Active Vision and Receptive Field Development in Evolutionary Robots

Dario Floreano   Dario.Floreano@epfl.ch
Mototaka Suzuki   Mototaka.Suzuki@epfl.ch
Claudio Mattiussi   Claudio.Mattiussi@epfl.ch
Laboratory of Intelligent Systems, Swiss Federal Institute of Technology Lausanne (EPFL), CH-1015 Lausanne, Switzerland

Abstract
In this paper, 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 grid sample 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.

1 Introduction

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

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 recent work (Floreano et al., 2004), we investigated the 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 input layer of the neural network consisted of an array of visual neurons spanning a rudimentary rectangular retina. The output layer was divided into two parts, one controlling the movement of the retina on the visual field and the other the behavior of the robotic system. Therefore, the network had to select the features necessary to perform a given task and at the same time control the vision system. We deliberately did not provide any hidden units to the neural network in order to keep its computing and representational abilities as simple as possible. 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. In a first set of experiments, we showed that sensitivity to very simple linear features is co-evolved with, and exploited by, active vision to perform size and position invariant shape discrimination. We also showed that such discrimination problem is very difficult for an architecturally similar neural system without active behavior. In a second set of experiments, we applied the same co-evolutionary method and architecture for driving a simulated car over mountain roads and showed that active vision is exploited to locate and fixate simple road features that are used to control the car steering and acceleration. In a third set of experiments, we applied once again the same co-evolutionary method and architecture to an autonomous robot equipped with a pan/tilt camera that was asked to navigate in an arena located in an office 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 article, we go one step further and investigate the *ontogenetic development*[1] 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 grid sample 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.

Figure 1: The Koala mobile robot with pan/tilt camera.

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 (Figure 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 (Figure 2). 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 (Figure 3). 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

---

[1]In this paper we will use the term *ontogenetic development* to indicate modifications during lifetime of an individual.
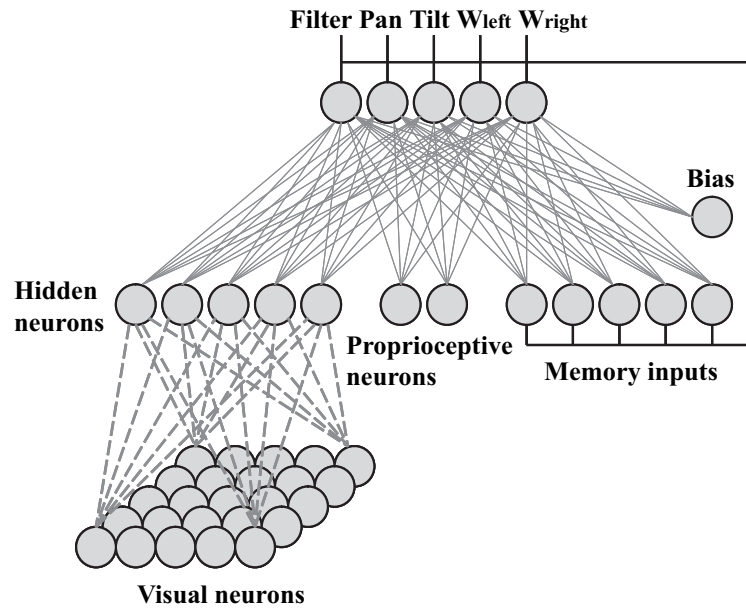
Figure 2: The architecture is composed of a grid of visual neurons with non-overlapping receptive fields whose activation is given by the grey level of the corresponding pixels in the image; a set of proprioceptive neurons that provide information about the movement of the camera with respect to the 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.

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 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.

Figure 3: Left: Snapshot from the robot camera. Center: Pixel average. Right: Pixel sample.

