

# ENACTIVE ROBOT VISION

THÈSE N° 3974 (2007)

PRÉSENTÉE LE 5 DÉCEMBRE 2007

À LA FACULTÉ DES SCIENCES ET TECHNIQUES DE L'INGÉNIEUR  
LABORATOIRE DE SYSTÈMES INTELLIGENTS  
PROGRAMME DOCTORAL EN SYSTÈMES DE PRODUCTION ET ROBOTIQUE

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

**Mototaka SUZUKI**

M.Sc in mechanical engineering, Waseda University, Tokyo, Japon  
et de nationalité japonaise

acceptée sur proposition du jury:

Prof. M.-O. Hongler, président du jury  
Prof. D. Floreano, directeur de thèse  
Prof. E. Di Paolo, rapporteur  
Dr F. Kaplan, rapporteur  
Prof. T. Ziemke, rapporteur



ÉCOLE POLYTECHNIQUE  
FÉDÉRALE DE LAUSANNE

Suisse  
2007



---

# Abstract

---

The complexity of today's autonomous robots poses a major challenge for Artificial Intelligence. These robots are equipped with sophisticated sensors and mechanical abilities that allow them to enter our homes and interact with humans. For example, today's robots are almost all equipped with vision and several of them can move over rough terrain with wheels or legs. The methods developed so far in Artificial Intelligence, however, are not yet ready to cope with the complexity of the information gathered through the robot sensors and the need for rapid action in partially unknown and dynamic environments.

In this thesis, I will argue that the apparent complexity of the environment and of the robot brain can be significantly simplified if perception, behavior, and learning are allowed to co-develop on the same time scale. In doing so, robots become sensitive to, and actively exploit, characteristics of the environment that they can tackle within their own computational and physical constraints.

This line of work is grounded on philosophical and psychological research showing that perception is an active process mediated by behavior. However, computational models of active vision are very rare and often rely on architectures that are pre-programmed to detect certain characteristics of the environment.

Previous work have shown that complex visual tasks, such as position and size invariant shape recognition as well as navigation, can be tackled with remarkably simple neural architectures generated by a coevolutionary process of active vision and feature selection. Behavioral machines equipped with primitive vision systems and direct pathways between visual and motor neurons were evolved while freely interacting with their environments.

I proceed further on this line of investigation and describe the application of this methodology in three situations, namely car driving with an omnidirectional camera, goal-oriented navigation of a humanoid robot, and cooperative tasks by two agents. I will show that these systems develop sensitivity to a number of oriented, retinotopic, visual features – oriented edges, corners, height – and a behavioral repertoire to locate, bring, and keep these features in sensitive regions of the vision system that allow them to accomplish their goals.

In a second set of experiments, I will show that active vision can be exploited by the robot to perform anticipatory exploration of the environment in a task that requires landmark-based navigation. Evolved robots exploit an internal expectation system that makes use of active exploration to check for expected events in the environment.

I will describe a third set of experiments where, in addition to an evolutionary process, the visual system of the robot can develop receptive fields by means of unsupervised Hebbian learning and show that these receptive fields are significantly affected by the behavior of the system and differ from those predicted by most computational models of visual cortex.

Finally, I will show that these robots replicate the performance deficiencies observed in experiments of motor deprivation with kitten when they are exposed to the same type of motor deprivations. Furthermore, the analyses of our robot brains suggest an explanation for the deficiencies observed in kitten that have not yet been fully understood.

Key words: active vision, enaction, mobile robots, neural networks, computer vision



---

## Version abrégée

---

La complexité des robots autonomes modernes représente un défi majeur pour l'Intelligence Artificielle. Ces robots sont équipés de capteurs sophistiqués et de capacités mécaniques leur permettant de s'installer dans nos maisons et d'interagir avec des humains. Par exemple, les robots d'aujourd'hui sont presque tous équipés de caméras et certains d'entre eux peuvent se déplacer sur des terrains difficiles à l'aide de chenilles ou de jambes. Néanmoins, les méthodes développées jusqu'à présent en Intelligence Artificielle ne permettent pas encore de gérer la complexité des informations récoltées par les senseurs d'un robot et de réagir rapidement en présence d'un environnement partiellement inconnu et dynamique.

Dans cette thèse, j'argumenterai que la complexité apparente de l'environnement et du cerveau du robot peuvent être simplifiées de manière significative si la perception, le comportement et l'apprentissage se développent conjointement suivant la même échelle temporelle. De cette manière, les robots deviennent sensibles aux caractéristiques de l'environnement auxquelles ils peuvent se confronter avec leurs propres contraintes computationnelles et physiques et les exploitent activement.

Cette ligne de recherche enracinée dans la philosophie et la psychologie montre que la perception est un processus actif en complète interaction avec le comportement. Néanmoins, les modèles computationnels de vision active sont très rares et se basent souvent sur des architectures programmées dans le but de détecter certaines caractéristiques prédéterminées de l'environnement.

Des recherches précédentes ont montré que des tâches visuelles complexes, comme la reconnaissance de formes insensible aux variations de taille et de position ainsi que la navigation, peuvent être approchées par des architectures neurales remarquablement simples générées par un processus d'évolution conjoint de vision active et de sélection de caractéristiques. Des machines comportementales équipées de systèmes de visions primitifs et de connexions directes entre les neurones moteurs et visuels ont été évoluées en interagissant librement avec leurs environnements.

Je poursuivrai sur cette ligne de pensée et décrirai l'application de cette technologie à trois situations, plus concrètement à la conduite d'une voiture équipée d'une caméra omnidirectionnelle, à la navigation d'un robot humanoïde, et à des tâches

nécessitant la coopération de deux agents. Je montrerai que ces systèmes développent une sensibilité à un nombre de caractéristiques visuelles orientées – angles orientés, coins, hauteur – et un répertoire comportemental pour localiser, amener et maintenir ces caractéristiques dans les régions sensibles du système visuel, leur permettant d’accomplir leurs buts.

Dans une seconde série d’expériences je montrerai que la vision active peut être exploitée par le robot pour réaliser une exploration anticipée de l’environnement dans une tâche de navigation à l’aide de points de repères. Les robots évolués exploitent un système interne de prédiction faisant usage d’exploration active afin de contrôler des événements attendus dans l’environnement.

Je décrirai une troisième série d’expériences où, en plus du processus d’évolution, le système visuel du robot peut développer des champs récepteurs à l’aide d’un apprentissage Hebbien non supervisé et montrer que ces champs récepteurs sont significativement affectés par le comportement du système et différent de ceux prédits par la plupart des modèles computationnels du cortex visuel.

Finalement, je montrerai que ces robots reproduisent les imperfections en terme de performance observées lors des expériences de privation motrice réalisées sur des chatons lorsqu’ils sont exposés aux mêmes types de privations motrices. De plus, l’analyse du cerveau de nos robots suggère une explication pour les déficiences observées chez les chatons et qui n’ont pas encore été complètement comprises.

Mots clés: vision active, énaction, robots mobiles, réseaux de neurones, vision assistée par ordinateur

---

# Acknowledgments

---

It is one of the most fortunate events in my entire life that I could work with Prof. Dario Floreano for four years. It was in August 2002 when I met and talked to him for the first time. Together with Dr. Takashi Gomi, we talked over a dinner in a nice restaurant in Edinburgh. After some simple questions, he immediately accepted me as a member of his lab from the next year and let me work on a fascinating research project on active vision, that indeed became the topic of this PhD thesis. One can never imagine how many things he has taught me on science and life over the past four years. I really cannot thank him enough for his advices, generosity and trust in myself. I have learned really a lot from him during this period. His secretary Anouk Hein and his ex-secretary Dominique Etienne have also helped me in my research.

I would also like to thank Dr. Takashi Gomi. He not only introduced me to Prof. Floreano, but also supported me since we met in Tokyo in the year 2001. I am especially grateful to him for his financial support during the first year of my stay in Switzerland. I could not survive and continue my research without him. The rest of my life in Lausanne and this thesis were made possible by the generous support of EPFL.

I owe a big thank you to Professors Tom Ziemke, Ezequiel A. Di Paolo, Frédéric Kaplan, who accepted to be on the jury of this thesis and to Prof. Max-Olivier Hongler for being the president of the jury. They also provided invaluable criticisms and suggestions on this thesis.

I would like to thank all my colleagues at the LIS and ex-colleagues at the ASL2, especially to Antoine Beyeler, Jesper Blynel, Michael Bonani, Michele Bongiovanni, Peter Dürr, Diego Federici, Yannick Fournier, Simon Harding, Sabine Hauert, Julien Hubert, Walter Karlen, Adam Klaptocz, Mirko Kovac, Severin Leven, Stéphane Magnenat, Daniel Marbach, Claudio Mattiussi, Sara Mitri, Andres Perez-Uribe, Fanny Riedo, James Roberts, Daniel Roggen, Timothy Stirling, Danesh Tarapore, and Jean-Christophe Zufferey for their friendship and help. A special thank you to Danesh Tarapore, Julien Hubert, Sabine Hauert and Fanny Riedo for constructive comments on this thesis and translating the abstract to french. Thanks also go to Guido de Croon, Marc Clapera, Jie (Roger) Luo, Chaitanya Ekanadham, Tommaso

Gritti, Ha Huy Hoang Le, Jacob van der Blij, Sylvain Quartier, and Antoine Béguin, who spent a lot of time and tackled challenging problems together with me. I have learnt a lot from the collaborations with them.

I also wish to thank my parents and family. They always worried more about my health in Switzerland rather than my PhD.

And most importantly, my darling girlfriend Kana encouraged me to do this thesis and supported me throughout. The final stage of my PhD would have been extremely difficult without her.

---

# Contents

---

<b>Abstract</b>	<b>i</b>
<b>Version abrégée</b>	<b>iii</b>
<b>Acknowledgments</b>	<b>v</b>
<b>Contents</b>	<b>vii</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Background and Motivation . . . . .	1
1.1.1 Evolution of an Eye: What is an eye for? . . . . .	1
1.1.2 Active Vision . . . . .	2
1.1.3 Active Perception in Animals . . . . .	3
1.1.4 Physiological and Psychological Evidence for Active Vision . . . . .	4
1.1.5 Active Vision Systems . . . . .	6
1.2 Aims and Originality . . . . .	8
1.3 Outline of the Thesis . . . . .	9
1.4 Published Work . . . . .	10
<b>2 Coevolution of Active Vision and Feature Selection</b>	<b>11</b>
2.1 General Framework . . . . .	11
2.2 Previous Work . . . . .	12
2.2.1 Shape Discrimination Task . . . . .	12
2.2.2 Simulated Car Driving . . . . .	13
2.2.3 Indoor Navigation . . . . .	14
2.3 Omnidirectional Active Vision . . . . .	15
2.4 Goal-Oriented Humanoid Robot Walking . . . . .	15
2.5 Active Perception in Cooperative Tasks . . . . .	17
2.6 Summary . . . . .	19

<b>3</b>	<b>Capturing Spatio-Temporal Relationships</b>	<b>21</b>
3.1	Introduction . . . . .	21
3.2	Method . . . . .	22
3.2.1	Experiment and Task . . . . .	23
3.2.2	Neural Architecture and Genetic Algorithm . . . . .	24
3.3	Results and Analysis . . . . .	26
3.4	Discussion . . . . .	29
3.5	Conclusion . . . . .	31
<b>4</b>	<b>Landmark-Based Navigation</b>	<b>33</b>
4.1	Introduction . . . . .	33
4.2	Method . . . . .	34
4.3	Results . . . . .	35
4.4	Discussion . . . . .	40
<b>5</b>	<b>Evolution and Ontogenetic Adaptation</b>	<b>43</b>
5.1	Introduction . . . . .	43
5.1.1	The Role of Active Behavior . . . . .	44
5.2	Method . . . . .	46
5.3	Experiments . . . . .	49
5.4	Behavioral analysis . . . . .	52
5.5	Receptive Field Development . . . . .	53
5.6	Discussion . . . . .	56
5.7	Conclusion . . . . .	58
<b>6</b>	<b>Computational Neuroethology</b>	<b>61</b>
6.1	Introduction . . . . .	61
6.2	Experiments . . . . .	64
6.2.1	Visual development during active or passive movements . . . . .	64
6.3	Analysis . . . . .	65
6.3.1	Lesion Studies . . . . .	65
6.4	Discussion . . . . .	70
6.5	Conclusion . . . . .	71
<b>7</b>	<b>Discussion and Outlook</b>	<b>73</b>
7.1	Discussion . . . . .	73
7.2	Future Directions . . . . .	75
<b>A</b>	<b>Omnidirectional Active Vision for Evolutionary Car Driving</b>	<b>77</b>
A.1	Introduction . . . . .	77
A.2	Experimental Setup . . . . .	78
A.3	Evolution of Neural Controllers . . . . .	81
A.4	Behavioral Analysis . . . . .	82
A.5	Discussion . . . . .	83
A.6	Conclusions . . . . .	85

<b>B Active Perception in Cooperative Tasks</b>	<b>87</b>
B.1 Introduction . . . . .	87
B.2 Method . . . . .	88
B.3 Shape localization task . . . . .	89
B.4 Shape discrimination task . . . . .	90
B.5 Conclusion . . . . .	92





---

# 1

---

## Introduction

### 1.1 Background and Motivation

Despite enormous progress in the past few decades, machine vision is still far from achieving the goal set by biological vision with respect to its speed and reliability. Many reasons can be considered, but we may still be missing something fundamental in the way that biological vision is organized. The origin of an eye therefore is still worth being reconsidered and may give us an insight and inspiration on the design of artificial vision systems.

#### 1.1.1 Evolution of an Eye: What is an eye for?

Impressive diversity of animal eye types is indeed the masterpiece of evolutionary process. Debates still go on about the early origins of eyes before the great Cambrian radiation event that gave us most of the eye types we see in animals today. Eyes can be defined as organs for *spatial vision* with the ability to compare intensity in different directions regardless of their precision (Land and Nilsson, 2002). In principle, the simplest way to produce spatial vision would be to have two light sensitive cells (photoreceptor cells) shielded so that they do not pick up light from exactly the same direction<sup>1</sup>.

*The Light switch theory* (Parker, 2003) suggests that the evolutionary invention of eyes fueled the Cambrian explosion. It is speculated that a few species of late Precambrian animals became large enough to acquire good spatial vision and improved mobility, and became the first visually guided predators. Such an ecological invention would have put a tremendous selection pressure on a large section of faunas, forcing other species to evolve protective methods such as body armor, shells, avoiding exposure by deep burrowing or developing good vision and mobility themselves. Thus the introduction of visually guided predation may have altered much of the ecological system.

---

<sup>1</sup>Interestingly, Gibson (1979) emphasized the importance of registering differences of intensities in different directions (i.e., the ambient array) to understand visual perception.

This scenario tells us two important aspects of biological vision. First, vision has been the most crucial sense for animals to survive in their environment. Thus, the invention of vision may have largely affected the evolution of animal ecology. More importantly, eyes have always been *mobile* since their origin; they have been forced to move as their owners move. Thus visually guided animals have always needed to process ever-changing visual inputs, rather than static ‘photographs’.

Another important notion in the design of animal eyes is *visual ecology* – the ways that eyes and visual processing are specifically adapted to the lifestyles of the animals that bear them. As humans we tend to think of vision as a general-purpose sense, supplying us with any kind of information we need. For most other animals, this is not so (Land and Nilsson, 2002, for a review). Each eye type and its visual processing have a specific purpose (e.g., mating, predation) in the lifestyle of each species.

Seeing vision and visual perception from these perspectives is far from the widely spread metaphor that eyes are static devices which register the scenes in front of them rather like surveillance cameras, recording the positions and motions of objects within a fixed field of view. Rather, the origin of vision seems to strongly support the philosophy of active vision, introduced in the next section.

### 1.1.2 Active Vision

Unfortunately seeing has often been –either explicitly or implicitly– assumed to be the passive reception of visual information. Since at least the time of Johannes Kepler some 400 years ago, vision research has been based on the image-forming eye. Methodologies in computer –and even biological– vision have been largely affected by the dogma of David Marr, who stated that “vision is the process of discovering from images what is present in the world and where it is” (Marr, 1982). Thus in contrast to the nature of biological vision, machine vision has long aimed at reconstructing the three dimensional visual world from two dimensional images (from pixels to predicates).

Active vision is a marked departure from the orthodoxy that has held sway for the last two decades in machine vision. Active vision emphasizes the role of vision as a sense for robots and other real-time perception-action systems (Bajcsy, 1988; Aloimonos et al., 1987; Aloimonos, 1990; Ballard, 1991; Findlay and Gilchrist, 2003). It picks out the properties of images which are necessary to perform the assigned tasks at hand and ignores the rest. In this context, there is no clear need for the sort of detailed reconstructions of the visible world. Active vision is the natural result of considering vision in the context of real-time perception-action systems and their ever-changing environment.

In a broad sense, active vision includes covert and overt forms of attention, although neural substrates which generate covert and overt attention are largely overlapped (see Findlay and Gilchrist, 2003, for a review). In this thesis we do not clearly distinguish these two processes in our artificial systems.

### 1.1.3 Active Perception in Animals

Active perception is ubiquitous in the animal kingdom. Examples include echolocation in bats and dolphins (Au, 1993; Thomas et al., 2004), active electrolocation in weakly electric fish (Bullock and Heiligenberg, 1986; Möller, 1995; Babineau et al., 2007), active touch in the rodent whisker system (Brecht et al., 1997; Hartmann, 2001), insect antennal systems (Horseman et al., 1997; Dürr et al., 2001), hydrodynamic imaging in blind cave fish (von Campenhausen et al., 1981; Hassan, 1989; Montgomery et al., 2001), active tracking with foveal vision in hoverflies (Collett and Land, 1975), edge detection with horizontally-moving eyes in jumping spiders (Land, 1971) and active depth-estimation in honey bees (Lehrer, 1997), locusts and praying mantises (Sobel, 1990). Nearly all animals with good vision, such as mammals, can move their eyes, and in most cases a mobile gaze is an essential feature of visual data acquisition. See (Land and Nilsson, 2002; Srinivasan and Venkatesh, 1997) for a review.

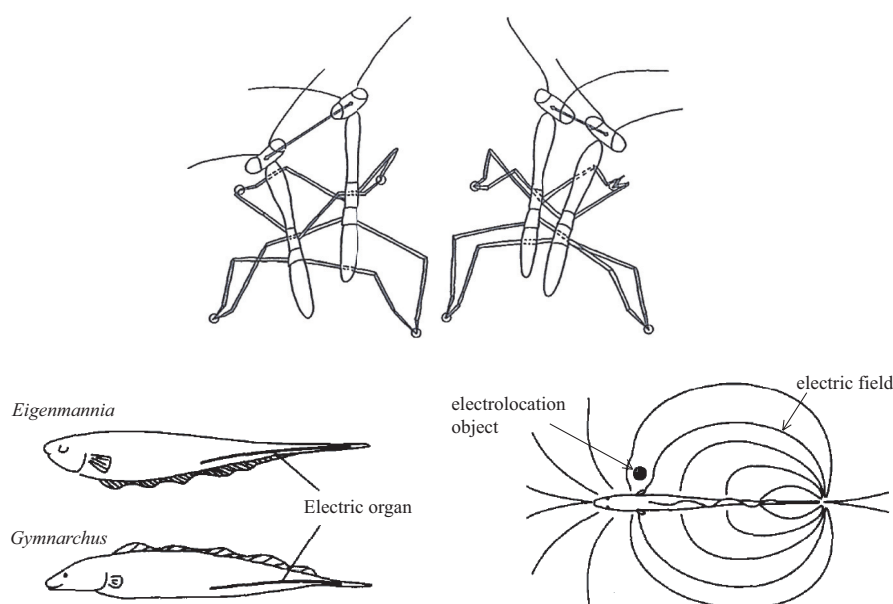


Figure 1.1: Top: An early instar praying mantis shows side-to-side scanning movements for depth estimation. Reproduced from (Land and Nilsson, 2002). Bottom: Weakly electric fish's poor spatial acuity is enhanced by the back-and-forth swimming. Adapted from Kawasaki (1997).

This ubiquity of active perception suggests that the ability to control sensory systems has been highly beneficial for these animals to survive in their environment. If active movements of sensory systems do cause a problem then one might have expected these animals to have evolved with the ability to lock their sensory systems. A possible reason of this ubiquity will be further discussed in Chapter 7.

### 1.1.4 Physiological and Psychological Evidence for Active Vision

Active vision has been intensively studied in human eye movements. Interestingly human eyes and those of other vertebrates are constructed in a surprisingly bad manner, contrary to our popular belief. Before reaching the photosensitive rods and cones, the light must first traverse not only a dense tangle of neural matter formed by the axons and layers of neurons that serve the first stages of visual computation, but also a vast web of blood vessels that irrigate the retina (Gibson, 1966). Both of these obscure the photosensitive layer and would be expected to impede vision. Other defects include the surprisingly large blind spot in each retina, the nonuniformity of cones even within the fovea, markedly worse color discrimination in peripheral vision (Noorlander et al., 1983; Mullen, 1991) and perturbations of the retinal image caused by saccadic eye movements (Findlay and Gilchrist, 2003).

Despite these significant defects, we seem to perceive a rich and stable world, which is extremely mysterious. However the idea of this detailed mental copy of our visual environment is increasingly recognized as an *illusion* (O'Regan, 1992; Blackmore et al., 1995; Noë, 2002, 2004). The radical skepticism known as the 'grand illusion' hypothesis, initiated by Dennett, claims that we are misled by our consciousness into thinking of visual experience as continuous and rich, whereas it is in fact discontinuous and sketchy. This view is apparently supported by striking phenomena such as *change blindness* and *inattention blindness* (Simons and Levin, 1997; Intraub, 1997). The phenomenon of change blindness happens when a change

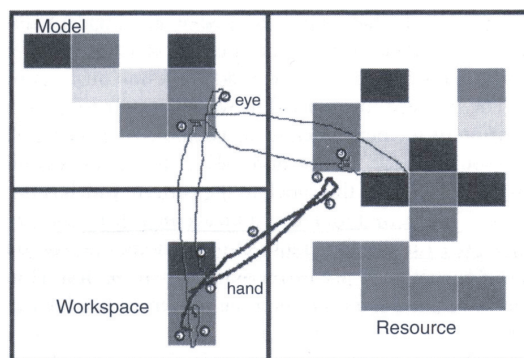


Figure 1.2: The 'blocks' task studied by Ballard et al. (1992). The subject operates on a computer display of colored blocks (denoted here by gray levels) and has the task of assembling a copy of the Model (top left) in the Workspace (bottom left). To do this it is necessary to operate with the mouse, obtaining blocks from the Resource (right) using a click and drag procedure. The traces show an episode of visual and motor activity. The heavy trace shows the mouse movement while the light trace shows the corresponding eye scan record. The subjects did not notice changes of the color of blocks if they were on their way to look at the Model before picking up a block, even though they had already looked at the Model many times.

in a central aspect of a scene remains undetected owing to attentional limitations (Rensink, 2002). A similar phenomenon is observed even when the changed feature is central to the agent’s activity (Ballard et al., 1992; Hayhoe, 2003) (Fig. 1.2).

These studies seem to reject –or at least diminish– the role of internal representation of the visual world<sup>2</sup>. The feeling that we seem to perceive a rich and stable world, instead arises from our ability to relocate our focus of attention to any part of the visual scene, as and when we wish, to ‘re-discover’ all the visual details in their entirety (Merleau-Ponty, 1968; Brooks, 1991; O’Regan, 1992). In other words, it is active vision that makes us feel as if we perceive the visual world with its complete richness and stability through such badly constructed eyes.

Active vision is significantly affected by cognitive context. In a classical psychological study Yarbus (1967) monitored eye movements during the presentation of a picture, e.g., Repin’s *An Unexpected Visitor* (Fig. 1.3). The observed gaze patterns were modulated by the question posed to the viewer in advance, suggesting that the informativeness of scene regions is task-dependent. This work clearly illustrates that eye movements are purposeful and significantly affected by the cognitive task at hand. In the following chapters we show how the embodiment<sup>3</sup>, the environment

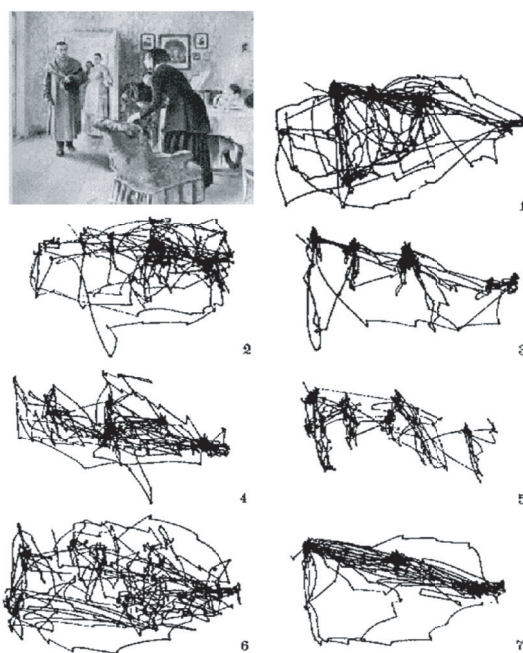


Figure 1.3: Seven records of eye movements by the same subject. 1) Free examination of the picture. 2–7) Gaze patterns were modulated by the question posed to the viewer in advance. Reproduced from Yarbus (1967).

<sup>2</sup>Debates still continue on this issue. Readers further interested in this topic may refer to (e.g., Simons and Rensink, 2005; Noë, 2005).

<sup>3</sup>In this thesis ‘embodiment’ means that systems are ‘structurally coupled’ to their environment (Ziemke, 2001).

and the task demand alter the vision control of behavioral systems.

