TOPOS: generalized Braitenberg vehicles that recognize complex real sounds as landmarks

Pablo González-Nalda[1] and Blanca Cases[2]
[1]Escuela Universitaria de Ingeniería de Vitoria-Gasteiz, UPV-EHU
[2]Facultad de Informática de Donostia-San Sebastián, UPV-EHU
pablo@si.ehu.es

Abstract

We deal with a generalization of Braitenberg (called *topos*) vehicles in a Evolutionary Robotics framework, to obtain a control system of a robot that can discriminate and reach the position of one of the complex sound sources, learned in a phylogenetic learning scheme.

Application TOPOS is a computational model for evolving populations of Khepera-like simulated robots whose control systems are weakly inter-connected symmetric spiking neural networks.

We take a strong biological referent, modeling the structures that perform the auditive processing of sound like the cochlea, through the Fourier Transform representation of complex signals. TOPOS *robots* are tested for recognizing and choosing one of two complex waves in an evolving population. In spite of the well-known difficulty of this problem, we found that Braitenberg vehicles that use spiking neural networks are very robust, obtaining high scores of success.

This work is an approach mainly but not only from the technical point of view to the problem of recognition of complex time signals in a navigation and embodied structure.

Further extensions of this work are proposed for discussion, as well as a deep analysis of the theoretical side.

Introduction: Braitenberg Vehicles and Evolutionary Robotics

Braitenberg vehicles (Braitenberg, 1984) are thought experiments based on tropism and taxis: the movements of plants and animals toward or away from a stimulus. The more interesting Braitenberg vehicles are symmetric devices composed of two frontal sensorial inputs together with a free wheel and two back wheels impelled by motors. The vehicle is governed by a circuit that makes a crossed connection from the left sensor to the right motor and from the right sensor to the left motor. If the left sensor is fired, the right motor speeds up and the vehicle runs turning to the left; when the left sensor receives the signal, the vehicle advances turning to the right, until the vehicle reaches the source of the stimulus.

In Artificial Life and Evolutionary Robotics (ER) (Beer and Gallagher, 1992) (Cliff et al., 1993) (Floreano and Mondada, 1994) (Floreano and Urzelai, 2000) (Harvey, 1995) (Harvey et al., 2005), the embodiment of Braitenberg vehicles in handcraft small robots or Khepera robots that move toward a white light source is very popular (Mobus, 1994) (Salomon, 1997), following the Braitenberg experiment. There are works dealing with sounds instead of white light, as cricket phonotaxis (Lund et al., 1997). This sound is a simulated male cricket song, formed by two *chirps* per second, where each chirp is four cycles of 25Hz square wave amplitude modulation of a 4.8kHz tone (Horchler et al., 2003).

The application presented in this paper receives the name of TOPOS (gr. *place*) in reference to navigation problem (Floreano and Mondada, 1996) based in recognition of complex signals. This is a way to represent a non structured environment, since the final objective of Autonomous Robotics (Brooks, 1991b) (Brooks, 1991a) (ER included) is to obtain systems robust enough to behave well in hard environmental conditions. Previous works only use simple constant signals as stimulus, and other references (Floreano and Mattiussi, 2001) put emphasis in the great difficulty of recognizing complex signals that vary in time.

The Dynamical Systems approach is also getting stronger in the world of ER (Beer, 1994) (Harvey et al., 2005). In TOPOS the motors-to-sensors external feedback is essential for the task of navigation, like the head movements we make to find the source of a sound. The time is processed in axon delays among neurons in a internal dynamical system, coupled with the outside through sensors and motors.

Generalizing Braitenberg vehicles to recognize the sounds of nature

The aim of this work is to enrich the simple structure of Braitenberg vehicles by increasing their perceptive capabilities to recognize and select between two sources of complex sound, like recorded ones. To have an idea of the hardness of the problem we face, think in a Braitenberg vehicle that recognizes ”disco-flashing lights”, when it perceives the dynamic pattern of some seconds of duration of two RGB-color-composed light sources.

