From embodied interaction to compositional referential communication: A minimal agent-based model without dedicated communication channels

Jorge I. Campos[1,2] and Tom Froese[2,3]

[1]Faculty of Higher Education Aragon, National Autonomous University of Mexico, Mexico City, Mexico
[2]Institute of Applied Mathematics and Systems Research, National Autonomous University of Mexico,Mexico City, Mexico
[3]Center for the Complex Sciences, National Autonomous University of Mexico,Mexico City, Mexico
jorgeivan.campos.bravo@gmail.com

Abstract

Referential communication is a ”representation-hungry” behavior, and the bee waggle dance is a classical example of referential communication in nature. We used an evolutionary robotics approach to create a simulation model of a minimalist example of this situation. Two structurally identical agents engage in embodied interaction such that one of them can find a distant target in 2D space that only the other could perceive. This is a challenging task: during their interaction the agents must disambiguate translational and communicative movements, allocate distinct behavioral roles (sender versus receiver), and switch behaviors from communicative to target seeking behavior. We found an evolutionary convention with compositionality akin to the waggle dance, correlating duration and angle of interaction with distance and angle to target, respectively. We propose that this behavior is more appropriately described as interactive mindshaping, rather than as the transfer of informational content.

Introduction

Communication in nature shows that there are certain crucial processes that require referential communication in order to keep the species alive. The best example in nature is the waggle dance of the bee (Figure 1), where an explorer bee goes out of the honey comb in order to search for a source of food. Once the explorer bee has found a good field the explorer bee comes back to the honey comb and the dance starts (Crist, 2004; Dornhaus and Chittka, 2004; Dyer, 2002; Seyfarth and Cheney, 2003).

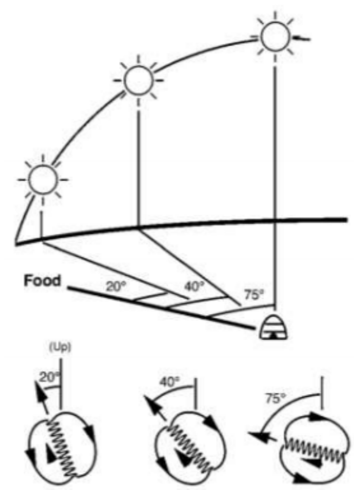

Figure 1: Waggle dance of the bee. The communication of the bee waggle dance depends on the angle at which the buzzing in the center of the figure of ”eight” with the sun is performed. In addition the duration of those buzzings correlate with the distance to where the food is located. Figure taken from (Dyer, 2002)

Through the waggle dance, the bee communicates the location of the source to the foraging bees. In the dance, the behavior of the bee corresponds with elements present in the environment that helps the foraging bees to go to the source of food. In animals, simple associations between world entities and signals are mostly innate, or they can be explained by mere mechanisms of rote learning and conditional learning (Cangelosi, 2001). In the case of the waggle dance, the angle with respect to the sun and the center of the dance is assumed to be a genetically determined association of a bee's movements with a certain state of affairs of the world.

In our previous model (Campos and Froese, 2017) we found an analogous communicative behavior performed by artificial agents that were evolved to solve a minimal referential communication task in a 1D environment. When we analyzed the performance of the agents we found that the best explanation involves an appeal to both agents as an extended system formed via their embodied interaction process: through their interaction the agents shaped the dynamical basis of each other's behavior into giving rise to an adequate solution to the task. This solution is better described as a non-representational form of ”mind shaping” (Zawidzki, 2013), rather than a traditional Theory of Mind mechanism. However, it could be argued that the task was not sufficiently ”representation-hungry,” because its solution did not depend on the principle of compositionality. Here we therefore extended the spatial dimension of the model to a 2D environment and analyzed the resulting behaviors and internal dynamics. We expect that some of the embodied strategies for achieving referential communication will take advantage of the possibility of separately responding to these two dimensions, and to do so in a combinatorial manner.

Compositionality in communication

In communication compositionality can be defined as the expression of a complex content in terms of a function of the expressions of its parts and their mode of composition.

In other words, the meaning of the whole is determined by the meanings of the parts and the mode of composition. Since the receiver knows the meanings of the simple parts, knows the semantic significance of a finite number of syntactic modes of composition and can parse the expression (i. e. recognize how it is built up out of simple parts) the interpreter can work out the meaning of the whole (Pagin, 2003). In this case we can attempt to explain compositional communication using both agents: how the receiver finds the right interpretation (meaning) and the sender finds an appropriate linguistic item (signal) where the receiver solves a task of interpretation and the sender solves a task of expression.

