



HOKKAIDO UNIVERSITY

Title	Evolution of communication through differentiation of communicative and goal-directed behaviors
Author(s)	Shibuya, Naohiro; Iizuka, Hiroyuki; Yamamoto, Masahito
Citation	Artificial life and robotics, 23(2), 225-234 https://doi.org/10.1007/s10015-017-0417-7
Issue Date	2018-06
Doc URL	https://hdl.handle.net/2115/74522
Rights	The final publication is available at Springer via https://doi.org/10.1007/s10015-017-0417-7
Type	journal article
File Information	main.pdf



Evolution of communication through differentiation of communicative and goal-directed behaviors

Naohiro Shibuya · Hiroyuki Iizuka · Masahito Yamamoto

Abstract Many studies have focused on the evolution of communication in artificial life, often presupposing dedicated communication channels such as voices or pheromones. However, it is plausible that only behavior existed when communication began and that some behaviors gradually evolved as a mean of communication. A study by Quinn has shown that communication can evolve purely through the behavior for two agents in a simulated environment. We extend this study to a more complex communication using three agents. In the simulation, the agents are equipped with two wheels and eight IR sensors and are controlled by identical neural networks. This network evolved by applying a genetic algorithm to a single population. After evolution, the agents' behaviors had differentiated into communicative and goal-directed types. We analyze the differences between these two types of behavior quantitatively, and show that the evolved behaviors have communication properties such as mutuality, information flow, and invariance of the motion gain.

Keywords Evolution of communication, Neural network, Genetic algorithm, Artificial life

1 Introduction

Evolutionary robotics is a powerful tool for reproducing phenomena in order to understand how they have been shaped over the long history of life on Earth[1–3]. The evolution

of communication is an important topic in the artificial life community, and communication models can be roughly classified into two groups, based on different modeling assumptions and ways of reproducing the evolution of communication. One group consists of models where signals or symbols are exchanged among agents. How to react to (or interpret) these symbols is not initially determined, i.e., it is random. As the better performing agents are selected, interdependencies among agents develop, and they start using these symbols as communication media. This approach has been tried not only in computer simulations[4–8] but also in human experiments[9, 10], studying how initially random signals in a predefined communication channel are self-organized into meaningful patterns. The channel is only for communication purposes and is separate from the agent's (or human's) usual movements. However, it is hard to imagine that such communication channels have existed since life began.

Other models take what is known as the embodiment approach and presume that only behavior existed when communication began, with some of these behaviors gradually evolving into a form of communication[11–13]. This is a plausible way of modeling the evolution of communication, but the distinction between communicative behaviors and others is not as clear as with communication through a predefined channel. The distinction is ambiguous because, without a predefined communication channel, what we can observe in these models, as an external observer, are just agents moving around, i.e., their behavior.

Following this idea, Quinn performed a simulation where communication emerged from behavior using two agents, evolved within a single population by a genetic algorithm [14]. The agents differentiated themselves into a leader and a follower through communicative behaviors in order to achieve a task. However, because of their ambiguity, the evolved communicative behaviors were shown to be qualitatively different from goal-directed ones. It was not clear whether communication really occurred through these interactions be-

N. Shibuya · H. Iizuka · M. Yamamoto
Graduate School of Information Science and Technology,
Hokkaido University, Kita 14, Nishi 9, Kita-ku, Sapporo,
Hokkaido, 060-0814, Japan
E-mail: shibuya@complex.ist.hokudai.ac.jp
H. Iizuka
E-mail: iizuka@complex.ist.hokudai.ac.jp
M. Yamamoto
E-mail: masahito@complex.ist.hokudai.ac.jp




Fig. 1 Khepera-like robot model.

Table 1 Parameters of the simulated robot.

Body radius : R	2.9 cm
Sensor range : s_{max}	5.0 cm
Wheel speed	8 cm/s
Δt	0.1
Wheel noise range	[0.9:1.1]
Distance between the wheels : L	5.2 cm

tween agents. It is possible that role allocation was innately embedded in the different neural controllers, which means that the task could have been achieved without the agents communicating with each other. In fact, Quinn mentioned that there were different types in the population and that mixed-type (leader/follower) interaction achieved a better performance than same-type interaction (Quinn, 2001, p.365, [14]).

