Learning and Evolution

Stefano Nolfi / Dario Floreano

- Evolution: slow, affects genotype
- Learning: fast, affects phenotype

reasons for studying interaction:

a) yields performance advantages?

b) understanding it in biological organisms

Outline

- „How learning can guide evolution“
- „Evolving individuals that learn a task different from what they are selected for“
- „Exploiting the regularities of the environment through learning“
- „Adaption to fast-changing environements“
- „Co-evolution of plastic neurocontrollers for competing robots“

Learning helps evolution I

the first hypothesis:

- learning helps individuals to adapt to changes in the environment
- thereby channels evolutionary search: evolution might choose individuals who have the (to be acquired) skills at birth [Baldwin(1896)], because...
- learning has costs

Learning helps evolution II

20 genes [0,1 (, ?)] – one successful configuration.

everything happens rather randomly.

Learning individuals build smooth curve because of unspecified genes.

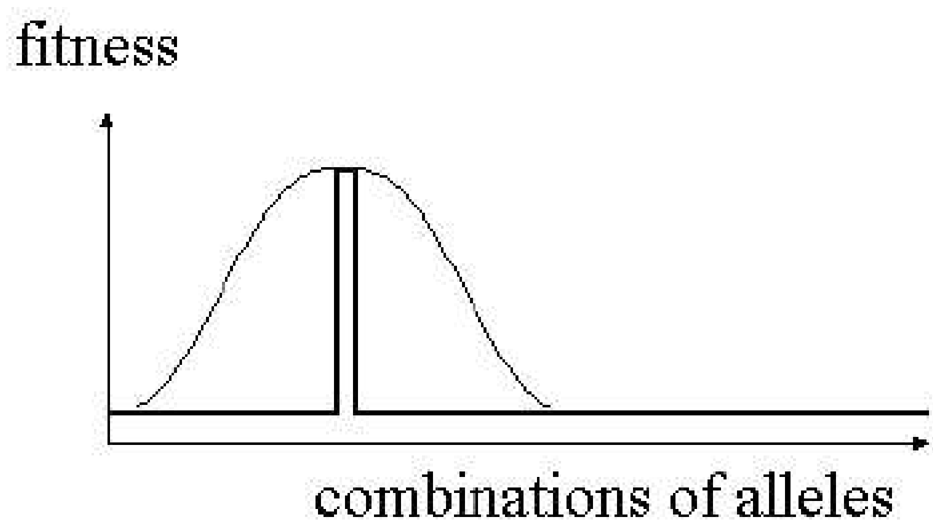

by Hinton and Nowlan (1987)

Learning helps evolution III

Another view: learning individuals moving in the search space increase probability for success: learning helps evolution to explore.

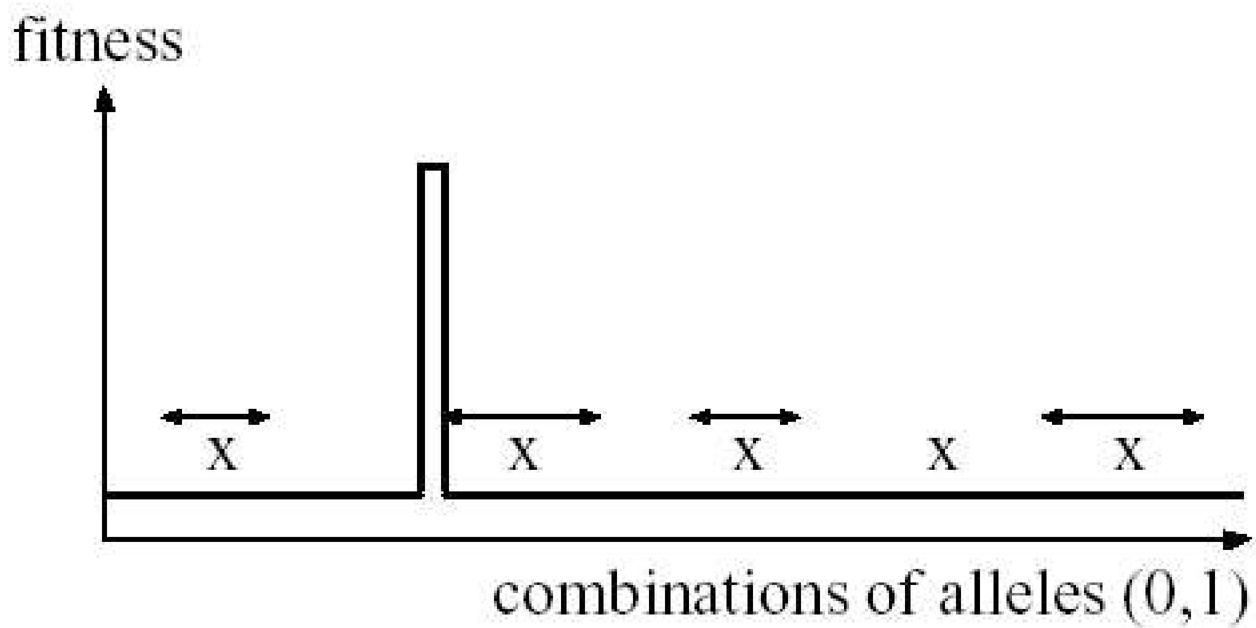

Experiment 1

by Hinton and Nowlan (1987)

different tasks I

- Evolution and learning have different tasks -> their search spaces are not the same (learning indirectly enhances fitness)
- here: predict good moves (learning) and eat (evolution)
- note: evolution and learning still operate on the same operators (strength of synapses)

