

Evolved Sensor Fusion and Dissociation in an Embodied Agent

Josh C. Bongard

Artificial Intelligence Laboratory

Department of Information Technology

University of Zürich

Winterthurerstrasse 190, Zürich, Switzerland

bongard@ifi.unizh.ch

Abstract

W. Grey Walter first demonstrated that an autonomous robot could follow an environmental gradient to its source. In this paper, neural networks are evolved that allow a simulated, embodied quadrupedal agent to sense and follow an environmental gradient—in this case, local chemical concentration—to its source. Through a series of ablation experiments performed *in silico*, it is shown how artificial evolution gradually integrates and dissociates the different sensor modalities available to the agent in order to produce chemotacting behaviour. This work builds on that of Walter by indicating that evolutionary methods automatically generate chemotaxis by modulating simpler behaviours (here, forward locomotion) using a sensor modality (chemosensors) separate from those driving the simpler behaviour. This suggests that evolutionary methods are well suited for automatically generating behaviours more complex than chemotaxis by using it in turn as a base behaviour.

1. Introduction

Since Grey Walter introduced his twin tortoises “Elmer” and “Elsie” in the late 1940’s (Grey Walter, 1950), behaviours such as light following (Braitenberg, 1986) and other related behaviours like stigmergy (Dorigo and Caro, 1999, Holland and Melhuish, 1999), chemotaxis (Grasso et al., 1996, Harvey et al., 1997, Ferree et al., 1997) and general gradient following (Kodjabachian and Meyer, 1998) have played a central role in the maturation of artificial intelligence, robotics, artificial life and adaptive animat research.

In this paper, we demonstrate the evolution of neural networks that control a quadrupedal agent to walk towards a chemical point source. The quadruped agent is simulated, but because it behaves within a physically-realistic simulated environment, and its behaviour is generated by sensor signals,

the agent is both situated and embodied, as were Walter’s tortoises. Evolutionary techniques have already been employed to generate sensory-based tracking in simulated, embodied agents: Reil (Reil and Massey, 2001) evolved a bipedal agent with sensors in its hips to track a light source; and Ijspeert (Ijspeert and Arbib, 2000) evolved an animat based on the salamander to track a moving object both in water and on the ground, in which vision modulates an underlying locomotor circuit. This paper furthers these results by demonstrating that artificial evolution can itself compartmentalize different behaviours using the different sensor modalities available to the agent.

Besides the phototaxis demonstrated by Walter’s tortoises, his experiments also hinted at the ease with which more complex behaviours could be generated through the aggregation or modulation of simpler behaviours. The simple trajectory of one tortoise became complex trajectories when two tortoises, each with a light source attached, were placed in proximity to each other (see Fig. 1). This in some way anticipated the subsumption architecture proposed by Brooks (Brooks, 1991), in which more complex behaviours are generated by combining and extending modular components in the robot’s controller in an intelligent manner. However, in the subsumption architecture, more complex behaviours are explicitly generated in the controller, whereas the complex trajectories observed for Walter’s tortoises were a result of unexpected behavioural changes in response to a more complex sensory signal (generated by a non-stationary light source). Here we provide evidence that artificial evolution adds more complex behaviours to a simpler one automatically: a simpler behaviour is generated by particular sensor modalities, which is then modulated to produce a more complex behaviour using an additional sensory modality. How, and to what extent, differing sensory modalities are cross-correlated in the brain is an important current research question in neuroscience (see, for example, (Shimojo and Shams, 2001)).

This property of the simulated evolutionary process is investigated here by systematically lesioning parts of the agent’s neural controller, and observing the change in be-

