

Stability of Coordination Requires Mutuality of Interaction in a Model of Embodied Agents

Tom Froese and Ezequiel A. Di Paolo

Centre for Computational Neuroscience and Robotics (CCNR),
Dept. of Informatics, University of Sussex, UK
{t.froese, ezequiel}@sussex.ac.uk

Abstract. We used an evolutionary robotics methodology to generate pairs of simulated agents capable of reliably establishing and maintaining a coordination pattern under noisy conditions. Unlike previous related work, agents were only evolved for this ability and not for their capacity to discriminate social contingency (i.e., a live responsive partner) from non-contingent engagements (i.e., a recording). However, when they were made to interact with a recording of their partner made during a successful previous interaction, the coordination pattern could not be established. An analysis of the system's underlying dynamics revealed (i) that stability of the coordination pattern requires ongoing mutuality of interaction, and (ii) that the interaction process is not only constituted by, but also constitutive of, individual behavior. We suggest that this stability of coordination is a general property of a certain class of interactively coupled dynamical systems, and conclude that psychological explanations of an individual's sensitivity to social contingency need to take into account the role of the interaction process.

Keywords: evolutionary robotics, social cognition, social contingency.

1 Introduction

Evolutionary robotics is typically employed to investigate simulation models of minimally cognitive behavior, namely the simplest behavior that raises issues of genuine cognitive interest [1, 8, 7]. Recently there have been some initial efforts to extend this methodology into the domain of social cognition [4, 9, 10, 11, 3, 14]. While these are important advances within the field of artificial life, they also generate insights of relevance that could form the basis for mutually informing collaborations with the empirical sciences such as psychology.

One promising target for such an endeavor is Murray and Trevarthen's [12] double TV monitor experiment. In this psychological study 2 month old infants were animated by their mothers to engage in coordination via a live double video link. However, when the live video of the mother was replaced with a video playback of her actions recorded previously, the infants became distressed or removed. These results, and those of a more rigorous follow-up study by Nadel and colleagues [13], indicate that 2 month old infants are sensitive to social contingency, i.e. the mutual responsiveness during an ongoing interaction, and that this sensitivity plays a

fundamental role in the unfolding of coordination. Traditional explanations of this sensitivity have focused on innate abilities. For example, Gergely and Watson [6] have postulated the presence of an innate cognitive module which enables the detection of social contingency, and Russell [15] hypothesizes that infants have an innate capacity to understand intentionality and to process agency.

Are these postulations of innate capacities on the part of the infant necessary in order to explain the empirical results? Iizuka and Di Paolo [9] used an evolutionary robotics approach to test whether simpler solutions could also emerge from the dynamics of the interaction process itself. In their simulation model the evolved agents successfully acquired the capacity to discriminate between ‘live’ (two-way) and ‘recorded’ (one-way) interaction. Moreover, an analysis of the resulting dynamics suggests that the interaction process itself plays an important role in enabling this behavior. Similar results were also found by the other simulation studies [4, 11, 10].

It could be argued that the result of Iizuka and Di Paolo’s [9] simulation study only represents a specific subset of the general solution space, in particular because they used evolutionary robotics to *explicitly* generate agents that terminate interaction when there is a lack of social contingency. We address this issue by testing *whether termination of interaction emerges under more general conditions*. Answering this question is important if the argument is made that these findings might apply more generally and in particular to human interactions. By changing the simulation setup in this manner we thus move closer to the original double TV monitor experiment: the infants presumably did not have the specific goal to detect whether they were dealing with a live video or just a recording. It is more likely that they were simply attempting to establish social coordination with their mothers but were unable to do so.

2 Methods

We implemented a minimal simulation model analogous to Murray and Trevarthen’s [12] double TV monitor experiment by building on work by Iizuka and Di Paolo [9]. A schematic of this simulation model is illustrated in Fig. 1.

The goal of the agents is to cross their sensors as far away from their starting positions as possible, a task which requires mutual localization, convergence on a target direction, and movement in that direction while not losing track of each other. This task is non-trivial since sensory stimulation only correlates with the overlapping of position (when the centers of the agents are less than 20 units of space apart); it does not convey the direction or speed of movement of the other agent.

