

The Role of Co-Representations in Joint Tracking

Ekaterina Sangati^{1,2} and Simon M. Hofmann³

¹Embodied Cognitive Science Unit, Okinawa Institute of Science and Technology Graduate University
1919-1 Tancha, Onna-son, Okinawa 904-0495, Japan

²Faculty of Philosophy, Theology and Religious Studies, Radboud University Nijmegen
Erasmusplein 1, 6525 HT Nijmegen, The Netherlands

³Department of Neurology, Max Planck Institute for Human Cognitive and Brain Sciences
Stephanstraße 1a, 04103 Leipzig, Germany
ekaterina.sangati@oist.jp

Abstract

In this paper we present a Minimal Cognitive Agent model of a joint action task. Pairs of agents realized as Continuous Time Recurrent Neural Networks are submitted to artificial evolution in the context of a task taken from psychological literature. In this task the agents are required to coordinate their complementary actions in order to jointly control the movement of a tracker and successfully follow a continuously moving target. It has been suggested that such a task requires a more complex type of cognitive mechanism than the types of processes postulated by the proponents of Embodied Embedded Cognition approach. Specifically, it might possibly require that the agents “co-represent” each other’s contributions to the common behavior. Our results show that simple agents with no such built-in co-representation mechanism are able to evolve a solution to the task. However, we also find emergent neural activity patterns that are consistent with it. In what sense these patterns can be said to be truly representational requires further study.

Introduction

The Embodied Embedded approach to cognition (EEC) proposes that rather than being an activity of computation over representations, cognition is primarily a matter of adaptive behavior emerging from the agent’s concrete bodily abilities, embeddedness in a particular environment, individual organismic needs and a history of agent-environment interaction. Importantly, EEC does not deny that the agent’s internal states also play a role in producing adaptive behavior. However, it questions whether conferring to such states a status of representations is explanatorily useful, once a full context of the agent’s resources is taken into account.

One way to test whether any specific cognitive achievement requires representations is via Minimal Cognitive Agents (MCAs) methodology (Beer, 1996, 2008). Agents with minimal brains (typically realized in neural networks with only a few nodes) are evolved to perform tasks that seemingly require representations (“representation-hungry problems”; Clark and Toribio, 1994), and then analyzed as to how successful behavior is actually produced in them

and whether that process relies on anything that can usefully qualify as a representation. Letting the solution to a cognitive task emerge, rather than building it in, allows one to question a priori assumptions about how such a task should be accomplished, including assumptions about specific types of cognitive mechanisms. This strategy can be applied both to generally defined cognitive capacities, such as perceptual categorization (Beer, 2003) and to tasks drawn from actual psychological experiments (Harvey et al., 2005).

In this paper we make preliminary steps in investigating the role of representational mechanisms posited in accounts of social behavior. We focus specifically on joint action, defined as a form of social interaction “whereby two or more individuals coordinate their actions in space and time to bring about a change in the environment” (Sebanz et al., 2006). The EEC camp has produced a number of MCA models of social interaction in the recent years. A large proportion of these, however, has been aimed at demonstrating a particular idea that social interaction is best understood as an emergent phenomenon, not reducible to the individuals that partake in it (De Jaegher and Di Paolo, 2007). For instance, several models (Di Paolo et al., 2008; Froese and Di Paolo, 2008, 2010) have been proposed for human capacity of agency detection, i.e. an ability to recognize when one is interacting with a live, responsive partner rather than an inanimate object (Murray and Trevarthen, 1985; Auvray et al., 2009). It has been shown that such recognition does not necessarily require an agency detection module inside an individual’s brain but can lie in interaction dynamics, i.e. the fact that interaction with a live partner is more stable, more self-reinforcing.

We believe these studies are interesting and important. However, if the goal is to provide a genuine challenge to the dominant paradigm in cognitive science, they are lacking for two reasons. First, to show that evolved agents are not in fact trading representations of (social) stimuli, a solid attempt at investigating the representational status of their internal states is required and this has not been done yet with respect

to the minimal social agents studies described above.¹ Second, we think the tasks that have been used so far are not sufficiently representation-hungry to convince a staunch proponent of a representation-based approach. This is because in the joint action literature it is generally acknowledged that a certain portion of social interactions (e.g. synchronization effects; Richardson et al., 2007) can be explained in dynamical non-representational terms, referred to as entrainment or emergent coordination. However, it is maintained that in more complex cases, all entrainment can do is make coordination more smooth and timely. Apparently, when a situation requires discrete actions, when no online information about co-actor’s movements is available or when prediction is necessary, people need to represent the co-actor’s task, their perceptual point of view or knowledge states, they need to use forward models to predict the co-actor’s upcoming actions and so on (Knoblich et al., 2011; Sebanz and Knoblich, 2009a; Vesper et al., 2010).

