

Toward Minimally Social Behavior: Social Psychology Meets Evolutionary Robotics

Tom Froese and Ezequiel A. Di Paolo

CCNR, University of Sussex, Brighton, UK
t.froese@gmail.com, ezequiel@sussex.ac.uk

Abstract. We report on a set of minimalist modeling experiments that extends previous work on the dynamics of social interaction. We used an evolutionary robotics approach to fine-tune the design of a recent psychological experiment, as well as to synthesize a solution that gives clues about how humans might perform under these novel conditions. In this manner we were able to generate a number of hypotheses that are open to verification by future experiments in social psychology. In particular, the results indicate some of the advantages and disadvantages of relying on social factors for solving behavioral tasks.

Keywords: Evolutionary robotics, social psychology, social interaction.

1 Introduction

Since its beginnings in the early 1990s evolutionary robotics (ER) has established itself as a viable methodology for synthesizing models of ‘minimally cognitive behavior’, namely the simplest behavior that raises issues of genuine cognitive interest [2]. Within this context there has been a growing interest in using this method to investigate the minimal dynamics of social interaction (cf. [5] for a review). As a specialization of the ER methodology, we can conceptualize this kind of modeling as investigations into the dynamics of ‘minimally *social* behavior’.

What is interesting about some of these recent advances in ER is that the synthetic method has been used to create models which are explicitly inspired by actual psychological experiments. Moreover, some of these models have been specifically designed to generate insights with the potential to generate mutually informing collaborations between the field of artificial life and the traditional empirical sciences, especially social psychology (e.g. [3, 4]). One promising target for this endeavor is Auvray, et al.’s [1] minimalist perceptual crossing experiment. This psychological study attempts to explore the most basic conditions necessary for participants to recognize each other through minimal technologically mediated interaction in a shared virtual space. Since this study will be the target of the modeling experiments presented in this paper we will describe the study in a bit more detail here. A schematic of the overall experimental setup is shown in Figure 1.

Two adult participants, acting under the same conditions, can move a cursor left and right along a shared 1-D virtual tape that wraps around. They are asked to indicate the presence of the other partner. The participants are blindfolded and all they

can sense are on/off tactile stimulations on a finger when their cursor crosses an object on the tape. Apart from each other, participants can encounter a static object on the tape, or a displaced ‘shadow image’ of the partner, which is strictly identical to the partner as regards to size and movement characteristics. There are thus three distinct types of objects (each is 4 pixels wide) which can be encountered by a participant, one of which is placed at a fixed location and two of which are moving within the 1-D space. The two mobile objects exhibit exactly the same movement, but only an overlap of the receptor fields of both participants gives rise to mutual sensory stimulation. Note that the difference between these three types of objects cannot be directly provided by the sensors, which in all cases can only produce a binary response depending on whether something is overlapping the receptor field or not.

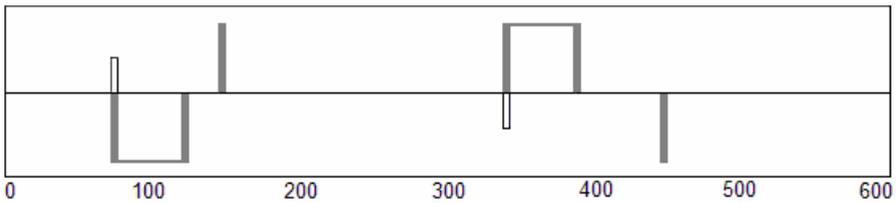


Fig. 1. The experimental setup: Two participants capable of horizontal movement face each other in a 600 unit 1-D ‘tape’ that wraps around at the edges. Note that a participant’s receptor field (white bar) can encounter three different objects (gray bars): a static object (located at 148 or 448, depending on the participant), the other participant’s avatar (coinciding with the location of its receptor field), and the other’s ‘shadow’ (attached to the other’s avatar via a rigid 48 unit link). Since all objects have the same width (4 units) they cannot be told apart by any difference between their appearances (they all give rise to an all-or-nothing tactile response).

The results of the psychological study show that, at least under the minimalist conditions of this experiment, the successful recognition of an ongoing interaction with another person is not only based on individual capacities. It is also based on certain properties that are intrinsic to the joint perceptual activity itself. The important issue is that the scanning of an object encountered will only stabilize in the case that both partners are in contact with each other —if interaction is only one-way, between a participant and the other’s shadow, the shadow will eventually move away, because the participant it is shadowing is still engaged in searching activity. Two-way mutual scanning is the only globally stable condition. Therefore, the solution to the task does not only rely on individuals performing the right kind of perceptual discrimination between different sensory patterns (the empirical data shows they cannot distinguish between the other and its shadow), but also emerges from the mutual perceptual activity of the experimental subjects that is oriented towards each other.