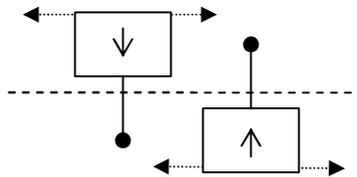


Fig. 1. A schematic view of the model adapted from Iizuka and Di Paolo [9]. The two identical agents are 40 units wide, only able to move in a horizontal direction, and equipped with a single on/off sensor at their centre. They face each other in an unlimited continuous 1-D space.

The agents are controlled by two identical continuous-time recurrent neural networks (CTRNNs), as described by Beer [1]. They were chosen to be clones because work by Iizuka and Ikegami [10] on a related task suggests that genetically similar agents are potentially better at coordination. They face each other in an unlimited continuous 1-D space (i.e. one agent faces ‘up’ and one agent faces ‘down’). Distance and time units are of an arbitrary scale. Each agent can only move horizontally. One on/off sensor is located in the centre of each agent. The sensor is activated (set to 1) when the agents cross each other, otherwise it is set to 0.

Noise is introduced into the simulation for 2 main reasons: (i) since the agents are identical they will need to make use of noise in order to break the symmetry of their movements and converge on a common target direction, and (ii) robustness against noise increases the ability of ‘live’ agents to cope with playback situations [10]. Accordingly, at each Euler time step there is a 5% probability that the current sensory state is flipped into its opposite state. We add a small perturbation to the motor outputs at each time step drawn from a Gaussian distribution ($\mu = 0$; $\sigma^2 = 0.05$). The noise is applied to the outputs before the application of motor gains.

In order to further increase the robustness of the behavioral strategies, the initial relative displacement between the agents varies (range [-25, 25]). Starting from any of these possible relative positions, the task for the agents is to coordinate their behavior such that they cross each other as far away from position 0 as possible. Since the agents are started in opposite orientation (‘up’ vs. ‘down’), it is not possible for the evolutionary algorithm to hard code any trivial solution (e.g. ‘always move left’).

2.1 Agents

The agents are 40 units wide, have an on/off sensor at their center, and can only move left or right by controlling the output of their left and right motor nodes (see Fig. 1). Agents are controlled by a CTRNN consisting of 3 fully-connected nodes with self-connections. The time evolution of the node activation follows:

$$\tau_i \dot{y}_i = -y_i + \sum_{j=1}^N w_{ji} z_j(y_j) + S I_i, \quad z_i(x) = 1/(1 + e^{-x-b_i}) \quad (1)$$

In this equation y_i represents the cell potential of node i , z_i is the firing rate as calculated by the standard sigmoid function, τ_i (range [1, 100]) is its time constant, b_i (range [-3, 3]) is a bias term, and w_{ji} (range [-8, 8]) is the strength of the connection from the neuron j to i . I_i represents the sensory input to node i and S is the sensor gain. The total number of nodes N is set to 3; there are no hidden nodes (all nodes receive sensory input). The sensory input is calculated by multiplying 1/0 (on/off) by an sensor gain parameter S (range [1, 100]), and this is applied to all nodes. There is one node, which only receives input and does not produce motor output and two actuator nodes for controlling movement; one for leftward and the other for rightward velocity. Each velocity is calculated by mapping the actuator output onto the range [-1, 1] and then multiplying it by an output gain parameter (range [1, 50]). The overall agent velocity is calculated as the difference between the left and right velocities. The time evolution of the simulation environment and each agent’s CTRNN controller is calculated by using Euler integration with a time step of 0.1.

Similar settings have already been successfully used by Iizuka and Di Paolo [9]. The main differences are that (i) the agents of the current study only have 3 nodes, (ii) the input is fed to all nodes instead of one dedicated sensory node, (iii) and each actuator node has its own gain parameter. The first difference was chosen to further minimize the conditions of the model and facilitate analysis; differences (ii) and (iii) were implemented to increase the evolvability of the solutions.