The crucial role of active vision in perceptual adaptation has been shown by Bach-y-Rita and colleagues, who have developed the tactile-vision sensory substitution (TVSS) (see Bach-y-Rita and Kercel, 2003, for a review). See Fig. 1.4 for an old version of the system. Visual stimulation received by a head-mounted camera is transduced to activate an array of vibrators on the back, thigh, or tongue of a blind subject. If the subject is free to move around and thus control tactile-motor dependencies, after a time she reports that she has the experience of objects arrayed in three-dimensional space. This study emphasizes the strong dependence of visual adaptation on the behavior and sensory-motor coordination of the subject. We show that the results from our robotic experiments described in Chapter 5 and 6 are consistent with this finding; we show how active behavior of a robot affects the development of visual receptive fields.




Figure 1.4: Tactile-vision sensory substitution (TVSS) shown in Bach-y-Rita et al. (1969). On the right, the 400 point two-dimensional tactile stimulator matrix array is shown mounted in the back of a dental chair for projecting mechanical television images on to the skin of the back of blind subjects. In the position shown, the camera permits subjects to examine hand held objects from a visual angle approximating that of the eyes. When placed in front of the subject the camera could be manipulated to examine various parts of an object. Reproduced from Bach-y-Rita et al. (1969).

### 1.1.5 Active Vision Systems

Active vision systems in the literature generally have mechanisms that can actively control camera parameters such as position, orientation, focus, zoom, aperture, and vergence (in a two-camera system) in response to the requirements of the task and external stimuli (Swain and Stricker, 1993). They may also have features such as



spatially variant (foveal) sensors (Sandini and Tagliasco, 1980). A review of active vision systems can be found in (Blake and Yuille, 1992; Christensen et al., 1993; Vieville, 1997).




Figure 1.5: Left: Sandini's spatially variant (foveal) sensors (Sandini and Tagliasco, 1980; Metta et al., 2004). Right: The head of a humanoid robot, Dynamic Brain. Each eye consists of two cameras, a 100 degrees field of view camera for peripheral vision and a 24 degrees field of view camera for foveal vision. Reproduced from Shibata and Schaal (2001).

Not only camera parameters, but visual processing algorithms could also be tied closely with the activities systems support (e.g., navigation, manipulation), allowing simplified control algorithms and scene representation. This is indeed the central topic of this thesis.

Notable examples of active vision systems include the following four systems. First, Dickmanns and Graefe (1988) have developed the vehicle navigation system with active vision. The system demonstrated road vehicle guidance at high speed including obstacle detection and monocular relative spatial state estimation. However, the search windows were fixed a priori based on a large amount of task-specific knowledge. Second, Franceschini (1992) described the two-dimensional (floor-bound) fly robot (Fig. 1.6, left), encircled by a planar compound eye with 24 elements. Direct couplings between the eye's elements and wheel motors lead to low-level behavior that is well adapted to obstacle avoidance and goal seeking. However, the vision-motor couplings were hard-wired in this robot. Third, Nolfi (1998) and Scheier et al. (1998) described evolved robots that exploit active perception to perform tasks that require perceptual discrimination (the robots are equipped with proximity sensors that indicate the distance to objects, not with vision). Lastly, Harvey et al. (1994) described evolution of sensory and neural morphology for a gantry robot asked to reach for a triangular shape while avoiding a rectangular shape painted on a wall. Evolved robot solves the problem exploiting only two visual neurons whose receptive fields are aligned with a lateral edge of the triangle. The sequential activation of

these neurons, caused by the sweeping of the image over the retina while the robot rotates, is sufficient to trigger the correct approaching, or avoidance, behavior.




Figure 1.6: Left: Franceschini’s fly robot capable of obstacle avoidance and goal seeking (Franceschini, 1992). Top right: A Khepera robot performing perceptual discrimination task (Nolfi, 1998). Bottom right: A gantry robot performing shape discrimination task (Harvey et al., 1994).

Computational models have also been developed to account for *feature selection* – the neural development of sensitivity to relevant features in the visual scene to which the system selectively responds (e.g., Linsker, 1988; Hancock et al., 1992; Olshausen and Field, 1996). However, these disembodied models were not allowed to freely interact with the environment and to choose their own sensory inputs. Moreover, the combination of these two processes, namely active vision and feature selection, is still largely unexplored.

## 1.2 Aims and Originality

In this thesis we proceed further on this line of research and describe a series of experiments on the interplay between active vision and neural development in behavioral systems. More precisely, this thesis addresses four specific questions: 1) how active control of visual systems adapts to different embodiments and simplifies apparently complex visual tasks (Chapter 2); 2) how robots capable of active vision capture spatio-temporal relationships between successive events (Chapter 3 and 4); 3) how active vision and receptive field development interact with each other (Chapter 5); 4) how the robotic experiments provide a new insight on the hypothesis raised in a psychological study carried out in the 60’s that the proprioceptive motor information resulting from self-generated actions was necessary for the development of



normal visually-guided behavior (Chapter 6).

To answer these questions, we consistently utilize behavioral systems or robots embedded in their environment. The crucial role of embodiment in understanding intelligence has been discovered and further understood by many researchers (e.g., Brooks, 1991; Pfeifer and Scheier, 1999; Sharkey and Ziemke, 1998; Nolfi and Floreano, 2000; Ziemke, 2001, 2004). This methodology is fully consistent with the *enactive approach to cognition* (Varela et al., 1991; O'Regan and Noë, 2001)<sup>4</sup>; enaction is a history of structural coupling that brings forth a world (Varela et al., 1991). Enactivism claims that sensory-motor activity and embodiment are crucial in perceiving the environment and that machine vision could be a much simpler business if considered in this context. However, the concept of enaction is still largely unexplored in robotic systems. As a concrete example, we show in Chapter 6 how evolved robots bring forth a world by developing sensitivity to a particular set of visual features and actively controlling their body and vision. In Chapter 7 we will also discuss how the results described in this thesis are different from those of disembodied models.

Another important characteristic of our methodology is that we intend to study deliberately simple neural controllers for each behavioral system and its cognitive task. There are two reasons behind this. First, we hypothesize that active control of visual system simplifies apparently complex visual tasks. Simple nervous systems can generate complex visual performances as a consequence of the complex interaction between body, brain and environment, and can efficiently solve complex visual tasks by exploiting the coherence, consistency, regularity or structure arising in the history of this interaction. The second reason is that simple control systems allow us to better understand the activity of the neural system of the behaving robot, which is impossible in current psychological and neurophysiological experiments with living brains. Readers may scale-up and design more complex neural systems according to their need at hand. It is hoped that the insight and knowledge obtained with simpler neural systems in this thesis will be useful for other applications as well. The following chapters will provide evidence that supports and enriches our argument.

### 1.3 Outline of the Thesis

Chapter 2 first illustrates the general framework which is consistently used throughout this thesis, namely coevolution of active vision and receptive fields. Then this chapter briefly reviews the robotic implementations of this concept, consisting of the previous work by my colleagues and three sets of new experiments by myself.

Chapter 3 addresses the question of how robots capture spatio-temporal relationships between successive events. To do so, we extend the neural architecture described in Chapter 2 and explore the development of internal expectation mechanism coupled with active vision. We show that active vision can be exploited by the

---

<sup>4</sup>In a less general context than Varela's theory of cognition, the term "enaction" has been used by Bruner (1968) to define a particular kind of representation.

robot to perform anticipatory exploration of the environment in a task that requires landmark-based navigation.

Promising results described in Chapter 3 suggest that active vision approach could be further applied to more realistic landmark-based navigation tasks. Chapter 4 aims at providing a computational model accounting for the insect visual navigation that has been found recently.

In Chapter 5, we introduce the ability of *ontogenetic adaptation* in the neural architecture and study how active vision and receptive field development interact with each other. The interaction between evolution and learning is another central issue to be discussed in this chapter.

Using the same architecture endowed with neural plasticity, Chapter 6 addresses the question raised in a psychological study carried out in the 60's; is the proprioceptive motor information resulting from self-generated actions necessary for the development of normal visually-guided behavior? This computational study provides us with a new vantage point of a behavioral system that interacts with its environment.

Chapter 7 discusses the results, the limitation and the scalability of this approach, the contribution of this thesis, and draws conclusions.

## 1.4 Published Work

Some parts of this thesis have already been published in journal articles and in conference proceedings. Section 2.3 is based on (Suzuki et al., 2006). Chapter 3, 5 and 6 are based on (Suzuki and Floreano, 2006), (Floreano et al., 2005) and (Suzuki et al., 2005a,b), respectively.

Section 2.4, 2.5 and Chapter 4 are also in preparation for submission. An abbreviated version of this thesis has been accepted and will appear with the same title in Adaptive Behavior Journal in early 2008.

---

# 2

## Coevolution of Active Vision and Feature Selection

---

*This chapter illustrates the general methodology that is consistently used in this thesis. The methodology is successfully applied to six sets of robotic experiments, which demonstrates the power and the scalability of this approach. Apparently-complex behaviors generated with deliberately simple neural controllers are a consequence of the history of robot-environment interactions. Section 2.3 is based on Suzuki et al. (2006) while Section 2.4 and 2.5 are in preparation for submission.*

**Abstract.** Enactivism claims that sensory-motor activity and embodiment are crucial in perceiving the environment and that machine vision can be a much simpler business if considered in this context. However, computational models of enactive vision are very rare and often rely on handcrafted control systems. In this chapter, we briefly describe results from experiments where evolved robots can choose whether to exploit sensory motor coordination in a set of vision-based tasks. These results suggest that complex visual tasks can be tackled with remarkably simple neural architectures generated by a co-evolutionary process of active vision and feature selection. We describe the application of this methodology in six sets of experiments.

### 2.1 General Framework

Co-evolution of active vision and feature selection could greatly simplify the computational complexity of vision-based behavior by facilitating each other's task. This hypothesis was investigated in a series of experiments on co-evolution of active vision and feature selection for behavioral systems equipped with a primitive moving retina and a deliberately simple neural architecture (Fig. 2.1). The neural architecture was composed of an artificial retina and of two sets of output units. One set of output units determined the movement and type of zooming factor of the retina, and the other set of units determined the behavior of the system, such as the response of a pattern recognition system, the control parameters of a robot, or the actions of a car driver. The neural network was embedded in a behavioral

system and its input/output values were updated every 300 ms while its fitness was computed. Therefore, the synaptic weights of this network were responsible for the visual features on which the system based its behavior, and for the motor actions necessary to search for the features.




Figure 2.1: The neural architecture of the active vision system is composed of A) a grid of visual neurons with non-overlapping receptive fields whose activation is given by B) the grey level of the corresponding pixels in the image; C) a set of proprioceptive neurons that provide information about the movement of the vision system; D) a set of output neurons that determine the behavior of the system (pattern recognition, car driving, robot navigation); E) a set of output neurons that determine the behavior of the vision system; F) a set of evolvable synaptic connections. The number of neurons in each sub-system can vary according to the experimental settings.

## 2.2 Previous Work

### 2.2.1 Shape Discrimination Task

In a first set of experiments (Kato and Floreano, 2001), the neural network was embedded in a simulated pan-tilt camera and asked to discriminate between triangles and squares of different size that could appear at any location of a screen (Fig. 2.2), a perceptual task similar to that explored with the gantry robot described in (Harvey et al., 1994). The visual system was free to explore the image for 60 seconds while continuously telling whether the current screen showed a triangle or a square. The fitness was proportional to the amount of correct responses accumulated over the 60 seconds for several screenshots containing various instances of the two shapes. Evolved systems were capable of correctly identifying the type of shape with 100%

accuracy after a few seconds despite the fact that this recognition problem is not linearly separable (see Floreano et al., 2004, for more detail) and that the neural network does not have hidden units, which in theory are necessary to solve non-linearly separable tasks. Indeed, a similar type of neural network with a number of hidden neurons presented with the same set of images and trained with supervised learning, but without the possibility to actively explore the scene, was not capable of solving the task. The evolved active vision system developed sensitivity to vertical edges, oriented edges and corners, and used its movement to search for these features in order to tell whether the shape was a triangle or a square. These features, which are found also in the early visual system of almost all animals, are invariant to size and location.




Figure 2.2: An evolved individual explores the screen searching for the shape and recognizes it by the presence of a vertical edge. The retina moves with respect to its top leftmost corner, here marked by a *dot*. The *dots* drawn after every retina movement are connected with a *line*. For graphical clarity, the values of the cells are not shown, only the retinal perimeter.

## 2.2.2 Simulated Car Driving

In a second set of experiments (Floreano et al., 2004), the neural network was embedded in a simulated car and was asked to drive over several mountain circuits (Fig. 2.3). The simulator was a modified version of a car race video game. The neural network could move the retina across the scene seen through the windscreen at the driver's seat and control the steering, acceleration, and braking of the car. The fitness was proportional to the distance covered across all circuits where the car was tested. Evolved networks completed all circuits with time laps competitive to those of well-practiced students controlling the car with a joystick. Evolved network started by searching for the edge of the road and tracked its relative position with respect to the edge of the windscreen in order to control steering and acceleration. This behavior was supported by the development of sensitivity to oriented-edges.




Figure 2.3: Search for the edge of the road at the beginning of a drive over a mountain road.

### 2.2.3 Indoor Navigation

In a third set of experiments (Marocco and Floreano, 2002), the neural network was embedded in a real mobile robot with a pan-tilt camera that was asked to navigate in a square arena with low walls located in an office (Fig. 2.4). The fitness was proportional to the amount of straight motion measured over two minutes. Robots that hit the walls because they watched people or other irrelevant features of the office, had lower fitness than robots that could perform long straight paths and avoid walls of the arena. Evolved robots tended to fixate the edge between the floor and the walls of the arena, and turned away from the wall when the size of its retinal projection became larger than a threshold. This combination of sensitivity to oriented edges and looming is found also in visual circuits of several insects and birds.



Figure 2.4: A mobile robot with pan-tilt camera is asked to move within the walled arena in the office environment.

## 2.3 Omnidirectional Active Vision

We investigated whether the same principle can be applied to a different optic device, omnidirectional camera that provides a 360 degrees field of view. In the fourth set of experiments (Suzuki et al., 2006), a 1/10 scale model car equipped with an omnidirectional camera (Fig. 2.5) was modeled in simulation and asked to drive over similar circuits to those used in the experiment described in Section 2.2.2.

Evolved systems were capable of quickly detecting the road area in front of the car so as to drive the car on the road. It took time for the retina to find the task-relevant feature in such a broad field of view, which explained the slightly lower fitness values of the best evolved individuals with the omnidirectional camera than those of the best evolved individuals with a conventional pan/tilt camera (Fig. 2.6). The evolved systems solved this problem by moving the car back and forth until the retina found the appropriate feature. Then the car started driving forward at full speed. These experiments have shown that the active vision architecture can adapt to the new optic device and accordingly develop a suitable visuo-motor coordination to solve its task. For more details of this experiment, see Appendix A.




Figure 2.5: A 1/10 scale model car equipped with omnidirectional camera.

## 2.4 Goal-Oriented Humanoid Robot Walking

Little attention has been paid to how the vision system is coordinated with the body dynamics of walking bipedal robots. Most researchers have tried to ‘compensate’ the perturbations caused by walking or neglected such perturbations and assumed that the vision system is stable. Although animals possess compensatory mechanism of eye movements, the body movements while walking still disturb the retinal image. One might be surprised to know the fact that saccadic eye movements create calamitous smearing and displacement of the retinal image; fixation accuracy during walking is far from perfect with retinal image slip attaining four degrees per second (Steinman and Collewijn, 1980).






Figure 2.6: Evolution of neural controllers of the car with the pan-tilt camera (left) and the omnidirectional camera (right) in an eight shaped circuit. Thick line shows the fitness values of the best individuals and thin line shows those of the population average. Each data point is the average of three evolutionary runs with different initializations of the population. Vertical lines show the standard error.

In a fifth set of experiments (Suzuki et al., 2008a), we explored how visual system can be coordinated with the body dynamics of a walking bipedal robot. From an engineering point of view, this ability would become really crucial if robots need, for example, to robustly detect and keep tracking an important visual target while their body is largely moving.

A humanoid robot HOAP-2 (Fujitsu Automation Co., Ltd.) was modeled in a physics-based simulation (Fig. 2.7, top left) and the neural network controlling the head camera movement and the walking behavior was developed by using artificial evolution. The robot was asked to reach a goal location by detecting the beacon (white window) while avoiding obstacles and walls (see Fig. 2.7, top right).

Our interest in this experiment did not reside in the bipedal walking itself, but in the visuo-motor coordination against the large movements of the body and head while walking. Therefore the neural architecture was designed to control the macroscopic behavior of the robot, such as walking straight, turning left/right, as well as the pan/tilt camera movement. For this purpose, the stable walking was realized by loading and following the joint trajectories which had been pre-calculated by Zero Moment Point walking algorithm (e.g., Takanishi et al., 1990). Evolution of stable walking has been studied in (Reil and Husbands, 2002) with a significantly simplified model of a bipedal robot.

Evolved robots accomplished the task by means of a close coordination between the retina, the head and the body dynamics (Fig. 2.7, bottom)<sup>1</sup>. Behavioral analysis of evolved robots revealed that under the experimental conditions studied, their

<sup>1</sup>A video clip showing the behavior of the best evolved bipedal robot is available at: <http://lis.epfl.ch/research/projects/ActiveVision/videos/>



‘eye’ movements did not compensate the body dynamics so as to stabilize the goal beacon within the retina, unlike human eye movements, but rather reminiscent of very weak compensatory eye movements of cats (Moeller et al., 2004). Further investigations on the neural architecture and the experimental condition are needed to draw conclusions on the compensatory eye movements of animals.




Figure 2.7: A bipedal humanoid robot with a pan-tilt camera mounted on its head is asked to reach the goal (white square) in the walled arena containing cylindrical obstacles in simulation. The location and orientation of the robot as well as the position of each obstacle are randomly initialized at the beginning of each test. The environment and the robot were simulated with Webots<sup>TM</sup> (<http://www.cyberbotics.com>). Bottom: A trajectory of the best evolved robot is drawn by the dotted line. Black disks show obstacles.

## 2.5 Active Perception in Cooperative Tasks

So far we have considered a single retina only for each task. However, in general multi-agent systems can complete tasks more quickly and reliably than single-agent systems in many practical applications. If systems can communicate with each other in an effective manner, their performance would be further improved.

To test this hypothesis, we performed a set of experiments (Suzuki et al., 2008b) where two active vision systems were asked to solve behavioral tasks together. The tasks were shape localization and discrimination. The former task and the experimental setup were exactly same as those described in Section 2.2.1. The latter task required both agents to fixate on the target shape. In order to endow the agents with the ability to communicate with each other, we slightly extended the original neural architecture by adding two sets of neurons, namely ‘hearing’ and ‘speaking’ neurons (Fig. 2.8, top). ‘Speaking’ neurons emitted sub-symbolic signals (i.e., output values of the two neurons) to the other agent while ‘hearing’ neurons received signals from the other agent every time step. Through these neurons the two agents could communicate with each other in a primitive way using continuous sub-symbolic signals.

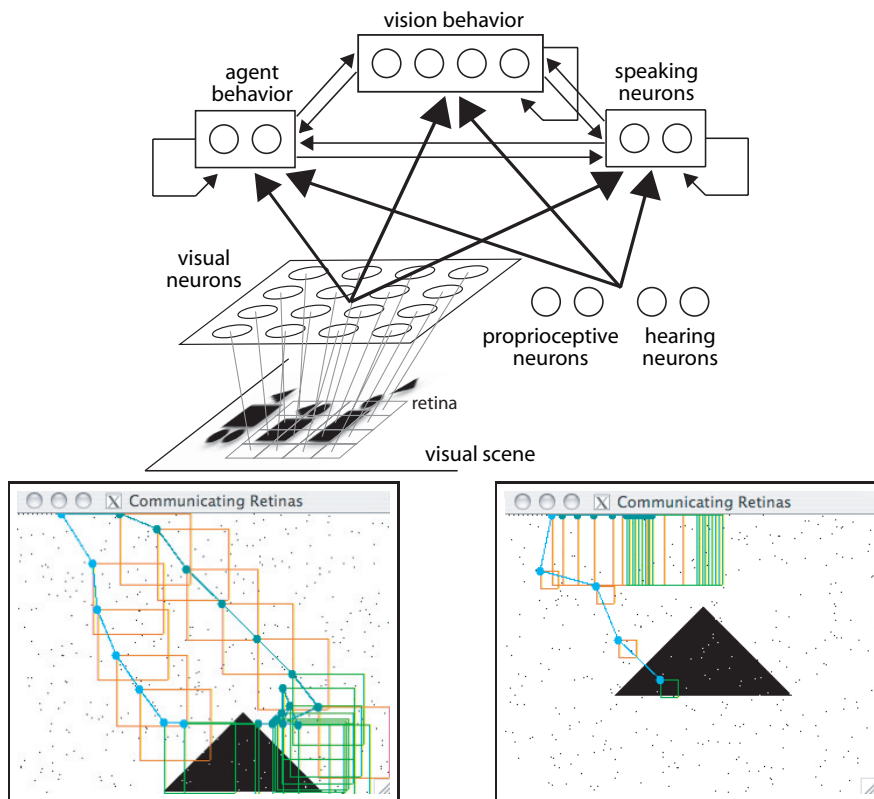


Figure 2.8: Top: An extended neural architecture for communicating active vision agents. Both agents have the same architecture. Two sets of neurons, namely ‘hearing’ and ‘speaking’ neurons are added to the original architecture shown in Fig. 2.1. ‘Speaking’ neurons emit signals (i.e., output values of the two neurons) to the other agent while ‘hearing’ neurons receive signals from the other agent every time step. Bottom left: The best evolved pair performing shape *localization* task. Bottom right: The best evolved pair performing shape *discrimination* task. The two agents exploit communication to perform the tasks efficiently.

The two active vision agents were coevolved by using a genetic algorithm and the best evolved pair exhibited cooperative behavior by means of an effective communication. In the shape discrimination task, the behavioral strategy of the best evolved pair was that they moved in parallel while maintaining a small distance between agents (see Fig. 2.8, bottom left). This cooperative strategy allowed them to wipe a larger area. The strategy also allowed one agent to immediately pull the other to its position when the agent found the target shape. This cooperation was made possible by the communication signals which transmitted the information of retina movements to each other. In the shape localization task, the strategy of the best evolved pair was very different from the one developed in the previous task; since they needed just to continuously emit correct responses (either square or triangle), they did not cooperate with each other. Instead, the ‘role specialization’ emerged by exploiting communication; one ‘active’ agent always looked for the target while the other ‘lazy’ agent just blindly trusted and copied the answer given by the active agent (see Fig. 2.8, bottom right).

In order to assess the advantages of cooperation and communication, we also measured the performance of the best evolved *single* retina and of the best evolved pair of retinas *without* communication in the same tasks. The best evolved pair of retinas capable of communication outperformed the best evolved single retina and the best evolved pair without communication. These results suggest that sub-symbolic communication greatly improves joint attention between agents. It is interesting to investigate how such sub-symbolic communication and cooperative behavior between two agents affect and interact with the emergence of grammar or language (Steels, 1997). For more details of this experiment, see Appendix B.

## 2.6 Summary

This chapter illustrated the basic framework that will be further extended in the following chapters. These experiments indicate that coevolution of visual features and behavior can address a variety of visual tasks by means of very simple architectures and computational abilities. Evolved individuals can solve position- and size-invariant tasks exploiting position- and size-variant receptive fields by actively searching and maintaining simple features of the visual scene over sensitive areas of the retina. The architecture can also be extended to various embodiments and multi-agent systems.

Active behavior affects, interacts with, and supports visual processing by selecting sensory experiences that can be dealt with by the system in a coherent manner. These experiments indicate that behavior is not only a variable to be considered in an ecological study of visual perception, but is also intimately related to the way in which vision mechanisms develop and are exploited by the system.



---

# 3

## Capturing Spatio-Temporal Relationships

---

*Cognitive systems are often faced with the necessity to capture not only spatial relationships (as those of the previous experiments), but also temporal relationships between successive events. For example, when navigating in complex environments, we actively search for landmarks and formulate expectations on what is ahead of us. In this chapter, this anticipatory mechanism is explored in our robotic setup. This chapter is based on Suzuki and Floreano (2006).*

**Abstract.** Active vision may be useful to perform landmark-based navigation where landmark relationship requires active scanning of the environment. In this chapter we explore this hypothesis by evolving the neural system controlling vision and behavior of a mobile robot equipped with a pan/tilt camera so that it can discriminate visual patterns and arrive at the goal zone. The experimental setup employed in this chapter requires the robot to actively move its gaze direction and integrate information over time in order to accomplish the task. We show that the evolved robot can detect separate features in a sequential manner and discriminate the spatial relationships. An intriguing hypothesis on landmark-based navigation in insects derives from the present results.

### 3.1 Introduction

Active vision may be useful to perform landmark-based navigation where landmark relationship requires active scanning of the environment. In this chapter we explore this hypothesis by evolving the neural system controlling vision and behavior of a mobile robot equipped with a pan/tilt camera so that it can discriminate visual patterns and arrive at the goal zone.

The experimental setup employed in this chapter has a notable characteristic: the visual landmarks are identical if the elevation of a robot's camera is fixed with the body. In that case, the robot could be unable to discriminate one from the other<sup>1</sup>.

---

<sup>1</sup>The use of a panoramic camera which provides larger field of view is discussed in section 3.4.

It needs to actively move its gaze direction and integrate information over time in order to differentiate these patterns. The sequential detection of spatially separate visual landmarks has been largely neglected in the literature. Instead most machine vision systems process an entire image of their large visual field every time step.

We show that the best evolved robots successfully perform the task by using an effective scanning strategy. The evolved active scanning trajectory covers only a small region of the entire visual field and, more importantly, consists of a sequence of feature-driven, anticipatory, and context-dependent gaze movements. We address the advantages of the present method and neural architecture in terms of algorithmic, computational and memory resources.

The rest of this chapter is organized as follows: the next section details the experimental setup, i.e. the environment, the simulated robot and the task for the robot. The neural network embedded in the robot and the genetic algorithm for developing the synaptic weights in the neural network are also described. Results and the analysis of the best evolved individual are described in Section 3.3. Finally an intriguing hypothesis on landmark-based navigation deriving from the present results and conclusions are discussed in Section 3.4 and 3.5 respectively.