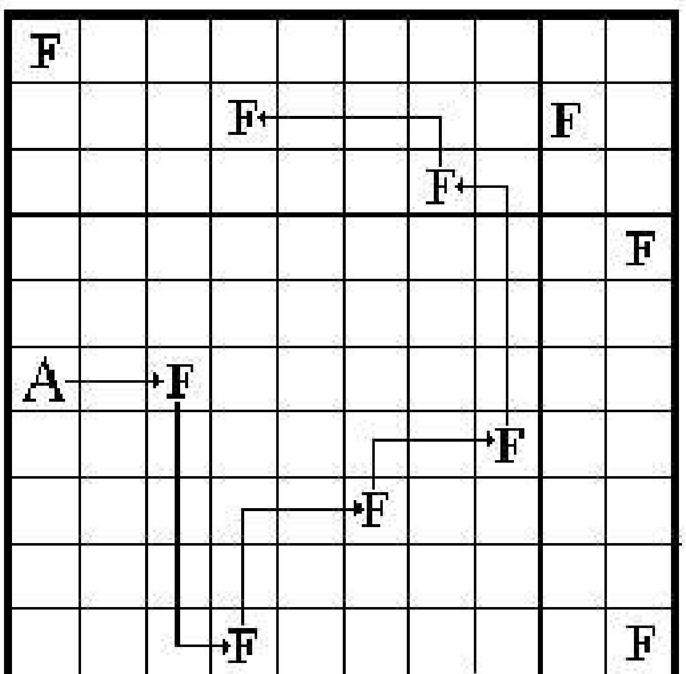

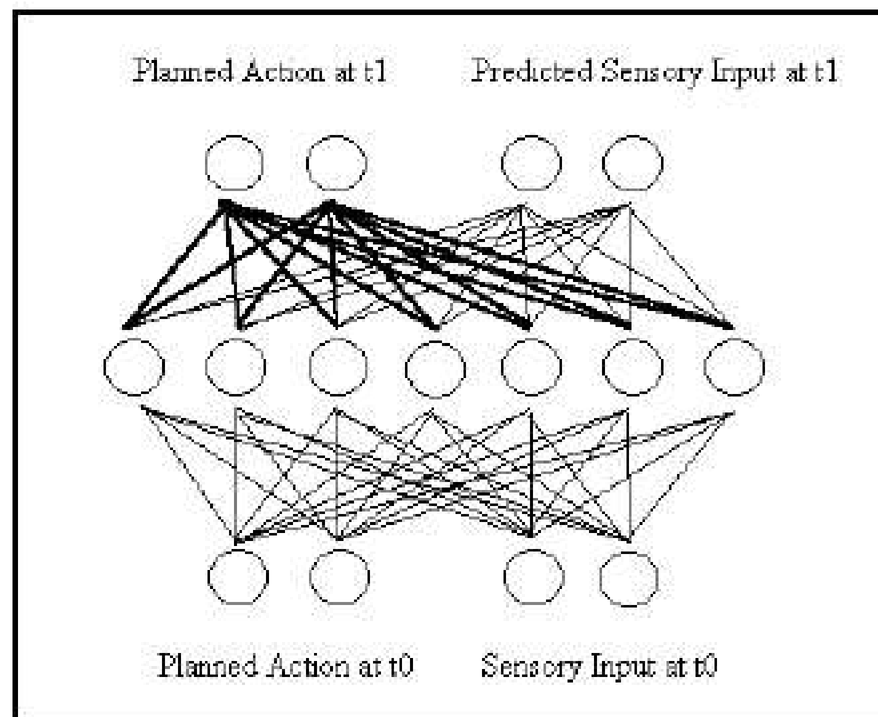

Experiment 2

by Nolfi, Elman and Parisi (1994)

different tasks II

- only the fittest are allowed to mutate and reproduce (here: copy themselves)