In our previous model (Campos and Froese, 2017), there is no compositionality present in the communicative behavior: the distance toward the target is simply proportional to the contact time between the agents, and this is sufficient to make possible their communication.

In order to get an intuitive idea about what compositional referential communication consists in, we can consider an alternative modeling approach. The Iterated Learning Approach (Kirby, 2000; Smith et al., 2003) is a framework to study the emergence of the language with compositionality in a sequential model that requires a couple of agents: a mature user of the language (speaker/mentor) and a new user of the language (hearer/learner), which will become the mature agent after learning to teach a new novice agent. During a series of mentor/learner iterations applied to a set of initially arbitrary meaning-signal pairs, a compositional language emerges mainly due to the pressure of keeping the language easy to learn. In other words, structured parts of the signal pick out structured parts of the meaning space.

In contrast, in our minimal referential communication model, there is no predefined meaning-signal pair set and no predefined communication channels. Instead, the "meaning" is the target's location in the environment and the "signal" is the movement of the agents in space. There is no sequential learning process and the pressure of finding a good referential communication system derives directly from the artificial evolution process. Nevertheless, by increasing the environment to two dimensions, we expect the agents to evolve a compositional behavior that involves each dimension of the space (i.e. distance and height) in order to solve the referential communication task. We can then analyze the dynamical basis of this behavior to see whether it involves internal compositional representations, or whether an alternative subpersonal explanation exists.

Interactive approach to referential communication

From the perspective of the dynamical approach to cognition, the components of an agent tend to be in continuous mutual interaction, and none of the components can be removed without thereby also modifying the behavior of the other components. In the case of a system consisting of two agents, this perspective implies that their social interaction process should be better conceived of as a collective property of a brain-body-environment-body-brain system as a whole (Froese et al., 2014). We applied this idea to study the phenomenon of referential communication. In addition, we required that the roles of sender and receiver are initially ambiguous and must be negotiated as part of a continuous flow of nonlinear interaction. Moreover, in contrast to a traditional broadcasting approach to communication, it cannot be assumed on an a priori basis that the agent that turned out to adopt the role of the "receiver" will not play a role in the successful realization of communication. We call this the interactive approach to referential communication.

The interactive approach has the virtue of being a broader perspective that includes the broadcasting approach as a special case, in which the dynamics of the sender are endogenously generated and sufficiently decoupled from its environment. We employ agent-based modeling to help us to develop this alternative theoretical framework.

Previous modeling work

In the field of artificial life there is a long tradition of modeling the evolution of communication and the compositionality of language (Cagnelosi and Parisi, 1998; Cangelosi, 2001; MacLennan and Burghardt, 1993; Williams et al., 2008; Manicka, 2012; Nolfi, 2005, 2013). In broad terms it can be said that communication occurs when the behavior of one agent modifies the future behavior of another agent in a task-relevant manner. Several researchers have analyzed the phenomenon as a special example of coordinated behavior between individuals Di Paolo (1997); Di Paolo (2000). But there is still a need to explore how the compositionality of communication can be explained from the dynamical systems approach, and we propose to do so by developing an interactive approach of the referential communication.

Our model is inspired by the work of Williams et al. (2008), who applied the minimal cognition modeling paradigm to the study of referential communication. In their model, they placed two embodied agents, a sender and a receiver, on a 1-D circle. The agents can perceive their own location on the circle and the presence of each other, but only the sender can also perceive the location of the target that the receiver must reach. Following the tradition of research of communication as a form of behavioral coordination (Ackley and Littman, 1994; Di Paolo, 1997; Di Paolo, 2000; Maturana and Varela, 1987), their model does not include dedicated communication channels and so the agents have to learn to distinguish communicative movements from translational movements (Quinn, 2001). The task is for the receiver to end up as close as possible to the target. They optimized the behavior of the two agents, each controlled by a structurally distinct continuous-time recurrent neural network (CTRNN) (Beer, 1995), using an evolutionary robotics approach (Beer et al., 1996; Harvey et al., 2005). They tested three types of conditions: one involving unconstrained interaction between the sender and the receiver, and two conditions involving constrained interaction. In brief, they found that the easiest solution to the task in the unconstrained condition was for both the sender and the receiver to move to the target together. In this way the sender could use the location of its own body to indicate to the receiver where the location of the target is. However, this solution does not seem to consist in referential communication. To force the evolutionary algorithm to find a solution that involves a form of referential communication they introduced a spatial constraint, such that the receivers position was limited to a subsection of the circle away from the target. They found that it was possible to evolve agents that coordinated their behavior in such a way that the receiver was able to locate a randomly chosen target of a small set of targets (discrete condition) as well as of a range of targets (continuous condition).