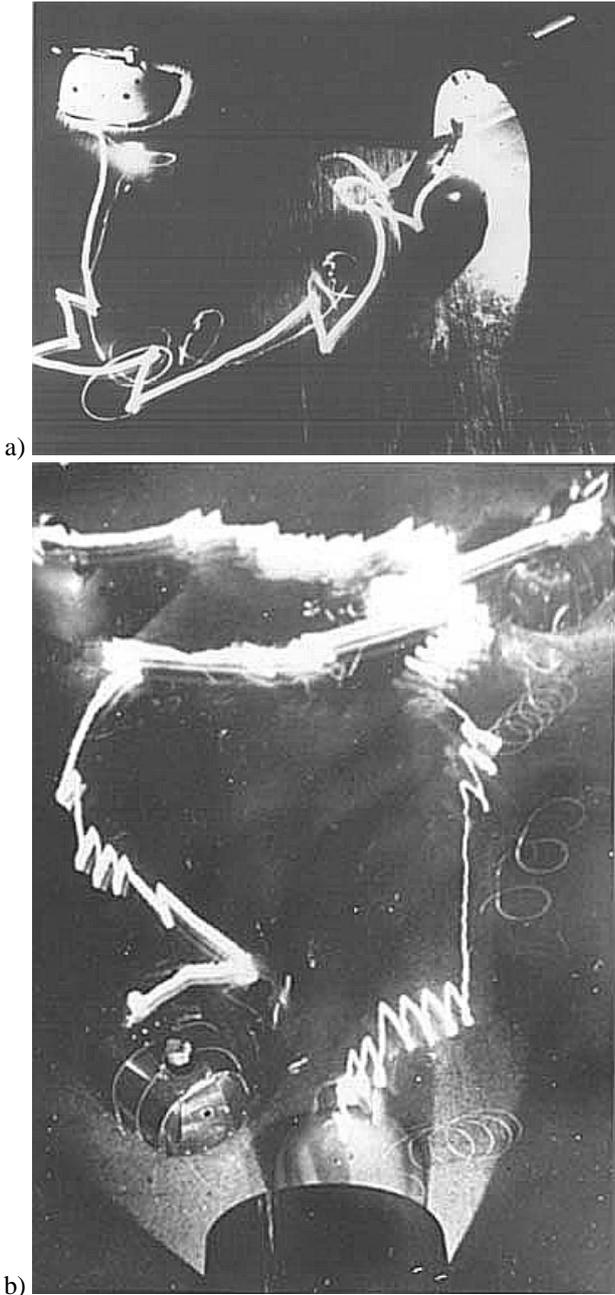


Figure 1: **The dance of the turtles.** In a), a single tortoise returns to its hutch. In b), two tortoises affect each other's movements due to a light source attached to each.

behaviour. Lesion studies have a long and respected tradition in neuroscience and evolutionary developmental biology, and have recently been proposed as a systematic method for understanding neural network behaviour (Aharonov et al., 2001).

2. Methods

Behaviours for a generic quadrupedal agent were evolved and analyzed in a physics-based, three-dimensional simulation environment¹. This environment simulates both the internal and external forces acting on the agent and objects in its environment, as well as various other physical properties such as contacts between the agent and the ground, and torque applied by the motors to the joints.

The agent, composed of 23 rigid components (12 spheres and 11 cylinders), is shown in Figure 2. It contains 8 one degree-of-freedom hinge joints, one in each of the knees (**J1**, **J2**, **J7** and **J8** in Fig. 2 c)) a pair in the shoulder sphere (**J3**, **J4**), and a pair in the pelvis (**J5**, **J6**). All of the joints have an axis of rotation lying in the horizontal plane. The two antennae are rigidly attached to the body, and thus cannot move independently of the body's motion. For simplicity, each body component has a mass of 1kg. The body spheres have radii of 20cm, and the antennal spheres radii of 10cm. The body cylinders have radii of 10cm and lengths of 50cm, and the antennal cylinders have a radii of 5cm, and lengths of 1.5m. Note that the lengths and masses do not approach those of any biological organism, but are important only in their magnitudes relative to each other, and the strength of the simulated, actuating motors.

The two-dimensional chemical gradient field² through which the agent moves is static; local chemical concentrations at each point do not change during the evaluation of an agent's behaviour. The gradient lies along a 12 by 12 meter square, and the agent is evaluated in four different chemical environments: ones in which a chemical point source is placed at four evenly spaced locations along the forward boundary of the gradient field (this can be seen most clearly in Fig. 5). The field is broken up into 400 discrete cells, 20 along each side. Within each cell, the chemical concentration can range between 0.0 (no chemical) and 1.0 (complete saturation). The cell containing the point source has a concentration of 1.0. All other cells contain a concentration of

$$1 - \frac{d}{\sqrt{2}s^2},$$

where d is the cell's distance from the point source, and s is the length of the gradient field. This ensures that there is a smooth, linear decay in chemical concentration out from the point source, and that the cell lying diametrically opposite to the point source (when the point source lies in one of the forward corners) has zero concentration.

Each agent contains a total of four touch sensors, four angle sensors, two chemosensors and eight actuated joints. One touch sensor is located in each of the feet, one angle sensor is located in each of the four shoulder and pelvic joints (but not in the knee joints), and the two chemosensors are placed in the left and right antennae tips.

¹Final beta release of MathEngine SDK; www.cm-labs.com

²Note that the gradient field is referred to as a chemical gradient, but could also be interpreted as a differential field of some other substance, such as light.

