values were used in the experiment: v_L and v_C for luminance and chrominance input respectively.

In this experiment, recognizing an object implies the recognition of co-located parts (when a set of local views describing the object are seen next to each other). This requires the "where" information in addition to the "what" information. In our model, the continuous exploration of the visual scene models eye saccades; while camera rotations correspond to head movements. By merging the eye and head directions, we are able to encode the orientation of the local view currently being seen in the idiothetic referential. Thus, the "where" information is represented in an azimuth vector a of size N_a where each neuron has a preferred direction around the yaw axis. We apply a lateral diffusion around the neuron coding for the direction of the current visual input to obtain gaussian bell-shaped activities with a standard deviation of σ_a.

Because of the continuous exploration of the visual scene, the system needs a short-term memory integrating the spatial position of the recognized local views. For the sake of simplicity, we implement the working memory as a collection of 3 dynamic neural fields (DNF) of size N_a; one field per object. During the learning phase, we operate a local view labeling while each object is presented separately. That is to say, the system learns to which of the objects the local views signatures belong. During the experimentation phase, these signatures-objects associations are used to implement a switch mechanism in order to direct the what-where input to the right field of the working memory. The combined what-where information is also a vector exhibiting a gaussian-bell shape centered on the direction of the current local view.

Figure 7.7 gives a more detailed illustration of the neural network that implement this model, with an emphasis on novelty detection, self-assessment and metacontrol. In the context of this recognition task, the Per matrix representing the system perceptions is the integration of the visual object categories Obj and the learned actions $ObjAct$ (respectively represented by Sen and Ac in equation (7.1)).

In this specific task, the 1st order novelty is sufficient to detect unexpected situations. This is due to the small number of objects and actions that are considered ((3 objects) $\times$ (3=(2+idle) actions)). Then, based on $\mathcal{N} = \mathcal{N}^{(1)}$ and the progress and regress deduced from it, self-assessment
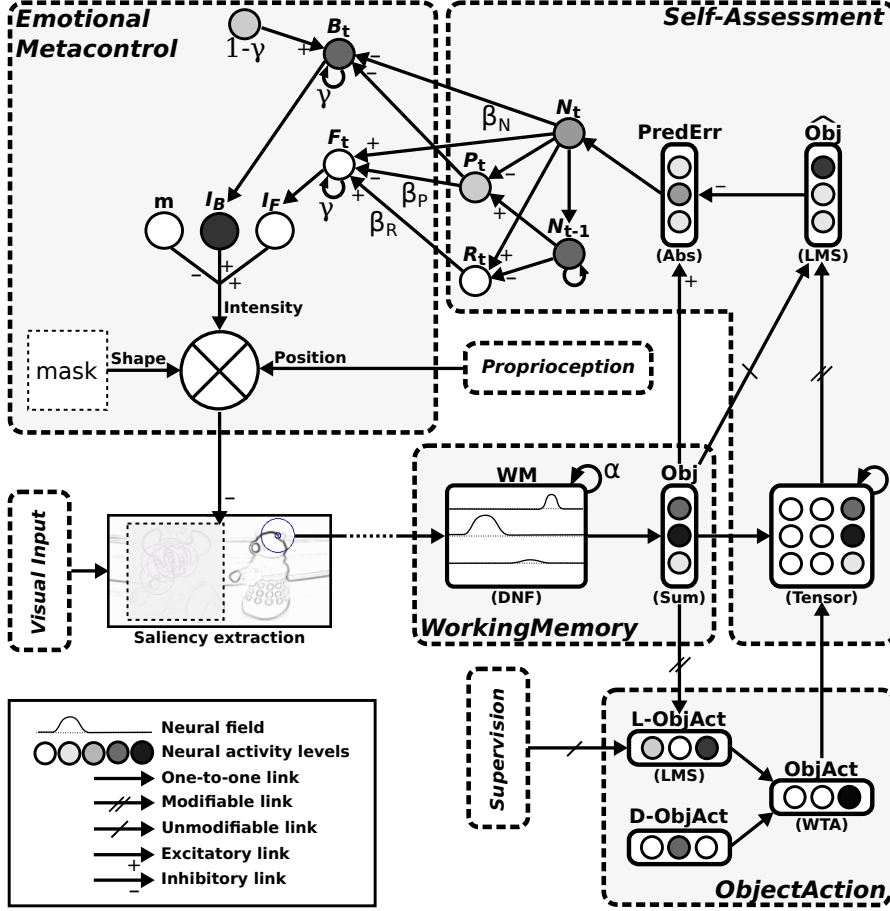
Figure 7.7: Self-Assessment and Emotional Metacontrol modules. LMS: Least Mean Square; DNF: Dynamic Neural Field; WTA: Winner-Takes-All; Abs: Absolute Value. WM: Working Memory. D-ObjAct: Default ObjectAction. L-ObjAct: Learned ObjectAction. N_t, P_t, R_t, F_t and B_t correspond to the novelty $\mathcal{N}$, progress $\mathcal{P}$, regress $\mathcal{R}$, frustration $\mathcal{F}$ and boredom $\mathcal{B}$ at time t in the equations. Some of the parameters of those equations like α and γ are also shown on this figure.

elicits frustration or boredom. Neither of these dual affects represent a positive experience. Instead, the system should seek an equilibrium between its skills and the task demands; a state in which it is neither frustrated nor bored. Therefore, we use the $\mathcal{F}$ and $\mathcal{B}$ levels as inputs to the *Emotional Metacontrol* mechanism that generates an inhibition potential I used to bias the bottom-up visual attention away from the current gaze direction. Thereby, the robot can avoid deadlock situations and explore the rest of the visual scene.

$$I(t) = I_{\mathcal{F}}(t) + I_{\mathcal{B}}(t) - \beta_m.m(t) \tag{7.12}$$

$$\text{with} \begin{cases} I_F(t) = g(\mathcal{F}(t)) \text{ and} \\ I_B(t) = g(\mathcal{B}(t)) \end{cases}$$

where $g(x) = 0.5 \times (1 + tanh(6x - 3))$ is the activation function of the top-down inhibition neurons. It remaps the values in [0,1] in a way that "pushes" low and high values toward the

extrema. By introducing this non-linearity, the effect of $\mathcal{F}$ and $\mathcal{B}$ is accentuated[1] . Besides, $m(t)$ is an estimation of the head movement to be generated. It is calculated from the angular difference between the current head direction and the position of the most salient object on which the system needs to focus its attention. This negative term avoids affect-based inhibition when the system is trying to focus on a new part of the visual scene. See Figure 7.7 and equations (7.10) and (7.11) for more details.

For this Emotional Metacontrol (based on inhibition I) to be efficient, it has to be maintained long enough to allow for the visual recognition of other elements of the scene. Otherwise, the camera would oscillate too rapidly between various positions. In this experiment, while local view recognition runs at 10 images/sec, the sub-network computing this inhibition is updated every 8 sec.

All the parameters described above are set to the values given in Table 7.2

7.9 Experiment 2: Deadlock detection and visual exploration

Figure 7.8 illustrates the experimental setup used in this search task. The robotic platform comprises a pan-camera and a 1-DoF (Degree of Freedom) arm. During task learning, objects are presented individually. In the experimentation phase, there are presented pairwise. The task is to search for (recognize) the target objects in these configurations and perform the corresponding actions to confirm the recognition.
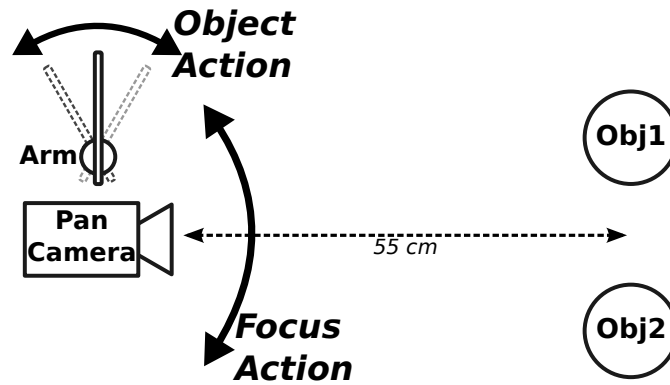
Figure 7.8: Experimental setup. The robotic platform consists in a pan-camera and an arm. In the experimentation phase, three objects – Target1, Target2 and Distractor – are presented by pairs (placed on Obj1/Obj2 positions). The robot can perform two kinds of actions: 1) FocusAction consists in moving the camera toward the most salient object in the scene; 2) ObjectAction consists in moving the arm to the left or the right whenever one of the targets is recognized. TargetX/ActionX associations are learned during the task learning phase.

7.9.1 Results

In this section, I will show results regarding the prototypical behaviors that are observed using the *Emotional Metacontrol* then statistically compare the results of this architecture to one that

---

[1]Although this was not implemented in Experiments 1.A and 1.B, it should have the same effect in the context of strategy selection.

is only feed-forward – i.e. does not perform the top-down attentional bias.

Prototypical behaviors: T1D configuration

In this configuration, the Distractor is more salient than Target1 (Figure 7.9, upper part, 2 sec). This means that with no metacontrol, the bottom-up attentional mechanism focuses the robot attention on the former (Figure 7.9, right part, 2 sec). Therefore, nearly no points of interest are detected on the target (Figure 7.9, 4 sec). Thus, the level of recognition of Target1 is not sufficient to trigger the corresponding action and fulfill the task. The inhibition of the Distractor area is required to allow enough points of interest to be detected on Target1 (Figure 7.9, 6 sec).

The neural activity plots in Figure 7.9 (subfigures) show the role of the top-down attentional bias in such situations. In the beginning of the trial, the recognition level of the Distractor is the highest and the robot starts turning the camera toward it. In the meantime, the novelty-based self-assessment increases the frustration level. Thus, a top-down inhibition of the area around the Distractor allows the recognition of Target1 and the performance of Action1. However, when the metacontrol module has to update the area to be inhibited, the robot head had moved away from the Distractor. Thus attention is again focused on the latter and the same situation as in the beginning of the trial is observed until the system is able to recognize Target1 once again. We note that during the periods Action1 is maintained, the novelty level decreases and boredom progressively takes over frustration.
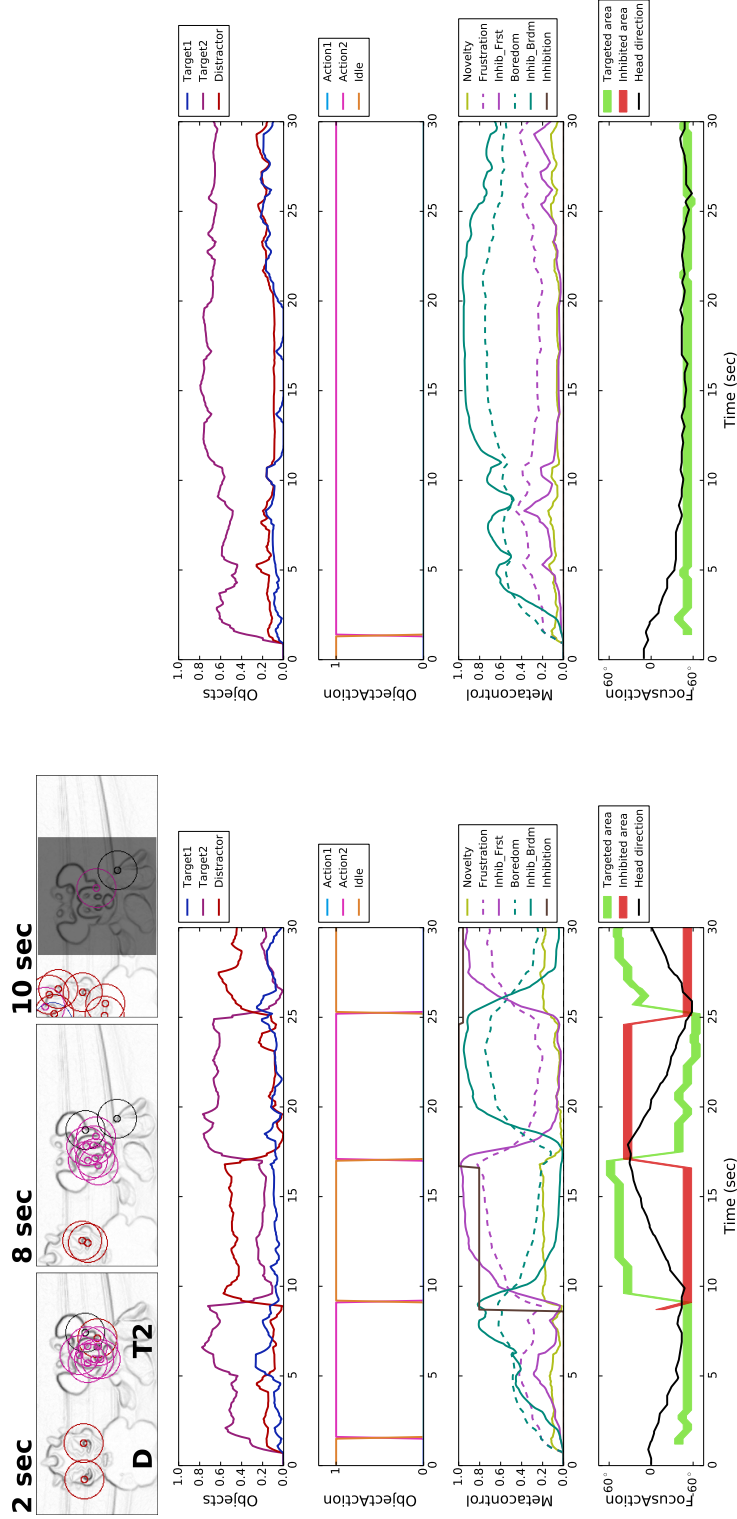
Prototypical behaviors: T2D configuration

As opposed to the previous configuration, here the target is more salient than the Distractor. Thus, the robot is quickly able to do the associated action and fulfill the task (Figure 7.10, upper part). Figure 7.10 illustrates the interest of boredom in the metacontrol mechanism in comparison with a feed-forward architecture. When Action2 is maintained, the self-assessment detects little novelty and generates boredom. Again, the current head direction is inhibited and the second object is recognized. In this configuration, it is a Distractor. So as soon as the robot is bored, the top-down attentional bias intervenes. Consequently, the robot stops performing the correct action and the frustration level increases. Since in this configuration, there is only one target that is relevant to the search task, the visual exploration achieved thanks to the boredom-driven metacontrol does not enhance the robot performance.

Prototypical behaviors: T1T2 configuration

The T1T2 configuration shows a situation in which the boredom-induced metacontrol can be relevant. Like in the previous configuration, Target2 is the most salient object in the visual scene (Figure 7.10, upper part). However, is this case, the boredom-driven metacontrol allows the robot to also recognize Target1 and perform the second learned action. Figure 7.10 shows the alternation between the two targets. Each time a different action is performed, the boredom level decreases a little until the novelty is too low again.
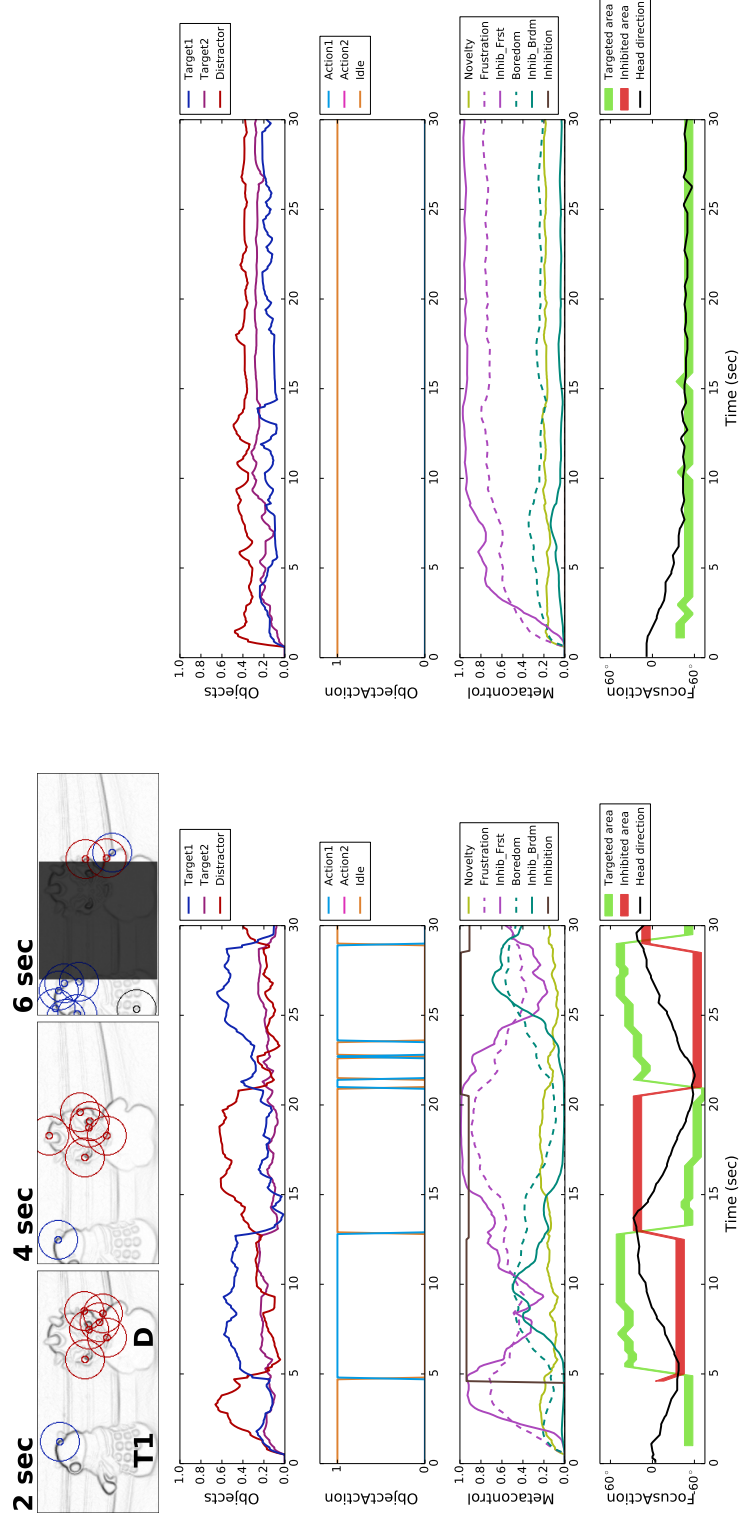
Figure 7.9: Prototypical behavior in the T1D configuration with the emotional metacontrol mechanism (LEFT) and without it (RIGHT). The top row gives three samples (beginning of the trial, then before and after the first top-down inhibition) of the points of interest detection represented on the gradient images – only shown in the case of emotional metacontrol. The subfigures respectively show the object recognition levels, the action performed by the robot arm, the self-assessment and metacontrol signals, and directions related to the focus of attention. With emotional metacontrol (LEFT), after the attention is first focused on the Distractor, frustration-driven metacontrol allows for recognizing the target and fulfilling the task. In contrast, without emotion metacontrol (RIGHT), the robot only focuses on the Distractor and no action is performed with the robot arm.