But why does artificial evolution tend to prefer solutions to multi-agent coordination problems that are based on joint embodied action rather than on referential communication? This open question makes it even more interesting to consider why bees evolved a complex waggle dance, given that it would be much simpler for the sender bee to directly guide the receiver bees to the target location by simply flying back to where it came from. In other words, it is likely that there is an important constraint that prevents this solution from being feasible. We will return to this open question in the discussion.

In our previous 1D version of this model (Campos and Froese, 2017), we showed that we can reduce the number of nodes in the agents' controllers (in that case to 3 nodes), and use structurally identical artificial neural networks for the two distinct roles (Izquierdo and Buhrmann, 2008), and still evolve the agents to make use of their movements in space as a means for referential communication about target locations of varying distance. Nevertheless, we still had to maintain the constraint of preventing the sender to move directly to the target location; otherwise referential communication would not emerge.

Also, we found an analogous strategy for the agents to solve the task, namely that the distance to the target is correlated with the contact time between the agents, which results from mutually coordinated movements. Therefore, this solution requires both agents to be active in the communicative process for the task to be achieved successfully. However, given that the solution consisted in the co-regulation of one continuous parameter, it did not exhibit compositionality.

The current contribution

The current contribution increases the complexity of our original model in a way that is intended to facilitate the emergence of compositional referential communication. We used the same modeling approach as before, but we expanded the environment from a 1D to a 2D space. We were interested in exploring the possible space of solutions to this 2D-version of the task. Accordingly, we did not pre-specify that the agents had to evolve to solve the task in a compositional manner, although we expected that compositional embodied behavior would be the most effective solution to this task. In other words, we evolved the embodied agents to solve a task that required referential communication between a sender and a receiver and ideally in a combinatorial manner, but without dedicated communication channels, without dedicated roles, and without dedicated signal components.

This kind of scenario is very far removed from traditional studies of communication based on information theory, which typically already assume a well-defined sender, receiver, channel, and symbol system. Our model can therefore serve as an inspiration for thinking about the origins of compositional referential communication in nature.

Methods

The model

Following our approach in Campos and Froese (2017), we created a model in which a pair of embodied agents needs to find a way for the agent named receiver to move through the environment to a target position, but only the agent termed sender knows the exact position of the target. The behavior of the agents was evolved using the methodology of evolutionary robotics (Harvey et al., 2005), using the structure of their artificial neural networks (a standard CTRNN) as their genome. Each artificial neural consisted of 6 nodes, and in line with related work on minimal cognition we used only one structural copy of the network for both agents (Izquierdo and Buhrmann, 2008). We now provide more details of the model.

The task

The task that both agents must fulfill is relatively straightforward: the so-called receiver must arrive at a target area, which will be one of four possibilities in the environment. The environment is a 2-unit-side square with center at (0,0), with a centered interaction zone of 0.6 units side where the sender is constrained to move. The only agent that can sense the position of a target is the so-called sender. The so-called receiver does not sense the target, and the agents only sense each other agent while both remain inside the interaction zone. Once the receiver leaves the interaction zone in order to find the target, there no longer is any possibility for interaction between the agents. Moreover, the receiver cannot sense the location of the target, and so must find the location of the target by relying solely on the history of embodied interaction between the agents (see Figure 2).

Each pair of agents was tested in four separate trials using four distinct targets in the environment. These targets were selected by arbitrarily varying the distance and angle to the interaction zone within fixed limits.

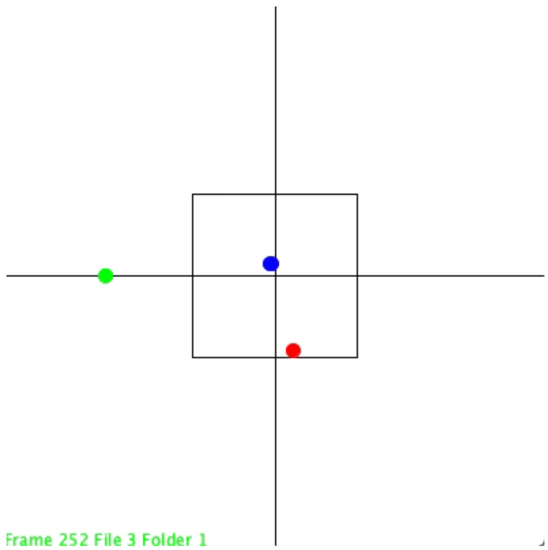

Figure 2: Simulation of the agents in the environment. The blue circle represents the sender while the red circle represents the receiver. The green circle is the target, which is in a random position along one of the two axes. The square in the middle represents the interaction zone. After interacting in the interaction zone, the red agent has to leave the zone and try to locate the target based only on its history of interaction with the other agent.

