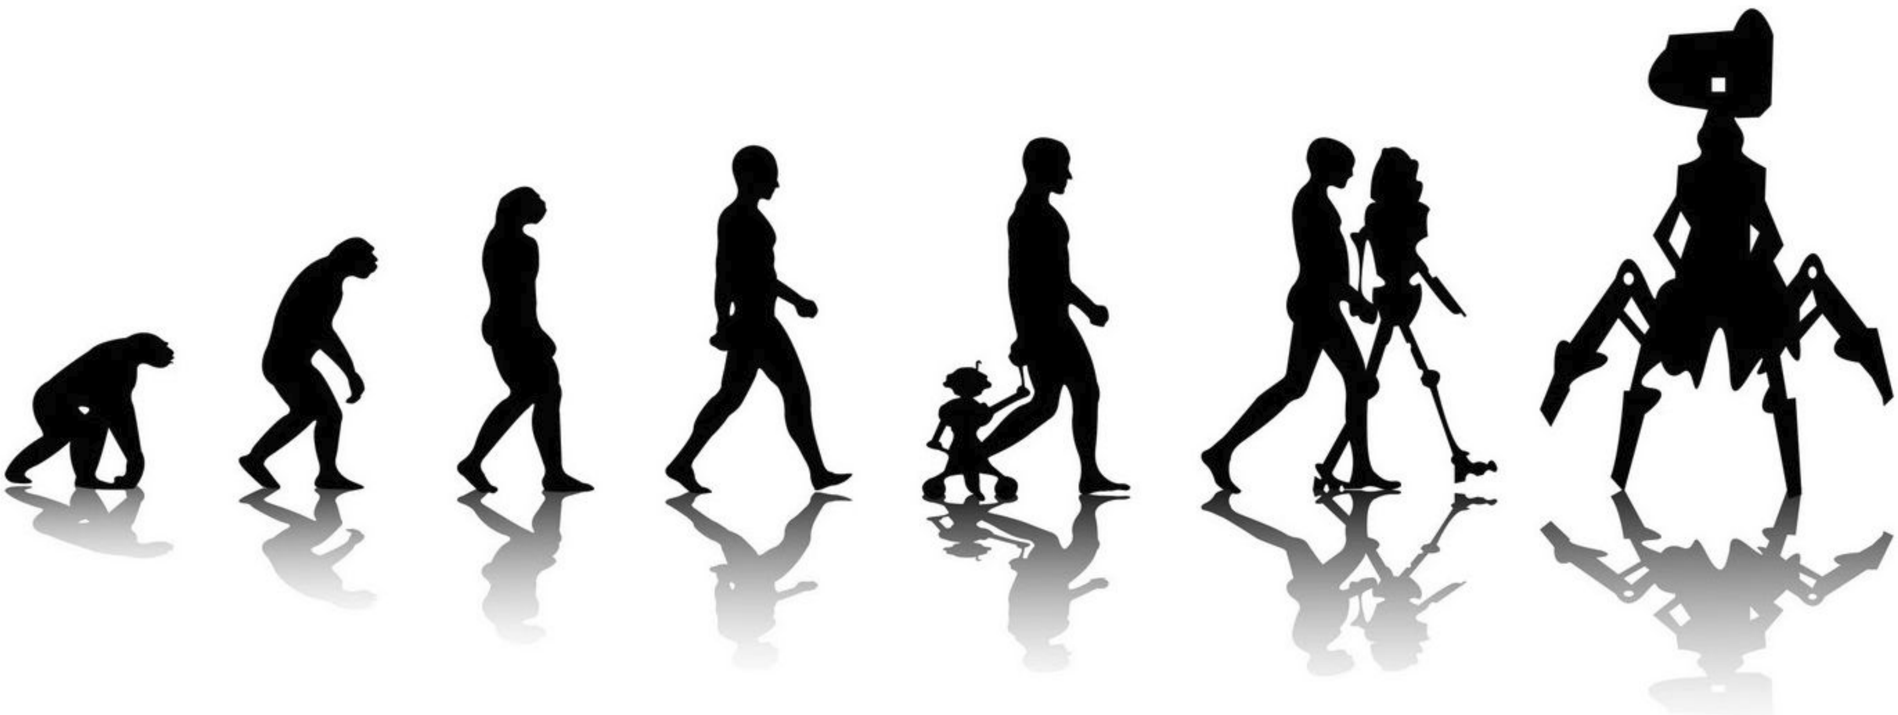

Modern Robotics: Evolutionary Robotics

COSC 4560 / COSC 5560

Professor Cheney

3/19/18

Learning and Evolution

Evolutionary Robotics

The Biology, Intelligence, and Technology of Self-Organizing Machines

Stefano Nolfi *and* **Dario Floreano**

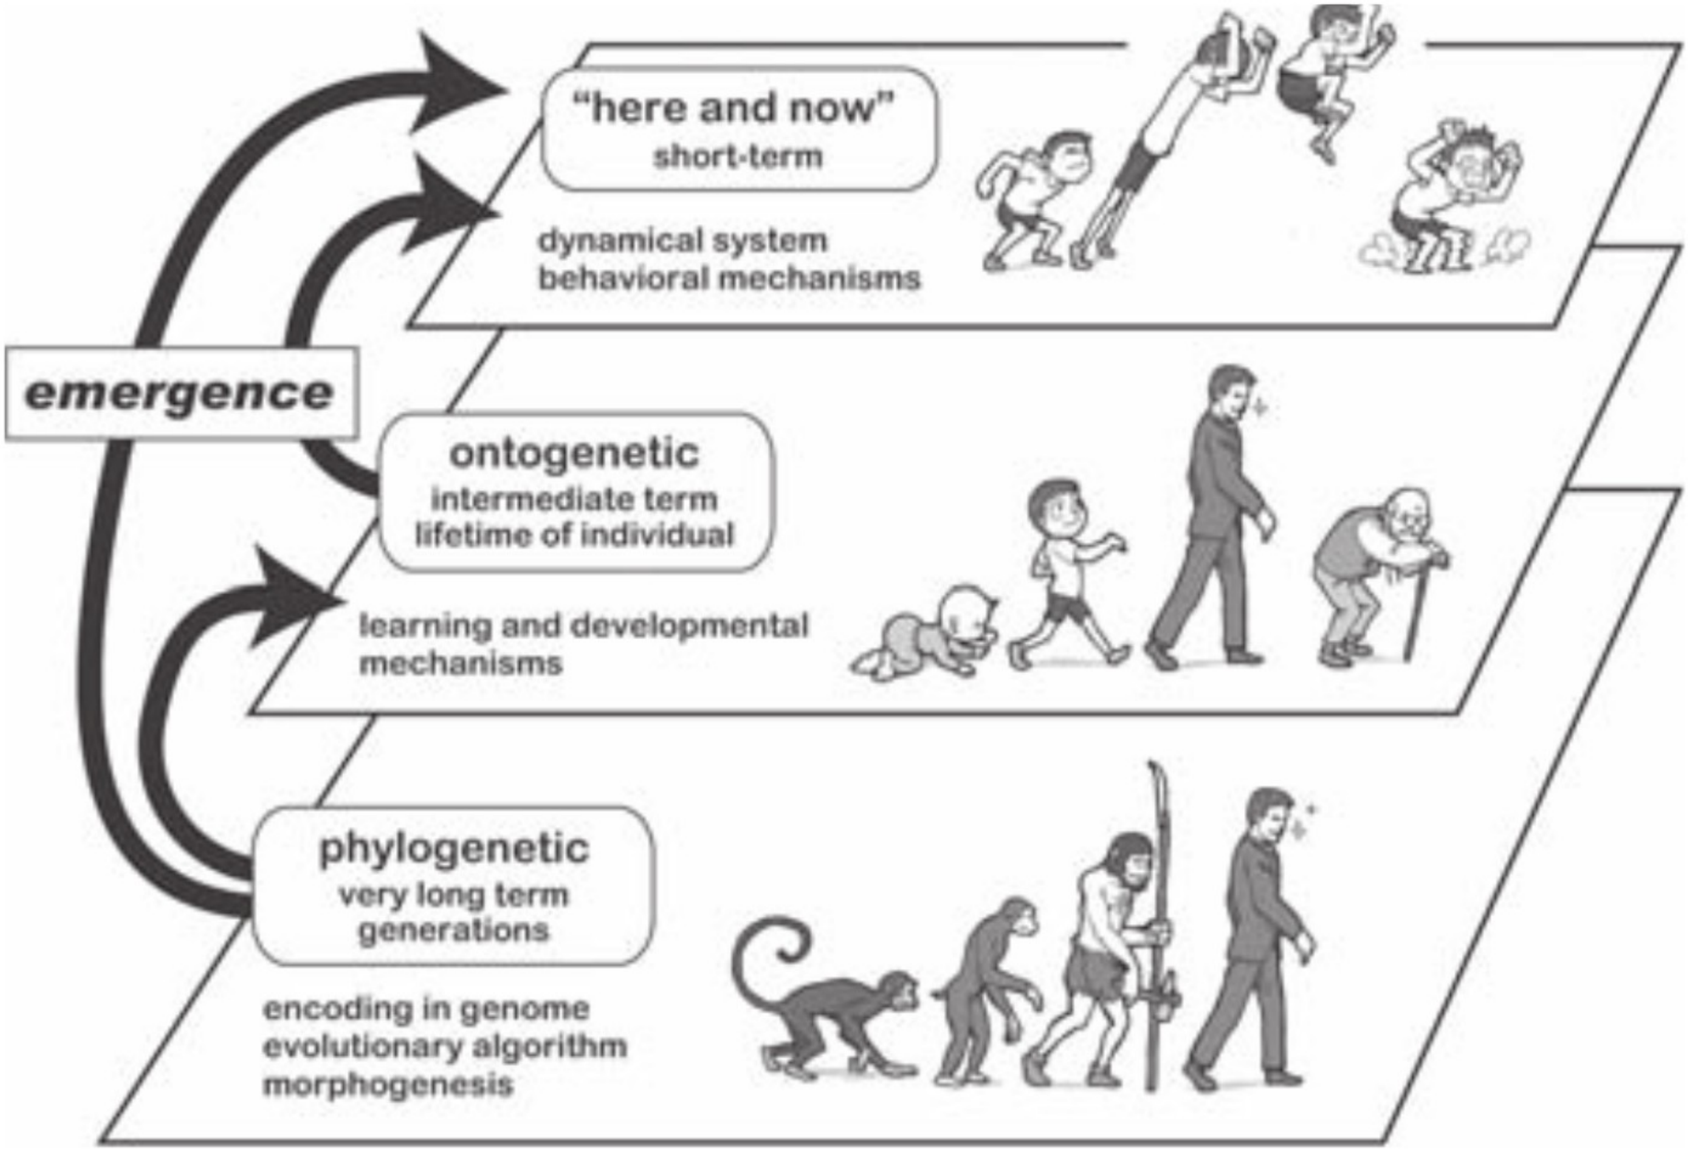
"here and now"
short-term
dynamical system
behavioral mechanisms
emergence
ontogenetic
intermediate term
lifetime of individual
learning and developmental
mechanisms
phylogenetic
very long term
generations
encoding in genome
evolutionary algorithm
morphogenesis