In light of such views, the MCA strategy needs to be applied to a task in which mere non-representational entrainment is believed to be insufficient. An example of such a task is a study by Knoblich and Jordan (2003), henceforth KJ. It is explicitly set up by the authors in such a way as to require (1) coordination with a discrete action alternative of the co-actor, not just its timing (2) coordination with a predicted action of the co-actor, not merely a preceived one.

The KJ task environment is a one-dimensional field which contains two objects that can overlap: target and tracker (see Figure 1). On any trial, one of the objects, the target, starts from the center of the screen and moves to the screen’s border (left or right), then reverses its direction, moves to the other border and does so again, completing 3 turns. The second object, the tracker, is controlled by the participant and their goal is to keep it on top of the target for as much time as possible by pressing two buttons: pressing the left button accelerates the tracker to the left, pressing the right button – to the right. When this task is performed by two people, each is responsible for tracker acceleration in only one of the directions.

The difficulty in this task is that once some amount of acceleration has been added to the tracker in one direction by pressing one of the buttons, the same amount of complementary button presses is required in order to first decelerate the tracker and then accelerate it in the opposite direction. This means that two strategies are possible in the region close to the screen border (the border region). Acting merely reactively would mean decelerating the tracker only once the target has reached the border and reversed its direction, which would result in adequate tracking before that moment but subsequent rapid accumulation of error (because the target moves at constant speed). A more sophisticated anticipatory strategy would mean starting to decelerate and accruing a

¹See, for instance, Mirolli (2012) for such an attempt applied to perceptual categorization.

small amount of error before the target reaches the border in order to minimize future error. In a joint setting this requires being able to coordinate with predicted actions of the co-actor controlling the complementary action alternative.




Figure 1: The joint tracking task environment (figure based on KJ and Sebanz et al. (2006)).

The details of the experimental setup and conditions do not concern us here. What is important is that good performance in the joint condition was shown to be related to the use of anticipatory strategy and a particular pattern of results obtained was interpreted by KJ and others as based on learning a forward model of the co-actor’s task behavior.² That is, it is thought that each participant learned to predict the timing of the other’s complementary action and the joint effects of combined responses and used these predictions to coordinate one’s own action.

In this paper we set aside the question of whether this interpretation of task behavior is plausible and we postpone the discussion of possible representational claims that underlie it to the Discussion. Our goals are more modest. First, we investigate whether MCAs can at all evolve to perform a version of the KJ representation-hungry tracking task and, if so, whether they are capable of adopting the anticipatory strategy. Second, we undertake some exploratory steps in answering the question whether the agents represent anything about their co-actor. Following terminology in the joint action literature we call this putative process *co-representation* and we define it minimally as *an internal state of one agent that correlates with future action of the other agent and functionally contributes to the generation of coordinated joint performance*.³ We examine whether such a state can be detected in the evolved agents’ brain activity.

²“[Learning to coordinate] involves the acquisition of new predictive models during joint practice. A study by Knoblich and Jordan (2003) provides evidence that, through training, it is possible to integrate predictions about one’s own actions and the actions of others” (Knoblich and Sebanz, 2006, p. 236). “Feedback led partners to develop a model of each other, allowing them to anticipate each other’s action timing” (Eskenazi et al., 2012, p. 102).

³There is considerable discussion as to the precise content of co-representations, e.g., Sebanz et al. (2005); Dolk et al. (2016).

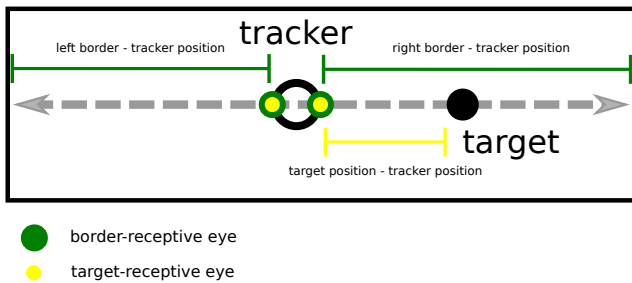


Figure 2: Agent environment.