The approach used to study the ability to recognize and reach sounds resembles a *Skinner box* (Skinner, 1938). It is commonly used with pigeons and rats that get food or an electrical discharge whether they do a task.

The phonotaxis behavior has been studied in ER (Di Paolo, 2003) (Scutt, 1994). As a example, in (Lund et al., 1997) a Khepera resembles a female cricket. This approach tries to build a close simulation of the cricket morphology and physics, including neural design that comes from the real nervous system of crickets, and can only recognize a concrete structure of chirps (two *chirps* per second, where each chirp is four cycles of 25Hz square wave amplitude modulation of a 4.8kHz tone).

Handcraft design is almost impossible (Salomon, 1997), so evolution is needed to obtain a whole controller from sensor to actuators (Chiel and Beer, 1997) (Tuci et al., 2002).

The agent is situated and embodied, its sensors have a shape that modifies the amplitude of the signal that passes to the neural network, and the place in the arena determines the intensity of sound (Chiel and Beer, 1997).

The agents also perform another modification of the sound that they receive to model the auditory system of mammals, as explained in next section.

Fourier analysis and frequency spectrum

Mammals process sound signals by means of the auditory system in conjunction with the nervous system (Handel, 1989) (Moore, 1997).

Mammalian auditory system has three parts that transform sound waves into input signals for nervous system. The external ear localizes the sound sources and funnels waves into auditory channel causing the vibration of the tympanic membrane. The middle ear transfers that movement of the membrane through a chain of bones (malleus, incus, and stapes) to the oval window of cochlea determining the dynamic range of the sound. Finally, oval window causes movement of fluid in cochlea, ultimately resulting in stimulation of cochlear hair cells which excite neurons of spiral ganglion, that send spike coded auditory signals to brain through cochlear nerve.

If we simplify non-linear effects, we can say that the cochlea performs a Fourier transform of sound waves, and this is the input received by the nervous system (Handel, 1989). The Fourier Transform takes as input a wave over time (e.g. a sound sample) and produces the vector of complex numbers representing the amplitude and phase of each interval of frequency (the spectrum). In sound-processing computer programs we can see sound in this type of representation (figure 1).

Sound and sensors in the model

In order to save computational resources we only use the amplitude data previously produced applying the Fourier Transform.

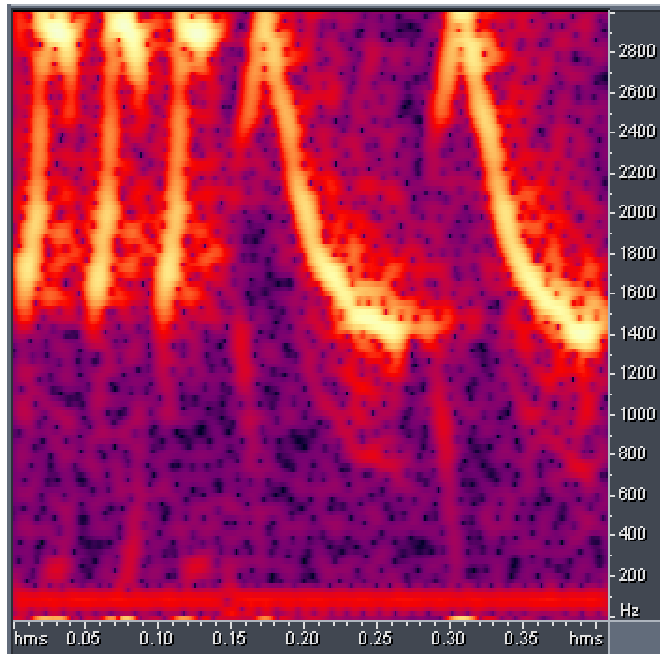

Figure 1: The frequency spectrum of a bird chirp. The second experiment prepared to test TOPOS system uses the chirps in the left (first 0.15 seconds) in a sound source and the others in the other source. Time in seconds (horiz), frequency in Hz (vert), and highest amplitudes in yellow.

The signal is a vector of real numbers presented to sensors every time step of 0.04 seconds using 64 strips in the range of frequency and 6000Hz as sampling frequency. To make a comparison, 0.1 seconds is the time necessary for human auditory system to differentiate two sound signals (echo threshold).