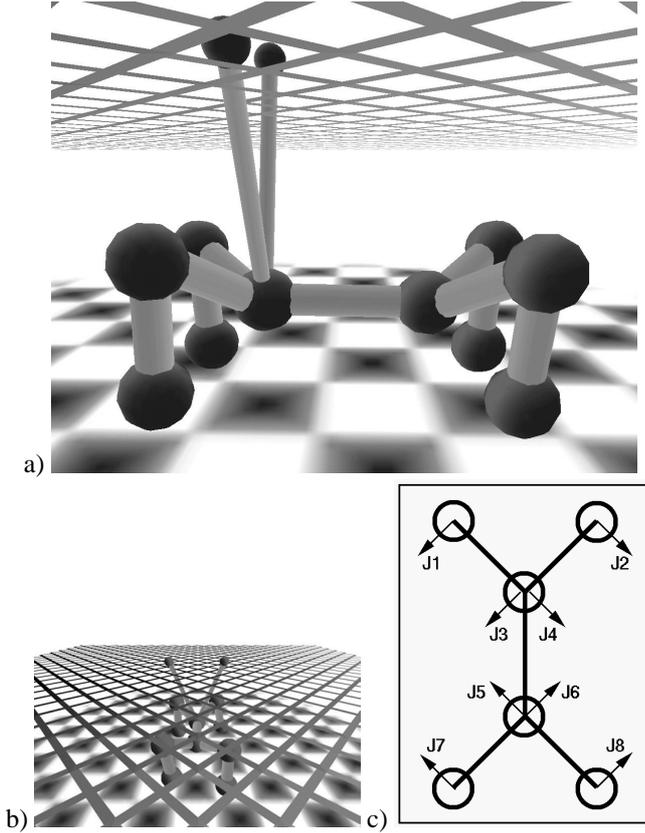


Figure 2: **The agent.** a) Side view. b) Top view. The two-dimensional gradient field is shown as a cross hatched pattern; darker lines indicate areas of higher concentration. In these images, the chemical point source lies in the front-left corner of the gradient field. c) The placement and axes of rotation for the eight actuated joints.

The touch sensors return a maximum positive signal if the body part in which they are contained is in contact with the ground plane, and return a maximum negative signal otherwise. The angle sensors return a signal commensurate with the joint's current angle. For example, the angle sensors emit a maximum negative signal when the joint to which they are attached is at maximum counterclockwise rotation, a zero value when the joint angle is equal to the original setting, and a maximum positive signal when the joint is at maximum clockwise rotation. The chemosensors return the chemical concentration found in the cell lying directly above or below them. This is necessary because although the agent moves in a three-dimensional environment, the gradient field is two-dimensional.

The joints can rotate between -30 and 30 degrees of their original setting. Each of these joints is actuated by a torsional motor, which receives desired angle settings from a neural controller, and exerts torque proportional to the difference between the current joint angle and the desired angle using

$$\tau_{t+1} = \max(I(\omega_t - k(\theta - \theta_d)), \tau_{max}),$$

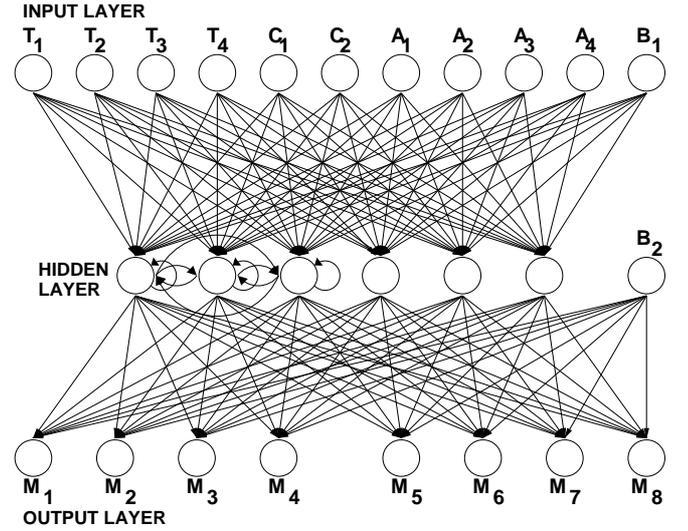


Figure 3: **The neural network architecture.** The four touch sensor signals are scaled and passed to input neurons T1–T4, the chemosensors are scaled and passed to input neurons C1–C2, and the angle sensors are scaled and passed to input neurons A1–A4. The output neuron values (M1–M8) are translated from desired angles into torque by the eight motors of the agent. Note that only the recurrent connections for the first three hidden neurons are shown.

where θ is the actual joint angle, θ_d is the desired joint angle, τ_{max} is the maximum torque ceiling, $\omega = \dot{\theta}$, and I is the inertia matrix.

All eight motors have the same maximum torque ceiling, as well as the same damping properties, which were tuned by hand to disallow extreme actions such as jumping or hopping. However, combined motor action was sufficient for walking, and in some cases dynamic gaits in which the agent's centre of mass passed outside of the support polygon created by its contacts with the ground plane emerged.

All of the agents are controlled by a partially recurrent neural network, the architecture of which is shown in Fig. 3. The input and output layers correspond to the sensor and motor array, respectively. There is an additional bias neuron at the input and hidden layers that outputs a constant signal of 1. The input layer is fully connected to a hidden layer containing six hidden neurons, and the hidden layer is fully connected to the output layer. In addition, the hidden layer is fully, recurrently connected. The additional recurrent synapses were added in order to allow for the generation of oscillatory signals partly or completely independent of the incoming sensor signals, if required.

At each time step of the simulation of an agent's behaviour, the eight sensor signals are scaled to floating-point values in $[-1.0, 1.0]$, and supplied to the input layer. The values are propagated to the hidden and output neurons. The hidden and output neurons scale their incoming values using the activation function

$$O = \frac{2}{1 + e^{-a}} - 1,$$

where a is the summed input to the neuron.

