

Embodiment and Perceptual Crossing in 2D

A Comparative Evolutionary Robotics Study

Marieke Rohde and Ezequiel Di Paolo

Centre for Computational Neuroscience and Robotics (CCNR)
University of Sussex, Brighton, BN1 9QH, UK
{m.rohde,ezequiel}@sussex.ac.uk

Abstract. We present the results from an evolutionary robotics simulation model of a recent unpublished experiment on human perceptual crossing in a minimal virtual two-dimensional environment. These experiments demonstrate that the participants reliably engage in rhythmic interaction with each other, moving along a line. Comparing three types of evolved agents with radically different embodiment (a simulated arm, a two-wheeled robot and an agent generating a velocity vector in Euclidean space), we identify differences in evolved behaviours and structural invariants of the task across embodiments. The simulation results open an interesting perspective on the experimental study and generate hypotheses about the role of arm morphology for the behaviour observed.

1 Introduction

In a recent unpublished study, Lenay et al. (personal communication) from the perceptual supplementation group (GSP) at the UT Compiègne tested human subjects on their capacity to discriminate static and mobile objects from other intentional sensing entities in a minimal two-dimensional virtual environment. This paper presents the results from an evolutionary robotics simulation model of this experiment that aims at exploring the space of possible behavioural strategies afforded by the experimental design, in order to help explain their results.

In the experiment, the blindfolded participants' task was to indicate via mouse-clicks whether an object encountered by moving the mouse (tactile stimulus to the finger) was another sensing entity (i.e. another participant) or not, where the only two other objects in the toroidal simulated two-dimensional environment were a mobile lure that shadows the other participants' movements and a static lure (see Sect. 2 for details of the task). This paradigm was a direct extension of a previous experiment on perceptual crossing in a one-dimensional simulated environment by the same group [1], to test whether the experimental results transfer qualitatively or quantitatively from a one-dimensional to a two-dimensional scenario. Some preliminary results from their study are that the results transfer qualitatively (i.e. 65% correct clicks), and that the behavioural strategies are strikingly and unexpectedly similar to those observed in the 1D version of the experiment. In particular, even though participants search for interaction exploring both dimensions, they move back and forth on a line once they encounter an interaction partner or object.

Having shown in our previous work [2, 3] that evolutionary robotics simulations of this kind of minimal sensorimotor experiments (in particular of the one-dimensional version of the same task [2]) can be a rich source of clarification, inspiration and illustration, we conducted the present simulation in order to understand the dynamical principles of the task and to generate hypotheses for the analysis and evaluation of the experimental data based on the analogy with the simpler and more controllable simulated agents.

Our intuition was that the rhythmic one-dimensional interaction is related to the morphology of the human arm. To explore this intuition, we compared a simple model of a human arm moving a mouse on a desk surface with two other types of agents that are controlled by the same control network, but have radically different bodies and sensorimotor dynamics: a two-wheeled robotic agent and an agent generating a velocity vector in Euclidean space with a vertical and a horizontal component (similar to a joystick; model described in Sect. 2).

The results (Sect. 3) allow to identify interesting differences and commonalities between the solutions evolved for these different types of agents and the results from the original experiment. Firstly, the *same behavioural patterns* (e.g., independent realisation of search and interaction, rhythmic oscillations) evolve for all agents. Secondly, the *realisation* of these behaviours varies a lot with morphology, which includes the production of near-linear rhythmic trajectories in arm agents, as hypothesised. Thirdly, the evolution of a *viable yet unintuitive strategy* (avoiding interaction with each other) in some of the agents. This refutation of our hypothesis that the results would qualitatively resemble those from the experiment is not a shortcoming: the existence of a counter-intuitive efficient strategy opens a different perspective on the experimental data (see Sect. 4).

Analysing sensorimotor data from experimental studies like the one we modelled is a difficult task because of the multitude of factors influencing human behaviour. Exploring the space of possible solutions in simulation (the task and the simulated environment used here are, apart from parametric details, identical with those used in the experiment), we hope that our results will directly facilitate the data analysis and interpretation for the researchers at the GSP.

2 The Model

The artificial agents evolved, just as the experimental participants, move around a virtual plane (200×200) that wraps around in both dimensions (i.e., a torus; see Fig. 1 (A)). In this plane, there are six different objects. Two circular simulated agents (circular objects in Fig. 1 (A)), two mobile lures that are attached to the agents and two fixed lures that are statically installed at (50, 50) and (150, 150) respectively. All objects are circular of diameter 20, even if the lures are represented as boxes in Fig. 1 (A) and other figures. The attached lures shadow the trajectories of each of the agents at a distance of 93 units, being attached in perpendicular directions (see Fig. 1 (A)).