The agents

The sensorimotor system we used has six sensors and two motors, which are the inputs and outputs of an agent's six-node CTRNN (see Figure 3).

Three of an agent's sensors serve to locate the other agent's position with respect to the self's location: a distance sensor providing a continuous input that varies linearly with the distance to the other agent, and two sensors providing the angle to the other agent, i.e. a sensor for $Sin(\theta)$ and another for $Cos(\theta)$). If the receiver moves outside the interaction zone then these three sensors are turned off in both agents. So the range of the sensor is [0, 1], where 0 is the maximum distance in the interaction zone (and also any further distances, which imply that the receiver is outside the zone), and 1 indicates that both agents are exactly on top of each other.

The other three of an agent's sensors serve to locate the target's position with respect to the self's location: a distance sensor providing a continuous input that varies linearly with the distance to the target, and two sensors providing the angle to the target, i.e. a sensor for $Sin(\theta)$ and another for $Cos(\theta)$). These three sensors are always turned on in the sender and always turned off (i.e. set to -1) in the receiver. Given that the sender is constrained to the interaction zone, the distance sensor is restricted to range [0.2, 0.8].

Each agent has a body in the shape of a circle of radius 0.05 that has two simulated wheels, one on each side, giving the agent a facing property. The velocity of a wheel is controlled by the output of a dedicated neuron, and these motor neurons were arbitrarily chosen to be the nodes also connected to the distance sensors. The agents control their overall velocity using differential steering control, which means that for an agent to move forward at a certain speed it is necessary that both wheels have the same velocity. Collisions are not modeled.

Artificial neural network

The behavior of each agent is modulated by a continuous-time neural network (CTRNN) (Beer, 1995) (Figure 3), whose equation is the following:

$$\tau_i \cdots_i = -s_i + \sum_{j=i}^{N}(\omega_{ij}\sigma(s_j + \theta_i)) + G_i I_i \qquad i = 1, ..., N \quad (1)$$

where is s the state of each neuron, τ is the time constant, ω_{ij} is the weight of the $j-th$ neuron to the $i-th$ neuron, θ is the bias, $\sigma(x) = \frac{1}{(1+e^{-x})}$ is the standard logistic activation function, G is the gain constant of the input of the neuron and I is the input of each neuron.

The CTRNNs of the two agents are structurally identical. The integration step size was set to 0.1.

Evolution of the agents

The parameters of the neural network were encoded using floating-point numbers and then were optimizing using an artificial evolution procedure. We perform 50 separate evolutionary runs and in each we used 100 individuals that were evolved for 1,000 generations. In each generation, the individuals were evaluated in the environment as follows:

At the start of a trial the agents are placed randomly inside the interaction zone within a circle of radius 0.3, and always facing rightward (0 degree). They are then allowed to behave for 300 units of time. At the end of the trial their fitness score for that trial is calculated with the following equation: $fitness = 1 - finalDistanceToTarget$. Every time we picked a target, we calculated the mean fitness over the population of pairs. Once the fitness increased by 20% from this initial value, we switch the target for a different randomly selected target location. This was done to allow local task optimization but without losing generality.

Each pair of agents is evaluated for 10 trials. Overall fitness of the pair is calculated as the inverse weighted average score of all 10 trials.

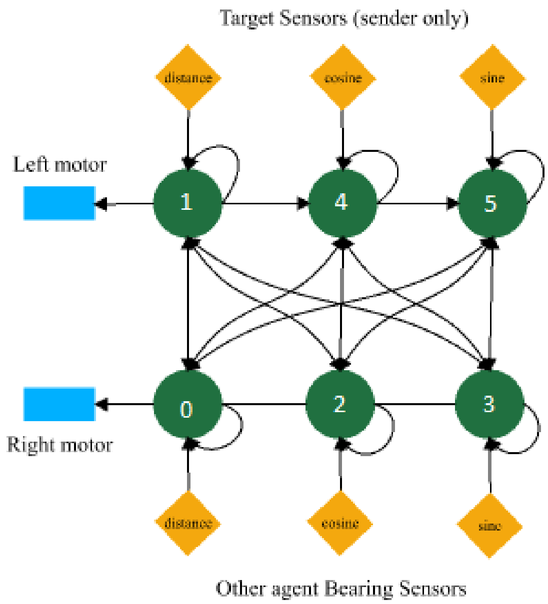

Figure 3: Illustration of artificial neural network structure. Green circles represent each node of the six-node network. Orange diamonds represent the six sensors, and the blue rectangles show the two motors.