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 Δw_{ij} depends solely on postsynaptic and presynaptic neuron activations y_i, x_j,

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

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 kth hidden neuron o'_k is computed by:

$$o'_k = o_k \times \frac{s_1}{s_k}, \qquad k = 1, 2, 3, 4, 5 \tag{2}$$

where o_k and s_k denote the current output value of kth 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 in the range $[-\sqrt{3/25}, \sqrt{3/25}]$ at the beginning of the life of each individual. 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 connections $\times$ 5 bits) in the learning condition and 950 bits (65 + 125 connections $\times$ 5 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).
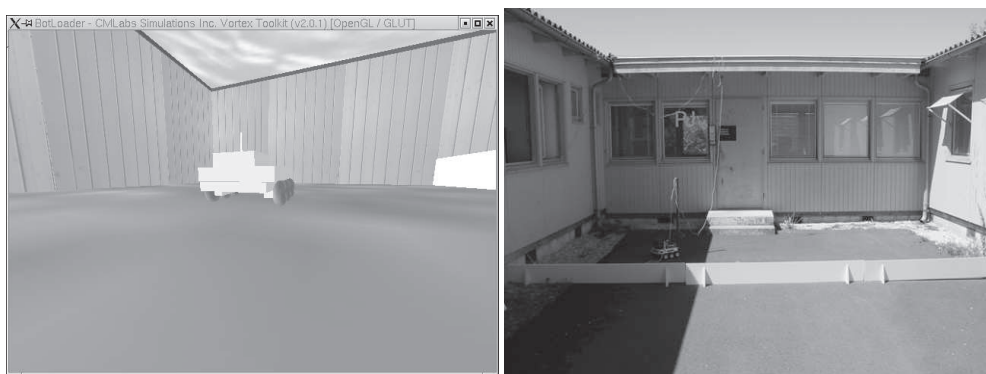
3 Experiments

Figure 4: Simulated and real environment.

Robots were evolved in physics-based simulation and tested both in simulated and real outdoor environments (Figure 4). 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 (Figure 5) 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 of 300 ms. 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:

$$F(S_{left}, S_{right}) = \frac{1}{E \times T} \sum_{e=0}^{E} \sum_{t=0}^{T'} \underbrace{(S_{left}^t + S_{right}^t) \times \left(1 - \sqrt{\frac{|S_{left}^t - S_{right}^t|}{2 \times S_{max}}} \right)}_{f(S_{left}, S_{right}, t)} \quad (3)$$

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 paper), 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 (Figure 5). In both cases, evolved individuals perform large, collision-free paths around the environment for the entire duration of a trial.

One way to observe the effects of receptive field development consists of freezing
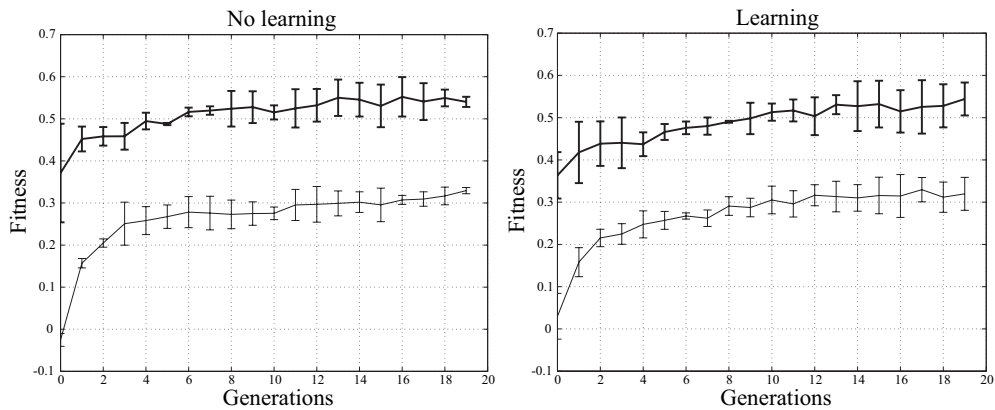
Figure 5: Population average (thin line) and best fitness (thick line) during evolution in physics-based simulations for "No learning" and "Learning" conditions. Each data point is the average of three evolutionary runs with different initializations of the population.