2.2 Evolutionary Algorithm

The agents are optimized by using a simple genetic algorithm (GA) which is based on the microbial GA, a steady-state GA with tournament selection [7]. Until some termination criterion is reached, two members of the population are chosen at random, both have their fitness evaluated, and while the ‘winner’ of the tournament remains unchanged in the population, the ‘loser’ is replaced by a slightly mutated copy of the ‘winner’. Each member is a clonal pair of agents. We define a generation as the number of tournaments required to generate a number of offspring equal to the population size. The population size is 40 and the run terminates at 5000 generations.

All CTRNN parameters and gains are genetically encoded by a real-valued vector which is initialized randomly. The mutation operator changes each gene by a random value drawn from a Gaussian distribution ($\mu = 0$; $\sigma^2 = 0.05$) with reflection at the gene boundaries. Before evaluation, each gene is decoded linearly to the corresponding range (except gains and time constants which are exponentially scaled).

During each fitness evaluation an agent is tested in 15 trials runs; to increase the robustness of the evolving solutions to noise and variations in initial conditions only the lowest score achieved in any of the trials is chosen as the overall score. Each trial run consists of 50 units of time (500 Euler time steps). At the start of each trial agents have their internal node activations set to small random values drawn from a standard Gaussian distribution. The initial distance between the agents varies; agent ‘down’ always gets placed at position 0, while agent ‘up’ starts at a different position for each trial (15 different positions evenly distributed across range [-25, 25]).

The fitness score of a trial run is calculated on the basis of a single factor, namely the absolute value of the final crossing position of the two agents divided by a factor of 10. Thus, in contrast to the work done by Iizuka and Di Paolo [9], these agents were not evolved to break off the interaction pattern when detecting a lack of social contingency. Instead, we aimed to generate a simulation model that under normal circumstances results in highly fit coordination behavior. Presumably, such behavior should be more robust when faced with the ‘playback’ condition.

3 Results

The GA was run 4 times. The fittest agent, with a score of 244.8, was produced during the 4th run in generation 3477. This solution was then tested extensively; agent ‘down’ was always placed at position 0, while agent ‘up’ starts at a different position for each trial (101 positions evenly distributed across range [-50, 50]). Each trial is repeated 150 times. The mean score across this range of initial conditions is plotted in Fig. 2 (left). The agents are able to generalize their behavior well beyond the range

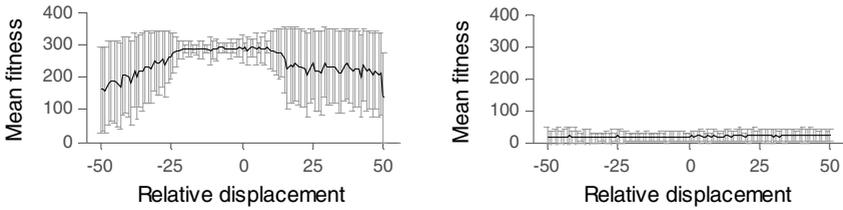


Fig. 2. *Left:* Mean score achieved by the fittest agent starting from various initial positions, with standard deviation. *Right:* Mean score by the fittest agent but this time interacting with non-responsive, recorded movements obtained from the original trials.

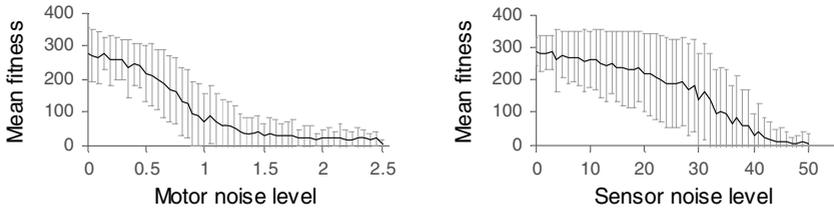


Fig. 3. Robustness to noise: mean fitness score achieved over 150 trials by the fittest evolved agent starting from position 11 for a range of noise levels, with standard deviation. Original noise strength during evolution is 0.05 for motor (left) and 5% for sensor noise (right).

that they were originally evolved to cope with. On average the best initial position for agent ‘up’ turned out to be at 11 (mean score: 292.9).