A fixed length, generational genetic algorithm is used to evolve behaviours for the agent. Genomes encode the 158 synaptic weights for the neural network as floating-point values, which can range between -1.00 and 1.00 . For the experiments reported in the next section, each evolutionary run was conducted using a population size of 200, and was run for 50 generations. At the end of each generation, strong elitism was employed: the 100 fittest genomes were copied into the next generation. Tournament selection, with a tournament size of 3, is employed to select genomes from the population to participate in mutation and crossover. Twenty-five pairwise one-point crossings produce 50 new genomes. The remaining 50 new genomes are mutated copies of genomes selected from the previous generation: an average of three point mutations are introduced into each of these new genomes, using random replacement.

Each genome was assigned a fitness using the following procedure. First, the synapses are labelled using the values encoded in the genome. The agent is then placed at the origin in the simulation, and allowed to behave for 500 time steps. During each time step of the evaluation, sensor readings are taken, the neural network is updated, and the motor commands are translated into torques. Also the body parts' positions, velocities and orientations are updated based on these torques as well as on external forces such as gravity, inertia, friction and collision or contact with the ground plane. At the end of this period, the agent is returned to the origin, and the chemical point source is moved and the gradient field is updated. The agent is then given another 500 time steps in which to behave. This procedure is repeated for each of the four point source locations, and the agent's fitness is given as the sum of the distances between the point source and the agent's centre of mass at the end of each of the four evaluations.

3. Results

Ten independent evolutionary runs were performed, starting with different, random starting populations. In all 10 runs, the agents were able to achieve successful chemotaxis: the final, most successful neural network from each population induced the agent to walk towards the four point sources in different locations. Fig. 4 shows the evolutionary curves for a typical population.

The most successful neural network produced by this run was then lesioned: the agent was evaluated again, but the signals returned by the chemosensors were suppressed, and zero values were returned instead. Fig. 5 shows the original trajectory of the agent for the four point sources, as well as the trajectory obtained from the lesioned network.

Two other evolved neural networks from the same evolutionary run were tested: the most successful network produced in generation 15, and generation 25. For all three networks, three lesion experiments were then performed. First only the left-hand chemosensor was lesioned, then only the

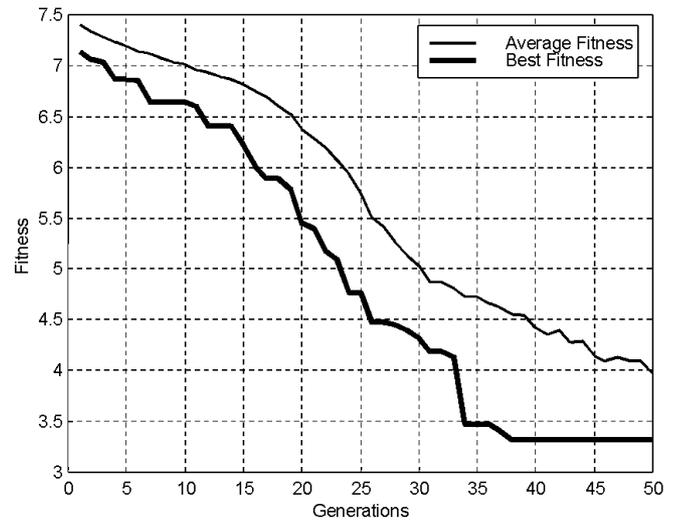


Figure 4: **Evolutionary change in a typical population.** The thin line indicates the average fitness of the population; the thick line indicates the fitness of the most successful neural network in the population at that generation.

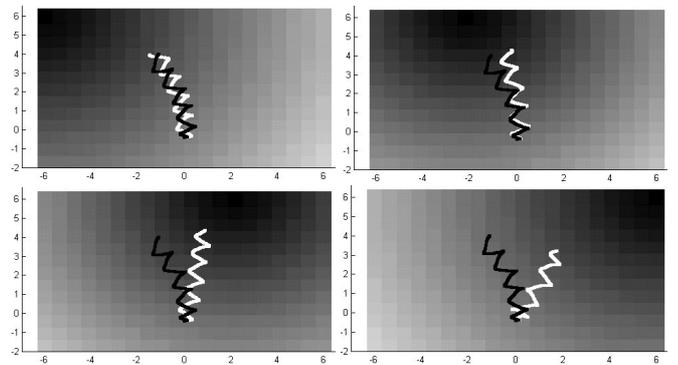


Figure 5: **Typical and lesioned trajectories.** The gradient field is shown: darker patches indicate higher chemical concentration. The white line indicates the trajectory of the evolved agent's centre of mass. The black line denotes the trajectory of the agent when the chemical sensory signals are suppressed. Note that only the horizontal component of the agent's trajectory is shown. The axes indicate the distance (in meters) from the agent's starting point.