2 Previous Work

Di Paolo, et al. [3] used ER to generate a simulation model of the perceptual crossing experiment, and successfully replicated the results while at the same time gaining some additional insights into the dynamics of the interaction process. For example,

the problems that their agents had with avoiding interactions with their respective static objects led them to predict similar difficulties for human participants. It turns out that repeated crossing of an object produces a pattern of stimulations located on the same spatial position. This is the same in the case of a fixed object and the other agent if it moves in coordinated anti-phase. Accordingly, it is difficult for the two patterns to be distinguished (see Fig. 3 for an example). This prediction was supported by the empirical data presented by Auvray and colleagues [1], but previously went unnoticed. Froese and Di Paolo [5] replicated these modeling results and introduced some variations, which resulted in further hypotheses about the stability of the interaction process in organizing the behavior of the interactors.

In both modeling studies [3, 5] it was demonstrated that the simulated agents make use of the *duration of contact* with objects in order to discriminate interactions with a static object (always same length of stimulation for same velocity) and the other agent (potentially shorter or longer stimulation, depending on whether the other passes by in in-phase or anti-phase movement). This is a reliable basis of distinction because the other agent is always moving. It is unlikely, however, that this is the main strategy employed by humans in the original psychological study. To be sure, during the training phase the participants were asked to interact with a four-pixel wide object in three conditions. The target object was either (i) static, (ii) moving at a constant speed of 15 units/second, or (iii) moving at a constant speed of 30 units/second, and each of these one min. training phases was announced as such. It could thus be possible that the participants learned the correlation between contact duration and whether an object is static or moving. In practice, however, the difference in duration is small enough such that it is unlikely to be the main strategy of the participants, though there is some evidence that contact duration made a difference, leading to 31.3% of clicking response (cf. event E6 [1, p. 40]). Still, we hypothesize that the successful behavior of the participants is based on different types of interactions afforded by the static object and the other active participant, rather than their differing durations of contact.

The question we want to address in this paper is: can we use ER to investigate the kinds of strategies that are available when such a duration-based strategy is excluded from the experimental design? One way to approach this is to make all objects (i.e. agents, shadow objects, and static objects) within the virtual environment infinitely small. This can be done by simply checking whether the sign of the difference of the locations (of the agent and some target object) has changed compared to the previous time step. If the sign has changed, then we activate the agent's receptor field. Since in this case all objects afford an equal duration of contact (i.e. 1 time step), it is no longer possible for the agents to trivially rely on the fact that other moving objects entail a shorter contact. Can we use ER to generate a strategy that enables the agents to successfully locate each other even under this more ambiguous situation?

3 Further Experiments in Perceptual Crossing

The simulation model includes two agents facing each other in a 1-D environment (i.e. one agent faces 'up' and one agent faces 'down'), which wraps around on itself after 600 units of space. In the simulation all distance and time units are of an arbitrary scale. There is no noise. Each model agent controls the horizontal movement of

its ‘body’, i.e. the position of its receptor field. The position of the agents is represented by continuous variables. The velocity of each agent is determined by taking the difference in output of two nodes of a continuous-time recurrent neural network (CTRNN), as described by Beer [2]. The CTRNN is fully inter-connected with self-connections. No symmetry is imposed on the network. The sensory input of an agent is activated (set to 1) when its receptor field passes another object (i.e. it crosses it between two time steps), otherwise the input remains off (set to 0). Each node receives sensory input which is multiplied by a specific input gain.

The GA evolved 100 solutions spread throughout 10 niches for several thousand generations. Each solution was evaluated for 100 trials lasting for 800 units of time, each with a time step of 0.1; the overall score was weighted toward the worse trials. Each solution coded for a clonal pair of CTRNNs (range of biases and weights $[-8, 8]$, time constants $[1, 200]$). For more details of the GA, see [5]. In contrast to the original experimental setup, each object in the environment, no matter whether it is moving or stationary, only activates the receptor field of a passing agent for 1 time step. To make the solutions more evolvable it was necessary to include a large range of input gains (range $[-1000, 1000]$) so as to compensate for the minimal period of stimulation; including a sensory delay of 5 units of time was also helpful. As in previous modeling experiments, the solutions were evaluated in terms of how close the two agents were to each other on average during a trial.