The sound signal is propagated from the sound sources with an intensity that is inversely proportional to the square of distance.

Each sensor reads information from a strip of the range of frequencies (characteristic frequency, CF) and takes the values from adjacent strips with less weight. The sensors also take into account of threshold and saturation level, and add some noise to signals, to provide flexibility to the system as says (Jakobi, 1998). The sensors are more sensitive to signals in the CF, but not only, so that they can react with close frequencies, depending on the genetic design. There are ”wide” and ”narrow” sensors, whether they have a high value in this parameter. In figure 2 we can see real data, and the neuron reacts when frequency and amplitude draw a point in the area of the solid line.

Each strip has a range of 47 Hz, so a deep bass sound is represented in the first strip. This sound will excite a sensor whose central frequency is the second strip, but with less intensity.

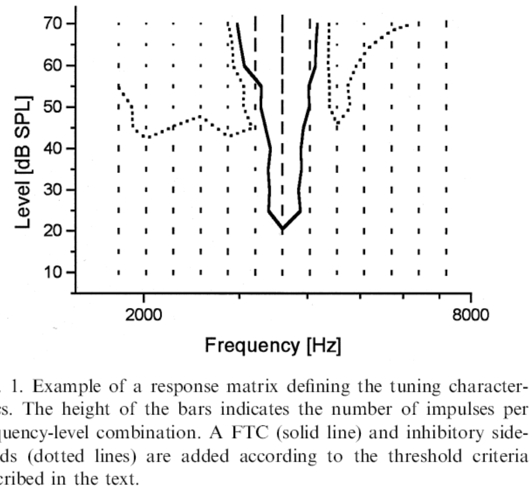

Figure 2: The activation threshold for each frequency, measured with electrodes implanted in starlings. Figure and caption taken from (Nieder and Klump, 1999)

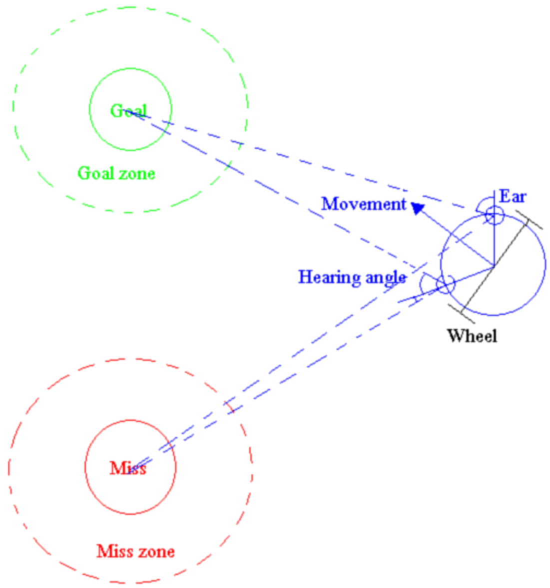

Figure 3: Process of hearing.

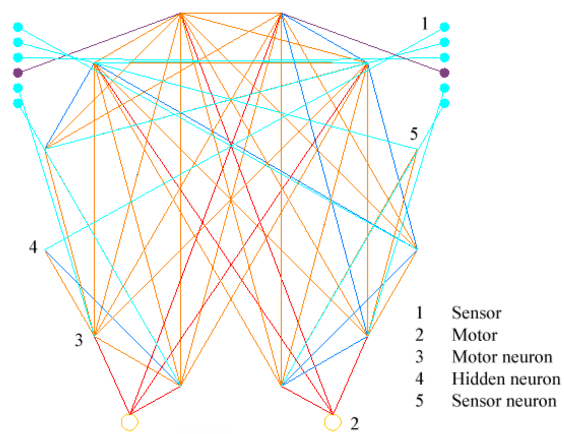

Figure 4: The internal structure of an evolved topo is shown. Connections to motors in red, axons in orange (+) and dark blue (-), and connections from sensors in cyan and magenta. Signs in brackets mean colors for positive and negative weights.