## 3.2 Method




Figure 3.1: Left: The original six-wheeled robot Koala equipped with a pan/tilt camera. Right: The robot’s perspective in a simulated environment. The robot can access the world with 5 by 5 retina at the center of the image.

The neural control system of a mobile robot equipped with a pan/tilt camera is evolved by means of a genetic algorithm to perform goal-directed navigation in an enclosed space using only visual information (Fig. 3.1). The evolutionary algorithm evaluates each neural controller with random mutations until an evolutionary stable control strategy is found (Nolfi and Floreano, 2000). In order to collect data from several independent runs and perform rigorous statistical analysis, we used fast, physics based simulations of the robot and its environment (Fig. 3.1).

We simulated the robot and the environment using physics-based Vortex libraries<sup>2</sup>. The robot has six wheels, but only the central wheel on each side is motorized. The robot base is 30cm(W) $\times$ 32cm(L) $\times$ 20cm(H). The pan and tilt angles of the camera are controlled by two separate and independent motors.

### 3.2.1 Experiment and Task




Figure 3.2: The arena (200cm $\times$ 600cm) and two visual patterns used in the simulation. The visual field of the robot can not cover both of the two black squares at any given moment. The difference of the two walls resides in the spatial relationship of the two squares (bottom). The position and direction of the robot are randomized at the beginning of each test.

Figure 3.2 shows the experimental setup where each of two facing white walls has two squares placed at different heights. The task of the robot is to visually discriminate one wall from the other in order to arrive at the goal zone at the end of each trial. There is no other identification of the goal than the visual patterns. Importantly this experimental setup is designed such that it does not allow the visual field of the robot to cover both of the two black squares at any given moment. Therefore the robot cannot discriminate the two walls by keeping the vertical angle of the camera constant because both walls have an identical black square in the same height. The difference of the two walls resides in the spatial relationship of the

<sup>2</sup><http://www.cm-labs.com>

two squares (Fig. 3.2, right). The robot needs to discriminate one pattern from the other by using active, sequential scanning of the two black squares of each pattern and integrating the information over time.

### 3.2.2 Neural Architecture and Genetic Algorithm




Figure 3.3: The neural architecture is composed of a grid of visual neurons with non-overlapping receptive fields whose activation is given by the gray level of the corresponding pixels in the image; an object detector unit that is activated when any visual neuron is strongly activated; a hidden unit with incoming synapses from visual neurons; a set of proprioceptive neurons that provide information about the movement of the camera with respect to the chassis of the robot; a set of output neurons that determine at each sensory motor cycle the filtering used by visual neurons, the new pan and tilt speeds of the camera and the rotational speeds of the two wheels of the robot; a set of associative neurons with recurrent connections. Solid arrows between layers represent fully connected synaptic weights. Dashed arrow represents a predetermined (non-evolvable) OR filter (see main text for more detail).

The neural network is characterized by a feedforward architecture with evolvable thresholds and discrete-time, fully recurrent connections at the associative layer (Fig. 3.3). A set of visual neurons, arranged on a grid, with non-overlapping receptive fields receives information about the gray level of the corresponding pixels in the image provided by the camera on the robot. The receptive field of each unit covers a square area of 48 by 48 pixels in the image. We can think of the total area spanned by all receptive fields (240 by 240 pixels) as the surface of an artificial retina. The activation of a visual neuron, scaled between 0 and 1, is given by the average gray level of all pixels spanned by its own receptive field or by the gray level



of a single pixel located within the receptive field. The choice between these two activation methods, or filtering strategies, can be dynamically changed by one output neuron at each time step. An object detector unit is activated when any visual neuron is strongly activated. Therefore the synaptic weights incoming into this unit can be seen as a predetermined (non-evolvable) OR filter. Two proprioceptive units provide input information about the measured horizontal (pan) and vertical (tilt) angles of the camera. These values are in the interval  $[-100, 100]$  and  $[-25, 25]$  degrees for pan and tilt, respectively. Each value is scaled in the interval  $[-1, 1]$  so that activation 0 corresponds to 0 degrees (camera pointing forward parallel to the floor). A set of memory units store the values of the associative neurons at the previous sensory motor cycle step and send them back to the associative units through a set of connections, which effectively act as recurrent connections among associative units (Elman, 1990). The bias unit has a constant value of  $-1$  and its outgoing connections represent the adaptive thresholds of associative, hidden and output neurons (Hinton and Sejnowski, 1999).

Associative, hidden and output neurons use the sigmoid activation function  $f(x) = 1/(1 + \exp(-x))$  in the range  $[0, 1]$ , where  $x$  is the weighted sum of all inputs. Output neurons encode the motor commands of the active vision system and of the robot for each sensory motor cycle. One neuron determines the filtering strategy used to set the activation values of visual neurons for the next sensory motor cycle. Two neurons control the movement of the camera, encoded as speeds relative to the current position. The remaining two neurons encode the direction and rotational speeds of the left and right motored wheels of the robot. Activation values above and below 0.5 stand for forward and backward rotational speeds respectively.

The present neural architecture has been incrementally developed based on our previous investigations (Floreano et al., 2004, 2005). The object detector unit incorporated in the architecture is explicitly designed to simplify the biological visual system capable of monitoring for change in the visual environment<sup>3</sup>. The hidden neuron is incorporated to equalize the contributions of the visual neurons, the object detector unit and the proprioceptive units to the activations of the associative neurons. The roles of the hidden and object detector units are further analyzed in section 3.3.

The neural network has 106 evolvable connections that are individually encoded on five bits in the genetic string (total length=530 bits). A population of 100 genomes is randomly initialized by the computer. Each individual genome is then decoded into the connection weights of the neural network and tested on the robot while its fitness is computed. The best 20% of the population (those with the highest fitness values) are reproduced, while the remaining 80% are discarded. Equal number of copies of the selected individuals are made to create a new population of the same size. The new genomes are randomly paired, crossed over with probability 0.1 per pair and mutated with probability 0.01 per bit. Crossover consists in swapping

---

<sup>3</sup>In our preliminary studies it seemed difficult to develop the visual system capable of significantly responding to the black squares detected at any location of the retina.

genetic material between two strings around a randomly chosen point. Mutation consists in toggling the value of a bit. Finally two copies of the best genomes of the previous generation are inserted in the new population at the places of the randomly chosen genomes (elitism) in order to improve the stability of the evolutionary process.

The fitness function was designed to select robots for their ability to arrive at the goal zone at the end of each life. Each individual is tested for six trials, each trial lasting for 300 sensory motor cycles. A trial can be truncated earlier if the operating system detects an imminent collision into the walls. The fitness criterion  $F$  is composed as follows:

$$F = F_{speed}(S_{left}, S_{right}) + F_{goal} \quad (3.1)$$

where  $F_{speed}(S_{left}, S_{right})$  is a function of the measured speeds of the left  $S_{left}$  and right  $S_{right}$  wheels and  $F_{goal}$  is a reward given if the robot reaches the goal at the end of its life<sup>4</sup>. More specifically  $F_{speed}(S_{left}, S_{right})$  is defined as follows:

$$F_{speed}(S_{left}, S_{right}) = \frac{1}{ET} \sum_{e=0}^E \sum_{t=0}^{T'} f(S_{left}, S_{right}, t) \quad (3.2)$$

$$f(S_{left}, S_{right}, t) = (S_{left}^t + S_{right}^t) \left(1 - \sqrt{|S_{left}^t - S_{right}^t| / 2S_{max}}\right) \quad (3.3)$$

where  $S_{left}$  and  $S_{right}$  are in the range  $[-8, 8]$  cm/sec and  $f(S_{left}, S_{right}, t) = 0$  if  $S_{left}$  or  $S_{right}$  is smaller than 0 (backward rotation);  $E$  is the number of trials (six in these experiments),  $T$  is the maximum number of sensory motor cycles per trial (300 in these experiments),  $T'$  is the observed number of sensory motor cycles;  $F_{goal}$  is 10 if the robot reaches the goal area at the end of the test, otherwise it is 0. The reward is given only if at least one lower square and one upper square are detected before reaching the goal. This stronger constraint on  $F_{goal}$  is to prevent selecting 'blind' individuals which arrive at the goal by chance without using visual patterns.

At the beginning of each trial the position and orientation of the robot are randomized in the interval  $[-50, 50]$  and  $[-20, 20]$  for the longitudinal and short axes respectively.

### 3.3 Results and Analysis

We performed six replications of the evolutionary run starting with different initial populations. In all cases the fitness reached stable values in less than 30 generations (Fig. 3.4), and the fitness value of the best evolved individual ranged from 40 to 60.

We analyzed the behavior of the best evolved individual which arrived at the goal six times out of six trials. Figure 3.5 shows the scanning strategy, the trajectory of

---

<sup>4</sup>One might think that the first term  $F_{speed}(S_{left}, S_{right})$  is not necessary, but in our preliminary study the fitness value remained zero without  $F_{speed}(S_{left}, S_{right})$ , meaning that evolution could not find the solution.




Figure 3.4: Evolution of neural controllers for the simple three dimensional landmark navigation. Fitness values of the population average (thin line) and the best individual (thick line) across 30 generations. Vertical bars show the standard deviation. The results are averaged over six evolutionary runs.

the robot, the camera movement with respect to the chassis of the robot and the activation of neuron 5 in the associative layer when the robot started in the face of pattern A and B. For clarity we show only the activation of neuron 5 because we found that it played the most significant role in the pattern discrimination.

The behavioral strategy of the best evolved robot can be illustrated as follows: 1. The robot searches for a lower square by moving the camera left-downward and turning its chassis counter-clockwise until it finds one; 2. Once it finds one, it points the camera right-upward to find an upper square; 3. If it finds an upper square after a short delay, it goes toward the goal while moving the camera left-downward. If not, it moves the camera left-downward again while turning its chassis counter-clockwise until it finds another lower square, and then goes back to step 2. Thus the robot always searches for pattern B to go toward the goal.

We studied the role of the hidden and object detector neurons by lesioning one at a time. Their operation was disrupted by clamping the activation value of the neuron to a constant value of 0.5 during behavior. Figure 3.6 (left) shows that both neurons significantly contribute to the successful performance. The best evolved individual while the object detector neuron was lesioned arrived at the goal zone five times out of 20 trials (10 in the face of pattern A plus 10 in the face of pattern B, at the beginning). However these successes were achieved only when the robot started facing pattern B. If the robot was facing pattern A it went in the opposite direction of the goal. In other words, the robot always goes right in the face of both pattern A and B. This result suggests the crucial role of the object detector neuron in the behavior selection or decision making. That neuron significantly contributes to measuring the time interval between looking right-upward and subsequent detection of an upper square. Without the object detector neuron the robot can not measure the time interval and therefore can not discriminate the patterns.

While the hidden neuron was lesioned, the best evolved individual never arrived

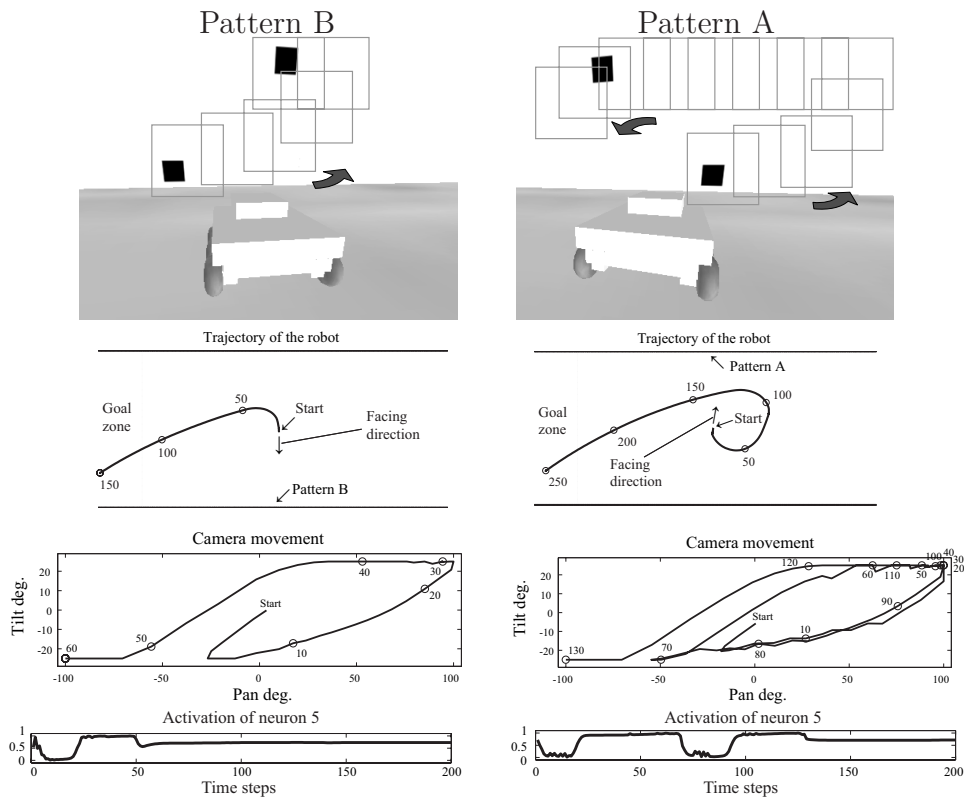


Figure 3.5: From top to bottom, the best evolved robot scanning two black squares of each pattern sequentially (gray squares depicting the trajectory of the retinal perimeter), the trajectory of the robot, the camera movement with respect to the chassis of the robot and the activation of neuron 5 in the associative layer during the behavior (shown only for the first 200 sensory motor cycles) when the robot started in the face of pattern B (left column) and A (right column).

at the goal. This result suggests that the individual uses not solely the temporal information given by the object detector neuron, but also the visual information given through the hidden neuron.

One might think that the scanning strategy is reactive, i.e. the detection of a lower or an upper square always activates a particular behavior, but it is not. For example, in Fig. 3.5 (right) the lower square of pattern A was detected for the second time in the left side of the robot around the 130th time step, but this event did not affect the behavior of the robot going toward the goal. Therefore it seems that the behavior had been ‘switched on’ before the event<sup>5</sup>. The decision might be made when the upper square of pattern B was detected shortly after looking right-upward. If an upper square is detected late after looking right-upward, the robot does not go toward the goal, but resumes searching for a lower square.

<sup>5</sup>The stable activation of the neuron 5 around 0.7 (see Fig 3.5, bottom) seems to reflect such a fixed behavior after the decision making.

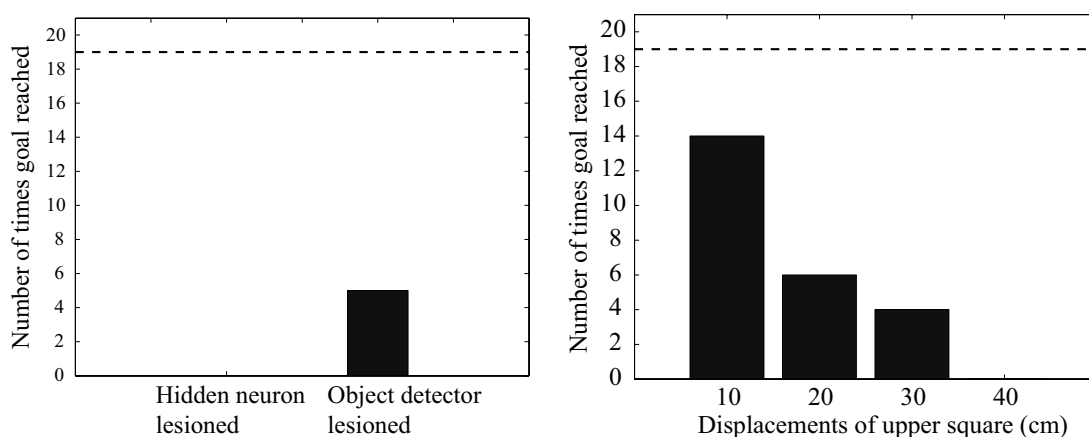


Figure 3.6: The number of successful arrivals at the goal is counted in each condition out of 20 trials. Left: Lesion test of the best evolved individual. Horizontal dotted line shows the score of the intact best evolved individual. Right: Test of the best evolved individual when upper squares are displaced. Horizontal dotted line shows the score when upper squares are *not* displaced.

This hypothesis was supported by another set of analyses where the upper square in each pattern was horizontally shifted toward the center (Fig. 3.6, right). The robot can not discriminate the two patterns any more if the upper square is shifted more than 20 cm.

The importance of the proprioceptive inputs is validated by another set of evolutionary runs without proprioceptive inputs (Fig. 3.7, left). Despite the shorter length of the genetic string (total length=480 bits), the best evolved individuals in all six evolutionary runs could reach the goal only three times out of six trials at maximum. Their behavioral analysis shows that these individuals always go left (or right depending on the evolutionary run) in the face of both pattern A and B. In other words they do not differentiate one pattern from the other.

One more set of evolutionary runs with another neural architecture which has fully recurrent connections at the output layer and does not have the associative layer shows worse fitness values than those with the original neural architecture (Fig. 3.7, right).

### 3.4 Discussion

We have shown that the evolved robot can detect two separate features in a sequential manner and discriminate the spatial relationships. If the system can perform active vision and sequentially store the events of visual feature detection, we do not need expensive computational power nor large memory storage capacity which would be required to resort to image memorization and matching. Although it has been shown that insects may indeed adopt such an image memorization and matching

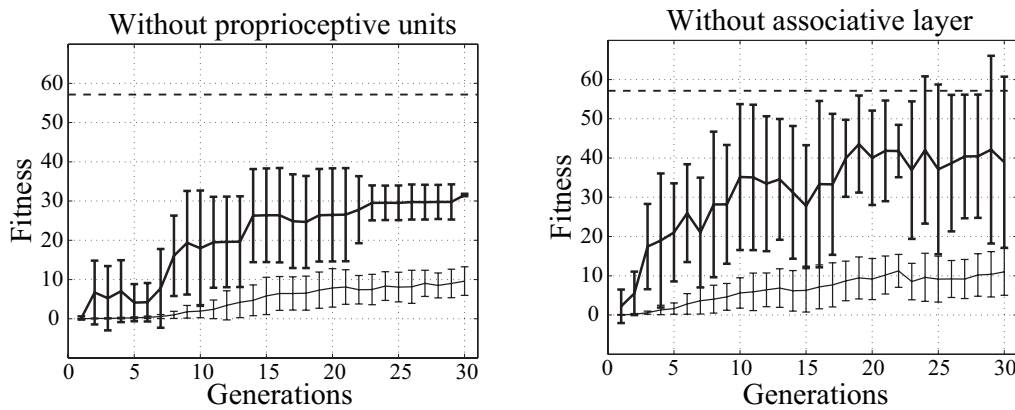


Figure 3.7: Left: Evolution *without* proprioceptive inputs encoding pan and tilt movements. Right: Evolution *without* the associative layer. Fitness values of the population average (thin line) and the best individual (thick line) across 30 generations. Vertical bars show the standard deviation. Averaged over six evolutionary runs. Horizontal dotted line shows the averaged fitness value of the best evolved individual with the original neural architecture (see Fig. 3.4).

strategy (Judd and Collett, 1998), it is tempting to speculate that their tiny brain with restricted memory capacity may favor a more economical strategy as shown in this chapter.

The evolved robot was able to effectively scan small regions of the broad visual field in an anticipatory manner in order to sequentially detect separate features. Such a characteristic of the evolved scanning strategy is in agreement with the evidence shown in (Land and Furneaux, 1997; Land and McLeod, 2000) that people direct their gaze to points of the scene where information is to be extracted. Land et al. recorded human eye movements while playing cricket and table tennis. The eyes are very active and their activity takes roughly the same path as the ball. Contrary to popular belief, they do not follow the ball, but work in an anticipatory way. For example, eyes anticipate the position of the ball before it bounces, making saccades to positions where there is, as yet, no visible stimulus. In this chapter we have shown a computational model capable of an anticipatory “eye” movement in a freely moving behavioral system.

In order to detect the spatially separate features, the evolved robot executes a particular scanning sequence in front of the visual patterns. That is, after detecting a lower square, the robot routinely directs its gaze right-upward. Such a scanning sequence might be reminiscent of the human ‘scanpath’ during facial recognition (Noton and Stark, 1971b,a): Noton and Stark claimed that when a particular visual pattern is viewed, a particular sequence of eye movements is executed and furthermore that this sequence is important in accessing the visual memory for the pattern. The evolved scanning strategy presented in this chapter is similar to the ‘scanpath’ in that the moving sequence is crucial to identify a particular pattern. However,

notice that the evolved scanning strategy is not for accessing the visual memory, but rather is tightly coordinated with the behavior of the robot.

From an engineering point of view one may argue that a panoramic camera could allow the robot to cover the entire visual field and discriminate the two patterns. However this approach would be computationally expensive if the entire image is to be uniformly processed in high resolution to extract tiny features out of a vast visual field as we have shown in this chapter. Active vision applied to an omnidirectional image is studied in a separate article (Suzuki et al., 2006).

Although the present neural architecture shown in Fig. 3.3 was investigated in the lesion test and additional evolutionary runs with modified neural architectures, further investigations must be done. We intend to identify the minimum components necessary for the neural controller of the robot to detect spatially separate features in the three dimensional visual environment.

## 3.5 Conclusion

In this chapter we have shown that active vision may help not only to locate important features of the environment, but also to capture spatio-temporal relationships between those features that could provide behaviorally relevant information.

From these results it can be hypothesized that landmark-based navigation in insects and robots could be mediated by similar mechanisms instead of resorting to image memorization and matching (Judd and Collett, 1998). We are currently exploring this hypothesis with simulated and physical robots.





---

# 4

## Landmark-Based Navigation

---

*Promising results described in the previous chapter suggest that active vision approach could be further applied to more realistic landmark-based navigation tasks. In this chapter, robots are evolved in similar environments to those used in the experiments of insect visual navigation.*

**Abstract.** Insects are impressive navigators despite the fact that they have far fewer neural resources than vertebrates. Recent studies of insect visual navigation suggest that visual landmarks might act primarily as ‘signposts’ that tell insects what particular action they need to perform, rather than telling them where they are. However the neural mechanism in insect brain allowing such landmark guidance is not well understood. In this chapter we explore this issue by evolving the neural controller of a mobile robot. We adopt a deliberately simple neural architecture, embed it in a mobile robot equipped with a low-resolution, active vision system, and develop its synaptic weights by a genetic algorithm. We show that the best evolved robot exhibits robust homing navigation by actively selecting relevant visual features and associating them to motor actions. The trajectories of the robots are strikingly similar to those of desert ants. Analysis of the neural activities recorded during behavior of the robot illustrates how robots perform such homing navigation. Although evolved controllers exploit visual signposts, they do so by actively searching for relevant visual features that trigger suitable motor actions rather than taking snapshots of the visual scene.

### 4.1 Introduction

Insects are impressive navigators despite the fact that they have far fewer neural resources than vertebrates (Healy, 1998). Much evidence suggests that many species of insects could do so without a rich spatial representation –or a map– of the environment (Collett et al., 2007). Instead, they mainly use path integration and visual landmark detection, although visual navigation is more dominantly used (Collett

and Collett, 2002).

Recent studies of insect visual navigation suggest that visual landmarks might act primarily as ‘signposts’ that tell insects what particular action they need to perform, rather than telling them where they are (Collett and Collett, 2002). Ants tend to ignore compass cues under certain conditions (Collett et al., 2001). However the neural mechanism in insect brain allowing such landmark guidance is not well understood.

We take a synthetic approach and explore possible neural mechanisms with a mobile robot capable of *active vision*. It picks out the properties of images which are necessary to perform its assigned task at hand, and ignores the rest (Bajcsy, 1988; Aloimonos et al., 1987; Aloimonos, 1990; Ballard, 1991; Findlay and Gilchrist, 2003). In this context, there is no clear need for the sort of detailed reconstructions of the visible world that have been an accepted, traditional goal of machine vision (Horn, 1986).

The ‘snapshot model’ of insect navigation, which assumes the storage of a photograph-like, retinotopically fixed image (Judd and Collett, 1998), is appealing to our intuition and there exists some evidence to support this hypothesis (Collett and Cartwright, 1983; Cartwright and Collett, 1987). However this model can be questioned in some respects; for instance, it requires increasingly large memory capacity as the journey of the insect become longer, as argued in (Srinivasan, 1998; Pratt et al., 2001). It is therefore worthwhile considering alternative, more economical, mechanisms.

In this chapter, we show that evolved robots with active vision exhibit robust homing navigation with a small amount of computational and memory resources. The trajectories and behaviors of robots are strikingly similar to those of desert ants in similar environments. To perform the task, robots search for edges and associate them with navigation directions instead of memorizing and matching a batch of photographic images.

## 4.2 Method

The robot and the genetic algorithm used in this experiment is the same as those described in the previous chapter. We study the evolution of homing strategies in two separate environmental conditions that differ in the relative positions of the robot, of the target location, and of a black wall (Fig. 4.1). The task of the robot is to go toward and stay at the nest position until the end of each trial. The nest is invisible and the only visible element is the black wall. This experimental setup is identical to the one used in experiment of desert ant navigation (Collett et al., 2001). The two environments are used for the purpose of checking if a common strategy is used in both of them, even if robots experience only one.