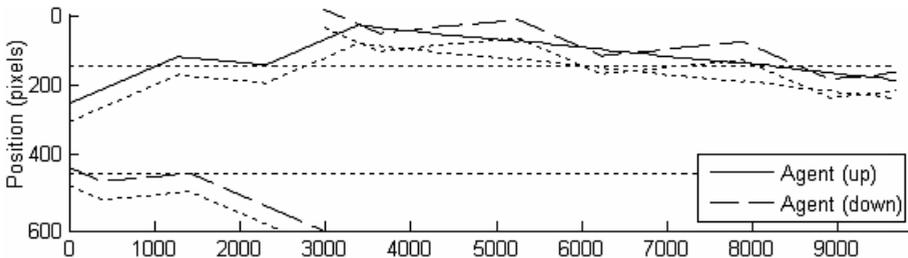


Fig. 2. A trial run showing the movement of the agents during the first 1000 time steps. The agents manage to avoid their respective static objects (after only 1 stimulation), eventually locate each other, and then continue to engage in perceptual crossing until the end of the trial.

Experiment 1: Making objects infinitely small.

We successfully evolved a 6-node CTRNNs to cope with this modified setup. To test the robustness of this solution we ran a comprehensive set of test trials; the average score is not significantly different from that of the original modeling setup. It turns out that the strategy of the agents is based on the close proximity of the two shadow objects. All the agents have to do is distinguish between one stimulation and two consecutive stimulations. This is a robust individual-based strategy to locate the other since: (i) passing the static object only causes *one* activation of the receptor, and (ii) passing the other agent with its attached shadow results in *two* activations. In other words, the evolution has found an individual-based solution that relies on an external factor, namely the relationship between the agents and their shadows (cf. Fig. 2).

Ironically, this behavioral strategy is robust because the shadow, which was meant to introduce an essential ambiguity into the experiment, has been appropriated to disambiguate the target from the static object. Is this a strategy that would be used by the human participants of the original study? Participants were indeed told about the experimental setup, including the three types of objects that they could encounter, but “the precise relation of the mobile lure yoked to the avatar was not explained” ([1], p. 38). Nevertheless, a large percentage of responses was preceded by a double stimulation (event E2, 32.3%, [1], p. 40), indicating that the shadow might have played a role in the positive empirical results. But what kind of strategies would be available if participants cannot take advantage of the agent-shadow relationship?

Experiment 2: Making shadows maximally distant.

We do not want to completely sever the link between the movements of the agents and their shadow objects, since this is an essential aspect of the experiment. Instead, we simply make the link between them maximally distant (150 units)¹. The rest of the setup remains the same as in the previous experiment. We evolved 6-node CTRNNs for several thousand generations and then chose the fittest solutions to run some test trials. The outcome of a typical trial is shown in Figure 3.

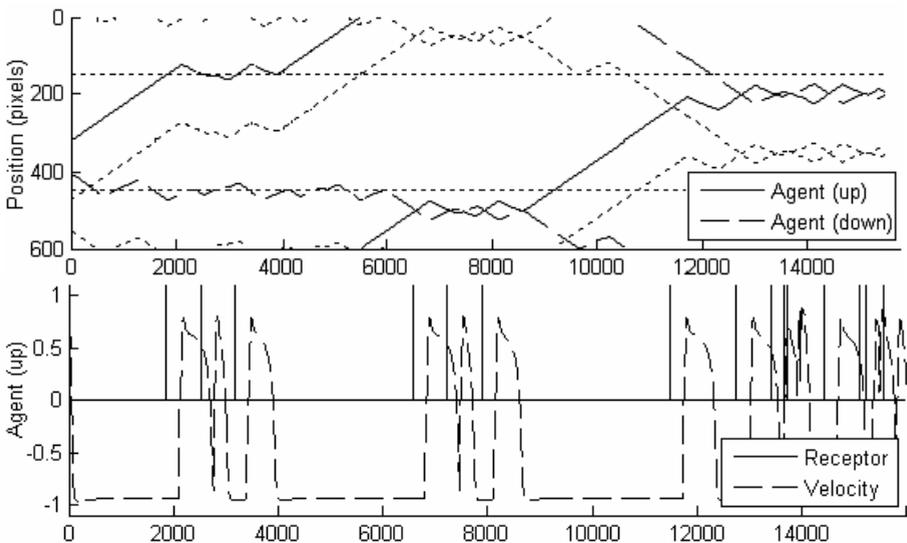


Fig. 3. Initial movement of the agents during a trial with maximally distant shadows. Note that it in this case it was impossible for agent ‘up’ to distinguish between the sensory-motor pattern produced by interacting with the static object (between ca. 2000-3000 time steps), and by interacting with the other agent (between ca. 7000-8000 time steps). Both encounters elicit the same three stimulations and the same motor response (bottom graph).