Figure 7.10: Prototypical behavior in the T2D configuration with the emotional metacontrol mechanism (LEFT) and without it (RIGHT). The top row gives three samples (beginning of the trial, then before and after the first top-down inhibition) of the points of interest detection represented on the gradient images – only shown in the case of emotional metacontrol. The subfigures respectively show the object recognition levels, the action performed by the robot arm, the self-assessment and metacontrol signals, and directions related to the focus of attention. With emotional metacontrol (LEFT), boredom-driven metacontrol inhibits the area around Target2 and allows for momentarily exploring the rest of the scene before attention is again focused on Target2. In contrast, without emotion metacontrol (RIGHT), the robot only focuses on the Target2 and does not explore the rest of the scene.
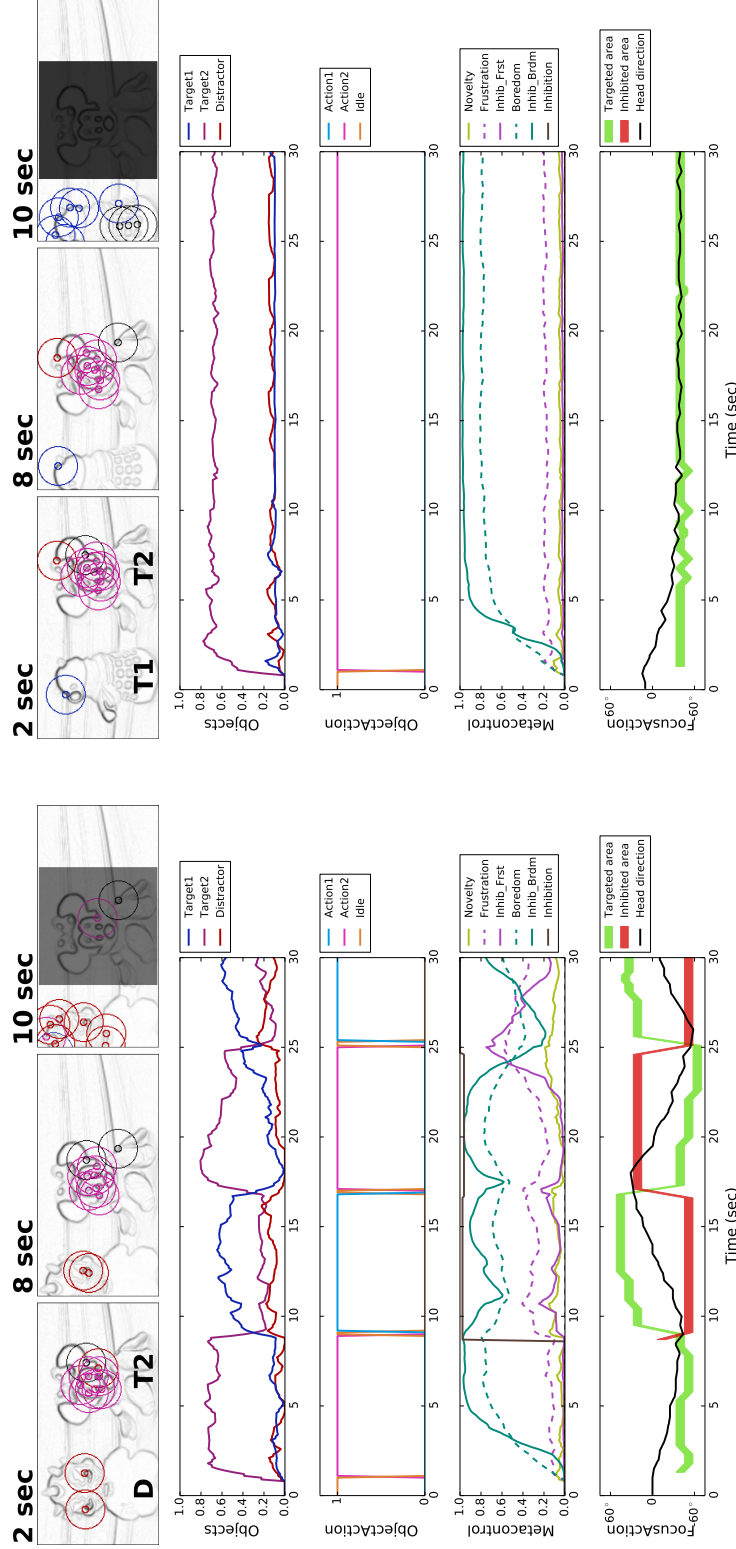
Figure 7.11: Prototypical behavior in the T1T2 configuration with the emotional metacontrol mechanism (LEFT) and without it (RIGHT). The top row gives three samples (beginning of the trial, then before and after the first top-down inhibition) of the points of interest detection represented on the gradient images – only shown in the case of emotional metacontrol. The subfigures respectively show the object recognition levels, the action performed by the robot arm, the self-assessment and metacontrol signals, and directions related to the focus of attention. With emotional metacontrol (LEFT), after Action2 is performed, boredom-driven metacontrol inhibits the area around Target2 and allows for recognizing Target1. The robot is thus able to find all targets in this configuration and alternates between the two learned actions. In contrast, without emotion metacontrol (RIGHT), the robot only focuses on the Target2 and does not explore the rest of the scene.

Statistical comparison: with and without Emotional Metacontrol

To compare the full architecture described so far to a feed-forward one, a total of 96 repetitions is performed, 16 for each configuration, with and without emotional metacontrol (resp. noted w/EM, w/oEM). Thus, 6 conditions arise from the combination of the *configuration* and *meta-control* factors (card({T1D,T2D,T1T2})=3 and card({w/oEM, w/EM})=2). A 30 sec timeout per trial is imposed to overcome deadlock situations. Given the dynamics of the system, this duration is long enough to recognize an object and thus defines an acceptable upper bound in case of deadlock. Also, small changes in the objects positions are voluntarily applied from a trial to another in order to introduce some variability.

Two measures are considered:

- *RT*: The Response Time corresponds to the delay between the beginning of a trial and the time when the robot performs an action associated with one of the target objects. First, it captures the difference between the 3 objects in terms of bottom-up saliency and could thus depend on the configuration. In addition, it indirectly characterizes the system success or failure.

- *QM*: The Quantity of head Movement characterizes the robot tendency to switch attention and explore the visual scene. It is defined as the total head rotations within a trial: $QM = \sum_{n=1}^{N} \theta_H(n) - \theta_H(n-1)$, where $\theta_H(n)$ is the orientation of the robot head in the n^{th} among N iterations.

After calculating the skewness and kurtosis of *RT* (skewness = 2.02; kurtosis = 2.35) and *QM* (skewness = 3.43; kurtosis = 17.77), it appears that none of these dependent variables follows a normal distribution. Therefore, the non-parametric Kruskal-Wallis (K-W) and Mann-Whitney (M-W) tests are performed [Corder and Foreman, 2009]. A K-W test shows a significant effect of the *configuration* on *RT* ($\chi^2 = 59.66$, $p < 0.01$) with a mean rank of 79.31 for T1D, 36.19 for T2D and 30.00 for T1T2. A M-W test confirms the difference between T1D and T2D ($U = -986.00$, $p < 0.01$) and between T1D and T1T2 ($U = -986.00$, $p < 0.01$) while no effect is found between T2D and T1T2 ($U = 397.00$, $p = 0.11$). Moreover, a M-W test does not show any statistical effect of the *metacontrol* on *RT* ($U = 1120.00$, $p = 0.81$).

The fact that I considered *RT* regardless of which action is performed does not allow us to conclude on the robot success or failure. For instance, doing Action2 in the T1D configuration means an erroneous recognition of Target2 but is not captured as such by *RT*. In addition, it does not allow for observing the effects of the *metacontrol* while taking into account the inherent difference of saliency between the objects. There is also a loss of information in the case of T1T2 configuration where both actions are possible. Therefore, it is important to analyze the effects of the independent variables on RT1 and RT2 (RT for Action1 and Action2 resp.) separately by omitting the configuration where the corresponding action is incorrect (e.g. Action2 is incorrect in T1D).

A M-W test again shows a significant effect of the *configuration* (T1D vs. T1T2) on *RT1* ($U = 358.00$, $p = 0.03$). But, this time we also observe a significant difference between the 'w/oEM' and 'w/EM' groups ($U = 91.00$, $p < 0.01$). On the other hand, no significant effect is observed on *RT2*, neither by the *configuration* ($U = 415.00$, $p = 0.19$) nor by the *metacontrol* ($U = 379.00$, $p = 0.07$). Indeed, since Target2 is the most salient object of the experiment, *RT2* only slightly varies across T1T2 and T2D configurations. In contrast, the *metacontrol* does
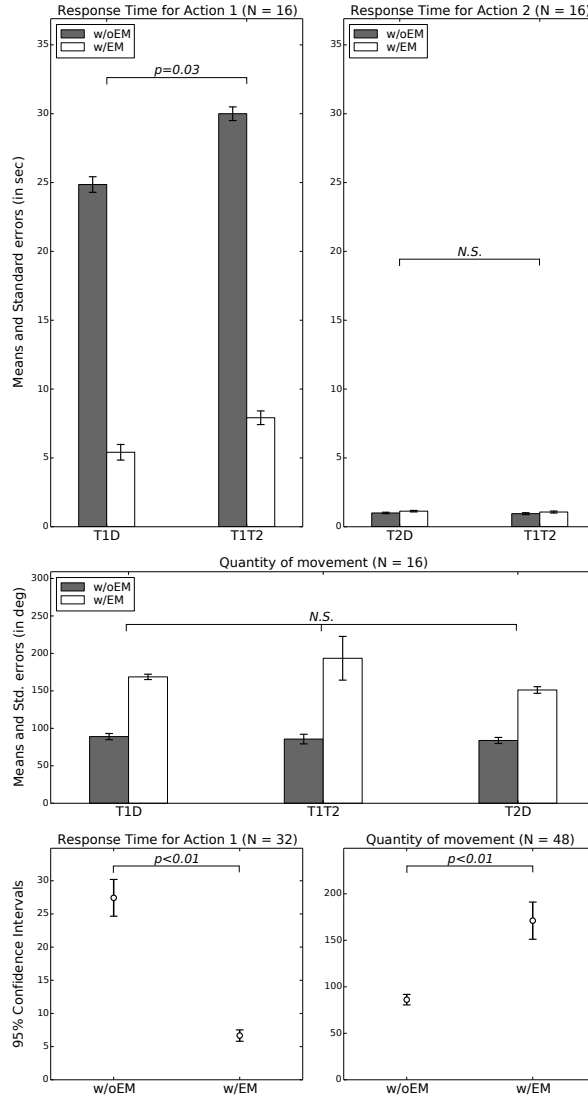
Figure 7.12: Statistical results. **TOP-LEFT**: Means and standard errors of the *RT1* observations in the T1D and T1T2 for the 'w/oEM' and 'w/EM' groups. *p* is the probability that the observations in T1D and T1T2 - regardless of the *metacontrol* variable - have the same distribution and is obtained by a Mann-Whitney test. **TOP-RIGHT**: Means and standard errors of the *RT2* observations in the T2D and T1T2 for the 'w/oEM' and 'w/EM' groups. **MIDDLE**: Means and standard errors of the *QM* for the 'w/oEM' and 'w/EM ' architectures in all configurations. **BOTTOM-LEFT**: Confidence intervals for *RT1* regardless of the configuration (T1D observations grouped with T1T2). **BOTTOM-RIGHT**: Confidence intervals for *QM* regardless of the configuration. The important gaps between the confidence intervals of the 'w/oEM' and 'w/EM' groups are consistent with the significant effects revealed by the M-W tests (*p* values reported on the figures). The *N* values in the all subfigures represent the sample sizes.

make a difference for *RT1* since Target1 is hardly recognized without the top-down attentional bias.

Lastly, a K-W test finds no effect of the *configuration* on *QM* ($\chi^2 = 1.56$, $p = 0.46$),

the mean ranks being 44.64, 47.66 and 53.20 for T1D, T2D, T1T2 respectively. However, a M-W again reveals a significant difference between the 'w/oEM' and 'w/EM' architectures ($U = 50.00$, $p < 0.01$).

These statistical results are confirmed by the analysis of the collected data. In particular, Figure 7.12 shows the means and standard errors of the *RT1*, *RT2* and *QM* independent variables. In addition, the confidence intervals demonstrate the magnitude of the observed effects. Indeed, the important gaps between the 'w/oEM' and 'w/EM' groups are consistent with the significant effects revealed by the M-W tests. Some of the *p*-values (significance results) obtained with the Kruskal-Wallis and Mann-Whitney tests are reported on Figure 7.12 for conciseness purpose – even though it is worth noting that these tests are rank-based and do not directly rely on the means and standard deviations of the samples.

7.9.2 Discussion

The results presented above highlight the interest of using self-assessment-based emotional signals to implement a second-order controller. The prototypical behaviors that are observed mainly describe two kinds of situations:

1) Situations of failure: *Frustration* reflects the system incapacity to perform the target recognition task;

2) Situation of long-lasting success: *Boredom* illustrates the lack of challenge in the task. Neither of these emotional experiences are considered positive. So, the *Emotional Metacontrol* is used to avoid these deadlock situations and try to maintain the system in a state of balanced skills and challenges. The statistical evaluation further validates the benefit of the *Emotional Metacontrol* as compared to a more basic feed-forward architecture. Indeed, the affect-based feedback loop both increases the robot performance (effect on *RT1*) and fosters the exploratory behavior to avoid deadlocks (effect on *QM*).

This experiment demonstrates the interest of the emotional modulation occurring at the level of sensory information. Indeed, objects compete for limited processing capacity. This is represented in our system by the limited number of local views that are extracted and encoded from each image. Since Target2 is much more salient than Target1, when both are presented concurrently, only 10-25% of the points of interests are initially detected on Target1. Then, as the camera turns toward Target2 due to the bottom-up focus of attention via head rotations, Target1 has even less chances to capture the robot attention. If the Emotional metacontrol applied the top-down bias after the objects recognition in the working memory, or even later at the action triggering level, most of the resources would still be used by Target2. Thus, the recognition of Target1 would be very low, indistinguishable from noisy activation of isolated, wrongly identified local views. In contrast, when the inhibition is done on the saliency map (i.e. the early stages of our visual system), it allows a more efficient detection of visual features out of the inhibited area. Besides, there is evidence that emotional stimuli are enhanced in the working memory [Gray et al., 2002] and consolidate long-term memory [McGaugh, 2015]. This further confirms the need for an early influence of emotional processes on cognitive ones.

For the purpose of a search task, a simple conditioning could also allow for learning the location of the interesting objects (e.g. subregion of the visual scene). Such a purely spatial learning would implement a rather bottom-up attentional bias. This mechanism would be efficient in certain situations. For instance, with a sufficiently high learning rate, the robot would be

able to predict the target position after a small number of repetitions of the same configuration; provided that the objects do not move. Nevertheless, in our experiments, the objects positions in front of the camera were alternated (e.g. in the T1D configuration, Target1 could be on the right or on the left). Therefore, the purely spatial attentional strategy would fail as soon as we move the objects.