TOPOS System

TOPOS application has been implemented in `java`; it evolves populations of generalized Braitenberg vehicles.

The use of a realistic biological referent for audition (described previously in *Fourier Transform*) is fundamental; but TOPOS is a model, that is, it doesn't reflect precisely the dynamics or shape of a concrete robot. Though, the experiment can easily be transformed to a simulation, since strong restrictions aren't imposed.

General description

TOPOS arena implements the physics of sound transmission and perception. The sound pressure (power) decreases with distance (as said in previous section). There is no echo nor sound reflection. The user puts two sound sources into the world selecting for each one if it is the goal or the miss. Braitenberg vehicles evolve to recognize and reach one of them (the goal), starting from a point at the same distance from both sources, and faced to them with a small amount of randomness in the starting angle (up to 30^o to the left or to the right).

TOPOS makes a generalization of the structure of Braitenberg vehicles: a neuronal network composed of two symmetrical sub-nets, shown in window figure 4. Each subnetwork has 8 internal neurons and some connections to the motors. Subnets are partially superposed and share 4 neurons, to ”pretend” the brain structure. Each sub-net is a fully connected spiking neural network(see next subsection).

Ears are composed of six sensors, each one connected to one of the input neurons as a special input. This neuron can be one of its side or of the opposite.

The physics of a real microphone is included in the model applying a cardioid-shape function to the received signal as a way to represent the pinna (external ear). The sound is attenuated multiplying the amplitude of every frequency (in the simplest scheme) proportionally to the hearing angle (multiplied by 0.0 if sound comes from the back, and by 1.0 if from the front). The emulation of a microphone will allow us the easy embodiment in a real robot.

A topo moves using two motors and wheels. Its difference of velocity produces rotation, the same mechanism used in Khepera-like robots and emulators. Motor neurons feed motors in a integration fashion, calculating the mean of time passed since motor neurons sent spike to motors. The lower the mean, the faster the wheel.

The spiking neural network

Beer's and others' work in ER (Beer and Gallagher, 1992) (Cliff et al., 1993) (Harvey, 1995) (Harvey et al., 2005) use Continuous Time Recurrent Neural Networks (CTRNN). We also use a recurrent model but with a spiking-type of neurons that can code the information in pulse trains and time delays and has a fixed strength spike when the potential of the neuron *overflows* the threshold.

TOPOS' neural network type has been used only recently in ER (Di Paolo, 2003) (Floreano and Mattiussi, 2001) (Maass, 1997), though it has been used in other works with the name *Pulsed Neural Networks* (Maass and Bishop, 1999).

Rich variety of behavior is shown (chaotic dynamics included) in (Beer, 1995). In Di Paolo's work: *"Spiking neural networks possess a number of attractive features."* (Di Paolo, 2003) The most important characteristics are that these networks are biologically plausible (Floreano et al., 2005), that they can integrate perceptions over time before actions (Yamauchi and Beer, 1994), that can process time information (delay among stimuli) and that are mathematically equivalent to sigmoid networks, stronger to noise, and sometimes requiring significantly fewer neurons and (Maass, 1997).

Spikes travel through axons with a speed and time delay till they reach the weighted synapses. Inhibitory synapses can block the receptor's spike.

Evolution

TOPOS has the structure of a typical ER system:

Genome: described by a vector of float numbers that determine the values of sensors, neurons, delays in axons, synapse weights, and speed of motors. Each *hemisphere* is a set of 8 fully interconnected neurons, and the two *hemispheres* share 4 neurons.

Morphogenesis: the translation of genome to a subnetwork is a direct one, but this information is mirrored to represent bilateral symmetry (Gallagher and Beer, 1999), connecting a low number of neurons between the two subnetworks in the same fashion as in mammals' brain.

Boolean values are represented by a real number, and when they reach a certain threshold get *true* value.

Fitness: each *topo* in the world is tested in five equivalent trials. A trial consists in performing a navigation task during a given time following the sound sources as landmarks.

The starting position of a topo in each trial is equidistant of the two sources (10.5 units of distance). The sources are randomly placed each trial, in the right or left side of the starting position.

