

Adaptation to Sensory Delays

An Evolutionary Robotics Model of an Empirical 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. Evolutionary robotics simulations can serve as a tool to clarify counterintuitive or dynamically complex aspects of sensorimotor behaviour. We present a series of simulations that has been conducted in order to aid the interpretation of ambiguous empirical data on human adaptation to delayed tactile feedback. Agents have been evolved to catch objects falling at different velocities to investigate the behavioural impact that lengthening or shortening of sensory delays has on the strategies evolved. A detailed analysis of the evolved model agents leads to a number of hypotheses for the quantification of the existing data, as well as to ideas for possible further empirical experiments. This study confirms the utility of evolutionary robotics simulation in this kind of interdisciplinary endeavour.

Key words: Evolutionary Robotics, Sensory Delays, Sensorimotor Adaptation.

1 Introduction

This paper presents results from an evolutionary robotics simulation implemented to conceptually guide the analysis of empirical data on adaptation to tactile delays. The use of evolutionary robotics simulations in close match with minimalist empirical methods is rather unexplored territory. This study is also an exploration into the utility of this kind of approach¹. The model, which is in itself rather simple, turns out to be an extremely rich source of hypotheses, both for analysis of the existing results and for future experiments.

The data under investigation stems from a recent empirical study (Rohde, Gapenne, Lenay, Aubert and Stewart (unpublished work)). In this study, we set out to investigate the dynamics of human adaptation to delayed sensory feedback in a minimal virtual environment. This experiment was inspired by the work of Cunningham et al. [2], who report an interesting negative aftereffect in subjects trained to solve a simple sensorimotor task with a 200ms visual delay:

¹ For conceptual clarity, in the following, the terms ‘empirical’ and ‘experiment’ are reserved for the real world experiments with humans, while the terms ‘simulated’ and ‘model’ are used to refer to the evolutionary robotics simulation of the task.

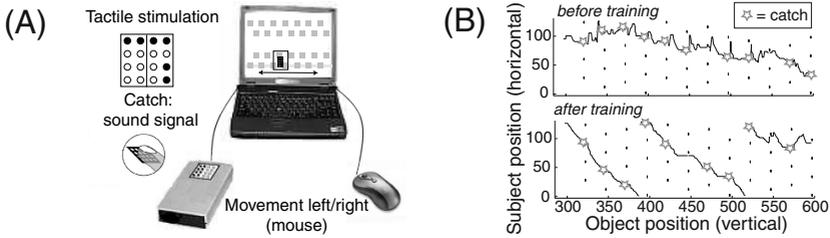


Fig. 1. (A) A diagram of the experimental set up for the empirical study. (B) Human trajectories before (top) and after (bottom) training, presented with a sequence of 16 lines of evenly spaced objects (black bars), both without delay: Even though the behaviour has clearly been transformed, the performance on the task is similar.

Subjects suffered a performance breakdown when returning to the normal no-delay condition, as a consequence of the adaptation to the sensory delay, an effect that had failed to occur in previous experiments with delays, which the authors suspect to be due to an *inherent time pressure* in the task.

Our experiment aimed to investigate the exact behavioural and dynamical conditions that produce such effects. We implemented a catch task in a minimal virtual environment, using the tactile feedback platform ‘TACTOS’ [4] by the GSP group in Compiègne (compare Fig. 1 (A)). Blindfolded subjects had to position themselves as fast as possible below falling objects in order to catch them. To include an inherent time pressure in our task, the objects fall down comparably fast. Subjects moved along one dimension, relying only on tactile stimuli to indicate the crossing of an object and a sound signal to indicate its velocity. We tested the subjects’ capacity to adapt to a $250ms$ sensory delay, monitoring all sensory (tactile display, sound signals) and motor (mouse movement) variables.

In terms of performance profiles, our experiments only produced a clear negative aftereffect in few of the 20 experimental subjects. Indeed, the results seem rather arbitrary at first glance: Some subjects had difficulties with the unperturbed condition already, or were not perturbed by the introduction of the delay, or simply did not suffer a negative aftereffect. A closer look at the trajectories, however, revealed regularities across subjects, as to how different strategies are impacted by the introduction or removal of a delay, and how strategies were modified to recover performance. These changes were not always well represented in performance profiles. Figure 1 (B) shows, e.g., how a subject’s initial strategy (fast oscillatory scanning) was transformed into a different but equally viable strategy (slowly drifting to the right and halting once an object is perceived).