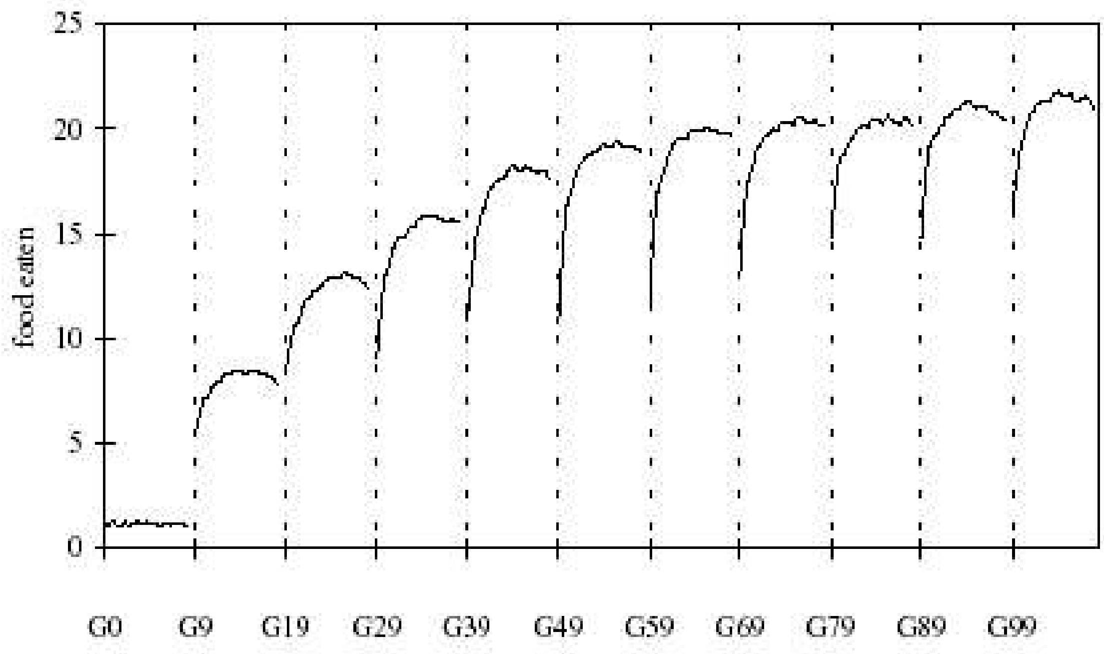

- note: the environment supervises the learning task

different tasks III

- Learning <u>forces</u> evolution to select individuals that improve their performance in both tasks