Methods

Task environment

The joint action paradigm created by Knoblich and Jordan (2003) was implemented in a simulated model as follows. The environment consisted in a 1D line 40 units wide, surrounded by borders. It contained a target and a tracker that could move horizontally between the borders in the increments determined by their speed and the time step granularity set to 0.01. The collaborating agents were located “inside” this environment, as if on top of the tracker and both in the exact same position (see Fig. 2).

Each agent pair in a single simulated generation was subjected to six trials that were a result of crossing three levels of target velocity and two initial target directions (left, right). Each trial started with both the target and the tracker in the middle. After an initial starting period (of $T = 100$ time steps) in which both objects remained stationary, the target would start moving away from the center at constant velocity. Upon reaching one of the borders, it changed direction. Each trial consisted of three target turns, ending again in the middle of the environment. The tracker movement was controlled by a pair of agents, where one member of the pair was responsible for the left and another for the right tracker velocity. When the tracker reached the border, its velocity was set to zero – even though the agents could still have non-zero output of their motor neurons. In this way, the tracker’s “virtual” non-zero velocity that forces the agents to decelerate in one direction before accelerating in the other is instantiated in our model in the motor outputs.

Agent architecture

Each agent received perceptual input and was producing motor output. The agents had four “eyes” directed towards left and right that received information about the absolute distance to the two borders and the target (see Fig. 2). That is, border-receptive eyes were constantly receiving input (since the agents were always in between the two borders) while target-receptive eyes were receiving input only if the target was located on the side for which a given eye was responsible (left or right). The distance was scaled linearly in such

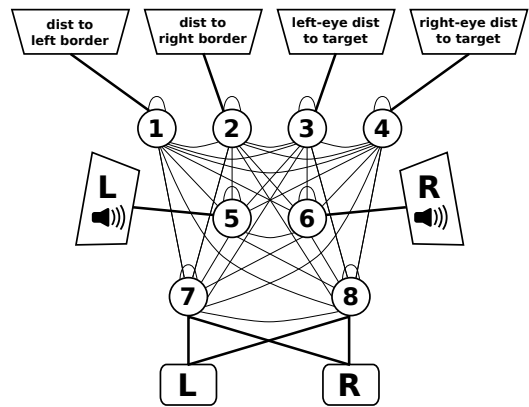


Figure 3: Agent morphology.

a way that the maximum visual input was 10, when the distance was 0 and minimum was 0, when the distance was maximal, i.e. 40. For the target-receptive eyes the input was also 0 when the target was on the opposite side to a given eye and it was maximal to both eyes when the tracker was directly on top of the target.⁴ The agents also received auditory input. Each sensor (“ear”) was activated by just one of the sounds produced by the motors: left or right. The sound was produced continuously by the motor activation and was directly proportional to its magnitude.

The motor output of the agents controlled the movement of the tracker. We have experimented with two modes of control: discrete button-based and direct velocity control. However, pilot studies showed that the former does not lead to reliable results and therefore we focus on the direct velocity control in the remainder of this paper (but see Discussion for possible limitations of this approach). Each agent had two motors (effectors), left and right, where the left motor velocity could only be negative and the right motor velocity could only be positive. Each agent contributed only one effector output to the tracker control and tracker velocity was a simple sum of motor activation of the two effective motors.

The agents in the present study were controlled by a network of 8 neurons as depicted in Figure 3. The 8 neurons were fully inter-connected and self-connected. There were additional visual and auditory weights (range $[-100, 100]$) from perceptual input to the neurons, as well as motor weights (range $[0, 10]$) from motor neurons to left and right motors. No symmetry was imposed on the network.

CTRNN update rule for each neuron was defined by the

⁴We have tested two basic ways of providing agents with visual input: as absolute target and tracker position and as relative distance. That is, in the first case the agents were “outside” the environment, looking at it as if presented on a screen, while in the second they were “inside”, as if moving on top of the tracker. We have observed that the second architecture led to faster and more reliable solutions, indicating perhaps that MCAs evolve more easily when the task is setup in a more embodied way.

following standard formula:

$$\dot{y}_i = \frac{1}{\tau_i}(-y_i + \sum_{j=1}^N w_{ji}\sigma(y_j + \theta_j) + I_i), i = 1, 2, \dots, N \quad (1)$$

where y is the state of each neuron, τ is its time constant (range [1, 100]), w_{ji} is the connection strength (the weight) from the j^{th} to the i^{th} neuron (range [-15, 15]), θ is a bias term (range [-15, 15]), σ is the standard logistic activation function and I_i represents an external (perceptual) input to the neuron. We applied the Euler integration method with a time step set to $h = 0.01$ to obtain the time evolution of the simulation.