The only sensory signal S that the agents receive is a touch signal, i.e. if the distance d between the agent and something else is $d < 20$, an input S_G




Fig. 1. Schematic diagram of the simulation model and control network. (A) The simulated environment with the two agents (circles), the attached lures (boxes attached with a line) and the fixed lures (boxes). Diagrams of the two-wheeled agent (B), the agent moving in Euclidean space (C) and the simulated arm agent, with the space in which they can act (D). The control network (E).

(sensory gain, evolved) is fed into the control network. Each agent can only perceive the other and one of each kind of lure, i.e., the dark agent can perceive all light objects in Fig. 1 (A), but not the dark ones, and vice versa, in order for interaction between the agents not to be mediated through another object.

Three different types of agents were evolved: 1.) A two-wheeled agent that generates the velocity $v_{l,r} \in [-20M_G, 20M_G]$ for each wheel (Fig. 1 (B); M_G is the evolved motor gain; velocities are specified in units/s). 2.) An agent that generates a horizontal and a vertical velocity vector $v_{h,v} \in [-30M_G, 30M_G]$ that are summed up ('Euclidean agent', Fig. 1 (C)). 3.) A simple simulated arm with two segments of length 400 units that is steered through angular velocity signals $\omega_{e,s} \in [-0.05M_G, 0.05M_G]$ to the elbow and the shoulder joint (see Fig. 1 (D)).

In order to approximate the dynamics of human mouse motion, the arm agent is restricted in its movements in two ways: through joint stops $\alpha_s \in [0.1\pi, 0.6\pi]$ and $\alpha_e \in [0.2\pi, \pi]$ and through the delimitation of movement to an area of 600×600 units that represents the 'desk' surface (i.e., the area within which a human participant would move the mouse), whose bottom left corner is fixed at $(-200, 200)$ taking the shoulder joint as the origin. The desk area is translated randomly with respect to both the desk area of the other agent and the simulated virtual environment (see Fig. 1 (D)) to avoid that agents evolve to meet in the middle of the desk.

For purpose of comparison, all three kinds of agent are controlled by structurally identical neural network controllers (Fig. 1 (E)), i.e., continuous-time recurrent neural network (CTRNN, see e.g., [4]) with one input neuron, four fully connected interneurons and 5 output neurons. Four output neurons regulate the two motor outputs ($M_1 = a_{M1} - a_{M2}, M_2 = a_{M3} - a_{M4}$) that are interpreted as $v_{l,r}, v_{h,v}$ or $\omega_{e,s}$ respectively. The fifth output neuron generates

the categorisation signal M_C . CTRNN dynamics are governed by

$$\tau_i \frac{da_i(t)}{dt} = -a_i(t) + \sum_{j=1}^N w_{ij} \sigma(a_j(t) + \theta_j) + I_i \quad (1)$$

where $\sigma(x) = 1/(1 + e^{-x})$ is the standard sigmoid function, θ_i a bias term, τ_i the activity decay constant and w_{ij} the strength of a connection from unit j to unit i . The parameters evolved (74 parameters) are: $S_G, M_G \in [1, 50]$. $\tau_i \in [20, 3000]$, $\theta_i \in [-3, 3]$ and $w_{i,j} \in [-6, 6]$.

The dynamics are simulated using the forward Euler method with a time step of 1 ms. All three kinds of agents were evolved with and without a 100 ms sensory delay. A modified version of the arm agent with three sensory neurons received the joint angles as additional proprioceptive inputs ($S_{2,3} = S_G \alpha_{e,s}$).