Score (fitness F) is the sum of the scores from each trial. Trial score is calculated from **minimal and last distance** in the trial, and cannot be a negative value. A **bonus** (+10) is added when a topo reaches a circle around the correct source of a given ratio ***(goal)***, while a **penalty** of 10 points is applied if the topo reaches the incorrect source ***(miss)*** when crossing its frointer. This limits are at 5.25 units, the distance between sources divided by 4.

$$Trialscore = 10 - mind - lastd + bonus - penalty$$

The more score, the better the vehicle has performed. If the vehicle gets a penalty its score would be negative. If we assign a 0 this trial and if it performs well in the other trials, the penalty will not affect global fitness.

If a vehicle goes in the wrong direction, its minimal distance is that of starting place, more than fixed value 10, so trial score will also be 0.

If a topo reaches the correct source (the limit is at 5.25 u), score is $10 - 5.25 - 5.25 + 10 - 0 = 9.5$

Fitness is an absolute value, not depending on environment or other topos' behavior. Red queen effect cannot appear with this kind of experiments.

Genetic algorithm and selection: an *élite* of the topos with more score in the test are selected for survival to next generation (25%) and the other places (75%) in the fixed size population are filled with crossover of two parents (two cutting points). Parents are chosen from all the past population (élite included) with a uniform probability.

When copying in crossover, all the genetic numbers (with a 5% of probability) are modified from 1 to 25% of its value.

Experimental results

Each experiment has these numbers: the population is of 100 topos, each topo is tested 5 times, with a duration of 20 seconds maximum each trial. If an individual does a goal or a miss, the trial is over. Last generation (200) is a special round of 100 trials to test the élite to obtain accurate data for the statistical analysis of the individuals with the best score of the run.

Instead of fitness, we use these two variables to measure the ability to distinguish the two sources:

- relative effectiveness **`efr`** = goals/(goals+misses)
- absolute effectiveness **`efa`** = goals/trials

Two slightly different passages in the beginning of Metallica's *"Nothing else Matters"* have been chosen. They have no voice, and are similar in music and length. Both sound simultaneously, and the vehicle has to move to face one and tell if it is the correct or has to spin to the other one, while it hears both sounds.

The **figures** 5 and 6 show the high scores in the task.

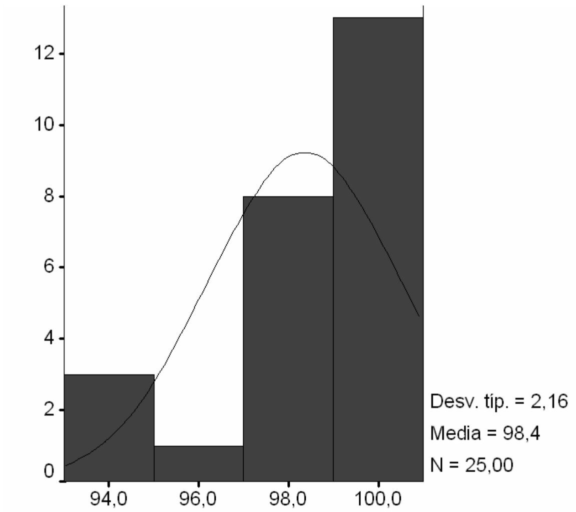

Figure 5: Histogram of Relative effectiveness (frequency in the x-axis vs number of individuals with that EFR in the y-axis).

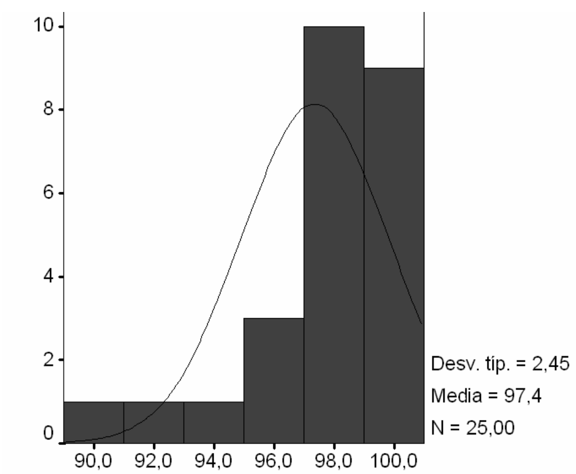

Figure 6: Histogram of Absolute effectiveness (frequency in the x-axis vs number of individuals in the y-axis).