Evolutionary algorithm

The behavior of the agents was evolved using a real-valued mixed genetic algorithm (GA) applied to their neural network parameters. The genomes were composed of all parameters (time constants, bias terms, weights of interneurons, visual sensors, auditory sensors and effectors) for the total of 90 positions separated into 6 indivisible modules.

The population size was $N = 100$ split equally in 2 populations for agents controlling left and right effectors. The population of agents was initialized with random parameters drawn from their full range. After each trial, agent fitness scores were calculated according to this formula:

$$F = 1 - \frac{\sum_{i=1}^T d_i}{D \cdot (T + S)} - P \quad (2)$$

$$P = \frac{1}{T} |\{v_i \mid v_i = 0 \text{ for } i = 1, 2, \dots, T\}| \quad (3)$$

where d_i is the absolute distance between target and tracker at a given simulation step, T is the simulation length ($T \in [3000, 4000, 6000]$, depending on target velocity), D is the maximum possible distance between target and tracker (equal to 40 for all simulations), S is the starting period length (equal to 100 for all simulations), P is the penalty function, which is defined as the average number of times the tracker velocity was 0 in a given trial. We added this penalty function to the fitness calculation to encourage movement of the agents. If penalty decreased the fitness score below 0, it was clipped to 0. The overall fitness score for every agent was calculated as harmonic mean over scores in all trials. The same fitness score was assigned to each member of the pair that controlled the tracker. Evolution operations were applied separately to left and right sub-populations of agents.

A new generation was created by copying the best agents without modification (5% of the new population), applying fitness proportionate selection to the whole population

(80%), with rank-based, stochastic universal sampling selection, mutation variance set to 10 and probability of crossover to 0.8 (applied at genetic module boundaries). The remaining 15% of the population was filled up with new randomly initialized agents. The evolution was stopped after 2000 generations or after 500 generations passed without improvements of the best agents' performance.

Results

Behavioral strategies

In analyzing trial behavior of successful agents we have observed two main strategies. The most prevalent task solution that evolved can be called an "independent strategy" in which one agent provides a continuous motor output of a given magnitude and the other agent varies its output in an oscillatory manner (Fig. 4A). This allows the pair to produce overall velocity that oscillates around 0 leading the tracker to continually switch directions staying close to the target. A second strategy involves both agents producing oscillating output and its combination allowing the tracker to follow the target (Fig. 4B-D). In further analysis we focus on the latter, given that it is a more interesting case of collaboration and a more comparable case to the human study⁵. We specifically examine the best pair of the last generation from one population that evolved such a strategy and that achieved a fitness score of $F = 0.966$ (range [0, 1]).

Figures 4B-D present behavior in three of the trials of this best pair. Overall, the patterns look similar in all 6 trials but are not completely symmetric between those in which the target starts by moving to the right and those in which it first moves to the left. Furthermore, motor oscillations (and the resulting velocity and movement oscillations) show a lower frequency in faster trials, which allows the agents to track the target by producing only longer period displacements, with less oscillations around the target. This is accomplished by slowing down of the left motor oscillations (Fig. 4D).

The left-right asymmetry is also noticeable in the tracking strategy of the agents. In particular, they exhibit something that looks like anticipatory strategy (reversing before reaching the border) only at the right border while at the left border they seem to adopt a more compensatory strategy. The reason for this asymmetry is not clear but the behavior at the right border gives us a proof of concept that the agents *can* adopt anticipatory strategy. Perhaps a stronger selection pressure would lead to its more consistent application.

Examining the movement patterns it seems that rather than adopting a division of labor in which one agent controls movement to the left and the other to the right, the division is more along the lines of the left agent (call it L) producing fast oscillations around the target position and the right agent (R) producing slow oscillations as the target moves in be-

⁵Recall that in the original KJ study both participants had to act in complementary ways at different time points.




Figure 4: Behavioral patterns in different trials (trial durations re-sampled to 3000 time steps).

tween the two borders. Successful tracking is accomplished by an alternation of the length of leftward and rightward displacement which is a result of the oscillation frequency of the left motor and the output strength of the right motor. That is, the predominantly negative velocity is produced by slow oscillation of the left motor whose output is higher than that of the right with either both highly active or both only slightly active motors. The predominantly positive velocity is a result of fast left motor oscillations combined with constant high output of the right motor which pulls the overall tracker velocity above 0 with only short negative drops. It can also be noted that the right motor is active in a nearly exact same way across the different speed trials and the only difference between trials in which target initially moves left or right is in a phase shift of its activity.