Each trial lasts $T \in [6000, 9000]$ ms. The task is to interact with something and correctly classify the object encountered as either of the lures ($M_C \leq 0.5$) or the other agent ($M_C > 0.5$). Agents are matched with clones of themselves in the task. The starting positions are random for the wheeled and the Euclidean agent and random within the centre area for the arm agent. The starting angle for the wheeled agents is random. For the arm agent and the Euclidean agent, the relative orientation of the agents to each other is random $\in \{-\frac{\pi}{2}, 0, \frac{\pi}{2}, \pi\}$. The fitness $F(i)$ of an individual i in each trial is given by the following function

$$F(i) = \begin{cases} 1 & \text{if } (d_s \leq D) \wedge (d_o > D) \wedge (M_C > 0.5) \text{ (true positive)} \\ 1 & \text{if } (d_s > D) \wedge (d_o \leq D) \wedge (M_C \leq 0.5) \text{ (true negative)} \\ 0.25 & \text{if } (d_o < D) \wedge (d_s < D) \text{ (ambiguity)} \\ 0.1 & \text{if false classification and } S > 0 \text{ (touch)} \\ 0 & \text{else} \end{cases} \quad (2)$$

where $D = 30$, d_o the distance to the closest of the two lures and d_s the distance to the other agent. Agents are tested on eight trials and fitness is averaged.

The search algorithm used to evolve the parameters of control networks (1000 generations) is a generational genetic algorithm (GA) using a population of 30 individuals with a real-valued genotype $\in [0, 1]$, truncation selection (1/3), vector mutation of magnitude $r = 0.6$ and reflection at the gene boundaries. The weights w_{ij} and the bias θ_i are mapped linearly to the target range, the sensor gain S_G , the motor gain M_G and the time constants τ_i are mapped exponentially.

3 Results

3.1 Evolvability

All agents evolved to a higher level of performance with delays than without (see Fig. 2 (A)). A similar counter-intuitive benefit of sensory delays had already been observed in the one-dimensional scenario [2]. As explained in the following section 3.2, this is because overshooting of the target bootstraps the evolution of active perceptual strategies.




Fig. 2. (A): Population fitness average \bar{F} (mean and maximum of 10 evolutionary runs). (B): Performance average across 100 evaluations for the best individual from the best evolution. Light: no delay, dark: 100 ms sensory delay. (C) Fitness for rhythmic solutions (white) is on average much higher than that for non-rhythmic solutions (grey). (No rhythmic activity was evolved for Euclidean or arm agents without delay; note that the measure for rhythmicity is an approximation as explained in section 3.2.)

The wheeled agent and the Euclidean agent evolve to a much higher level of performance (see Fig. 2 (A)), with the best individual from the best evolutionary run achieving nearly perfect performance (Fig. 2 (B)), whilst even the best evolved arm agent stays well below the 50 % performance level. As explained in Sect. 3.2, this is because evolution of search strategies is much more difficult for the arm agent than for the other two.

3.2 Evolved Behaviours

Two large classes of behaviour dominate the fitness landscape for the evolved task, irrespective of embodiment. The more successful strategy (1) is to always output ‘no’ ($M_C \leq 0.5$) and seek contact with the fixed lure, avoiding any interaction with any mobile entity (up to perfect fitness). This perfectly viable yet slightly ‘autistic’ behaviour differs a lot from the participants’ behaviour, who avoid the fixed lure and seek interaction with each other. This discrepancy is interesting, not lamentable, because it illustrates alternative solutions afforded by the paradigm. Replicating human behaviour was not the main objective. The other predominating strategy (2) is to interact with anything indiscriminately and constantly output ‘yes’ ($M_C > 0.5$) and yields a fitness of up to ca. 40%. The arm agents nearly exclusively evolve strategy (2), whereas the Euclidean and wheeled agents evolve strategy (1), frequently passing during evolution through a phase of strategy (2). What evolved, therefore, were preferences rather than discriminatory capacity: only four agents (one arm, one wheeled, two Euclidean) evolved additionally a classification output contingent on stimulation.

Both strategy (1) and strategy (2) can in principle be realised by rhythmical interaction with the target or simply by halting. It appears that rhythmic behaviour is more effective: as an approximation, let rhythmic behaviour be activity confined to a radius of $d = 50$ around an entity with least five inversions of

sensory state during the last second of a trial. Within each condition for which both rhythmic and non-rhythmic solutions evolved, the rhythmic ones were on average 9% more successful (see also Fig. 2 (C); some rhythmically interacting arm agents that were not captured by the approximate measure). The reason for this advantage is that an agent that evolves to simply stop will be clueless upon unexpected cessation of the stimulus, e.g. when crossing the object at an unfortunate angle, and therefore re-start the search for sensation. An agent that interacts with an object rhythmically and stays in touch with its boundary can reverse the effect of actions that make stimulation disappear and return to where it last had been stimulated. This minimal *spatialisation* increases the probability to re-encounter a lost object.

