ELSEVIER

The contribution of active body movement to visual development in evolutionary robots ☆

Mototaka Suzuki[a,*], Dario Floreano[a], Ezequiel A. Di Paolo[b]

[a]Laboratory of Intelligent Systems, Swiss Federal Institute of Technology (EPFL) CH-1015 Lausanne, Switzerland
[b]Centre for Computational Neuroscience and Robotics, University of Sussex, Brighton, BN19QH, UK

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.
© 2005 Elsevier Ltd. All rights reserved.

Keywords: Active Vision; Visual Receptive Fields; Artificial Evolution; Learning; Neural Networks; Mobile Robots

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 (Kohler, 1964; Stratton, 1896, 1897; 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. 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.

A host of experiments have 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 & Cooper, 1970). Additionally, *active vision*, i.e. the sequential and
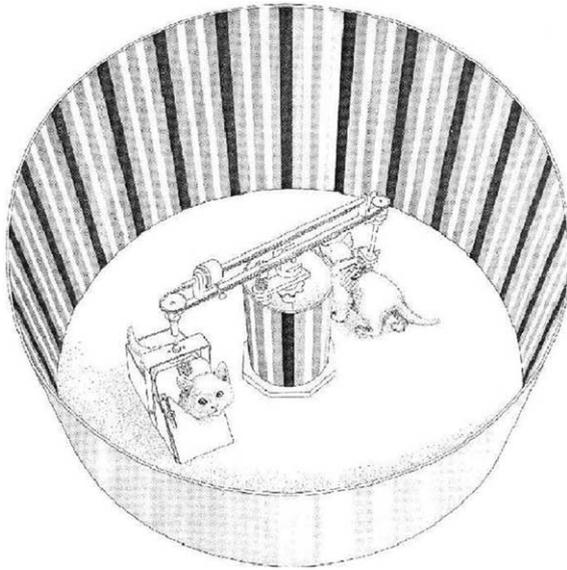

Fig. 1. The original apparatus in Held and Hein (1963), where the gross movements of a kitten moving almost freely were transmitted to a second kitten that was carried in a gondola. Both kittens were allowed to move their head. They received essentially the same visual stimulation because of the unvarying pattern on the walls and the center post of the apparatus. Reproduced, with permission, from (Held, 1965).

interactive process of electing and analyzing parts of a visual scene, selects the subset and sequence of images that the visual system perceives (Aloimonos, 1990; Aloimonos, Weiss & Bandopadhay, 1987; Bajcsy, 1988; 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, Einhäuser, Körding, and König (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 & Kercel, 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, Chen-Bee & Frostig, 1999; Polley, Kvašňák & Frostig, 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 (Baddeley & Smith, 1992; Field, 1994; Hancock, Olshausen & Field, 1996; Rao & 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 investigated the co-development of active vision and receptive fields within the same time scale using behavioral robotic systems (Floreano, Kato, Marocco & Sauser, 2004). We showed 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 article, 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. This article is an extended version of the conference paper (Suzuki, Floreano, D., & Di Paolo, 2005) to provide a more complete and detailed experimental setup and behavioral analyses.

2. Methods

We use a Koala (K-Team S.A.) wheeled robot equipped with a pan/tilt camera (Sony EVI-D31), shown in Fig. 2. The robot has six wheels, but only the central wheel on each side is motorized. The pan and tilt angles of the camera are controlled by two separate and independent motors.