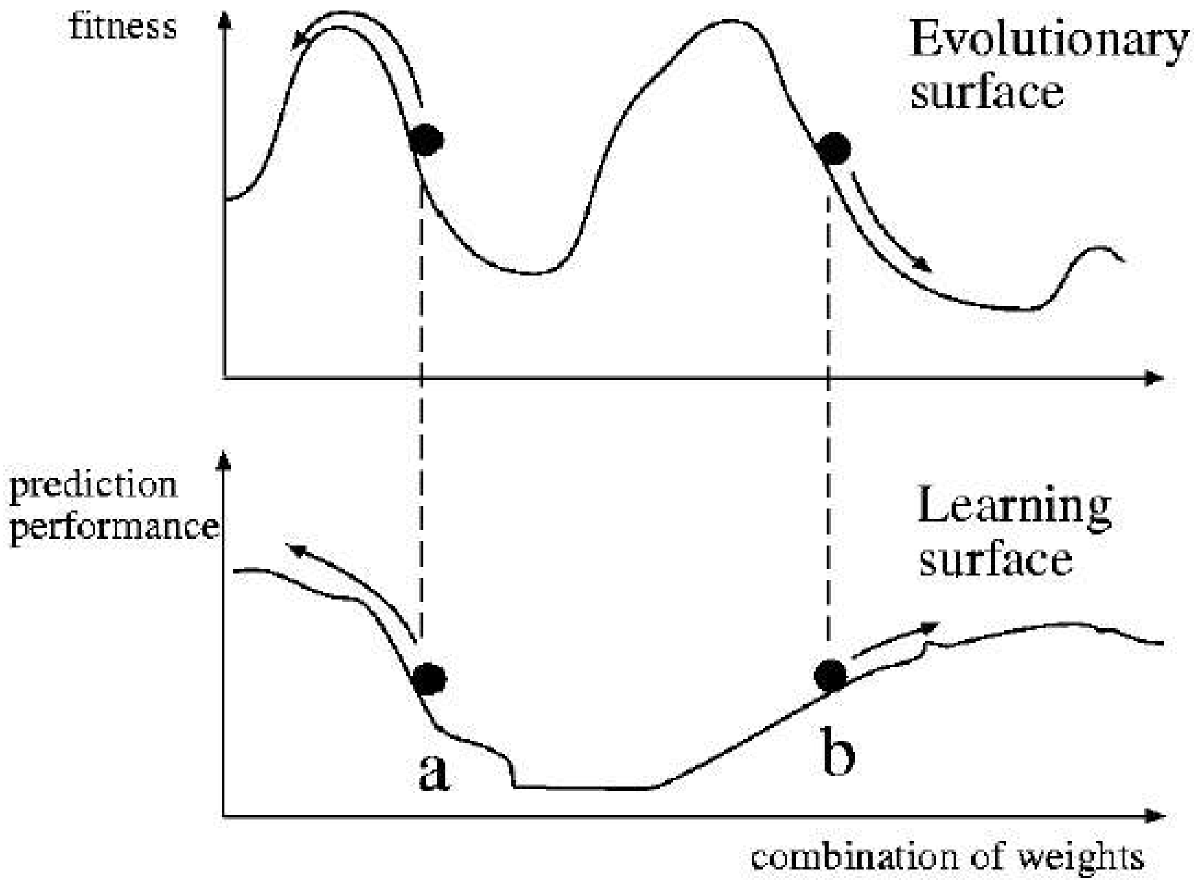

Using The Environment

Navigating
in a „maze“

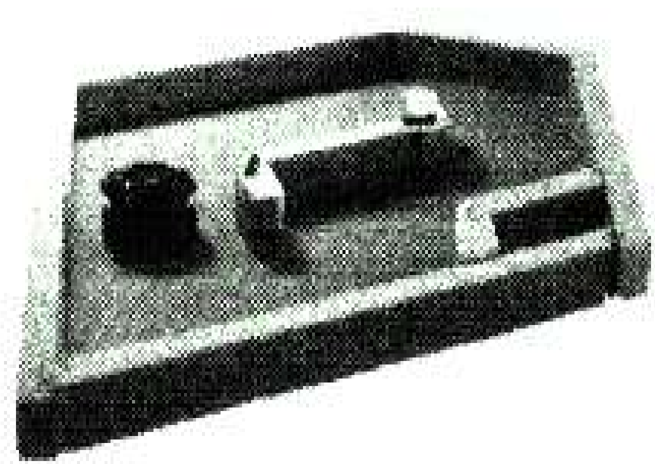

- Genome only encodes meta-properties
  - learning rule
  - learning rate
  - excit./inhibit.
- Phenotype determined by learning
  (to large extent)

Navigating

Fitness is measured by ability to navigate

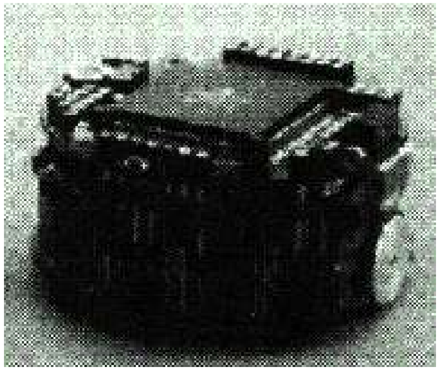

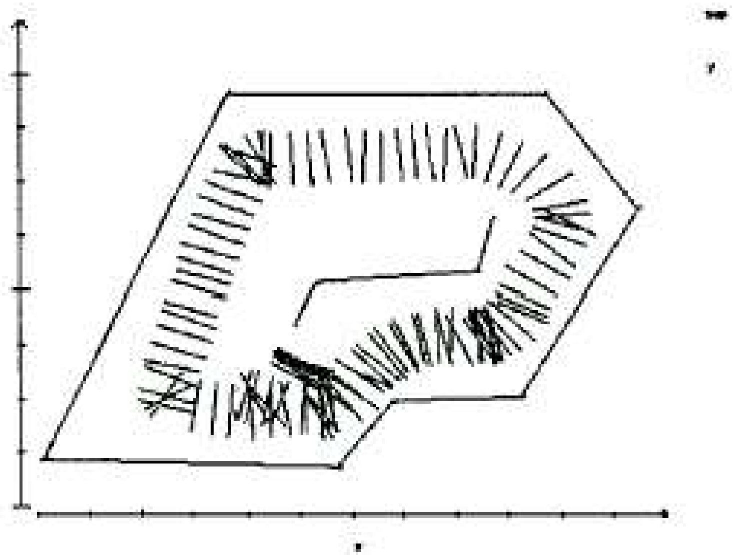

Synapses continue changing, even after behavior stabilizes!

Experiment 3

Floreano & Mondada (1996)

Dynamic Equilibrium

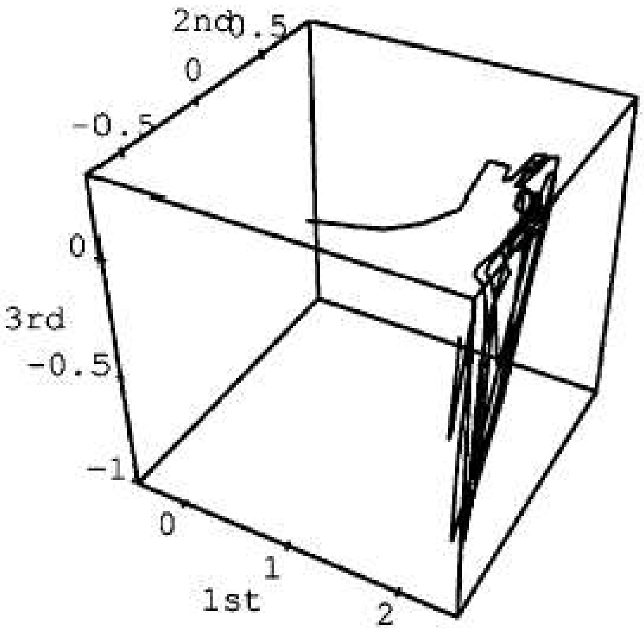

PCA

Floreano & Mondada (1996)