In order to investigate the evolution of communication more quantitatively, we reproduce Quinn's model and extend the original task to communication among three agents equipped with identical neural controllers. We increase the number of agents in order to make it more difficult to encode the roles into their neural structures and also to capture the emergence of complex communication more easily. To avoid the possibility of completing the task without communication, the agents have identical neural controllers. The resulting behaviors are evaluated both qualitatively and quantitatively. In particular, they are investigated in terms of two types of joint attention, namely, participatory and instrumental joint attentions. Mutuality, information flow, invariance against the motion gain are used as measures of these two joint attentions. Our study shows that communication can evolve simply through the behavior of three identical agents and that the evolved behaviors demonstrate these three communication characteristics.

2 Models

2.1 Khepera-like robot

We simulated three mobile robots interacting with each other in an arena. The agents were modeled as Khepera-like robots, equipped with two wheels and eight IR sensors that mea-




Fig. 2 Structure of the recurrent neural network.

sured the distances to the other agents (Fig. 1). The agents did not have functional communication organs, only actuators to move and IR sensors to perceive, so they could only communicate through behavior.

The sensor's maximum detection range s_{max} was 5cm. If an object is closer than 5cm, the sensor outputs a value between 0 and 1, which is converted linearly into a distance. The sensor values I_i^k of sensor i on robot k are given by the distance $\text{Dis}(Ps_i^k, Pa_j)$ between the sensor position Ps_i^k and the agent position Pa_j as follows:

$$I_i^k = \sum_{j=1, j \neq k}^3 \frac{\max\left(0, \left(s_{max} - \text{Dis}(Ps_i^k, Pa_j)\right)\right)}{s_{max}}. \quad (1)$$

When the sensors perceived other agents, random Gaussian noise (mean 0, standard deviation 0.25) was added to the sensor values. When nothing was in sensor range, uniform random noise in the range [0:0.1] was added instead. The sensor values for each agent were normalized to the range [0:1.0], so the maximum value was 2.0.

The translational and angular velocities, v and ω , of each robot were updated according to the following equations:

$$v = \frac{v_r + v_l}{2}, \quad \omega = \frac{v_r - v_l}{L}. \quad (2)$$

Here, v_l and v_r represent the velocities of the left and right wheels, respectively. The velocities v_l and v_r were multiplied by a uniform random noise factor. Table 1 shows the robot parameters used in our experiment.

2.2 Task

The task for the agents was to move as far away as possible from their initial positions in 30 sec while staying close together as a group and not colliding with each other. If each

agent were to choose a movement direction independently, they would not complete the task because they would neither stay close together nor move in the same direction. This means that they must communicate with each other at the start of the task about which direction they should move in and how to move while keeping close together. Furthermore, this communication must be encoded in their behavior, which must be recognized by the other agents as being different from goal-directed behavior.

2.3 Recurrent neural network

The agents were controlled by a recurrent neural network. The structure of the neural controller that we used was the same as that of the one used in Quinn's model [14,15]. To encourage the agents to evolve communication, they were given identical neural networks. In addition, the initial positions of the agents were symmetric to prevent them from encoding strategies for breaking the symmetry into their physical differences. This means that any agreement among the three agents must be the result of interaction because there is no inherent structural difference in their neural networks. Their neural networks consisted of eight sensor nodes connected to the IR sensors, 20 hidden nodes, and four wheel nodes, i.e., a forward node and a backward node, for each wheel (Fig. 2). The layers were fully connected.

The neurons' internal states were updated as follows:

$$m_j^t = \begin{cases} (1 - \gamma_{Aj})m_j^{t-1} + \sum_{i=1}^8 w_{ji}I_i & (h_j^{t-1} = 0) \\ (1 - \gamma_{Bj})m_j^{t-1} + \sum_{i=1}^8 w_{ji}I_i & (h_j^{t-1} = 1) \end{cases}. \quad (3)$$

Here, m_j^t is the internal state of the j th neuron, I_i is the i th sensor input, w_{ji} is the weight from the i th to the j th neuron at time t , and γ_{Aj} and γ_{Bj} are the decay constants in the range [0:1.0]. This model retains past internal states m_j^{t-1} so that the agents can determine their current behavior based on previous behavior. In addition, h_j^t is a hidden-layer output of the hidden layers, calculated from m_j^t and the neuron's threshold θ_j as follows:

$$h_j^t = \begin{cases} 1 & (m_j^t \geq \theta_j) \\ 0 & (m_j^t < \theta_j) \end{cases}. \quad (4)$$

If m_j^t exceeds θ_j , then $h_j^t = 1$; otherwise, $h_j^t = 0$.