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 (Figure 6). 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 (Figure 7). 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 (Figure 6, 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 real environment that the learning mechanism had to integrate.
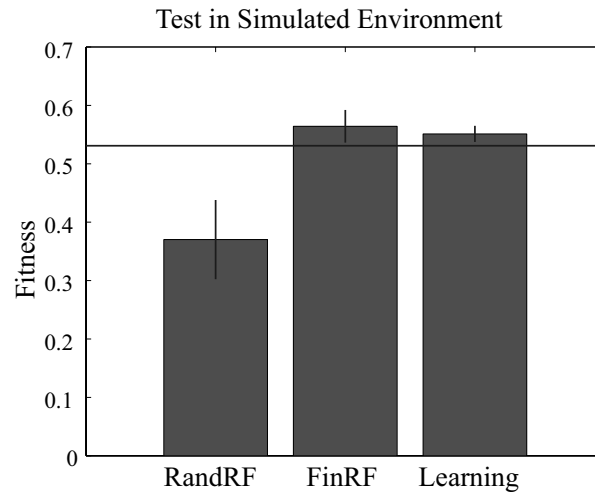
Test in Simulated Environment

Figure 6: 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.

4 Behavioral analysis

Figure 8 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 towards 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 (Figure 9). A possible explanation for this difference is that "Learning" individuals become sensitive to a smaller number of environmental features than "No learning" 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.
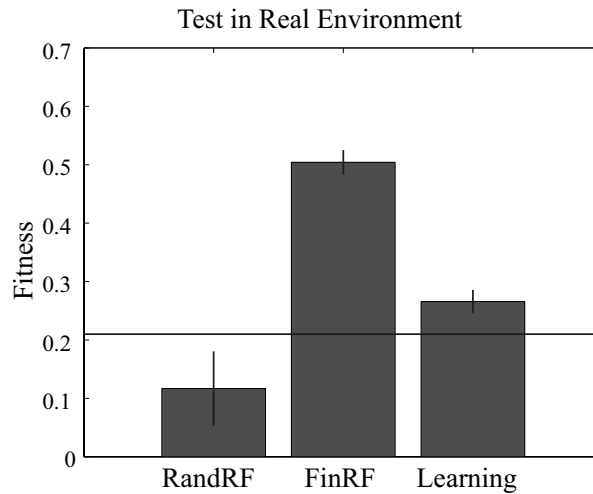
Figure 7: Performance of the best individuals evolved in the simulated environment (three for each evolutionary condition) tested each five times in the *real environment* from different initial positions and orientations. See legend of figure 6 for explanation at column labels.

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 a grid sample of all images that the robot could possibly gather in the outdoor environment. Grid sample 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 (Figure 10). 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 1. Synaptic weight values stabilized approximately after 200 applications of the learning rule, which is less than the trial time of evolutionary robots.

   Figure 11 shows the final receptive fields of the network trained on grid sample 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 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
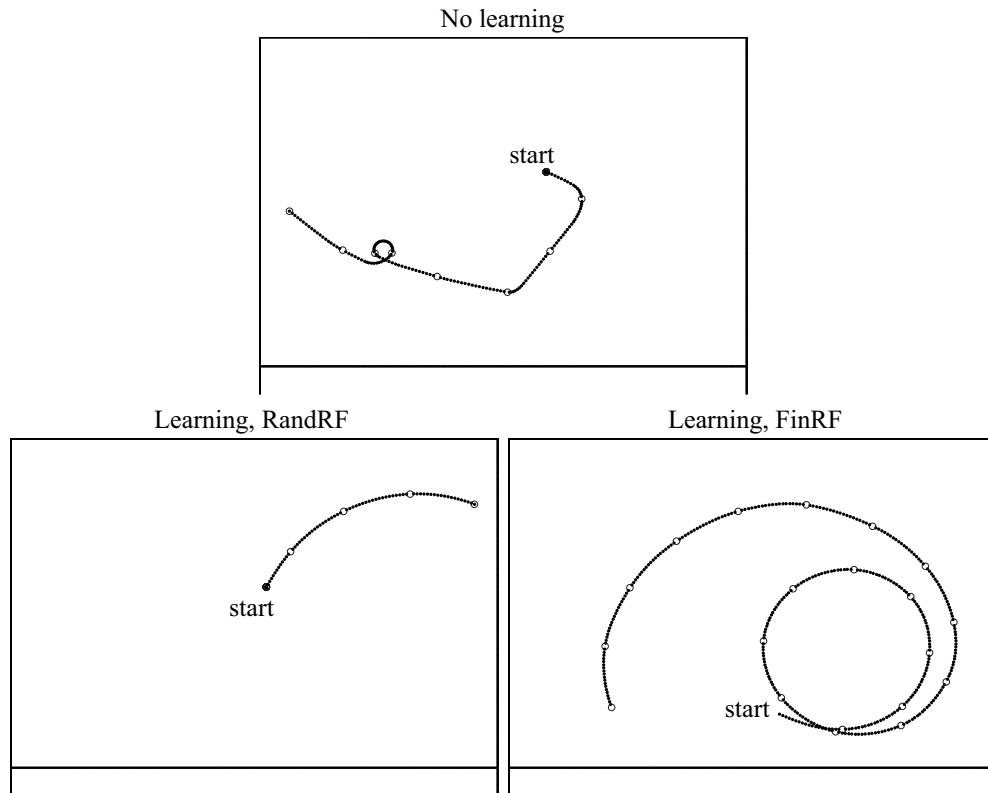
No learning