The neural control system of the robot is evolved by means of a genetic algorithm to perform collision-free navigation in an enclosed space using only visual information. The algorithm evolves the neural controller through random mutations, crossover, and a competitive selection process until an evolutionary stable control strategy is found (Nolfi & Floreano, 2000). The details of the evolutionary algorithm are described at the end of this section. The visual receptive fields of the neural network are modified online while the robot is evaluated. The modification consists of a Hebbian learning rule that tends towards the principal components of the input image set. At the end of the evaluation phase, the resulting receptive fields are not memorized in the genotype of the neural controller.
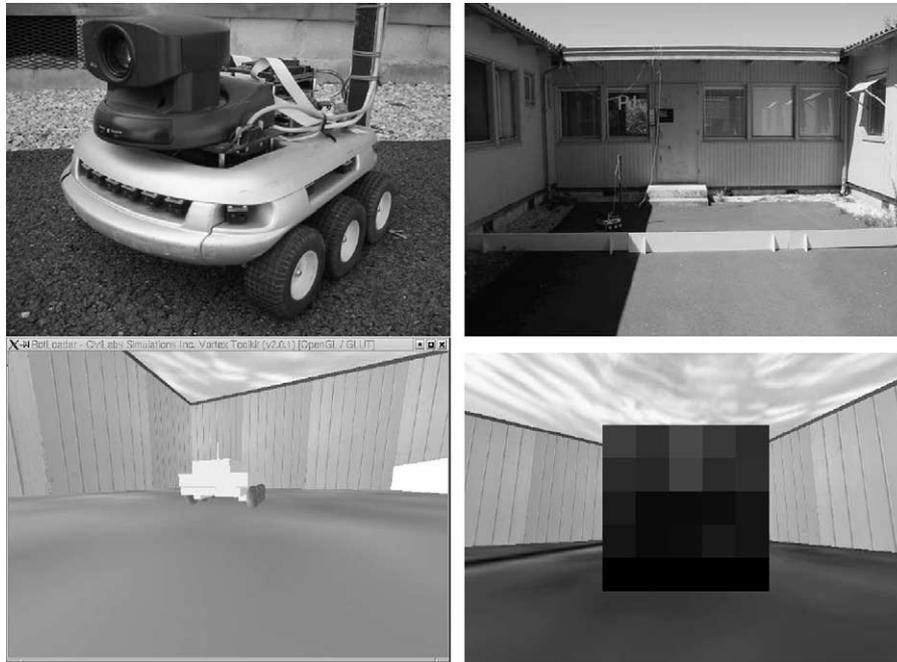
Fig. 2. Top left: The Koala mobile robot by K-Team™ with pan/tilt camera by Sony™. Top right: The real environment. Bottom left: simulation of the robot and the environment. Bottom right: view from the simulated camera. The robot is capable of visually accessing the 5×5 pixels in the center of the image.

The interactions between evolution and learning are described in a separate article (Floreano, Suzuki, M., & Mattiussi, submitted for publication). In this article we study the development of receptive fields in evolved robots, that is in robots that are capable of displaying collision-free navigation. In particular, we analyze the formation of receptive fields while the robot is let free to behave according to the evolved sensory motor pathways and while the robot is constrained to move according to a variety of externally imposed motor commands. We call the first case 'active behavior' and the second 'passive behavior' in accordance to the methodology used in the kittens study.

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. 2, bottom). The texture of the surfaces in the simulated environment was generated from pictures taken in the real outdoor environment.

The neural architecture consists of a simple feedforward network with evolvable thresholds and discrete-time, fully-recurrent connections at the output layer (Fig. 3). The input layer is an artificial retina of 5 by 5 visual neurons that receive input from a gray level image of 240 by 240 pixels. Visual neurons have non-overlapping receptive fields that receive information from a 48 by 48 pixels (240/5) area of the image. The state of visual neurons is determined by the activation of the *filter* output unit of the network. For activation values below 0.5, the state of each visual neuron corresponds to average gray level (scaled in the range [0, 1]) of the corresponding image patch (Fig. 4, center). For activation values equal to