The neural architecture adopted in this experiment (Fig. 4.2) is very similar to the one described in the previous chapter (see Fig. 3.3); the only difference is that the present architecture does not have an object detector unit, but have simple

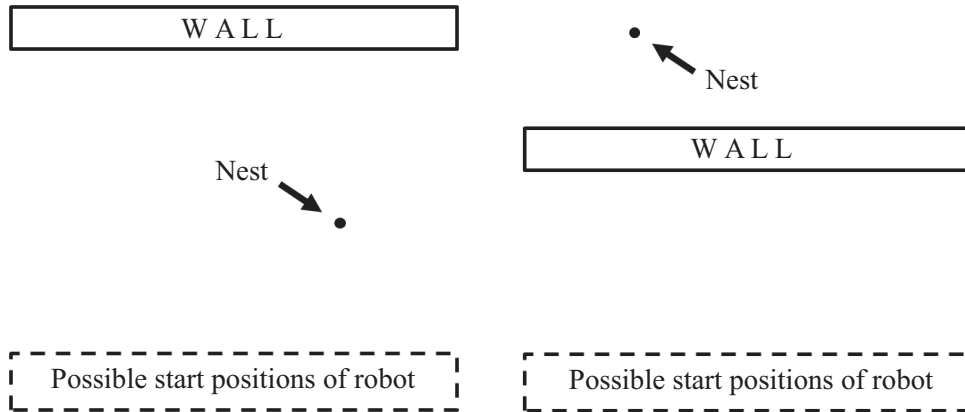


Figure 4.1: The evolutionary environment 1 (left) and 2 (right) containing a black wall. The size of the wall is  $600\text{cm}(W) \times 60\text{cm}(L) \times 300\text{cm}(H)$  in both environments. The position and orientation of the robot are randomized at the beginning of each trial.

hidden neurons instead.

In each environment we performed six replications of the evolutionary run starting with different initial populations. The fitness function was designed to select robots for their ability to arrive at the nest position at the end of each life (see Appendix for more detail). Robots are always evaluated in a single environment, either 1 or 2, during the evolutionary run. Each individual is tested for four trials, each trial lasting for 450 sensory motor cycles. A trial can be truncated earlier if the operating system detects an imminent collision into the wall. At the beginning of each trial the position and orientation of the robot are randomized (see Fig. 4.1).

### 4.3 Results

Figure 4.3 shows the fitness values across generations for the two environmental conditions. In several evolutionary runs, fitness values of the best individual reached the maximum value in both environments, which means that the best evolved robot could successfully reach the nest position in all four trials. Comparison of the two fitness graphs shows that Environment 2 is more demanding than Environment 1 for the robot because the robot must avoid and circumvent the wall.

The best evolved individual in Environment 1 shows robust homing navigation from different positions and orientations (Fig. 4.4, left). Figure 4.5 illustrates how the robot utilizes the wall to pinpoint the nest position. The behavioral strategy can be described as follows: 1) the robot searches for the top edge of the wall while keeping the camera left-upward; 2) once the robot finds the edge, the robot moves straight while keeping the edge at a certain location of the retina; 3) the robot starts turning on the spot once it loses sight of the edge. This spot is the expected nest position. The camera was always pointing left-upward with respect to the chassis




Figure 4.2: The neural architecture is composed of a grid of visual neurons with non-overlapping receptive fields whose activation is given by the gray level of the corresponding pixels in the image; two hidden units with incoming synapses from visual neurons; a set of proprioceptive neurons that provide information about the movement of the camera with respect to the chassis of the robot; a set of output neurons that determine at each sensory motor cycle the filtering used by visual neurons, the new pan and tilt speeds of the camera and the rotational speeds of the two wheels of the robot. Solid arrows between layers represent fully connected synaptic weights.




Figure 4.3: Evolution of neural controllers in Environment 1 across 40 generations (left) and in Environment 2 across 100 generations (right). Fitness values of the population average are shown with thin lines and the best individual with thick lines. The results are averaged over six evolutionary runs. Vertical bars show the standard deviation.




Figure 4.4: The best evolved individual in Environment 1 was tested 40 times in the same environment. The robot successfully reached the nest location from random positions and orientations (left). The yellow diamond shows the nest position and red disks show the end points of each trajectory. Right: The view of the robot near the goal position.

of the robot.

As in (Collett and Collett, 2002), the best evolved individual in Environment 1 was tested 40 times in an environment where the wall and the nest were rotated 45 degrees counterclockwise. The robot started from the same random positions and orientations used during evolution. The robot could successfully pinpoint the new nest position, thus replicating the performance of desert ants (Fig. 4.6) (Collett et al., 2001). Notice that the robot has never experienced the rotation of the wall during evolution.

The best individual evolved in Environment 2 was also tested in the same environment 40 times (Fig. 4.7). The robot successfully reached the nest location from random orientations and positions. Figure 4.8 illustrates how the robot utilizes the wall to pinpoint the nest position. The behavioral strategy can be described as follows: 1) the robot turns clockwise to search for the left edge of the wall while keeping the camera upward; 2) once the robot finds the edge, the robot turns right slightly while keeping the edge at a certain location of the retina; 3) after passing the wall, the robot goes straight toward the nest while detecting the top-right corner; 4) the robot starts turning on the spot once it loses sight of the corner. This spot is the expected nest position. The camera was always pointing right-upward with respect to the chassis of the robot.

The best evolved individual in Environment 2 was also tested 40 times in an environment with the wall rotated 45 degrees counterclockwise (Fig. 4.9). The robot started from the same random positions and orientations used during evolution. The trajectories of the robot are strikingly similar to those of desert ants shown in (Collett




Figure 4.5: Analysis of pinpointing behavior in Environment 1 during which all neural activities were recorded. The camera of the robot was always pointing left-upward with respect to the chassis of the robot.




Figure 4.6: The best evolved individual in Environment 1 was tested 40 times in a new environment with the wall rotated 45 degrees counterclockwise. The robot successfully reached the new nest location from random positions and orientations which were not rotated 45 degrees. The yellow diamond shows the new nest position and red disks show the end points of each trajectory. Notice that the robot has never experienced the rotation of the wall during evolution.




Figure 4.7: The best evolved individual in Environment 2 was tested 40 times in the same environment. The robot successfully reached the nest location from random positions and orientations (left). The yellow diamond shows the nest position and red disks show the end points of each trajectory. Right: The view of the robot near the goal position.




Figure 4.8: Analysis of pinpointing behavior in Environment 2 during which all neural activities were recorded. The camera of the robot was always pointing right-upward with respect to the chassis of the robot.




Figure 4.9: The best evolved individual in Environment 2 was tested 40 times in a new environment with the wall rotated 45 degrees counterclockwise. The robot started from random positions and orientations which were not rotated 45 degrees. The yellow diamond shows the ‘new’ nest position and red disks show the end points of each trajectory. Notice that the robot has never experienced the rotation of the wall during evolution.

and Collett, 2002). Notice again that the robot has never experienced the rotation of the wall during evolution.

## 4.4 Discussion

The strategies evolved in the two environments are different because Environment 2 requires robots to avoid and circumvent the wall whereas Environment 1 does not. However both strategies share a common characteristic; evolved robots actively search for relevant visual features that trigger suitable motor actions rather than taking snapshots of the visual scene (Judd and Collett, 1998) nor computing homing vectors (Möller et al., 1998).

The evolutionary experiments seem to lend support for a landmark-based navigation strategy that is based on locating edges and advancing so to keep those edges at a relative distance and orientation. The evolved strategies resemble the



model suggested by (Pratt et al., 2001); Pratt *et al.* showed that ants (*Leptothorax albipennis*) could be guided by mere edges along the route to their nest, rather than retinotopic snapshots. Their model was very simple and therefore did not require a lot of memory capacity as long as it was able to robustly extract edges. Further experiments are necessary to check if ants can also utilize edges that are not parallel to the route, as shown in our experiments with mobile robots.

In this chapter we hypothesized and tested a neural mechanism capable of active vision in order to explain the landmark guidance behavior (Collett and Collett, 2002). We described a set of experiments showing that a mobile robot endowed with a simple, yet active vision system coupled with a simple recurrent neural network can perform robust homing navigation as desert ants do in similar environments. Although evolved controllers exploit visual signposts, they do so by actively searching for relevant visual features that trigger suitable motor actions rather than taking snapshots of the visual scene. This strategy does not seem to require the memorization of images (Judd and Collett, 1998) nor the computation of homing vectors (Möller et al., 1998). It agrees with the experimental data on horizontal edge detection by desert ants (Pratt et al., 2001).

## Appendix: Fitness Function

The fitness criterion  $F$  is intended to select robots for their ability to arrive at the nest position at the end of each life. More precisely, it is composed as follows:

$$\begin{aligned}
 F &= \frac{1}{E} \sum_i^E f_i & (4.1) \\
 f_i &= \begin{cases} 1 - d/D_{max} & : \text{ if } d < D_{max} \\ 0 & : \text{ otherwise.} \end{cases}
 \end{aligned}$$

where  $E$  is the number of trials (four in these experiments);  $d$  is the distance between the end position of the robot and the nest position;  $D_{max}$  is the maximum value of  $d$  (10 meters in these experiments).



---

# 5

## Evolution and Ontogenetic Adaptation

---

*We further investigated the “ontogenetic development” of receptive fields in an evolutionary mobile robot with active vision. In contrast to our previous experiments where synaptic weights for both receptive field and behavior were genetically encoded and evolved on the same time scale, here the synaptic weights for receptive fields developed during the life of the individual. The results clearly illustrate the strong interplay between active vision and receptive field formation. This chapter is based on Floreano et al. (2005).*

**Abstract.** In this chapter, we describe the artificial evolution of adaptive neural controllers for an outdoor mobile robot equipped with a mobile camera. The robot can dynamically select the gazing direction by moving the body and/or the camera. The neural control system, which maps visual information to motor commands, is evolved online by means of a genetic algorithm, but the synaptic connections (receptive fields) from visual photoreceptors to internal neurons can also be modified by Hebbian plasticity while the robot moves in the environment. We show that robots evolved in physics-based simulations with Hebbian visual plasticity display more robust adaptive behavior when transferred to real outdoor environments as compared to robots evolved without visual plasticity. We also show that the formation of visual receptive fields is significantly and consistently affected by active vision as compared to the formation of receptive fields with uniformly sampled images in the environment of the robot. Finally, we show that the interplay between active vision and receptive field formation amounts to the selection and exploitation of a small and constant subset of visual features available to the robot.

### 5.1 Introduction

Biological vision systems filter, compress, and organize the large amount of optical stimulation as electrical signals proceed from the retina toward deeper structures of the brain. This data reduction is achieved by a layered, distributed, and topo-

logically organized set of neurons that individually respond to specific aspects of the optical stimulus. In mammals, for example, neurons in the early stage of the visual cortex selectively respond to particular features of the environment, such as oriented edges (Hubel and Wiesel, 1968), that are linear combinations of the pattern of retinal activations. Neurons in later stages of the visual cortex respond to more complex patterns that also take into account the direction of movement of the stimulus and cannot easily be reduced to a linear combination of lower-level features (Wandell, 1995).

The features that trigger the response of a neuron represent the *receptive field* of that neuron. The receptive fields of cortical visual neurons are not entirely genetically determined, but develop during the first weeks of the newborn baby and there is evidence that this process may already start before birth. Studies of newborn kitten raised in boxes with only vertical texture show that these animals do not develop as many receptive fields for horizontal features as kitten raised in normal environments (Blakemore and Cooper, 1970) and therefore see the world in a different way. The development of visual receptive fields occurs through Hebbian synaptic plasticity, an adaptive process based on the degree of correlated activity of pre- and post-synaptic neurons (Singer, 1987, e.g.). This amounts to a bottom-up, data-driven, and self-organizing process that captures the statistics of the environment where the animal lives. Simple computational models, in the form of feed-forward neural networks with Hebbian learning, develop receptive fields that resemble those found in the early stages of the mammalian visual cortex when exposed to input signals taken from uniform contrast distributions (Linsker, 1988) or from large sets of natural images (Hancock et al., 1992).

The ability to become sensitive only to a subset of all the possible features that could be derived from the huge array of optical stimulation allows brains with limited computational and storage resources to exploit at best their information processing capacity. Since the receptive fields are not entirely genetically determined, but depend on the visual stimulation to which the animal is exposed during its life, the developmental process plays an adaptive role for the animal. Although the details of the features accounted for by the receptive fields, and their role in perception, is still a topic of debate (Field, 1994, e.g.), it is generally accepted that animals and models of early-vision stages extract the dominant statistical features of their visual environment.

### 5.1.1 The Role of Active Behavior

Within that perspective, the visual system is a passive, albeit plastic, device shaped by the environment. However, all behavioral systems, biological and artificial, are capable of actively selecting –at least in part– their own sensory stimulation. They do so by moving in the environment and exploring the visual field by means of head and eye movements. The sequential and interactive process of selecting and analyzing parts of a visual scene is known as *active vision* (Bajcsy, 1985, 1988; Ballard, 1991). Active vision can simplify the computation involved in vision processing by

selecting only characteristics of the visual scene that are relevant for the task to be solved (Yarbus, 1967; Aloimonos, 1993), thus reducing the information load on the neural system. Indeed, a deeper understanding of the visual cortex can be gained by correlating neurophysiological measurements with the optical stimulation recorded from a camera placed on the head of a freely moving animal in its own natural environment (Betsch et al., 2004).

Nonetheless, the interaction between receptive field formation and active vision has been largely neglected in the biological and computational literature. This void most likely stems on the one hand from the difficulty of measuring neural dynamics in freely behaving animals and on the other hand from the fact that computational models are often designed to operate on predefined input pattern distributions (Parisi et al., 1990). Mobile robots are therefore an ideal tool to investigate adaptive processes that take place in a behavioral context because such robots can autonomously select the sensory stimulation by moving in the environment (Pfeifer and Scheier, 1999; Nolfi and Floreano, 2000).

In Chapter 2 we described the experiments on co-development of active vision and receptive fields within the same time scale using behavioral robotic systems equipped with a primitive retinal system and deliberately simple neural architectures. The synaptic strengths of the network were encoded in a binary string and evolved with a genetic algorithm while the robotic system was free to move in the environment. Evolved neural controllers exploited active vision and simple features to direct their gaze at invariant features of the environment and perform collision-free navigation.

In this chapter, we go one step further and investigate the *ontogenetic development*<sup>1</sup> of receptive fields in an evolutionary mobile robot with active vision. In contrast to our previous work where synaptic weights for both receptive field and behavior were genetically encoded and evolved on the same time scale, here the synaptic weights for receptive fields develop during the life of the individual. The synaptic weights of the neural network are genetically encoded and evolved, but the synaptic weights from visual photoreceptors to internal neurons (receptive fields) can also be modified by Hebbian synaptic plasticity while the robot moves in the environment. The Hebbian mechanism and architecture is one of those used in the literature for modeling receptive field formation (Hancock et al., 1992). In these experiments, behavioral abilities and receptive fields develop on two different temporal scales, phylogenetic and ontogenetic respectively. The evolutionary experiments are carried out in physics-based simulation and the evolved controllers are tested on the physical robot in an outdoor environment. We show that robots evolved in simulation with Hebbian visual plasticity display more robust adaptive behavior when transferred to real outdoor environments as compared to robots evolved without visual plasticity. We also show that the development of visual receptive fields is significantly and consistently affected by active vision as compared to the development of receptive fields with uniformly sampled images in the environment of the robot. Finally, we show that the interplay between active vision and receptive field

---

<sup>1</sup>In this thesis we will use the term *ontogenetic development* to indicate modifications during lifetime of an individual.

formation amounts to the selection and exploitation of a small and constant subset of visual features available to the robot.

## 5.2 Method

We use a Koala (K-Team S.A.) wheeled robot equipped with a pan/tilt camera (Sony EVI-D31) and infrared proximity sensors distributed around the body of the robot (Fig. 3.1). Infrared sensors are used only by the operating system to detect collisions and reposition the robot between trials; their activation values are not given to the neural controller. The robot has three wheels on each side, but only the central wheel (which is slightly lower) is motorized (the remaining two wheels rotate passively using a system of gears). The wheels have rotation encoders that are used for reconstructing the trajectories of the robot during behavioral analysis. The pan and tilt angles of the camera are independently controlled by two motors. The robot is equipped with an onboard computer (PC-104), hard disk, and operating system (Linux). The cables are used only for powering the robot through rotating contacts and to download/upload data through Ethernet at the beginning and end of an experiment.

The neural network has a feed-forward architecture with evolvable thresholds and discrete-time, fully-recurrent connections at the output layer (Fig. 5.1). A set of visual neurons, arranged on a five by five grid, with non-overlapping receptive fields receive information about the grey level of the corresponding pixels in the image provided by a camera on the robot. The receptive field of each neuron covers a square area of 48 by 48 pixels in the image. We can think of the total area spanned by all receptive fields (240 by 240 pixels) as the surface of an artificial retina. The activation of a visual neuron, scaled between 0 and 1, is given by the average grey level of all pixels spanned by its own receptive field or by the grey level of a single pixel located within the receptive field (Fig. 5.2). The choice between these two activation methods can be dynamically changed by one output neuron at each time step. Two proprioceptive neurons provide information about the measured horizontal (pan) and vertical (tilt) angles of the camera. These values are in the interval  $[-100, 100]$  and  $[-25, 25]$  degrees for pan and tilt, respectively. Each value is scaled in the interval  $[0, 1]$  so that activation 0.5 corresponds to 0 degrees (camera pointing forward parallel to the floor). Memory input units store the values of the output neurons at the previous sensory motor cycle step and send them back to the output units through a set of connections, which effectively act as recurrent connections among output units (Elman, 1990). The bias unit has a constant value of  $-1$  and its outgoing connections represent the adaptive thresholds of output neurons (Hertz et al., 1991, ch. 5).

Hidden and output neurons use the sigmoid activation function  $f(x) = 1/(1 + \exp(-x))$ , where  $x$  is the weighted sum of inputs to the neuron. Output neurons encode the behavior of the active vision system and of the robot at each sensory motor cycle. One neuron determines the filtering method used to activate visual




Figure 5.1: The architecture is composed of a grid of visual neurons with non-overlapping receptive fields whose activation is given by the grey level of the corresponding pixels in the image; a set of proprioceptive neurons that provide information about the movement of the camera with respect to the chassis of the robot; a set of output neurons that determine at each sensory motor cycle the filtering used by visual neurons, the new pan and tilt speeds of the camera, and the rotational speeds of the left and right wheels of the robot; a set of memory units whose outgoing connection strengths are equivalent to recurrent connections among output units; and a bias neuron whose outgoing connection weights represent the thresholds of the output neurons. Dashed connection lines can be modified by Hebbian learning in some experimental conditions.

neurons and two neurons control the movement of the camera, encoded as angles relative to the current position. The remaining two neurons encode the direction and rotational speeds of the left and right motored wheels of the robot. Activation values above –and below– 0.5 stand for forward –and backward– rotational speeds, respectively.

The connection strengths between visual neurons and hidden neurons are either genetically determined (“No learning” condition) or they can be modified by means of a Hebbian learning rule (“Learning” condition), which has been shown to produce connection strengths that approximate the eigenvectors corresponding to the principal eigenvalues of the correlation matrix of the input patterns (Sanger, 1989). In other words, this learning rule implements an approximate Principal Component Analysis of the input images (Jolliffe, 1986). The modification of connection strength  $\Delta w_{ij}$  depends solely on postsynaptic and presynaptic neuron activations  $y_i, x_j$ ,



Figure 5.2: Left: Snapshot from the robot camera. Center: Pixel average. Right: Pixel sample.

$$\Delta w_{i,j} = y_i \left( x_j - \sum_{k=1}^i w_{kj} y_k \right) \quad (5.1)$$

where  $k$  is a counter that points to postsynaptic neurons up to the neuron whose weights are being considered. The new connection strengths are given by  $w^{t+1} = w^t + \eta \Delta w_{ij}^t$ , where  $0 < \eta \leq 1$  is the learning rate, which in these experiments starts at 1.0 and is halved every 80 sensory motor cycles. This learning rule has been used in previous computational models of receptive field development (Hancock et al., 1992) and is intended to capture a system-level property of visual plasticity, not the precise way in which biological synaptic strengths are modified in the visual cortex. Among the several available models of synaptic plasticity (Hinton and Sejnowski, 1999, for a review), we opted for this one because it can be applied online while the robot moves in the environment and because it implements an approximate Principal Component Analysis, which is a widely used technique for image compression.

The robotic system and neural network are updated at discrete time intervals of 300 ms. At each time interval, the following steps are performed: 1. the activations of the visual and proprioceptive neurons are computed from the values provided by the robot, the values of the memory units are set to the values of the output units at the previous time step (or to zero if the individual starts its “life”); 2. the activations of the hidden units are computed and normalized (in experimental conditions where learning is enabled); 3. the activations of the output units are computed; 4. the wheels of the robot are set at the corresponding rotational speed for 300 ms while the camera is set to its new position; 5. in the experimental conditions where learning is enabled, the connection weights from visual neurons to hidden neurons are modified using the current neuron activation values. In step 2 the activation of each hidden unit is normalized to have the same magnitude in order to equalize the contributions of hidden units to activations of the output units. The normalized output value of the  $k$ th hidden neuron  $o'_k$  is computed by:

$$o'_k = o_k \times \frac{s_1}{s_k}, \quad k = 1, 2, 3, 4, 5 \quad (5.2)$$

where  $o_k$  and  $s_k$  denote the current output value of  $k$ th hidden neuron and the standard deviation of the stored output values by the current time step, respectively.



The strengths of the synaptic connections are encoded in a binary string that represents the genotype of the individual. In experimental conditions where learning is enabled, the connections between visual neurons and hidden neurons are *not* genetically encoded. In that case, they are randomly initialized at the beginning of the life of each individual (for more detail, see Appendix). Genetically encoded connection strengths can take values in the range  $[-4.0, 4.0]$  and are encoded on 5 bits. The genotype totals 325 bits ( $65 \text{ connections} \times 5 \text{ bits}$ ) in the learning condition and 950 bits ( $65 + 125 \text{ connections} \times 5 \text{ bits}$ ) in the no-learning condition. A population of 100 genomes is randomly initialized by the computer and evolved for 20 generations. Each genome is decoded into the corresponding neural network and tested for four trials during which its fitness is computed. The best 20% individuals (those with highest fitness values) are reproduced, while the remaining 80% are discarded, by making an equal number of copies so to create a new population of the same size. These new genomes are randomly paired, crossed over with probability 0.1 per pair, and mutated with probability 0.01 per bit. Crossover consists in swapping genetic material between two strings around a randomly chosen point. Mutation consists in toggling the value of a bit. Finally, a copy of the best genome of the previous generation is inserted in the new population at the place of a randomly chosen genome (elitism).

### 5.3 Experiments




Figure 5.3: Simulated and real environment.

Robots were evolved in physics-based simulation and tested both in simulated and real outdoor environments (Fig. 5.3). We simulated the robot and the environment using *Vortex* libraries, a commercially-available software package that models gravity, mass, friction, and collisions (<http://www.cm-labs.com>). Texture was generated from pictures taken in the outdoor environment. However, even simulations that include physical quantities and realistic textures as in this case, are quite different from the corresponding physical world. Consequently, robots designed or evolved in simulated environments do not operate properly when transferred to real

environments. The reason for evolving robots in simulation and testing the best controllers in the real environment was to assess the adaptive role of ontogenetic visual plasticity.

We have carried out two sets of evolutionary experiments (Fig. 5.4) to investigate whether ontogenetic development of receptive fields provides an adaptive advantage in new environmental conditions –namely when transferred from a simulated to a real outdoor environment– and to investigate the interactions between evolution and learning. The evolutionary experiments are carried out in simulations and the best evolved individuals are tested in the real environment. In the first condition (“No learning”), which serves as a control condition, all synaptic connections are genetically encoded and evolved without learning. In the second condition (“Learning”), learning is enabled for the connections from visual neurons to hidden neurons which are not genetically encoded, but initialized to small random values. Connection strengths developed during learning are not transmitted to offspring.

The fitness function selected robots for their ability to move straight forward as long as possible for the duration of the life of the individual. This is quantified by measuring the amount of forward rotation of the two motorized wheels of the robot. Each individual is decoded and tested for four trials, each trial lasting 400 sensory motor cycles<sup>2</sup>. A trial can be truncated earlier if the operating system detects an imminent collision with infrared distance sensors. The fitness function  $F(S_{left}, S_{right})$  is a function of the measured speeds of the left  $S_{left}$  and right  $S_{right}$  wheels:

$$F(S_{left}, S_{right}) = \frac{1}{E \times T} \sum_{e=0}^E \sum_{t=0}^{T'} \underbrace{(S_{left}^t + S_{right}^t)}_{f(S_{left}, S_{right}, t)} \times \left( 1 - \sqrt{\frac{|S_{left}^t - S_{right}^t|}{2 \times S_{max}}} \right) \quad (5.3)$$

where  $S_{left}$  and  $S_{right}$  are in the range  $[-8, 8]$  cm/sec and  $f(S_{left}, S_{right}, t) = 0$  if  $S_{left}$  or  $S_{right}$  is smaller than 0 (backward rotation);  $E$  is the number of trials (four in this chapter),  $T$  is the maximum number of sensory motor cycles per trial (400 in these experiments),  $T'$  is the observed number of sensory motor cycles (for example, 34 for a robot whose trial is truncated after 34 steps to prevent collision with a wall). At the beginning of each trial the robot is relocated in the real outdoor environment at a random position and orientation by means of a motor procedure during which the robot moves forward and turns in a random direction for 10 seconds. In simulation, position and orientation are instantly randomized. In the learning condition, the synaptic weight values are re-initialized to random values at the beginning of each trial.

The results of the two experimental conditions, without learning and with learning, do not show significant differences and reach comparable performance levels (Fig. 5.4). In both cases, evolved individuals perform large, collision-free paths around the environment for the entire duration of a trial.

---

<sup>2</sup>Preliminary experiments reported in Floreano et al. (2005) show that less than 300 updates are necessary to stabilize the plastic weights from visual to hidden neurons.




Figure 5.4: Population average (thin line) and best fitness (thick line) during evolution in physics-based simulations for “No learning” and “Learning” conditions. Each data point is the average of three evolutionary runs with different initializations of the population.