Our next step will be to classify the strategies and to quantify the transitions they underwent across training. It is, however, not easy to formally characterise classes of strategies and the adaptation effects described. Many variables (such as velocity, acceleration, number of crossings, distance, velocity of the objects, movement direction, etc.) and any cross-correlation or dependency between them

in any subgroup of subjects could be used. Apart from the fact that the data is relatively noisy, a major problem we face in our classification is that, despite our minimalist approach, our data is very high dimensional.

The evolutionary robotics simulations presented in this paper have been conducted in order to facilitate this classification. We want to stress that the task as it is solved by the agents, even though it is very simple, is very close to the real experimental set-up. This combination of minimal sensorimotor experiments and evolutionary robotics models proved to be fruitful before [3] and the present study confirms its usefulness. We analyse how delays impact on *reflex-like*, *reactive* and *anticipatory* behaviour, and we identify movement velocity, the number of movement inversions and the exact displacement from the object centre as variables with a large potential to shed light on our empirical data.

2 The Model

The evolved artificial agents, just as the human subjects, can act by moving left or right in an infinite one dimensional space, i.e., a looped tape (See Fig. 2). Evenly spaced objects (size 1×4 , horizontal distance: 28 units) fall down from a vertical distance of 25 at one of seven constant velocities ($v_o \in \{0.004, 0.006, 0.008, 0.010, 0.012, 0.014, 0.016\} \text{units/s}$), each trial consists of a sequence of 32 objects at variable velocities. The agents have a perceptive field of 16×8 units. We decided not to model the exact tactile input patterns the subjects received, but to simply feed a continuous input signal representing the horizontal distance from the centre when an object entered the receptive field ($I_1 = |d_h|/6$ if $|d_h| \leq 6 \wedge d_v \leq 16$). The subjects additionally received auditory pulsed signals to indicate the velocity of falling objects. Such pulses (I_2) are fed into a second input neuron. A third input signal I_3 is a reward signal (auditory in the real experiment), in case an object is caught (rectangular input for 100ms). An object is caught if it is in the centre region of the agent's receptive field when reaching the bottom line ($|d_h| < 4 \wedge d_v = 0$).

All three input signals are fed into the control network scaled by the sensory gain S_G and with a temporal delay. In the 'no delay condition' (*NDC*), there is a minimal processing delay (on average 35ms) in the real experiment, which is prolonged by 250ms to 285ms in the 'delay condition' (*DC*). The same values (i.e., 35 and 285ms) are used in the simulation. The agents are controlled by a continuous time recurrent neural network (CTRNN, see e.g., [1]), whose dynamics are governed by

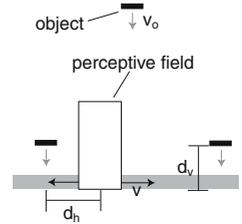


Fig. 2. Illustration of the experimental set-up

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

where $\sigma(x) = 1/(1 + e^{-x})$ the standard sigmoid function, $a_i(t)$ the activation of unit i at time t , θ a bias term, τ the activity decay constant and w_{ij} the strength of a connection from unit j to unit i . The three input neurons feed forward into a fully connected layer of six hidden neurons, which feed the two non-recursively coupled output neurons. The dynamics are simulated using the forward Euler method, with a time step of $7ms$. The basic velocity output v calculated by the network is $v = \text{sign}(\sigma(a_{M1}) - 0.5) \cdot M_G \cdot \sigma(a_{M2})$, so one neuron controls velocity and another one direction, the motor gain M_G scales the output.

The search algorithm used to evolve the parameters of the control network 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 parameters evolved (145 parameters) are: $S_G \in [1, 50]$, $M_G \in [0.001, 0.1]$, $\tau_i \in [25, 2000]$, $\theta_i \in [-3, 3]$ and $w_{i,j} \in [-6, 6]$. 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. The fitness $F(i)$ of an individual i in each trial is given by the proportion of objects caught $F(i) = o_i/O$, O is the number of objects exposed (usually 32). The catch task is in some ways similar to the classification task investigated by Beer in [1], which inspired us to adopt some of his visualisation tools to describe the sensorimotor dynamics.

3 Results

We evolved agents to solve the task under the *DC* and *NDC*, in order to get an idea of the strategy space resulting from our model (10 evolutionary runs with 1000 generations for each condition). Discarding one evolution from both series in which simply nothing evolved, we find that solutions for both conditions evolved to a high level of performance (see Fig. 3 (A)).