In order to demonstrate the general robustness of the evolved agents under this initial condition, we ran another set of trials with agent ‘up’ starting from position 11 while varying noise levels. The motor noise was varied while the sensor noise remained constant at evolutionary strength (5%), and sensor noise was varied while motor noise remained constant ($\sigma^2 = 0.05$). At each noise level we tested the agents for 150 trials. As shown in Fig. 3, the agents are able to cope with a wide range of perturbations. Indeed, their overall performance degrades gracefully until the sensor and motor signals are completely swamped by noise. In the case of sensor noise, for example, average performance only approaches 0 just before reaching the 50% mark (at which point sensory activation becomes completely arbitrary). This demonstrates that the agents are able to produce highly robust coordination behavior.

Finally, another 150 trials were conducted with agent ‘up’ at position 11 (under normal noise conditions). The movement of agent ‘down’ during the best trial (score: 321) was recorded for playback. Another 150 trials were then run under *playback* conditions: the initial conditions reflect those of the recorded best trial run (agent ‘up’ always starts at position 11 and with the same initial internal activation), and the movement of agent ‘down’ replicate those which it produced during the recording. While the sensorimotor noise for agent ‘up’ was different during each of these trials, no additional noise was added to the recorded movement of agent ‘down’.

The results are striking: whereas the original 150 trials of mutual (two-way) interaction were highly successful (mean score: 268), the 150 trials of playback

(one-way) interaction were a drastic failure (mean score: 19). The severity of this failure is surprising since under normal conditions the active agent is robust against various forms of noise, and able to cope effectively with a wide range of initial conditions. Moreover, during the playback condition its ‘partner’ performs what had previously been a highly fit behavioral repertoire. Still, the active agent is unable to adapt to the situation of interacting with a non-responsive ‘partner’. It could be argued that this result is unique to the chosen situation. However, this is not the case: when testing agent ‘up’ with each of the original trials we get the same result (see Fig. 2, right).

3.1 Behavioral Analysis

The behavior of the agents under normal conditions can be broken down conceptually into three important aspects: (i) localization, (ii) alignment, and (iii) coordination. We will briefly discuss the first two aspects and then focus on the third. The activity during the first time steps of the best trial run is shown in Fig. 4.

Initially the agents have no knowledge of how their own position relates to that of their partner. Moreover, they have no way of gaining that information except when changing their sensory input by engaging in movement. However, it turns out that one stereotypical behavioral pattern is sufficient to solve the non-trivial problem of reliable localization. First, each agent moves rightwards for a few units of time, and then starts moving leftwards. This sweeping behavior usually takes up to 5 units of time and under evolved conditions always enables the agents to locate each other. In the case of negative initial displacement they will encounter each other during their rightward sweep; otherwise they will cross their positions during their leftward return. Interestingly, the agents always end up with positive *relative* displacement after their initial localization. With this clever maneuver the agents have significantly reduced the complexity of their coordination task: while sensory input is ambiguous (there is

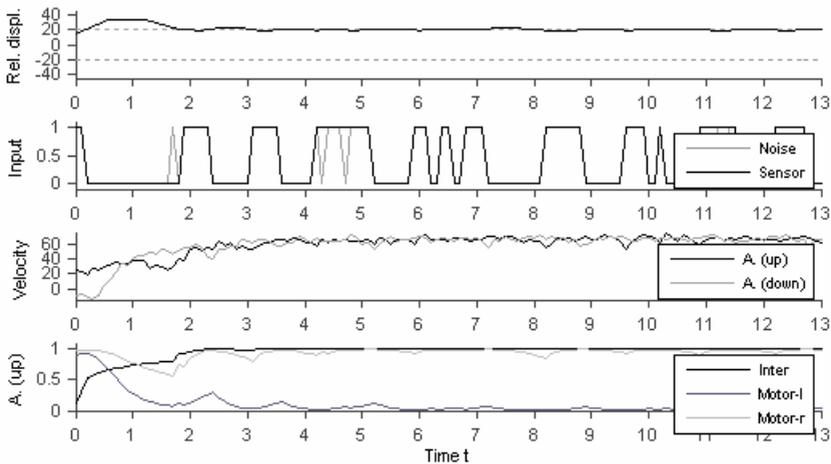


Fig. 4. Initial activity of the two agents during the best trial run. From top to bottom the traces show the evolution over time of (i) their relative displacement, (ii) their noisy input signal and actual sensory contact, (iii) their velocity, and (iv) the CTRNN node outputs of agent ‘up’.

no indication about the direction or speed of the other agent's movement), it has now been *co-arranged* as a 'touching on the left' indicator! This change of the sensory meaning is possible because the CTRNN controllers are not symmetric.