One way to observe the effects of receptive field development consists of freezing the receptive fields at different periods of life and measuring the corresponding performance of individuals in the environment. We have carried out a series of such tests in the simulated environment for all best evolved individuals with learning and compared their performance to that of best evolved individuals without learning (Fig. 5.5). The initial random synaptic strengths (RandRF) of visual neurons do not provide sufficient information to let the robot move around properly, but the final connection strengths (FinRF) do provide the structure necessary for the rest of the network to achieve a performance comparable to that of best evolved individuals without learning. The absence of difference between performance measured with well-formed receptive field and performance measured during the entire learning phase (400 sensory motor cycles) indicates that learning does not produce a significant cost for the individuals – see (Mayley, 1996) for a discussion of the cost of learning in evolutionary individuals. The development of sufficiently good receptive fields is very fast probably because the relative uniformity and predictability of the simulated environment can be quickly captured by the Hebbian learning mechanism.

The advantage of receptive field development becomes evident when the best evolved robots from simulation are tested in the real outdoor environment (Fig. 5.6). In that case, the performance of best individuals evolved without learning drops significantly (horizontal line in graph) whereas the performance of “Learning” individuals with receptive fields that developed in the new environment (FinRF) is almost the same as that measured in simulation tests (Fig. 5.5, FinRF). The large difference between performance with initial random receptive fields (RandRF) and final receptive fields (FinRF), as well as the significant performance cost incurred during learning (Learning), is an indication of the discrepancies and novelties in the




Figure 5.5: Performance of best individuals evolved in the “Learning” condition (three for each evolutionary condition) and tested each five times in the *simulated environment* from different initial positions and orientations. The horizontal line represents the test performance of the best evolved individuals evolved in the “No learning” condition. Individuals are tested with random and fixed receptive fields (RandRF), with final and fixed receptive fields (FinRF), and with learning enabled. The duration of the test is the same as that used during evolution.

real environment that the learning mechanism had to integrate.

## 5.4 Behavioral analysis

Figure 5.7 plots the trajectories in the outdoor environment of the typical best individuals evolved in the two evolutionary conditions. In order to discount trajectory adjustments due to online development of the receptive fields in the “Learning” condition, the synaptic weight values have been fixed while the trajectory was recorded. The individual from the “No learning” condition cannot always avoid collision with the environment. Instead, the trajectories of the individual from the “Learning” condition change during its life. With initial random receptive fields, it performs straight trajectories that, although taking it closer to obstacles, increase temporarily the fitness score during evolution. When the receptive fields are not yet developed though, the robot cannot avoid a collision with the wall. However, when the receptive fields are in their final form, they support collision-free navigation in the outdoor environment.

In both experimental conditions, best evolved individuals tilt the camera toward the floor where they can see the edge between the ground and the walls of the building. Individuals evolved in the “Learning” condition display higher use of pan movements of the camera than individuals evolved in the “No learning” condition (Fig. 5.8). A possible explanation for this difference is that “Learning” individuals become sensitive to a smaller number of environmental features than “No learning”




Figure 5.6: Performance of the best individuals evolved in the simulated environment (three for each evolutionary condition) tested each five times in the *real environment* from different initial positions and orientations. See legend of figure 5.5 for explanation at column labels.

individuals, and they use the camera more actively in order to track those features while they navigate in the environment. We will come back to this hypothesis in the next section. Also, best individuals evolved in the two experimental conditions converge on the use of pixel average to activate input visual neurons.

## 5.5 Receptive Field Development

In order to understand the effect of behavior on receptive field formation, we compared the final receptive fields of evolved robots tested in the outdoor environment with the receptive fields obtained by training visual neurons on uniformly sampled images that the robot could possibly gather in the outdoor environment. Uniformly sampled images were collected by placing the robot in the center of the outdoor environment and taking snapshots with the camera pointed at different directions, tilt, and time of the day (Fig. 5.9). We chose 36 directions by letting the robot perform a full rotation in place, three tilt values ( $-25$ ,  $0$  and  $25$  degrees), and three times of the day to account for different lighting conditions and shadows that varied also during the behavioral tests with evolved robots. The complete set was composed of 324 images. Images were filtered using pixel averaging (the strategy chosen by all best evolved individuals) and fed as input to the visual part of the neural architecture in random order. The synaptic weight values were initialized to small random values within the same range used during evolution and modified after each image presentation using the learning rule described in equation 5.1. Synaptic weight values stabilized approximately after 200 applications of the learning rule, which is less than the trial time of evolutionary robots.

Figure 5.10 shows the final receptive fields of the network trained on uniformly

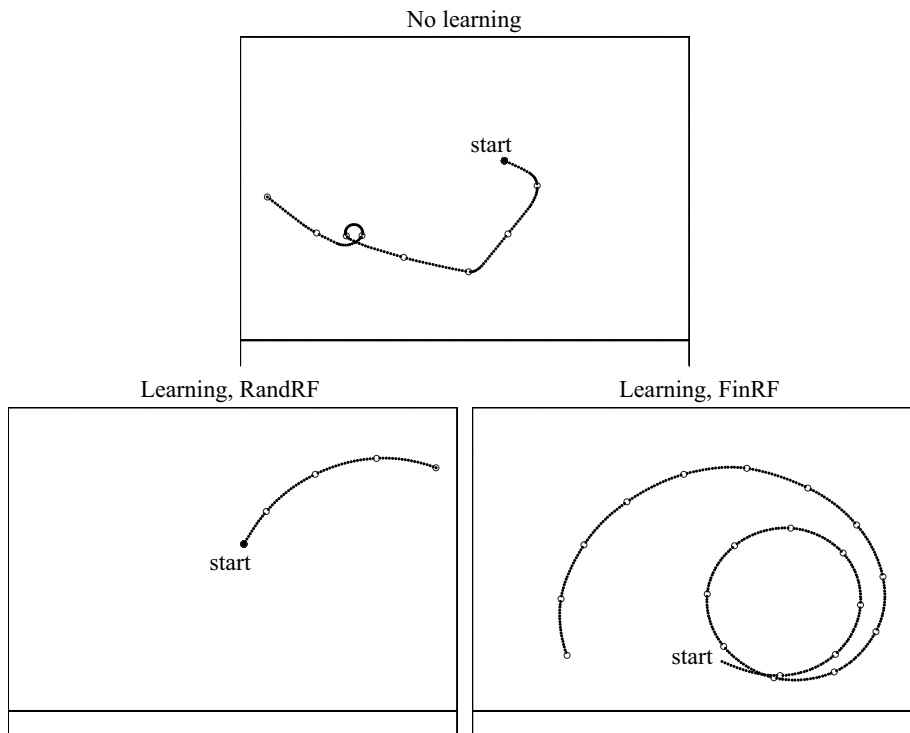


Figure 5.7: Trajectories of robots in real outdoor environment reconstructed from wheel odometric information. A dot is plotted every 20 sensory motor cycles. The maximum duration is 400 sensory motor cycles. Top: “No learning”. Bottom: “Learning”, with fixed random receptive fields and final receptive fields.

sampled images of the network of the best individual evolved in “Learning” condition and of the network of the best individual evolved in “No learning” condition for the sake of completeness. As we explained in section 5.2 above, the synaptic weight values developed with the Hebbian learning rule used here approximate the principal components of the correlation matrix of the images. The leftmost receptive field in the figure corresponds to the first principal component, that is the component that explains most of the image variance which represents the most important feature in the visual environment. The second-from-left receptive field corresponds to the second principal component, that is the component that explains most of the residual image variance, which represents the second most important feature in the visual environment, and so forth.

The most significant difference emerging from this comparison is that the dominant receptive field for evolved learning individuals is an horizontal edge on the lower portion of the visual field whereas the dominant receptive field for the network trained on uniformly sampled images resembles a light central blob against a dark background combined with a weak edge in the upper portion of the visual field. The horizontal edge appears only as the second most important receptive field for the network trained on uniformly sampled images. This difference is sup-




Figure 5.8: Pan angles of the camera during the trajectory in the outdoor environment for the “No learning” and “Learning” individual with fully formed receptive fields (FinRF). Angle values are mapped from 0.5 to 1 for rotations to the left of the robot and from 0.5 to 0 for rotations to the right of the robot. The trial of the “No learning” individual is terminated after 155 sensory motor cycles (see also Fig. 5.7).



Figure 5.9: Examples of images taken at different times of the day (9am, 3pm, and 9pm) from the same location.

ported by the fact that the subspace of the distribution of snapshots collected by evolved learning robots with well formed receptive fields is significantly different from that of the distribution of uniformly sampled snapshots (Fig. 5.11). The comparisons become more difficult for the remaining three receptive fields, but this is not surprising because these receptive fields capture any small residual variance of the image distribution that may have been caused by different starting positions and/or trajectory differences of evolved individuals. The receptive fields of the “No learning” individual cannot be easily interpreted as in the other two cases because our genetic representation and evolutionary method is not forced to represent visual information in any predefined order. We can however speculate that these receptive fields may contain a linearly scrambled version of the first principal component observed in the “Learning condition” because both evolved robots perform similar trajectories and point the camera toward the edge between the ground and the wall.

The behavioral and receptive field analysis indicate that evolved learning robots






Figure 5.10: Visual Receptive fields. Each 5 by 5 matrix displays the incoming synaptic strengths for one of the five visual neurons. Gray levels correspond to the sign and strength of the corresponding synapse (white = maximum value; black = minimum value). Leftmost receptive field approximates first principal component of the visual environment.

pay attention to specific visual features of the environment, notably the edge between the ground and the walls, in order to navigate in the environment. This phenomenon can be further investigated by measuring the diversity of snapshots captured by evolved robots as compared to uniformly sampled snapshots. For sake of clarity, we project the 25 dimensional space of the input images (given by the five by five visual neurons) onto the 3 dimensional subspace of the first three principal components of the image distribution (here the principal component analysis is performed using the exact mathematical procedure (Jolliffe, 1986)). Each snapshot is represented by a dot in the three dimensional space. Figure 5.11 shows that the distribution of snapshots collected by evolved learning robots with well formed receptive fields is significantly more compact and is composed of a smaller number of different snapshots than the distribution of uniformly sampled snapshots. These data indicate that such robots self-select through their behavior (body movements and camera movements) a significantly smaller and consistent subset of visual features that are useful for performing their survival task of collision-free straight navigation. This holds also for individuals of the “No learning” condition evolved and tested in simulated environments (data not shown), but not when such individuals are tested in the real environment, as shown in Fig. 5.11.

## 5.6 Discussion

The results described in this chapter agree with those obtained in previous work (Floreano et al., 2004) where synaptic weights for both receptive field and behavior were genetically encoded and evolved on the same time scale to the extent that






Figure 5.11: Distribution of snapshots projected onto the three dimensional space of the first three principal components (computed using the exact mathematical procedure). In the last two conditions, snapshots have been recorded during trajectories of evolved individuals. The number of plots in each graph is 324, 155, 400, respectively.

receptive fields become sensitive to linearly separable features of the environment, such as oriented edges, that can be actively pursued in order to support the navigation task. But, in addition to that, the results described here show that evolved active behavior significantly affects the development of visual receptive fields during the life of our robotic individual as compared to receptive fields obtained from uniformly sampled images. These experiments suggest that the statistics of the environment may not be sufficient to explain the visual receptive fields (and consequently the perception) of behavioral systems even without invoking top-down mechanisms that may guide perception. Active behavior, here instantiated as the coordinated movement of the robotic body and camera, can affect the formation of receptive fields by actively selecting a subset of the visual world during the early plasticity period. A similar pattern of interaction between perception and behavior has been recently documented using a different neural architecture embedded in a mobile robot (Verschure et al., 2003).

The type and organization of receptive fields developed in these experiments are

most likely specific to the choice of learning rule and neural architecture. A different learning rule and architecture may have resulted in different receptive fields, but not necessarily into different patterns of interactions between behavior and development as long as the learning system belongs to the class of data-driven, Hebbian instances of unsupervised learning (Hinton and Sejnowski, 1999) and is computationally sufficient to support evolutionary maximization of the fitness criterion used in the experiments.

The evolutionary experiments have been carried out in simulations and the best evolved controllers have been tested in the real world. As mentioned earlier in the chapter, even the most recent type of physics-based and visually realistic simulation methods include approximations that adaptive robots may exploit to develop efficient strategies. For that reason, control strategies evolved in simulated environments most often do not transfer well to real environments. Still, simulations offer several advantages, such as that of “being faster” than evolving robots in real environments and of avoiding the shortcomings of mechanical failures. Moving back and forth from simulation to reality could therefore open several new possibilities for evolutionary robots if one could ensure a way to guarantee a smooth transfer. The results described in this chapter indicate that evolution of adaptive –instead of fixed– control systems is a promising solution to this issue because the differences between real and simulated environments can be taken into account by ontogenetic adaptation of evolved individuals. The successful transfer of evolved adaptive individuals from simulation to reality described here provide further support for the benefits of incorporating unsupervised learning in evolutionary control systems that are transferred from simulated to real environments (Urzelai and Floreano, 2001). This holds as long as the difference between simulated and real environments do not violate the consistency of the world (however, see (Di Paolo, 2000) for a case of ontogenetic adaptation to inverse worlds) or introduce completely new situations that lifelong adaptation mechanisms alone cannot cope with.

As in our previous work on evolution and learning, we have not genetically encoded the initial synaptic strengths of connections that are subject to learning. This prevents the indirect assimilation of learned features into the genetic code, whereby genetically encoded strengths tend to converge toward values that are produced by the learning mechanism, thus gradually reducing the space for adaptation. This indirect genetic assimilation, a facet of the so-called Baldwin effect (Baldwin, 1896), can happen when the environment does not change or changes significantly slower than generational time (Mayley, 1996; Belew and Mitchell, 1996; Nolfi and Floreano, 1999, for a review).

## 5.7 Conclusion

The experimental results described in this chapter indicate that the interaction between learning and behavior within an evolutionary context brings a number of synergetic phenomena: a) behavior affects learning by selecting a subset of learning

experiences that are functional to the survival task; b) learning affects behavior by generating selection pressure for actions that actively search for situations that are learned; c) learning contributes to the adaptive power of evolution (as long as the parameters subject to learning are not also genetically encoded) by coping with change that occurs faster than evolutionary time, as is the case of transfer from simulation to reality. These results are promising for scalability and potential applications of evolutionary robotics in real-world situations where robots cannot possibly be evolved in real-time in the real world, but may have to evolve at least partly in simulations. They are also an indication that complex behavior can be generated by relatively simple control architectures with active behavior and local unsupervised learning that can be implemented in low-power, low-cost micro-controllers. The significant role of behavior in receptive field formation during learning is being increasingly recognized in the neuroscience community (Polley et al., 1999, 2004) too. Within this context, Evolutionary Robotics is very well positioned to address specific questions raised by theories and models of learning in a behavioral context because it does not require hardwiring of behavior, but only set up of the environment and of the performance criterion where evolution will operate.

The experimental settings, neural architecture, genetic encoding, and learning mechanisms used in these experiments have been kept deliberately simple to facilitate the analysis and comparison with previous results, but they also represent a limitation for what concerns the generalization and scalability of the system. For example, the learning component of the network (architecture, rule, neuron model) is such that only linearly separable, and static, features of the image can be detected. It is thus unlikely that this specific system will scale up to more complex situations where time-dependent and non-linear combinations of features are functionally related to the task at hand. Similarly, in the context of an open-ended evolutionary experiments where robots would evolve in a dynamic and increasingly more complex environment, the one-to-one encoding of synaptic parameters used here is most likely not suitable. Our current work is aimed at both these aspects as part of our effort to pave the ground for a methodology where complex robotic artifacts continuously evolve without human intervention and pre-defined performance criteria.

## Appendix

The connections between visual neurons and hidden neurons are randomly initialized in the range  $[-\sqrt{3/25}, \sqrt{3/25}]$  at the beginning of the life of each individual. We derive the value  $\sqrt{3/25}$  from the fact that the number of pixels is 25 and that synaptic weights are randomly initialized with uniform probability distribution.

We have  $n$  synapses that we want to initialize randomly (i.e., with uniform probability distribution) in the range  $[-A, A]$ . We would like to choose the value of  $A$  in order to meet the requirement  $\|w\| = 1$ , but of course we can do it only probabilistically. Hence, we ask the expected value of  $\|w\|$  to meet the requirement, that is,  $E[\|w\|^2] = 1$ .

Since the probability distribution is uniform in the range  $[-A, A]$ , the value of the probability density function is a constant  $1/(2A)$  in the interval, and zero outside. It follows that

$$E[w_1^2 + w_2^2 + \dots + w_n^2] = n \int_{-A}^A \frac{w^2}{2A} dw. \quad (5.4)$$

By evaluating the definite integral we obtain  $E[w_1^2 + w_2^2 + \dots + w_n^2] = (n/3) \times A^2$ . The condition  $E[|w|^2] = 1$  can thus be written as  $(n/3) \times A^2 = 1$ , from which the result  $A = \sqrt{3/n}$  follows (in our case  $n = 25$ ).

---

# 6

## Computational Neuroethology

---

*Our methodology can also be used to investigate open questions in neuroscience and cognitive science (Harvey et al., 2005) because it offers the vantage point of a behavioral system that interacts with its environment (Cliff, 1991). Although the results should be carefully considered when drawing analogies with biological organisms, our methodology can generate and test hypotheses that could be further investigated with mainstream neuroscience methods. This chapter is based on Suzuki et al. (2005a,b).*

**Abstract.** Inspired by the pioneering work by Held and Hein (1963) on the development of kitten visuo-motor systems, we explore the role of active body movement in the developmental process of the visual system by using robots. The receptive fields in an evolved mobile robot are developed during active or passive movement with a Hebbian learning rule. In accordance to experimental observations in kittens, we show that the receptive fields and behavior of the robot developed under active condition significantly differ from those developed under passive condition. A possible explanation of this difference is derived by correlating receptive field formation and behavioral performance in the two conditions.

### 6.1 Introduction

Perception in the natural case is a process dependent as much on the sensory systems available to the organism as on its motor activity. This is not only because the stream of sensory inputs is directly affected by movements and adjustments (such as scanning, focusing, orienting, positioning) but often also because perceptual invariants are built upon correlations between sensory and motor dynamics. Evidence of this double dependence is provided by classical experiments on visual adaptation to distortion of the visual field (Stratton, 1896, 1897; Kohler, 1964; Taylor, 1962), where perceptual adaptation only takes place after many days of the subject actively engaging in different behaviors. The adaptations thereby achieved are typically non-transferable to other behaviors—which themselves must be enacted in order

to adapt—and they do not take place if the subject is passive or moved externally.

A similar dependence is found in the process of perceptual development. For instance, Held and Hein (1963) have shown that normal visual development depends not only on movement of the body relative to the environment, but also on self-actuated movement. The authors performed an experiment (Fig. 6.1) in which the gross movements of a kitten moving almost freely (*active kitten*) were transmitted to a second kitten that was carried in a gondola (*passive kitten*). Consequently, they received identical visual stimulation, but only one of them received that stimulation as a result of self-movement. Importantly, only the active kitten developed normal behavior in several visually guided tasks, such as paw extension on approaching horizontal surface from above and blinking at object put in front of its eyes, while the passive one failed. The authors concluded that visual stimulation correlated with self-actuated movement was necessary for the development of the visual control of behavior. However, it is still not clear how the active body movement of the kitten enabled it to develop such visually guided behaviors.




Figure 6.1: The original apparatus in Held and Hein (1963), where the gross movements of a kitten moving almost freely were transmitted to a second kitten that was carried in a gondola. Both kittens were allowed to move their head. They received essentially the same visual stimulation because of the unvarying pattern on the walls and the center post of the apparatus. Reproduced, with permission, from (Held, 1965).

A host of experiments has shown that the characteristics of biological and artificial adaptive systems strongly depend on the type of inputs they receive during the developmental process (e.g. Blakemore and Cooper, 1970). Additionally, *active vision*, i.e., the sequential and interactive process of selecting and analyzing parts of a visual scene, selects the subset and sequence of images that the visual system perceives (Bajcsy, 1988; Aloimonos et al., 1987; Aloimonos, 1990; Ballard, 1991).

Thus, it is tempting to speculate that the way of scanning of the visual scene may significantly alter the development of the visual system of the animal.

Indeed, recent experimental results suggest that free exploration of the visual field may impact the development of the visual system. Betsch et al. (2004) showed that the exploration strategy and the difference of vantage point of animals significantly altered the statistics of natural scenes. This is fully consistent with the lesson from the studies of visuo-tactile interfaces for blind people (see Bach-y-Rita and Kerckel, 2003, for a review). The authors showed that human subjects could develop visually-guided behavior capabilities only if they were allowed to manipulate the camera by themselves, and that if someone moves the camera for them, they just sense a noisy stream of inputs, but no “visual” perception. Furthermore, it was demonstrated that active exploration in a naturalistic environment had a powerful impact on the expression of plasticity in whisker-deprived adult rats (Polley et al., 1999, 2004). These results suggest that the development of the visual system largely depends not only on the characteristics of the visual field but also on the behavior of the perceivers.

Recent advances in computational neuroscience have shown that relatively simple models of developmental visual systems are capable of developing qualitatively similar properties to those found in the early stages of visual processing in cats and monkeys (Hancock et al., 1992; Field, 1994; Olshausen and Field, 1996; Rao and Ballard, 1999). However, those models often use images from publicly available databases or photographs taken in a natural environment as visual stimuli, and do not allow the system to freely interact with the environment and choose those sensory events.

Previously we have investigated the co-development of active vision and receptive fields within the same time scale using behavioral robotic systems (Floreano et al., 2004). We have shown that co-evolved feature selection and active vision can address a variety of visual tasks that range from complex shape discrimination to navigation in complex environments by means of very simple mechanisms. However, the system investigated in those experiments could not change during the life of the “organism”.

In this chapter, we go one step further and explore the role of active body movement in the formation of the visual system by studying the development of visual receptive fields and behavior of robots under active and passive movement conditions. The receptive fields in an evolved mobile robot are developed during active and passive movement with a Hebbian learning rule. We show that the receptive fields and behavior of robots developed under active condition significantly differ from those developed under passive condition. Our analyses show that the coherence of receptive fields developed in active condition plays an important role in the performance of the robot.

## 6.2 Experiments

Methods used in this experiment, such as the robot, its environment, the fitness functions, the genetic algorithm, the neural architecture with visual plasticity, are exactly same as described in the previous chapter.

As I explained in the previous chapter, at the beginning of each trial the position and orientation of the robot are instantly randomized and the synaptic weight values are re-initialized to random values. We performed these replications of the evolutionary run starting with different genetic populations. In all cases the fitness reached stable values in less than 20 generations (Figure 6.2) which corresponded to collision-free trajectories. Notice that the fitness can never be one because the robot must rotate in order to avoid walls.




Figure 6.2: Evolution with synaptic plasticity. Left: Population average (thin line) and best fitness (thick line) during evolution in physics-based simulations. Each data point is the average of three evolutionary runs with different initializations of the population. Vertical lines show the standard error. Right: An example of trajectory of the best individual in the last generation while synaptic plasticity is active. A dot is plotted every 20 sensory motor cycles.

### 6.2.1 Visual development during active or passive movements

After evolution the receptive field formation of the best evolved individuals were studied in two behavioral conditions: one where the evolved robot was free to control the movements of its wheeled platform and of the camera, and another where the movement of the wheeled platform was constrained (but not that of the camera). First, we let the evolved robot move freely while the receptive fields were developed (we label the resulting receptive fields in active movement condition  $RF_a$ ). In the



second condition, the same evolved robot was constrained to move according to four pairs  $(S_{left}, S_{right})$  of wheel speeds while the receptive fields were developed.

- Behavior 1:  $(S_{left}, S_{right}) = (S_{max}, -S_{max})$
- Behavior 2:  $(S_{left}, S_{right}) = (0.4 \times S_{max}, -0.4 \times S_{max})$
- Behavior 3:  $(S_{left}, S_{right}) = (S_{max}, 0)$
- Behavior 4:  $(S_{left}, S_{right}) = (S_{max}, 0.2 \times S_{max})$

where  $S_{max}$  denotes the maximum speed of the wheels (8 cm/s). We call these four behaviors “passive” to highlight that the evolved neural network can not control the wheels<sup>1</sup> and label the resulting receptive fields  $RF_{p1}$ ,  $RF_{p2}$ ,  $RF_{p3}$ , and  $RF_{p4}$ . Behavior 1 and 2 correspond to ‘turning-on-the-spot’ while behavior 3 and 4 produce small circular behaviors with different radii. The camera could be freely controlled by the evolved neural controller in all four passive conditions.

In both conditions, the robot was located randomly at the beginning of each test and allowed to move for 400 sensory motor cycles while the visual receptive fields were developed from initial random weights. The test was repeated ten times for each condition starting from different random weights and locations. Figure 6.3 shows the receptive fields resulting from active and passive behaviors of one trial. We could not measure any statistical difference or distance between the five sets of receptive fields.

After development in the active and four passive conditions the corresponding receptive fields  $RF_a$ ,  $RF_{p1}$ ,  $RF_{p2}$ ,  $RF_{p3}$  and  $RF_{p4}$  were fixed and the performance of the robot was evaluated while the robot moved freely for maximum 400 sensory motor cycles. Figure 6.4 shows that the performance obtained with receptive fields developed during active behavior ( $RF_a$ ) is significantly better than those with receptive fields developed during passive behavior ( $RF_{p1-4}$ ). A typical trajectory of the robot with fixed  $RF_a$  and that of the robot with fixed  $RF_{p2}$  are shown in Fig. 6.5. The other trajectories corresponding to the receptive fields developed under the remaining three passive conditions,  $RF_{p1}$ ,  $RF_{p3}$  and  $RF_{p4}$ , are similar to that of  $RF_{p2}$ .

## 6.3 Analysis

### 6.3.1 Lesion Studies

This section describes a variety of behavioral analyses to understand why the performance of  $RF_a$  differs from that of  $RF_{p1-4}$ . First, we investigated the role of  $RF_a$

---

<sup>1</sup>Passive behavior was accomplished by simply neglecting the output values  $(W_{left}, W_{right})$  of the neural controller and reading one of the four pairs  $(S_{left}, S_{right})$  of wheel speeds instead. However note that the output values  $(W_{left}, W_{right})$  were not overwritten by  $(S_{left}, S_{right})$  but copied to the memory units so that passive behavior of the robot would be analogous to that of the kitten carried in a gondola in that they could move their wheels or legs freely without any contribution to the actual movement of their whole bodies.