Learning, RandRF

Learning, FinRF

Figure 8: 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.

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 a grid sample 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 grid sample images. This difference is supported 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 grid sample snapshots (Figure 12). 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
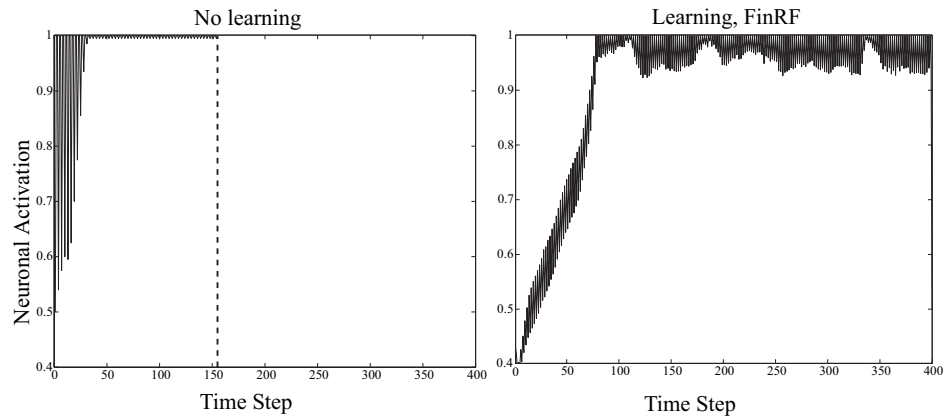
Figure 9: Pan angles of the camera during the trajectory in the outdoor environment for the "No learning" and "Learning" individual with fully formed receptive fields (FinRF). Angle values are mapped from 0.5 to 1 for rotations to the left of the robot and from 0.5 to 0 for rotations to the right of the robot. The trial of the "No learning" individual is terminated after 155 sensory motor cycles (see also Figure 8).

Figure 10: Examples of images taken at different times of the day (9am, 3pm, and 9pm) from the same location.

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 towards the edge between the ground and the wall.

The behavioral and receptive field analysis indicate that evolved learning robots 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 grid sample 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
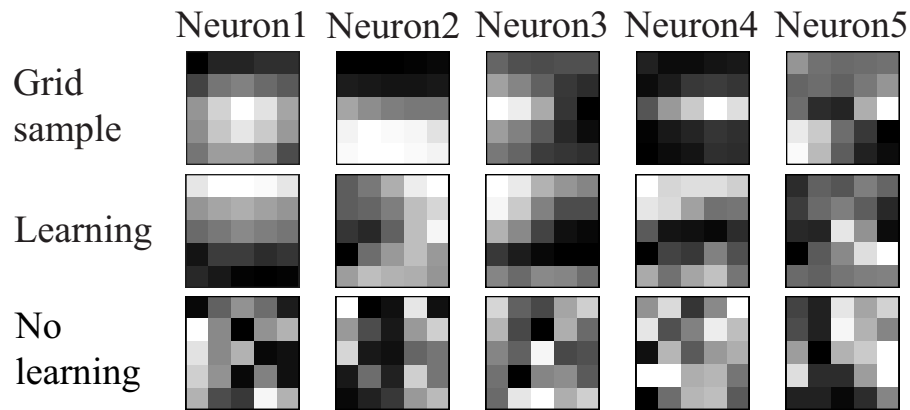
Figure 11: Visual Receptive fields. Each 5 by 5 matrix displays the incoming synaptic strengths for one of the five visual neurons. Gray levels correspond to the sign and strength of the corresponding synapse (white = maximum value; black = minimum value). Leftmost receptive field approximates first principal component of the visual environment.