To compare, we can use these results (see **figures** 7 and 8) from a test experiment that has the same sound in both sources. Solution is to goal or fail, but there is no information to choose, so is 50% of effectiveness.

For a second test, we have prepared two sets of bird chirps (say A and B sounds), that are parts of a song from one bird, from the same recording. Sound A has a chain of chirps like the ones in the first 0.15 seconds of the **figure** 1, and B like

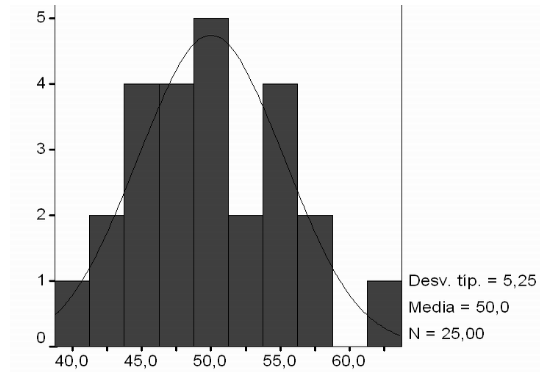

Figure 7: Test experiment: EFR.

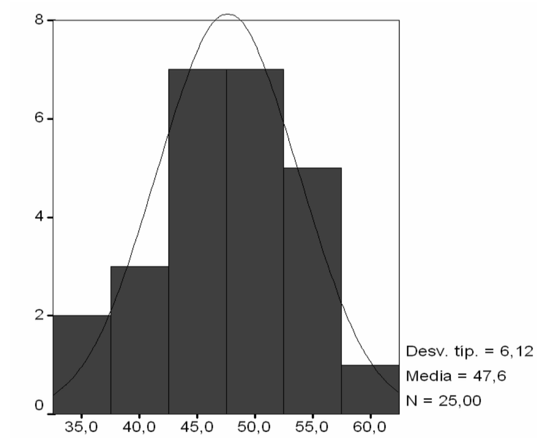

Figure 8: Test experiment: EFA.

the ones in the end of the figure.

Sound A has half a second of chirps and half a second of silence. Sound B has first silence and then the sound, so agents can hear each sound separately. Sounds continue at the point last trial ended, to avoid vehicles use a short-cut. For a new *topo* the sound begins at a random point, and comes from the correct or from the wrong source.

We name experiments *AB*, *BA* and *AA*, being the first letter the sound that rewards and second the miss. Double AA means that we have the same sound in both sources, and topos cannot choose the right one. See **table** 1.

Exp	EFR $\bar{x}$	EFR σ	EFA $\bar{x}$	EFA σ
AB	99.1	1.76	90.5	4.21
BA	100	0	97.8	2.33
AA	34.8	37.6	2.6	2.86

Table 1: Mean $\bar{x}$ and standard deviation σ of EFR and EFA data. Each line has one experiment.

Discussion and Future Work

First, we can say that "it works". To our knowledge, there isn't any work that performs such a complex task of recognition of a wide and general type of signals using this framework of Evolutionary Robotics.

We have tried with hand-made sounds, like white noise with different band-pass filters creating two rhythms, but it seems that artificial sounds are too difficult, and that it is easier to get and recognize the characteristics of natural sounds.

Yet, we can think that the choice of a pair of sounds could have the trick, like the presence of a strong frequency in only one of the sounds. Nevertheless, the preliminary results of an ongoing experiment show that the a sound in a source and just the same sound with a different order in its parts (a simple "cut and paste" modification) can be differentiated, so whether the sound has strong parts should not be a advantage.