Figure 6.3: Receptive fields of five hidden neurons developed in active and passive conditions. Small shaded squares represent the connection strengths from visual neurons, scaled so to fill the gray scale from black (minimum value) to white (maximum value). The leftmost receptive field in each row corresponds to the first principal component of the visual input experienced by the robot. A receptive field is the pattern of synaptic strengths to a neuron, here plotted as a gray level matrix.




Figure 6.4: Performances of the robot with receptive fields developed in active ( $RF_a$ ) and passive ( $RF_{p1-4}$ ) conditions. The fitness values are averaged over ten tests. Vertical lines show the standard error.

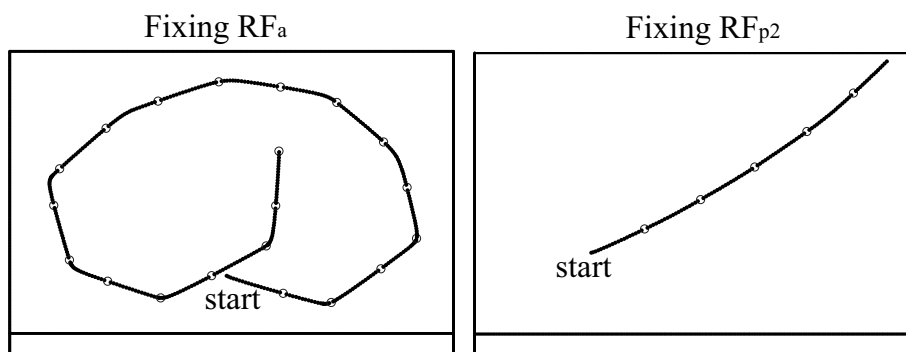


Figure 6.5: Trajectory of the robot with fixed receptive fields after development. Left: With  $RF_a$ , the receptive fields developed during active behavior. Right: With  $RF_{p2}$ , the receptive fields developed during passive behavior 2. A dot is plotted every 20 sensory motor cycles. The trajectories corresponding to receptive fields developed under the remaining three passive conditions,  $RF_{p1}$ ,  $RF_{p3}$  and  $RF_{p4}$ , are similar to that of  $RF_{p2}$ .

by lesioning hidden units one at a time and testing the lesioned controller in the environment ten times for a duration of 400 sensory motor cycles each. Lesion was performed by clamping the activation value of the neuron to a constant value of 0.5 (approximately equal to the average activation level). During these tests the receptive fields were not allowed to change.

Figure 6.6 shows that lesions of the first and second units (units 1-2) affects performance most significantly in the case of  $RF_a$ . This finding was validated by another set of tests where simultaneous lesion of the first two units significantly reduced the robot’s performance, but simultaneous lesion of the last three units did not.<sup>2</sup>

Then, we noted that the receptive fields of the first two units developed in passive condition 2 ( $RF_2$ ) were similar to those developed in active condition, but that the performance of that neural controller was one of the worst observed.

A possible explanation of the performance difference between neural controllers developed in active and passive conditions is that the neurons that capture statistically less dominant features (neurons 3, 4, and 5) may develop sensitivity to “interfering” features in the passive conditions. To test the validity of this hypothesis, we lesioned simultaneously neurons 3, 4, and 5 in the passive conditions and tested the performance of the robot. Figure 6.7 shows that the performances of the robot were, as expected, improved by lesioning units 3, 4, and 5. These neurons may interfere with the first two neurons by capturing information that “distracts” or con-

<sup>2</sup>These results can not be simply explained by the larger variance attributed to the first two units by the learning algorithm because, as described in section 5.2, the magnitudes of the output of the five hidden units are normalized so that each hidden unit can equally contribute to firing of the output units.




Figure 6.6: Performance with lesioned  $RF_a$ . ‘L1’ denotes the performance of the robot when the first hidden neuron was lesioned; ‘L345’ when units 3, 4, and 5 were lesioned simultaneously. Fitness values are averaged over ten tests. This figure shows that the first and second hidden neurons play an important role for the performance of the robot. Horizontal dashed line represents the fitness value of the robot with intact receptive fields.

trasts the information provided by the first two neurons, which encode statistically dominant features of the environment.

Furthermore, if the coherence of the receptive fields is at least as important as the actual information encoded, then substituting receptive field developed in passive condition with those developed in active condition should not restore the performance of the robot fully. In a first set of tests, the receptive fields of units 3, 4, and 5 of neural controller developed in passive conditions were substituted by those developed in active condition (Fig. 6.8, gray bars). The performances of the robot with modified  $RF_{p1-4}$  were not consistently better as when lesioning units 3, 4, and 5. A notable exception is the case of  $RF_{p2}$ . The performance is very close to that with  $RF_a$  because the receptive fields of the first two units are very similar. In the second set of tests, the substitution concerned the receptive fields of the first two units (Fig. 6.8, white bars). Also in this case, the performance of the robot was not as good as that obtained by the neural controller developed in active condition.

The last analysis concerns how the evolved learning robots sample the visual input in active and passive learning conditions. The distribution of the entire set of snapshots (25 pixels) was projected onto the three dimensional space of the first principal components. Figure 6.9 shows that the snapshots taken in the active learning condition are distributed in more structured manner than those taken in the passive learning condition because the constraints on body movement did not allow the robot to freely sample the images. Indeed, the distribution of snapshots




Figure 6.7: Performance with lesioned receptive fields. Dark gray bar shows the performance with five intact receptive fields, whereas light gray bar with three lesioned neurons. The fitness values are averaged over ten tests. The performances with all of  $RF_{p1-4}$  were improved by lesioning units 3, 4, and 5.




Figure 6.8: Performance in the substitution test. Fitness value of the robot was computed when units 3, 4, and 5 of  $RF_{p1-4}$  were substituted by those of  $RF_a$  (gray bars) and when units 1 and 2 of  $RF_{p1-4}$  were substituted by those of  $RF_a$  (white bars). Black bar shows the performance with five intact receptive fields for the sake of comparison. Horizontal dashed line represents the fitness value of the robot with intact  $RF_a$ . The fitness values are averaged over ten tests.

taken in the passive condition is close to that of uniformly sampled images, that was previously shown in Floreano et al. (2005). The distributions obtained in the other three passive conditions,  $RF_{p1}$ ,  $RF_{p3}$ , and  $RF_{p4}$ , are similar to that of  $RF_{p2}$ .




Figure 6.9: Distribution of snapshots taken during active and passive learning conditions. These plots are projected onto the three dimensional space of the first three principal components. The number of plots in each figure is 400 (=maximum sensory motor cycles per trial). The distributions obtained in the remaining three passive conditions,  $RF_{p1}$ ,  $RF_{p3}$  and  $RF_{p4}$ , are similar to that of  $RF_{p2}$ .

## 6.4 Discussion

Using an experimental setup similar to that used for kittens (Held and Hein, 1963), we have explored the correlation between receptive field formation and behavior in two conditions. The present results suggest that constraints on body movement disturb the development of “healthy” visual receptive fields. Although we can not see any significant difference in the level of receptive fields themselves, they caused a significant difference in behavior. Furthermore we have shown that the coherence of receptive fields developed in active condition plays an important role in the good performance of the robot.

Although the arrangement and relative importance of the receptive fields described depend on the specific learning rule used in these experiments, the results suggest that during passive movement the developing system incorporates sensory stimulation that is not functional for normal behavior. In other words, freely behaving systems select a subset of stimuli that coherently support the generation of behavior itself.

One would say that we could conceivably have evolved a robot that would also produce correct behavior under conditions p1-4 if these were presented during evolution, and thus we have only demonstrated that the robot is not good at doing something which it was not evolved to do. This criticism would miss the point of the

study which is to demonstrate how motor activity affects development. Evolution is free to pick up a convenient pattern of motor activity that facilitates development. If it were easier to ignore motor activity and perform some sort of non-historical image analysis on every visual input so as to extract the necessary information for navigation, evolution would have very likely found that solution or something close, but that is not the case.

It would be good to come back here to the bigger picture that was set at the start of the chapter: the point that not only visually-guided behavior depends non-trivially on motor activity (active vision) but that its development relies on it as well. This thesis has a stronger and a weaker version. The weaker says: to the extent that sensory input is dependent on movement, and the development of receptive fields dependent on sensory input, then this development also depends on movement. If you change the pattern of allowed movements, you will affect development. This is what the present experiments have shown in this chapter.

However there is a stronger version that includes the former but adds the following: there is also a direct dependence of development on how movement is registered by the system, i.e., on proprioceptive activity, or efferent copies or similar mechanisms for distinguishing self-generated movement from non-self-generated movement. For this stronger version, even if one manages to replicate the precise sensory input (thus removing this indirect dependence on movement), development will also be impaired, because it lacks another fundamental component, the information of how visual input and movement (through proprioception) are coordinated.

This stronger version is what the original kittens' experiment demonstrated in Held and Hein (1963). If we accept that the device effectively "copies" the active kitten's sensory input into the passive kitten's, then the latter's lack of visual development can only be attributed to its lack of the temporal correlation—and the resulting association—between a measure of actual body movement and the corresponding proprioceptive input (barring other factors such as stress, etc). This situation is not quite the same as the one currently reproduced with the robot as there is only camera proprioception. To support this stronger version of the argument, one should carry out further experiments with an extended sensory system measuring actual body movement by means of accelerometers or gyroscope.

## 6.5 Conclusion

We carried out a set of robotic experiments to study the contribution of active body movement to the development of visual system in a mobile robot. Although the present experimental setup is not exactly same as the one shown in Held and Hein (1963), the essence of the original experiment was extracted and reproduced in an artificial manner by means of physics-based simulation. A Hebbian learning rule performing PCA was implemented for the development of visual receptive fields of the robot.

We have firstly shown that the receptive fields and performance of the robot

developed in active condition are significantly different from those developed in four passive conditions. An explanation of this difference is that the coherence of receptive fields developed in active condition plays a vital role in the performance of the robot. This hypothesis is given support by a set of analyses performed on the neural controller and robot behavior.

Our current work aims at extending the analysis to the integration of different information modalities. A new set of experiments where the additional sensory information of actual body movement is available for the neural network of the mobile robot may allow us to explore the role of each modality or inter-modal correlations on the development of the visual system.



---

# 7

## Discussion and Outlook

---

*Specific experiments and results are thoroughly discussed in each chapter. In this chapter, we summarize the contribution of this entire thesis and also discuss promising directions for further research.*

### 7.1 Discussion

Researchers have proposed computational models of active vision and visual attention (see e.g., Frintrop (2005) for a review). In these models attention is driven solely by ‘bottom-up’ signals (e.g., Itti and Koch (2001)) or by the interaction between ‘bottom-up’ signals and ‘top-down’ bias generated by the visual inputs received over time (e.g., Tsotsos et al. (1995)). Notice however that these disembodied models are not allowed to actively choose their visual inputs. Instead they are usually tested with images from publicly available databases or photographs taken in a natural environment as visual stimuli. However recent findings in neurobiology (e.g., Betsch et al., 2004) suggest the significant dependence of visual perception on the behavior of the perceiver. In other words, behavior of the perceiver significantly affects the attentional selection and consequently the statistics of visual input. These models overlook the importance of this sensory-motor loop integrating vision with behavior (Arbib, 1981).

Unlike these models, a distinguishing feature of our methodology is that the system is embodied and allowed to freely interact with its environment and actively select the next sensory input. This is indeed what the enactive approach stressed (Varela et al., 1991): the structural coupling where agents on the one hand select properties in the physical world that are relevant to their structure (e.g., body-size, morphology, sensory-motor capacities) and on the other hand environments select sensory-motor capacities in the agent and thereby constrain its activity (Thompson et al., 2002). This tight coupling and interaction were made possible by the Evolutionary Robotics framework (Nolfi and Floreano, 2000). It further allowed us to study the contribution of active vision to the development of an internal expectation system (Chapter 3) and of visual receptive fields (Chapter 5 and 6),

which have been largely neglected in the biological and computational literature. This void most likely stems on the one hand from the difficulty in measuring neural dynamics in freely behaving animals and on the other hand from the fact that computational models are often designed to operate on predefined input pattern distributions (Parisi et al., 1990). Mobile robots are therefore an ideal tool to investigate adaptive and cognitive processes that take place in a behavioral context because such robots can autonomously select the sensory stimulation by moving in the environment (Pfeifer and Scheier, 1999; Nolfi and Floreano, 2000).

Earlier attempts in this research direction can be found in the literature. In addition to the work reviewed in Section 1.1.5, pioneering work by Wilson (Wilson, 1985) aimed at evolving exploratory behavior of a simple animat situated in a grid world. Nolfi (1996) has demonstrated that sensory-motor coordination of behaving robots can significantly simplify hard perceptual problems. Cliff and Bullock (1993) further extended Wilson's approach and demonstrated that the variations of the sensory field had a significant effect on the animat's external observable behavior. In this thesis we went one step further; as the sensory field affected the behavior of the perceiver, the behavior also altered the development of the sensory field. Instead of handcrafting a bank of sensory fields, we allowed the system to develop its own receptive fields together with the control strategy of its vision and behavior. We have shown that this reciprocal dependence plays a crucial role in the development of visually guided behavior, an internal expectation system and visual receptive fields.

A limitation of our approach comes from the fact that the neural architecture, or topology of the neural network, must be carefully designed for each task, although we deliberately kept the architecture as simple as possible in the present experiments. A promising research direction could therefore be to evolve neural architecture together with synaptic weights. This approach would be very promising in applications where we do not know a priori what type of neural architecture could solve the assigned task. To do this, several promising algorithms have already been proposed (e.g., Stanley and Miikkulainen, 2002; Mattiussi and Floreano, 2007).

Finally, let us bring back the question raised in Section 1.1.3; why do many animals choose to develop strategy to effectively control badly constructed sensors instead of developing well-constructed, high-resolution, and static sensors? As briefly reviewed in Section 1.1.3 and 1.1.4, biological eyes are badly constructed contrary to our belief. It is tempting to speculate that it is more evolutionary expensive for animals to develop elaborate static sensors, than to develop behavioral strategies to make up for bad sensors. Interestingly, our robotic experiments showed a similar tendency; while in many experiments we let robots freely choose the strategy to solve their task, they did choose to actively control sensors, instead of developing static elaborate filters. In these experiments active vision played an essential role in their performance. It may seem that the concept of *behavior makes up for bad sensors* is a general design principle of biological systems<sup>1</sup>. Although a machine may

---

<sup>1</sup>Although a similar idea has been explored in Spier (2004), the approach taken in his article is largely different from ours; the robot controller was hand-coded and the camera was fixed to the robot's body. Instead, the central topics of this thesis are the development of visuo-motor

not be subject to the same limitations as biological systems, behaviorally-enhanced perceptual ability might still be of great interest in many real-time perception-action systems. Perceptual ability goes beyond sensory ability. For further understanding of this principle, we shall describe possible directions to proceed.

## 7.2 Future Directions

We have used simple patterns or shapes in our shape discrimination (Section 2.2.1 and 2.5) and landmark navigation tasks (Chapter 3 and 4). However, to further validate the scalability of the approach, the active vision system needs to be applied to more computationally demanding tasks, e.g., more complex pattern discrimination. The experiments with blindfold people conducted by MacKay (1952) suggest that the concept of a shape, say triangle, is invariantly related with and can be defined by the sequence of elementary responses necessary in the act of replicating the outline of the triangle. Thus the problem of shape recognition is solved by the process of *active replication* of the stimuli perceived and not of passive, one-shot reception of the entire shape<sup>2</sup>. The problem of recognizing complex patterns thus turns into the problem of learning to make complex replicas, which is a problem of sensory-motor coordination, such as writing or walking. We speculate that this enactive approach may greatly contribute to complex pattern recognition problems in computer vision research.

As described in Section 1.1.5, many computational models have been developed to explain the functional role of specific brain areas, e.g., the primary visual cortex (e.g., Linsker, 1988; Hancock et al., 1992; Olshausen and Field, 1996). These disembodied models have significantly contributed to understanding one aspect of neural development in the brain, namely how the bottom-up input shapes or canalizes the neural development. However, this may not be sufficient. A large amount of evidence suggests that neural development in many brain areas is largely affected by top-down processes or the behavioral context of animals (e.g., Polley et al., 1999, 2004). Also most brain areas are interconnected and interacting with each other; therefore we cannot understand each area in isolation. Integrating these disembodied models in the sensory-motor loop allows us to study how these models interact with the environment and how behavior alters the development of these neural systems. Indeed we provided one such example (described in Chapter 5 and 6) where plastic receptive fields which were developed according to a Hebbian learning rule were integrated in the robot controller. The receptive field formation was consistently and significantly affected by the behavior of the robot, and the behavior of the robot was regulated as the receptive fields developed. Other disembodied models of the primary visual cortex, the hippocampus, the parietal or the prefrontal cortex,

---

coordination and the interdependence between behavior and neural development.

<sup>2</sup>The notion of *active replication* is fully consistent with the enactive approach and the finding in Bach-y-Rita's TVSS studies described in Section 1.1.4; blind subjects need to be able to freely control their new 'eye' (i.e., a camera) at will in order to develop the ability to 'see' the world via their skin or tongue.

may also be understood eventually in the behavioral context. Studying brain models in robotic experiments may provide new insight into the interaction between neural development in these brain areas, the body and the environment.

---

# A

# Omnidirectional Active Vision for Evolutionary Car Driving

---

*This appendix describes in detail the experiments briefly overviewed in Section 2.3 and is based on Suzuki et al. (2006).*

**Abstract.** We describe a set of simulations to evolve *omnidirectional active vision*, an artificial retina scanning over images taken via an omnidirectional camera, being applied to a car driving task. While the retina can immediately access features in any direction, it is asked to select behaviorally-relevant features so as to drive the car on the road. Neural controllers which direct both the retinal movement and the system behavior, i.e., the speed and the steering angle of the car, are tested in three different circuits and developed through artificial evolution. We show that the evolved retina moving over the omnidirectional image successfully detects the task-relevant visual features so as to drive the car on the road. Behavioral analysis illustrates its effective strategy in algorithmic, computational, and memory resources.

## A.1 Introduction

The omnidirectional camera is a relatively new optic device that provides a 360 degrees field of view, and it has been widely used in many practical applications including surveillance systems and robot navigation (Matsumoto et al., 1999; Paletta et al., 2001; Stratmann, 2002). However, in most applications visual systems uniformly process the entire image, which would be computationally expensive when detailed information is required. In other cases the focus is determined for particular uses by the designers or users. In other words, the system is not allowed to freely interact with the environment and selectively choose visual features.

Contrarily, all vertebrates and several insects –even those with a very large field of view– share the steerable eyes with a foveal region (Land and Nilsson, 2002), which means that they have been forced to choose necessary information from a vast visual field at any given time so as to survive. Such a sequential and interactive process of selecting and analyzing behaviorally-relevant parts of a visual scene is

called *active vision* (Bajcsy, 1988; Aloimonos et al., 1987; Aloimonos, 1990; Ballard, 1991). Our previous work has demonstrated that it can also be applied to a variety of real world problems (Floreano et al., 2004).

In this chapter we explore *omnidirectional active vision* applied to a car driving task: Coupled with an omnidirectional camera, a square artificial retina can immediately access any visual feature located in any direction, which is impossible for the conventional pan-tilt camera because of the mechanical constraints. It is challenging for the artificial retina to select behaviorally-relevant features in such a broad field of view so as to drive a car on the road. This study may offer a computationally effective methodology to many engineering applications

Omnidirectional active vision is not biologically plausible. But as Langton claimed in (Langton, 1989), it is interesting to study visual systems from a broader point of view which contains those that have never been available in biology. Also, there are several engineering applications that could benefit from omnidirectional vision. The present study may offer a computationally effective methodology to these applications. Some promising ones are discussed in section A.5.

A 1/10 scale model car equipped with an omnidirectional camera and three different circuits are modeled in simulation. We show that the evolved retina moving over the omnidirectional image successfully detects the task-relevant visual features so as to drive the car on the road. Behavioral analysis illustrates its effective strategy in algorithmic, computational, and memory resources. In comparison to the results obtained with a pan-tilt camera mounted on the same car, we show that omnidirectional active vision performs the task very robustly in spite of more difficult initial conditions.

## A.2 Experimental Setup

Figure A.1 shows the real and simulated model cars as well as the views through the real and simulated omnidirectional cameras. The omnidirectional camera consists of a spherical mirror and a CCD camera. It is mounted on a 1/10 scale model car (Kyosho<sup>TM</sup>) which has four motorized wheels. We simulated the car and the circuits using Vortex<sup>1</sup> libraries, a commercially available software package that models gravity, mass, friction, and collisions. Additionally we used a vision software for modeling the view from the omnidirectional camera, which had originally been developed in the Swarm-bots project<sup>2</sup>. Figure A.2 shows the three circuits; ellipse, banana, and eight shaped, used in the present evolutionary experiments.

An artificial retina actively moves over the omnidirectional view<sup>3</sup>. Figure A.3 illustrates the unwrapping process from the polar view to the Cartesian view and the retina overlaid on each image. In order to evaluate the performance of the

---

<sup>1</sup><http://www.cm-labs.com>

<sup>2</sup><http://www.swarm-bots.org/>

<sup>3</sup>A similar approach has been taken for evolving flocking behavior of three simulated robots independently in (Lanza, 2004), inspired by our previous work (Floreano et al., 2004).




Figure A.1: The real 1/10 scale 4WD model car (Kyosho<sup>TM</sup>) with an omnidirectional camera mounted on the roof of the car (top left) and the simulated one (top right). The car base is 19.5 cm (W), 43 cm (L), and 13.5 cm (H). View through the real omnidirectional camera (bottom left) and one through the simulated camera (bottom right).

omnidirectional active vision system, we also simulate a pan-tilt camera mounted on the same car and compare the results obtained in the same experimental condition.

The neural network is characterized by a feedforward architecture with evolvable thresholds and discrete-time, fully recurrent connections at the output layer (Fig. A.4). The input layer is an artificial retina of five by five visual neurons that receive input from a gray level image of 240 by 240 pixels. The activation of a visual neuron, scaled between 0 and 1, is given by the average gray level of all pixels spanned by its own receptive field or by the gray level of a single pixel located within the receptive field. The choice between these two activation methods, or filtering strategies, can be dynamically changed by one output neuron at each time step. Two proprioceptive neurons provide input information about the measured position of the retina with respect to the chassis of the car: the radial and angular coordinates for the omnidirectional camera; or the pan and tilt degrees for the pan-tilt camera. These values are in the interval  $[retina\_size/2, radius - retina\_size/2]$  pixels ( $retina\_size = 240$  pixels,  $radius = 448$  pixels in these experiments) and  $[0, 360]$  degrees for radial and angular coordinates respectively. The values for the pan-tilt camera are in the interval  $[-100, 100]$  and  $[-25, 25]$  degrees respectively. Each value is scaled in the interval  $[0, 1]$  and encoded as a proprioceptive input. A set of memory units stores the values of the output neurons at the previous sensory motor cycle






Figure A.2: Three circuits where the robot car is tested: from left to right, ellipse, banana, and eight shaped circuits. Each circuit is designed such that a 8 m×8 m room accommodates it. The width of the roads in all circuits is 50 cm. In the present experiments the sidelines are omitted.




Figure A.3: Left: The polar image taken by the omnidirectional camera and the overlaid retina. Right: The corresponding unwrapped image and the retina in Cartesian coordinate.

step and sends them back to the output units through a set of connections, which effectively act as the recurrent connections among output units (Elman, 1990). The bias unit has a constant value of  $-1$  and its outgoing connections represent the adaptive thresholds of the output neurons (Hertz et al., 1991).

Output neurons use the sigmoid activation function  $f(x) = 1/(1 + \exp(-x))$  in the range  $[0, 1]$ , where  $x$  is the weighted sum of all inputs. They encode the motor commands of the active vision system and of the car for each sensory motor cycle. One neuron determines the filtering strategy used to set the activation values of visual neurons for the next sensory motor cycle. Two neurons control the movement of the retina (or camera), encoded as speeds relative to the current position. The remaining two neurons encode the directional and rotational speeds of the wheels of the car. Activation values above 0.5 stand for left (directional) and forward (rotational) speeds whereas activation values below 0.5 stand for right and backward speeds, respectively.

The neural network has 165 evolvable connections that are individually encoded on five bits in the genetic string (total length=825). A population of 100 individuals is evolved using truncated rank-based selection with a selection rate of 0.2 (the best






Figure A.4: The architecture is composed of a grid of visual neurons with nonoverlapping receptive fields whose activation is given by the gray level of the corresponding pixels in the image; a set of proprioceptive neurons that provide information about the movement of the retina with respect to the chassis of the car; a set of output neurons that determine at each sensory motor cycle the filtering used by visual neurons, the new angular ( $V_{ang}$ ) and radial ( $V_{rad}$ ) speeds of the retina (or pan and tilt speeds of the pan-tilt camera), and the directional ( $D$ ) and rotational ( $S$ ) speeds of the wheels of the car; a set of memory units whose outgoing connection strengths represent recurrent connections among output units; and a bias neuron whose outgoing connection weights represent the thresholds of the output neurons. Solid arrows between neurons represent fully connected layers of weights between two layers of neurons. Dashed arrows represent 1:1 copy connections (without weights) from output units to memory units, which store the values of the output neurons at the previous sensory motor cycle step.

20 individuals make four copies each) and elitism (two randomly chosen individuals of the population are replaced by the two best individuals of the previous generation). One-point crossover probability is 0.1 and bit-toggling mutation probability is 0.01 per bit.