The evolutionary algorithm employed the hill-climbing method: each solution generates a new solution by applying Gaussian mutation with variance of 0.2 to each parameter, and this descendant replaces the parent solution if its performance is better.

Results

Once the 50 independent evolutionary runs finished, we saved the best pair of agents of each run. With these resulting 50 pairs of agents we performed several tests to see the performance of the agents for different initial conditions.

Behavior analysis

By looking at the 50 best agents we noticed that this is a challenging task. Most of the runs failed to give rise to strategies involving referential communication. Instead the agents relied on suboptimal strategies, such as getting to a certain target spawning position. In some cases the agents apparently got too entrained in the interaction and the receiver never left to find the target in the first place. In other cases the interaction between the agents resulted in complex behavior of the receiver, but without getting close to the target.

Just one evolutionary run out of the 50 runs accomplished the task successfully: the two agents interacted at the beginning of the trial and then the receiver moved to where the target is. The rest of our analysis focuses on this successful pair of agents.

To confirm that their strategy relied on referential communication rather than on direct influence on each other's behavior via the agent sensors, we ran a test trial where we removed the sender as soon as the receiver first goes outside the interaction zone. Using this setup, most trials still achieved the task. A few failed, but this was because we broke the interaction too soon for the receiver to fulfill the task (i.e. the receiver had returned into the interaction zone but the sender was no longer there).

To improve the success of these trials we introduced this pair of agents in a new optimization process. This time we generated a population of 300 individuals that begin with the same CTRNN, and we change the task by adding the constraint that once the sender has left the interaction zone he can't enter again and interact with the sender. We let the population evolve through 300 generations.

At the end of this new evolutionary run, the best pair of agents found the target on all of the trials. We further tested these agents as described in the following subsections.

Overall performance

We performed 1,000 test trials with these best agents, placing random targets all over the environment. More specifically, at the start of each trial a target was randomly placed inside a space defined by x-range [-1, 1] and y-range [-1, 1], but not inside the interaction zone (ranges [-0.6, 0.6]). We got the following distribution of final distances between the receiver and the target (Figure 4).

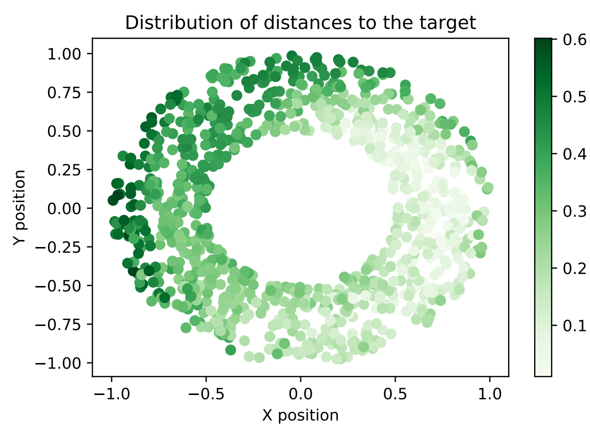

Figure 4: Distribution of 1,000 randomly placed test target locations. Shading represents distances between the receiver and the target at the end of a trial. Light green represents closer to the target and darker represents further away.

For the 1,000 trials we got an average distance to the target of 0.262094, where the most of the trials with the target at the right of the agent's interaction zone the distances range between 0.15 and 0.05. The worst case was located at $[-0.8858, -0.4051]$ with a distance of 0.6014 one of the right most target appearances. We can say that the agents generalized well the task for most of these targets, because the agents were evolved to achieve the task only for 4 different angle locations.

In general, the agents perform the task with the following behavior. First, they spend some time interacting inside the interaction zone. This typically involves spiraling in ever larger circles. Then, the receiver eventually is on a trajectory that takes it out from the interaction zone, and then it starts its movement trajectory to the target to end up as close as possible based on the history of interaction with the sender.

Strategy per target position

We are interested in finding the strategy that the agents perform to achieve the task.

First, we compared the final interaction distances between the sender and receiver, i.e. their distance at the moment when the receiver leaves the interaction zone, with the distance from the center of the space to the target location. We wanted to see if there is a correlation between their distance and target distance. And indeed, we find such a correlation, at least for some areas of the space (see Figure 5).