The solutions evolved for both scenarios generally involve halting abruptly once the object is encountered, frequently slightly overshooting the target, to then invert velocity and slowly move back to place the object in the centre of the receptive field (see Fig. 4). This strategy is rather trivial, which is probably

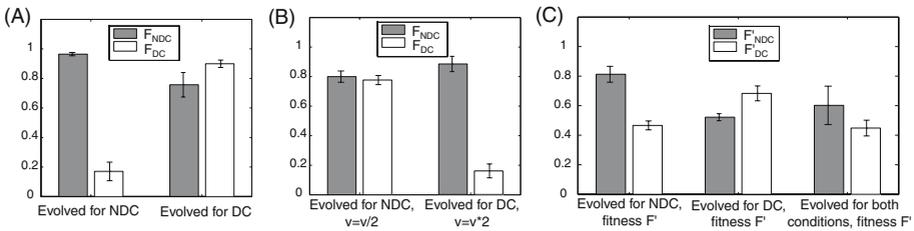


Fig. 3. Performance profile averaged over 9 runs. (A) Unperturbed condition. (B) Scaled velocities. (C) Modified fitness function F' (50% performance chance level).

related to tight temporal constraints and the coarseness of the fitness function that does not well capture the subtleties of sensorimotor perturbation and adaptation (see following analysis). As a stand-alone model, this model would, at best, have been a pilot study for more interesting simulations. However, there are reasons to believe that the discrepancy between performance measure and behaviour are *actual problems in the experiment and that subjects frequently use these kinds of trivial strategies* in response to the inherent limitations of the experimental set-up, just as the agents do. Therefore, a close analysis of the problems with the model and subtleties in the agents' trajectories can be useful for the data analysis and an improvement of the experimental set-up.

If we look at the solutions in terms of agent performance only, a striking difference is that most of the solutions to the *DC* are robust to the removal of the delay, while most of the solutions evolved for the *NDC* suffer a drastic breakdown in performance below chance level once the delay is introduced (see Fig. 3 (A)). If succeeding in the *DC* in many cases subsumes succeeding in the *NDC* for this task, and this is not just a property of the model, this issue could be crucial in explaining the failure to produce a negative aftereffect in the subjects, despite the transformation of sensorimotor strategies.

Another difference is that the velocity at which the object is first touched is on average twice as high for the *NDC* ($\bar{v} = 0.025$) than it is for the *DC* ($\bar{v} = 0.014$). Do maybe agents use the same strategy for both conditions, only varying their velocity? And are furthermore solutions generally able to deal with shorter, but not with longer delays? A very crude test for this hypothesis is to scale the velocity generated by evolved controllers, i.e. to double v in the networks evolved for the *DC* and to divide it by two in the networks evolved for the *NDC*. This test seems to confirm this hypothesis, because by virtue of this scaling, the performance profiles of the solutions evolved for both conditions are inverted (compare Fig. 3 (A) and (B)): Networks evolved under the *NDC* become generalists that are good at both conditions, whereas sped up networks evolved under the *DC* lose their capacity to perform the task with delay and become specialists for the *NDC*. A closer look at the sensorimotor dynamics, however, shows that things are not quite this simple.

As a first step in the analysis, it can be established that all evolved controllers seem to function independently of the reward signal and the pace at which the objects fall (I_2 and I_3), agents simply try to put objects as quickly as possible into the centre of the perceptive field. Therefore, agents produce the same trajectories for different object velocities that are just cut off at different points in time, and we will disregard these inputs in our following analysis.

Systematic displacements. Figure 4 depicts trajectories from different starting positions for two example individuals, one evolved for the *DC* (A) and one evolved for the *NDC* (B), and how they differ when tested under the *NDC* (top) and the *DC* (bottom). Both agents achieve to locate the object in the centre of their receptive field for most possible starting positions in the respective condition they have been evolved for (bottom left for *DC*, top right for *NDC*). Taking a look at how the behaviour is altered by the introduction/removal of a

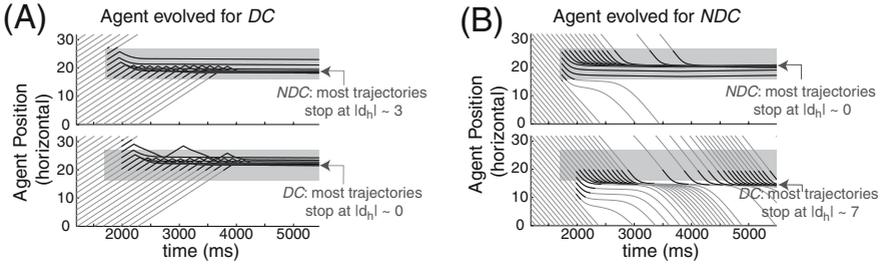


Fig. 4. Trajectories for different agent starting positions across time, presentation of a single object. Crossing the object (grey region) produces a (delayed) input stimulus I_1 (trajectories black during stimulation). Top: *NDC*, bottom: *DC*. (A) an agent evolved for the *DC* (B) an agent evolved for the *NDC*.