Dynamic Equilibrium

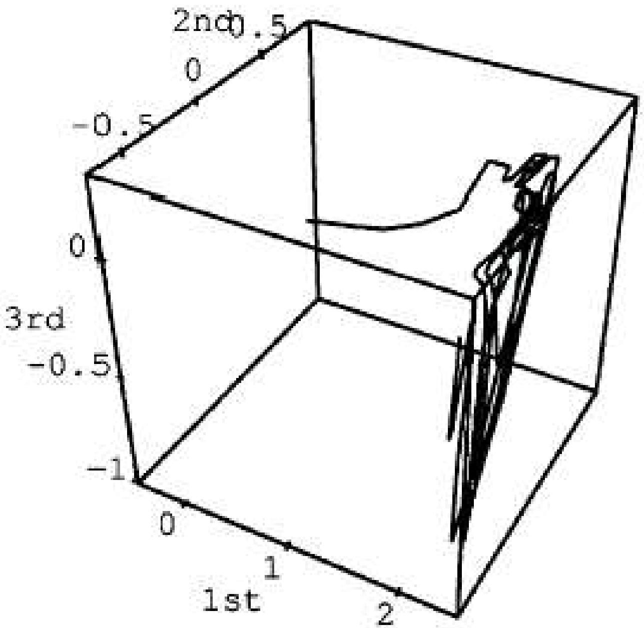

PCA

Comparison to preceding experiments:

- Evolution and learning operate on different synaptic entities (meta-properties vs. weights)
- Learning task is evolved
- Exploiting the regularities of the environment

changing environments I

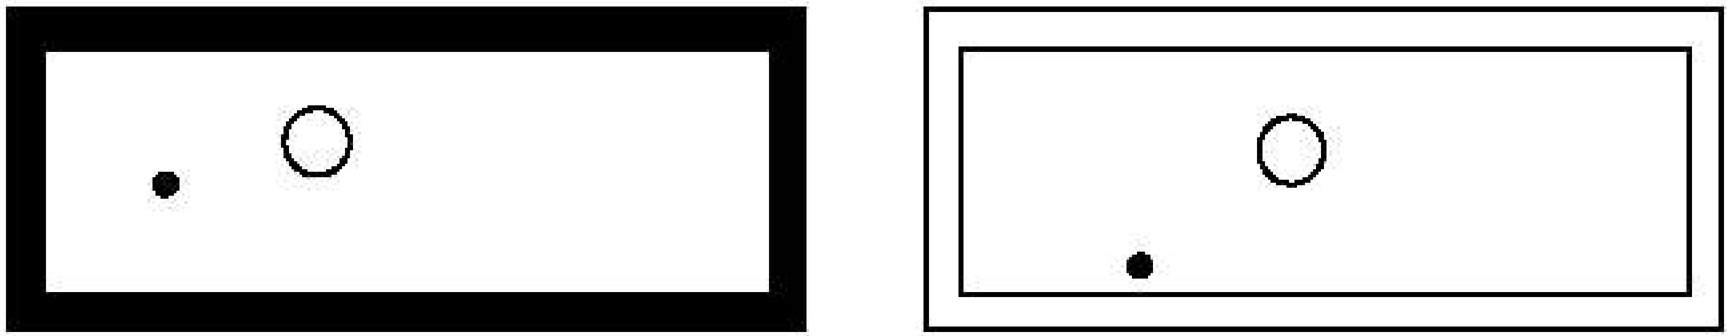

Khepera bots explore 60x20 cm arena, environment can be one of two cases:

a) dark walls – move carefully
b) bright walls – can move faster

fitness function: find the target (it is not visible for the robots).

changing environments II

Two distinct networks: „teaching“ (thin lines) and „standard“ (thick lines).

The output of the two teaching units is used as inputs for the two motor units.

All weights are inherited, but only the „standard“ network changes in lifetime => evolution determines how the environment modifies behavior.