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 12 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 grid sample 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 Figure 12.

6 Discussion

The results described in this paper 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 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 grid sample 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
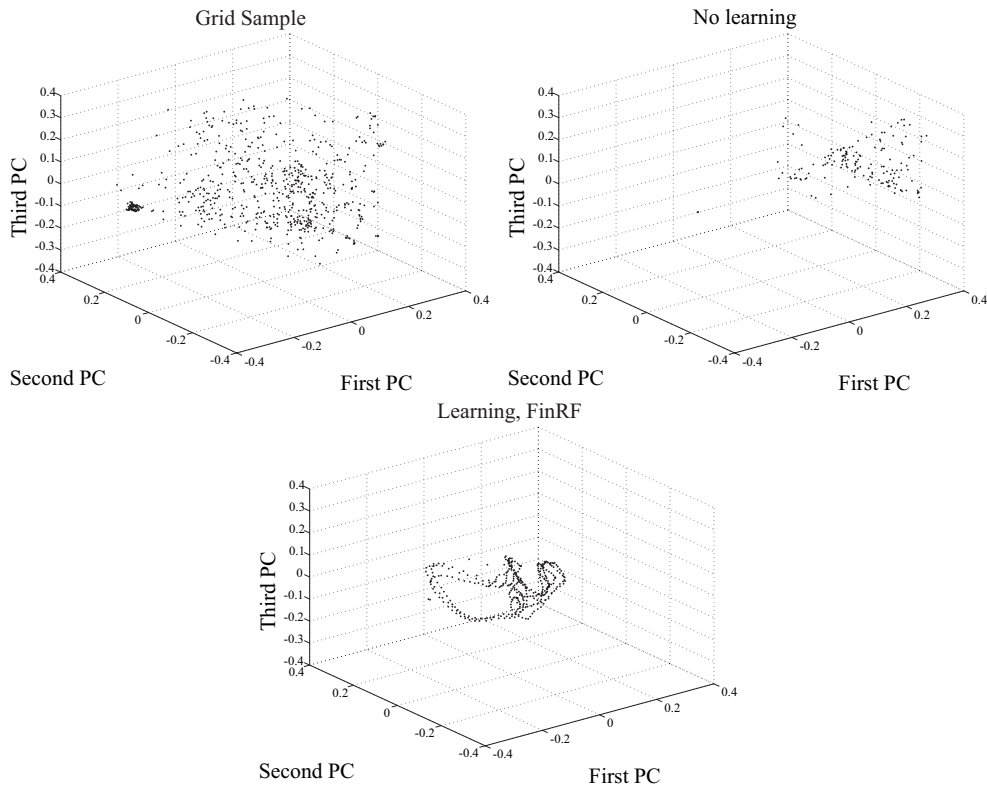
Figure 12: Distribution of snapshots projected onto the three dimensional space of the first three principal components (computed using the exact mathematical procedure). In the last two conditions, snapshots have been recorded during trajectories of evolved individuals. The number of plots in each graph is 324, 155, 400, respectively.

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 paper, 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 paper 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 (DiPaolo, 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 towards 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).

7 Conclusion

The experimental results described in this article 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.

Acknowledgements

References

Aloimonos, Y. (Ed.), 1993. Active Perception. Lawrence Erlbaum, Hillsdale, NJ.

Bajcsy, R., 1985. Active perception versus passive perception. In: Proceedings of the 3rd IEEE workshop on Computer Vision. IEEE Press, Los Alamitos, CA.

Bajcsy, R., 1988. Active Perception. Proceedings of the IEEE 76, 996–1005.

Baldwin, J. M., 1896. A new factor in evolution. American Naturalist 30, 441–451.