delay, it can be seen that in both cases the trajectories are systematically displaced from the centre of the perceptive field: The *NDC* agent overshoots its goal when exposed to a prolonged delay (bottom right), while the *DC* agent stops too early if the delay is removed (top left). The magnitude of these displacements is proportional to the agents' velocities. As the *DC* agent is much slower, its systematic displacement is small enough for trajectories to stay in the viable centre zone ($|d_h| < 4$) of the receptive field, as defined by the fitness function $F(i)$, while for the *NDC* agent, the displacement takes trajectories outside its receptive field. Such systematic displacement of trajectories can be observed for most agents. Since *DC* agents are on average half as fast, trajectory displacement is on average smaller for them. This seems to explain their robustness towards removal of the delay, which would then be due only to the coarseness of the fitness function that does not detect or punish micro displacements.

In order to further investigate these matters we evolved a new set of agents with a spatially more exact fitness function $F'(i) = \frac{1}{\mathcal{O}} \sum_{\mathcal{O}} 1 - \sqrt{d_h}/4$. With this modification, solutions to the *DC* cease to be robust to the removal of the delay (see Fig. 3 (C)), which confirms the hypothesis that *DC* agent robustness is due to the immunity of $F(i)$ to micro displacements. Interestingly, the experimental subjects were on average slower after training with delays. Therefore, these insights about systematic displacements, their relation to velocity and their effect on catch events (performance) may help in explaining the failure to produce a negative aftereffect in the experiment.

The modified fitness function also allows us to explore *agent adaptation to delays* by evolving agents under both conditions (Fig. 3 (C)). In the original task, this was not interesting, as many *DC* agents are able to solve both conditions anyway, so there was no selection pressure to go beyond robustness. Unfortunately, despite longer evolution, no adaptive adjustment of strategies evolved, only fixed strategies that compromised between the two conditions. We decided not to explore these possibilities further, as these are purely theoretical exercises.

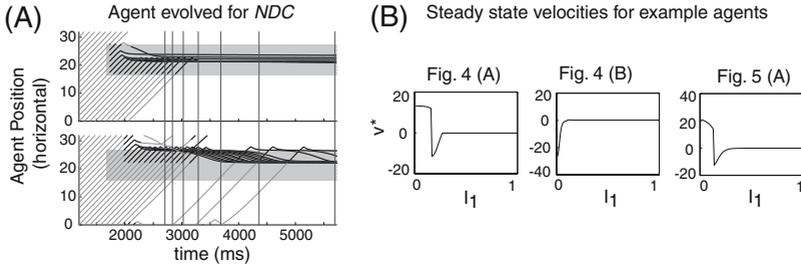


Fig. 5. (A) Trajectories for different agent starting positions across time for an agent evolved for the *NDC*, presentation of a single object. Crossing the object (grey region) produces a (delayed) input stimulus I_1 (trajectories black during stimulation). Top: *NDC*, bottom: *DC*. Vertical lines: cut off for different v_o . (B) Steady state velocities v^* for different I_1 for the agents in Fig. 4 (A) and (B) and Fig. 5 (A).

As future empirical experiments, such set-ups are neither possible, nor desirable, as discussed in the conclusion.

Reactive and reflex-like behaviour. Looking at the steady state velocities for varying I_1 in evolved agents, (Fig. 5 (B)), we generally find that $v^* = 0$ for values of I_1 that exceed a certain rather low threshold value of I_1 . Behaviourally, this means that the agents are only sensitive to the onset of the stimulation when an object enters the receptive field, which triggers a rapid decay of v to 0, sometimes preceded by a movement direction inversion (negative peaks in v^* in Fig 5 (B) left and right, compare corresponding trajectory graphs). The variation in signal magnitude, as an agent moves to the exact position to stop, however, is without effect on agent behaviour. This is why, e.g., the agent depicted in Fig. 4 (A) in the *NDC* remains in its location displaced from the centre of the receptive field, rather than to actively search for the exact centre. Such strategies are *reflex-like* in that they produce stereotyped trajectories.