How does the final oscillatory coordination pattern emerge out of the relative movements of the agents? Before analyzing the behavior of the agents in more detail it is necessary to briefly describe the evolved CTRNN controller. Most importantly, the sensory input excites all of the nodes with a gain of $S = 10.9$, and the right output gain (44.5) is almost twice as high as the left output gain (24.9). The two motor nodes are inhibited by the non-motor node and they also inhibit each other while hardly affecting the non-motor node. As an example, we can see that the output of the right motor node (z_3) of agent 'up' starts to slightly decrease just before time $t = 8$, due to lack of sensory stimulation. This shift in velocity entails that agent 'down' catches up with agent 'up' and they remain in contact ($I_i = 1$) until just before $t = 9$. During this contact agent 'up' regains its previous rightward velocity due to sensory stimulation. After separating again ($I_i = 0$) the firing of the left motor node goes down followed by the right motor node which eventually leads to the behavioral pattern being reinitiated. Accordingly, agent 'up' should be able to engage with a playback recording. The activity during the *playback* trial run is shown in Fig. 5.

At first the 'live' agent aligns itself with the 'playback' agent as in the original situation (Fig. 4). During mutual (two-way) interaction agent 'down' would always respond to contact by moving away slightly. However, in the playback situation this co-regulation is prevented from occurring. Accordingly, every encounter results in a slight decrease of relative displacement between the two agents, thereby in turn making it more likely that there will be another sensory stimulation. Up to about $t = 3$, agent 'up' is still able to partially regulate this displacement on its own by adjusting the output of its right motor node. However, from that point onwards the right motor node saturates at $z_3 = 1$, and thereafter remains unaffected by further sensory

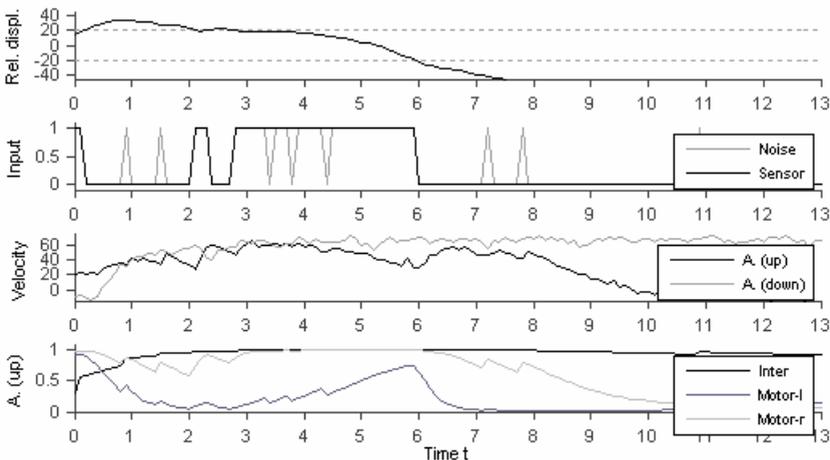


Fig. 5. Initial activity during a *playback* trial run in which the movements of agent 'down' are the same as in Fig. 4. From top to bottom the traces show the evolution over time of (i) the relative displacement between the agents, (ii) their noisy input signal and the actual moments of sensory contact, (iii) their velocities, and (iv) the CTRNN node outputs of agent 'up'.

stimulation. Finally, at around $t = 6$ the positive feedback loop between increasing sensory stimulation and mounting leftward velocity becomes unstable in such a way that recovery from breakdown is impossible. The live agent falls behind the playback agent, and heads in the opposite direction.

Why does this breakdown of coordination not occur when both agents engage in ‘live’ interaction? The simple answer provided by this model is that the *stability of ongoing coordination requires mutuality of interaction*. After the initial alignment we find that coordinated movement in one direction consists of continuous co-regulated oscillatory behavior. Agents control their respective velocities such that they cross their sensors at relatively regular intervals. This iterative reaction chain constitutes an ongoing pattern of turn-taking; noise perturbations get amplified in a way that requires continuous co-regulated re-establishment of the interaction [11].