From the perspective of visual exploration, a simpler way to make the system seek novel stimuli is to rely on a habituation mechanism. For instance, in the active vision model proposed by Breazeal et al. [2001], it prevents the robots from fixating highly salient background stimuli and forces the human caretaker to constantly engage in interactions with slightly different stimulations. In contrast with this kind of solutions that operate like an ad-hoc timeout mechanism, our *Emotional Metacontrol* generates an inhibition potential according to the performance of the system. Thus, it has the potential to adapt to situations where different objects imply actions of different complexity or duration – e.g. attention must be focused on Target1 longer that on Target2 in order to perform Action1 – as long as it is assessed by the system as a normal situation. Besides, instead of forcing the human caretaker to maintain balanced novelty, it can allow the robot to communicate deadlocks and ask for help only when needed as seen in Experiment 1.A.

Unlike other models of top-down attentional bias, the proposed solution only inhibits the current gaze direction on the bottom-up saliency map. So, its success highly depends on the fact that the exploration of close regions of the visual scene would help the robot focus on relevant objects. This solution could be too simple in more complex tasks. An alternative could be to rely on a prior knowledge about the task to implement a top-down mechanism that biases attention toward the visual cues that are relevant for the current sub-task [Breazeal et al., 2001]. Another option is to use reinforcement learning in order to avoid such pre-wired connections between the visual system and the tasks demands. Ognibene and Baldassare [2015] provide a good insight into the use of actor-critic models for the implementation of goal-directed top-down attentional bias in active vision architectures.

7.10   General discussion

The main interest of the way I obtain the internal model of 'normal', 'usual' experiences is that it is generic and adaptable to any sensorimotor architecture. Theoretically, one can use a unique perception tensor that integrates all sensations and actions. It should be able to build expectations about separate sensations as long as the predictor sees enough situations to filter invariants. For instance, in Experiment 2, I could include the head direction vector with the tensor inputs. But, the robot would have to learn the sensation–action combinations by placing the object in various positions to decrease the influence of this input. This solution also seems quite inefficient and computationally costly. I would rather argue for a distributed encoding of perception, based on several tensors combining couples of sensations and actions.

Another issue is the possible combinatorial explosion due to the integration of very large input vectors in the perception tensors. The simplest way to overcome this issue is to somehow compress the raw input (e.g. strong discretization or intermediate categorization layers to build abstract representations).

In addition, a biological source of inspiration lies in the conjunctive cells found in the entorhinal cortex Sargolini et al. [2006]. This particular type of neuron provides compressed codes combining position and direction information. An interesting property is that only encountered

patterns are encoded - instead of all the possible combinations. The hypothesis is that the entorhinal cortex take advantage from the strong connectivity with other cortical regions to provide the hippocampus with compressed codes that represent multimodal conjunctions of states (possibly all sorts of modalities)[Gaussier et al., 2007; Jauffret et al., 2015]. Moreover, the hippocampus has been implicated in novelty detection through its role in memory formation [Knight et al., 1996; Lisman and Otmakhova, 2001].

In the literature, novelty has been shown to be an important factor in the perceptual processing of emotional stimuli for different modalities [Ethofer et al., 2009; Wendt et al., 2011]. It is considered essential in the appraisal theories of emotions [Scherer, 2001; Grandjean and Peters, 2011]. According to Scherer [2001], novelty is evaluated in terms of suddenness, familiarity and predictability of the stimuli. These appraisals respectively occur at the sensorimotor level (novel sensory input), the schematic level (adequacy with learned preferences) and the conceptual level (ability to predict the input). However, little is said about the processes leading from a level to another. Our work gives an insight on how a novelty detection mechanism that is rooted in the sensorimotor experience of an embodied and situated robot can be a key intrinsic information for the self-assessment of the skills/challenges compatibility.

The study conducted by Beer et al. [2006] with patients with damaged OFC is in line with the emotion–cognition integration in self-evaluation processes as proposed in this work. In the experiment of Beer et al. [2006], the patients were aware of social norms. Yet, they only showed emotional reactions related to embarrassment and shame when they actively engaged in self-monitoring but not during the social interactions in which they had had inappropriate behaviors. Thus, OFC might be a key region in which high-level appraisal – i.e. evaluating complex situations, as opposed to stimulus-based appraisal in the amygdala for instance – occurs and triggers emotional responses.

Personal publications

Jauffret, A., Belkaid, M., Cuperlier, N., Gaussier, P., and Tarroux, P. (2013b). Frustration: a generic mechanism to improve autonomy in robotics. In *Proceedings of the IEEE Joint International Conference on Development and Learning and Epigenetic Robotics*, pages 1–7.

Jauffret, A., Belkaid, M., Cuperlier, N., Gaussier, P., and Tarroux, P. (2013d). Frustration: a generic mechanism to improve autonomy in robotics. In *Symposium on Biology of Decision Making*.

Belkaid, M., Cuperlier, N., and Gaussier, P. (2016b). De la détection de nouveauté au métacontrôle en passant par l'ennui et la frustration. In *Actes du Workshop Affect Compagnon Artificiel et Interaction*.

Belkaid, M., Cuperlier, N., and Gaussier, P. (2016c). Emotional Metacontrol: top-down emotional modulation of sensorimotor processes during active object perception. *Submitted*, pages 1–13.

Table 7.1: Experiments 1.A and 1.B: Parameter values

	Value	Description
α	0.95	Weight of recurrent link in Per
γ	0.999	Weight of recurrent link in $\mathcal{F}$
$\beta_{\mathcal{N}}$	0.0015	Weight of novelty $\mathcal{N}$ link in $\mathcal{F}$
$\beta_{\mathcal{R}}$	0.01	Weight of regress $\mathcal{R}$ in $\mathcal{F}$
$\beta_{\mathcal{P}}$	0.01	Weight of progress $\mathcal{P}$ in $\mathcal{F}$
β_h	0.001	Weight of human intervention h in inhibition I
N_{vp}	5	Nb. of vanishing point neurons

Table 7.2: Experiment 2: Parameter values

	Value	Description
α	0.9	Weight of recurrent link in Per
γ	0.99975	Weight of recurrent link in $\mathcal{F}$ and $\mathcal{B}$
$\beta_{\mathcal{N}}$	0.001	Weight of novelty $\mathcal{N}$ link in $\mathcal{F}$ and $\mathcal{B}$
$\beta_{\mathcal{R}}$	0.01	Weight of regress $\mathcal{R}$ in $\mathcal{F}$ and $\mathcal{B}$
$\beta_{\mathcal{P}}$	0.01	Weight of progress $\mathcal{P}$ in $\mathcal{F}$ and $\mathcal{B}$
β_m	0.01	Weight of head movement m in inhibition I
m_{size}	149	Size of inhibition mask (in pixels)
p_y	427	Elevation of objects in the image
N_{PoI}	10	Nb. of PoI extracted per image
r_{small}	4	Small disk radius in local views (in pixels)
r_{big}	24	Big disk radius in local views (in pixels)
W_{LP}	18	Width of the visual descriptor
H_{LP}	60	Height of the visual descriptor
v_L	0.9	Vigilance for luminance descriptor
v_C	0.95	Vigilance for chrominance descriptor
N_a	61	Size of azimuths vector : $\sim 360^o/6$
σ_a	8	Std. dev. of azimuths diffusion (in degrees)
σ_{DoG_1}	10	Std. dev. of 1^{st} DoG gaussian (in degrees)
σ_{DoG_2}	20	Std. dev. of 2^{nd} DoG gaussian (in degrees)

Conclusion

Summary

In this dissertation, I presented my research work on the interactions between cognitive and emotional processes through the prism of robotics. In terms of positioning, my work lies in the intersection between connectionism and enactivism. The programs that control the robot behavior were implemented on artificial neural networks and embodied in robotic platforms. They were designed by taking inspiration from biological systems – i.e. what we know about humans and animals from studies in other cognitive science disciplines. Sensorimotor architectures were proposed, based on the idea that complex, yet stable, behaviors can emerge from the sensation–action loop in an embodied and situated agent. Given the ability to learn and adapt, these robotic systems could benefit from a rich physical and social interaction with their environment.

The first part of the dissertation considered spatial navigation as a framework to study biological and artificial cognition. In that matter, I started Chapter 1 with an overview of some crucial neuroscientific findings implicating the hippocampal region in navigation-related functions. Of particular interest for my work are place cells: a particular type of neurons that exhibits localized firing in familiar environments. Previous works in the lab proposed a biologically plausible model of place recognition that mimics these cells.

In Chapter 2, I took a closer look at visual place recognition in the outdoors. This type of environment is challenging because it is much less structured and controlled than the indoors. A particular issue to which the existing model of place cells is subject is the efficiency and robustness over long distances. The approach that I followed was to examine the type of visual information that is used for localization. In the literature, methods using local features generally perform better. However, global visual information is also informative and its computation may be less costly. To confirm that, I presented an experiment comparing seven different global descriptors to the local one used in the existing LPMP model. As expected, the LPMP showed the best overall performance. However, the results also suggest that holistic features convey significant information in the case of navigation. Consistent with this idea, there is evidence that human are able to capture the gist of a scene from its global aspect based on low-resolution information. In addition, the notion of global precedence indicates that these visual contexts facilitates a more refined recognition level that is based on local visual details. Hence, I proposed a model for context-based place recognition inspired by this literature. This solution increased the performance of the existing model. I also showed promising results in terms of navigation in wide environment and over long distances. More precisely, using fast holistic recognition as a contextual bias to limit the search space on the local-vision level not only enhanced the

localization but also maintained a quasi-constant computational cost.

It was also shown that, an implementation of the place cells model on robots (LPMP model) could be used to teach them navigation trajectories based on sensorimotor coupling. In Chapter 3, I took an interest in the learning of such navigation tasks through non-verbal human–robot communication based on low-level visuomotor signals. In particular, I worked on a model aimed to take into account human sensitivity to rhythmical interactions. Therefore, synchrony detection is used for partner selection. Visual recognition and tracking of human partners allowed the robot to re-engage and maintain interactions during learning phases. The results I presented showed that this architecture was able to learn and perform a sensorimotor navigation task.

The second part of the dissertation addresses the central question of emotion–cognition interactions. Chapter 4 reviewed emotion research as a cross-disciplinary enterprise. After an overview of the main schools of thoughts and examples of influential psychological theories, I reported neuroscientific findings related to brain regions and neurochemical signals involved in emotion. Then I presented some examples of computational models implemented either on robots or virtual conversational agents. For a roboticist, a functional view of emotions make them interesting to model for two reasons: enhancing robot–robot and human–robot interaction, and increasing robot autonomy and adaptation capabilities. On the other hand, computational models allow researchers to test hypotheses. Going from theoretical concepts to concrete implementations also reveals ambiguities and forces more explicit definition.

As highlighted in Chapter 5, converging findings suggest that emotion and cognition are not easily separable. In fact, they appear to continuously influence each other. Processes like perception, attention, memory and reasoning, which are associated to cognition, influence emotional experiences through appraisal and affect elicitation. In return, they are also modulated by emotional signals. Biologically, these interactions seem to be represented by the bi-directional influence between neural activities and neurochemical modulations. Computationally, dynamical systems theory appears to provide a good framework to study such a coupling. On a higher level of abstraction, I proposed a conceptual model of emotion–cognition interaction. This model provides a complementary view to neural network-based descriptions by representing connections between processes – similarly to a workflow diagram for example. First, this allows to easily map the model to descriptions that are also generally process-oriented. Second, it provides a generic view of the concept regardless to any implementation or application. In a nutshell, the model illustrates how the system continuously:

(a) appraises events from the body and the world with a particular interest to emotionally relevant stimuli,

(b) processes emotionally modulated signals and reintegrates them in the information processing flow for the purpose of higher order processing and appraisal.

In Chapter 6, I presented a first instantiation of the abstract model. Thereby, I addressed the effect of emotion on the perception of the peripersonal space; the space where one can safely and comfortably interact with the physical and social environment. For humans and animals, the ability to build such a representation is of fundamental importance, even in most basic behaviors involving approach and avoidance. Besides, there is evidence that emotions have a strong influence on the perception of the surrounding space. In the model I proposed, emotionally valenced sensory and physiological signals modulated the information processing flow at the sensory and the motor levels. I presented experiments showing that emotional modulation of peripersonal space can serve as a paradigm for reactive, behavior-based robot control. Additional results sug-

gested that the model was able to account for the impact of objects desirability on the perception of reachability. Moreover, in a social context, robots exhibited fearful or aggressive behavior depending on internal interactions between emotional signals. This provided an illustration of how internal emotional processes impact external aspects, more related to physical and social interactions.

Another example of emotion–cognition interaction is given in Chapter 7. I introduced the concept of Emotional Metacontrol as the use of emotional signals elicited by self-assessment to modulate various computational processes for the purpose of behavior regulation. Self-assessment is the capacity to evaluate one's own abilities and performance. In the proposed model, it is done on the fly, based on the evaluation of sensorimotor contingencies and the detection of novel situations. In the literature, frustration and boredom characterize the incompatibility between skills and task demands. Consequently, in my model, frustration was elicited by signals capturing regress and stagnation in high level of novelty. In contrast, boredom was generated by the absence of novelty and progress. In my model, these dual affects were both used to trigger a top-down modulation mechanism that regulate the robot behavior. In the first experiment of the chapter, frustration was used to monitor and modulate independent sub-behaviors of a mobile robot in a navigation task. First, I showed that the model provided an efficient means to communicate failure and the need for help to a human partner/supervisor. Second, I demonstrated that it also allowed strategy selection in a task that alternated road following and a learned sensorimotor trajectory. The second experiment addressed object recognition in a search task. This time, frustration and boredom were used to bias attention. They served as a way to escape from deadlocks, either due to failure or to long-lasting success. My results demonstrated that, as compared to a feed-forward architecture, emotional metacontrol increased the robot performance and fostered exploratory behaviors.

Perspectives

Although seemingly not at the core of my thesis works, social interactions were repeatedly evoked throughout the dissertation. This was mostly the case when discussing the issue of communicating internal, emotional states to peers and competitors (robots) as well as teachers and observers (humans). However, social cues and emotional expressions *per se* were very poorly addressed. When it was not voluntarily non-verbal and based on robot behavior (e.g. Experiment 2 in Chapter 6), communication relied on ad-hoc signals (e.g. sound, Experiment 1.A in Chapter 7). Yet, in biological systems, external and internal aspects of emotion (resp. sociality and autonomy) are two sides of the same coin. Internal regulatory processes of emotions are expressed, which improves communication and in return fosters adaptation through social interactions. My work would certainly benefit from richer expressiveness (e.g. facial expressions and body posture) in order to account for the social part of this dynamical system.

The experiment I presented in Chapter 3 explicitly engaged in the issue of using human–robot interactions as a catalyst for learning. I believe this work opens broad prospects for further research bridging robot autonomy and social interactions. For example, the question of "selective" learning was addressed: rudimentary novelty detection prevented the robot from learning sensorimotor associations that were too similar to what it had already learned. In Chapter 7, a more elaborated self-assessment mechanism was proposed, taking the dynamics of the robot behavior into account. The latter could obviously replace the former for a more adaptive solu-

tion. More interestingly, if the human teacher is walking by the robot side, gait synchronization should also provide a continuous feedback on the quality of learning and task performance. Thus, synchrony detection could serve as an input to another type of self-evaluation mechanism.

Additionally, future work should further study self-improvement; that is to say behavior and skill learning based on self-assessment. According to the literature on intrinsic motivation, the proposed model is knowledge-oriented since it relies on a comparison between current situation and learned representation. But it allows for indirectly evaluating skills because learning normal/usual situations occurs during task learning under human (semi-)supervision. In this framework, competence and knowledge are intertwined, both grounded in the robot sensorimotor and social experience. Moreover, related works emphasize the interest of detecting novelty and skills/task incompatibility in terms of learning capabilities. Making a step in that direction could consist in triggering an exploratory behavior to discover new sensorimotor patterns. Also, curiosity could make the robot seek situations in which learning progress can be maximized. Such a feature could allow it to engage in increasingly challenging activities: once simple skills are learned, an exploratory behavior allows for discovering more complex ones by avoiding 'unlearnable' situations.

One of the main ideas I wanted to defend in my thesis is that, in autonomous systems, emotion and cognition cannot be separated. Emotions emerge from physical and social interactions with the environment. They also filter information related to survival and well-being and modulate other brain activities accordingly. Thus, I believe that to be meaningful in artificial systems, emotions have to be integrated in the whole architecture through bidirectional influences with sensory, attentional, decisional and motor processes. Otherwise, they are mere responses to external stimuli that the designer labels as emotional. Throughout the second part of this dissertation, I attempted to provide evidence of how this approach to internal emotional processes could foster efficient physical and social interactions with the environment. By "efficient" I mean with regards to goals that are related to survival and well-being.