We have said that "it works", but we need to study why it does. It is a difficult task like the ones done by neurobiologists do with *C. elegans* and *Aplysia*, animals that have the same role than *Drosophila* for Genetics. We yet have to do the hard work of isolating the behavior of each sensor, neuron and even axon, because the delay thas causes can be the clue to recognition, or maybe dynamical systems are the way to understand how the vehicle works.

We have two different ways of developing this framework. First, on the technical side, we can say that this structure seems powerful enough to achieve more complex problems in navigation through landmark recognition. Nevertheless, there remain some weaknesses like the absence of morphogenesis when we evolve neural networks. Another problem is the huge amount of computational resources of this kind of models (about an hour per generation in a PC).

Ontogenetical learning and applying the Hebbian rule (Hebb, 1949) is another development that must be faced to improve results and to open to problems other than Skinner-box ones.

The ability of the *topos* to emit sounds could open a way to communication in a ecosystem scheme like Artificial Worlds (*Polyworld* is a very interesting example (Yaeger, 1994)), instead of independent trials in a traditional genetic algorithm.

Acknowledgments

This work was supported by UPV project 9/UPV 00003.230-15840/2004.

We also thank people from *Logic and Philosophy of Science* Department at the UPV/EHU for their help.

References

Beer, R. (1994). A dynamical systems perspective on agent-environment interaction. In Smithers, T. and Moreno, A., editors, *On the role of Dynamics and Representation in Adaptive Behaviour and Cognition III International Workshop on Artificial Life DRABC94*. Euskal Herriko Unibertsitatea, Universidad del País Vasco.

Beer, R. (1995). On the dynamics of small continuous-time recurrent neural networks. *Adaptive Behavior*, 3(4):469–509.

Beer, R. and Gallagher, J. C. (1992). Evolving dynamical neural networks for adaptive behavior. *Adaptive Behavior*, 1(1):91–122.

Braitenberg, V. (1984). *Vehicles. Experiments in Synthetic Psychology*. MIT Press, MA.

Brooks, R. (1991a). Artificial life and real robots. In Varela, F. and Bourgine, P., editors, *Toward a Practice of Autonomous Systems, First European Conference on Artificial Life (ECAL)*, pages 3–10. MIT Press, Cambridge, MA.

Brooks, R. (1991b). Intelligence without representation. *Artificial Intelligence*, 47:139–159.

Chiel, H. and Beer, R. (1997). The brain has a body: Adaptive behavior emerges from interactions of nervous system, body and environment. *Trends in Neurosciences*, 20:553–557.

Cliff, D., Harvey, I., and Husbands, P. (1993). Explorations in evolutionary robotics. *Adaptive Behavior*, 2(1):71–108.

Di Paolo, E. A. (2003). Evolving spike-timing dependent plasticity for single-trial learning in robots. *Philosophical Transactions of the Royal Society A.*, 361:2299–2319.

Floreano, D., Epars, Y., Zufferey, J., and Mattiussi, C. (2005). Evolution of spiking neural circuits in autonomous mobile robots. *International Journal of Intelligent Systems.*, In press.

Floreano, D. and Mattiussi, C. (2001). Evolution of spiking neural controllers for autonomous vision-based robots. In Gomi, T., editor, *Evolutionary Robotics IV*, pages 3–10. Berlin, Springer-Verlag.

Floreano, D. and Mondada, F. (1994). Automatic creation of an autonomous agent: Genetic evolution of a neural-network driven robot. In Cliff, D., Husbands, P., Meyer, J., and Wilson, S., editors, *From Animals to Animats III: Proceedings of the Third International Conference on Simulation of Adaptive Behaviour SAB'94*. MIT Press-Bradford Books, Cambridge, MA.

Floreano, D. and Mondada, F. (1996). Evolution of homing navigation in a real mobile robot. *IEEE Transactions on Systems, Man and Cybernetics*, 26(3):396–407.

Floreano, D. and Urzelai, J. (2000). Evolutionary robotics: The next generation. In Gomi, T., editor, *Evolutionary Robotics III*. Ontario (Canada), AAI Books.