Hebbian Learning

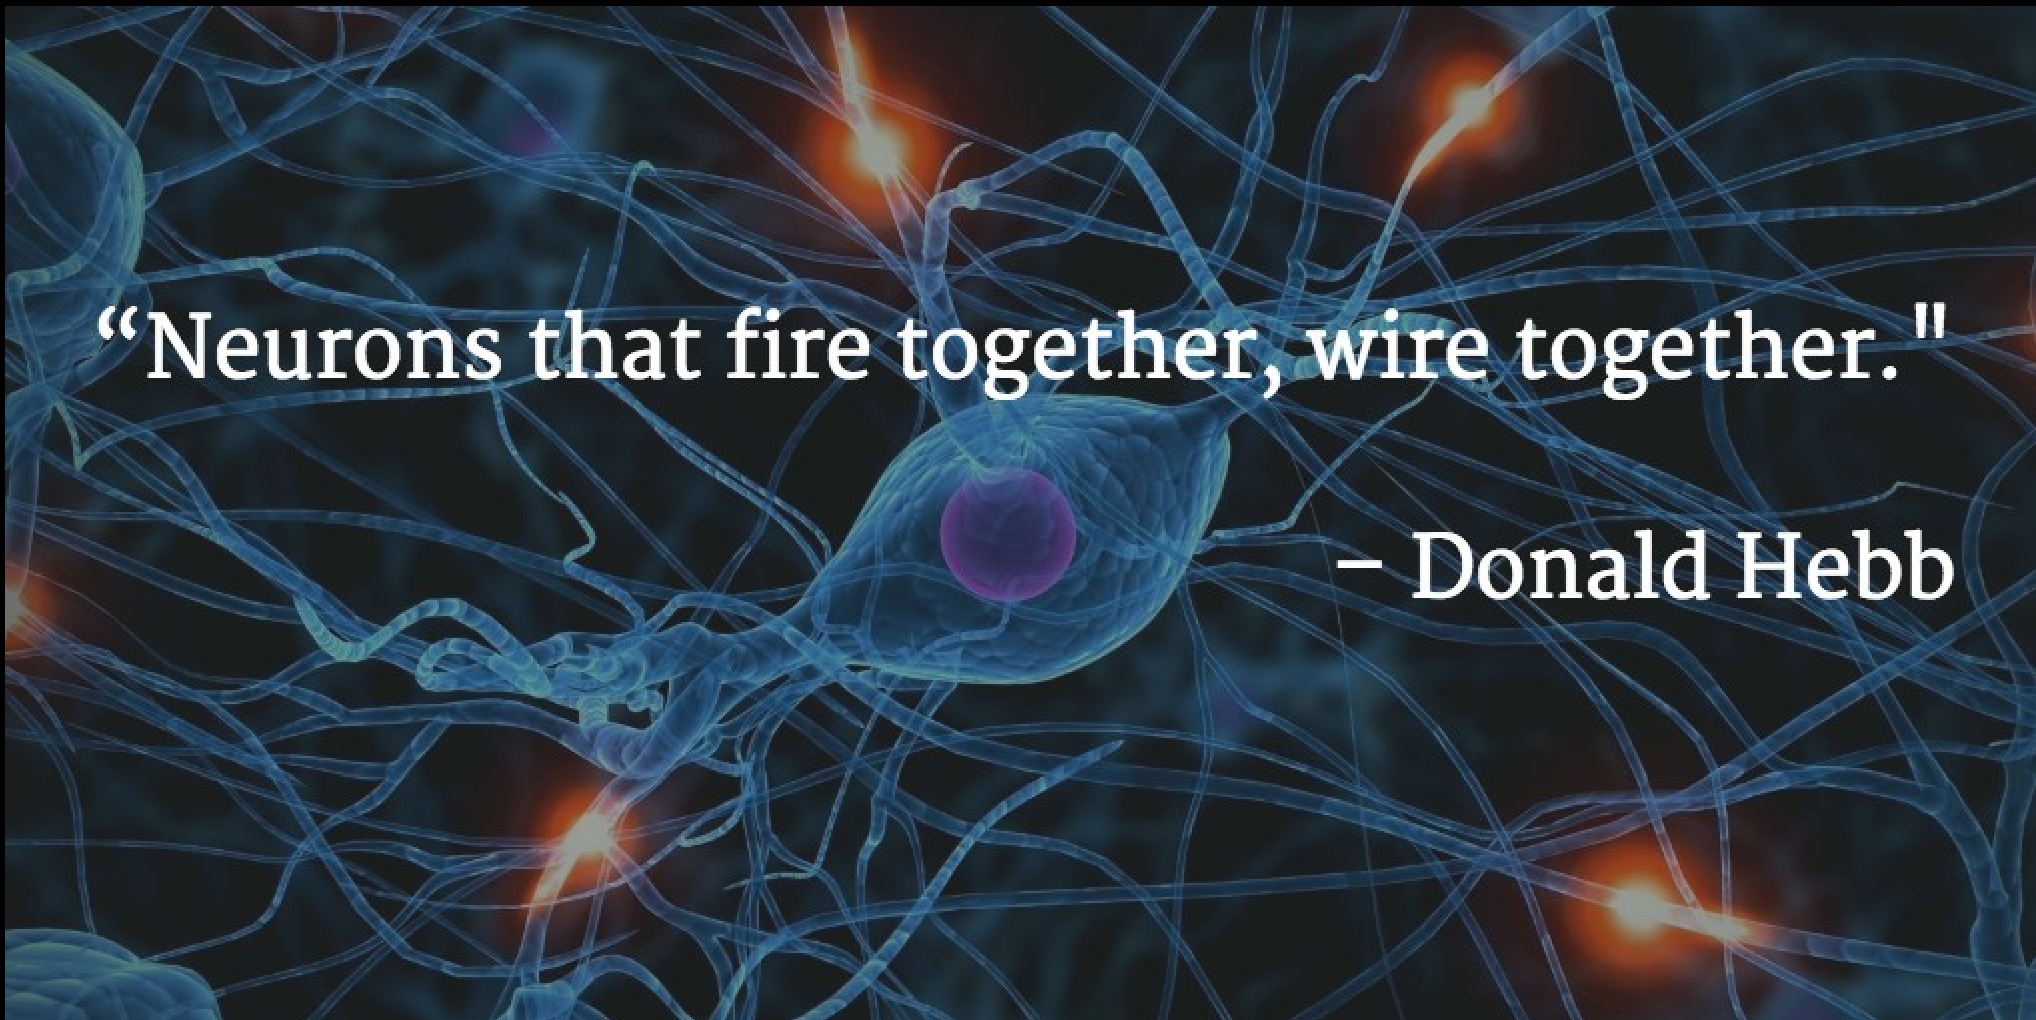
“Neurons that fire together, wire together."
– Donald Hebb