Sensory delays seem to be crucially involved in the evolution of rhythmic behaviour, accounting for their evolutionary advantage: We compared the 60 best individuals from all evolutionary runs (across bodies). Only two of the 30 individuals evolved without delays behaved rhythmically at least once in 10 trials (opposed to 16 out of 30 evolved with delay). Sensory delays lead the agents to overshoot an object, which triggers the evolution of return trajectories. This, in turn, facilitates the evolution of rhythmic interaction as effective and active perceptual strategy, and thus helps to overcome the local fitness maximum to stop upon any stimulation and start the search anew if stimulation ceases.

Despite these abstract commonalities, the exact realisation and behavioural dynamics vary quite significantly between the different agents. The following subsections analyse in detail the strategies evolved (with delays) to explore the strategy space and how it is constrained by different embodiments.

The Two-Wheeled Agent. Two behavioural phases that can be variably realised independent of each other can be identified: search and interaction. Wheeled agents evolved a variety of strategies to establish contact: some shoot off straight, others drive around in large circles, yet others in arches or spirals. Interaction is initiated immediately on contact, or, otherwise, the agent backs off and comes back to see if the stimulating object is still there. All wheeled agents evolved to drive in circles (of variable size) around the encountered entity, most of them aiming at a distance from the object that makes stimulation rhythmically appear and disappear.

Figure 3 depicts a sample behaviour of the best agent evolved ($F(i) = 0.92$). Whilst agent 1 is in stable interaction with the fixed lure, agent 2 is momentarily trapped in an interaction with agent 1's attached lure ($t = [500, 1500]$), eventually abandons it, passes the other agent twice and then finds the fixed lure (Fig 3 (A)). Stimulation received by either of the mobile stimuli is not long and rhythmic enough for interaction to stabilise. Even if interaction with the attached lure is maintained over a number of crossings, the irregularity and intermittence of the sensation, which becomes amplified through gradual modification of return trajectories, eventually allows the agent to move on. This strategy only fails in very exceptional cases in which interaction with a mobile entity is phase-locked in a way that resembles interaction with a fixed lure.

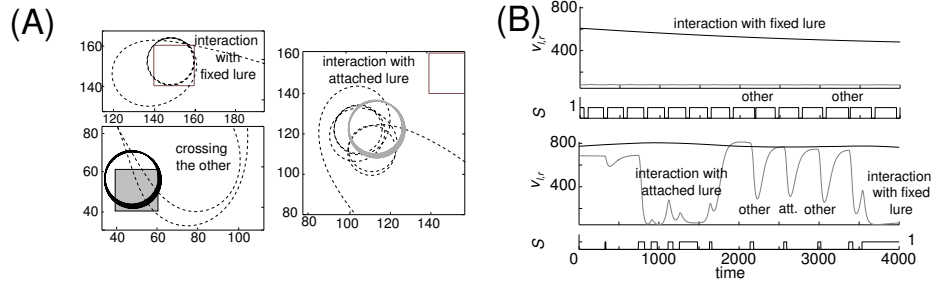


Fig. 3. Best evolved agent. (A) Trajectories (dotted line) of agent 2 during interaction with the fixed lure $t > 3500$, crossing the other $t = [2000, 3000]$ and interaction with attached lure (grey) $t = [750, 1900]$. (B) Corresponding sensorimotor values: $v_{l,r}$ and S (rectangular) for agent 1 (top) and agent 2 (bottom).

The Euclidean Agent. The Euclidean agents evolved mostly scan the space by infinitely going straight around the torus in a direction that produces slightly inclined thin stripes (see best agent with $F(i) = 0.96$ in Fig. 4). This extremely effective search strategy is made possible by the fact that their motor outputs define absolute direction in space. Only two agents evolved to start search in a large curve. About half of the agents evolve to simply stop upon stimulation (hence the slight across population disadvantage compared to the wheeled agents), while the other half engages in rhythmic interaction along one dimension, making stimulation continually appear and disappear. Some of the agents evolved to slowly grind past objects encountered, or to move a bit further away with each oscillation. With such strategies, interaction with the fixed lure is not permanently stable, even if it lasts much longer than interaction with a mobile object (thus avoiding mobile objects). Due to the efficient search behaviour, the chances that to re-encounter the fixed object are still very high. This strategy, again, is very effective and fails only in exceptional cases. Fig. 4 shows how the best agent is hardly perturbed by encountering the fixed lure of the stabilised other because both of them move so fast that stimulation is too short to induce an actual return to the locus of stimulation.