To go beyond the description of experimental trials that the agents were trained on, we first checked the ability of the best pair to generalize to four new trial types: target starting from a set of random points within $[-3, 3]$ coordinates, new target speeds, borders increased to $[-30, 30]$ and the number of turns increased to 5. We found that the agents perform reasonably well in all these conditions, with an average score of $F = 0.955$.

Next, we conducted a series of lesion studies in which we disabled some perceptual input to the agents and observed the resulting behavior. We determined that auditory perception does not seem to be critical to the performance in this task. A lack of possibility to hear the co-actor’s motor activation leads to less precise tracking but it is still sufficiently good throughout the trial. Visual perception, on the other hand, predictably, does seem critical. If either perception of the border or the target is removed from the beginning of the trial, performance is severely disrupted. If border input is removed around the middle of the trial, from the time of the lesion the agents manage to complete faster trials but struggle on the slower ones. Such delayed loss of target information is slightly more detrimental – soon after the input is removed the left agent stops its oscillations leading the tracker off course. Interestingly, also here in the fastest trial adequate behavior seems to be maintained. This, together with

generalization results, confirms that the agents rely on on-line information about the target’s position and the tracker’s position relative to the borders in accomplishing the task, rather than simply memorizing the required behavior.

Finally, we checked whether the contribution of both agents is in fact required to accomplish the task and run the trials with just L or R agent acting alone and having control of both motors. With only L agent in control, its motors become maximally active and the tracker drifts to the right border and stays there immobile. With R agent acting alone, both of its motors start slow oscillations that results in the tracker moving to the right border and wiggling there without ever returning to the middle. Thus, it would seem that the agents trained on joint action cannot perform the task alone. However, following the EEC idea that interaction is what matters to social behavior, we performed a further test of the actual “jointness” of agent behavior: by running trials in which only one agent was actually performing the task and the contribution of the other agent was played back from previous runs of the same trials. We found that this makes virtually no difference to the overall behavior of the tracker, suggesting that the agents do not evolve to be sensitive to the live interaction partner but rather to merely control the tracker independently given the background conditions of an assumed co-actor’s reliable contribution. We return to this observation in the Discussion.

Cross-agent relationships

We have defined co-representation as “an internal state of one agent that correlates with future action of the other agent” which in our model translates into a correlation between the state of the nodes of one agent and the motor output of the other agent with a time lag. In order to capture this relationship we employed an information dynamics approach, which allows one to assess how information flows through the network (Williams and Beer, 2010; Bosso-maier et al., 2016).⁶ This approach can be used to investigate statistical relationships between particular variables but

⁶It is important to emphasize that ‘information’ here is meant in a strictly statistical sense, not in semantic sense.

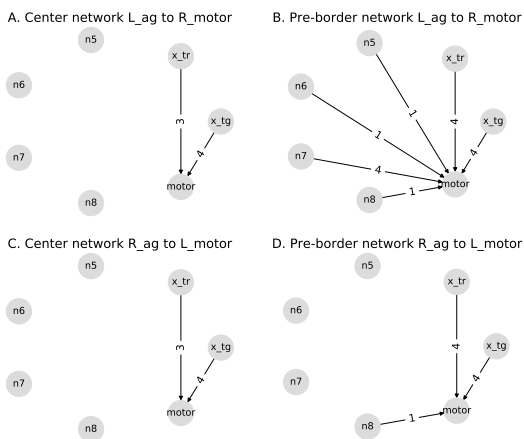


Figure 5: Effective networks across agents in 2 regions of interest.

also to infer effective networks in a given system that explain the observed time series. Such effective network inference (ENI) procedure typically involves measuring pairwise transfer entropy (TE) between all pairs of variables, pruning the connections above a certain p -value and correcting for multiple comparisons (Lizier and Rubinov, 2012).

In this study we employed the ENI procedure to investigate whether there are detectable relationships between two agents by conceptualizing a potential effective network as a combined agent-agent system. We included time-series data from a non-visual subset of the neurons (neurons 5-8) as potential TE sources and an output of the motor effectively controlled by the other agent as TE target. In order to minimize the possibility of finding spurious correlations we included experimental target and tracker positions as conditional variables. That is, given the constrained and regular nature of the task, it is conceivable that one could predict motor activity just based on the location of the target and tracker. Since we wanted to know whether there is additional predictive power from the neural activity of one of the agents, over and above task context, we conditioned on these positions. The time-series data from all trials was used as separate replications of the process. We used an open-source software (Wollstadt et al., 2017) to run ENI with an estimator most suited for non-linear continuous data (the Kraskov algorithm; Kraskov et al., 2004).