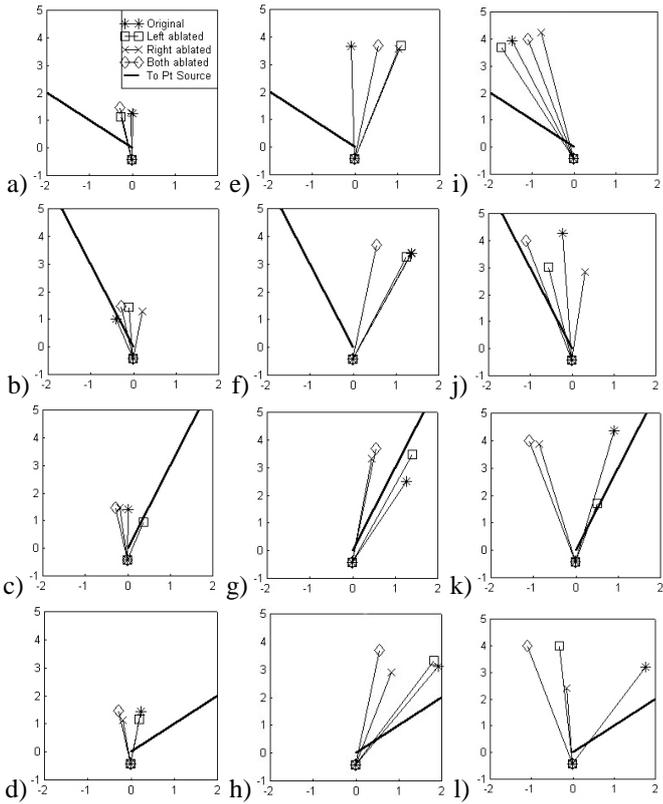


Figure 6: **Chemosensor lesions.** a), b), c) and d) indicate the trajectories induced by the best network taken from generation 15. e), f), g) and h) indicate the trajectories induced by the network taken from generation 25. i), j), k) and l) indicate the trajectories induced by the best network taken from the final generation. The thick lines point towards the chemical point source. The axes indicate distance (in meters) away from the agent's initial position.

right-hand chemosensor was lesioned, and finally both were lesioned together. The resulting trajectories are reported in Fig. 6, but are represented as vectors; the origin of the vector indicates the start point of the agent, and the end point indicates the agent's final position.

A second set of lesion experiments were then performed on these three networks, in which each sensory modality was lesioned in turn. First the entire set of touch sensors were lesioned together, then the entire set of angle sensors were lesioned together, and finally, again, the two chemosensors were lesioned together. The resulting trajectories are shown in Fig. 7.

Finally, a third set of lesion experiments was conducted. Each of the six hidden neurons was lesioned in turn. That is, for each time step of evaluation, the actual value output by the lesioned hidden neuron is suppressed, and a zero value is output instead. The resulting trajectories are reported in Fig. 8.

The most successful neural networks were then taken from two other successful evolutionary runs, and were lesioned. Fig. 9 shows the resulting trajectories when first the left-

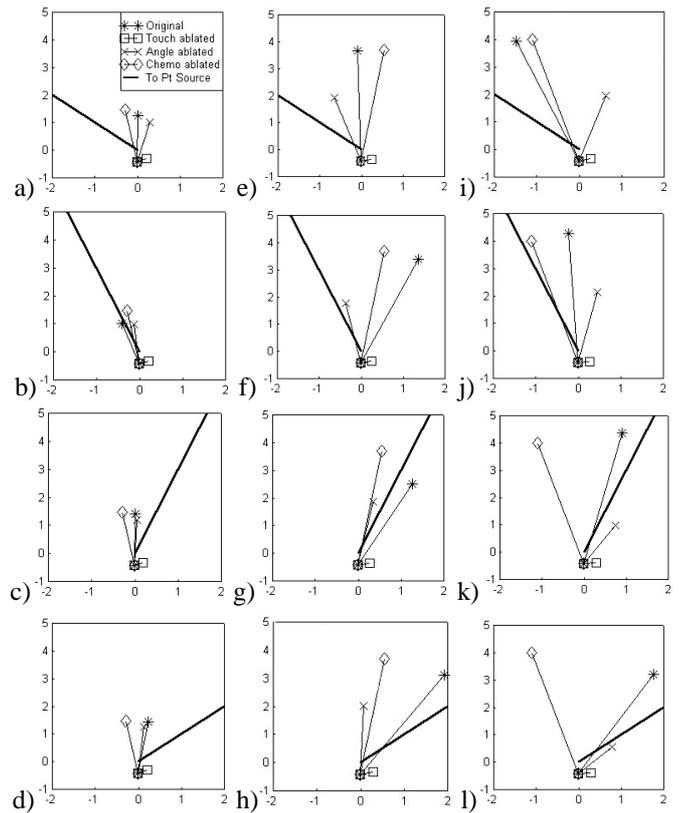


Figure 7: **Sensor modality lesions.** a), b), c) and d) indicate the trajectories induced by the best network taken from generation 15. e), f), g) and h) indicate the trajectories induced by the network taken from generation 25. i), j), k) and l) indicate the trajectories induced by the best network taken from the final generation.

hand, then the right-hand, and finally both chemosensors are lesioned in these two networks. Fig. 10 reports the trajectories when the three different sensor modalities are lesioned in these two networks.