The output O_k^t of the neural controller was calculated from h_j^t and the weights v_{kj} as follows:

$$O_k^t = \begin{cases} 1 & \left(\sum_{j=1}^N v_{kj}h_j^t > 0 \right) \\ 0 & \left(\sum_{j=1}^N v_{kj}h_j^t \leq 0 \right) \end{cases}. \quad (5)$$

When the sum is positive, then the output is 1; otherwise, it is 0. The output for each wheel was found by subtracting the value of the backward node from that of the forward node, which means that there were three possible patterns: $\{-1, 0, 1\}$. The speeds of the left and right wheels, v_l and v_r , were calculated by multiplying the wheel outputs by the constant wheel speed and the wheel noise factor.

2.4 Genetic algorithm (GA)

The agents' neural networks evolved using a genetic algorithm [16]. The parameters for each neural network, that is, the weights w_{ji} and v_{kj} , the decay constants γ_{Aj} and γ_{Bj} , and the threshold θ_j , were encoded into a single real-valued vector. The population size was 180, and they evolved for 2,000 generations.

Each real-valued vector was used to create three identical agents. The agents interacted freely with each other during each trial, and the behaviors that they produced were evaluated. The fitness function considered how far they moved during each trial while staying close together and how often they managed not to collide with each other. Each trial ran for $T (= 300)$ steps. The fitness function can be expressed as follows:

$$F = p \cdot \left(\sum_{t=1}^T \left[g^t \cdot \left(1.0 + 2.0 \cdot \tanh\left(\frac{s^t}{8.0 \cdot R}\right) \right) \right] \right), \quad (6)$$

$$p = \max\left(1.0 - \frac{c}{1 + c_{max}}, 0\right), \quad (7)$$

$$g^t = \begin{cases} 0 & (d^t \leq D^{t-1} \text{ or } D^{t-1} \geq 40.0) \\ d^t - D^{t-1} & (\text{otherwise}) \end{cases}, \quad (8)$$

$$s^t = \sum_{i=1}^3 \left(4.0 \cdot R - \text{Dis}(Pc^t, Pa_i^t) \right). \quad (9)$$

Here, p is the collision penalty, and c is the number of collisions during the trial. If c exceeds $(1 + c_{max})$, p is set to 0, and the fitness is also 0. In addition, g^t evaluates how far the agents have traveled from their initial position, d^t is the distance between the center of the agents' positions at step t and the initial position, and D^{t-1} is the maximum d^t score up to step $t - 1$. g^t is set to 0 if the agents have not improved on their best distance or if they have already achieved the maximum distance, which is 40cm in this paper. By evaluating the distance at every step, we could evaluate the farthest distance that the agents managed to travel from their initial position during the trial. In addition, s^t evaluates how close together the agents are: if the distance from the center of the agents' positions Pc^t to each agent Pa_i is less than $4.0 \cdot R$ cm, s^t is high, and the closer agents becomes, the higher its value is.




Fig. 3 Initial positions of the three agents.




Fig. 4 Best fitness values for each generation.

Each new generation was created as follows. First, the two best agents were left in the population (*elitism*). Then, 20% of the new population was produced by a uniform *crossover* between pairs of agents, chosen by a roulette strategy where the agents were ranked by their normalized fitness values. The rest of population was produced by *mutation*, where uniform random values in the range $[-0.05:0.05]$ were added to the real-valued vectors.

The initial layout of the three agents was fixed, as shown in Fig. 3. Groups of three identical agents, based on the real-valued vector, were tested over two sets of five trials, and they started each trial facing in a different direction, namely, $\alpha = \{0, \frac{2\pi}{5}, \frac{4\pi}{5}, \frac{6\pi}{5}, \frac{8\pi}{5}\}$. The fitness values were averaged over the 10 trials.