In order to get a handle on the functionality of any potential correlation, we separately tested the relationships for two different regions within the task environment: (1) pre-border region when the target was within 200 time steps before reaching the border and (2) center region when the target was around 200 time steps within the center of the environment. We reasoned that since predicting the co-actor’s action is more relevant to controlling the tracker before the border (when anticipatory strategy is required) than control-

ling it when it is around the center, statistical relationship in question should be more pronounced there.

The network plots of Figure 5 show ENI results for the two regions for both agents. In these plots, x_{tr} and x_{tg} represent the variables that capture the tracker and target positions; *motor* is the left or right motor activation and $n5$ to $n8$ are the activation of neurons 5 to 8 of the agents’ brains. The edges between the nodes are statistically significant links between the variables they represent while the numbers on those edges are the time lags for which (the strongest) significant link was found (all $p < 0.003$).

The plots suggest that indeed in the areas around the center of the environment, the only predictive links are between the positions of target and tracker and the motor activity of both agents. On the other hand, before the border region there is a statistical dependence between all considered neurons of the left agent (responsible for movement around the target) and the motor output of the right agent (in charge of movement between the borders) and a statistical dependence between the 8-th neuron of the right agent and the motor output of the left agent. This seems consistent with a minimal notion of ‘co-representation’ we adopted. However, whether the relationship is really functional in producing behavior and whether it is most useful to interpret the network activity as a *representation* requires further investigation

Discussion

In this paper we have presented a study of joint action between minimal cognitive agents that evolved to follow a target while controlling a common tracker in a complementary manner. We have found that the agents can emerge that solve this task successfully, that the solution relies on their timely coordination and that part of the solution resembles a more complex anticipatory strategy employed by human participants in the original KJ study. This result could mean that the task is actually not as representation-hungry as assumed in joint action literature, i.e., it can be solved without (co-)representations. Another possibility is that the task does indeed require representations and success indicates that MCAs are capable of representation-based activity. Which of the two is the case would need an agreement on a precise definition of representation, which is notoriously difficult to come by (Haselager et al., 2003).

We have adopted one possible definition and operationalization of “co-representation” as a state in one of the agents that carries (statistical) information about something in the other agent and contributes to the production of joint behavior. We based this definition on a common idea that representation is a state that stands in for certain external features (entities, properties, events) and plays some function in further cognitive processing and producing adaptive behavior (Dretske, 1988; Millikan, 1984). Our results delivered some support for the existence of such a state in the evolved agents. However, one could argue that our definition is both

too liberal – in that it risks proliferation of representations beyond cognitive systems (Hutto and Myin, 2013) – and too simplistic in not capturing what is actually intended by the term “co-representation”.

Specifically, what many of the appeals to representations in social behavior are supposed to capture is that good performance requires that one person has an inner model of their partner, uses it to predict what they will do, how, when and where (Sebanz and Knoblich, 2009b) and, thereby, adjust one’s own action to this and the joint outcome prospectively rather than after the fact. The content of this model can vary in abstraction, from simulations of the co-actor’s movements to their intentions (Rizzolatti and Sinigaglia, 2010) or the way the joint task is divided (Vesper et al., 2010). What is typically common to such varied possible representations is that they instantiate functional similarity between two individuals and it is in virtue of this similarity that coordinated joint action can ensue. That is, it is not just that agents have certain internal states that are caused by and/or correlate with, say, a particular characteristic of the co-actor. Rather, it is that agents possess complex internal states that instantiate some structural isomorphism between them, which is then used to understand, predict and coordinate behavior.

What remains unclear at present is if MCAs are also capable of this type of structural representation and how its presence and function could be established. Presumably, it would require some way to measure the similarity between the evolved internal structure and what it represents and a way to determine whether the similarity itself plays a functional role in the cognitive process. To our knowledge, no such work has been carried out yet – at least with respect to MCAs.⁷ However, it is a promising avenue to clarify the cognitive requirements of different social behaviors and the meaning of constructs used in their explanations.

Apart from this large conceptual question, there is a number of smaller but still relevant open issues that need to be addressed in future research. First, the difficulty with evolving successful solution with button-based velocity control makes our findings vulnerable to a common complaint from the opponents of the EEC approach that it is unable to deal with discrete behaviors. In the KJ task the fact that participants are required to produce discrete complementary actions is seen as a crucial component of its representation-hungriness. By switching to direct velocity control we are both not facing the challenge heads-on and allowing for solutions to evolve along the lines not conceived by the original study. This is interesting on its own but makes our results less comparable to the study we aimed at replicating.