### A.3 Evolution of Neural Controllers

The fitness function was designed to select cars for their ability to move straight forward as long as possible for the evaluation time of the individual. Each individual is decoded and tested for three trials, each trial lasting 500 sensory motor cycles. A trial can be truncated earlier if the operating system detects a drop of the height of the center of the car, which corresponds to going off-road. The fitness criterion  $F$  is a function of the measured speed  $S_t$  of the four wheels and the steering direction

$D_t$  of the front wheels:

$$F = \frac{1}{E \times T \times S_{max}} \sum_{e=0}^E \sum_{t=0}^{T'} f(S_t, D_t, t) \quad (\text{A.1})$$

$$f(S_t, D_t, t) = S_t \times (1 - \sqrt{|D_t|/D_{max}}) \quad (\text{A.2})$$

where  $S_t$  and  $D_t$  are in the range  $[-8.9, 8.9]$  cm/sec and  $[-45, 45]$  degrees, respectively;  $f(S_t, D_t, t) = 0$  if  $S_t$  is smaller than 0 (backward rotation);  $E$  is the number of trials (three in these experiments),  $T$  is the maximum number of sensory motor cycles per trial (500 in these experiments),  $T'$  is the observed number of sensory motor cycles (for example, 34 for a robot whose trial is truncated after 34 steps if the car goes off-road). At the beginning of each trial the position and orientation of the car as well as the position of the retina within the image are randomly initialized. We performed these replications of the evolutionary run starting with different genetic populations. In all cases the fitness reached stable values in less than 20 generations (Fig. A.5) which correspond to successful on-road driving. The fitness values both of the best individuals and of the population average obtained with the pan-tilt camera were higher than those with the omnidirectional camera in all three circuits. Notice that the fitness can never be one because the car must steer in corners so as not to go off-road.




Figure A.5: Evolution of neural controllers of the car with the pan-tilt camera (left) and the omnidirectional camera (right) in the eight shaped circuits. Thick line shows the fitness values of the best individuals and thin line shows those of the population average. Each data point is the average of three evolutionary runs with different initializations of the population. Vertical lines show the standard error.

## A.4 Behavioral Analysis

The behavioral strategy of the best evolved car equipped with the pan-tilt camera is as follows: At the beginning of the trial the car immediately points the camera

downward and to its right (or left depending on the evolutionary run), and it steers so to maintain the edge between the road and the background within the retina.

Due to the lack of space, we show only the behavioral analysis of the best individual with the omnidirectional camera evolved in the eight shaped circuit because the circuit possesses all of the geometrical characteristics which the ellipse and banana shaped circuits have. It also has an intersection which would disturb the car’s perception if simple edge detection is the strategy developed by the evolved individual, which is nevertheless sufficient for driving in the banana and ellipse shaped circuits. Indeed, since the best evolved individuals with the pan-tilt camera all adopted this strategy, they were disturbed more largely at the intersection and went off-road in several trials. Our preliminary tests also confirmed that the best individuals evolved in the eight shaped circuit were general in the sense that they could drive successfully in the other two circuits as well.

Figure A.6 shows the strategy of the best individual evolved in the eight shaped circuit: During the starting period the evolved car moves back and forth, and then starts going forward at full speed once the retina finds the front road arena. Another effective strategy acquired by the best individuals in other evolutionary runs is that the car moves forward very slowly until the retina finds the road area, and starts driving at full speed immediately after finding it (data not shown). Both behaviors serve to spare time for the retina to find the road area during the most critical period at the beginning of the trial.

In the intersection although the perception of the car is disturbed by the crossing road, which corresponds to the slight deviation of the trajectory, the evolved car manages to find the straight road beyond the intersection by moving the retina upward in the radial direction, which corresponds to “looking farther”, and corrects its trajectory (Fig. A.6, right). After passing the intersection, the retina moved downward again and maintained the straight road area in sight. The rotational speeds of the wheels and the angular position of the retina did not change significantly while passing the intersection.

## A.5 Discussion

The slightly lower fitness values of the evolved individuals with the omnidirectional camera than those with the pan-tilt camera are due to two main reasons: 1) It is harder to find the front road area out of the omnidirectional camera view than out of the pan-tilt camera view<sup>4</sup>; 2) Evolved individuals can find the area, but it takes them some time because during the road search the car does not move much.

Despite this more difficult initial condition, we have shown that artificial evolution selects the individuals capable of “waiting” until the retina finds the appropriate location and of driving at full speed after that. Therefore the slightly lower fitness values of the best evolved individuals with the omnidirectional camera do not mean

---

<sup>4</sup>Notice that the movement of the pan-tilt camera is limited in the interval  $[-100, 100]$  and  $[-25, 25]$  degrees respectively.




Figure A.6: Behavioral and neural analysis of the evolved car with the omnidirectional camera driving in the eight shaped circuit. Top left: Motor activations of the wheel speed (thick line) and the steering angle (thin line) during the first 150 time steps. Speed values above –and below– 0.5 stand for forward –and backward– rotational speeds. Angle values are mapped from 0.5 to 1 for rotations to the left of the car and from 0.5 to 0 for rotations to the right. Dotted vertical line denotes the moment when the retina fixated its position on the front road area. Bottom left: The corresponding trajectory of the car (shown only for the first 70 time steps). Top right: The radial position (elevation) of the retina. Shaded period corresponds to while passing the intersection. The angular position of the retina remained stable and did not change even while passing the intersection. Bottom right: The corresponding trajectory of the car around the intersection.

that those individuals are inferior to those with the pan-tilt camera. The lower fitness is caused by the waiting behavior at the beginning of each test in order to robustly start driving at full speed afterward. Such a behavior has never been observed in any evolutionary run with the pan-tilt camera. Once the road area is detected, the best evolved car with the omnidirectional camera perfectly drives as that with the pan-tilt camera does. The advantages of the present methodology and neural architecture are lower algorithmic, computational, and memory resources.

For comparison with the results obtained with the pan-tilt camera, we did not implement the zooming function in the present experimental setup. However it enables the system to select visual features more specifically by choosing an appropriate resolution for each feature. Indeed, our previous work has demonstrated that in a non-panoramic view the active zooming function played a crucial role in their

performance (Floreano et al., 2004), which encourages us to further apply it to the current system.

The current setup also allows us to lay multiple retinas within a single omnidirectional camera image so that they each specialize in detecting separate features simultaneously and perform behaviors. In real circuits, there are a number of features to which car drivers must pay attention (e.g., sidelines, signals, other cars). Indeed, Dickmanns et al. developed a multi-window system with a conventional non-panoramic camera for detecting such features during real car navigation (Dickmanns et al., 1990), but the position and shape of each window were predetermined by the designers. Active feature selection by multiple retinas which are moving and zooming independently over the omnidirectional view may display significant advantages over an active single retina or the fixed multi-window system in several tasks, e.g., navigation in dynamic, unknown environments. Further investigations must be done to validate this hypothesis.

The present method may also offer a significant advantage in landmark navigation of mobile robots because of its fast, parallel searching for multiple landmarks in any direction. Instead of processing the 360 degrees field of view, the system could actively extract only small task-related regions from the visual field, which would dramatically lighten the computation and the memory consumption.

## A.6 Conclusions

In this chapter we have explored *omnidirectional active vision* applied to a car driving task. The present simulations have shown that the evolved artificial retina moving over the omnidirectional view successfully detects the behaviorally-relevant feature so as to drive the car on the road.

Although it costs time to find the appropriate feature in such a broad field of view during the starting period, the best evolved car overcomes the disadvantage by moving back and forth or moving slowly until the retina finds the appropriate feature. Then the car starts driving forward at full speed. The best evolved car equipped with the omnidirectional camera performs robust driving on the banana, eight, and ellipse shaped circuits in spite of the more difficult initial condition. The advantages of the present methodology and neural architecture are lower algorithmic, computational and memory resources.

Currently we aim to implement an active zooming function of the retinas, to generalize the neural controllers in different-featured circuits (e.g., backgrounds, sidelines, textures), and to transfer the evolved neural controller into the real car shown in Fig. A.1.



---

# B

## Active Perception in Cooperative Tasks

---

*This appendix describes in detail the experiments briefly overviewed in Section 2.5 and is based on Suzuki et al. (2008a).*

**Abstract.** In this chapter we study the coevolution of two active-vision agents in cooperative tasks. Each agent consists of a small artificial retina which is controlled by a simple recurrent neural network. The two agents can “communicate” with each other in a primitive way via continuous signals. These signals are emitted from and received by a particular set of neurons in their neural controller. The two agents are coevolved by a genetic algorithm in two different tasks. We show that the agents with communication outperform the agents without communication in both tasks. Analysis of the best evolved agents reveals how they communicate with each other while performing the task successfully.

### B.1 Introduction

In general multi-agent systems can complete tasks more quickly and reliably than single-agent systems in many practical applications. If systems can communicate with each other in an effective manner, their performance would be further improved. Therefore efficient and reliable communication between agents is of great importance in many multi-agent systems.

Steels and colleagues explored the emergence of language among communicating robots and carried out “talking heads” experiments (see Steels, 2003, for a review). They argued that language bootstrapped or emerged from the interaction between two talking robots. However a problem in their experiments resides in the assumption they made; they assumed that the “hearer” robot could effortlessly inspect the region that the “speaker” robot previously indicated. This process of joint attention (Moore and Dunham, 1995) may not be trivial for robots and rather may be an important prerequisite for the emergence of language.

In this chapter we study the coevolution of two active-vision agents in cooperative

tasks. Each agent consists of a small artificial retina which is controlled by a simple recurrent neural network. The two agents can “communicate” with each other in a primitive way via continuous signals. These signals are emitted from and received by a particular set of neurons (“speaking” and “hearing” neurons, respectively) in their neural controller.

The two agents are coevolved by a genetic algorithm in two different tasks. We used a similar experimental setup to the one used in (Kato and Floreano, 2001; Floreano et al., 2004) in order to compare the performance of cooperative agents with the performance of single agents. We show that the agents capable of communication outperform the agents without communication in both tasks. Behavioral analysis of the best evolved agents reveals how they communicate with each other while performing the task successfully.

## B.2 Method

The neural architecture described in Chapter 2 is extended by adding “speaking” and “hearing” neurons (Fig. B.1). The neural controllers of the two agents have the same architecture while their synaptic weights are developed separately by a genetic algorithm. The “hearing” neurons of one agent directly receive and encode the neural activations of the “speaking” neurons of the other agent.

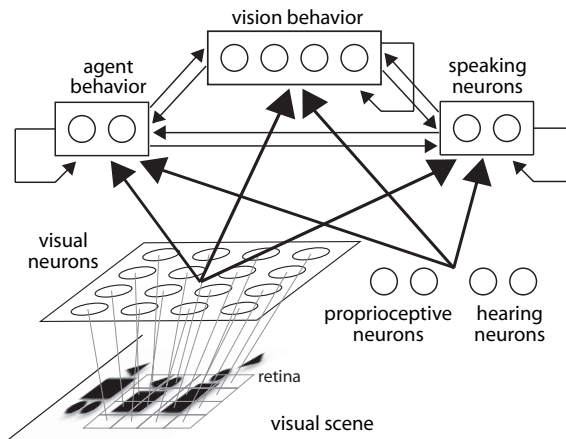


Figure B.1: Neural architecture of one active-vision agent capable of communication. It is largely based on the architecture presented in Floreano et al. (2004), but extended by adding “speaking” neurons, which send signals to the other agent, and “hearing” neurons, which receive signals from the other agent. As in Kato and Floreano (2001), the rest of the architecture is composed of a grid of visual neurons with non-overlapping receptive fields whose activation is given by the grey level of the corresponding pixels in the image; a set of proprioceptive neurons that provide information about the position of the retina; a set of output neurons that determine at each sensory motor cycle the filtering used by visual neurons, the zooming size, the direction and the speed of retina movement.



The synaptic weights of the two neural controllers are developed in two ways; they are all encoded in a single genome or two neural controllers are separated into different populations and cooperatively coevolved (Potter and De Jong, 2000). However, in this chapter we do not describe the results of the cooperative coevolution because the results were not significantly different. A population of 100 genomes is randomly generated, and tested for a number of trials during which its fitness value is computed. The best 20% individuals are reproduced, while the remaining are discarded. These new genomes are randomly mated, crossed over with probability 0.1 per pair, and mutated with probability 0.001 per bit. Crossover consists in swapping genetic material between two strings around a randomly chosen point. For each of the two areas (agents) in the genome, the crossover point is chosen separately (i.e., two-point crossover). Mutation consists in toggling the value of a bit. Finally two copies of the best genomes of the previous generation are inserted in the new population at the places of the randomly chosen genomes (elitism) in order to improve the stability of the evolutionary process.

### B.3 Shape localization task

In the first set of experiments, we ask the two retinas to localize or fixate on the target object together. The target can be either a square or a triangle that can appear at random locations in the visual scene (320x240 pixels) and can take a random size between 20 and 100 pixels in height. Visual scene presents only one image during one trial. Also, noise is added to the entire image by inverting the value of each pixel (black to white or vice versa) with a probability of 0.005 per pixel.

Each team of agents is presented with 20 images, 10 images for each shape. The location and size of the shapes are randomly computed for each new image. Whenever a new image is presented, the values of the output units are reset to zero and retinas are positioned at the center of the image, after which they are free to move and change their zooming size and sampling strategy for 50 time steps. Positive fitness values are given every time step when both retinas find and stay on the target. In order to see the advantage of communication between agents, we also conducted a control experiment under the same setup, but where the agents were not able to communicate with each other.

For both setups, a population of 100 individuals was evolved for 300 generations. Five evolutionary runs were performed, each starting with a different random initialization. Figure B.2 shows that the agents with communication outperformed those without communication. The best evolved agents were able to find the target together in less than 20 steps on average.

Evolved strategies exploiting communication varied slightly across the five evolutionary runs, but all shared some basic features. Evolved agents always tried to move in parallel, while maintaining a small distance between themselves (close to the largest retinal side), as if there was a gravity between them. This strategy al-




Figure B.2: Fitness data of the coevolving active-vision agents performing shape localization task. Fitness values of the population average (thin line) and the best individual (thick line) across 300 generations. *Solid line*: with communication. *Dashed lines*: without communication. The agents with communication performed slightly yet significantly better than those without communication.

lowed one agent to pull the other to its position, when the agent found the target. This cooperative strategy allowed the agents to wipe a larger area. They also used the maximum retina size for most of the time. Their trajectories were diagonal in most cases, the two retinas started from the top left corner of the scene and moved toward the bottom right corner. Figure B.3 shows the trajectories of the evolved agents. Once one of the agents found the target, the retina stayed on the target and ‘pulled’ the other agent to the shape.

On the other hand, evolved agents without communication moved just randomly and therefore it was rarely possible that both agents found and stayed on the target object.

## B.4 Shape discrimination task

In the second set of experiments, two active-vision agents were asked to discriminate the shape (either square or triangle) of the target together. Fitness value was given to both agents when they answered the shape correctly every time step. The fitness function  $F$  was defined as follows:

$$F_{dis} = \frac{1}{I * S} \sum_{i=1}^I \sum_{s=1}^S R_s^i \quad (\text{B.1})$$

where  $R_s^i$  is 1 if both agents gave the correct responses at step  $s$  for image  $i$ , 0.1 if only one agent gave the correct response and 0 otherwise.  $S$  is the time steps per trial




Figure B.3: Four representative trajectories of the best coevolved agents performing shape localization task. The agents always zoomed-in and moved with respect to their top leftmost corner, here marked by a dot. The *dots* drawn after every retina movement are connected with a *line*. For graphical clarity, the values of the cells are not shown, only the retinal perimeter. If the agent made a correct response, the outline of the retina was shown in green, otherwise it was red. Evolved retinas always started from the top left corner and began to scan the screen. They moved in parallel while maintaining a reachable distance between themselves. As soon as one of the agents found the target, it immediately signaled to the other agent.

(50 in this experiment), and  $I$  is the total number of images (20 in this experiment). However, notice that since the agents were asked to provide a discrimination response at every time step and the probability of being presented with a certain shape was 50%, it was very easy to obtain a fitness value of 0.5 by always giving the same response together. As in the previous experiment, we also compared the results with and without communication.

Figure B.4 shows the best evolved agents with communication outperformed those without communication. This result suggests that communication was significantly beneficial in this task too. Analysis of the best evolved agents tells that one of the agents always became inactive or lazy (i.e., staying still, or moving slowly along the border) and just blindly trusted and copied the answer given by the active agent (see Fig. B.5). Notice that the lazy agent has never tried to make its own guess, even when it met the target during the lazy moves (see Fig. B.5, bottom right). In many evolutionary runs evolved agents seemed to develop this strategy; it might be easier for agents to share the answer than to discriminate the shape individually. Further experiments and analysis are required to draw conclusions on this issue.




Figure B.4: Fitness data of the coevolving active vision systems performing shape discrimination task. Fitness values of the population average (thin line) and the best individual (thick line) across 300 generations. *Solid line*: with communication. *Dashed lines*: without communication. Agents with communication outperformed those without communication.

## B.5 Conclusion

In the two experiments described above, active vision agents developed several ways of communication with each other by using simple signals. In the shape localization task, evolved agents moved in a cooperative manner by exploiting the communica-

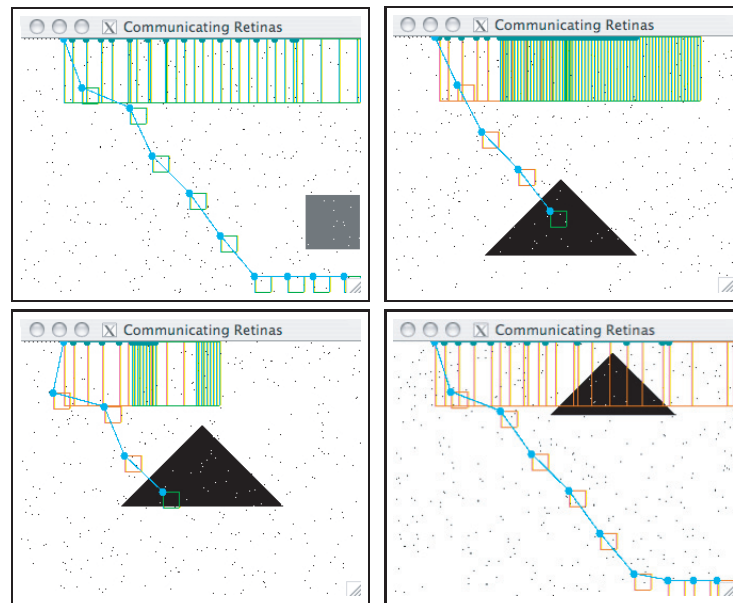


Figure B.5: Four representative trajectories of the best coevolved agents performing shape discrimination task. One agent was always inactive and just moved slowly along the border. As soon as the active agent found the target and discriminated it, the other agent also started emitting the correct response. Notice that the inactive agent never try to make its own guess, even when it meets the target during the lazy moves (the bottom right figure).

tion signals to perform the task. In the shape discrimination task, the strategy of the evolved agents was very different from the one developed in the previous task; since they needed just to continuously emit correct responses (either square or triangle), they did not cooperate with each other. Instead, the ‘role specialization’ emerged by exploiting communication; one ‘active’ agent always looked for the target while the other ‘lazy’ agent just blindly trusted and copied the answer given by the active agent. Although the active vision agents do not cooperate in an intelligent way as animals do, the present experiments show how a primitive, yet effective communication and cooperative behaviors of agents can be developed in cognitive tasks. It is interesting to investigate how such sub-symbolic communication and cooperative behavior between two agents affect and interact with the emergence of grammar or language (Steels, 1997).



---

## Bibliography

---

- Aloimonos, J. (1990). Purposive and qualitative active vision. In *Proceedings of International Conference on Pattern Recognition*, volume 1, pages 346–360.
- Aloimonos, J., Weiss, I., and Bandopadhyay, A. (1987). Active vision. *International Journal of Computer Vision*, 1(4):333–356.
- Aloimonos, Y. (1993). *Active Perception*. Lawrence Erlbaum, Hillsdale, NJ.
- Arbib, M. A. (1981). Perceptual structures and distributed motor control. In *Handbook of Physiology: The Nervous System II. Motor Control*, pages 1449–1480. American Physiological Society, Bethesda, MD.
- Au, W. W. L. (1993). *The Sonar of Dolphins*. Springer, Berlin Heidelberg New York.
- Babineau, D., Lewis, J. E., and Longtin, A. (2007). Spatial acuity and prey detection in weakly electric fish. *PLoS Computational Biology*, 3(3):1–10.
- Bach-y-Rita, P., Collins, C. C., Saunders, F. A., White, B., and Scadden, L. (1969). Vision substitution by tactile image projection. *Nature*, 221:963–964.
- Bach-y-Rita, P. and Kercel, S. W. (2003). Sensory substitution and the human-machine interface. *Trends in Cognitive Sciences*, 7:541–546.
- Bajcsy, R. (1985). Active perception versus passive perception. In *Proceedings of the 3rd IEEE workshop on Computer Vision*, Los Alamitos, CA. IEEE Press.
- Bajcsy, R. (1988). Active perception. *Proceedings of the IEEE*, 76:966–1005.
- Baldwin, J. M. (1896). A new factor in evolution. *American Naturalist*, 30:441–451.
- Ballard, D. H. (1991). Animate vision. *Artificial Intelligence*, 48(1):57–86.

- Ballard, D. H., Hayhoe, M. M., Li, F., and Whitehead, S. D. (1992). Hand-eye coordination during sequential tasks. *Philosophical Transactions of the Royal Society, Series B*, 337:331–339.
- Belew, R. K. and Mitchell, M. (1996). *Adaptive Individuals in Evolving Populations: Models and Algorithms*. Addison-Wesley, Redwood City, CA.
- Betsch, B. Y., Einhäuser, W., Körding, K. P., and König, P. (2004). The world from a cat’s perspective – statistics of natural videos. *Biological Cybernetics*, 90(1):41–50.
- Blackmore, S. J., Brelstaff, G., Nelson, K., and Troscianko, T. (1995). Is the richness of our visual world an illusion? Transsaccadic memory for complex scenes. *Perception*, 24:1075–1081.
- Blake, A. and Yuille, A. (1992). *Active Vision*. The MIT Press, Cambridge, MA.
- Blakemore, C. and Cooper, G. F. (1970). Development of the brain depends on the visual environment. *Nature*, 228:477–478.
- Brecht, M., Preilowski, B., and Merzenich, M. M. (1997). Functional architecture of the mystacial vibrissae. *Behavioural Brain Research*, 84:81–97.
- Brooks, R. A. (1991). Intelligence without representation. *Artificial Intelligence*, 47:139–159.
- Bruner, J. (1968). *Processes of Cognitive Growth: Infancy*. Clark University Press, Worcester, MA.
- Bullock, T. H. and Heiligenberg, W. (1986). *Electroreception*. Wiley, New York.
- Cartwright, B. A. and Collett, T. S. (1987). Landmark maps for honeybees. *Biological Cybernetics*, 57:85–93.
- Christensen, H. I., Bunke, H., and Bowyer, K. (1993). *Active Robot Vision: Camera Heads, Model Based Navigation and Reactive Control*. World Scientific Publishing, River Edge, NJ, USA.
- Cliff, D. and Bullock, S. (1993). Adding “foveal vision” to Wilson’s animat. *Adaptive Behavior*, 2(1):49–72.
- Cliff, D. T. (1991). Computational neuroethology: A provisional manifesto. In Meyer, J.-A. and Wilson, S. W., editors, *From Animals to Animats: Proceedings of the First International Conference on the Simulation of Adaptive Behavior*. The MIT Press, Cambridge, MA.
- Collett, T. S. and Cartwright, B. A. (1983). Eidetic images in insects: Their role in navigation. *Trends in Neuroscience*, 6:101–105.



- Collett, T. S. and Collett, M. (2002). Memory use in insect visual navigation. *Nature Reviews Neuroscience*, 3:542–552.
- Collett, T. S., Collett, M., and Wehner, R. (2001). The guidance of desert ants by extended landmarks. *The Journal of Experimental Biology*, 204:1635–1639.
- Collett, T. S., Graham, P., and Harris, R. A. (2007). Novel landmark-guided routes in ants. *The Journal of Experimental Biology*, 210:2025–2032.
- Collett, T. S. and Land, M. F. (1975). Visual control of flight behaviour in the hoverfly, *Syrirta pipiens* L. *Journal of Comparative Physiology*, 99:1–66.
- Di Paolo, E. A. (2000). Homeostatic adaptation to inversion of the visual field and other sensorimotor disruptions. In Meyer, J.-A., Berthoz, A., Floreano, D., Roitblat, H., and Wilson, S., editors, *From Animals to Animats VI: Proceedings of the Sixth International Conference on the Simulation of Adaptive Behavior*, pages 440–449. The MIT Press, Cambridge, MA.
- Dickmanns, E. and Graefe, V. (1988). Dynamic monocular machine vision. *Machine Vision and Applications*, 1:223–240.
- Dickmanns, E. D., Mysliwetz, B., and Christians, T. (1990). An integrated spatio-temporal approach to automatic visual guidance of autonomous vehicles. *IEEE Transactions on Systems, Man, and Cybernetics*, 20(6):1273–1284.
- Dürr, V., König, Y., and Kittman, R. (2001). The antennal motor system of the stick insect *carausius morosus*: Anatomy and antennal movement pattern during walking. *Journal of Comparative Physiology A*, 187:131–144.
- Elman, J. (1990). Finding structure in time. *Cognitive Science*, 14:179–211.
- Field, D. (1994). What is the goal of sensory coding? *Neural Computation*, 6:559–601.
- Findlay, J. M. and Gilchrist, I. D. (2003). *Active Vision: The Psychology of Looking and Seeing*. Oxford University Press, Oxford.
- Floreano, D., Kato, T., Marocco, D., and Sauser, E. (2004). Coevolution of active vision and feature selection. *Biological Cybernetics*, 90(3):218–228.
- Floreano, D., Suzuki, M., and Mattiussi, C. (2005). Active vision and receptive field development in evolutionary robots. *Evolutionary Computation*, 13(4):527–544.
- Franceschini, N. (1992). Real-time visual motor control: From flies to robots. *Proceedings of the Royal Society of London. Series B*, 337:282–293.
- Frintrop, S. (2005). *VOCUS: A Visual Attention System for Object Detection and Goal-Directed Search*. Springer Berlin / Heidelberg.