Reflex-like behaviour evolved in all runs except one. The agent in Fig. 5 (A) is the only agent evolved for the *NDC* that maintains a relatively high level of performance when exposed to the *DC*. This is because the overshooting of the target and the first movement inversion bring it back to the margin of the perceptive field, where sensory activation is sufficiently low to trigger another movement inversion, which brings the trajectory to the centre. In that sense, the behaviour is more *reactive*, because it is sensitive to changes in magnitude of the signal caused by ongoing behavioural dynamics (Fig. 5 (A) top vs. bottom). Why did such reactive focussing strategies not evolve more systematically? The answer is simple: Because the deliberate inherent time pressure in the task does not allow for them to be beneficial. The cut off time for trials with the top three velocities is 1000, 1142 and 1333ms after the objects become perceptible, which corresponds to the cut off points (vertical lines) $t = 2701, 2843$ and 3033 in Fig. 5 (A). The reactive mechanism that brings back the overshooting trajectories

needs more time to come into effect. Therefore it is important that the agents induce the right behaviour immediately when the object is perceived, the time window is just big enough to execute a reflex, not for reactive correction.

Why do faster solutions evolve for the *DC* than for the *NDC*? The intuitive answer to this question is that slowing down seems to be the obvious way of coping with a delay. This is, however, only directly true for reactive strategies, in which ongoing behaviour correction is informed by and has to wait for the delayed signal representing the effect of one's own previous actions. For the execution of a reflex, however, there is no real disadvantage to a delay other than a shortening of the absolute time window in which to execute it, which is what makes it more mysterious that *DC* agents are slower. We investigated three plausible hypotheses: 1.) A maximum period of drifting back related to the upper limit of 2000ms for τ . 2.) An optimisation of velocity to a minimal network reaction time $t(d) = t_r + d$. 3.) An advantage in absolute time for *NDC* agents due to the 250ms longer time window. We didn't find evidence for any of these possible explanations. We are still unsure why faster solutions evolve for the *NDC*, there seems to be no simple answer. The answer may well be a combination of several of these simple factors, but, at the moment, we can give no obvious pressing reasons for why *DC* agents should not be fast.

4 Conclusion

We present an evolutionary robotics study that is a very close model of a minimal empirical study on adaptation to sensory delays in human subjects. The simulation model in itself is simple and not behaviourally sophisticated. However, we hypothesise that these problems of the simulation are analogous to problems in the empirical experiment, whose results we do not yet fully understand. Since evolutionary robotics simulations are scientific tools that can generate new hypotheses and conceptual clarifications [5], the model merits closer inspection.

A substantial number of hypotheses about the existing data set can be derived from the model. In the evolved solutions, sensorimotor behaviour is independent of the reward signal and object velocity (apart from different cut off points), and the behaviour within one object presentation is largely independent of previous object presentations. If this is similar in the empirical data, dimensionality could be reduced immensely. Another property of the evolved solutions is that, unless halting leads to a cessation of the input signal, the exact position of the object in the receptive field is not adjusted, which leads to another hypothesis to be explored for data simplification. On the other hand, insights were gained about the systematic displacement of trajectories from the object centre and its relation to velocity. This gives us a starting point to search for evidence in our data that a negative aftereffect could indeed have been produced, even if this is not reflected in the performance profile. As a general case, the visualisation of steady states and neighbouring trajectories that we adopted from Beer's work [1] has led to a helpful pre-theoretical understanding of evolved strategies. We will explore whether such techniques are, to an extent, transferable to the empirical data.

These insights about the discrepancy between sensorimotor adaptation and task performance should in principle enable us to design a better experiment in which these two factors concur. By using a fitness function that is spatially more exact, we can easily modify the simulation for this to be the case. Unfortunately, this modification is not feasible for the real experiment. Both, the temporal sampling rate and the spatial resolution are coarser in the experimental virtual environment than in the simulation, a coarseness that is necessary for the experiment to work in real time. There are fierce limitations to the transferability of such simulations to real experimental designs, an issue that can be seen as parallel to issues of simulation vs. real robots in classical robotics research.