Second, our post-hoc playback analysis suggests that the agents are not really acting jointly in this “joint action” setting if jointness or, more generally, sociality is understood as

⁷There is related work on categorization in neural networks (Laakso and Cottrell, 2000) that could be taken as an inspiration.

something more than merely performing the same task. That is, at least from an EEC perspective (Froese et al., 2011) interacting with another social agent is different than interacting with, say, an inanimate tool because another agent is herself a locus of adaptive behavior, goal-directedness and normativity, which have to be reckoned with in any joint endeavor. This might well require planning and prediction but also sensitivity to the ongoing social dynamics. The fact that our evolved agents behave largely the same with a live vs playback partner suggests that the task does not actually require co-regulation and can be accomplished individually while acting in parallel. However, this brings up the question whether the original KJ study involved a genuine joint action as well, highlighting the potential for artificial simulations to inform further psychological experiments. A follow-up study could be designed with a heavier demand for co-regulation, for instance, by requiring agents (or human participants) to coordinate with different individuals or by introducing other sources of uncertainty such as sensorimotor noise or delay (Rohde and Di Paolo, 2007).

Finally, we have employed the information dynamics approach to a study of a very constrained part of the full extended system operational in this task, i.e. the agents’ brains and actions. Future work should employ it to investigate the contributions of other parts of this system – the agents’ anatomy and their environment – to deliver a deeper understanding of how the solution to the task is produced and to situate putative co-representations’ role in this process. It might turn out that once such an analysis is complete, co-representations are merely an emergent pattern of statistical dependencies without an actual explanatory import.

Acknowledgements

This work was supported by the Netherlands Organization for Scientific Research (NWO), research project PGW-12-32. We also want to acknowledge the contribution of Jelle Zuidema, George Kachergis, Dimitri Coelho Mollo and Tom Froese at various stages of this project.

References

- Auvray, M., Lenay, C., and Stewart, J. (2009). Perceptual interactions in a minimalist virtual environment. *New Ideas in Psychology*, 27(1):32–47.
- Beer, R. D. (1996). Toward the evolution of dynamical neural networks for minimally cognitive behavior. In Maes, P., Mataric, M., Meyer, J. A., Pollack, J., and Wilson, S., editors, *From animals to animats*, volume 4, pages 421–429, Cambridge, MA. MIT Press.
- Beer, R. D. (2003). The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, 11(4):209–243.
- Beer, R. D. (2008). The dynamics of brain-body-environment systems: A status report. In Calvo, P. and Gomila, T., editors, *Handbook of cognitive science: An embodied approach.*, pages 99–120. Elsevier, Oxford.