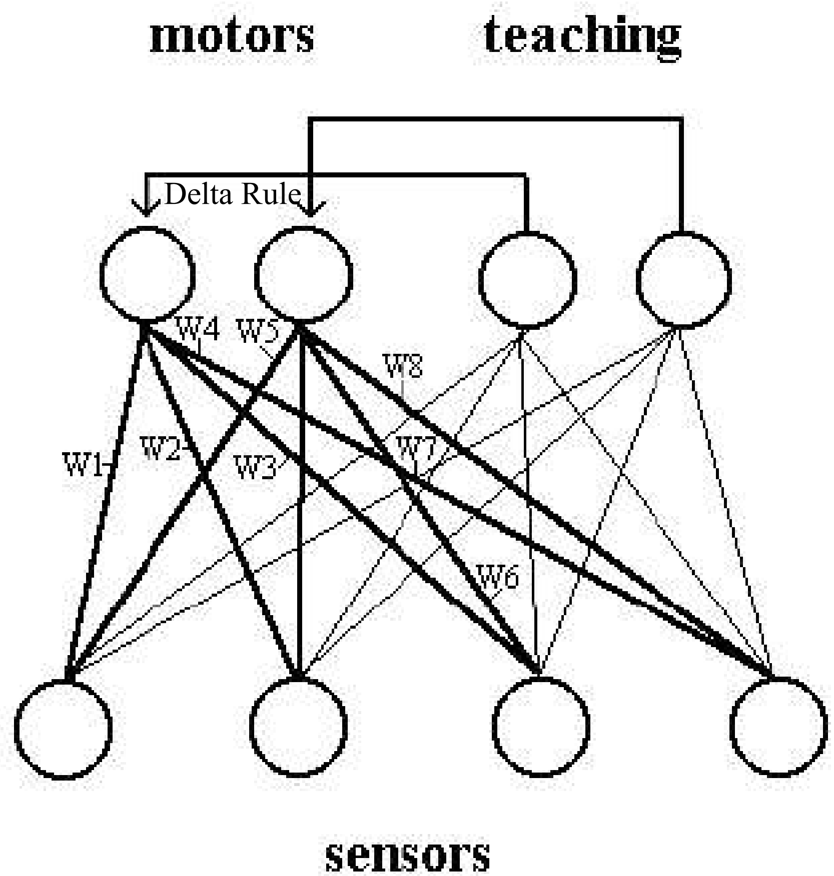

Results:

a) lifetime learning outperforms non-learning

b) acquired characters are adapted to the specific environment - how?

Experiment 4

N olfi & Parisi (1997)

changing environments III

- Activation level of sensors are very different and affect what values the „teaching" network computes. Discrepancies are held high at the beginning of learning (see below).
- The „standard" network performs poor when no learning is allowed.

=> both encode a *predisposition to learn to behave* efficiently.

„behavior is the *emergent* result of the interaction between standard weights, teaching weights and the environment"

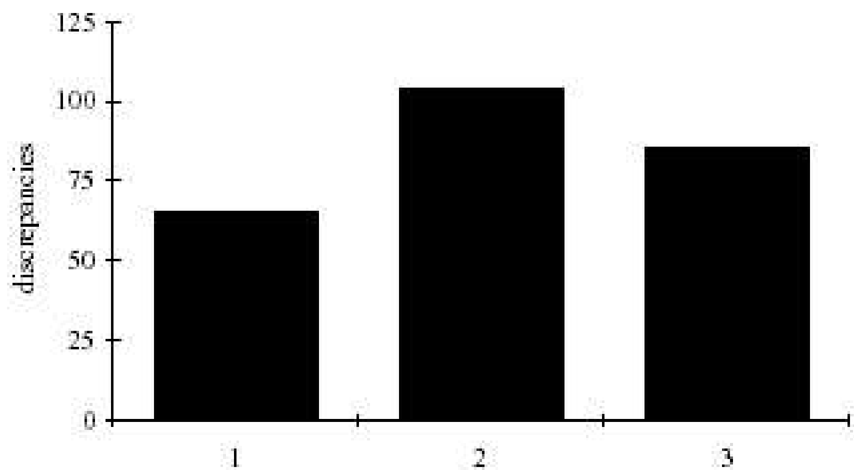

Experiment 4

Nolfi & Parisi (1997)