or above 0.5, the state of each visual neuron corresponds to the gray level (scaled in the range [0, 1]) of the top leftmost pixel of the corresponding image patch (Fig. 4, right). In addition, two proprioceptive input neurons encode the measured horizontal (pan) and vertical (tilt)
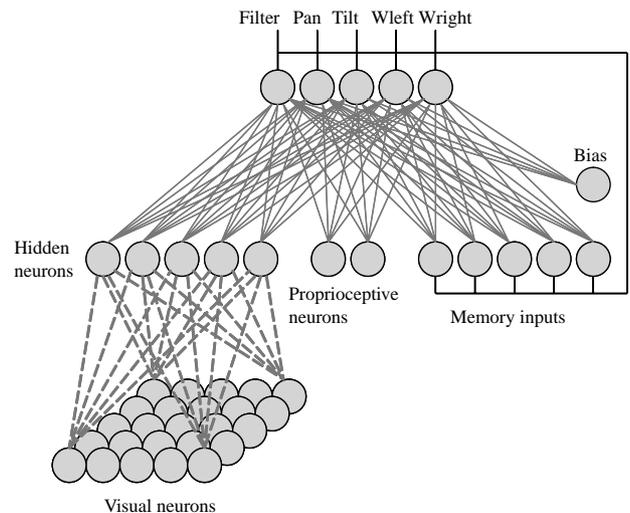

Fig. 3. The architecture is composed of a grid of visual neurons with non-overlapping receptive fields whose activation is given by the gray level of the corresponding pixels in the image; a set of proprioceptive neurons that provide information about the movement of the camera with respect to the chassis 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 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.

Fig. 4. The two filtering strategies of visual input. Left: snapshot from the robot camera. Center: pixel average. Right: pixel sample. The choice between these two strategies can be dynamically changed by one output neuron at each time step.

angles of the camera. These values are in the interval $[-100, 1001]$ 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).

The activations of the output units are passed through the sigmoid function. Two output units determine the speeds of the wheels of the robot. Activation values above 0.5 stand for forward rotational speed whereas activation values below 0.5 stand for backward rotational speed. Two output units encode the speed of the motor of the camera on the horizontal (pan) and vertical (tilt) planes in the same way described above. If the camera has reached a maximum allowed position ($[-100, 100]$ and $[-25, 25]$ degrees for pan and tilt, respectively), output speeds in the same direction have no effect. The remaining output unit encodes the filtering strategy, as described above. Recurrent connections are implemented by adding a set of memory units that encode a copy of the output unit activations at the previous time step (Elman, 1990).

The connections between visual neurons and hidden neurons are randomly initialized at the beginning of the life of each individual (for more detail, see Appendix). These connections are modified by means of a Hebbian learning rule, 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 approximates Principal Component Analysis (PCA) of the input images (Jolliffe, 1986). The modification of connection strength Δw_{ij} depends solely on postsynaptic and presynaptic neuron activations y_i, x_i,

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

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 widely used in previous computational models of receptive field development (e.g.

Hancock et al., 1992). Among the several available models of synaptic plasticity (see Hinton & 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 is equivalent to a widely used technique for image compression.

The neural network is updated at discrete time intervals of 300 ms. At each time interval (sensory motor cycle), the following steps are performed: 1. the activations of the visual and proprioceptive neurons are computed, 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; 3. the activations of the output units are computed; 4. the camera and wheels of the robot are set at the corresponding rotational speed for one sensory motor cycle; 5. the connection weights from visual neurons to hidden neurons are modified using the current neuron activation values.

In step 2 the activations of five hidden units are normalized to operate within the same range of variability in order to equalize the contributions of hidden units to activations of the output units. Notice that this procedure is necessary to prevent the output units of the neural network to depend mainly on the activation of the first one or two hidden units. Our preliminary study showed that the principal five eigenvalues of the correlation matrix of the input vector corresponding to the variances of activation of the hidden neurons have different magnitudes (for more detail, see Suzuki et al., 2005). This means that the activations of the hidden neurons will display different ranges of variability. The normalized output value of the kth hidden neuron o'_k is computed by: $o'_k = o_k \times s_1/s_k$ where o_k and s_k denote the current output value of kth hidden neuron and the standard deviation of all stored output values up to the current time step (s_1 when $k=1$).