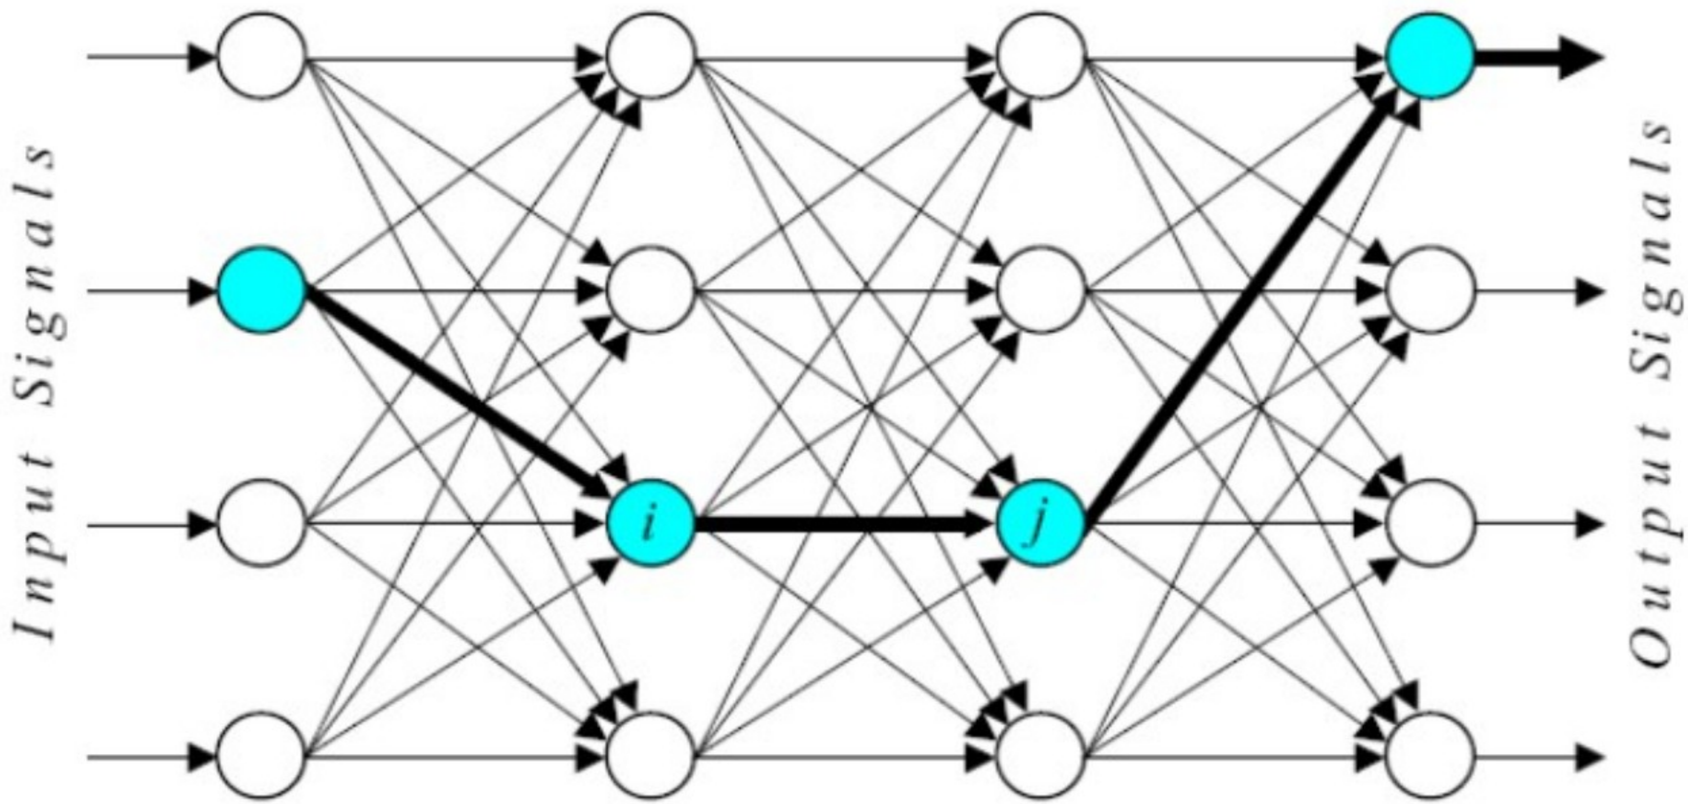

$$\Delta w_i = \eta x_i y$$

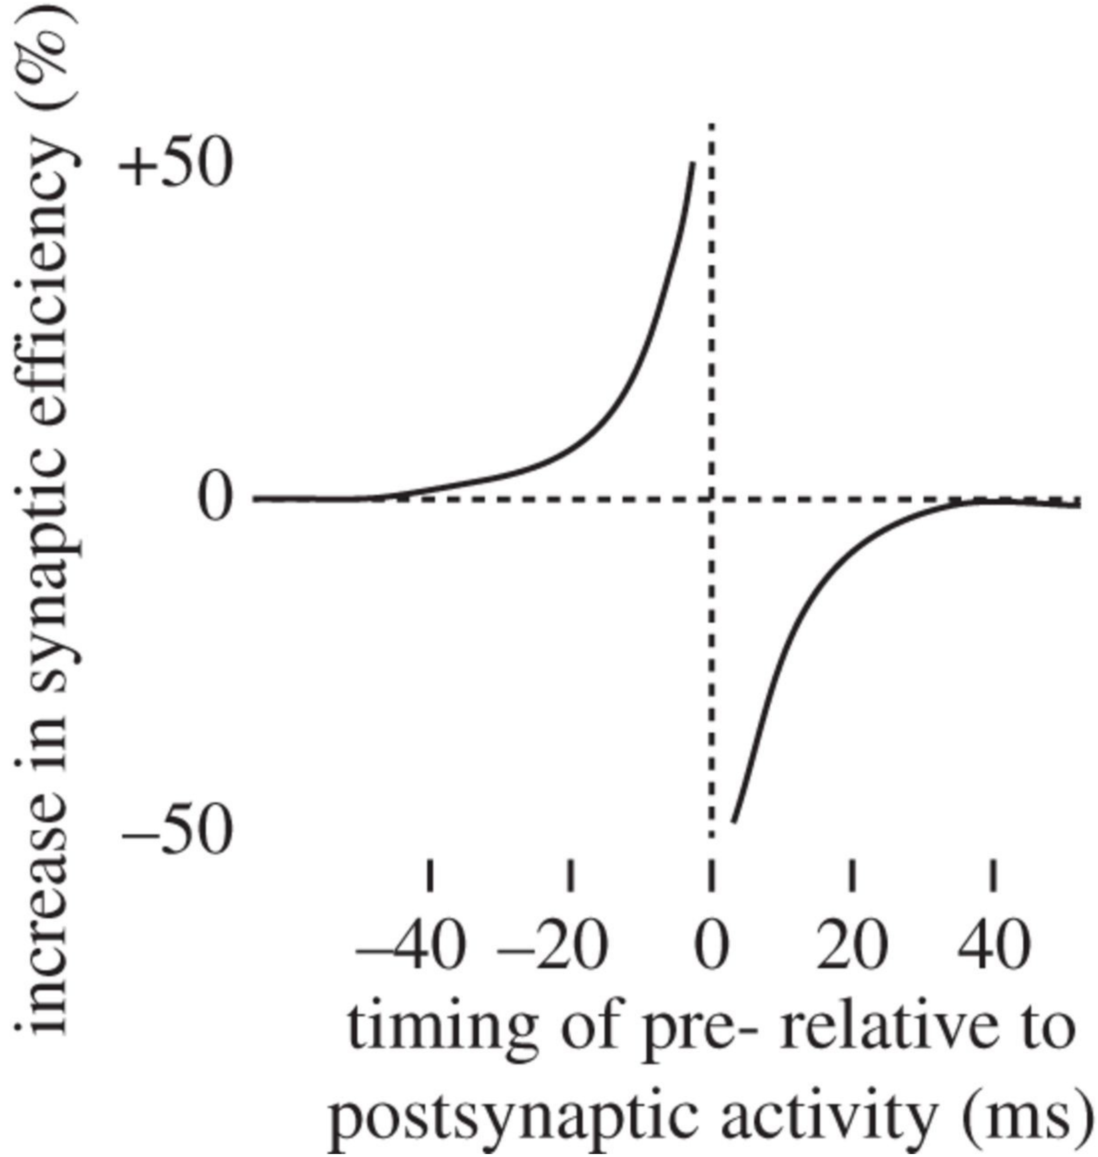
increase in synaptic efficiency (%)
+50
0
–50
–40
–20
0
20
40
timing of pre- relative to
postsynaptic activity (ms)

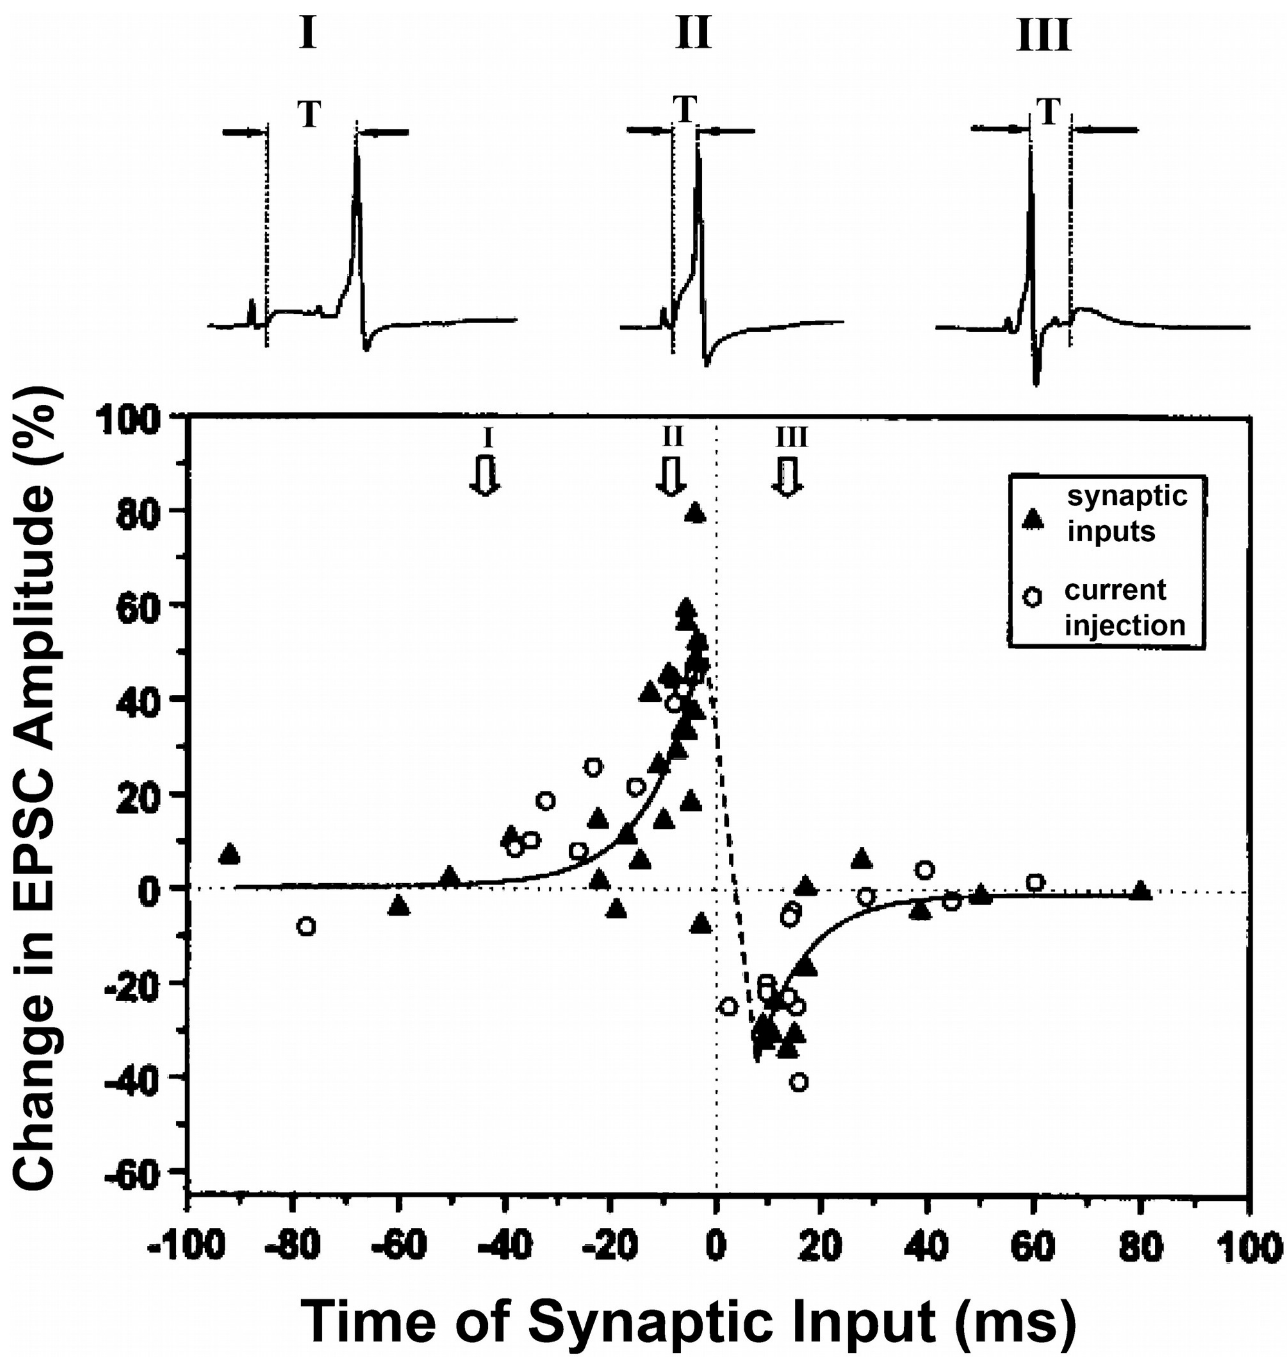

I
II
III
T
T
T
Change in EPSC Amplitude (%)
100
80
60
40
20
0
-20
-40
-60
I
II
III
synaptic inputs
current injection
-100
-80
-60
-40
-20
0
20
40
60
80
100
Time of Synaptic Input (ms)

Evolution of Plastic Control Networks

DARIO FLOREANO
Institute of Robotic Systems, Swiss Federal Institute of Technology (EPFL), Lausanne, Switzerland

JOSEBA URZELAI
ELCA Informatique SA, Lausanne, Switzerland

Evolutionary Robots with On-line Self-Organization and Behavioral Fitness

Dario Floreano[1] and Joseba Urzelai[2]

The genetic string encodes the architecture and a set of Hebbian rules, but *not* the synaptic strengths. Every time a genotype is decoded into a neural controller, its synaptic values are *randomly initialized* (always, from the first to the last generation) and are let free to adapt for ever using the genetically-specified Hebbian rules while the robot operates in the environment.

1) *Plain Hebb* rule increments the sinaptic strength proportionally to the correlated activity of the pre- and postsynaptic neurons; 2) *Postsynaptic rule* changes the synaptic strength only if the postsynaptic neuron is active, incrementing it if the presynaptic neuron is active, otherwise decrementing it; 3) *Presynaptic rule* is similar to postsynaptic, but changes the synaptic strength only if the pre-synaptic neuron is active; 4) *Covariance rule* increments and decrements synaptic strength depending whether the correlated activity of the two neurons is above or below a threshold. Details of the genetic encoding and Hebbian rules are to be found in Floreano and Urzelai (2000).

1. *Plain Hebb rule*: can only strengthen the synapse proportionally to the correlated activity of the pre- and post-synaptic neurons.

$$\Delta w = (1 - w)\, xy \tag{1}$$

2. *Postsynaptic rule*: behaves as the plain Hebb rule, but in addition it weakens the synapse when the postsynaptic node is active but the presynaptic is not.

$$\Delta w = w\,(-1 + x)\, y + (1 - w)\, xy \tag{2}$$

3. *Presynaptic rule*: weakening occurs when the presynaptic unit is active but the postsynaptic is not.

$$\Delta w = wx\,(-1 + y) + (1 - w)\, xy \tag{3}$$

4. *Covariance rule*: strengthens the synapse whenever the difference between the activations of the two neurons is less than half their maximum activity, otherwise the synapse is weakened. In other words, this rule makes the synapse stronger when the two neurons have similar activity and makes it weaker otherwise.

$$\Delta w = \begin{cases} (1 - w)\mathcal{F}(x, y) & \text{if } \mathcal{F}(x, y) > 0 \\ (w)\mathcal{F}(x, y) & \text{otherwise} \end{cases} \tag{4}$$

where $\mathcal{F}(x, y) = \tanh(4(1 - |x - y|) - 2)$ is a measure of the difference between the presynaptic and postsynaptic activity. $\mathcal{F}(x, y) > 0$ if the difference is bigger or equal to 0.5 (half the maximum node activation) and $\mathcal{F}(x, y) < 0$ if the difference is smaller than 0.5.