Co-Evolution

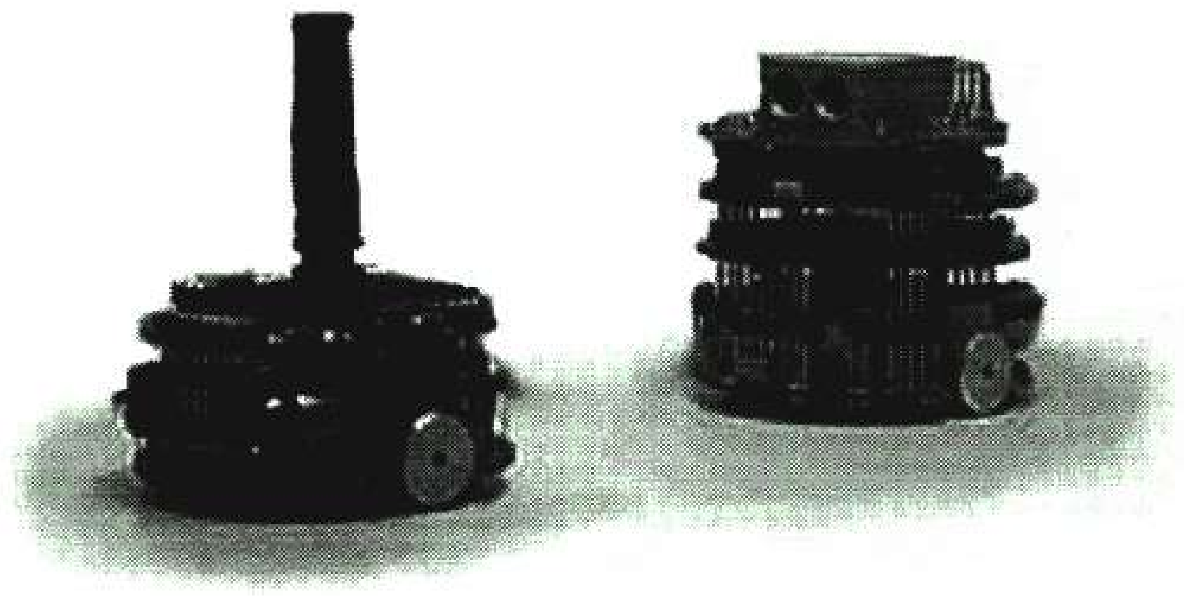

fast prey

slow predator with camera

Floreano & Nolfi (1999)

Cyclic, non-general Strategies

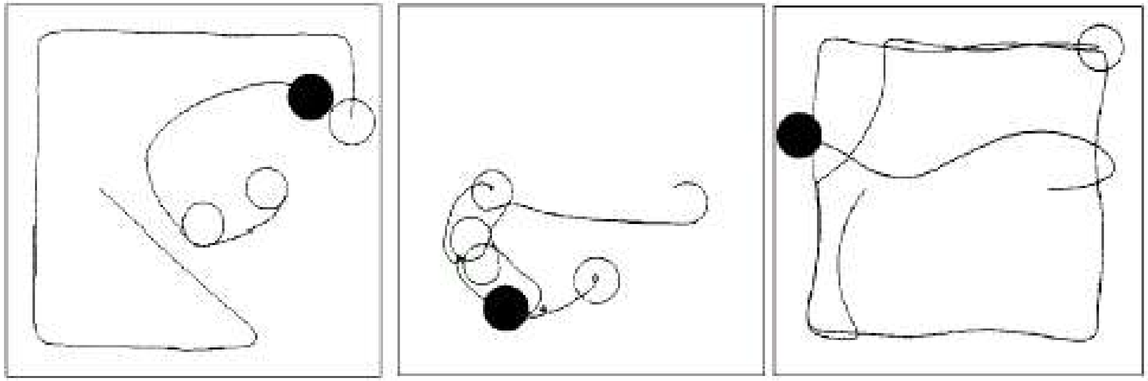

Experiment 5

Floreano & Nolfi (1999)