The neural network has 65 evolvable connections (including bias connections) that are individually encoded on five bits in the genetic string (total length=325). A population of n genomes is randomly initialized by the computer. Each genome is decoded into the corresponding neural network and tested for a number of trials during which its fitness is computed. Each individual genome is
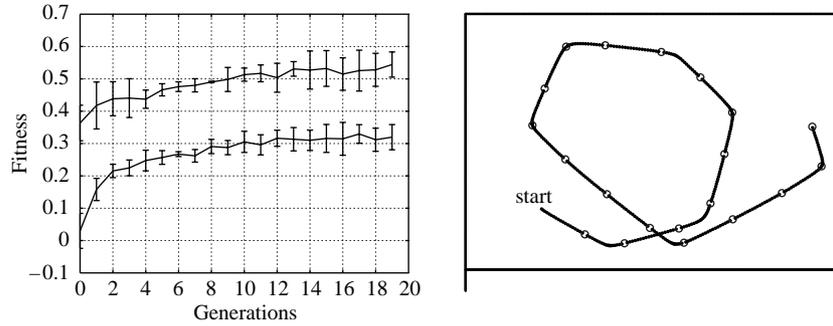
Fig. 5. Evolution with synaptic plasticity. Top: population average (thin line) and best fitness (thick line) during evolution in physics-based simulations. Each data point is the average of three evolutionary runs with different initializations of the population. Bottom: an example of trajectory of the best individual in the last generation while synaptic plasticity is active. A dot is plotted every 20 sensory motor cycles. Vertical lines show the standard error.

then decoded into the connection weights of the neural network (except for the connections from visual neurons to hidden neurons, which are randomly initialized) and tested on the robot while its fitness is computed. A population of 100 individuals is evolved using truncated rank-based selection with a selection rate of 0.2 (the best 20 individuals make four copies each) and elitism (a randomly chosen individual of the population is replaced by the best individual of the previous generation). One-point crossover probability is 0.1 and bit-toggling mutation probability is 0.01 per bit.

3. Simulations and results

3.1. Evolution with synaptic plasticity

In this experiment the robots are evaluated in the context of collision-free navigation. The fitness function was designed to select robots for their ability to move straight forward as long as possible for the duration of life of the individual. This corresponds to 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.[1] A trial can be truncated earlier if the operating system detects an imminent collision.

The fitness criterion $F(S_{\text{left}}, S_{\text{right}})$ is a function of the measured speeds of the left S_{left} and right S_{right} wheels:

$$F(S_{\text{left}}, S_{\text{right}}) = \frac{1}{E \times T} \sum_{e=0}^{E} \sum_{t=0}^{T'} f(S_{\text{left}}, S_{\text{right}}, t) \qquad (2)$$

$$f(S_{\text{left}}, S_{\text{right}}, t) = (S_{\text{left}}^{t} + S_{\text{right}}^{t}) \times \left(1 - \sqrt{\frac{|S_{\text{left}}^{t} - S_{\text{right}}^{t}|}{2 \times S_{\text{max}}}}\right)$$

$$(3)$$

where S_{left} and S_{right} are in the range $[-8, 8]$ cm/sec and

---

[1] Preliminary experiments reported in Floreano et al. (submitted for publication) show that less than 300 updates are necessary to stabilize the plastic weights from visual to hidden neurons.

$f(S_{\text{left}}, S_{\text{right}}, t) = 0$ if S_{left} or S_{right} is smaller than 0 (backward rotation); E is the number of trials (four in these experiments), 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 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 (Fig. 5) which corresponded to collision-free trajectories. Notice that the fitness can never be one because the robot must rotate in order to avoid walls.

Previously we have shown that the resulting fitness values are comparable to those of the evolved robot without learning—all of the synaptic weights in the same architecture were genetically determined and not allowed to change during behavior (Floreano et al., submitted for publication). Furthermore our preliminary study suggests that not only the generalized Hebbian algorithm Eq. (1) that is consistently used in this article, but another learning rule, Oja's *M*-unit rule (Oja, 1989), can also be applied, that is, the fitness achieves the comparable values.