Through my research, I advocate the notion that the natural bias in favor of emotionally relevant information cannot only intervene at the level of action selection and decision-making. In Chapters 6 and 7, I showed examples of emotional influence on attentional and sensory processes, which also indirectly had an impact on short-term, working memory. From the functional perspective, it is essential that emotional modulation occur at earlier stages. Indeed, objects compete for limited processing capacity. Therefore, brain resources – memory and computation – must be recruited appropriately. There is still a lot to investigate in that matter. For instance, a more precise modeling of the visual cortex would open the way for numerous experiments linked to emotion influence on low-level perception. Further research in this direction not only implies changes in the image processing algorithms (and corresponding neural network implementations) but also implies different hardware solutions.

To conclude, this thesis provides a framework and various perspectives for a deeper understanding of the role of emotion in autonomous systems. For instance, a more detailed mapping of the brain structures involved in the processes that I studied might bring new insights or suggest new hypotheses on the functioning of biological emotions. But, the goal must not be to model one-to-one mapping between brain areas and their functions. On the contrary, the notion of connected networks should be emphasized. Thereby, internal dynamics can contribute to the interactions between agents and their physical and social environments.

Bibliography

Adam, C., Herzig, A., and Longin, D. (2009). A logical formalization of the occ theory of emotions. *Synthese*, 168(2):201–248. 81

Alexander, A. S. and Nitz, D. A. (2015). Retrosplenial cortex maps the conjunction of internal and external spaces. *Nature neuroscience*, 18(8):1143–1151. 55

Alexander, G. E., DeLong, M. R., and Strick, P. L. (1986). Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annual review of neuroscience*, 9(1):357–381. 77, 90

Alter, A. L. and Balcetis, E. (2011). Fondness makes the distance grow shorter: Desired locations seem closer because they seem more vivid. *Journal of Experimental Social Psychology*, 47(1):16–21. 112

Amari, S.-i. (1977). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological cybernetics*, 27(2):77–87. 103, 174

Ancona, N. and Poggio, T. (1995). Optical flow from 1-d correlation: Application to a simple time-to-crash detector. *International Journal of Computer Vision*, 14(2):131–146. 62

Anderson, A. K. (2005). Affective influences on the attentional dynamics supporting awareness. *Journal of Experimental Psychology: General*, 134(2):258. 88

Anderson, M. I. and Jeffery, K. J. (2003). Heterogeneous modulation of place cell firing by changes in context. *The Journal of neuroscience*, 23(26):8827–8835. 46, 54

Andry, P., Blanchard, A., and Gaussier, P. (2011). Using the rhythm of nonverbal human–robot interaction as a signal for learning. *IEEE Transactions on Autonomous Mental Development*, 3(1):30–42. 59

Angela, J. Y. and Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron*, 46(4):681–692. 80

Arbib, M. A. and Fellous, J.-M. (2004). Emotions: from brain to robot. *Trends in cognitive sciences*, 8(12):554–561. 70, 91

Arkin, R. C. and Ulam, P. D. (2009). An ethical adaptor: Behavioral modification derived from moral emotions. 118

Arleo, A. and Gerstner, W. (2000). Spatial cognition and neuro-mimetic navigation: a model of hippocampal place cell activity. *Biological cybernetics*, 83(3):287–299. 31

Arleo, A. and Rondi-Reig, L. (2007). Multimodal sensory integration and concurrent navigation strategies for spatial cognition in real and artificial organisms. *Journal of integrative neuroscience*, 6(03):327–366. 31

Arleo, A., Smeraldi, F., and Gerstner, W. (2004). Cognitive navigation based on nonuniform gabor space sampling, unsupervised growing networks, and reinforcement learning. *IEEE Transactions on Neural Networks*, 15(3):639–652. 31

Arnold, M. B. (1960). *Emotion and personality: Neurological and physiological aspects*, volume 2. Columbia University Press. 86

Asada, M., Hosoda, K., Kuniyoshi, Y., Ishiguro, H., Inui, T., Yoshikawa, Y., Ogino, M., and Yoshida, C. (2009). Cognitive developmental robotics: a survey. *IEEE Transactions on Autonomous Mental Development*, 1(1):12–34. 22

Avila-garcía, O. and Cañamero, L. (2004). Using hormonal feedback to modulate action selection in a competitive scenario. In *From Animals to Animats: Proceedings of the 8th International Conference of Adaptive Behavior*. Citeseer. 83

Avila-garcía, O. and Cañamero, L. (2005). Hormonal modulation of perception in motivation-based action selection architectures. In *Proceedings of the Symposium on Agents that Want and Like*. SSAISB. 83

Balcetis, E. and Dunning, D. (2009). Wishful seeing: more desired objects are seen as closer. *Psychological science*. 88, 98, 111, 112

Baldwin, A. E., Sadeghian, K., and Kelley, A. E. (2002). Appetitive instrumental learning requires coincident activation of nmda and dopamine d1 receptors within the medial prefrontal cortex. *The Journal of neuroscience*, 22(3):1063–1071. 79

Balleine, B. W. and Killcross, S. (2006). Parallel incentive processing: an integrated view of amygdala function. *Trends in neurosciences*, 29(5):272–279. 76

Bar, M. (2004). Visual objects in context. *Nature Reviews Neuroscience*, 5(8):617–629. 38, 46, 55

Bard, P. (1934). On emotional expression after decortication with some remarks on certain theoretical views: Part I. *Psychological Review*, 41(4):309. 76

Barrett, L. F. and Bar, M. (2009). See it with feeling: affective predictions during object perception. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 364(1521):1325–1334. 77, 78

Barto, A. G., Singh, S., and Chentanez, N. (2004). Intrinsically motivated learning of hierarchical collections of skills. In *Proc. 3rd Int. Conf. Development Learn*, pages 112–119. 118

Bassareo, V. and Di Chiara, G. (1999). Modulation of feeding-induced activation of mesolimbic dopamine transmission by appetitive stimuli and its relation to motivational state. *European Journal of Neuroscience*, 11(12):4389–4397. 79

Bay, H., Ess, A., Tuytelaars, T., and Van Gool, L. (2008). Speeded-up robust features (SURF). *Computer vision and image understanding*, 110(3):346–359. 38

Bechara, A. and Damasio, A. R. (2005). The somatic marker hypothesis: A neural theory of economic decision. *Games and economic behavior*, 52(2):336–372. 77

Beer, J. S., John, O. P., Scabini, D., and Knight, R. T. (2006). Orbitofrontal cortex and social behavior: integrating self-monitoring and emotion-cognition interactions. *Journal of cognitive neuroscience*, 18(6):871–879. 117, 139

Beer, J. S., Lombardo, M. V., and Bhanji, J. P. (2010). Roles of medial prefrontal cortex and orbitofrontal cortex in self-evaluation. *Journal of Cognitive Neuroscience*, 22(9):2108–2119. 117, 118

Berridge, K. C. (2012). From prediction error to incentive salience: Mesolimbic computation of reward motivation. *European Journal of Neuroscience*, 35(7):1124–1143. 76, 79

Berridge, K. C. and Kringelbach, M. L. (2008). Affective neuroscience of pleasure: reward in humans and animals. *Psychopharmacology*, 199(3):457–480. 77, 79, 99

Beurle, R. L. (1956). Properties of a mass of cells capable of regenerating pulses. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 240(669):55–94. 174

Blanchard, A. and Canamero, L. (2005). Using visual velocity detection to achieve synchronization in imitation. In *Procs 3rd Int Symposium on Imitation in Animals and Artifacts*. SSAISB. 59

Bolles, R. C. and Fanselow, M. S. (1982). Endorphins and behavior. *Annual review of psychology*, 33(1):87–101. 80

Bostock, E., Muller, R. U., and Kubie, J. L. (1991). Experience-dependent modifications of hippocampal place cell firing. *Hippocampus*, 1(2):193–205. 54

Botvinick, M. M., Cohen, J. D., and Carter, C. S. (2004). Conflict monitoring and anterior cingulate cortex: an update. *Trends in cognitive sciences*, 8(12):539–546. 78, 118

Boucenna, S., Anzalone, S., Tilmont, E., Cohen, D., and Chetouani, M. (2014). Learning of social signatures through imitation game between a robot and a human partner. *IEEE Transactions on Autonomous Mental Development*, 6(3):213–225. 71

Boucenna, S., Cohen, D., Meltzoff, A. N., Gaussier, P., and Chetouani, M. (2016). Robots learn to recognize individuals from imitative encounters with people and avatars. *Scientific reports*, 6. 71

Bouton, M. E., Kenney, F. A., and Rosengard, C. (1990). State-dependent fear extinction with two benzodiazepine tranquilizers. *Behavioral neuroscience*, 104(1):44. 54

Bouton, M. E., Westbrook, R. F., Corcoran, K. A., and Maren, S. (2006). Contextual and temporal modulation of extinction: behavioral and biological mechanisms. *Biological psychiatry*, 60(4):352–360. 54

Braitenberg, V. (1984). *Vehicles: Experiments in synthetic psychology*. MIT press. 22, 28

Breazeal, C. (2003). Emotion and sociable humanoid robots. *International Journal of Human-Computer Studies*, 59(1):119–155. 83

Breazeal, C. and Brooks, R. (2005). Robot emotion: A functional perspective. In Arbib, M. A. and Fellous, J.-M., editors, *Who Needs Emotions? The Brain Meets the Robot*, pages 271–310. Oxford University Press Oxford. 70, 83

Breazeal, C., Edsinger, A., Fitzpatrick, P., and Scassellati, B. (2001). Active vision for sociable robots. *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, 31(5):443–453. 138

Brooks, R. (1986). A robust layered control system for a mobile robot. *IEEE Journal on Robotics and Automation*, 2(1). 22, 115

Brooks, R. A. (1991). Intelligence without reason. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence - Volume 1*, IJCAI'91, pages 569–595. 22, 28

Bullmore, E. and Sporns, O. (2009). Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature reviews. Neuroscience*, 10(3):186–98. 112, 174

Burgess, N., Donnett, J. G., Jeffery, K. J., and John, O. (1997). Robotic and neuronal simulation of the hippocampus and rat navigation. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 352(1360):1535–1543. 31

Burgess, N., Recce, M., and O'Keefe, J. (1994). A model of hippocampal function. *Neural networks*, 7(6-7):1065–1081. 31

Bush, G., Luu, P., and Posner, M. I. (2000). Cognitive and emotional influences in anterior cingulate cortex. *Trends in cognitive sciences*, 4(6):215–222. 78

Cajal, S. R. (1911). Histologie du système nerveux de l'homme et des vertébrés. *Maloine, Paris*. 29

Cañamero, D. (1997). Modeling motivations and emotions as a basis for intelligent behavior. In *Proceedings of the first international conference on Autonomous agents AGENTS 97*, number Abbott 1884, pages 148–155. ACM Press. 83

Cañamero, L. (2014). Bridging the gap between HRI and neuroscience in emotion research: robot as models. In *HRI: a Bridge between Robotics and Neuroscience (Workshop at HRI)*. 70

Cañamero, L. and Gaussier, P. (2005). Emotion understanding: robots as tools and models. In Nadel, J. and Muir, D., editors, *Emotional Development*. 70, 71

Cardinal, R. N., Parkinson, J. A., Hall, J., and Everitt, B. J. (2002). Emotion and motivation: The role of the amygdala, ventral striatum, and prefrontal cortex. *Neuroscience and Biobehavioral Reviews*, 26(3):321–352. 76

Carpenter, G. A. and Grossberg, S. (1988). The art of adaptive pattern recognition by a self-organizing neural network. *Computer*, 21(3):77–88. 20, 115, 172, 174

Casey, K. L. (2000). Concepts of pain mechanisms: the contribution of functional imaging of the human brain. *Progress in brain research*, 129:277–287. 78

Chun, M. M. (2000). Contextual cueing of visual attention. *Trends in cognitive sciences*, 4(5):170–178. 38, 46

Coello, Y. (2005). Spatial context and visual perception for action. *Psicologica: International Journal of Methodology and Experimental Psychology*, 26(1):39–59. 98

Coello, Y., Bartolo, A., Amiri, B., Devanne, H., Houdayer, E., and Derambure, P. (2008). Perceiving what is reachable depends on motor representations: evidence from a transcranial magnetic stimulation study. *PLoS One*, 3(8):e2862. 98

Coello, Y., Bourgeois, J., and Iachini, T. (2012). Embodied perception of reachable space: how do we manage threatening objects? *Cognitive processing*, 13(1):131–135. 88, 98, 111, 112

Cohen, H. S. (2000). Vestibular disorders and impaired path integration along a linear trajectory. *Journal of Vestibular Research*, 10(1):7–15. 29

Comoli, E., Das Neves Favaro, P., Vautrelle, N., Leriche, M., Overton, P. G., and Redgrave, P. (2012). Segregated anatomical input to sub-regions of the rodent superior colliculus associated with approach and defense. *Frontiers in neuroanatomy*, 6:9. 102

Cools, R., Roberts, A. C., and Robbins, T. W. (2008). Serotoninergic regulation of emotional and behavioural control processes. *TRENDS in Cognitive Sciences*, 12(1). 75, 80

Coppin, G. and Sander, D. (2012). Contemporary theories and concepts in the psychology of emotions. In Pelachaud, C., editor, *Emotion-Oriented Systems*, pages 1–31. Wiley Online Library. 70, 71, 86

Corder, G. W. and Foreman, D. (2009). *Nonparametric statistics for non-statisticians: a step-by-step approach*. Wiley. 106, 135

Courgeon, M. and Clavel, C. (2013). Marc: a framework that features emotion models for facial animation during human–computer interaction. *Journal on Multimodal User Interfaces*, 7(4):311–319. 82

Courgeon, M., Clavel, C., and Martin, J.-C. (2009). Appraising emotional events during a real-time interactive game. In *Proceedings of the International Workshop on Affective-Aware Virtual Agents and Social Robots*, page 7. ACM. 82

Courgeon, M., Martin, J.-C., and Jacquemin, C. (2008). User's gestural exploration of different virtual agents' expressive profiles. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 3*, pages 1237–1240. International Foundation for Autonomous Agents and Multiagent Systems. 82

Cressant, A., Muller, R. U., and Poucet, B. (1997). Failure of centrally placed objects to control the firing fields of hippocampal place cells. *The Journal of Neuroscience*, 17(7):2531–2542. 29

Csikszentmihalyi, M. (1991). *Flow: The Psychology of Optimal Experience*. Harper Perennial. 116, 117, 120

Cuperlier, N., Gaussier, P., and Quoy, M. (2008). Interest of spatial context for a place cell based navigation model. In *From Animals to Animats 10*, pages 169–178. Springer. 46, 49

Cuperlier, N., Quoy, M., and Gaussier, P. (2006). Navigation and planning in an unknown environment using vision and a cognitive map. In *European Robotics Symposium 2006*, pages 129–142. Springer. 32, 176

Cuperlier, N., Quoy, M., and Gaussier, P. (2007). Neurobiologically inspired mobile robot navigation and planning. *Frontiers in neurorobotics*, 1:3. 32, 57

Damasio, A. R. (2003). *Looking for Spinoza: Joy, sorrow, and the feeling brain*. 70, 73, 74, 77, 78, 99

Damasio, A. R., Everitt, B. J., and Bishop, D. (1996). The somatic marker hypothesis and the possible functions of the prefrontal cortex [and discussion]. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 351(1346):1413–1420. 77

Damasio, A. R., Grabowski, T. J., Bechara, A., Damasio, H., Ponto, L. L., Parvizi, J., and Hichwa, R. D. (2000). Subcortical and cortical brain activity during the feeling of self-generated emotions. *Nature neuroscience*, 3(10):1049–1056. 78

Damasio, H., Grabowski, T., Frank, R., Galaburda, A. M., and Damasio, A. R. (1994). The return of Phineas Gage: clues about the brain from the skull of a famous patient. *Science*, 264(5162):1102–1105. 77

Dastani, M. and Lorini, E. (2012). A logic of emotions: from appraisal to coping. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pages 1133–1140. International Foundation for Autonomous Agents and Multiagent Systems. 81

Deci, E. L. and Ryan, R. M. (1975). *Intrinsic motivation*. Wiley Online Library. 116

Delarboulas, P., Gaussier, P., Quoy, M., and Caussy, R. (2014). Robustness study of a multimodal compass inspired form HD-cells and Dynamic neural fields. In *From Animals to Animats 13*. Springer. 33

Delevoye-Turrell, Y., Vienne, C., and Coello, Y. (2011). Space boundaries in schizophrenia. *Social Psychology*. 98

Denton, D., Shade, R., Zamarippa, F., Egan, G., Blair-West, J., McKinley, M., Lancaster, J., and Fox, P. (1999). Neuroimaging of genesis and satiation of thirst and an interoceptor-driven theory of origins of primary consciousness. *Proceedings of the National Academy of Sciences*, 96(9):5304–5309. 78

Deriche, R. (1987). Using Canny's criteria to derive a recursively implemented optimal edge detector. *International Journal of Computer Vision*, 1(2):167–187. 32, 39, 127

Di Chiara, G. (2002). Nucleus accumbens shell and core dopamine: differential role in behavior and addiction. *Behavioural brain research*, 137(1):75–114. 77

Diamond, A. (2013). Executive functions. *Annual review of psychology*, 64:135. 85

Dosey, M. A. and Meisels, M. (1969). Personal space and self-protection. *Journal of personality and social psychology*, 11(2):93. 98

