Active Vision and Receptive Field Development in Evolutionary Robots

by Floreano et al. (2005)

CPSC 644, Spring 2007

Presented by Yoonsuck Choe

Overview

- Mobile robot with pan/tilt camera.
- Neural controller mapping visual input to motor output.
- Combined GA and Hebbian learning (in visual system).
- Addition of Hebbian learning results in robust adaptive behavior.
- Active vision gives better RFs than from randomly sampled images.
- Interplay of active vision and RF formation amounts to selection and exploitation of a small and constant subset of visual features available to the robot.

Motivation

- "... it is generally accepted that animals and models of early-vision stages extract the dominant statistical features of their visual environment. Within that perspective, the visual system is a passive, albeit plastic, device shaped by the environment."
- "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, thus reducing the information load on the neural system."
- "... the interaction between receptive field formation and active vision has been largely neglected in the biological and computational literature."

Neural Architecture

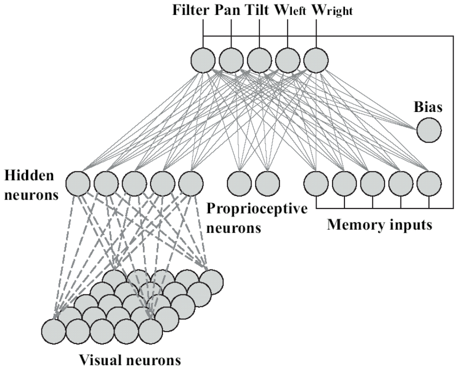

- Network weights evolved using GA.
- 5×5 array of visual neurons, each with 48×48 receptive field.
- Visual-to-hidden RF weights learned through Hebbian learning.
- Sigmoidal units: hidden and output.

Two Visual Activation Modes

- Left: Input
- Middle: visual neuron's activity = mean of all pixels in RF.
- Right: visual neuron's activity = one pixel sample in RF.
- **Note**: visual neuron's RF is not adaptive; visual neuron to hidden neuron weights are adaptive.

Adaptation and Evolution

- Standard sigmoid of weighted sum.
- Weight adaptation follows $x_j \longrightarrow y_i$:

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

- 100 genomes, evolved for 20 generations.
- Best 20% individuals reproduces. Crossover prob. 0.1 per pair and mutation prob. 0.01 per bit.
- Fitness: ability to move straight forward.

Operating Environment

- Simulated or real.
- Trained in simulation (ontogenetic development), tested in real environment.
  - Simulation to real transition is not always easy.
  - Does learning make the transition easier?

Learning in Simulated Environment

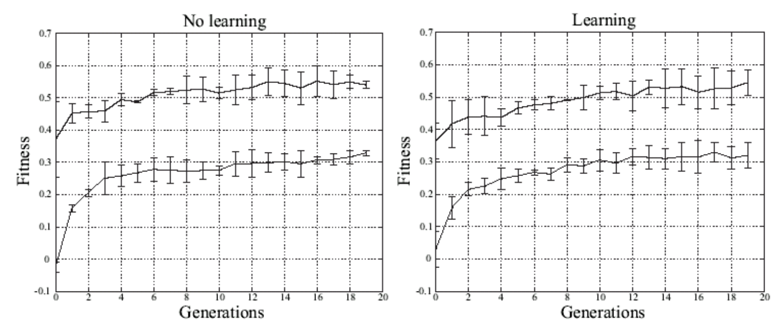

- No difference between learning vs. no learning (evolution only).

Learning vs. Fixed (Nearly Optimal) RF

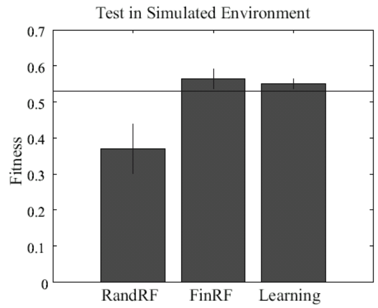

- RandRF: Random visual to hidden connections.
- FinRF: Visual to hidden weights fixed during lifetime (to well-formed RF)
- Learning: Initially random, learn weights during individual's life.
- No difference between FinRF and Learning conditions: no overhead for learning, in terms of fitness.

Simulation to Real Environment

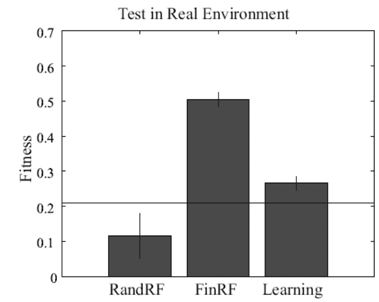

- RandRF: Random visual to hidden connections.
- FinRF: Visual to hidden weights fixed during lifetime (to well-formed RF in new env)
- Learning: Initially random, learn weights during individual's life.
- Overhead in learning, due to novelties in environment.