4. Discussion

As can be seen from Fig. 5, forward locomotion is maintained, but directional locomotion towards the chemical point source is lost when both chemosensors are lesioned. This indicates that either one or both of the chemosensors modify the agent's direction of travel, but do not themselves drive the locomotory gait. Fig. 6 shows then when either of the chemosensors is lesioned in the most successful network (i), j), k) and l)), the trajectory deviates from the original one, which reveals that both chemosensors play a role in determining the agent's direction. Further, this behavioural effect can be seen in the two other ancestor networks, taken from generations 15 and 25. This suggests that either historical accident involved both chemosensors in changing the agent's direction early on, or that for this particular experimental regime both chemosensors are necessary for changes in direction.

Fig. 7 presents a slightly different picture. Here, lesioning

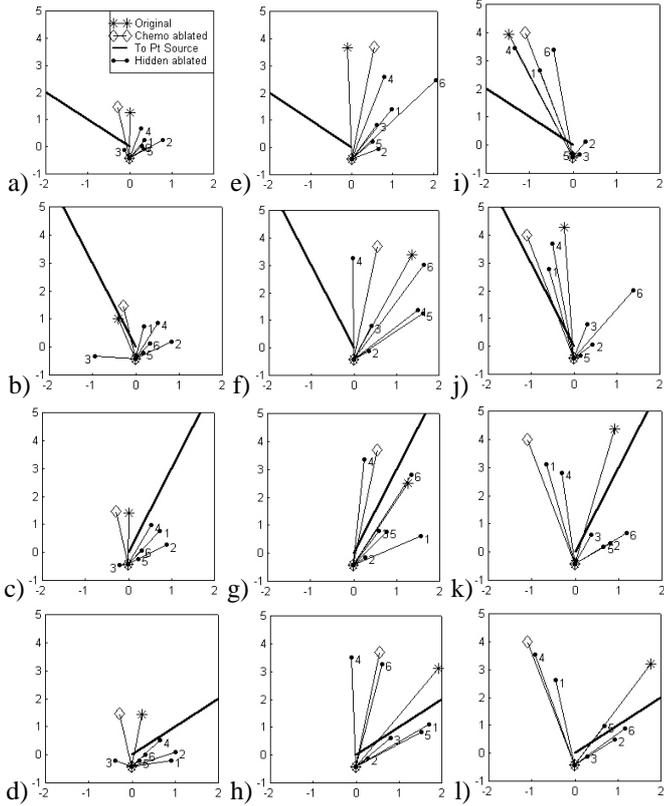


Figure 8: **Hidden neuron lesions.** a), b), c) and d) indicate the trajectories induced by the best network taken from generation 15. e), f), g) and h) indicate the trajectories induced by the network taken from generation 25. i), j), k) and l) indicate the trajectories induced by the best network taken from the final generation. The numbered vectors indicate the trajectory produced when the corresponding hidden neuron was lesioned (i.e., vector '1' is the trajectory produced when the first hidden neuron is lesioned).

of the touch sensors completely disrupts locomotion, for all three networks. Lesioning of the angle sensor group, though, partially, but not completely, degrades locomotion for the networks taken from generations 25 and 50, and has no appreciable effect on the behaviour generated by the network taken from generation 15. This indicates that touch sensors were from the beginning the main generators of locomotion in this population, but that angle sensor signals were only gradually appropriated over evolutionary time to improve locomotion, and has less of a role to play than touch sensor signals.

By lesioning the hidden neurons, it is possible to gain some insight into how the evolved networks combine and dissociate information arriving from the sensors. Fig. 8 shows that for the most successful network, lesioning of the first or the fourth hidden neuron produces trajectories very similar to those obtained by lesioning both chemosensors. This suggests that these two hidden neurons either separately, or in concert, modulate the underlying forward locomotion behaviour, which is presumably controlled by the other four hid-