3.2 Dynamical Analysis

Can we account for the oscillating pattern in dynamical terms? Since the output of the non-motor node $z_l(y_l) = 1$ during coordination, it can be treated as a fixed parameter. The rest of the system consists only of the two motor nodes. The parameters are $\tau_2 = 1.6$, $b_2 = 2.6$, $w_{12} = -3.7$, $w_{22} = 1.0$, $w_{32} = -7.9$, and $\tau_3 = 1.1$, $b_3 = 2.9$, $w_{13} = -5.8$, $w_{23} = -5.8$, $w_{33} = 2.3$. If agents are not in contact with each other ($I_i = 0$), there is a globally attracting stable equilibrium point at $(-3.4, -7.5)$. Being in this state effectively slows down rightward velocity. Because of this the agents eventually make contact. When $I_i = 1$ the equilibrium point is shifted to $(0.3, 1.9)$. This effectively speeds up the rightward velocity of the agent.

Interestingly, under normal conditions the dynamical system never reaches either of these two equilibrium points, because their existence is made transitory through the ongoing interaction. This is illustrated in Fig. 6 (left) in terms of the motor node firing rates for agent ‘up’ over a whole run (50 units of time). The trajectory settles down into an oscillatory pattern that traces the corner near point $(0, 1)$, in the middle of the two equilibrium points (located at $(0.95, 1)$ when $I_i = 1$, and at $(0.30, 0)$ when $I_i = 0$). The state trajectory for the playback situation of the same run is displayed in Fig. 6 (right). At first the trajectory moves into the same region of state space but then, during the period of prolonged contact, the left motor node gets saturated while the right motor node remains at 1. This continues until the system almost reaches the equilibrium point at $(0.95, 1)$, but it eventually causes agent ‘up’ to slow down too much thereby breaking out of the coordination pattern.

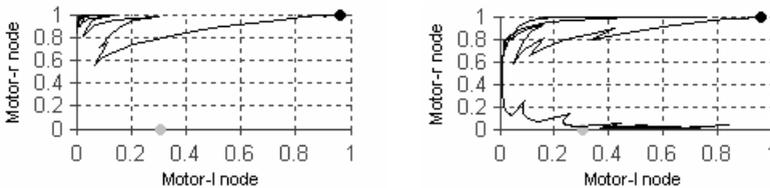


Fig. 6. State trajectory of the outputs for the 2 motor nodes of agent ‘up’ during 50 units of time. The trace starts at the top right corner of each graph. The gray and black dot represent the globally attracting stable equilibrium point when sensory input $I = 0$ and $I = 1$, respectively. *Left:* mutual (two-way) interaction. *Right:* playback (one-way) interaction.

After agent ‘up’ slows down enough such that the playback movement of agent ‘down’ overtakes it, its input I_i stays at 0. This causes its motor system to settle near the equilibrium point at (0.30, 0), from where it is occasionally perturbed by sensor noise. Thus, without the responsive help of the other agent, agent ‘up’ is unable to regulate its behavior such as to avoid falling into this attractor, an event which limits its further behavior to mere leftward movement.

4 Concluding Remarks

With our simulation study it was found that stable and robust coordination can be reliably established between simulated agents. While the agents were only selected on the basis of this coordination ability (rather than their capacity to detect social contingency), coordination still breaks down when a ‘live’ agent is forced to interact with a playback of movements from a previous, successful trial. Agents interacting with such a non-responsive ‘partner’ do not have the capacity to generate and sustain the kind of oscillatory behavior necessary for coordination. Thus, what at first appears to be a *behavioral capacity of the individual agent* emerges out of a combination of the internal dynamics *as well as the interaction process*. This role of the interaction process is also consistently demonstrated in previous models [9, 4, 11, 3].

We are thus faced with a peculiar situation in which the behavior of the individual agents brings forth the interaction process, and that interaction process enables the behavior of the individual agents. This makes a reduction of the coordination breakdown to an individual agent’s capacity to detect social contingency impossible. Moreover, it points to the *autonomy* of the interaction process, as postulated by an enactive approach to social cognition [2]. A more detailed analysis of the dynamics of the interaction process in this context is desirable, especially in terms of an artificial life investigation of the systemic basis of “constitutive autonomy” [5].