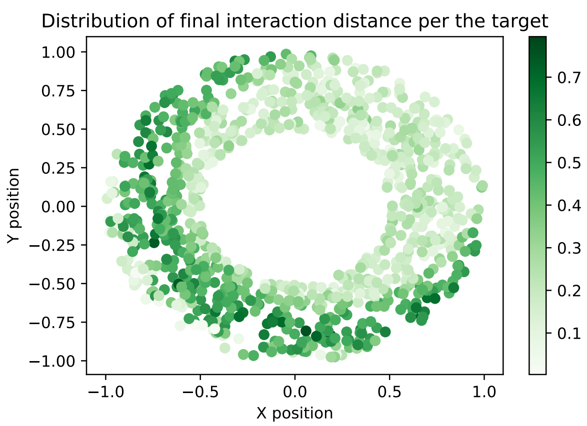

Figure 5: Distribution of 1,000 randomly placed test target locations. Shading represents the distance between the agents before their interaction breaks off, i.e. before the receiver leaves the zone of interaction. Light corresponds to a shorter and dark corresponds to a longer distance. There is a tendency for targets that are further away from the interaction zone to be associated with greater distances between the agents (darker dots).

Next, we analyzed the angles exhibited by the agents when their interaction ends (Figure 6). What we found is that there is also a relation between the angle between the agents at the end of their interaction and the angle from the center of the interaction zone to the target position. This relation complements the previous relation.

The agents found a strategy, in general, using two components: their distance to each other and their angle with respect to each other. This is the compositionality we tried to find, similar to the compositionality of the bee's dance, albeit in our model there was not such a clear demarcation of the components of the referential communication given that everything is embedded in a parallel manner in a continuously spiraling mutual interaction.

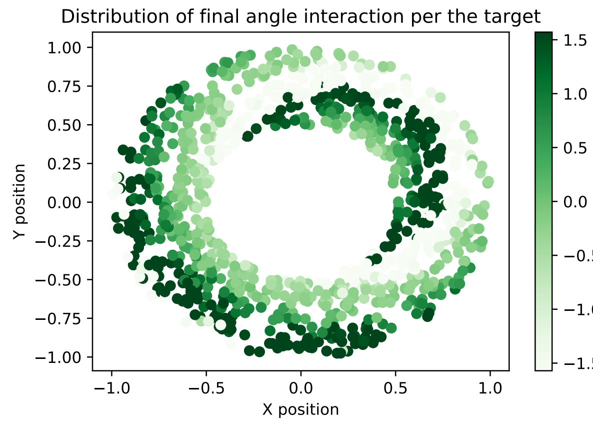

Figure 6: Distribution of 1,000 randomly placed test target locations. Shading represents the difference between the angle between the agents at their final point of contact and the angle from the center of the interaction zone toward the target's location. The shades of green represent the angle difference going from $-\frac{\pi}{2}$ to $\frac{\pi}{2}$, which is the possible range of angle differences for the quadrant in which the target is located.

Attractor analysis of the artificial neural network

In order to get an intuitive grasp of the state space configuration of the CTRNN of the best pair of agents, we performed 10 tests with different initial conditions with the isolated neural network. We found that the trajectories in CTRNN state space converged toward a single attractor located at coordinates [-7.10114 , 4.76134 , -6.69269 , 10.3879 , -8.68615 , 0.199544]. The fact that the agent's dynamical system contains a single attractor means that the complexity of the behavioral solution of the task is strongly dependent on the interaction between the agents in the environment. It is a solution that is better conceived of as a property of the system as a whole than as a property of an agent's neural network activity. This is consistent with what we found in our previous 1D model.

In other words, it appears that increasing the complexity of the referential communication task has only increased the complexity of the collective system but not the internal complexity of the agents. This increase in interaction complexity without concomitant increase in internal complexity is consistent with findings reported from another multi-agent system modeling study that investigated this relationship more systematically (Candadai et al., 2019).

In Figure 7 we show the states of three neurons of the isolated CTRNN starting from arbitrary initial states, where in this 3D-space all the states converge on the same attractor.

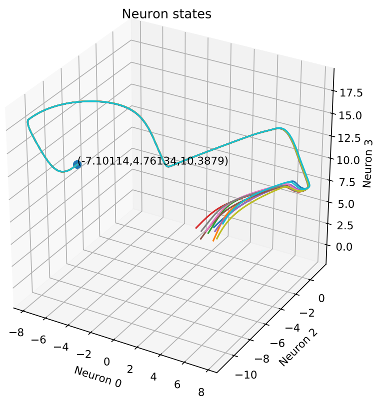

Figure 7: Neural states of 3 representative neurons of the agent's decoupled CTRNN during 10 trials with different initial conditions. The colored lines represent each trial and the point represents the attractor of the system

Discussion

This model illustrates the interactive approach to referential communication, where the agents tend to be in continuous mutual interaction. We chose a task that cannot be accomplished by the receiver alone, because the information about the target is only available by the sender. However, the sender also cannot behave as expected without the presence of a receiver that interacts in the right manner, giving rise to a mutual interaction process.