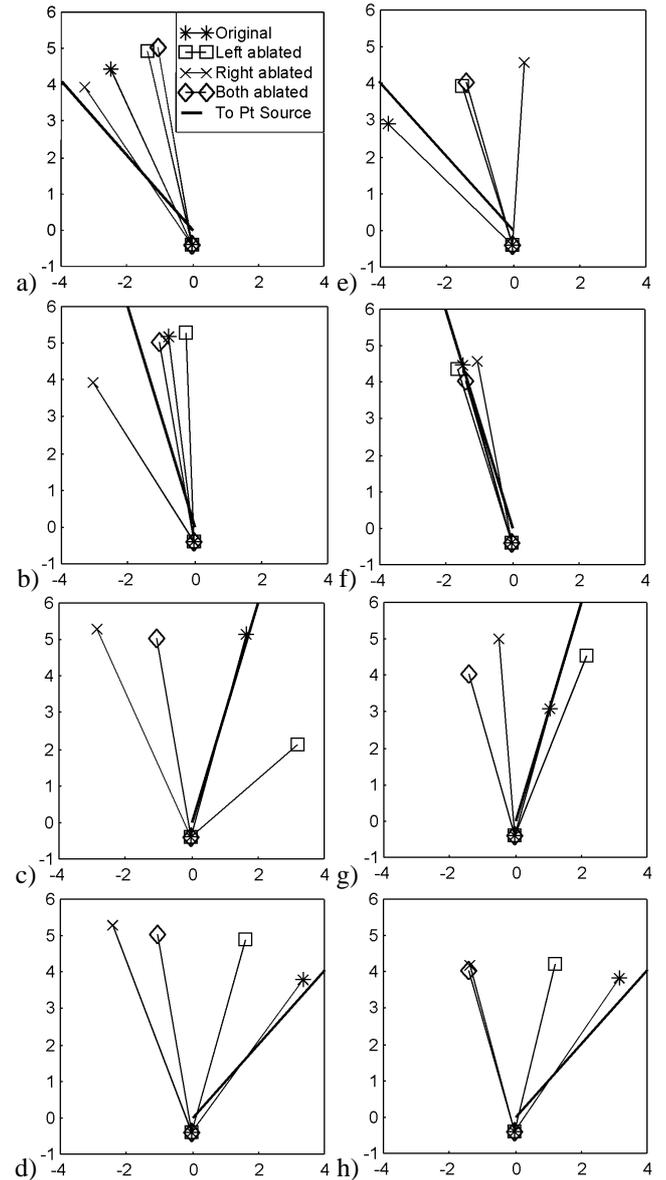


Figure 9: **Chemosensor lesioning in other evolved populations.** The effects of lesioning individual and both chemosensors in the most successful networks produced by two other evolutionary runs. a), b), c) and d) show the trajectories for one evolutionary run, and e), f), g) and h) show the trajectories for the other run. Note the increase in axes length, compared to those in the previous three figures to accommodate the longer trajectories of these more successful runs.

den neurons. Further, it can be seen that for the network taken from generation 15, there is no similarity between the resulting trajectories when any hidden neuron is lesioned, suggesting that no hidden neuron has yet specialized to process the incoming chemical sensory signals. For the network taken from generation 25, there is a weak similarity between the trajectories produced by lesioning hidden neuron 4, and lesioning the chemosensors, indicating that hidden neuron 4 was first appropriated to handle incoming chemical sensory

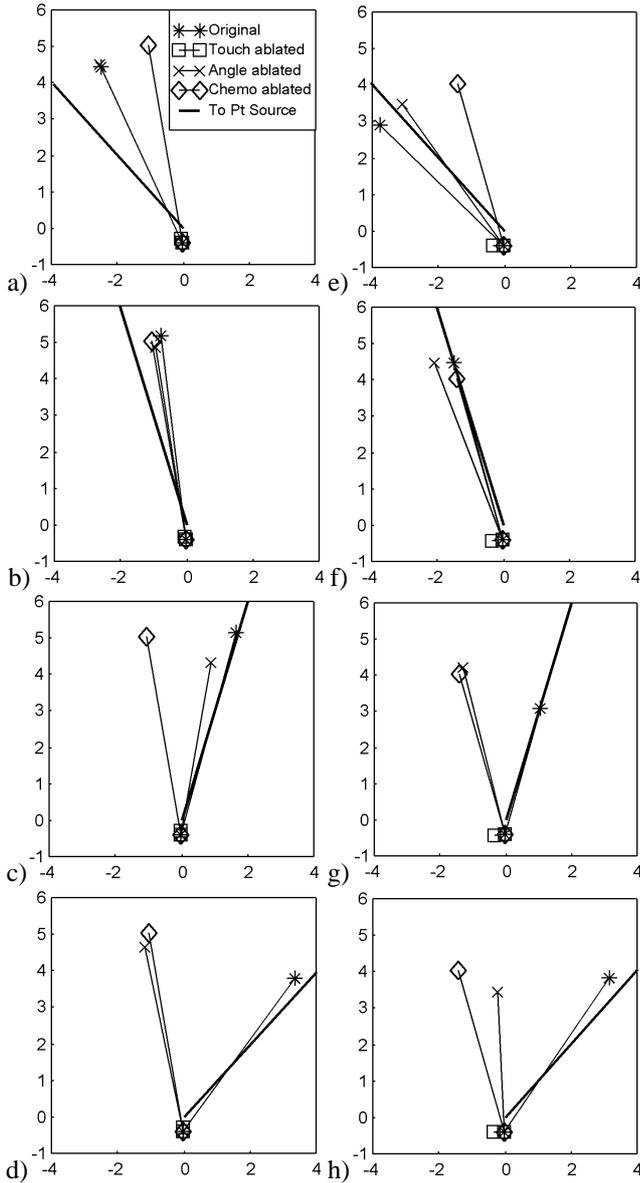


Figure 10: **Sensor modality lesioning in other evolved populations.** The effects of lesioning entire sensor modalities separately in the most successful networks produced by two other evolutionary runs. a), b), c) and d) show the trajectories for one evolutionary run, and e), f), g) and h) show the trajectories for the other run.

signals.