Gallagher, J. and Beer, R. (1999). Evolution and analysis of dynamical neural networks for agents integrating vision, locomotion and short-term memory. In Banzhaf, W., Daida, J., Eiben, A., Garzon, M., Honavar, V., Jakiela, M., and Smith, R., editors, *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-99)*, pages 1273–1280. Morgan Kaufmann.

Handel, S. (1989). *Listening: An Introduction to the Perception of Auditory Events*. The MIT Press, Cambridge, MA.

Harvey, I. (1995). *The artificial evolution of adaptive behaviour. D. Phil. Thesis*. University of Sussex.

Harvey, I., Di Paolo, E., Wood, R., Quinn, M., and Tuci, E. A. (2005). Evolutionary robotics: A new scientific tool for studying cognition. *Artificial Life*, 11(1-2):79–98.

Hebb, D. (1949). *The organization of behavior: A neuropsychological theory.* New York: John Wiley and Sons.

Horchler, A., Reeve, R., Webb, B., and Quinn, R. (2003). Robot phonotaxis in the wild: a biologically inspired approach to outdoor sound localization. In *Proceedings of 11th International Conference on Advanced Robotics (ICAR'2003)*, pages 18(8):801–816. VSP.

Jakobi, N. (1998). *Minimal Simulations for Evolutionary Robotics. PhD thesis.* COGS, University of Sussex.

Lund, H. H., Webb, B., and Hallam, J. (1997). A robot attracted to the cricket species gryllus bimaculatus. In Husbands, P. and Harvey, I., editors, *Proceedings of IV European Conference on Artificial Life ECAL97*, pages 246–255. MIT Press/Bradford Books, MA.

Maass, W. (1997). Networks of spiking neurons: the third generation of neural network models. *Neural Networks*, 10:1659–1671.

Maass, W. and Bishop, C. M. e. (1999). *Pulsed Neural Networks*. MIT Press.

Mobus, G. (1994). Toward a theory of learning and representing causal inferences in neural networks. In Levine, D. and Aparicio, M., editors, *Neural Networks for Knowledge Representation and Inference*, pages Chapter 13, 339–374. Lawrence Erlbaum Associates, Hillsdale, New Jersey.

Moore, B. C. J. (1997). *An Introduction to the Psychology of Hearing.* 4th Ed. Academic Press, London.

Nieder, A. and Klump, G. (1999). Adjustable frequency selectivity of auditory forebrain neurons recorded in a freely moving songbird via radiotelemetry. *Hearing Research*, 127:41–54.

Salomon, R. (1997). The evolution of different neuronal control structures for autonomous agents. *Robotics and Autonomous Systems*, 22:199–213.

Scutt, T. (1994). The five neuron trick: Using classical conditioning to learn how to seek light. In Cliff, D., Husbands, P., Meyer, J., and S.W. Wilson, S., editors, *From Animals to Animats III: Proceedings of the Third International Conference on Simulation of Adaptive Behaviour SAB'94*, pages 364–370. MIT Press-Bradford Books, Cambridge, MA.

Skinner, B. F. (1938). *The behavior of organisms: An experimental analysis.* New York: Appleton-Century.

Tuci, E., Harvey, I., and M., Q. (2002). Evolving integrated controllers for autonomous learning robots using dynamic neural networks. In Hallam, B., Floreano, D., Hallam, J., Hayes, G., and Meyer, J.-A., editors, *From Animals to Animats 7: Proceedings of The 7th International Conference on the Simulation of Adaptive Behavior (SAB'02)*, pages 282–291. MIT Press, Edinburgh, UK.

Yaeger, L. (1994). Integrating reactive, sequential, and learning behavior using dynamical neural networks. In Langton, C., editor, *Artificial Life III,Vol. XVII*, pages 263–298. Addison-Wesley.

Yamauchi, B. and Beer, R. (1994). Integrating reactive, sequential, and learning behavior using dynamical neural networks. In Cliff, D., Husbands, P., Meyer, J., and S.W. Wilson, S., editors, *From Animals to Animats III: Proceedings of the Third International Conference on Simulation of Adaptive Behaviour SAB'94*, pages 382–391. MIT Press-Bradford Books, Cambridge, MA.