The artificial neural network of both the sender and receiver is structurally the same. That is why the agents must negotiate their respective roles during their embodied interaction before one of them attempts to go to the target. Once the interaction begins they must coordinate their behavior to break their interaction at the right moment. This means that, in accordance with the interactive approach, successful referential communication is co-created by both agents: their individual behaviors give rise to the interaction process but the interactive process shapes their individual behaviors.

As we can see in figures 5 and 6 we can analyze their mutual interaction process into separate components that together help the receiver to reach the target position in space. This can be seen as the compositionality of referential communication that was absent in our previous 1D version of the model.

Interestingly, the complex behavior of the agents does not require equally complex internal organization of their CTRNNs. As Figure 7 shows, the agents succeed with a CTRNN that in its decoupled mode only exhibits a single attractor. In other words, the complexity of the behavior is outsourced into the complexity of the interaction itself. This finding is in line with the enactive approach to social interaction, which holds that social cognition can be constituted by social interaction (De Jaegher and Froese, 2009), and that this can lead to performance that would otherwise be outside the agents' reach in practice and even in principle (Froese et al., 2013).

Future work

In order to be able to evolve agents that were able to perform the 2D version of the referential communication task, we had to increase the number of CTRNN nodes to six nodes from the three nodes we used for the 1D version. It is possible that this increased internal dimensionality is needed for the receiver to sufficiently cope after being decoupled from the sender, as was also found in related modeling work (Fernandez-Leon and Froese, 2010). This requires further analysis, but it seems that more nodes may facilitate the production of more robust internal dynamics for the agents.

The best-evolved strategy also needs to be analyzed in more detail. While we found that two components of the relationship between the agents at the end of their mutual interaction, namely their distance and angle, tended to correspond to the distance and angle to the target, this correspondence was perhaps not as consistent or uniform as would normally be expected from a genetically evolved referential communication strategy such as the bee's dance. In future work we would therefore like to apply other measures.

Another interesting direction for future research is to increase the number of agents in the model. For instance, if several senders had to compete with each other in the interaction zone (similar to the chaotic situation of multiple concurrent dances in a bee's hive), it would not be convenient for the sender to simply return to the target's location with other agents in tow. In other words, perhaps the externally imposed restriction of the sender to the interaction zone would no longer be necessary under such more realistic conditions.

Acknowledgements

This work was supported by CONACYT scholarship for J.I.C.

References

Ackley, D. H. and Littman, M. L. (1994). Altruism in the evolution of communication. In Brooks, R. A. and Maes, P., editors, *Artificial life IV: Proceedings of the Fourth International*

Workshop on the Synthesis and Simulation of Living Systems, pages 40–48, Cambridge, MA. MIT Press.

Beer, R. D. (1995). On the dynamics of small continuous-time recurrent neural networks. *Adaptive Behavior*, 3(4):469–509.

Beer, R. D. et al. (1996). Toward the evolution of dynamical neural networks for minimally cognitive behavior. In Maes, P., Mataric, M. J., Meyer, J.-A., Pollack, J., and Wilson, S. W., editors, *From animals to animats: Proceedings of the Fourth International Conference on Simulation of Adaptive behavior*, volume 4, pages 421–429, London,England. MIT Press.

Cagnelosi, A. and Parisi, D. (1998). The emergence of a 'language' in an evolving population of neural networks. *Connection Science*, 10(2):83–97.

Campos, J. I. and Froese, T. (2017). Referential communication as a collective property of a brain-body-environment-body-brain system: A minimal cognitive model. In *2017 IEEE Symposium Series on Computational Intelligence (SSCI)*, pages 863–870, Honolulu, Hawaii, USA. IEEE Press.

Candadai, M., Setzler, M., Izquierdo, E. J., and Froese, T. (2019). Embodied dyadic interaction increases complexity of neural dynamics: A minimal agent-based simulation model. *Frontiers in Psychology*, 10:540.

Cangelosi, A. (2001). Evolution of communication and language using signals, symbols, and words. *IEEE Transactions on Evolutionary Computation*, 5(2):93–101.

Crist, E. (2004). Can an insect speak?: The case of the honeybee dance language. *Social Studies of Science*, 34(1):7–43.