Cyclic, non-general Strategies

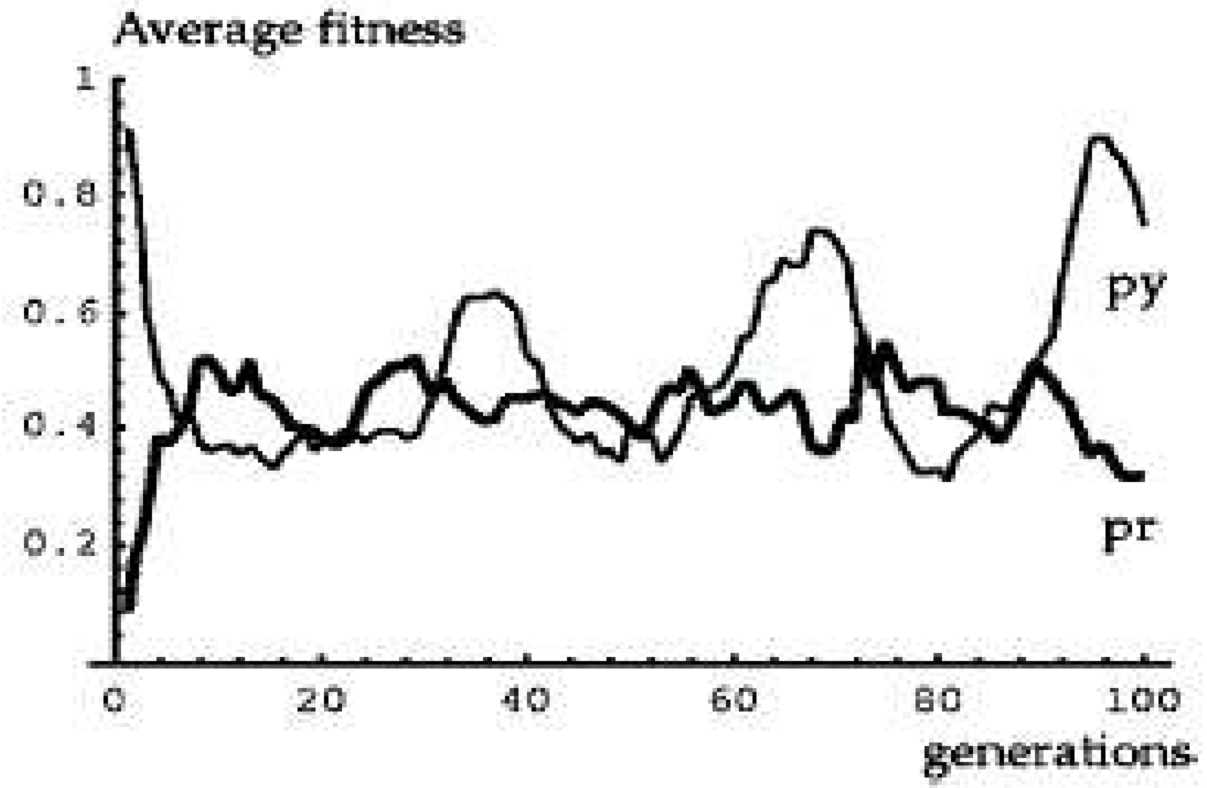

Experiment 5

Floreano & Nolfi (1999)

There is no general strategy!

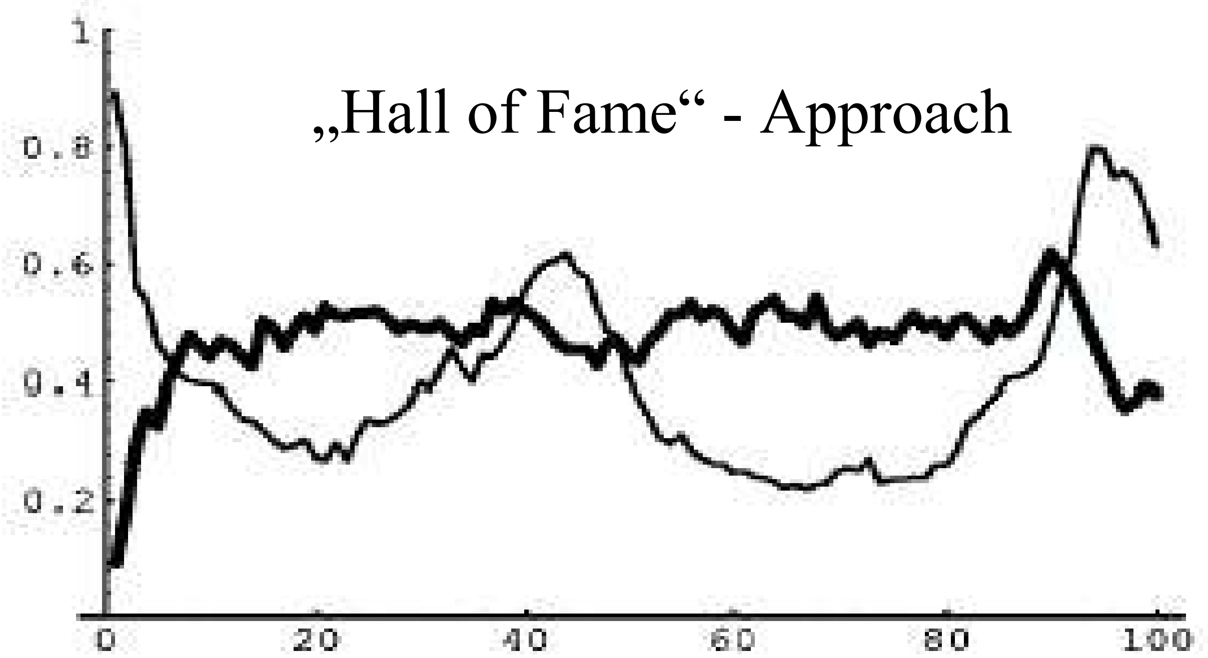

Figure 18. Average fitness across generations. Thick line, predator; thin line, prey.

Experiment 5

Floreano & Nolfi (1999)

Conclusions

- Evolution guides learning
  - Predisposition to learn
    - Channel lifetime learning into the right direction.
    - Inherit a tendency „to make diverse experiences" (perceptual differences)
  - The environment acts as „supervisor"
    - Individuals capable of translating environmental information into teaching input
- Learning can aid evolution (Baldwin-Eff.)
  - Enhancing performance & channeling search
  - simplicity & generality in one:
    - Plastic-general individuals

Want it all?

They made a book of it:

- [19] Nolfi, S., and Floreano, D. **Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines**.

  MIT Press, Cambridge, 2000.

  „This book describes the basic concepts and methodologies of evolutionary robotics and the results achieved so far. An important feature is the clear presentation of *a set of empirical experiments of increasing complexity*. Software with a graphic interface, freely available on a Web page, will allow the reader to replicate and vary (in simulation and on real robots) most of the experiments.“

Questions?