3.2. 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 $(S_{\text{left}}, S_{\text{right}})$ of wheel speeds while the receptive fields were developed.
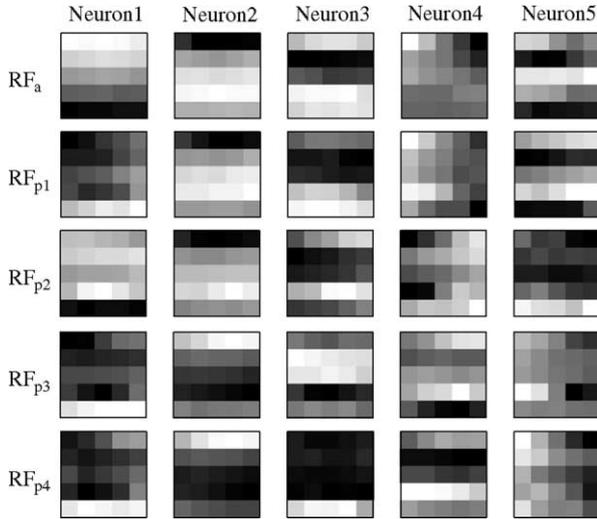
Fig. 6. Receptive fields of five hidden neurons developed in active and passive conditions. Small shaded squares represent the connection strengths from visual neurons, scaled so to fill the gray scale from black (minimum value) to white (maximum value). The leftmost receptive field in each row corresponds 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.

- 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[2] 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 10 times for each condition starting from different random weights and locations. Fig. 6 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},

---

[2] 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}, W_{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.
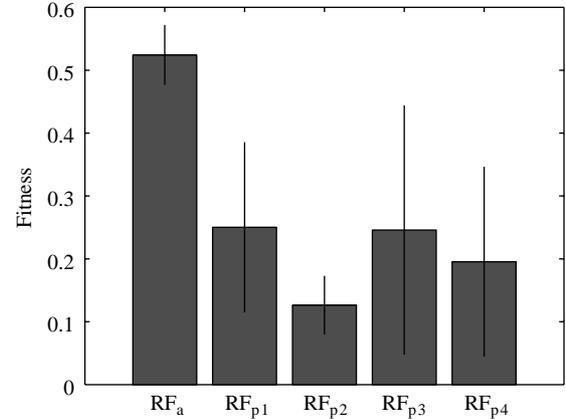

Fig. 7. Performances of the robot with receptive fields developed in active (RF_a) and passive (RF_{p1-4}) conditions. The fitness values are averaged over ten tests. Vertical lines show the standard error.

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. Fig. 7 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. 8. 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}.

4. Analysis

4.1. 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 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.

Fig. 9 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.[3]

---

[3] 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 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.
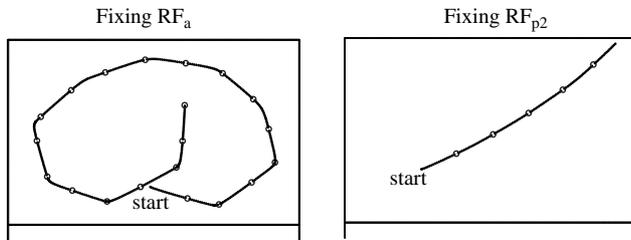
Fixing RF_a

Fixing RF_{p2}

start

start

Fig. 8. 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}.