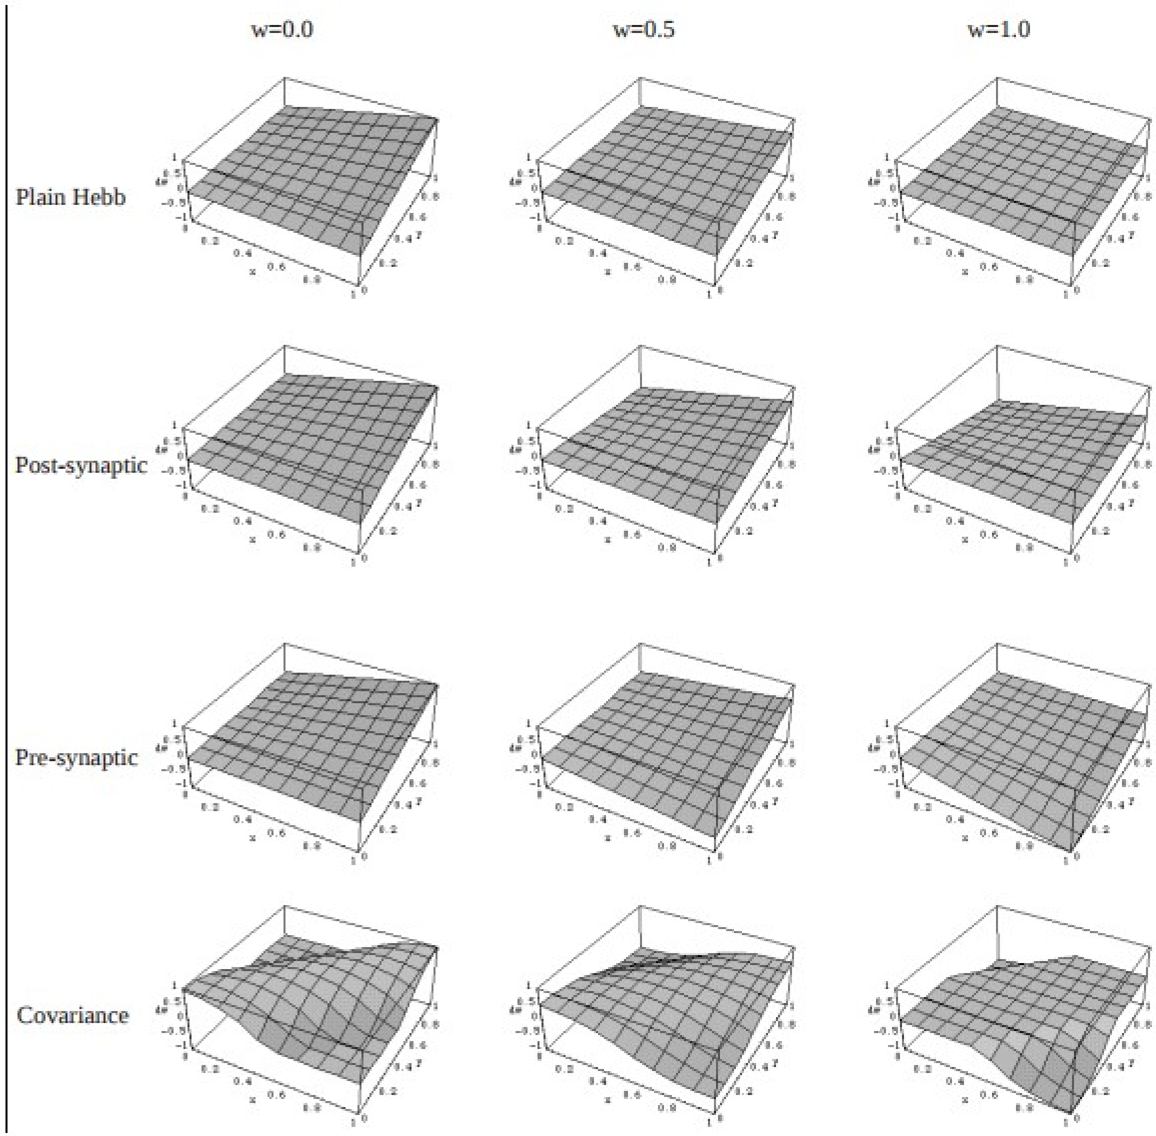

Figure 1: Synaptic change for each of the four Hebb rules. Notice that this is the *amount of change* Δw added to the synapses, not the synaptic strength. Each graph indicates the amount of change as a function of presynaptic x and postsynaptic y activity. The amount of change also depends on the current strength w of the synapse so that synapses are always bound between 0 and 1. Three graphs are shown for each rule, in the case of current strength 0.0, 0.5, and 1.0 respectively.

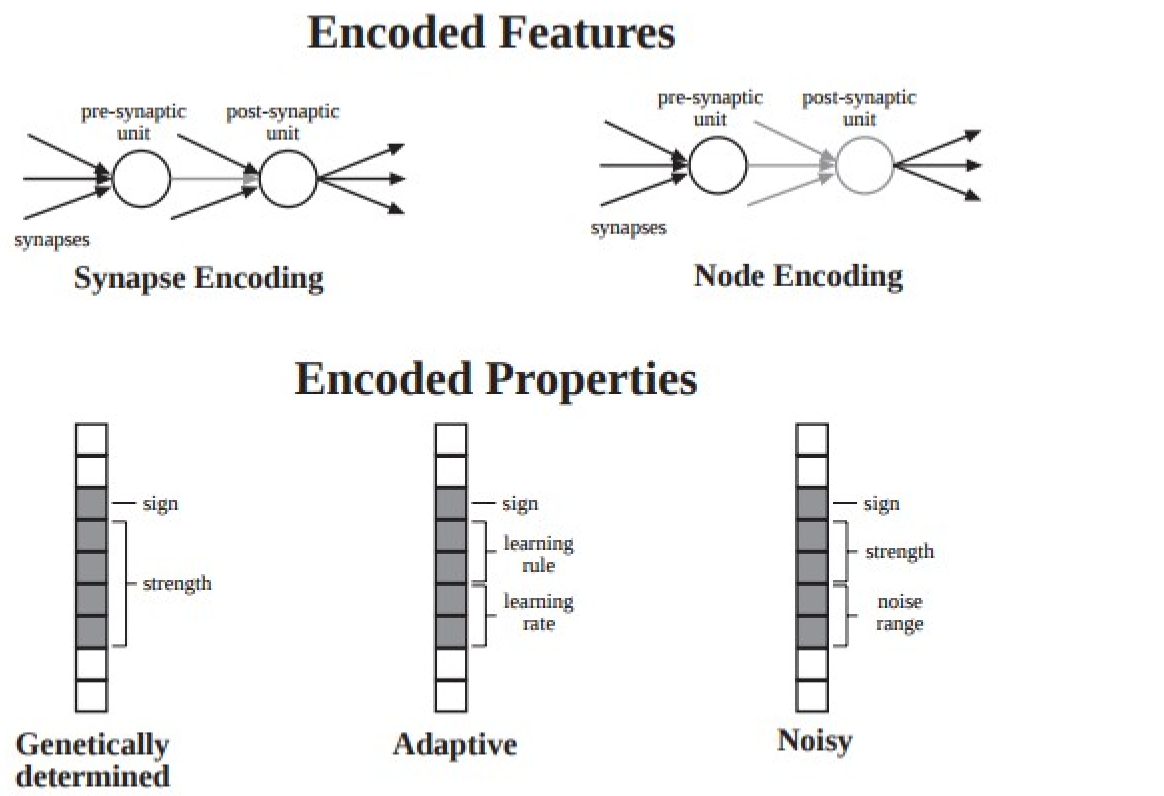

Figure 2: Different type of Genetic Encoding. **Top**: Two feature levels. The genetic string can either encode the properties of each individual synapse in the network (Synapse Encoding) or encode the properties of each node only (Node Encoding). In the latter case, the encoded properties are applied to all incoming synapses to that node. Node Encoding results in shorter genetic strings. **Bottom**: Three types of properties. Genetically-determined properties specify the connection sign and strengths of synapses. Adaptive properties specify the sign, the adaptation rule, and the learning rule of the synapses. Noisy properties specify the sign, weight strength, and a noise range that is continuously applied to the synapse. Properties are applicable to both Synapse and Node encoding, but in the latter case all incoming synapses will be have the same properties.

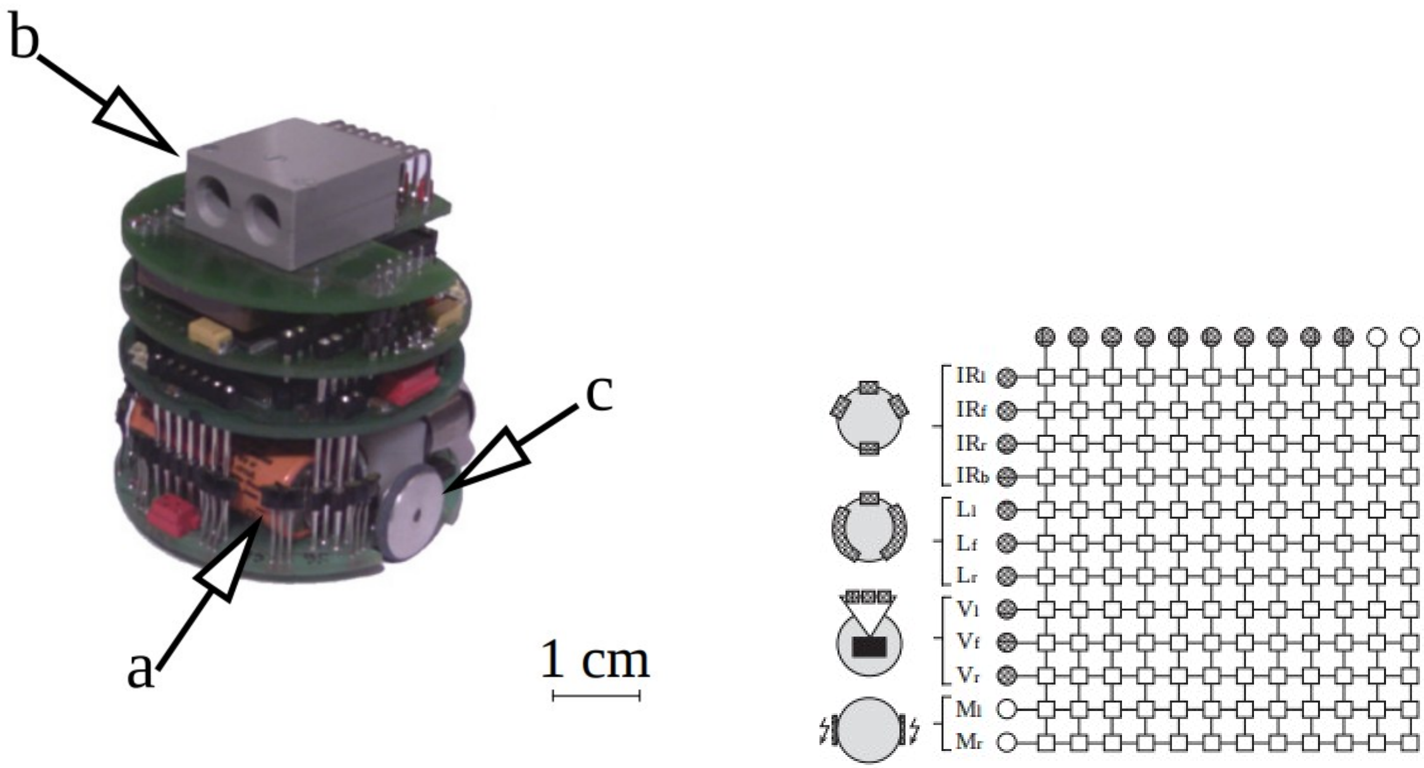

Figure 5: The neural controller is a fully-recurrent discrete-time neural network composed of 12 neurons giving a total of 12 x 12= 144 synapses (here represented as small squares of the unfolded network). 10 sensory neurons receive additional input from one corresponding pool of sensors positioned around the body of the robot shown on the left (l=left; r=right; f=front; b=back). $\vec{IR}$=Infrared Proximity sensors; $\vec{L}$=Ambient Light sensors; $\vec{V}$=vision photoreceptors. Two motor neurons $\vec{M}$ do not receive sensory input; their activation sets the speed of the wheels ($M_i > 0.5$ forward rotation; $M_i < 0.5$ backward rotation)