The chemosensor lesion experiments for two other evolved neural networks (shown in Fig. 9) seems to suggest that both chemosensors are required for chemotaxis (as formulated in these experiments), and is not the result of historical accident during the evolutionary process. An alternative hypothesis is that chemotaxis driven by paired chemosensor readings is easier for artificial evolution to discover than chemotaxis driven by differential readings of a single chemosensor over time. Although the recurrent links at the hidden layer do allow for comparisons against sensor readings taken during pre-

vious time steps, a more generic network that can more easily gauge temporal changes in sensory readings might produce chemotaxis that relies on only a single chemosensor. Also, because only the direction of travel is affected by lesioning both chemosensors in both networks, not the distance travelled, this seems to suggest that the genetic algorithm invariably evolves networks in which the underlying locomotory gait is driven by sensors other than the chemical sensors.

When the different sensory modalities were lesioned for these two networks (see Fig. 10), it can be seen that again, loss of touch sensor signals completely disrupts locomotion. However, lesioning of the angle sensors does not seem to impede locomotion; the agent walks just as far as when it is driven by the non-lesioned network. This suggests that for the original population, the fusing of touch and joint angle sensory signals for driving locomotion was a historical accident (i.e., it does not appear in every evolutionary run), and is not necessary for the achievement of the underlying behaviour of forward locomotion. However, the change in direction when the angle sensors are lesioned suggests that at the hidden or output layer, joint angle and chemical information is somehow combined; the reasons for this are not immediately clear, but are worthy of further study.

5. Conclusions

This paper has documented how artificial evolution can be used to produce gradient-following behaviours, a type of behaviour first studied in autonomous agents by Grey Walter over 50 years ago. Here, through lesion experiments, we have investigated how artificial evolution uses the sensory modalities made available to it to produce such behaviours. This investigation has revealed an interesting dynamic, namely that artificial evolution produces neural networks that modulate basal behaviours (here, forward locomotion) using sensory modalities separate from those that drive the basal behaviours. By lesioning hidden neurons, it was found that this dissociation between different sensory modalities extends to the hidden layer as well.

Secondly, we have shown that a behaviour that is driven by a single sensor modality early during evolution (here, locomotion driven by touch sensors), can come to be driven by a combination of more than one modality (here, touch and joint angle sensors).

Future experiments are planned in which the chemical environment is extended to three dimensions, and the gradient field is animated by the simulation of hydrodynamics and/or turbulence. Also, by subjugating more aspects of the agent to evolutionary control—such as its neural architecture, sensory apparatus and body shape—it may be possible to study how different biological species evolved to exploit chemical gradients in their environments. Finally, it would be useful to perform these experiments on different types of agents in different task environments in order to learn whether, and how, automatic sensor fusion and dissociation generalizes beyond gradient-following behaviours.

References

- Aharonov, R., Meilijson, I., and Ruppin, E. (2001). Understanding the agent's brain: A quantitative approach. In Keleman, J. and Sosik, P., (Eds.), *Sixth European Conference on Artificial Life*, pages 216–225.
- Braitenberg, V. (1986). *Vehicles*. MIT Press.
- Brooks, R. A. (1991). New approaches to robotics. *Science*, 253:1227–1232.
- Dorigo, M. and Caro, G. D. (1999). The ant colony optimization meta-heuristic. In Corne, D., Dorigo, M., and Glover, F., (Eds.), *New Ideas in Optimization*, pages 11–32. McGraw-Hill.
- Ferree, T. C., Marcotte, B. A., and Lockery, R. (1997). Neural network models of chemotaxis in the nematode *C. elegans*. In *Advances in Neural Information Processing Systems*, volume 9, pages 55–61, Colorado, USA. MIT Press.
- Grasso, F., Consi, T., Mountain, D., and Atema, J. (1996). Locating odor sources in turbulence with a lobster inspired robot. In Maes, P., Mataric, M., Meyer, J.-A., Pollack, J., and Wilson, S., (Eds.), *Proceedings of SAB'96, From Animals to Animats 4*, pages 104–112, Cape Cod, USA. MIT Press.
- Grey Walter, W. (1950). An imitation of life. *Scientific American*, 182(5):42–45.
- Harvey, I., Husbands, P., Cliff, D., Thompson, A., and Jakobi, N. (1997). Evolutionary robotics: the Sussex approach. *Robotics and Autonomous Systems*, 20:205–224.
- Holland, O. and Melhuish, C. (1999). Stigmergy, self-organization, and sorting in collective robotics. *Artificial Life*, 5:173–202.
- Ijsspeert, A. J. and Arbib, M. (2000). Visual tracking in simulated salamander locomotion. In Meyer, J. A. and Berthoz, A., (Eds.), *Proceedings of SAB'00, From Animals to Animats 6*, pages 88–97, Paris, France.
- Kodjabachian, J. and Meyer, J.-A. (1998). Evolution and development of neural controllers for locomotion, gradient-following and obstacle-avoidance in artificial insects. *IEEE Transactions on Neural Networks*, 9(5):796–812.
- Reil, T. and Massey, C. (2001). Biologically inspired control of physically simulated bipeds. *Theory in Biosciences*, 120:1–13.
- Shimojo, S. and Shams, L. (2001). Sensory modalities and not separate modalities: plasticity and interactions. *Current Opinion in Neurobiology*, 11:505–509.