De Jaegher, H. and Froese, T. (2009). On the role of social interaction in individual agency. *Adaptive Behavior*, 17(5):444–460.

Di Paolo, E. A. (1997). Social coordination and spatial organization: Steps towards the evolution of communication. In Husbands, P. and Harvey, I., editors, *Proceedings of the Fourth European Conference on Artificial Life (ECAL'97)*, pages 464–473. MIT Press.

Di Paolo, E. A. (2000). Behavioral coordination, structural congruence and entrainment in a simulation of acoustically coupled agents. *Adaptive Behavior*, 8(1):27–48.

Dornhaus, A. and Chittka, L. (2004). Why do honey bees dance? *Behavioral Ecology and Sociobiology*, 55(4):395–401.

Dyer, F. C. (2002). The biology of the dance language. *Annual Review of Entomology*, 47(1):917–949.

Fernandez-Leon, J. A. and Froese, T. (2010). What is the relationship between behavioral robustness and distributed mechanisms of cognitive behavior? In *IEEE Congress on Evolutionary Computation*, pages 1–8. IEEE.

Froese, T., Gershenson, C., and Rosenblueth, D. A. (2013). The dynamically extended mind. In *2013 IEEE Congress on Evolutionary Computation*, pages 1419–1426. IEEE.

Froese, T., Iizuka, H., and Ikegami, T. (2014). Embodied social interaction constitutes social cognition in pairs of humans: a minimalist virtual reality experiment. *Scientific reports*, 4:3672.

Harvey, I., Paolo, E. D., Wood, R., Quinn, M., and Tuci, E. (2005). Evolutionary robotics: A new scientific tool for studying cognition. *Artificial life*, 11(1-2):79–98.

Izquierdo, E. and Buhrmann, T. (2008). Analysis of a dynamical recurrent neural network evolved for two qualitatively different tasks: Walking and chemotaxis. In Bullock, S., Noble, J., Watson, R., and Bedau, M. A., editors, *Artificial life XI: Proceedings of the eleventh international conference on the simulation and synthesis of living systems*, pages 257–264. MIT Press.

Kirby, S. (2000). Syntax without natural selection: How compositionality emerges from vocabulary in a population of learners. In Knight, C., Studdert-Kennedy, M., and Hurford, J., editors, *The Evolutionary Emergence of Language: Social Function and the Origins of Linguistic Form*, pages 303–323, Cambridge, MA. Cambridge University Press.

MacLennan, B. J. and Burghardt, G. M. (1993). Synthetic ethology and the evolution of cooperative communication. *Adaptive Behavior*, 2(2):161–188.

Manicka, S. (2012). Analysis of evolved agents performing referential communication. In Adami, C., Bryson, D. M., Ofria, C., and Pennock, R. T., editors, *ALIFE 2012: The Thirteenth International Conference on the Synthesis and Simulation of Living Systems*, pages 393–400. MIT Press.

Maturana, H. R. and Varela, F. J. (1987). *The tree of knowledge: The biological roots of human understanding.* New Science Library/Shambhala Publications, Boston, MA.

Nolfi, S. (2005). Emergence of communication in embodied agents: co-adapting communicative and non-communicative behaviours. *Connection Science*, 17(3-4):231–248.

Nolfi, S. (2013). Emergence of communication and language in evolving robots. In Lefebvre, C., Comrie, B., and Cohen, H., editors, *New perspectives on the origins of language*, pages 533–554. John Benjamins, Amsterdam,Philadelphia.

Pagin, P. (2003). Communication and strong compositionality. *Journal of Philosophical Logic*, 32(3):287–322.

Quinn, M. (2001). Evolving communication without dedicated communication channels. In Kelemen, J. and Sosk, P., editors, *Advances in Artificial Life: 6th European Conference on Artificial Life, ECAL 2001*, pages 357–366, Berlin,Heidelberg. Springer.

Seyfarth, R. M. and Cheney, D. L. (2003). Signalers and receivers in animal communication. *Annual Review of Psychology*, 54(1):145–173.

Smith, K., Kirby, S., and Brighton, H. (2003). Iterated learning: A framework for the emergence of language. *Artificial Life*, 9(4):371–386.

Williams, P. L., Beer, R. D., and Gasser, M. (2008). Evolving referential communication in embodied dynamical agents. In Seth Bullock, Jason Noble, R. W. and Bedau, M. A., editors, *Artificial life XI: Proceedings of the eleventh international conference on the simulation and synthesis of living systems*, pages 702–709. MIT Press.

Zawidzki, T. W. (2013). *Mindshaping: A New Framework for Understanding Human Social Cognition*. MIT Press, Cabridge, MA.