The Arm Agent. For the arm agent, scanning the environment is much more difficult than for either of the other agents. Without proprioceptive feedback it has no way of telling where it is and whether it is still moving or has run up to a joint stop or the edge of the ‘desk’. Most of the agents evolved to sweep across the surface just once (either by running up to the desk edge in a large arch and then grinding down or by running up to the desk edge in a large arch and then grinding down or by running up to a joint stop and then back in a large arch, using slow and fast τ_s). If during their sweep no object is touched, they will indefinitely remain immobile and receive $F = 0$, as there is no environmental feedback to guide or inform further action. This enters an element of chance




Fig. 4. Best evolved agent. (A) Trajectories of agent 1 (solid) and agent 2 (dotted) leading up to stabilised interaction with the fixed lure. Crossing the attached lure (grey) stimulates for a short time only and thus negligibly slows down the agents (kinks in trajectories). (B) Corresponding sensorimotor values $v_{h,v}$ and S (rectangular) for agent 1 (top) and agent 2 (bottom).

into fitness evaluation, which makes arm evolution noisy and probably accounts for the performance disadvantage. The only way to reliably sustain motion is to evolve a neural oscillator as central pattern generator (CPG), as it was found in one of agents (Fig. 5 (A)), which, by virtue of this CPG, is the second best agent, despite just stopping when being stimulated.

There are, in principle many possibilities to mitigate this problem. A modified evolution with proprioception (joint angles), for instance, immediately produced much higher results (population average/best after 1000 generations in 10 runs: 0.33/0.70; see Fig. 5 (3) for example behaviour). However, as stated earlier, it is not the primary objective of the paper to reproduce human behaviour but to explore dynamical principles given the simulation set-up. Even in the sub-optimal solutions evolved, nearly all arm agents evolve to rhythmically interact with any entity encountered (which is not always recognised by the criterion specified in Sect. 3.2), making the sensory stimulation constantly appear and disappear (see Fig. 5 (2) and (3)). This rhythmic activity is realised by varying one of joint angles that control the arm, which leads to the generation of one-dimensional near straight trajectories, similar to those observed in the human participants. Our hypothesis that this kind of behaviour can result from arm morphology is thus supported by the present results.

4 Discussion

The behavioural strategies that the different kinds of evolved agents employ, even though they differ a lot in immediate appearance and quantitative aspects, follow the same dynamical principles. The behaviour evolved for all three agents is dominated by the two local maxima in the fitness landscape (i.e., ‘say yes and indiscriminately interact’ and ‘say no and seek out the fixed lure’). Also,

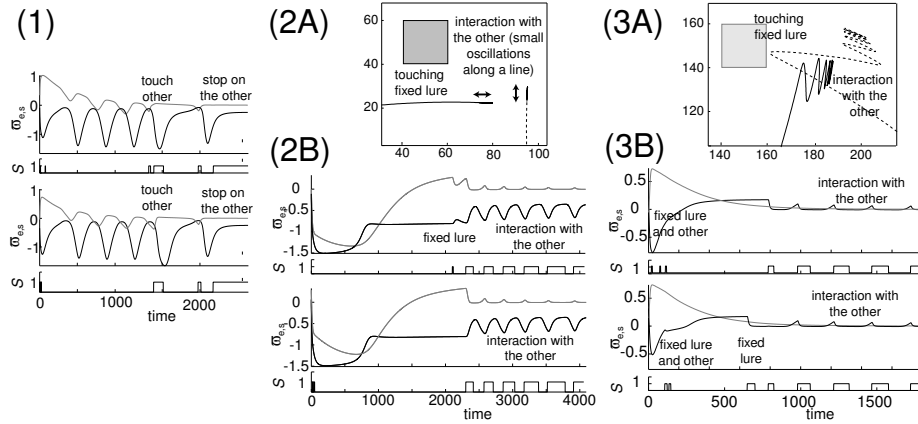


Fig. 5. (1) Agent with neural oscillator as CPG. (2) Best evolved agent. (3) Example agent with proprioception. Trajectories (2A, 3A): Both diagrams depict the rhythmic linear interaction of both agents. In (2A) these oscillations are not visible, because moving just one joint leads to the *exact* inversion of the path just made. Sensorimotor diagrams (1, 2B, 3B): $\omega_{e,s}$ and S (rectangular) for agent 1 (top) and agent 2 (bottom).

for all embodiments, a scanning strategy and an interaction strategy can be distinguished and independently realised. Furthermore, all three types of agents have a tendency to evolve rhythmic interactions that make sensations appear and disappear, which is more likely to evolve with delayed sensory feedback and implies an advantage in fitness, irrespective of preference for the other or for the fixed lure. This is because agents that stay in touch with the boundary of their stimulant are more likely to adjust interaction patterns to a certain degree and maintain interaction against perturbations.