Eacott, M. J. and Norman, G. (2004). Integrated memory for object, place, and context in rats: a possible model of episodic-like memory? *The Journal of Neuroscience*, 24(8):1948–1953. 54

Ekman, P. and Friesen, W. V. (1971). Constants across cultures in the face and emotion. *Journal of personality and social psychology*, 17(2):124. 71, 72, 82, 83, 115

Ekman, P., Levenson, R. W., and Friesen, W. V. (1983). Autonomic nervous system activity distinguishes among emotions. *Science*, 221(4616):1208–1210. 71

Erdem, U. M., Milford, M. J., and Hasselmo, M. E. (2015). A hierarchical model of goal directed navigation selects trajectories in a visual environment. *Neurobiology of learning and memory*, 117:109–121. 31, 35

Ethofer, T., Kreifelts, B., Wiethoff, S., Wolf, J., Grodd, W., Vuilleumier, P., and Wildgruber, D. (2009). Differential influences of emotion, task, and novelty on brain regions underlying the processing of speech melody. *Journal of cognitive neuroscience*, 21(7):1255–1268. 139

Etienne, A. S. and Jeffery, K. J. (2004). Path integration in mammals. *Hippocampus*, 14(2):180–192. 30

Etienne, A. S., Maurer, R., and Saucy, F. (1988). Limitations in the assessment of path dependent information. *Behaviour*, 106(1):81–110. 29

Fails, J. A. and Olsen, D. R. (2003). Interactive machine learning. In *Proceedings of the 8th international conference on Intelligent User Interfaces*, IUI '03, pages 39–45. ACM Press. 58, 65

Fanselow, M. S., DeCola, J. P., De Oca, B. M., and Landeira-Fernandez, J. (1995). Ventral and Dorsolateral Regions of the Midbrain Periaqueductal Gray (PAG) Control Different Stages of Defensive. *Aggressive Behavior*, 21:63–77. 75

Fellous, J.-M. (2004). From human emotions to robot emotions. *Architectures for Modeling Emotion: Cross-Disciplinary Foundations, American Association for Artificial Intelligence*, pages 39–46. 70, 71, 78, 90, 93, 99, 116

Franz, M. O. and Mallot, H. A. (2000). Biomimetic robot navigation. *Robotics and autonomous Systems*, 30(1):133–153. 28

Frings, C., Englert, J., Wentura, D., and Bermeitinger, C. (2010). Decomposing the emotional Stroop effect. *The quarterly journal of experimental psychology*, 63(1):42–49. 88

Gaussier, P., Banquet, J., Sargolini, F., Giovannangeli, C., Save, E., and Poucet, B. (2007). A model of grid cells involving extra hippocampal path integration, and the hippocampal loop. *Journal of integrative neuroscience*, 6(03):447–476. 32, 103, 139

Gaussier, P., Joulain, C., Banquet, J.-P., Leprêtre, S., and Revel, A. (2000). The visual homing problem: an example of robotics/biology cross fertilization. *Robotics and autonomous systems*, 30(1):155–180. 32

Gaussier, P., Joulain, C., Zrehen, S., Banquet, J.-P., and Revel, A. (1997). Visual navigation in an open environment without map. In *Intelligent Robots and Systems, 1997. IROS'97., Proceedings of the 1997 IEEE/RSJ International Conference on*, volume 2, pages 545–550. IEEE. 32

Gaussier, P., Prepin, K., and Nadel, J. (2004). Toward a cognitive system algebra: Application to facial expression learning and imitation. In *Embodied Artificial Intelligence*, pages 243–258. Springer. 119

Gaussier, P., Revel, A., Banquet, J.-P., and Babeau, V. (2002). From view cells and place cells to cognitive map learning: processing stages of the hippocampal system. *Biological cybernetics*, 86(1):15–28. 32, 35, 38, 47

Gaussier, P. and Zrehen, S. (1995). PerAc: A neural architecture to control artificial animals. *Robotics and Autonomous Systems*, 16(2-4):291–320. 22, 31, 32, 35, 38, 173

Gebhard, P. (2005). Alma: a layered model of affect. In *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pages 29–36. ACM. 81

Gibson, J. J. (1979). *The ecological approach to visual perception*. Psychology Press. 21

Giovannangeli, C. and Gaussier, P. (2007). Orientation system in robots: Merging allothetic and idiothetic estimations. In *13th International Conference on Advanced Robotics*, pages 349–354. 33

Giovannangeli, C. and Gaussier, P. (2008). Autonomous vision-based navigation: Goal-oriented action planning by transient states prediction, cognitive map building, and sensory-motor learning. In *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 676–683. IEEE. 32

Giovannangeli, C. and Gaussier, P. (2010). Interactive Teaching for Vision-Based Mobile Robots: A Sensory-Motor Approach. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, 40(1). 58, 65, 121

Giovannangeli, C., Gaussier, P., and Banquet, J. P. (2006). Robustness of visual place cells in dynamic indoor and outdoor environment. *International Journal of Advanced Robotic Systems*, 3(2):115–124. 32, 33, 34, 35, 38, 39, 47, 48

Gogate, L. J. and Bahrick, L. E. (1998). Intersensory redundancy facilitates learning of arbitrary relations between vowel sounds and objects in seven-month-old infants. *Journal of experimental child psychology*, 69(2):133–149. 59

Gold, P. E., Hankins, L., Edwards, R. M., Chester, J., and McGaugh, J. L. (1975). Memory interference and facilitation with posttrial amygdala stimulation: Effect on memory varies with footshock level. *Brain research*, 86(3):509–513. 89

Goodrich, M. A. and Schultz, A. C. (2007). Human-Robot Interaction: A Survey. *Foundations and Trends® in Human-Computer Interaction*, 1(3):203–275. 57, 58, 59

Grand, C., Mostafaoui, G., Hasnain, S. K., and Gaussier, P. (2014). Synchrony detection as a reinforcement signal for learning: Application to human robot interaction. In *International Conference on Timing and Time Perception*. 59, 61, 65

Grandjean, D. and Peters, C. (2011). Novelty processing and emotion: conceptual developments, empirical findings and virtual environments. In *Emotion-Oriented Systems*, pages 441–458. Springer. 115, 119, 139

Grandjean, D. and Scherer, K. R. (2008). Unpacking the cognitive architecture of emotion processes. *Emotion*, 8(3):341. 87, 93

Gray, J. R. (2001). Emotional modulation of cognitive control: Approach–withdrawal states double-dissociate spatial from verbal two-back task performance. *Journal of Experimental Psychology: General*, 130(3):436. 89

Gray, J. R., Braver, T. S., and Raichle, M. E. (2002). Integration of emotion and cognition in the lateral prefrontal cortex. *Proceedings of the National Academy of Sciences*, 99(6):4115–4120. 89, 137

Graziano, M. S., Hu, X. T., and Gross, C. G. (1997). Coding the locations of objects in the dark. *Science*, 277(5323):239–241. 100

Graziano, M. S. A. and Cooke, D. F. (2006). Parieto-frontal interactions, personal space, and defensive behavior. *Neuropsychologia*, 44(6):845–859. 98

Gross, J. J. (1998). The emerging field of emotion regulation: an integrative review. *Review of general psychology*, 2(3):271. 87

Haber, S. N. (2008). Functional anatomy and physiology of the basal ganglia: non-motor functions. In Tarsy, D., Vitek, J. L., Starr, P., and Okun, M., editors, *Deep Brain Stimulation in Neurological and Psychiatric Disorders*, pages 33–62. Springer. 77, 90

Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., and Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436(7052):801–806. 30

Hanoune, S. (2015). *Vers un modèle biologiquement plausible de la sélection de l'action pour un robot mobile* . PhD thesis. 29

Hanoune, S., Banquet, J. P., Gaussier, P., and Quoy, M. (2015). Cooperation/supervision of a habit by a cognitive strategy in a goal-directed navigational paradigm. In *Annual Computational Neuroscience Meeting CNS*. 57

Harnad, S. (1990). The symbol grounding problem. *Physica D: Nonlinear Phenomena*, 42(1-3):335–346. 20

Hasnain, S. K., Gaussier, P., and Mostafaoui, G. (2012). "Synchrony" as a way to choose an interacting partner. In *Proceedings of the 2nd Joint IEEE International Conference on Development and Learning and on Epigenetic Robotics (ICDL-EpiRob)*, pages 1–6. 59, 61, 65

Hasson, C. (2011). *Modélisation des mécanismes émotionnels pour un robot autonome : perspective développementale et sociale*. PhD thesis. 102

Hasson, C. and Gaussier, P. (2010). Path integration working memory for multi tasks dead reckoning and visual navigation. In *Proceedings of the 11th international conference on Simulation of adaptive behavior: from animals to animats*, pages 380–389. Springer-Verlag. 101

Hasson, C., Gaussier, P., and Boucenna, S. (2011). Emotions as a dynamical system: the interplay between the meta-control and communication function of emotions. *Paladyn*, 2(3):111–125. 70, 71, 92, 119

Hayman, R. M. and Jeffery, K. J. (2008). How heterogeneous place cell responding arises from homogeneous grids—a contextual gating hypothesis. *Hippocampus*, 18(12):1301–1313. 55

Hebb, D. O. (1949). *The organization of behavior: A neuropsychological theory*. Wiley and Sons. 20, 172

Hodgkin, A. L. and Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of physiology*, 117(4):500. 21

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8):2554–2558. 20

Iachini, T., Coello, Y., Frassinetti, F., and Ruggiero, G. (2014). Body space in social interactions: a comparison of reaching and comfort distance in immersive virtual reality. *PloS one*, 9(11):e111511. 98, 100, 111

Iachini, T., Coello, Y., Frassinetti, F., Senese, V. P., Galante, F., and Ruggiero, G. (2016). Peripersonal and interpersonal space in virtual and real environments: Effects of gender and age. *Journal of Environmental Psychology*, 45:154–164. 98, 111

Inzlicht, M., Bartholow, B. D., and Hirsh, J. B. (2015). Emotional foundations of cognitive control. *Trends in Cognitive Sciences*, 19(3):126–132. 78, 89, 99

Iriki, A., Tanaka, M., and Iwamura, Y. (1996). Coding of modified body schema during tool use by macaque postcentral neurones. *Neuroreport*, 7(14):2325–2330. 97, 98

Issartel, J., Marin, L., and Cadopi, M. (2007). Unintended interpersonal co-ordination: "can we march to the beat of our own drum?". *Neuroscience Letters*, 411(3):174 – 179. 59

Izard, C. E. (2010). The many meanings/aspects of emotion: Definitions, functions, activation, and regulation. *Emotion Review*, 2(4):363–370. 69, 70, 85

Jahanshahi, M., Obeso, I., Rothwell, J. C., and Obeso, J. A. (2015). A fronto-striato-subthalamic-pallidal network for goal-directed and habitual inhibition. *Nature Reviews Neuroscience*. 77, 90, 91

James, W. (1890). *The principles of psychology*. New York: Dover. 69

Jauffret, A. (2014). *De l'auto-évaluation aux émotions - approche neuromimétique et bayésienne de l'apprentissage de comportements complexes impliquant des informations multimodales* . PhD thesis. 32, 58, 174, 175

Jauffret, A., Cuperlier, N., and Gaussier, P. (2015). From grid cells and visual place cells to multimodal place cell: a new robotic architecture. *Frontiers in neurorobotics*, 9. 32, 35, 139

Jauffret, A., Cuperlier, N., Tarroux, P., and Gaussier, P. (2013a). From self-assessment to frustration, a small step toward autonomy in robotic navigation. *Frontiers in neurorobotics*, 7(16). 119

Jauffret, A., Grand, C., Cuperlier, N., Gaussier, P., and Tarroux, P. (2013b). How can a robot evaluate its own behavior? a neural model for self-assessment. In *Neural Networks (IJCNN), The 2013 International Joint Conference on*, pages 1–8. IEEE. 120, 121

Jones, B. E. (2005). From waking to sleeping: neuronal and chemical substrates. *Trends in pharmacological sciences*, 26(11):578–586. 80

Jung, M. and McNaughton, B. (1993). Spatial selectivity of unit activity in the hippocampal granular layer. *Hippocampus*, 3(2):165–182. 28

Kearns, M. J., Warren, W. H., Duchon, A. P., and Tarr, M. J. (2002). Path integration from optic flow and body senses in a homing task. *Perception*, 31(3):349–374. 29

Kelley, A., Bakshi, V., Haber, S., Steininger, T., Will, M., and Zhang, M. (2002). Opioid modulation of taste hedonics within the ventral striatum. *Physiology & behavior*, 76(3):365–377. 80

Kelley, A. E. (2005). Neurochemical networks encoding emotion and motivation. In Arbib, M. A. and Fellous, J.-M., editors, *Who Needs Emotions? The Brain Meets the Robot*, pages 31–77. Oxford University Press Oxford. 78

Kelley, S. A., Brownell, C. A., and Campbell, S. B. (2000). Mastery motivation and self-evaluative affect in toddlers: Longitudinal relations with maternal behavior. *Child development*, 71(4):1061–1071. 117

Kennedy, D. P., Gläscher, J., Tyszka, J. M., and Adolphs, R. (2009). Personal space regulation by the human amygdala. *Nature neuroscience*, 12(10):1226–1227. 88, 98

Kihl, O., Picard, D., and Gosselin, P.-H. (2015). A Unified framework for local visual descriptors evaluation. *Pattern Recognition*, 48(4):1174–1184. 39

Kim, C., Kroger, J. K., and Kim, J. (2011). A functional dissociation of conflict processing within anterior cingulate cortex. *Human brain mapping*, 32(2):304–312. 118

Knight, R. T. et al. (1996). Contribution of human hippocampal region to novelty detection. *Nature*, 383(6597):256–259. 139

Knox, W. B., Stone, P., and Breazeal, C. (2013). Training a Robot via Human Feedback: A Case Study. In *Proceeding of the 5th International Conference on Social Robotics*, pages 460–470. Springer. 58, 65

Koch, M. and Schnitzler, H. U. (1997). The acoustic startle response in rats–circuits mediating evocation, inhibition and potentiation. *Behavioural brain research*, 89(1-2):35–49. 93

Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological cybernetics*, 43(1):59–69. 20, 172

Kohonen, T. (1993). Physiological interpretation of the self-organizing map algorithm. *Neural Networks*, 6(7):895–905. 115

Kohonen, T. and Oja, E. (1976). Fast adaptive formation of orthogonalizing filters and associative memory in recurrent networks of neuron-like elements. *Biological Cybernetics*, 21(2):85–95. 115

Kong, H., Audibert, J.-Y., and Ponce, J. (2009). Vanishing point detection for road detection. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 96–103. IEEE. 120

Krichmar, J. L. (2008). The neuromodulatory system: a framework for survival and adaptive behavior in a challenging world. *Adaptive Behavior*, 16(6):385–399. 84

Krichmar, J. L. (2012). Design principles for biologically inspired cognitive robotics. *Biologically Inspired Cognitive Architectures*, 1:73–81. 22

Krichmar, J. L. (2013). A neurorobotic platform to test the influence of neuromodulatory signaling on anxious and curious behavior. *Frontiers in neurorobotics*, 7(February):1. 70, 84

Kunde, W. and Mauer, N. (2008). Sequential modulations of valence processing in the emotional stroop task. *Experimental psychology*, 55(3):151–156. 88

Kwang, T. and Swann, W. B. (2010). Do people embrace praise even when they feel unworthy? a review of critical tests of self-enhancement versus self-verification. *Personality and Social Psychology Review*. 117, 118

LaBar, K. S., Gatenby, J., Gore, J. C., LeDoux, J. E., and Phelps, E. A. (1998). Human Amygdala Activation during Conditioned Fear Acquisition and Extinction: a Mixed-Trial fMRI Study. *Neuron*, 20(5):937–945. 76

Lagarde, M., Andry, P., and Gaussier, P. (2008). Distributed real time neural networks in interactive complex systems. In *CSTST*, pages 95–100. 169

Lazarus, R. S. (1991). *Emotion and adaptation*. Oxford University Press. 86

LeDoux, J. (1996). *The emotional brain: The mysterious underpinnings of emotional life*. Simon and Schuster. 76

LeDoux, J. (2003). The emotional brain, fear, and the amygdala. *Cellular and molecular neurobiology*, 23(4-5):727–738. 76

LeDoux, J. (2012). Rethinking the emotional brain. *Neuron*, 73(4):653–676. 76

Leonard, J. J. and Durrant-Whyte, H. F. (1991). Simultaneous map building and localization for an autonomous mobile robot. In *Intelligent Robots and Systems' 91.'Intelligence for Mechanical Systems, Proceedings IROS'91. IEEE/RSJ International Workshop on*, pages 1442–1447. Ieee. 27

Lepora, N. F., Verschure, P., and Prescott, T. J. (2013). The state of the art in biomimetics. *Bioinspiration & biomimetics*, 8(1). 22

Lester, L. S. and Fanselow, M. S. (1985). Exposure to a cat produces opioid analgesia in rats. *Behavioral neuroscience*, 99(4):756. 80

Levitt, T. S. and Lawton, D. T. (1990). Qualitative navigation for mobile robots. *Artificial intelligence*, 44(3):305–360. 35