¹ Strictly speaking, in a 600 unit-wide circular world, being apart 300 units would be maximally distant. However, in this case there is another stable perceptual crossing situation, in which the agents can interact at a distance by stimulating each other with their shadow objects.

First, the agents explore their static objects for some time, then proceed to explore the rest of the space, then encounter each other and engage in some initial perceptual crossing. This mutual interaction breaks down after a while, and they continue exploring until they re-establish perceptual crossing at another location. Such break-downs occur more often than in the original setup, because agents are more likely to miss each other with infinitely small object sizes. Indeed, the possibility of coordination break-down could be a first indication that the agents have to be much more active and responsive in their interaction in order to disambiguate the situation. They cannot make use of persistent and reliable external factors to assess the viability of their behavior, and thus they are more open to commit errors and mistaken responses. The behavior of the agents during the trial run thus looks much more lifelike than that of previous solutions. This modeling experiment leads us to the prediction that the performance of human participants under these modified conditions would not be significantly different than from the original setup.

However, there still remains a problem in terms of this model. When the agents meet without receiving different stimulation beforehand, they engage on the basis of identical controllers (same structure and same internal state) such that they will mirror their behavior perfectly. This produces the same sensory-motor correlation as if they were oscillating around their static object. And since agents are more likely to encounter each other, evolution produces solutions which treat the occurrence of this sensory-motor pattern as always being due to the other rather than to the static object (a good choice, given the circumstances). However, occasionally this will result in both agents getting stuck on their static objects for the whole of a trial, giving rise to what looks like truly pathological behavior. It appears that such interactions do not break down when agents become entrained in too close proximity, thereby always making another contact with the object on their return path.

Experiment 3: Coordinated behavior.

How can we use ER to generate solutions that are better at distinguishing the other agent from the static object? As a first step, we remove the possibility of functionally identical CTRNNs encountering each other by simply activating the receptor field of a randomly chosen agent at the start of the trial, thus ensuring a minimal difference in individual histories. Moreover, in [4] we showed that sensitivity to social contingency can emerge from the interaction process if agents are required to coordinate their behaviors in a way that forces them to break the symmetry of their interactions. We therefore introduce an additional requirement into the fitness function by rewarding solutions in which agents travel together while continuing to engage in perceptual crossing. Since the agents are clones, this coordinated activity will require them to break the symmetry of their interactions in order to succeed. Moreover, since it is impossible to coordinate with the static object, we have emphasized the possibility of distinguishing the other in terms of its responsiveness.

We evolved agents with this modified fitness function that are able to coordinate their behavior so as to travel together while interacting (cf. Fig. 4). While engaging in perceptual crossing, the agents eventually start to drift together horizontally. In other words, even though the agents are structurally identical, have minimally different histories (internal states), and are not affected by noise during the trial, they are

nevertheless able to regulate the interaction such that the symmetry of their individual behaviors is broken. In fact, when one of the agents encounters its static object during this coordination process, the agents are able to re-negotiate the direction of drift and return the other way, much like what was found in the pioneering work by Quinn, et al. (2003). Future work could analyze in more detail the precise manner in which this symmetry breaking is realized in the dynamics of the interaction process.

Nevertheless, it is still the case that this strategy is not as robust as the solutions which we have excluded from the experimental design. If both agents happen to encounter their static objects at the start of the trial, they are likely to continue oscillating around it until the trial is terminated, even if the possibility of functional identity has been removed in principle. At first sight this appears to be evidence that the agents are not sensitive to the social contingency of their interaction; otherwise they would presumably notice the lack of responsiveness of the static object and eventually move away. But this way of looking at the problem essentially demands an *individual* response alone, i.e. detecting lack of social contingency when there is none to be detected. In contrast, perhaps the existence of this pathological behavior is an indication of the truly *social* nature of the evolved solution? Indeed, if only one agent becomes trapped it will eventually be freed by the other agent, which entrains it in an interaction process such that they move away together.