Ballard, D. H., 1991. Animate vision. Artificial Intelligence 48, 57–86.

Belew, R. K., Mitchell, M. (Eds.), 1996. Adaptive Individuals in Evolving Populations: Models and Algorithms. Addison-Wesley, Redwood City, CA.

Betsch, B. Y., Einhauser, W., Koerding, K. P., Koenig, P., 2004. The world from a cat's perspective – statistics of natural videos. Biological Cybernetics 90, 41–50.

Blakemore, C., Cooper, G. F., 1970. Development of the brain depends on the visual environment. Nature 228, 477–478.

DiPaolo, E. A., 2000. Homeostatic adaptation to inversion of the visual field and other sensori-motor disruptions. In: Meyer, J., Berthoz, A., Floreano, D., Roitblat, H., Wilson, S. (Eds.), From Animals to Animats VI: Proceedings of the Fifth International Conference on Simulation of Adaptive Behavior. MIT Press-Bradford Books, Cambridge, MA, pp. 440–449.

Elman, J. L., 1990. Finding Structure in Time. Cognitive Science 14, 179–211.

Field, D. J., 1994. What is the goal of sensory coding? Neural Computation 4, 559–601.

Floreano, D., Kato, T., Marocco, D., Sauser, E., 2004. Coevolution of active vision and feature selection. Biological Cybernetics 90(3), 218–228.

Hancock, P. J., Baddeley, R. J., S., S. L., 1992. The principal components of natural images. Network 3, 61–70.

Hertz, J., Krogh, A., Palmer, R. G., 1991. Introduction to the theory of neural computation. Addison-Wesley, Redwood City, CA.

D. Floreano, M. Suzuki, C. Mattiussi

Hinton, G. E., Sejnowski, T. J. (Eds.), 1999. Unsupervised Learning. MIT Press, Cambridge, MA.

Hubel, D. H., Wiesel, T. N., 1968. Receptive fields and functional architecture of monkey striate cortex. Journal of Physiology 195, 215–243.

Jolliffe, I. T., 1986. Principal Components Analysis. Springer Verlag, New York.

Linsker, R., 1988. Self-Organization in a Perceptual Network. Computer 3, 105–117.

Mayley, G., 1996. Landscapes, learning costs and genetic assimilation. Evolutionary Computation 4(3), 213–234.

Nolfi, S., Floreano, D., 1999. Learning and evolution. Autonomous Robots 7(1), 89–113.

Nolfi, S., Floreano, D., 2000. Evolutionary Robotics: Biology, Intelligence, and Technology of Self-Organizing Machines. MIT Press, Cambridge, MA.

Parisi, D., Cecconi, F., Nolfi, S., 1990. Econets: Neural networks that learn in an environment. Network 1, 149–168.

Pfeifer, R., Scheier, C., 1999. Understanding Intelligence. MIT Press, Cambridge, MA.

Polley, D. B., Chen-Bee, C. H., Frostig, R. D., 1999. Two directions of plasticity in the sensory-deprived adult cortex. Neuron 24, 623–637.

Polley, D. B., Kvašňák, E., Frostig, R. D., 2004. Naturalistic experience transforms sensory maps in the adult cortex of caged animals. Nature 429, 67–71.

Sanger, T. D., 1989. Optimal unsupervised learning in a single-layer feedforward neural network. Neural Networks 2, 459–473.

Singer, W., 1987. Activity-dependant self-organisation of synaptic connections as a substrate of learning. In: Changeux, J. P., Konishi, M. (Eds.), The Neural and Molecular Bases of Learning. Wiley, London.

Urzelai, J., Floreano, D., 2001. Evolution of adaptive synapses: Robots with fast adaptive behavior in new environments. Evolutionary Computation 9, 495–524.

Verschure, P. F. M. J., Voegtlin, T., Douglas, R. J., 2003. Environmentally mediated synergy between perception and behaviour in mobile robots. Nature 425, 620–624.

Wandell, B. A., 1995. Foundations of Vision. Sinauer Associates, Inc., Sunderland, MA.

Yarbus, A. L., 1967. Eye movements and vision. Plenum, New York.