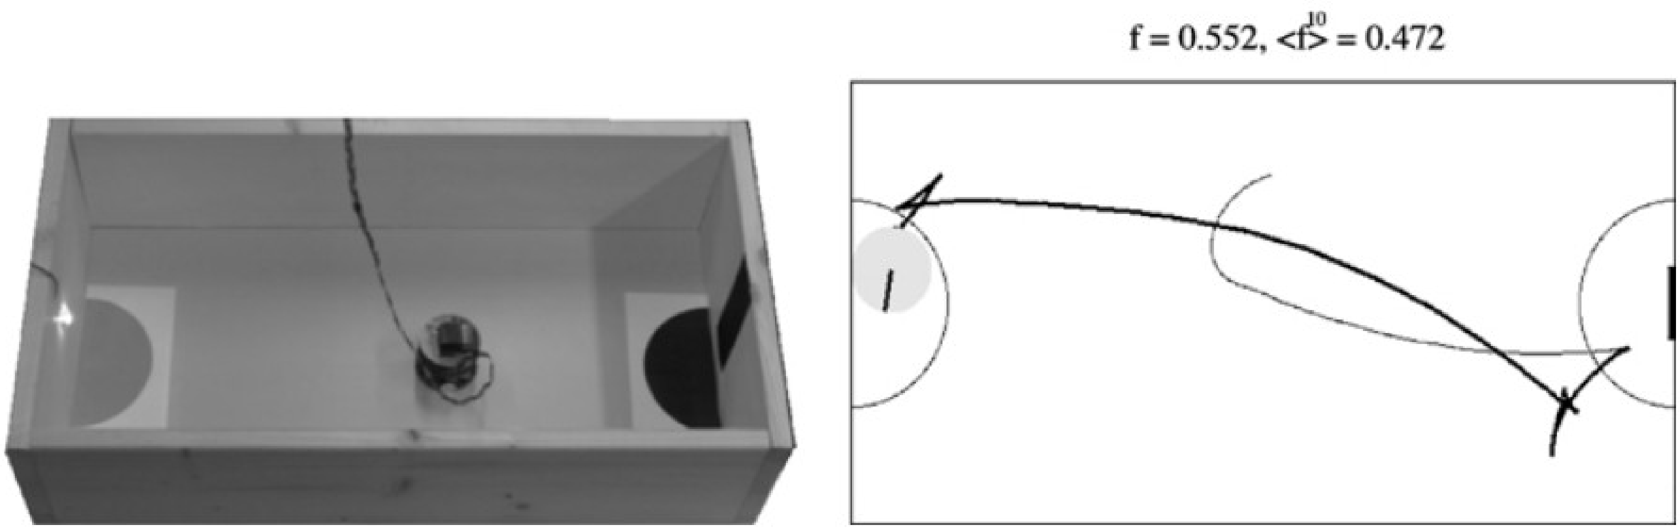

Figure 1. Left: A Khepera robot equipped with a vision module can gain fitness points only when it is sitting on the light (grey zone on the left) *when the light is on*. Initially the light is off, but the robot can switch it on by going over the black area on the right. No fitness points are given for the light switching behavior. Right: Trajectory of an evolved robot that adapts its connection strengths on the fly using the genetically-specified Hebbian rules. The figures on top show the fitness of this trajectory and the average fitness of this individual when tested ten times with different synaptic random initializations and starting positions.

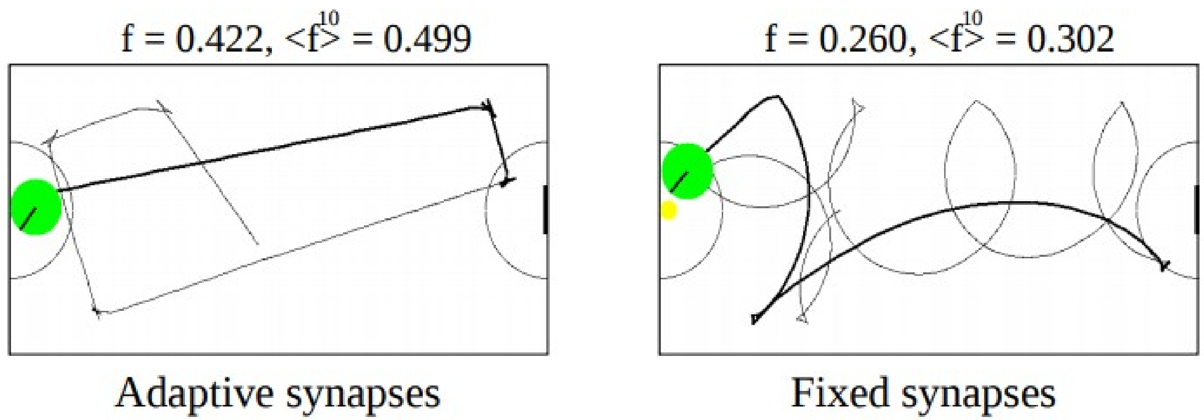

Figure 7: Behaviors of two best individuals (from last generation) with adaptive synapses and Node Encoding (*left*) and with genetically-determined synapses and Synapse Encoding (*right*). When the light is turned on, the trajectory line becomes thick. The corresponding fitness value is printed on the top of each box along with the average fitness of the same individual tested ten times from different positions and orientations.