- Gibson, J. J. (1966). *The Senses Considered as Perceptual Systems*. Houghton Mifflin, Boston.
- Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Lawrence Erlbaum Associates.
- Hancock, P. J., Baddeley, R. J., and Smith, L. S. (1992). The principal components of natural images. *Network*, 3:61–70.
- Hartmann, M. J. (2001). Active sensing capabilities of the rat whisker system. *Autonomous Robots*, 11:249–254.
- Harvey, I., Husbands, P., and Cliff, D. T. (1994). Seeing the light: Artificial evolution, real vision. In Cliff, D. T., Husbands, P., Meyer, J.-A., and Wilson, S., editors, *From Animals to Animats 3: Proceedings of the Third International Conference on the Simulation of Adaptive behavior*, pages 392–401. The MIT Press, Cambridge, MA.
- Harvey, I., Paolo, E. A. D., Wood, R., Quinn, M., and Tuci, E. (2005). Evolutionary robotics: A new scientific tool for studying cognition. *Artificial Life*, 11(1/2):79–98.
- Hassan, E. S. (1989). Hydrodynamic imaging of the surroundings by the lateral line of the blind cave fish *Anoptichthys jordani*. In Coombs, S., Peter, G., and Heinrich, M., editors, *The Mechanosensory Lateral Line: Neurobiology and Evolution*, pages 217–227. Springer, Berlin Heidelberg New York.
- Hayhoe, M. M. (2003). What guides attentional selection in natural environments? In *Proceedings of the Fifth Workshop on Active Vision*, University of Sussex.
- Healy, S. (1998). *Spatial Representation in Animals*. Oxford University Press, Oxford.
- Held, R. (1965). Plasticity in sensory-motor systems. *Scientific American*, 213(5):84–94.
- Held, R. and Hein, A. (1963). Movement-produced stimulation in the development of visually guided behavior. *Journal of Comparative and Physiological Psychology*, 56(5):872–876.
- Hertz, J., Krogh, A., and Palmer, R. G. (1991). *Introduction to the Theory of Neural Computation*. Addison-Wesley, Redwood City, CA.
- Hinton, G. E. and Sejnowski, T. J. (1999). *Unsupervised Learning*. The MIT Press, Cambridge, MA.
- Horn, B. (1986). *Robot Vision*. McGraw-Hill, New York.

- Horseman, B. G., Gebhardt, M., and Honegger, H. W. (1997). Involvement of the suboesophageal and thoracic ganglia in the control of antennal movements in crickets. *Journal of Comparative Physiology A*, 181:195–204.
- Hubel, D. H. and Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology*, 195:215–243.
- Intraub, H. (1997). The representation of visual scenes. *Trends in Cognitive Sciences*, 1:217–222.
- Itti, L. and Koch, C. (2001). Computational modeling of visual attention. *Nature Reviews Neuroscience*, 2(3):194–203.
- Jolliffe, I. T. (1986). *Principal Component Analysis*. Springer Verlag, New York.
- Judd, S. P. D. and Collett, T. S. (1998). Multiple stored views and landmark guidance in ants. *Nature*, 392:710–714.
- Kato, T. and Floreano, D. (2001). An evolutionary active-vision system. In *Proceedings of the Congress of Evolutionary Computation*, volume 1, pages 107–114.
- Kawasaki, M. (1997). Sensory hyperacuity in the hamming avoidance response of weakly electric fish. *Current Opinion in Neurobiology*, 7:473–479.
- Kohler, I. (1964). The formation and transformation of the perceptual world. *Psychological Issues*, 3:1–173.
- Land, M. F. (1971). Orientation by jumping spiders in the absence of visual feedback. *The Journal of Experimental Biology*, 54(1):119–139.
- Land, M. F. and Furneaux, S. (1997). The knowledge base of the oculomotor system. *Philosophical Transactions of the Royal Society of London, Series B*, 352:1231–1239.
- Land, M. F. and McLeod, P. (2000). From eye movements to actions: How batsmen hit the ball. *Nature Neuroscience*, 3(12):1340–1345.
- Land, M. F. and Nilsson, D.-E. (2002). *Animal Eyes*. Oxford University Press, Oxford.
- Langton, C. G. (1989). Artificial life. In Addison-Wesley, editor, *Artificial Life*, pages 1–48.
- Lanza, S. (2004). Active vision in a collective robotics domain. Master’s thesis, Faculty of Engineering, Technical University of Milan.
- Lehrer, M. (1997). Active acquisition of depth information by honeybees. In Srinivasan, M. V. and Venkatesh, S., editors, *From Living Eyes to Seeing Machines*, pages 37–51.

- Linsker, R. (1988). Self-organization in a perceptual network. *Computer*, 3:105–117.
- MacKay, D. M. (1951–1952). Mindlike behaviour in artefacts. *British Journal of Philosophy of Science*, 2:105–121.
- Marocco, D. and Floreano, D. (2002). Active vision and feature selection in evolutionary behavioral systems. In *From Animals to Animats 7: Proceedings of the Seventh International Conference on the Simulation of Adaptive Behavior*, pages 247–255.
- Marr, D. (1982). *Vision*. Freeman, San Francisco, CA.
- Matsumoto, Y., Ikeda, K., Inaba, M., and Inoue, H. (1999). Visual navigation using omnidirectional view sequence. In *Proceedings of International Conference on Intelligent Robots and Systems*, pages 317–322.
- Mattiussi, C. and Floreano, D. (2007). Analog genetic encoding for the evolution of circuits and networks. *IEEE Transactions on Evolutionary Computation*, 11(5):596–607.
- Mayley, G. (1996). Landscapes, learning costs and genetic assimilation. *Evolutionary Computation*, 4(3):213–234.
- Merleau-Ponty, M. (1968). *The Visible and the Invisible*. Northwestern University Press, Evanston, IL.
- Metta, G., Gasteratos, A., and Sandini, G. (2004). Learning to track colored objects with log-polar vision. *Mechatronics*, 14:989–1006.
- Moeller, G. U., Kayser, C., Knecht, F., and König, P. (2004). Interactions between eye movement systems in cats and humans. *Experimental Brain Research*, 157:215–224.
- Möller, P. (1995). *Electric Fishes: History and Behavior*. Chapman & Hall, London.
- Möller, R., Lambrinos, D., Pfeifer, R., Labhart, T., and Wehner, R. (1998). Modeling ant navigation with an autonomous agent. In Pfeifer, R., Blumberg, B., Meyer, J.-A., and Wilson, S. W., editors, *From Animals to Animats 5: Proceedings of the Fifth International Conference on the Simulation of Adaptive Behavior*, pages 185–194, Cambridge, MA. MIT Press.
- Montgomery, J. C., Coombs, S., and Baker, C. F. (2001). The mechanosensory lateral line system of the hypogean form of *astyanax fasciatus*. *Environmental Biology of Fishes*, 62:87–96.
- Moore, C. and Dunham, P. J. (1995). *Joint Attention: Its Origins and Role in Development*. Lawrence Erlbaum, Hillsdale, NJ.

- Mullen, K. T. (1991). Colour vision as a post-receptoral specialization of the central visual field. *Vision Research*, 31:119–130.
- Noë, A. (2002). *Is the Visual World a Grand Illusion?* Imprint Academic, Thorverton.
- Noë, A. (2004). *Action in Perception*. The MIT Press, Cambridge, MA.
- Noë, A. (2005). What does change blindness teach us about consciousness? *Trends in Cognitive Sciences*, 9(5):218.
- Nolfi, S. (1996). Adaptation as a more powerful tool than decomposition and integration. In Fogarty, T. and Venturini, G., editors, *Proceedings of the Workshop on Evolutionary Computing and Machine Learning*, Bari, Italy.
- Nolfi, S. (1998). Evolutionary robotics: Exploiting the full power of self-organization. *Connection Science*, 10:167–183.
- Nolfi, S. and Floreano, D. (1999). Learning and evolution. *Autonomous Robots*, 7(1):89–113.
- Nolfi, S. and Floreano, D. (2000). *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-organizing Machines*. The MIT Press, Cambridge, MA.
- Noorlander, C., Koenderink, J. J., Den Ouden, R. J., and Edens, B. W. (1983). Sensitivity to spatiotemporal colour contrast in the peripheral visual field. *Vision Research*, 23:1–11.
- Noton, D. and Stark, L. (1971a). Scanpaths in eye movements during pattern perception. *Science*, 171:308–311.
- Noton, D. and Stark, L. (1971b). Scanpaths in saccadic eye movements while viewing and recognizing patterns. *Vision Research*, 11:929–942.
- Olshausen, B. A. and Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381:607–609.
- O’Regan, K. J. (1992). Solving the “real” mysteries of visual perception: The world as an outside memory. *Canadian Journal of Psychology*, 46(3):461–488.
- O’Regan, K. J. and Noë, A. (2001). A sensorymotor approach to vision and visual consciousness. *Behavioural and Brain Sciences*, 24(5):939–973.
- Paletta, L., Frintop, S., and Hertzberg, J. (2001). Robust localization using context in omnidirectional imaging. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 2072–2077.
- Parisi, D., Cecconi, F., and Nolfi, S. (1990). Econet: Neural networks that learn in an environment. *Network*, 1:149–168.

- Parker, A. (2003). *In the Blink of an Eye: How Vision Sparked the Big Bang of Evolution*. Perseus Publishing.
- Pfeifer, R. and Scheier, C. (1999). *Understanding Intelligence*. The MIT Press, Cambridge, MA.
- Polley, D. B., Chen-Bee, C. H., and Frostig, R. D. (1999). Two directions of plasticity in the sensory-deprived adult cortex. *Neuron*, 24:623–637.
- Polley, P. B., Kvašňák, E., and Frostig, R. D. (2004). Naturalistic experience transforms sensory maps in the adult cortex of caged animals. *Nature*, 429:67–71.
- Potter, M. A. and De Jong, K. A. (2000). Cooperative coevolution: An architecture for evolving coadapted subcomponents. *Evolutionary Computation*, 8(1):1–29.
- Pratt, S. C., Brooks, S. E., and Franks, N. R. (2001). The use of edges in visual navigation by the ant *Leptothorax albipennis*. *Ethology*, 107:1125–1136.
- Rao, R. and Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2:79–87.
- Reil, T. and Husbands, P. (2002). Evolution of central pattern generators for bipedal walking in a real-time physics environment. *IEEE Transactions on Evolutionary Computation*, 6(2):159–168.
- Rensink, R. A. (2002). Change detection. *Annual Review of Psychology*, 53:245–277.
- Sandini, G. and Tagliasco, V. (1980). An anthropomorphic retina-like structure for scene analysis. *Computer Graphics and Image Processing*, 14(4):365–372.
- Sanger, T. D. (1989). Optimal unsupervised learning in a single-layer feedforward neural network. *Neural Networks*, 2:459–473.
- Scheier, C., Pfeifer, R., and Kuniyoshi, Y. (1998). Embedded neural networks: Exploiting constraints. *Neural Networks*, 11:1551–1596.
- Sharkey, N. E. and Ziemke, T. (1998). Life, mind, and robots: The ins and outs of embodied cognition. In *Hybrid Neural Systems*, pages 313–332.
- Shibata, T. and Schaal, S. (2001). Biomimetic gaze stabilization based on feedback-error-learning with nonparametric regression networks. *Neural Networks*, 14:201–216.
- Simons, D. J. and Levin, D. T. (1997). Change blindness. *Trends in Cognitive Sciences*, 1:261–267.
- Simons, D. J. and Rensink, R. A. (2005). Change blindness: Past, present, future. *Trends in Cognitive Sciences*, 9:16–20.



- Singer, W. (1987). Activity-dependant self-organisation of synaptic connections as a substrate of learning. In Changeux, J. P. and Konishi, M., editors, *The Neural and Molecular Bases of Learning*. Wiley, London.
- Sobel, E. C. (1990). The locust's use of motion parallax to measure distance. *Journal of Comparative Physiology A*, 167:579–588.
- Spier, E. (2004). Behavioural categorisation: Behaviour makes up for bad vision. In Pollack, J., Bedau, M., Husbands, P., Ikegami, T., and Watson, R., editors, *Artificial Life IX: Proceedings of the Ninth International Conference on the Simulation and Synthesis of Life*, pages 133–138. The MIT Press.
- Srinivasan, M. V. (1998). Animal navigation: Ants march as they march. *Nature*, 392:660–661.
- Srinivasan, M. V. and Venkatesh, S. (1997). *From Living Eyes to Seeing Machines*. Oxford University Press, Oxford New York Tokyo.
- Stanley, K. O. and Miikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10(2):99–127.
- Steels, L. (1997). The synthetic modeling of language origins. *Evolution of Communication*, 1(1):1–34.
- Steels, L. (2003). Evolving grounded communication for robots. *Trends in Cognitive Sciences*, 7(7):308–312.
- Steinman, R. M. and Collewijn, H. (1980). Binocular retinal image motion during active head rotation. *Vision Research*, 20:415–429.
- Stratmann, I. (2002). Omnidirectional imaging and optical flow. In *Proceedings of the IEEE Workshop on Omnidirectional Vision*, pages 104–114. IEEE Computer Society.
- Stratton, G. M. (1896). Some preliminary experiments of vision without inversion of the retinal image. *Psychological Review*, 3:611–617.
- Stratton, G. M. (1897). Vision without inversion of the retinal image. *Psychological Review*, 4:341–360.
- Suzuki, M. and Floreano, D. (2006). Evolutionary active vision toward three dimensional landmark-navigation. In Nolfi, S., Baldassarre, G., Calabretta, R., Hallam, J., Marocco, D., Miglino, O., Meyer, J.-A., and Parisi, D., editors, *From Animals to Animats 9: Proceedings of the Ninth International Conference on the Simulation of Adaptive Behavior*, pages 263–273.

- Suzuki, M., Floreano, D., and Di Paolo, E. A. (2005a). Constraints on body movement during visual development affect behavior of evolutionary robots. In *Proceedings of the International Joint Conference on Neural Networks*, pages 2778–2783, Montréal, Canada.
- Suzuki, M., Floreano, D., and Di Paolo, E. A. (2005b). The contribution of active body movement to visual development in evolutionary robots. *Neural Networks*, 18(5/6):656–665.
- Suzuki, M., Gritti, T., and Floreano, D. (2008a). Active vision for goal-oriented humanoid bipedal walking. In *Creating Brain-like Intelligence: Challenges and Achievements*. Springer Verlag, Berlin / Heidelberg. in preparation for publication.
- Suzuki, M., Luo, J. R., and Floreano, D. (2008b). Sub-symbolic communication improves dynamic joint attention between two active-vision agents. in preparation for submission.
- Suzuki, M., van der Blij, J., and Floreano, D. (2006). Omnidirectional active vision for evolutionary car driving. In *Proceedings of the Ninth International Conference on Intelligent Autonomous Systems*, pages 153–161, Tokyo, Japan.
- Swain, M. J. and Stricker, M. A. (1993). Promising directions in active vision. *International Journal of Computer Vision*, 11(2):109–126.
- Takanishi, A., Takeya, T., Karaki, H., Kumeta, M., and Kato, I. (1990). A control method for dynamic walking under unknown external force. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 795–801, Tsuchiura, Japan.
- Taylor, J. G. (1962). *The Behavioral Basis of Perception*. Yale University Press, New Haven.
- Thomas, J. A., Moss, C. F., and Vater, M. (2004). *Echolocation in Bats and Dolphins*. University of Chicago Press, Chicago.
- Thompson, E. T., Palacios, A., and Varela, F. J. (2002). Ways of coloring: Comparative color vision as a case study for cognitive science. In Noë, A. and Thompson, E. T., editors, *Vision and Mind: Selected Readings in the Philosophy of Perception*. The MIT Press, Cambridge, MA.
- Tsotsos, J. K., Culhane, S. M., Wai, W. Y. K., Lai, Y., Davis, N., and Nufflo, F. (1995). Modeling visual attention via selective tuning. *Artificial Intelligence*, 78(1/2):507–545.
- Urzelai, J. and Floreano, D. (2001). Evolution of adaptive synapses: Robots with fast adaptive behavior in new environments. *Evolutionary Computation*, 9:495–524.



- Varela, F. J., Thompson, E. T., and Rosch, E. (1991). *The Embodied Mind*. The MIT Press, Cambridge, MA.
- Verschure, P. F. M. J., Voegtlin, T., and Douglas, R. J. (2003). Environmentally mediated synergy between perception and behaviour in mobile robots. *Nature*, 425:620–624.
- Vieville, T. (1997). *A Few Steps Towards 3D Active Vision*. Springer-Verlag New York, Inc., Secaucus, NJ, USA.
- von Campenhausen, C., Riess, I., and Weissert, R. (1981). Detection of stationary objects by the blind cave fish *Anoptichthys jordani* (Characidae). *Journal of Comparative Physiology A*, 143:369–374.
- Wandell, B. A. (1995). *Foundations of Vision*. Sinauer Associates, Inc., Sunderland, MA.
- Wilson, S. W. (1985). Knowledge growth in an artificial animal. In *Proceedings of the First International Conference on Genetic Algorithms*, pages 16–23. Lawrence Erlbaum Associates, Inc., Mahwah, NJ.
- Yarbus, A. L. (1967). *Eye Movements and Vision*. Plenum, New York.
- Ziemke, T. (2001). Are robots embodied? In Balkenius, P., Zlatev, J., Brezeal, C., Dautenhahn, K., and Kozima, H., editors, *Proceedings of the First International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, volume 85, pages 75–83, Lund, Sweden.
- Ziemke, T. (2004). Embodied AI as science: Models of embodied cognition, embodied models of cognition, or both? In *Embodied Artificial Intelligence*, volume 3139, pages 27–36. Springer Berlin / Heidelberg.



---

# Curriculum Vitae

## Mototaka Suzuki

---


### PERSONAL DETAILS

Gender: Male

Date of birth: 15 October, 1978

Place of birth: Kagoshima, Japan

Present Citizenship: Japanese



### Education

04/2004–12/2007 PhD student at Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland.

Project title: *Enactive Robot Vision*; Supervisor: Prof. Dr. Dario Floreano

04/2002–03/2004 Master in Mechanical Engineering at Waseda University, Tokyo, Japan

Specialization: Mobile Robots, Computer Vision, Neural Networks, Evolutionary Robotics

Thesis: *Adaptive Neural Controllers for Autonomous Outdoor Navigation*; Supervisor: Prof. Dr. Dario Floreano (EPFL) and Prof. Dr. Shigeki Sugano (Waseda U.)

04/1998–03/2002 Undergraduate Studies in Mechanical Engineering at Waseda University, Tokyo, Japan

Specialization: Humanoid Robots, Grasping, Computer Vision

Thesis: *Designing the Object Handling Mechanism of a Humanoid Robot*; Supervisor: Prof. Dr. Shigeki Sugano

## Working Experience

- Since 03/2008     **Postdoctoral Fellow** at Mahoney-Keck Center for Brain & Behavior Research, Department of Neuroscience, Columbia University, New York, USA.
- 04/2004–12/2007     **Research and Teaching Assistant** at Laboratory of Intelligent Systems, Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland.
- 04/2003–01/2004     **Research Assistant** at Autonomous Systems Lab 2, Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland.
- 04/2002–03/2003     **Research Assistant:** Laboratory for Advanced Brain Signal Processing, RIKEN Brain Science Institute, Saitama, Japan  
*Overview:* to develop software to make visual stimuli to show examinees for EEG experiment, data analysis, and experimental assist.
- 04/2002–03/2003     **Teaching Assistant:** Department of Mechanical Engineering, Waseda University, Tokyo, Japan
- 03/2002             **Internship** at Applied AI Systems, Inc., Ottawa, Canada  
*Overview:* to develop software of an octopod walking robot which was a commercial product.

## Professional Activity

**Editorial Board:** Frontiers in Neurorobotics (review editor, 11/2007–)

**Program Committee Member:** The 4th International Workshop on Attention in Cognitive Systems, January 8, 2007, Hyderabad, India.

### Paper Reviewer

I did/will serve as a reviewer of papers submitted to the following conferences/journals.

*Conferences:*

- International Joint Conference on Neural Networks, Orland, Florida, 12–17 August 2007.

- International Conference on Architecture of Computing Systems, Zurich, Switzerland, 12–15 March 2007.
- The Fourth International Workshop on Attention and Cognitive Systems, Hyderabad, India, 8 January 2007.
- IEEE/RSJ International Conference on Intelligent Robots and Systems, Beijing, China, 9–15 October 2006.

*Journals:*

- Robotics and Autonomous Systems
- Adaptive Behavior
- Frontiers in Neurorobotics

## Awards and Fellowships

- |                 |   |
|-----------------|---|
| Since 03/2008   | <b>Fellowship for Prospective Researchers</b> , Swiss National Science Foundation, Bern, Switzerland.                                 |
| 03/2006         | <b>Excellent Paper</b> , The Ninth International Conference on Intelligent Autonomous Systems, Tokyo, Japan.                          |
| 08/2005         | <b>Excellent Paper</b> , International Joint Conference on Neural Networks, Montreal, Canada.   |
| 04/2004–03/2005 | <b>Fellowship for Outstanding PhD Students</b> , Doctoral Program of Manufacturing Systems and Robotics, EPFL, Lausanne, Switzerland. |
| 03/2002         | <b>Waseda University Student Cultural Award</b> , Waseda University, Tokyo, Japan.  |

## Publications

*Journal Papers*

1. Suzuki, M. and Floreano, D. (2008). **Enactive robot vision**. *Adaptive Behavior*. in press
2. Suzuki, M., Noda, K., Suga, Y., Ogata, T. and Sugano, S. (2006). **Dynamic perception after visually guided grasping by a human-like autonomous robot**. *Advanced Robotics*, 20(2) pp. 233–254.

3. Suzuki, M., Floreano, D. and Di Paolo, E.A. (2005). **The contribution of active body movement to visual development in evolutionary robots.** *Neural Networks*, 18(5/6) pp. 656–665.
4. Floreano, D., Suzuki, M. and Mattiussi, C. (2005). **Active vision and receptive field development in evolutionary robots.** *Evolutionary Computation*, 13(4) pp. 527–544.
5. Suzuki, M. and Floreano, D. **Active perception for landmark-based navigation.** in preparation for submission.
6. Suzuki, M., Luo, J., Floreano, D. and Ikegami, T. **Communication helps joint attention between active-vision agents.** in preparation for submission.

*Peer-Reviewed Conference Papers*

1. Suzuki, M. and Floreano, D. (2006). **Evolutionary active vision toward three dimensional landmark-navigation.** In *From Animals to Animats 9: Proceedings of the Ninth International Conference on the Simulation of Adaptive Behavior*, pp. 263–273.
2. Suzuki, M., van der Blij, J. and Floreano, D. (2006). **Omnidirectional active vision for evolutionary car driving.** In *Proceedings of the Ninth International Conference on Intelligent Autonomous Systems*, pp. 153–161.
3. Suzuki, M. and Floreano, D. (2006). **Active vision for neural development and landmark navigation.** In *Proceedings of the 50th Anniversary Summit of Artificial Intelligence*, pp. 247–248.
4. Floreano, D. and Suzuki, M. (2006). **Active vision and neural development in animals and robots.** In *Proceedings of the Seventh International Conference on Cognitive Modeling*, pp. 10–11.
5. Suzuki, M., Floreano, D. and Di Paolo, E.A. (2005). **Constraints on body movement during visual development affect behavior of evolutionary robots.** In *Proceedings of the International Joint Conference on Neural Networks*, pp. 2778–2783.
6. Noda, K., Suzuki, M., Tsuchiya, N., Suga, Y., Ogata, T. and Sugano, S. (2003). **Robust modeling of dynamic environment based on robot embodiment.** In *Proceedings of the 2003 IEEE International Conference on Robotics and Automation*, pp. 3565–3570.

*Review Article*

1. Suzuki, M. and Floreano, D. (2006). **Trends in dynamic and embodied cognition.** *International Neural Network Society Newsletter*, 4(2) pp. 9–10.

## Presentations in International Conferences

1. “Evolutionary Active Vision toward Three Dimensional Landmark-Navigation”, From Animals to Animats 9: The Ninth International Conference on the Simulation of Adaptive Behavior, Rome, Italy (October 2006).
2. “Omnidirectional Active Vision for Evolutionary Car Driving”, International Conference on Intelligent Autonomous Systems, Tokyo, Japan (March 2006). **Excellent Paper Award**
3. “Constraints on Body Movement Affect Behavior of Evolutionary Robots”, International Joint Conference on Neural Networks, Montreal, Canada (July 2005). **Excellent Paper Award**
4. “Robust Modeling of Dynamic Environment Based on Robot Embodiment”, The 2003 IEEE International Conference on Robotics and Automation, Taipei, Taiwan (September 2003).

## Supervision of Student Projects

1. Luo, J. (2007). **Dynamic Joint Attention between Two Active-Vision Agents**, doctoral course mini-project, June 2007.
2. Clapera, M. (2007). **Active Vision and Online Adaptation to Different Environments**, semester project, March-June 2007.
3. Ekanadham, C. (2006). **Active Vision and Sequential Landmark Detection in Mobile Robot Navigation**, stage project, June-September 2006.
4. Gritti, T. (2006). **Active Vision for Bipedal Walking in Rough Terrains**, stage project, June-October 2006.
5. Le, H. H. H. (2006). **Visually Guided Flocking Behavior by Physical Robots**, semester project, March-June 2006.
6. van der Blij, J. (2005). **Omnidirectional Active Vision in Evolutionary Car Driving**, diploma project, March-August 2005.
7. Quartier, S. (2005). **Active Vision and Independent Visual Feature Extraction**, semester project, October 2004-February 2005.
8. Béguin, A. (2005). **Active Vision with Omnidirectional Camera**, semester project, October 2004-February 2005.