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. Fig. 10 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 contrasts 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
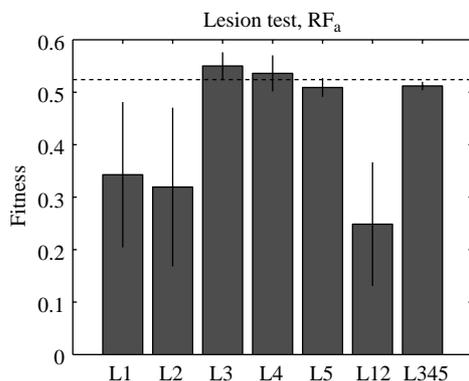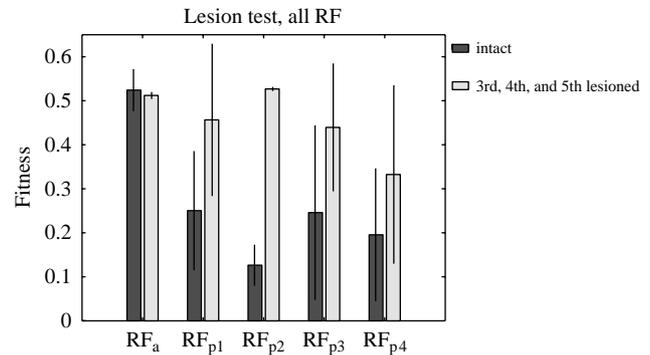
Lesion test, all RF

Fig. 10. 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 10 tests. The performances with all of RF_{p1-4} were improved by lesioning units 3,4, and 5.

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. 11, 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. 11, 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. Fig. 12 shows that the snapshots taken in the active learning condition are distributed in more structured manner than those taken in

Lesion test, RF_a

Fig. 9. 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.
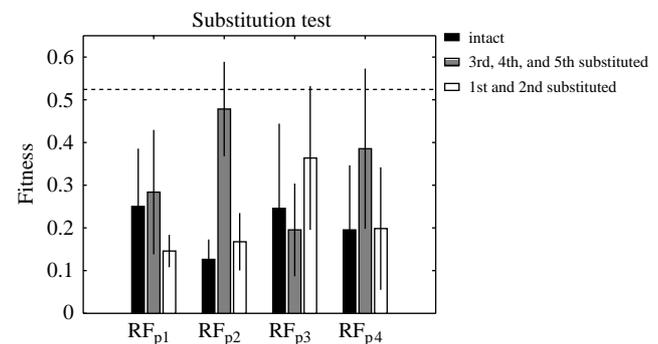
Substitution test

Fig. 11. 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 10 tests.

During learning, RF$_a$
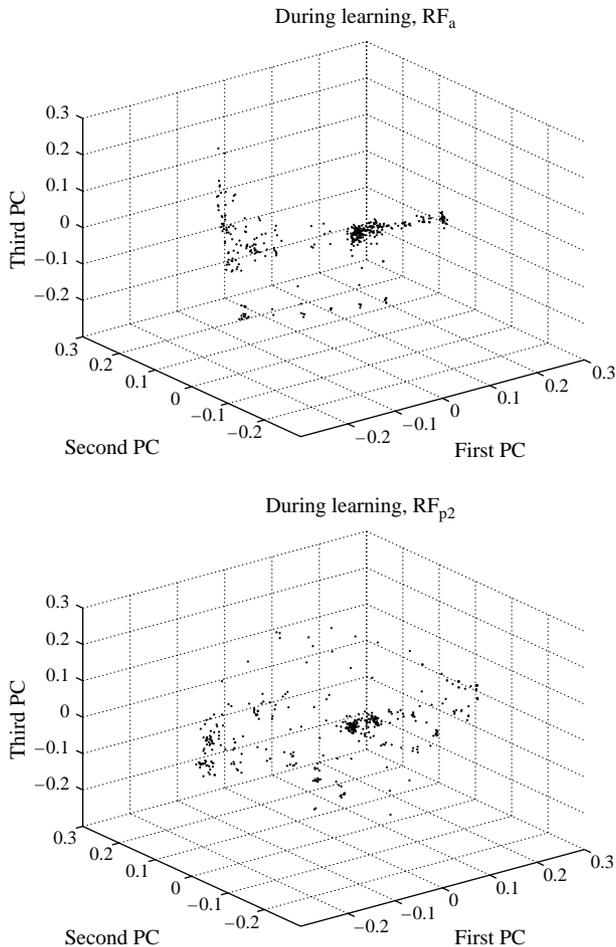

During learning, RF$_{p2}$

Fig. 12. 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$_{p2}$.