Node Encoding and adaptive synapses

Synapse Encoding and genetically-determined synapses

50 individuals

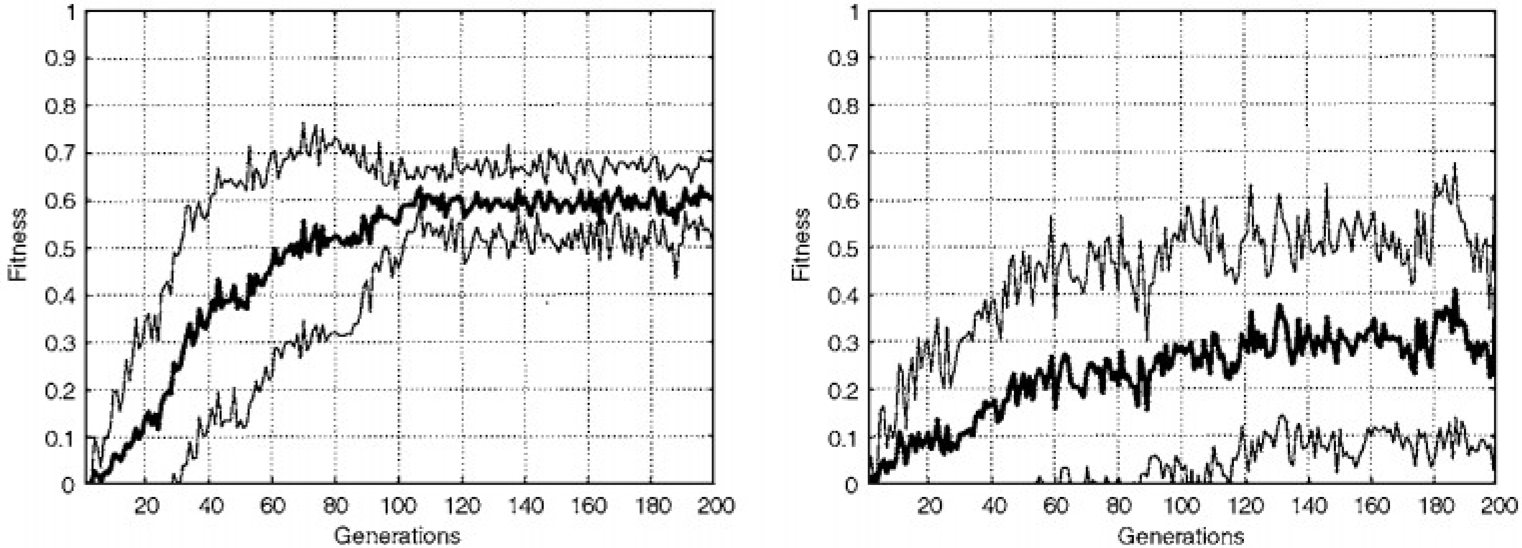

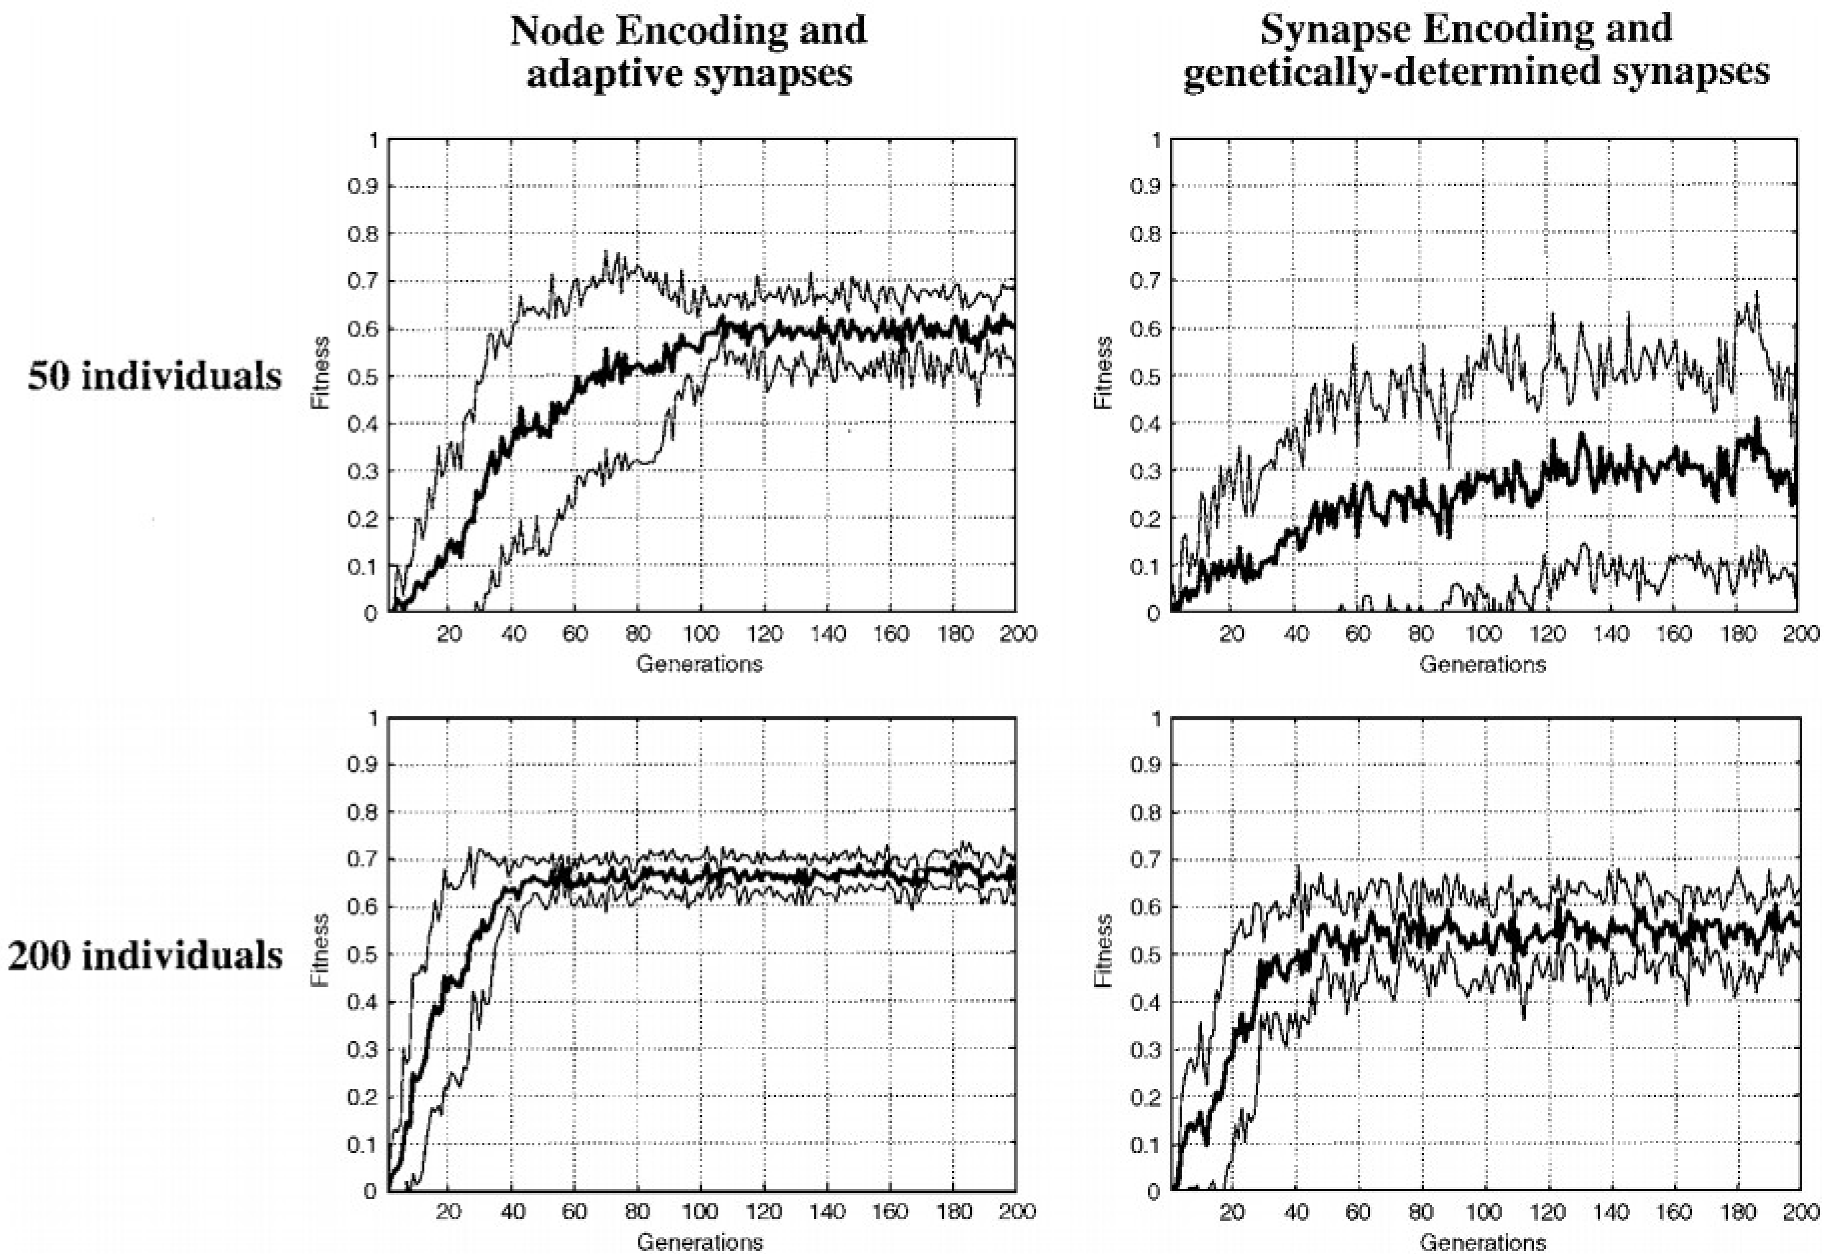
Node Encoding and adaptive synapses
Synapse Encoding and genetically-determined synapses
50 individuals
200 individuals
Fitness
Generations
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
20
40
60
80
100
120
140
160
180
200

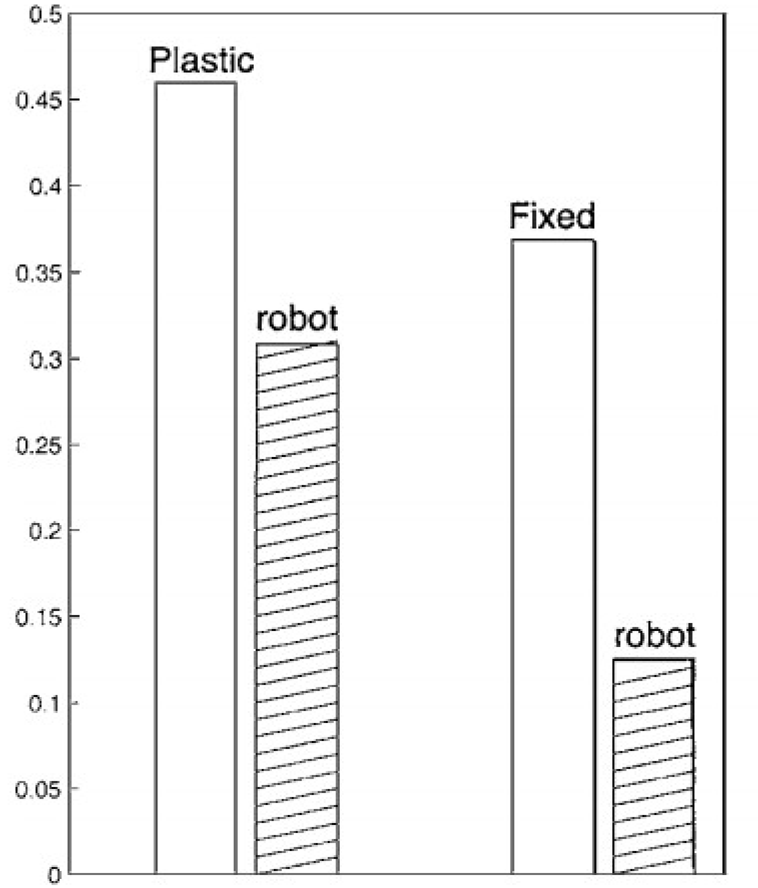

Figure 3. Left: Performance of best controller evolved in simulation and transferred to the physical robot for the two conditions: plastic networks and genetically-determined and fixed networks. Each bar is the average of 10 tests of the best individuals of 10 evolutionary runs

Phenotypic Plasticity in Evolving Neural Networks

Stefano Nolfi[+] **Orazio Miglino**[*] **Domenico Parisi**[+]
+Institute of Psychology, National Research Council, Rome, Italy.
*Department of Psychology, University of Palermo, Italy

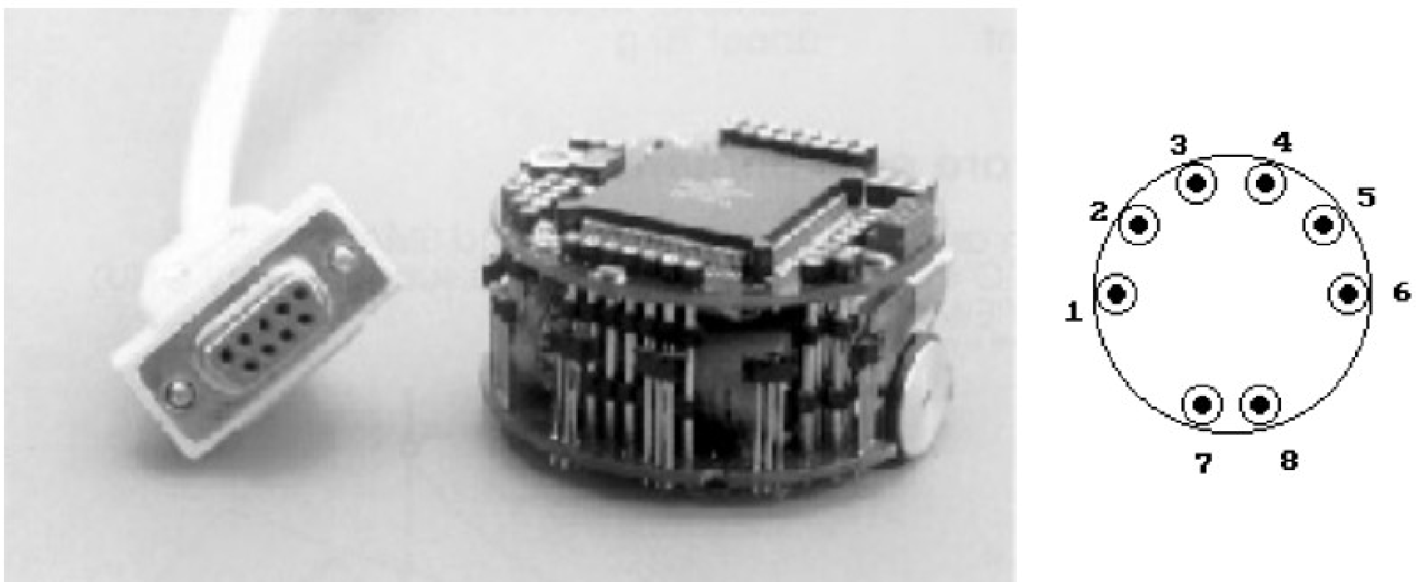

Figure 1. The picture on the left side shows Khepera, the miniature mobile robot. The picture on the right side shows how the 8 sensors are distributed on the robot's body. The large empty circles represents the ambient light sensors and the small full circles represents the obstacle detectors.

Our goal is to develop a neural network that controls an organism trying to reach a target area in its environment. The environment is a 60x35 cm arena surrounded by walls made of naturally colored wood. The organism should find as quickly as possible a target area of circular shape with a diameter of 2 cm which is randomly positioned in the arena. The target area may or may not be illuminated by a 1-watt lightbulb positioned right above the area.

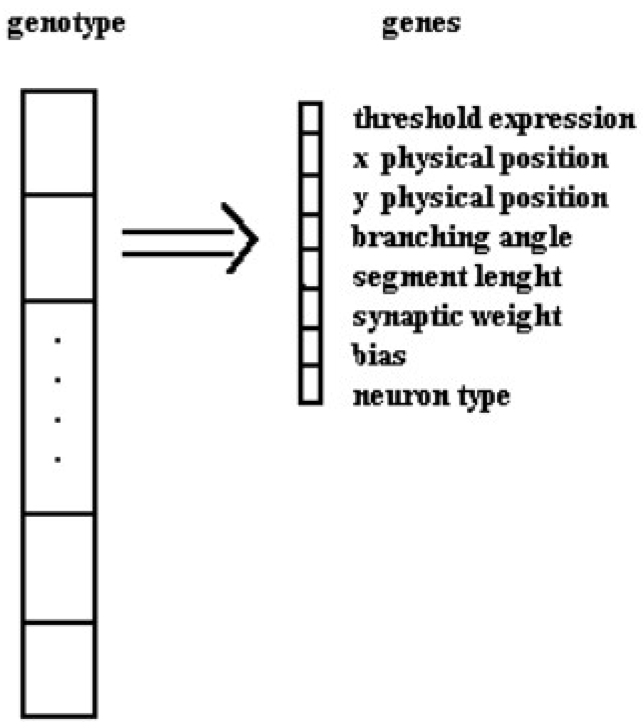

Figure 2. Developmental instructions specified in an organism's genotype.

Each organism includes both a genotype and a phenotypic neural network. The genotype contains developmental instructions for the construction of the phenotypic network. The genotype has a fixed length and is divided up into blocks, each block corresponding to a single neuron (Figure 2).