Lewis, M. D. (2005). Bridging emotion theory and neurobiology through dynamic systems modeling. *Behavioral and brain sciences*, 28(02):169–194. 76, 92, 99

Lindquist, K. A., Siegel, E. H., Quigley, K. S., and Barrett, L. F. (2013). The hundred-year emotion war: are emotions natural kinds or psychological constructions? Comment on Lench, Flores, and Bench (2011). 70

Lisin, D., Mattar, M., Blaschko, M. B., Learned-Miller, E. G., and Benfield, M. C. (2005). Combining local and global image features for object class recognition. In *Computer Vision and Pattern Recognition-Workshops, 2005. CVPR Workshops. IEEE Computer Society Conference on*, page 47. IEEE. 38, 52, 53

Lisman, J. E. and Otmakhova, N. A. (2001). Storage, recall, and novelty detection of sequences by the hippocampus: elaborating on the socratic model to account for normal and aberrant effects of dopamine. *Hippocampus*, 11(5):551–568. 139

Little, K. B. (1965). Personal space. *Journal of Experimental Social Psychology*, 1(3):237–247. 97

Lloyd, D. M. (2009). The space between us: A neurophilosophical framework for the investigation of human interpersonal space. *Neuroscience & Biobehavioral Reviews*, 33(3):297–304. 97, 98

Lones, J., Lewis, M., and Cañamero, L. (2014). Hormonal modulation of interaction between autonomous agents. In *2014 IEEE 4th Joint International Conference on Development and Learning and Epigenetic Robotics, ICDL-Epirob 2014*. 83

Longo, M. R. and Lourenco, S. F. (2006). On the nature of near space: Effects of tool use and the transition to far space. *Neuropsychologia*, 44(6):977–981. 97, 98

Lourenco, S. F. and Longo, M. R. (2009). The plasticity of near space: Evidence for contraction. *Cognition*, 112(3):451–456. 97, 98

Lourenco, S. F., Longo, M. R., and Pathman, T. (2011). Near space and its relation to claustrophobic fear. *Cognition*, 119(3):448–453. 98, 111

Lowe, D. G. (2004). Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60(2):91–110. 38

Lowry, S., Sünderhauf, N., Newman, P., Leonard, J. J., Cox, D., Corke, P., and Milford, M. J. (2016). Visual place recognition: A survey. *IEEE Transactions on Robotics*, 32(1):1–19. 37

Ma, B., Lakshmanan, S., and Hero, A. O. (2000). Simultaneous detection of lane and pavement boundaries using model-based multisensor fusion. *IEEE Transactions on Intelligent Transportation Systems*, 1(3):135–147. 120

Maass, W. (1997). Networks of spiking neurons: the third generation of neural network models. *Neural networks*, 10(9):1659–1671. 21

Maass, W. (2000). On the computational power of winner-take-all. *Neural computation*, 12(11):2519–2535. 174

MacLean, P. D. (1949). Psychosomatic Disease and the" Visceral Brain": Recent Developments Bearing on the Papez Theory of Emotion. *Psychosomatic medicine*, 11(6):338–353. 74

MacQueen, J. (1967). Some methods for classification and analysis of multivariate observations. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, volume 1, pages 281–297. Oakland, CA, USA. 172

Mann, J. J., Malone, K. M., Sweeney, J. A., Brown, R. P., Linnoila, M., Stanley, B., and Stanley, M. (1996). Attempted suicide characteristics and cerebrospinal fluid amine metabolites in depressed inpatients. *Neuropsychopharmacology*, 15(6):576–586. 80

Marin, L., Issartel, J., and Chaminade, T. (2009). Interpersonal motor coordination: From human-human to human-robot interactions. *Interaction Studies*, 10(3):479–504. 59

Marsella, S., Gratch, J., and Petta, P. (2010). Computational models of emotion. In Scherer, K. R., Bänziger, T., and Roesch, E., editors, *A Blueprint for Affective Computing-A sourcebook and manual*, pages 21–46. 70, 71

Marsland, S. (2003). Novelty detection in learning systems. *Neural computing surveys*, 3(2):157–195. 115

Marsland, S., Nehmzow, U., and Shapiro, J. (2005). On-line novelty detection for autonomous mobile robots. *Robotics and Autonomous Systems*, 51(2):191–206. 116

Matsumoto, D. and Ekman, P. (2004). The relationship among expressions, labels, and descriptions of contempt. *Journal of personality and social psychology*, 87(4):529. 71, 72

Maturana, H. R. and Varela, F. J. (1980). *Autopoiesis and cognition: The realization of the living*, volume 42. 21

McCulloch, W. S. and Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, 5(4):115–133. 19, 20, 172

McGaugh, J. L. (2000). Memory–a century of consolidation. *Science*, 287(5451):248–251. 89

McGaugh, J. L. (2015). Consolidating memories. *Annual review of psychology*, 66:1–24. 89, 137

McKenna, F. P. and Sharma, D. (2004). Reversing the emotional stroop effect reveals that it is not what it seems: the role of fast and slow components. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30(2):382. 88

Mehlman, P., Higley, J., Faucher, I., Lilly, A., Taub, D., Vickers, J., Suomi, S., and Linnoila, M. (1994). Low csf 5-hiaa concentrations and severe aggression and impaired impulse control in nonhuman primates. *American Journal of Psychiatry*, 151:1485–1491. 80

Metcalfe, T. and Gresty, M. (1992). Self-controlled reorienting movements in response to rotational displacements in normal subjects and patients with labyrinthine disease. *Annals of the New York Academy of Sciences*, 656(1):695–698. 29

Middleton, F. A. and L, S. P. (2001). A revised neuroanatomy of frontal-subcortical circuits. In Lichter, D. G. and Cummings, J. L., editors, *Frontal-subcortical circuits in psychiatric and neurological disorders*, page 44. Guilford Press. 77, 90

Milford, M., Scheirer, W., Vig, E., Glover, A., Baumann, O., Mattingley, J., and Cox, D. (2014). Condition-invariant, top-down visual place recognition. In *ICRA 2014*, pages 5571–5577. IEEE. 37, 38, 39, 40, 52

Milford, M. J. and Wyeth, G. F. (2008). Mapping a suburb with a single camera using a biologically inspired SLAM system. *Robotics, IEEE Transactions on*, 24(5):1038–1053. 31, 35, 38, 40, 50

Milford, M. J., Wyeth, G. F., and Prasser, D. (2004). Ratslam: a hippocampal model for simultaneous localization and mapping. In *Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on*, volume 1, pages 403–408. IEEE. 31

Milham, M., Banich, M., Webb, A., Barad, V., Cohen, N., Wszalek, T., and Kramer, A. (2001). The relative involvement of anterior cingulate and prefrontal cortex in attentional control depends on nature of conflict. *Cognitive Brain Research*, 12(3):467–473. 118

Miller, A. M. P., Vedder, L. C., Law, L. M., and Smith, D. M. (2014). Cues, context, and long-term memory: the role of the retrosplenial cortex in spatial cognition. *Frontiers in human neuroscience*, 8. 55

Minsky, M. (1986). *Society of mind*. Simon and Schuster. 83

Minsky, M. and Papert, S. (1969). Perceptrons. 20

Mirolli, M. and Baldassarre, G. (2013). Functions and mechanisms of intrinsic motivations. In *Intrinsically Motivated Learning in Natural and Artificial Systems*, pages 49–72. Springer. 118

Mittelstaedt, H. and Mittelstaedt, M.-L. (1982). Homing by path integration. In *Avian navigation*, pages 290–297. Springer. 29

Mizumori, S. J., Smith, D. M., and Puryear, C. B. (2007). Hippocampal and neocortical interactions during context discrimination: electrophysiological evidence from the rat. *Hippocampus*, 17(9):851–862. 54, 55

Moran, J. M., Heatherton, T. F., and Kelley, W. M. (2009). Modulation of cortical midline structures by implicit and explicit self-relevance evaluation. *Social neuroscience*, 4(3):197–211. 117

Morgado, L. and Gaspar, G. (2004). Unifying Emotion and Cognition: The Flow Model Approach. In *FLAIRS Conference*, pages 245–250. 82

Morgado, L. and Gaspar, G. (2007). Abstraction level regulation of cognitive processing through emotion-based attention mechanisms. In *Attention in Cognitive Systems. Theories and Systems from an Interdisciplinary Viewpoint*, pages 59–74. Springer. 82

Muller, R. U. and Kubie, J. L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *The Journal of Neuroscience*, 7(7):1951–1968. 29

Nadel, J., Carchon, I., Kervella, C., Marcelli, D., and Reserbat-Plantey, D. (1999). Expectancies for social contingency in 2-month-olds. *Developmental Science*, 2(2):164–173. 59

Nair, R., Tambe, M., and Marsella, S. (2005). The role of emotions in multiagent teamwork. In Arbib, M. A. and Fellous, J.-M., editors, *Who Needs Emotions? The Brain Meets the Robot*, pages 311–329. Oxford University Press Oxford. 70

Navon, D. (1977). Forest before trees: The precedence of global features in visual perception. *Cognitive psychology*, 9(3):353–383. 38, 46

Nicolescu, M. and Mataric, M. J. (2006). Task learning through imitation and human-robot interaction. *Models and Mechanisms of Imitation and Social Learning in Robots, Humans and Animals: Behavioural, Social and Communicative Dimensions*, pages 407–424. 58, 65

Nieuwenhuis, S., Gilzenrat, M. S., Holmes, B. D., and Cohen, J. D. (2005). The role of the locus coeruleus in mediating the attentional blink: a neurocomputational theory. *Journal of Experimental Psychology: General*, 134(3):291. 89

Oatley, K. and Johnson-Laird, P. N. (1987). Towards a cognitive theory of emotions. *Cognition and emotion*, 1(1):29–50. 86

Ochsner, K. N. and Gross, J. J. (2005). The cognitive control of emotion. *Trends in cognitive sciences*, 9(5):242–249. 87

Ognibene, D. and Baldassare, G. (2015). Ecological Active Vision: Four Bioinspired Principles to Integrate Bottom–Up and Adaptive Top–Down Attention Tested With a Simple Camera-Arm Robot. *Autonomous Mental Development, IEEE Transactions on*, 7(1):3–25. 138

Öhman, A., Flykt, A., and Esteves, F. (2001). Emotion drives attention: detecting the snake in the grass. *Journal of experimental psychology: general*, 130(3):466. 88, 93

O'keefe, J. and Conway, D. (1978). Hippocampal place units in the freely moving rat: why they fire where they fire. *Experimental Brain Research*, 31(4):573–590. 29

O'Keefe, J. and Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain research*, 34(1):171–175. 28

Oker, A., Prigent, E., Courgeon, M., Eyharabide, V., Urbach, M., Bazin, N., Amorim, M.-A., Passerieux, C., Martin, J.-C., and Brunet-Gouet, E. (2015). How and why affective and reactive virtual agents will bring new insights on social cognitive disorders in schizophrenia? an illustration with a virtual card game paradigm. *Frontiers in human neuroscience*, 9. 71

Oliva, A. and Torralba, A. (2006). Building the gist of a scene: The role of global image features in recognition. *Progress in brain research*, 155:23–36. 38, 46, 47

Öngür, D. and Price, J. L. (2000). The organization of networks within the orbital and medial prefrontal cortex of rats, monkeys and humans. *Cerebral cortex*, 10(3):206–219. 77

O'Regan, J. K. and Noë, A. (2001). What it is like to see: A sensorimotor theory of perceptual experience. *Synthese*, 129(1):79–103. 21

Ortony, A., Clore, G. L., and Collins, A. (1988). *The Cognitive Structure of Emotions*. 70, 72, 81, 86

Oudeyer, P.-Y., Kaplan, F., and Hafner, V. V. (2007). Intrinsic motivation systems for autonomous mental development. *Evolutionary Computation, IEEE Transactions on*, 11(2):265–286. 116, 118, 120

Padmala, S. and Pessoa, L. (2008). Affective learning enhances visual detection and responses in primary visual cortex. *The Journal of Neuroscience*, 28(24):6202–6210. 88

Panksepp, J., Herman, B., Vilberg, T., Bishop, P., and DeEskinazi, F. (1980). Endogenous opioids and social behavior. *Neuroscience & Biobehavioral Reviews*, 4(4):473–487. 80

Papez, J. W. (1937). A proposed mechanism of emotion. *Archives of Neurology & Psychiatry*, 38(4):725–743. 74, 76

Paton, J. J., Belova, M. A., Morrison, S. E., and Salzman, C. D. (2006). The primate amygdala represents the positive and negative value of visual stimuli during learning. *Nature*, 439(7078):865–870. 76

Pavlov, I. P. (1927). Conditioned reflexes: An investigation of the physiological activity of the cerebral cortex. *Annals of Neurosciences*. 173

Peciña, S. and Berridge, K. C. (2005). Hedonic hot spot in nucleus accumbens shell: where do μ-opioids cause increased hedonic impact of sweetness? *The Journal of neuroscience*, 25(50):11777–11786. 77, 80

Perlstein, W. M., Elbert, T., and Stenger, V. A. (2002). Dissociation in human prefrontal cortex of affective influences on working memory-related activity. *Proceedings of the National Academy of Sciences*, 99(3):1736–1741. 89

Pessoa, L. (2008). On the relationship between emotion and cognition. *Nature Reviews Neuroscience*, 9(2):148–158. 75, 89, 90

Pessoa, L. (2010). Emergent processes in cognitive-emotional interactions. *Dialogues in clinical neuroscience*, 12(4):433. 76, 89

Pessoa, L. and Adolphs, R. (2010). Emotion processing and the amygdala: from a'low road'to'many roads' of evaluating biological significance. *Nature reviews neuroscience*, 11(11):773–783. 77

Pfeifer, R., Lungarella, M., and Iida, F. (2007). Self-organization, embodiment, and biologically inspired robotics. *science*, 318(5853):1088–1093. 22

Pfeifer, R. and Pitti, A. (2012). La révolution de l'intelligence du corps. 9

Pfeifer, R. and Scheier, C. (1995). From perception to action: The right direction? In Gaussier, P. and Nicoud, J., editors, *From perception to action*. IEEE Computer Society Press. 22

Phelps, E. A., Ling, S., and Carrasco, M. (2006). Emotion facilitates perception and potentiates the perceptual benefits of attention. *Psychological science*, 17(4):292–299. 88

Piaget, J. (1951). *Play, dreams and imitation in childhood*, volume 25. Routledge. 21

Picard, R. W. (1997). *Affective computing*, volume 252. MIT press Cambridge. 70

Posner, J., Russell, J. A., and Peterson, B. S. (2005). The circumplex model of affect: an integrative approach to affective neuroscience, cognitive development, and psychopathology. *Development and psychopathology*, 17(3):715–734. 72

Prepin, K. and Gaussier, P. (2010). How an agent can detect and use synchrony parameter of its own interaction with a human? In *Development of Multimodal Interfaces: Active Listening and Synchrony*, pages 50–65. Springer. 59

Previc, F. H. (1998). The neuropsychology of 3-d space. *Psychological bulletin*, 124(2):123. 112

Quirk, G. J., Muller, R. U., and Kubie, J. L. (1990). The firing of hippocampal place cells in the dark depends on the rat's recent experience. *The Journal of Neuroscience*, 10(6):2008–2017. 29

Quirk, G. J., Muller, R. U., Kubie, J. L., and Ranck, J. (1992). The positional firing properties of medial entorhinal neurons: description and comparison with hippocampal place cells. *The Journal of Neuroscience*, 12(5):1945–1963. 28

Quoy, M., Moga, S., and Gaussier, P. (2003). Dynamical neural networks for planning and low-level robot control. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 33(4):523–532. 176

Quoy, M., Moga, S., Gaussier, P., and Revel, A. (2000). Parallelization of neural networks using pvm. In *European Parallel Virtual Machine/Message Passing Interface Users' Group Meeting*, pages 289–296. Springer. 169

Raleigh, M. J. and McGuire, M. T. (1991). Bidirectional relationships between tryptophan and social behavior in vervet monkeys. In *Kynurenine and Serotonin Pathways: Progress in Tryptophan Research*, pages 289–298. Springer. 80

Ranck Jr, J. (1984). Head-direction cells in the deep cell layers of dorsal presubiculum in freely moving rats. In *Soc Neurosci Abstr*, volume 10. 30

Rebai, K., Azouaoui, O., and Achour, N. (2014). HS Combined Histogram for Visual Memory Building and Scene Recognition in Outdoor Environments. *ICRA 2014*. 38, 40

Redish, A. D. and Touretzky, D. S. (1997). Cognitive maps beyond the hippocampus. *Hippocampus*, 7(1):15–35. 31

Revel, A. (1997). *Contrôle d'un robot autonome par approche neuro-mimétique*. PhD thesis. 120

Revel, A. and Andry, P. (2009). Emergence of structured interactions: From a theoretical model to pragmatic robotics. *Neural networks*, 22(2):116–125. 61

Riesenhuber, M. and Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nature neuroscience*, 2(11):1019–1025. 38

Rizzolatti, G., Fadiga, L., Fogassi, L., and Gallese, V. (1997). The space around us. *Science*, 277(5323):190–191. 97, 98