It is worth emphasizing that we do not claim that our model instantiates the phenomenon which we are investigating or that the baby-mother interaction studied by Murray and Trevarthen [12] is reducible to such a simple system. The model is purely conceptual in that it shows at work a possible explanation that may later be considered and tested in specific empirical cases. Thus, by generating simple models which do not presuppose the methodological individualism which prevails in social cognitive science and psychology, we can re-conceptualize the space of possible explanations [4]. In particular, the model presented in this paper suggests that the capacity for social behavior is strongly dependent on the existence of an appropriate social context, one whose stability is in turn dependent on the active and responsive engagement of the participants. On this basis we propose that an explanation for the distressed reaction, observed when confronting infants with a video recording rather than a live stream of their mother, also needs to take into account the role of the interaction process. Of course, this does not mean that the infants cannot detect social contingency or that they cannot develop this ability, but it does open up the possibility for explanations that do not suppose any necessity for innate behavioral capabilities and/or a complex perceptual strategy on the part of the infant.

References

1. Beer, R.D.: The dynamics of adaptive behavior: A research program. *Robotics and Autonomous Systems* 20, 257–289 (1997)
2. De Jaegher, H., Di Paolo, E.A.: Participatory sense-making: An enactive approach to social cognition. *Phenomenology and the Cognitive Sciences* 6, 485–507 (2007)
3. Di Paolo, E.A.: Behavioral coordination, structural congruence and entrainment in a simulation of acoustically coupled agents. *Adaptive Behavior* 8, 25–46 (2000)
4. 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* (in press)
5. Froese, T., Virgo, N., Izquierdo, E.: Autonomy: a review and a reappraisal. In: Almeida e Costa, F., et al. (eds.) *ECAL 2007. LNCS (LNAI)*, vol. 4648, pp. 455–464. Springer, Heidelberg (2007)
6. Gergely, G., Watson, J.: Early social-emotional development: contingency perception and the social biofeedback model. In: Rochat, P. (ed.) *Early social cognition: Understanding others in the first months of life*, pp. 101–137. Lawrence Erlbaum, Hillsdale (1999)
7. Harvey, I.: Artificial Evolution: A Continuing SAGA. In: Gomi, T. (ed.) *ER-EvoRob 2001. LNCS*, vol. 2217, pp. 94–109. Springer, Heidelberg (2001)
8. Harvey, I., Di Paolo, E.A., Wood, R., Quinn, M., Tuci, E.A.: Evolutionary Robotics: A new scientific tool for studying cognition. *Artificial Life* 11, 79–98 (2005)
9. Iizuka, H., Di Paolo, E.A.: Minimal Agency Detection of Embodied Agents. In: Almeida e Costa, F., et al. (eds.) *ECAL 2007. LNCS (LNAI)*, vol. 4648, pp. 485–494. Springer, Heidelberg (2007)
10. Iizuka, H., Ikegami, T.: Adaptability and diversity in simulated turn-taking behavior. *Artificial Life* 10, 361–378 (2004)
11. Ikegami, T., Iizuka, H.: Turn-taking interaction as a cooperative and co-creative process. *Infant Behavior and Development* 30, 278–288 (2007)
12. Murray, L., Trevarthen, C.: Emotional regulations of interactions between two-month-olds and their mothers. In: Field, T.M., Fox, N.A. (eds.) *Social perception in infants*, pp. 177–197. Ablex Publishing, Norwood (1985)
13. Nadel, J., Carchon, I., Kervella, C., Marcelli, D., Réserbat-Plantey, D.: Expectancies for social contingency in 2-month-olds. *Developmental Science* 2, 164–173 (1999)
14. Quinn, M.: Evolving communication without dedicated communication channels. In: Kelemen, J., Sosík, P. (eds.) *ECAL 2001. LNCS (LNAI)*, vol. 2159, pp. 357–366. Springer, Heidelberg (2001)
15. Russel, J.: *Agency: Its role in mental development*. Taylor & Francis, Hove (1996)