There are three types of neurons, sensory neurons, internal neurons, and motor neurons. Genotypes are 32 blocks in length, i.e. each organism can have a maximum number of 32 neurons in its nervous system. The first 10 blocks in the genotype correspond to sensory neurons, the last 5 blocks to motor neurons, and the 17 intermediate blocks to internal neurons. Internal neurons can be arranged in a maximum of 7 layers. The standard logistic function is used to compute the activation value of the neurons.

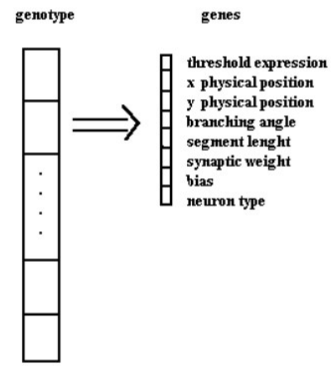

(a) The "expression threshold gene" determines at what condition a neuron is allowed to grow its axon and, therefore, to possibly establish connections with other neurons. For a neuron to grow its axon the variability in a sequence of 10 successive activation states of the neuron must exceed the threshold value specified by this gene. If the gene codifies a 0 value the corresponding neuron is allowed to grow immediately.

(b) The two "physical position genes" specify the Cartesian coordinates of the neuron in the bidimensional nervous system of the organism.

(c) The "branching angle gene" and the "segment length gene" determine the angle of branching of the neuron's axon and the length of the branching segments. While angle of branching and segment length can vary from one neuron to another in the same network, all the branchings of the axon of a given neuron have the same branching angle and the same length. The growing axon of all neurons branches a fixed number of times (five) with exclusively binary branchings. In order to obtain feedforward networks without recurrent connections, the branching angles are so restricted that all units send their growing axons in the direction from input units (bottom layer) to output units (top layer).

(d) The "synaptic weight gene" determines the synaptic coefficients of all the connections that will

(e) The "bias gene" represents the activation bias (threshold) of the neuron itself.

(f) The "neuron type gene" specifies, in the case of a sensory neuron, to which of Khepera's 8 sensors the corresponding neuron should be clamped and if it should detect obstacles or the ambient light, and in the case of a motor neuron, if the corresponding motor neuron determines the speed level of the left or the right motor. If more than a single output neuron codifying the same motor information exist, the actual motor output is taken to be the average of the activation levels of the different neurons.

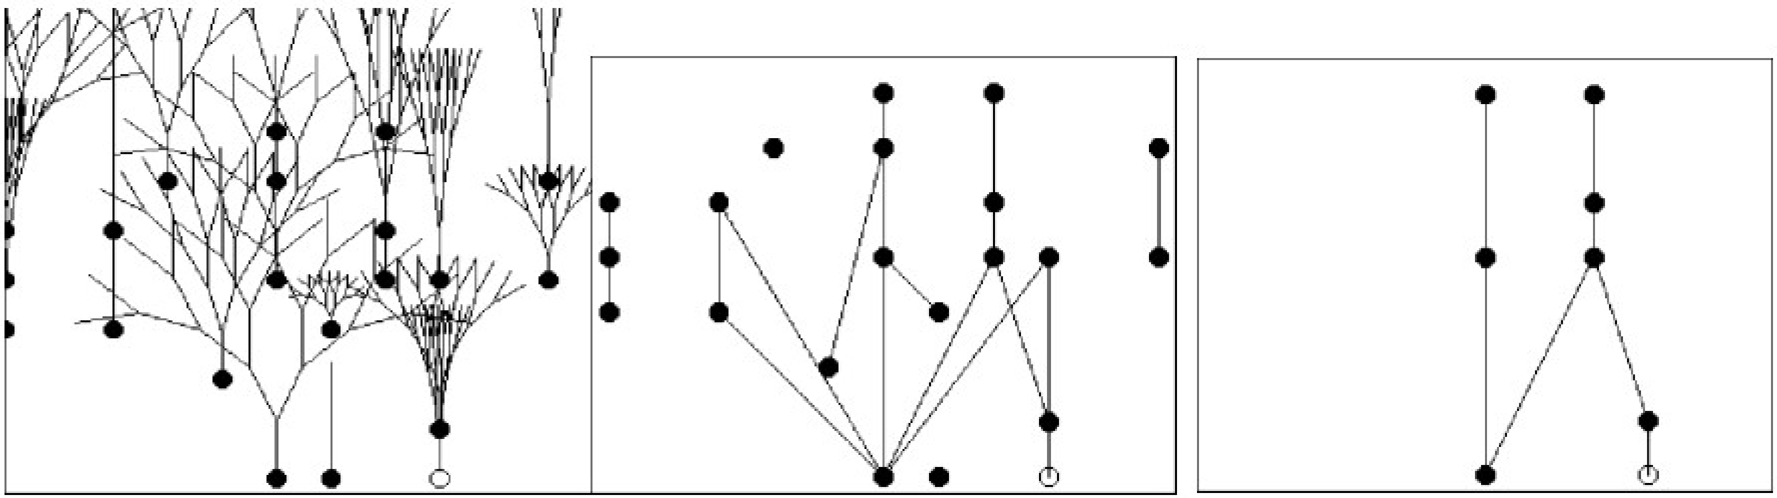

Figure 3. Development of the neural network of an evolved organism. The growing and branching process of axons is shown on the left side of the figure. The resulting neural network after removal of nonconnecting branches is shown in the center. The functional network after elimination of isolated (nonfunctional) neurons and groups of interconnected neurons is shown on the right side of the figure.

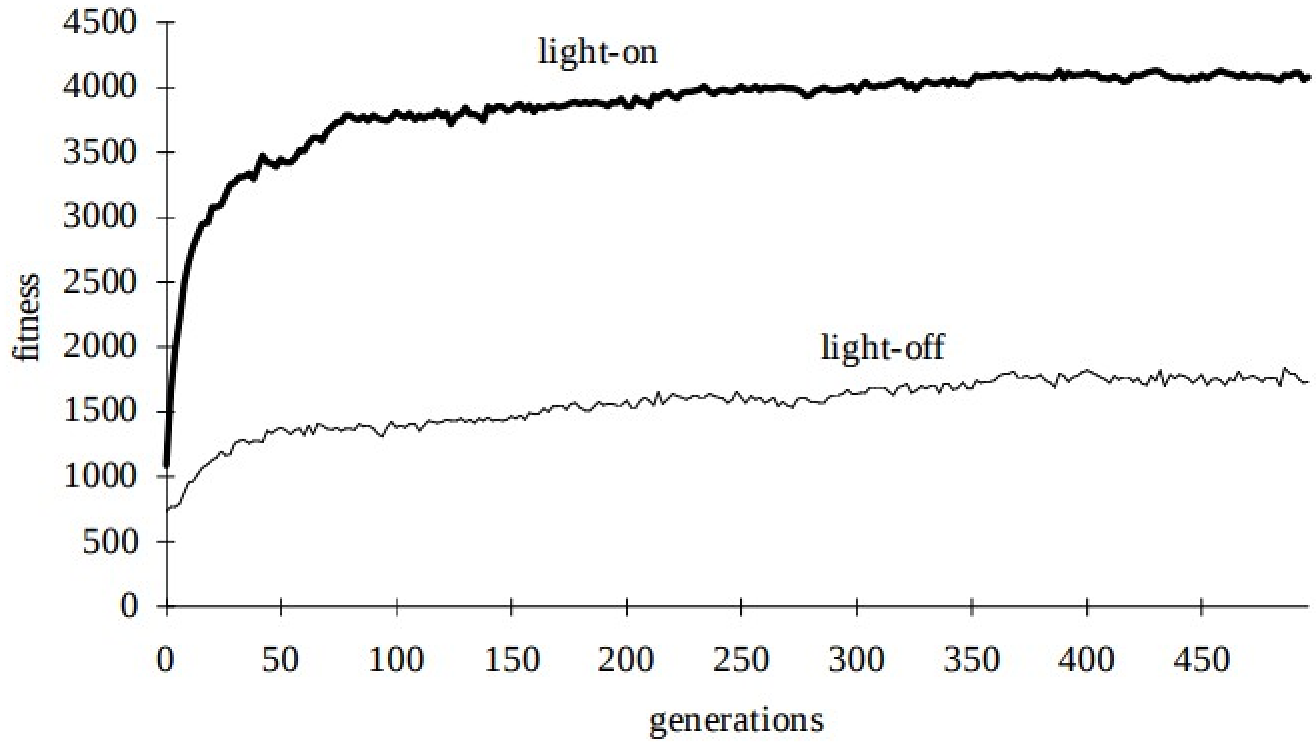

Figure 4. Fitness of the 20 best individuals in each of the 500 generations. The fat curve represents the performance of the even generations (those exposed to the "light" environment), the thin curve represents the performance of the odd generations (those exposed to "dark" environment). Borth curves are the average results of 10 simulations starting from initial populations with different randomly generated genotypes.

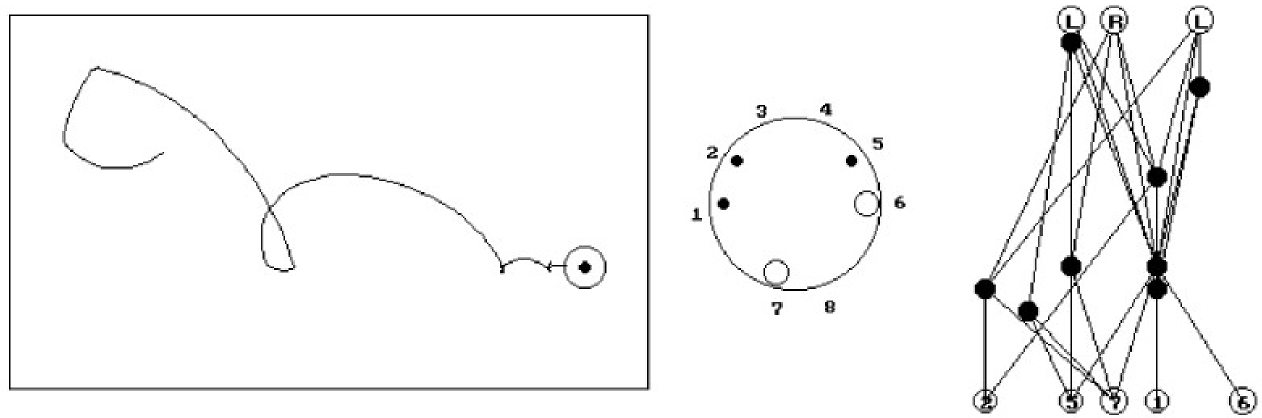

Figure 5. An evolved genotype developing in a "light" environment. The picture on the left side shows the organism's behavior. The environment and the target area are represented as a rectangular box and as a circle, respectively. The picture shows the organism's trajectory from the starting position to the target area. The central picture represents the sensors that have developed in the organism. Large empty circles represents light detectors and small full circles represent wall (obstacle) detectors. Finally, the picture on the right side represents the organism's entire nervous systems. The bottom layer circles labeled with progressive numbers represents sensory neurons, the upper layer circles labeled with "L" and "R" represents the left and right motor neurons, and the remaining layers represent internal neurons.