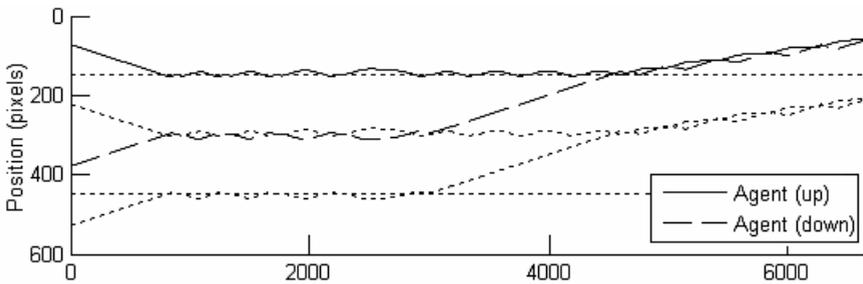


Fig. 4. Initial movements during a trial run showing coordinated behavior. First, both agents get stuck: agent ‘up’ on its static object and agent ‘down’ on the other’s shadow. Then agent ‘down’ manages to break free and continues searching until the agents meet. This interaction frees agent ‘up’ from its entrainment and they start moving together until the end of the trial.

4 Discussion

We used ER as a process of fine-tuning the simulated setup to avoid the evolution of trivial solutions. At the same time it enabled the formulation of novel empirical predictions, in particular that those elements of the original experimental setup which were used as the basis for trivial solutions are not essential to the general results of that original study. Thus, we can hypothesize that the overall outcome of the original study will not be significantly altered when making the objects infinitely small and displacing the shadow object by 150 units. Indeed, we can venture a further hypothesis that part of the reason why the original study found less response to the static object, when compared to the two mobile objects, was that entrainment with this object was often broken by the actions of the other participant.

Another thing that we can learn from the design process is just how difficult it is to evolve a behavioral strategy that is only based social interaction. One important factor is that basing a behavioral strategy on the responsiveness of the other introduces an inherent risk factor in to the situation. What happens when the presumed 'other' does not react to your behavior in a suitable manner but your individual discriminatory ability depends on a certain kind of interaction? This appears to be the case in the 'pathological' behavior of the agents. More generally, the other's behavior can be influenced by your own actions, but it evades your direct control in principle. Another factor is that detecting another's responsiveness *as such* during an interaction is a much more demanding task than detecting simple environmental cues (e.g. difference in stimulus duration, difference in number of contacts, difference in noise, etc.).

Finally, the modeling experiments presented here point to the possibility of future dialogue between ER and cognitive science. The fact that the agents in the final experiment are unable to distinguish between the static object and the other agent individually, but can do so when that other agent is present, deserves further study, especially in relation to empirical findings. For example, studies of rehabilitation after brain damage have shown that patients often (i) find it impossible to individually achieve sensory-motor tasks in an abstract context, (ii) have difficulty with them in a pragmatic context, and (iii) can function normally in socially situated circumstances (cf. Gallagher & Marcel 1999). A modeling hypothesis we can draw from these empirical findings is that discrimination of the static object will be possible for individual agents with more complex CTRNN controllers. Indeed, some preliminary experiments with 10-node CTRNNs have indicated this to be the case.

References

1. Auvray, M., Lenay, C., Stewart, J.: Perceptual interactions in a minimalist virtual environment. *New Ideas in Psychology* 27(1), 32–47 (2009)
2. Beer, R.D.: Toward the evolution of dynamical neural networks for minimally cognitive behavior. In: Maes, P., et al. (eds.) *Proc. of the 4th Int. Conf. on Simulation of Adaptive Behavior*, pp. 421–429. The MIT Press, Cambridge (1996)
3. Di Paolo, E.A., Rohde, M., Iizuka, H.: Sensitivity to social contingency or stability of interaction? Modelling the dynamics of perceptual crossing. *New Ideas in Psychology* 26(2), 278–294 (2008)
4. Froese, T., Di Paolo, E.A.: Stability of coordination requires mutuality of interaction in a model of embodied agents. In: Asada, M., Hallam, J.C.T., Meyer, J.-A., Tani, J. (eds.) *SAB 2008. LNCS (LNAI)*, vol. 5040, pp. 52–61. Springer, Heidelberg (2008)
5. Froese, T., Di Paolo, E.A.: Modeling social interaction as perceptual crossing: An investigation into the dynamics of the interaction process. *Connection Science* (in press)
6. Gallagher, S., Marcel, A.J.: The self in contextualized action. *Journal of Consciousness Studies* 6(6), 4–30 (1999)
7. Quinn, M., Smith, L., Mayley, G., Husbands, P.: Evolving controllers for a homogeneous system of physical robots: structured cooperation with minimal sensors. *Phil. Trans. R. Soc. Lond. A* 361, 2321–2343 (2003)