Behavior in Real Environment

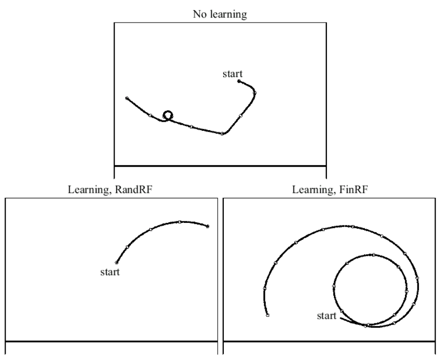

- Evolution only: collision, tilt camera down
- Learning, initial: collison
- Learning, final: no collision, pan camera (possibly due to smaller number of features encoded)

Pan Angle of Learning vs. No Learning Agents

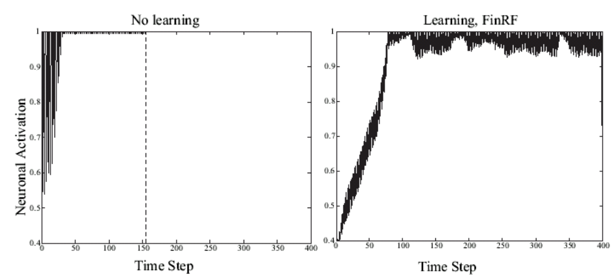

- Learning condition give more variability in panning angle.

RF Training

- Systematically sampled outdoor image samples for training RF, compared to Hebbian learning.
- Typical input scenes shown above.

RF Development

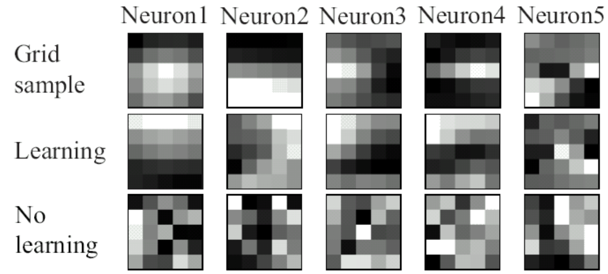

- Grid sample: RF learned off-line.
- Learning: RF learned on-line.
- No learning: Random RF of best no-learning case.
- "...robots self-select ... a significantly smaller and consistent subset of visual features ... for performing ..."

PCA of Observed Scenes

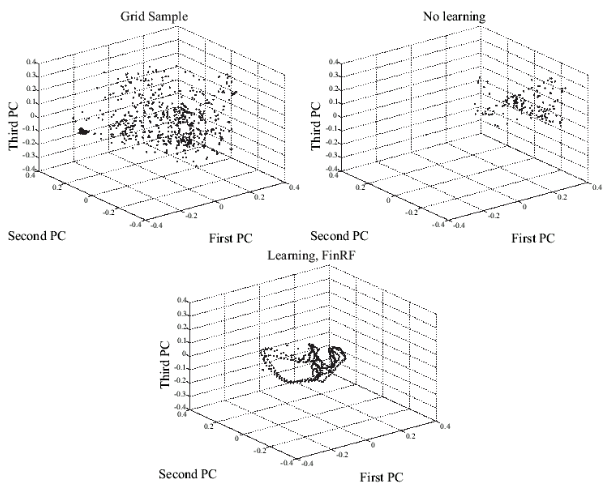

- Scenes gathered via grid sampling, or learning robot's final RF, or no-learning robot's observation.
- PCA of the scenes.

Conclusion

- Behavior affects learning by selecting a subset of learning experiences that are functional to the survival task.
- Learning affects behavior by generating selection pressure for actions that actively search for situations that are learned.
- Learning contributes to the adaptive power of evolution by coping with change that occurs faster than the evolutionary timescale.

Discussion (YC)

- “Interplay of active vision and RF formation amounts to selection and exploitation of a small and constant subset of visual features available to the robot.”
  - That “constant subset” may not be that of “visual features” but of **motor primitives**.
- The fitness “ability to move straight forward” seems too arbitrary.
- The main conclusions are a bit too far reaching.

References

Floreano, D., Suzuki, M., , and Mattiussi, C. (2005). Active vision and receptive field development in evolutionary robots. *Evolutionary Computation*, 13:527–544.