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

5. Discussion

Using an experimental setup similar to that used for kittens (Held & 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 cannot 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 p$_{1-4}$ 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 article: 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 article.

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 re-moving 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. Conclusions

We carried out a set of robotic experiments to study the contribution of active body movement to the development of the visual system in the mobile robot. Although the present experimental setup is not exactly same as that shown in Held and Hein (1963), the essence of the original experiment was 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 in 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—the coherence of receptive fields developed in active condition plays an important role in the good performance of the robot—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.

Acknowledgements

Appendix A. 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

$$E[w_1^2 + w_2^2 + \cdots + w_n^2] = n \int_{-A}^{A} \frac{w^2}{2A} \mathrm{d}w.$$

By evaluating the definite integral we obtain $E[w_1^2 + w_2^2 + \cdots + w_n^2] = (n/3) \times A^2$. 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 $n=25$).

References

Aloimonos, J. (1990). Purposive and qualitative active vision. In *Proceedings of international conference on pattern recognition*, Vol. 1. (pp. 346–360).
Aloimonos, J., Weiss, I., & Bandopadhay, A. (1987). Active vision. *International Journal of Computer Vision*, 1(4), 333–356.
Bach-y–Rita, P., & Kercel, S. W. (2003). Sensory substitution and the human-machine interface. *TRENDS in Cognitive Sciences*, 7, 541–546.
Bajcsy, R. (1988). *Active perception Proceedings of the IEEE*, Vol. 76 (pp. 996–1005).
Ballard, D. H. (1991). Animate vision. *Artificial Intelligence*, 48(1), 57–86.
Betsch, B. Y., Einhäuser, W., Körding, K. P., & König, 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.
Elman, J. (1990). Finding structure in time. *Congnitive Science*, 14, 179–211.
Field, D. (1994). What is the goal of sensory coding? *Neural Computation*, 6, 559–601.
Floreano, D., Kato, T., Marocco, D., & Sauser, E. (2004). Coevolution of active vision and feature selection. *Biological Cybernetics*, 90(3), 218–228.
Floreano, D., Suzuki, M., & Mattiussi, C. (submitted for publication). Active vision and receptive field development in evolutionary robots. *Evolutionary Computation*.
Hancock, P. J., Baddeley, R. J., & Smith, L. S. (1992). The principal components of natural images. *Network*, 3, 61–70.
Held, R. (1965). Plasticity in sensory-motor systems. *Scientific American*, 213(5), 84–94.
Held, R., & Hein, A. (1963). Movement-produced stimulation in the development of visually guided behavior. *Journal of Comparative and Physiological Psychology*, 56(5), 872–876.
Hinton, G. E., & Sejnowski, T. J. E. (1999). *Unsupervised learning*. Cambridge, MA: MIT Press.
Jolliffe, I. T. (1986). *Principal component analysis*. New York: Springer Verlag.
Kohler, I. (1964). The formation and transformation of the perceptual world. *Psychological Issues*, 3, 1–173.
Nolfi, S., & Floreano, D. (2000). *Evolutionary robotics: Biology, intelligence, and technology of self-organizing machines*. Cambridge, MA: MIT Press.

Oja, E. (1989). Neural networks, principal components, and subspaces. *International Journal of Neural Systems*, *1*, 61–68.

Olshausen, B. A., & Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, *381*, 607–609.

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

Rao, R., & Ballard, D. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, *2*, 79–87.

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

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 (see also pp. 463–481).

Suzuki, M., Floreano, D., & 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*, July 31–August 4, 2005, Montréal, Canada.

Taylor, J. G. (1962). *The behavioral basis of perception*. New Haven: Yale University Press.