But, even if these technical limitations could be mitigated, there are more profound lessons to be learned from the analysis of the simulation, lessons that relate to our discussion of *reactive* and *reflex-like* behaviour and that propose a more radically different experiment. In purely reactive sensorimotor loops, a sensory delay has an effect similar to an increase in inertia: Ongoing behaviour has to slow down, because what to do next depends on previous movements and the (delayed) perception of their effects. To an extent, this had already been recognised by Cunningham et al. [2], who hypothesise that negative aftereffects failed to occur in previous studies because they always provided the possibility to compensate for the delay by slowing down – hence the emphasis on the inherent time pressure in their task, which makes compensation impossible and forces real adaptation. This convincing argument had led us to include a strong time pressure in our simulation (and in our experiment) by choosing high object velocities. As desired, this hindered the evolution of reactive solutions, in which subjects can slow down to compensate. However, evolution found a different, faster ‘cognitive shortcut’: Agents use constant velocities, such that a fixed sensory delay corresponds to a fixed spatial displacement, which forms part of the stereotyped reflex-like trajectories that take objects to the centre of an agent’s receptive field. In this kind of sensorimotor loop, we find a systematic error for both shortening and lengthening of delays, which makes it in principle possible to yield a negative aftereffect.

However, a more profound and complex adaptation to the delay seems to have taken place in Cunningham et al.’s experiment, something that goes beyond the tuning of reflex-like trajectories in our agents (and possibly in our subjects). This teaches us that there is more than just time pressure to their experimental design. Their visual task forces subjects to *produce fast variable velocity motor sequences*, in which a sensory delay can be neither understood as an increase in inertia, nor as a fixed displacement. However, in order to be able to usefully employ a fast motor sequence, it has to be possible to anticipate consequences of actions far into the future, i.e., the signal needs to be complex enough to feature cohesion between momentary signal structure, own movement possibilities, and future signal structure over a longer time course, regularities that make *anticipatory* behaviour possible. When designing a task analogous to Cunningham et al.’s within our minimalist approach, the complexity of the visual sensation has been traded for a more tractable but less rich tactile sensory

signal. Thereby, long term predictability was lost, a loss that we were aware of, but that we did not previously consider essential. The lack of predictability, however, lead to the evolution of reflex-like behaviour. Reactive behaviour could not evolve because of the inherent time pressure in the task, and anticipatory behaviour could not evolve because of the poverty of the input signal. We now hypothesise that only anticipatory behaviour will produce interesting negative aftereffects to adaptation to delay and plan to design a series of new minimalist experiments with high predictability, to explore this hypotheses and elaborate our ideas about reactive, reflex-like and anticipatory sensorimotor strategies.

This study is also an exploration into the usefulness of evolutionary robotics simulations that closely match minimalist empirical experiments. The amount of hypotheses derived from the model on different levels of conceptual analysis speaks for itself; we consider this model a full success. However, it also became clear that some of the experimental mistakes would probably not have been made if the simulation had been run *before* the experiment. The design of an empirical experiment goes through a cycle of generate-and-test, which costs a lot of time and effort, both for the experimenter and for the piloting subjects, and frequently is ended pre-maturely. To transfer this generate-and-test phase as far as possible to simulation would be much less resource expensive. Therefore, we recommend the use of evolutionary robotics models not just after, but also before and alongside the conduction of empirical experiments.

In our analysis, we have focussed on one particular kind of agent behaviour, which we saw as model for a salient and characteristic, but nonetheless particular strategy adopted by the subjects. Both the empirical and the evolved solutions are much more versatile than that. Whether or not the generated hypotheses help in the classification and description of the data remains to be shown. In case these results lead the analysis into a dead end, this is not a reason to put the head in the sand. The possibilities of learning from evolutionary robotics models about the investigated scenario are not nearly exhausted yet.

References

1. Beer, R.: The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior* 11(4), 209–243 (2003)
2. Cunningham, D., Billock, V., Tsou, B.: Sensorimotor adaptation to violations of temporal contiguity. *Psychological Science* 12(6), 532–535 (2001)
3. Di Paolo, E., Rohde, M., Iizuka, H. (Forthcoming): Sensitivity to social contingency or stability of interaction? Modelling the dynamics of perceptual crossing. *New Ideas in Psychology*, Special issue on Dynamics and Psychology
4. Gapenne, O., Rovira, K., Ali Ammar, A., Lenay, C.: Tactos: Special computer interface for the reading and writing of 2D forms in blind people. In: Stephanidis, C. (ed.) *Universal Access in HCI: Inclusive Design in the Information Society*, pp. 1270–1274. Lawrence Erlbaum Associates, London (2003)
5. Harvey, I., Di Paolo, E., Wood, R., Quinn, M., Tuci, E.A.: Evolutionary Robotics: A new scientific tool for studying cognition. *Artificial Life* 11(1-2), 79–98 (2005)