Robinson, T. E. and Berridge, K. C. (1993). The neural basis of drug craving: an incentive-sensitization theory of addiction. *Brain research reviews*, 18(3):247–291. 79

Rochat, P. and Striano, T. (2000). Perceived self in infancy. *Infant Behavior and Development*, 23(3 - 4):513 – 530. 59

Rolls, E. (1999). *Emotion and the Brain*. Oxford: Oxford University Press. 73

Rolls, E. (2005). What are emotions, why do we have emotions, and what is their computational basis in the brain. In Arbib, M. A. and Fellous, J.-M., editors, *Who Needs Emotions? The Brain Meets the Robot*, pages 117–146. Oxford University Press Oxford. 70, 73, 116

Rolls, E. T. (2004). The functions of the orbitofrontal cortex. *Brain and cognition*, 55(1):11–29. 78

Rosenblatt, F. (1958). The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological review*, 65(6):386. 20, 172

Rostami, V., Ramli, A. R., and Sojodishijani, O. (2014). Integration of Global and Local Salient Features for Scene Modeling in Mobile Robot Applications. *Journal of Intelligent & Robotic Systems*, 75(3-4):443–456. 38, 52

Rumelhart, D. E., Hinton, G. E., and Williams, R. J. (1985). Learning internal representations by error propagation. pages 318–362. 20

Rumelhart, D. E. and McClelland, J. L. (1986). *Parallel distributed processing*, volume 1. MIT Press. 20

Rumelhart, D. E. and Zipser, D. (1985). Feature discovery by competitive learning. *Cognitive science*, 9(1):75–112. 174

Russell, J. A. (1980). A circumplex model of affect. *Personality and Social Psychology*, pages 1161–1178. 72

Russell, J. A. and Barrett, L. F. (1999). Core affect, prototypical emotional episodes, and other things called emotion: dissecting the elephant. *Journal of personality and social psychology*, 76(5):805. 72, 99

Russell, J. A. and Mehrabian, A. (1977). Evidence for a three-factor theory of emotions. *Journal of research in Personality*, 11(3):273–294. 72, 81, 82, 83

Sander, D., Grandjean, D., Pourtois, G., Schwartz, S., Seghier, M. L., Scherer, K. R., and Vuilleumier, P. (2005). Emotion and attention interactions in social cognition: brain regions involved in processing anger prosody. *Neuroimage*, 28(4):848–858. 87

Sargolini, F., Fyhn, M., Hafting, T., McNaughton, B. L., Witter, M. P., Moser, M.-B., and Moser, E. I. (2006). Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science*, 312(5774):758–762. 30, 138

Saunders, J., Nehaniv, C. L., and Dautenhahn, K. (2006). Teaching Robots by Moulding Behavior and Scaffolding the Environment. In *Proceedings of the 1st ACM Conference on Human-robot Interaction*, pages 118–125. 58, 65

Schembri, M., Mirolli, M., and Baldassarre, G. (2007). Evolving internal reinforcers for an intrinsically motivated reinforcement-learning robot. In *2007 IEEE 6th International Conference on Development and Learning*, pages 282–287. IEEE. 118

Scherer, K. R. (1984). On the nature and function of emotion: A component process approach. In Scherer, K. R., , and Ekman, P., editors, *Approaches to emotion*, volume 2293, page 317. 73, 82

Scherer, K. R. (2001). Appraisal considered as a process of multilevel sequential checking. *Appraisal processes in emotion: Theory, methods, research*, 92:120. 73, 82, 87, 93, 99, 139

Scherer, K. R. (2009). Emotions are emergent processes: they require a dynamic computational architecture. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, 364(1535):3459–3474. 70, 71, 73, 86, 115, 117

Scherer, K. R. (2010). The Component Process Model: architecture for a comprehensive computational model of emergent emotion. In Scherer, K. R., Bänziger, T., and Roesch, E., editors, *A Blueprint for Affective Computing*, pages 47–70. 73, 86, 93

Schmidhuber, J. (1991). Curious model-building control systems. In *International Joint Conference on Neural Networks*, pages 1458–1463. IEEE. 116, 118

Schöner, G., Dose, M., and Engels, C. (1995). Dynamics of behavior: Theory and applications for autonomous robot architectures. *Robotics and autonomous systems*, 16(2):213–245. 100, 103, 124, 174, 176

Schultz, W. (2004). Neural coding of basic reward terms of animal learning theory, game theory, microeconomics and behavioural ecology. *Current opinion in neurobiology*, 14(2):139–147. 79

Schwabe, L., Merz, C. J., Walter, B., Vaitl, D., Wolf, O. T., and Stark, R. (2011). Emotional modulation of the attentional blink: the neural structures involved in capturing and holding attention. *Neuropsychologia*, 49(3):416–425. 88

Schwartz, E. L. (1977). Spatial mapping in the primate sensory projection: analytic structure and relevance to perception. *Biological cybernetics*, 25(4):181–194. 39, 128

Scott, D. J., Heitzeg, M. M., Koeppe, R. A., Stohler, C. S., and Zubieta, J.-K. (2006). Variations in the human pain stress experience mediated by ventral and dorsal basal ganglia dopamine activity. *The Journal of Neuroscience*, 26(42):10789–10795. 79

Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and brain sciences*, 3(03):417–424. 20

Sedikides, C. and Strube, M. J. (1995). The multiply motivated self. *Personality and Social Psychology Bulletin*, 21(12):1330–1335. 117, 118

Shackman, A. J., Salomons, T. V., Slagter, H. A., Fox, A. S., Winter, J. J., and Davidson, R. J. (2011). The integration of negative affect, pain and cognitive control in the cingulate cortex. *Nature Reviews Neuroscience*, 12(3):154–167. 78

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, pages 379–423. 19

Siagian, C. and Itti, L. (2007). Rapid biologically-inspired scene classification using features shared with visual attention. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 29(2):300–312. 39, 40, 46

Siagian, C. and Itti, L. (2009). Biologically inspired mobile robot vision localization. *Robotics, IEEE Transactions on*, 25(4):861–873. 35, 38, 39, 40, 45, 46, 52

Smith, D. M. (2008). The hippocampus, context processing and episodic memory. *Handbook of behavioral neuroscience*, 18:465–630. 46, 54, 55

Smith, D. M., Barredo, J., and Mizumori, S. J. (2012). Complimentary roles of the hippocampus and retrosplenial cortex in behavioral context discrimination. *Hippocampus*, 22(5):1121–1133. 28, 55

Smith, R. C. and Cheeseman, P. (1986). On the representation and estimation of spatial uncertainty. *The international journal of Robotics Research*, 5(4):56–68. 27

Spackman, M. P. (2004). Can machines adequately simulate human emotion? a test of four theories of emotion. *Theory & Psychology*, 14(6):755–776. 70

Sparbert, J., Dietmayer, K., and Streller, D. (2001). Lane detection and street type classification using laser range images. In *Intelligent Transportation Systems, 2001. Proceedings. 2001 IEEE*, pages 454–459. IEEE. 120

Stackman, R. W. and Herbert, A. M. (2002). Rats with lesions of the vestibular system require a visual landmark for spatial navigation. *Behavioural brain research*, 128(1):27–40. 29

Stipek, D., Recchia, S., McClintic, S., and Lewis, M. (1992). Self-evaluation in young children. *Monographs of the society for research in child development*, pages i–95. 117

Sutton, R. S. and Barto, A. G. (1998). *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge. 118

Tajadura-Jiménez, A., Pantelidou, G., Rebacz, P., Västfjäll, D., and Tsakiris, M. (2011). I-space: The effects of emotional valence and source of music on interpersonal distance. *PLoS ONE*, 6(10). 88, 97, 98, 111, 112

Taube, J. S. (2007). The head direction signal: origins and sensory-motor integration. *Annu. Rev. Neurosci.*, 30:181–207. 30

Taube, J. S., Goodridge, J. P., Golob, E. J., Dudchenko, P. A., and Stackman, R. W. (1996). Processing the head direction cell signal: a review and commentary. *Brain research bulletin*, 40(5):477–484. 30

Taube, J. S., Muller, R. U., and Ranck, J. B. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats. i. description and quantitative analysis. *The Journal of neuroscience*, 10(2):420–435. 112, 174

Taylor, S. E., Neter, E., and Wayment, H. A. (1995). Self-evaluation processes. *Personality and Social Psychology Bulletin*, 21(12):1278–1287. 117

Thompson, L. and Best, P. (1990). Long-term stability of the place-field activity of single units recorded from the dorsal hippocampus of freely behaving rats. *Brain research*, 509(2):299–308. 28

Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59(236):433–460. 19

Ungerstedt, U. (1971). Adipsia and aphagia after 6-hydroxydopamine induced degeneration of the nigro-striatal dopamine system. *Acta Physiologica Scandinavica*, 82(S367):95–122. 79

Valdés-Conroy, B., Román, F. J., Hinojosa, J. A., and Shorkey, S. P. (2012). So far so good: Emotion in the peripersonal/extrapersonal space. *PLoS one*, 7(11):e49162. 88, 98, 111, 112

Van de Sande, K. E. A., Gevers, T., and Snoek, C. G. M. (2008). A comparison of color features for visual concept classification. In *Proceedings of the 2008 international conference on Content-based image and video retrieval*, pages 141–150. ACM. 38, 39, 40, 45

Vankov, A., Hervé-Minvielle, A., and Sara, S. J. (1995). Response to novelty and its rapid habituation in locus coeruleus neurons of the freely exploring rat. *European Journal of Neuroscience*, 7(6):1180–1187. 80

Varela, J. F., Evan, T., and Eleanor, R. (1992). *The embodied mind: cognitive science and human experience*. MIT Press. 21

Velásquez, J. D. (1997). Modeling emotions and other motivations in synthetic agents. In *AAAI/IAAI*, pages 10–15. 83

Wagner, D. D., Haxby, J. V., and Heatherton, T. F. (2012). The representation of self and person knowledge in the medial prefrontal cortex. *Wiley Interdisciplinary Reviews: Cognitive Science*, 3(4):451–470. 117

Wallace, S. A. (2005). S-assess: a library for behavioral self-assessment. In *Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pages 256–263. ACM. 118

Walter, W. G. (1950). An imitation of life. *Scientific American*. 19, 27

Wang, Y., Teoh, E. K., and Shen, D. (2004). Lane detection and tracking using b-snake. *Image and Vision computing*, 22(4):269–280. 120

Wehner, R. and Srinivasan, M. V. (1981). Searching behaviour of desert ants, genuscataglyphis (formicidae, hymenoptera). *Journal of comparative physiology*, 142(3):315–338. 29

Wendt, J., Weike, A. I., Lotze, M., and Hamm, A. O. (2011). The functional connectivity between amygdala and extrastriate visual cortex activity during emotional picture processing depends on stimulus novelty. *Biological psychology*, 86(3):203–209. 139

Weymar, M. and Schwabe, L. (2016). Amygdala and Emotion: The Bright Side of It. *Frontiers in neuroscience*, 10:224. 76

Widrow, B. and Hoff, M. E. (1960). Adaptive switching circuits. In *IRE WESCON convention record*, volume 4, pages 96–104. New York. 62, 109, 119, 173

Wiener, N. (1948). *Cybernetics: Control and communication in the animal and the machine*. Wiley New York. 19

Williams, J. M. G., Mathews, A., and MacLeod, C. (1996). The emotional Stroop task and psychopathology. *Psychological bulletin*, 120(1):3. 88

Wilson, H. R. and Cowan, J. D. (1973). A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. *Kybernetik*, 13(2):55–80. 174

Wilson, M. and Mcnaughton, B. (1993). Dynamics of the hippocampal ensemble code for space. *Science*. 28

Wolpert, L. (2002). Love is a many-moleculed thing. *The Observer*, 47. 27

Wood, E. R., Dudchenko, P. A., Robitsek, R. J., and Eichenbaum, H. (2000). Hippocampal neurons encode information about different types of memory episodes occurring in the same location. *Neuron*, 27(3):623–633. 54

Wystrach, A. and Graham, P. (2012). What can we learn from studies of insect navigation? *Animal Behaviour*, 84(1):13–20. 27

Yan, S., Wang, H., Tang, X., and Huang, T. (2007). Exploring feature descritors for face recognition. In *Acoustics, Speech and Signal Processing, 2007. ICASSP 2007. IEEE International Conference on*, volume 1, pages I—-629. IEEE. 39

The neurocybermagic hat: hardware, software and algorithmic tricks

A.1   Robotic platforms

Figure A.1 shown the three mobile robots that I used during my thesis. They are customized versions of robotic platforms made by Robosoft©. Robulab is the indoor platform (LEFT and CENTER) and Roburoc is the robot used in the outdoors (RIGHT). In all cases, the platforms are equipped with embedded computers – Intel©i7, 6-core 12-thread 3.33GHz CPU and 16 GB of RAM – used to run the programs designed for the experiments. These embedded computers communicate with the mobile base provided by Robosoft ©via Ethernet connection. Depending on the experiments, the platforms are also equipped with a variety of sensors (e.g. camera, laser, compass) and actuators (e.g. servomotor) used for physical and social interactions with the environment.

A.2   Robot simulator

Some of the experiments in Chapters 6 and 7 were performed on Webots simulator developed by Cyberbotics ©. Webots allows for simulating environments with realistic physics. The indoor platform described above was modeled by former doctoral students of the ETIS lab so that experiments could be run on the simulator regardless of the use of real or simulated robots. Figure A.2 shows the simulated robot and an example of physical environment.

## A.3   Software environment for neural network simulations

A.3.1   Prométhée

The *Prométhée* software developed by Gaussier and colleagues at the ETIS laboratory is a neural network simulator that allows distributed execution of independent computational nodes (groups of neurons) [Quoy et al., 2000; Lagarde et al., 2008]. These nodes are functions written in C language and implementing a variety of features: neural algorithms, arithmetic operations, image processing algorithms, I/O interfaces, etc. The most important module of *Prométhée* is a
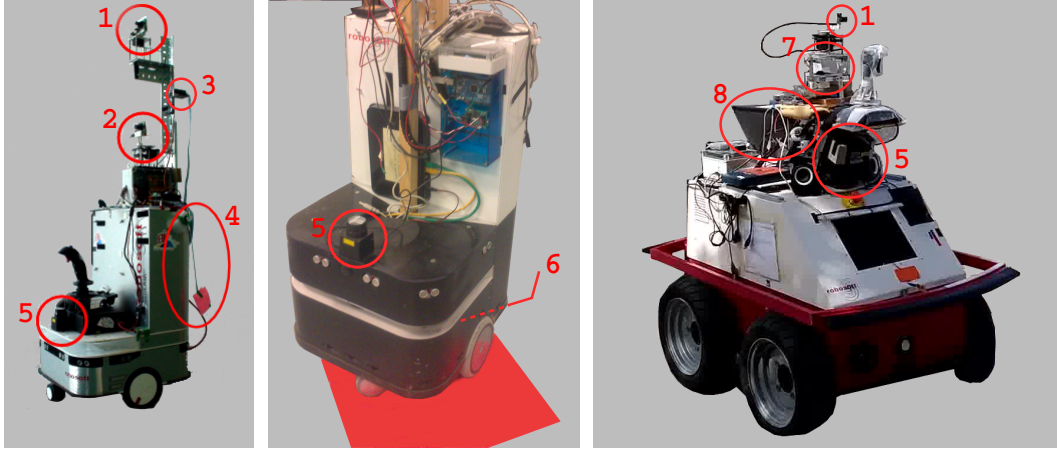
Figure A.1: Robotic platforms. **LEFT**: Indoor platform used in Chapter 3. **CENTER**: Indoor platform used in Chapter 6. **RIGHT**: Outdoor platform used in Chapter 2. These robot are embedded with various devices: (1) and (2) pan-camera, (3) magnetic compass, (4) servomotor-based oscillating arm, (5) laser sensor, (6) light sensor (color detector placed underneath the robot), (7) compass and stabilizing platform. The laptop (8) on the outdoor platform is used to launch and debug the programs running on the embedded computer.
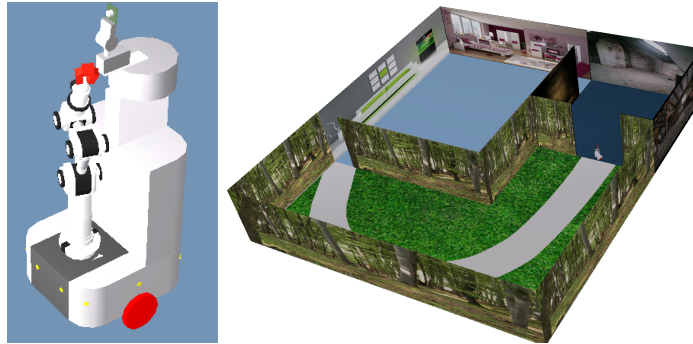
Figure A.2: Webots robot simulator. **LEFT**: The simulated Robulab with 2 driving wheels. **RIGHT**: The environment used in experiments of Chapter 7.

scheduler handling the execution of these nodes. Much like a *Petri* net, the scheduling is based on *tokens* transmitted from a node to its successors once it is executed. Each operation of the information processing flow can be computed as soon as the information from previous modules is updated. Independent modules are executed in parallel (i.e. in separate threads). It is also possible to use *real-time tokens* that guarantee that activities inside the network are updated at a specific frequency.