Comparing the different sensorimotor diagrams, there are some remarkable similarities as to how such rhythmic interaction is realised. Irrespective of embodiment, once a stimulation is encountered, one of the motor signals is frozen (statistically significant $p \ll 0.01$ difference between variance in M_1 and M_2 during the last second of a trial across conditions), in order to keep interaction in one place, and only becomes active again if stimulation ceases for too long a period of time. The motor neurons that generate this ‘frozen’ output also tend to have slower time constants (difference not statistically significant). The other motor signal is used to implement local motion and quickly reacts to changes in the rhythmic inputs, actively maintaining interaction. In the wheeled agent, implementing this principle results in small circular trajectories, whereas in the Euclidean and the arm agents, it results in oscillation along a line, just like those observed in human participants. The simulation results thus support our hypothesis that arm morphology plays a role in constituting the one dimensional rhythmic interaction, as the arm-specific implementation of a more general dy-

namical principle of dimensionality reduction during rhythmic interaction. These simulation results generate the hypotheses that the direction in which the participants oscillate should be orthogonal to the orientation of the arm, and that dimensionality reduction should serve rhythmic interaction.

Two phenomena already observed in the model for the one-dimensional version of the experiment have been found to occur again: the role of delays in the evolution of oscillatory scanning behaviour and the distinction between the fixed lure and the other agent on the basis of integrated stimulation time, which propose the investigation of dependencies between a) latencies in sensorimotor action and frequency of oscillations and b) the variation in integrated stimulation time due to anti-phase co-ordination and its role in behavioural preference.

In our earlier simulation model [2], we had evolved agents to seek interaction with one another, presuming a preference for live interaction. Leaving artificial evolution to determine an agent's preference, however, favours 'autistic' agents that avoid any interaction with moving entities, because stable interaction with a static object is easier established and more reliably maintained. This opens an interesting perspective on the previous experiment and simulation, for which we had concluded that perceptual crossing is a nearly inevitable result from the mutual search for each other. In the light of the present simulation results, it becomes clear that the dynamics of the task alone (i.e., leaving aside motivational factors such as boredom) do not favour perceptual crossing, but much rather interaction with the static lure, and that perceptual crossing is established *despite* this strong basin of attraction. Interestingly, a recent modified replication of the one-dimensional experiment with humans (Di Paolo, personal communication) appears to produce results that resemble more the 'autistic' behaviour reported here than the social behaviour observed in the original study in some subjects.

The results presented identify dynamical principles in all evolved solutions and variations in how these principles are realised across different embodiments. Our simulation experiments have generated a number of hypotheses for analysis of the empirical results and thus reaffirm our conviction that evolutionary robotics simulation models are a rich source of illustration and proofs of concept to aid minimalist experimental research on human sensorimotor dynamics.

References

1. Auvray, M., Lenay, C., Stewart, J.: Perceptual interactions in a minimalist virtual environment. *New Ideas in Psychology* (forthcoming)
2. Di Paolo, E., Rohde, M., Iizuka, H.: Sensitivity to social contingency or stability of interaction? modelling the dynamics of perceptual crossing. *New Ideas in Psychology* (forthcoming) Special Issue on Dynamical Systems approaches.
3. Rohde, M., Di Paolo, E.: Adaptation to sensory delays: An evolutionary robotics model of an empirical study. In Almeida e Costa, F., Rocha, L., Costa, E., Harvey, I., Coutinho, A., eds.: *Proceedings of the 9th European Conference on Artificial Life. Lecture Notes in Artificial Intelligence*, Berlin, Heidelberg, Springer (2007) 193–202
4. Beer, R.: Toward the evolution of dynamical neural networks for minimally cognitive behavior. In Maes, P., Mataric, M., Meyer, J., Pollack, J., Wilson, S., eds.: *From Animals to Animats 4*, MIT press (1996) 421–429