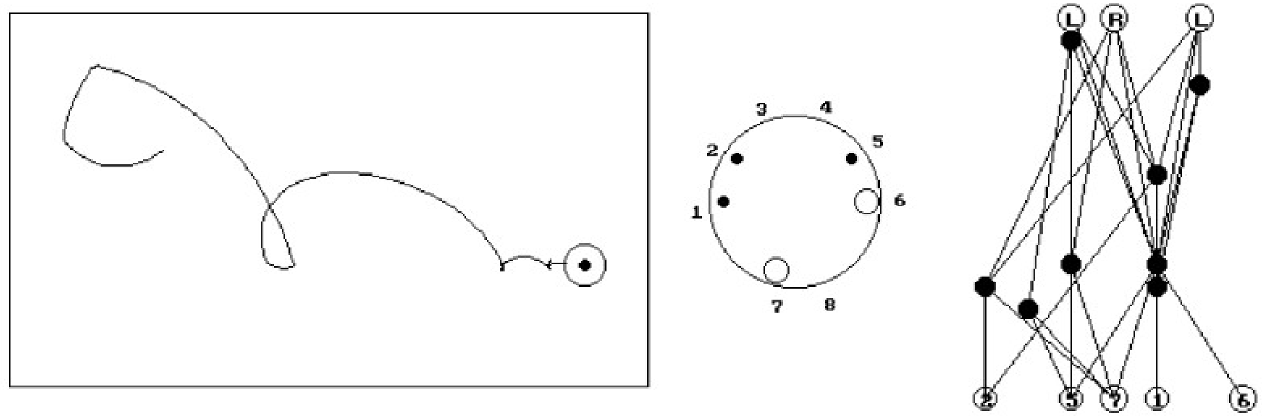

Figure 5. An evolved genotype developing in a "light" environment. The picture on the left side shows the organism's behavior. The environment and the target area are represented as a rectangular box and as a circle, respectively. The picture shows the organism's trajectory from the starting position to the target area. The central picture represents the sensors that have developed in the organism. Large empty circles represents light detectors and small full circles represent wall (obstacle) detectors. Finally, the picture on the right side represents the organism's entire nervous systems. The bottom layer circles labeled with progressive numbers represents sensory neurons, the upper layer circles labeled with "L" and "R" represents the left and right motor neurons, and the remaining layers represent internal neurons.

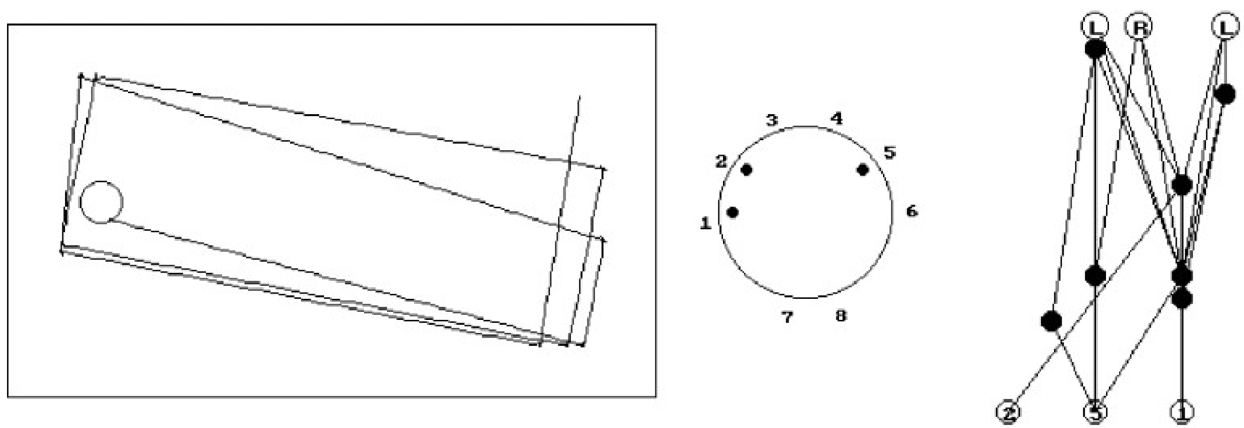

Figure 6. The same genotype of Figure 5 developing in a "dark" environment.

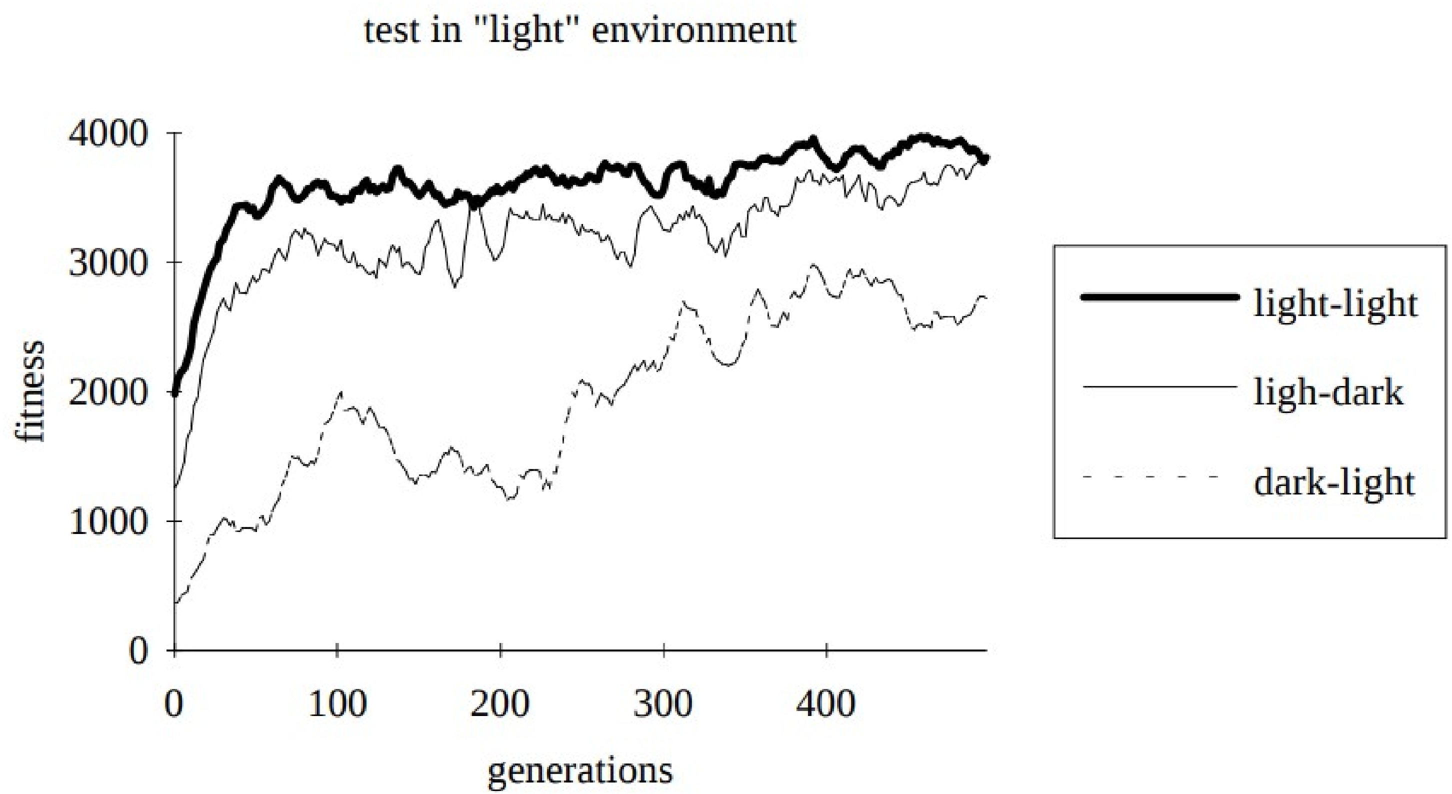

Figure 7. Fitness of the best organisms of each generation tested in a "light" environment for 10 epochs after completion of development. The fat curve represents the performance of organisms that have been selected in a "light" environment (i.e. organisms of even generations) and have developed in the same type of environment (light-light). The thin curve represents the performance of organisms that have been selected in a "dark" environment (i.e. organisms of odd generations) and have developed in a "light" environment (dark-light). The dashed curve represents the performance of organisms that have been selected in a "light" environment (i.e. organisms of even generations) and have developed in a "dark" environment (light-dark). Each curve represents the 7 point moving average of 10 different simulations.

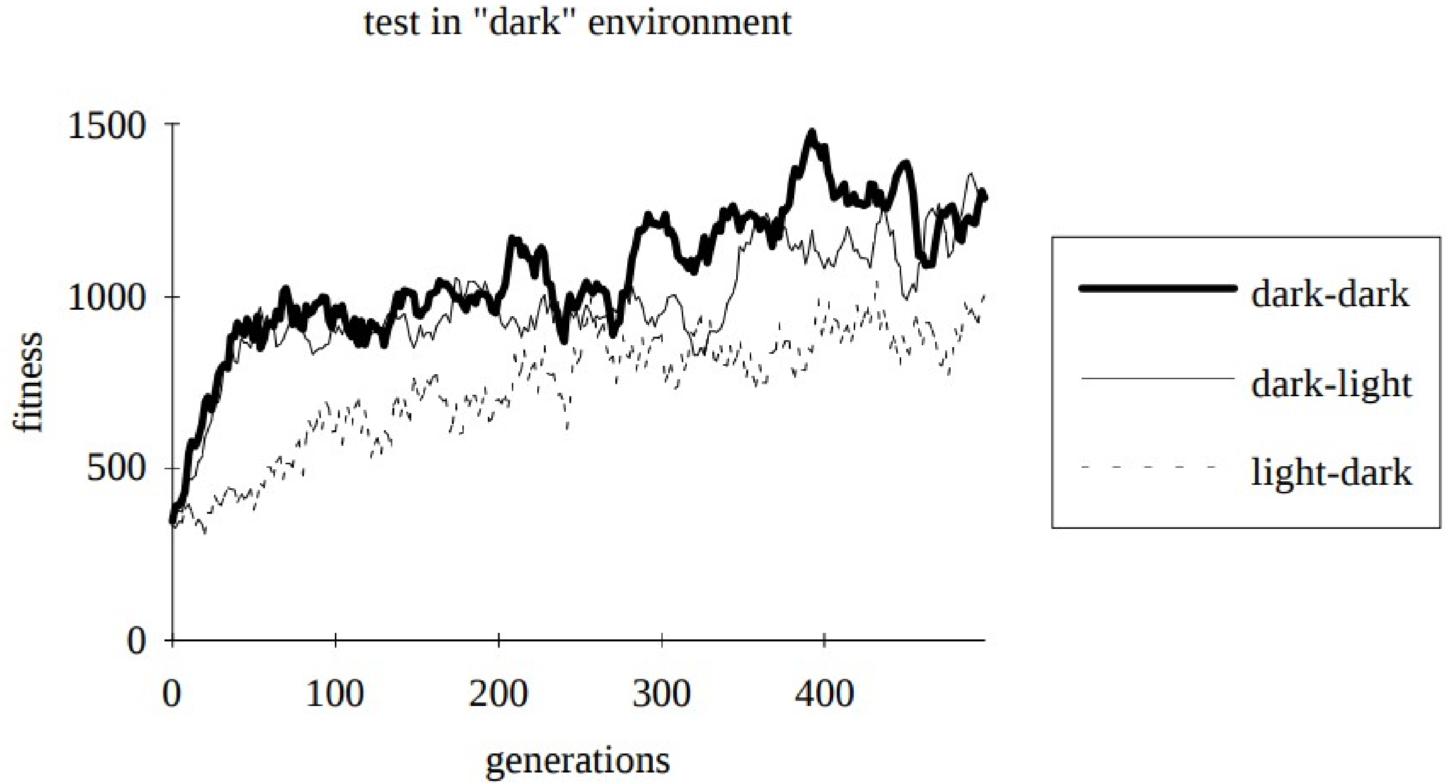

Figure 8. Fitness of the best organisms of each generation tested in a "dark" environment for 10 epochs after completion of development. The fat curve represents the performance of the organisms that have been selected in a "dark" environment and have developed in the same type of environment (dark-dark). The thin curve represents the performance of organisms that have been selected in a "light" environment and have developed in a "dark" environment (light-dark). The dashed curve represents the performance of organisms that have been selected in a "dark" environment and have developed in a "light" environment (dark-light). Each curve represents the 7 point moving average of 10 different simulations.