A.3.2 Coeos

The neural networks executed by Prométhée are organized in the form of *scripts* specifying the topology of connections between the computational nodes. *Coeos* is a GUI allowing users to design such scripts. Figure A.3 shows an example of network as it can be seen on this interface.
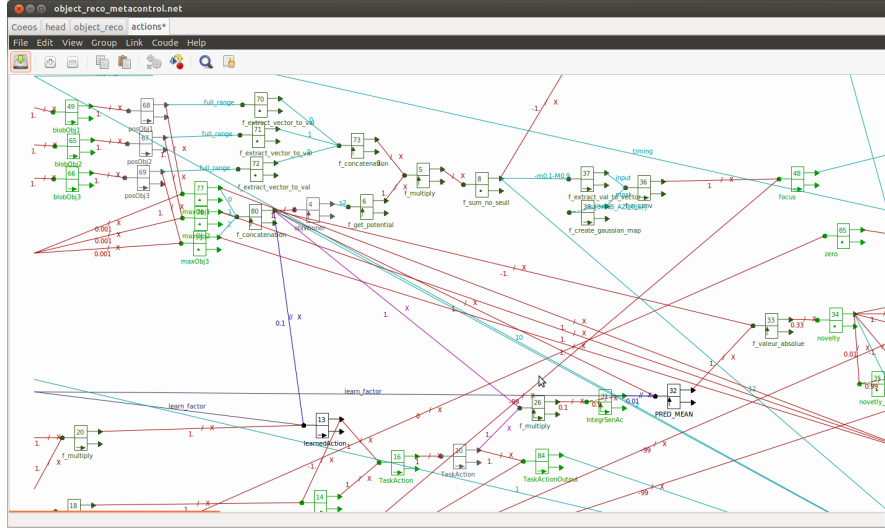
Figure A.3: Coeos GUI for the design of neural networks executed by Prométhée.

A.3.3 Themis

An entire application can be decomposed in several scripts. These scripts can communicate with each other, thus forming a global network. *Themis* is a GUI allowing users to easily manage such applications. Its gives access to all parameters related to the communication between scripts. Also, it controls the execution of these scripts by Prométhée.
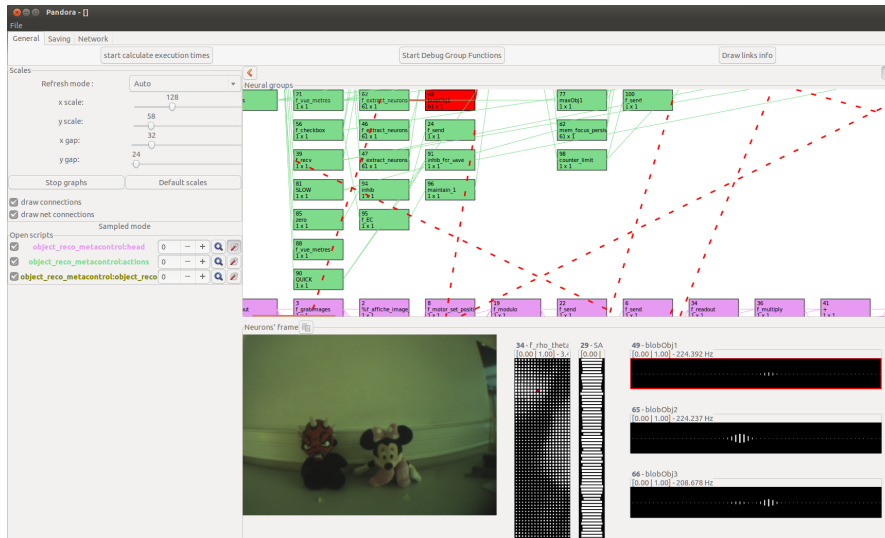
A.3.4 Pandora

Figure A.4: Pandora GUI for the display of neural networks executed by Prométhée..

Pandora is a GUI allowing users to display the activities of specific nodes in real-time during

the execution of Prométhée. It also possible to communicated with certain I/O nodes in order to change the values of some neurons in the network on the fly. Figure A.4 shows an example of the information displayed by Pandora.

A.4 Fundamental neural algorithms

A.4.1 Classification

The *perceptron* proposed by Rosenblatt [1958] is one of the most famous models of artificial neural networks. In its most basic form, it consists in a single elementary computational unit that mimics the behavior of a biological neuron as suggested by [McCulloch and Pitts, 1943]. Such a neuron x basically calculates a weighted sum of its N inputs y_i:

$$x = H(\sum_{i=1}^{N} w_i y_i - b)$$ (A.1)

where w_i is the weight of the input y_i, b a constant bias, and $H(z)$ the Heaviside function equal to 1 if $z \geq 1$ and 0 otherwise. The latter is called "activation function" since it determines whether the neuron accumulated enough "potential" based on its inputs to respond. The bias serves as a threshold that the weighted sum of the inputs have to exceed in order to "activate" the output neuron. Like all the neurons I used in my work, this equation models an average firing rate as a simplification of spiking activities in biological neurons.

The *perceptron* is a linear classifier. As shown in Figure A.5, it can model the logical operators AND and OR but fails in the case of a XOR[1]. Indeed, when represented in a 2D space, the two former a linearly separable problems, but not the latter. For non-linearly separable problems, more than one layer neurons is required. This is called a *multilayer perceptron*. Each layer decomposes the problem linearly so that the next layer use intermediate output. It is also worth noting that a large variety of activation functions can be used to introduce other kinds of non-linearities in the neuron response – e.g. sigmoid or hyperbolic tangent functions.

The notion of learning in a classifier consists in drawing the separation line[2]. This is done by modifying the weights of the links between a neuron and its input; much like neuromodulatory mechanisms strengthen or weaken synaptic connections between biological neurons. A famous example is the Hebbian learning rule, named after Donald Hebb. The idea of the Hebbian theory is that "cells that fire together, wire together" [Hebb, 1949]. In artificial neurons, this can be translated in the following learning rule:

$$\frac{dw_{ij}}{dt} = \alpha.x_j.y_i$$ (A.2)

where $w_i j$ is the link between x_j and y_i and α the learning rate. The latter parameter determines the speed at which the connection weights are modified.

There exists a large variety of classifiers in the literature [MacQueen, 1967; Kohonen, 1982; Carpenter and Grossberg, 1988].

---

[1]Reminder: for two binary inputs A and B, i) A AND B $= 1$ only if both A and B equal 1, ii) A OR B $= 1$ if at least A or B is 1, and iii) A XOR B $= 1$ if and only if either A or B is 1, not both.

[2]The separation line becomes a hyperplane in multidimensional problems
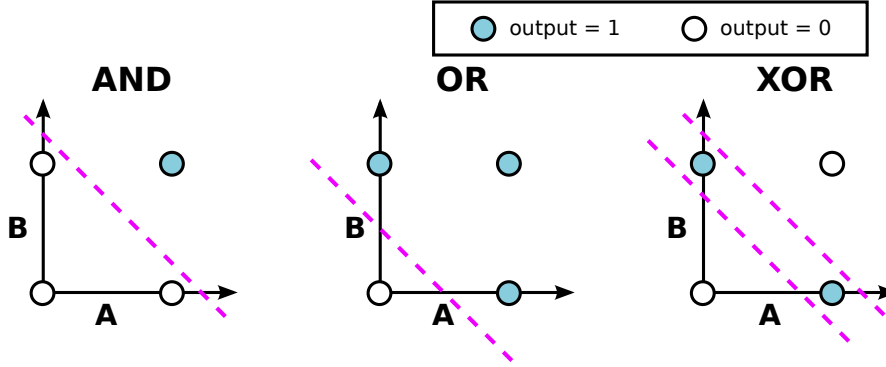
Figure A.5: Linear classification using a perceptron. In the case of an AND or OR operation between the inputs A and B, a single line can separate the two classes of responses exhibited by the perceptron, i.e. 1 or 0. Hence, they are linearly separable problems. In contrast, a single line is not sufficient for the XOR operation.

A.4.2 Conditioning

Conditioning is a form of associative learning that is observed in humans and animals. *Classical conditioning* refers to the notion made famous by Pavlov [1927]. In his experiment, Pavlov presented a auditory stimulus right before feeding dogs. After a few repetitions, the dogs began to salivate in response to the stimulus even without food. To generalize, the concept consists in pairing a conditional, neutral stimulus (e.g. sound) to an unconditional, biologically potent one (e.g. food). As a results, the unconditioned, reflex response is also associated to conditional stimulus as conditional response.

Least Mean Square (LMS) is an optimization methods that can computationally model this principle. It can be implemented in artificial neural networks by means of the Widrow and Hoff [1960] rule. Thereby, the connection weights are modified as follows:

$$\frac{dw_{ij}}{dt} = \alpha . y_i (\widehat{x}_j - x_j) \tag{A.3}$$

where $w_i j$ is the link between the output neuron x_j and the input neuron y_i, $\widehat{x}_j$ the desired output and α the learning rate. Here, y_i, x_j and $\widehat{x}_j$ respectively represent the neutral stimulus, the conditioned response and the unconditioned response. This gradient descent reduces the error between the actual and the desired output based on the input patterns. Notably, the difference between this model and biological Pavlovian learning is that in the latter both the conditional and unconditional stimuli trigger the response. Thus an additional layer summing x_j and $\widehat{x}_j$ is required to fully model classical conditioning.

Such a technique is used in PerAc model to build sensorimotor associations [Gaussier and Zrehen, 1995]. In the context of navigation, these associations can consist in place–action coupling so that the desired direction shown by the human teacher is associated with the current place.

A.4.3   Competition

Competition mechanisms consist in inhibitory lateral connections among a population of neurons. As a results, those having the highest activities – e.g. thanks to excitatory inputs – can reduce the response of their neighbors. This idea is illustrated in Figure A.6.
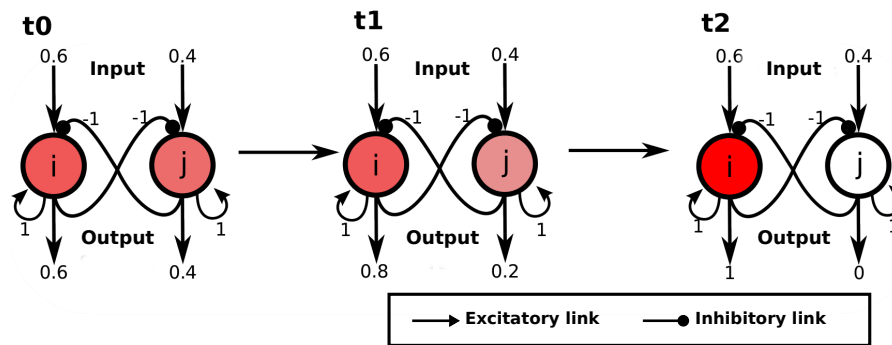
Figure A.6: Dynamics of competition between two neurons. Each of them have an excitatory input, an excitatory recurrent connection and an inhibitory connection with the other. Starting from different input values (0.4 and 0.6) at t_0, the most activated neuron wins the competition after a few steps. The figure is borrowed from Jauffret [2014]

This principle has been widely used in artificial neural networks to model a decision-making mechanism for instance [Rumelhart and Zipser, 1985; Carpenter and Grossberg, 1988]. Classically, a *winner-takes-all* competition is obtained in the case only one neuron remains active after a few iterations – like in Figure A.6. However, it is possible to make the competition less strict by changing the activation functions [Maass, 2000]

A.4.4   Competition/Cooperation in Dynamic Neural Fields

The concept of neural field is based on the fact that the brain uses *population coding* [Taube et al., 1990; Bullmore and Sporns, 2009]. Indeed, neighboring cells generally have similar activities. Rather than having the information encoded in a single cell – the so-called *grandmother cells* – it is distributed over a population of neurons. By this principle of redundancy, each neuron have a portion of the information. Thereby, population coding ensures a relative robustness to neurodegeneration for instance.

This notion of continuous neural space can be attributed to Beurle [1956]. Wilson and Cowan [1973] then proposed the concept of excitatory and inhibitory neurons to model interactions in populations of neurons. Today, one of the most commonly used model was proposed by Amari [1977]. It is based on the idea that excitatory interactions are local (proximal) while inhibitory ones are distal. This is modeled by an interaction kernel in the shape of a "mexican hat", i.e. consisting in a Laplacian of Gaussian or a Difference of Gaussian. Based on Amari's work, Schöner et al. [1995] proposed the theory of Dynamic Neural Fields (DNF) as a framework for robotic control.

The model of Amari [1977] have interesting properties in the sense of dynamical systems (bifurcation, fusion, hysteresis, memory). This allows dynamical filtering implementing distal competition and proximal cooperation. Let us consider a 1D neural field of size N. The potential
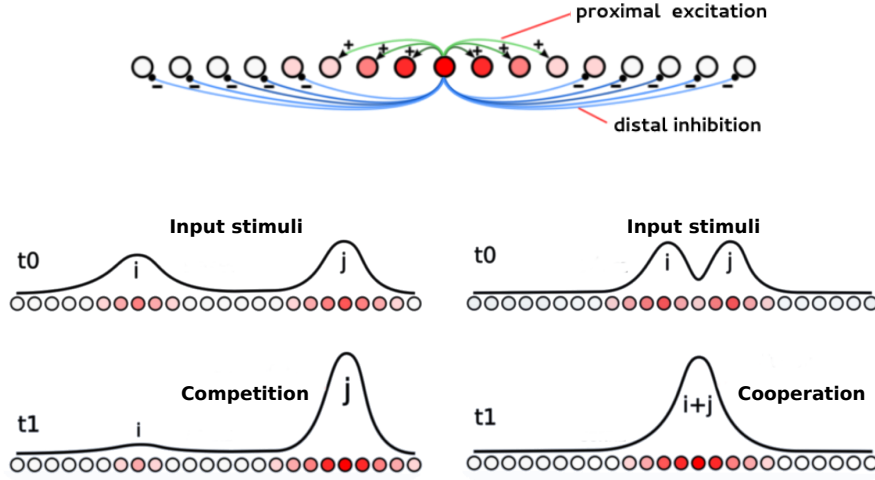
Figure A.7: Competition/cooperation in 1D Dynamic Neural Field. **TOP**: The DoG kernel that governs the interaction between neurons of the field. For clarity, only connections projecting from the central neuron are shown. **BOTTOM**: Examples of competition (LEFT) and cooperation (RIGHT) between input stimuli. In general, proximal stimuli reinforce each other and inhibit distant ones. The figure is slightly readapted from Jauffret [2014]

u of the j^{th} neuron of the field is updated as follows:

$$\tau . \frac{u_j(t)}{dt} = -u_j(t) + X_j(t) + c + \int_{z \in V_z} w(z).f(u_{j-z}(t)).dz \qquad \text{(A.4)}$$

where X is the input, $f(x) = tanh(x)$ is the activation function, τ is the time constant, c is a constant inhibition potential, w is an interaction kernel and V_z the interaction neighborhood of size N. Using a Difference of Gaussian (DoG) as an interaction kernel allows the dynamical filtering described above as illustrated in Figure A.7.
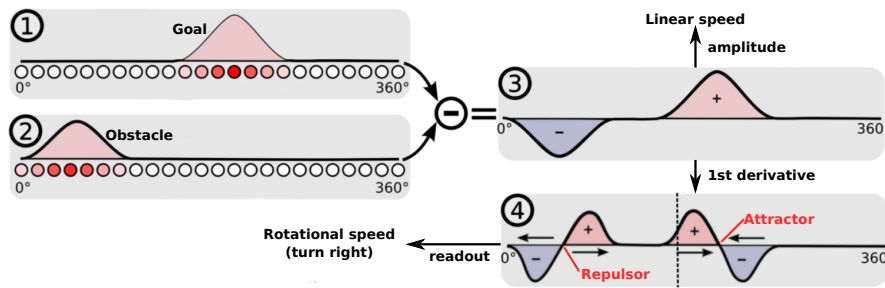
Figure A.8: DNF-nased mobile robot control. The neural populations represent the 360^o field surrounding the robot. Two types of inputs are represented in (1) and (2), respectively representing an approach stimulus (e.g. goal) and an avoidance stimulus (e.g. obstacle). From the resulting DNF in (3), two information are extracted: i) linear speed, depending on the amplitude, ii) rotational speed, based on a readout of the first derivative. Indeed, points where the derivative equals zero form attractors and repulsors that trigger a decision to turn left or right. In the example illustrated here, the closest attractor with respect to the current direction (dotted line) indicates that the robot has to turn right. The figure is readapted from Jauffret [2014]

Thanks to the above mentioned properties, DNF theory is a good candidate for robotic control Schöner et al. [1995]; Quoy et al. [2003]; Cuperlier et al. [2006]. The example of 1D field given above can be exploited in navigation. In such case, controlling the behavior of a mobile robot basically consists in giving motor orders corresponding to linear and rotational speed. If there are several sub-behaviors (suggested by different navigation strategies), they can be considered as weighted inputs of a DNF that determines the action to perform. This principle is illustrated in Figure A.8.