- Bossomaier, T., Barnett, L., Harré, M., and Lizier, J. T. (2016). *An Introduction to Transfer Entropy: Information Flow in Complex Systems*. Springer, Berlin.
- Clark, A. and Toribio, J. (1994). Doing without representing? *Synthese*, 101(3):401–431.
- De Jaegher, H. and Di Paolo, E. (2007). Participatory sense-making. *Phenomenology and the Cognitive Sciences*, 6(4):485–507.
- Di Paolo, E. A., Rohde, M., and Iizuka, H. (2008). Sensitivity to social contingency or stability of interaction? modelling the dynamics of perceptual crossing. *New Ideas in Psychology*, 26(2):278–294.
- Dolk, T., Prinz, W., Obhi, S. S., and Cross, E. S. (2016). What it takes to share a task: sharing versus shaping task representations. *Shared Representations: Sensorimotor Foundations of Social Life*, page 1.
- Dretske, F. (1988). *Explaining Behavior*. Bradford, MIT, Cambridge.
- Eskenazi, T., Van Der Wel, R., and Sebanz, N. (2012). Mechanisms of skilled joint action performance. *Skill Acquisition in Sport: Research, Theory and Practice*, page 229.
- Froese, T. and Di Paolo, E. (2008). Stability of coordination requires mutuality of interaction in a model of embodied agents. In Asada, M., Hallam, J. C. T., Meyer, J.-A., and Tani, J., editors, *From Animats to Animals 10, The 10th International Conference on the Simulation of Adaptive Behavior*, pages 52–61, Berlin. Springer-Verlag.
- Froese, T. and Di Paolo, E. (2010). Modelling social interaction as perceptual crossing: An investigation into the dynamics of the interaction process. *Connection Science*, 22(1):43–68.
- Froese, T., Paolo, Di, and Ezequiel, A. (2011). The enactive approach: Theoretical sketches from cell to society. *Pragmatics & Cognition*, 19(1):1–36.
- Harvey, I., Di Paolo, E., Wood, R., Quinn, M., and Tuci, E. (2005). Evolutionary robotics: a new scientific tool for studying cognition. *Artificial Life*, 11(1-2):79–98.
- Haselager, P., de Groot, A., and van Rappard, H. (2003). Representationalism vs. anti-representationalism: A debate for the sake of appearance. *Philosophical Psychology*, 16(1):5–24.
- Hutto, D. D. and Myin, E. (2013). *Radicalizing enactivism: Basic minds without content*. MIT Press, Cambridge.
- Knoblich, G., Butterfill, S., and Sebanz, N. (2011). Psychological research on joint action. In *Advances in Research and Theory*, volume 54 of *Psychology of Learning and Motivation*, pages 59–101. Elsevier.
- Knoblich, G. and Jordan, J. S. (2003). Action coordination in groups and individuals: learning anticipatory control. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 29(5):1006–1016.
- Knoblich, G. and Sebanz, N. (2006). The social nature of perception and action. *Current Directions in Psychological Science*, 15:99–104.
- Kraskov, A., Stögbauer, H., and Grassberger, P. (2004). Estimating mutual information. *Physical Review E, Statistical, Nonlinear, and Soft Matter Physics*, 69(6 Pt 2):066138.
- Laakso, A. and Cottrell, G. (2000). Content and cluster analysis: Assessing representational similarity in neural systems. *Philosophical Psychology*, 13(1):47–76.
- Lizier, J. and Rubinov, M. (2012). Multivariate construction of effective computational networks from observational data.
- Millikan, R. (1984). *Language, Thought and Other Biological Categories*. MIT Press, Cambridge.
- Mirolli, M. (2012). Representations in dynamical embodied agents: re-analyzing a minimally cognitive model agent. *Cognitive Science*, 36(5):870–895.
- Murray, L. and Trevarthen, C. (1985). Emotional regulations of interactions between two-month-olds and their mothers. In Field, T. M. and Fox, N. A., editors, *Social perception in infants*, pages 177–197. Ablex Publishing, Norwood.
- Richardson, M. J., Marsh, K. L., Isenhower, R. W., Goodman, J. R. L., and Schmidt, R. C. (2007). Rocking together: dynamics of intentional and unintentional interpersonal coordination. *Human Movement Science*, 26(6):867–891.
- Rizzolatti, G. and Sinigaglia, C. (2010). The functional role of the parieto-frontal mirror circuit: interpretations and misinterpretations. *Nature Reviews: Neuroscience*, 11(4):264–274.
- Rohde, M. and Di Paolo, E. (2007). Adaptation to sensory delays. In Almeida e Costa, F., Rocha, L. M., Costa, E., Harvey, I., and Coutinho, A., editors, *Advances in Artificial Life*, pages 193–202, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Sebanz, N., Bekkering, H., and Knoblich, G. (2006). Joint action: Bodies and minds moving together. *Trends in Cognitive Sciences*, 10(2):70–76.
- Sebanz, N. and Knoblich, G. (2009a). Jumping on the ecological bandwagon? mind the gap! *European Journal of Social Psychology*, 39(7):1230–1233.
- Sebanz, N. and Knoblich, G. (2009b). Prediction in joint action: what, when, and where. *Topics in Cognitive Science*, 1(2):353–367.
- Sebanz, N., Knoblich, G., and Prinz, W. (2005). How two share a task: corepresenting stimulus-response mappings. *Journal of Experimental Psychology. Human Perception and Performance*, 31(6):1234–1246.
- Vesper, C., Butterfill, S., Knoblich, G., and Sebanz, N. (2010). A minimal architecture for joint action. *Neural Networks*, 23(8-9):998–1003.
- Williams, P. L. and Beer, R. D. (2010). Information dynamics of evolved agents. In Doncieux, S., Girard, B., Guillot, A., Hallam, J., Meyer, J.-A., and Mouret, J.-B., editors, *From animats to animats II, Lecture notes in computer science*, pages 38–49, Berlin. Springer.
- Wollstadt, P., Lizier, J. T., Vicente, R., Finn, C., Martinez-Zarzuela, M., Lindner, M., Martinez-Mediano, P., and Wibral, M. (2017). IDTxl - the information dynamics toolkit xl: a python package for the efficient analysis of multivariate information dynamics in networks.