3 Evolved behaviors and their differentiation

3.1 Evolved behaviors

Figure 4 shows the resulting fitness values, which increased for 1,500 generations and then converged to 0.96 at around 2,000 generations. Since the maximum fitness score is 1, the evolved behaviors were good enough to achieve the task. The best evolved agents were then tested in the following experiments.

Figure 5 shows the example trajectories for the best evolved agents, and Fig. 5(b) shows a magnified view near their initial position. The agents appeared to show a communica-




Fig. 5 Example of the best evolved agents' behavior.

tive behavior near the initial position during the early stages of the trial, with their behavior changing to be more goal-directed in the latter stages. During the communicative behavior phase, they seemed to make agreements as to the direction in which they should move. During the goal-directed behavior phase, they aligned with each other and moved in the decided direction while staying close together. Figure 6 shows examples of their interaction during the communicative behaviors phase. Initially, all three agents rotated slowly in the same way to face each other (Fig. 6(a)). This happened because the initial positions and direction were symmetric and identical neural networks were used. The agents then started oscillating back and forth (Fig. 6(b)). In Figs. 6(a) and 6(b), the symmetry of the three agents was maintained, but it started breaking down in Fig. 6(c). Here, agent 0 and agent 1 lined up in parallel, whereas agent 2 faced perpendicularly to them, forming a U-shape. Among the agents lined up in parallel, the agent facing outward then took the front role and started moving forward (Fig. 6(d)). This action, namely, the front agent moving forward, was the signal that the direction had been determined. After a while, they all started moving in the same direction (Fig. 6(e)).

The goal-directed behaviors of the agents could be separated into three roles, namely, *front*, *center* and *rear*, depending on their alignment and the need to maintain formation during the goal-directed phase (Fig. 7). Once the roles had been allocated, they were fixed for the direction of the trial. The agents therefore decided their roles as well as their direction of motion before the goal-directed behavior appeared.

The directions in which the agent team moved and the way the roles were allocated changed between trials. The fact that both their initial configuration (positions and angles) and their neural networks were identical indicates that the roles were allocated through embodied interactions among the agents. This means that the behaviors were somehow being used to communicate with each other.




Fig. 6 Example of interaction behaviors during the communicative phase.

3.2 Differentiation into communicative and goal-directed behaviors

The results in the previous section show that the agents were able to decide both the movement direction and their roles through embodied interactions. In this section, we examine whether the behaviors used by the agents for communication were different from their goal-directed behaviors and when this differentiation occurred.

If the roles were allocated to the agents by communication until a certain time step, these allocated roles would not change even if many different simulations (with different noise sequences) were performed starting from that time step because, at that point, role allocation was complete. However, this would not be the case before that time step because communication would still be going on. In order to clarify when the roles were determined by communication, we tested this hypothesis. First, we recorded all information about the agents' behaviors and the internal states of their neural controllers during a single trial. Next, we performed many different simulations, starting from the stored states at each time step, to see whether or not the roles ultimately allocated to the agents would change. For this analysis, we performed 100 different simulations for each time step.

Figures 8(a), (b), and (c) show the ratios with which each agent was finally allocated each role, where the horizontal axis shows the time step at which the simulations started. The original roles in the recorded trial were front, center, and rear for agents 0, 1, and 2, respectively. After time step 75, the roles did not change, even when different noise sequences were applied, and these roles were the same as those for the original trial. This means that the process of role allocation by communication was complete by time step 75. On the other hand, before that time step, the noise sequence affected the role allocation result, which means that they were still communicating. This distinction threshold also appeared in other tests using other recorded trials, showing qualitatively that the behaviors of the agents had differentiated into communicative and goal-directed types.




Fig. 7 Formation during the goal-directed phase.

4 Quantitative analysis of evolved behaviors

Although we have shown that the goal-directed and communicative behaviors become distinct at a certain time step, it is not yet clear whether the behavior that we have seen is really communicative, a point that was not clearly shown in Quinn's original paper either. As discussed in [17–20], two kinds of joint attention play an important role in developing proto-linguistic communication: participatory joint attention, whose purpose is to maintain social interaction and share experiences, and instrumental joint attention, which communicates information between the speaker and the hearer. Now, we evaluate our communicative behavior quantitatively in terms of these joint attentions. We use mutuality as a measure of participatory joint attention, as it is regarded as a bidirectional interaction in [21,22]. For instrumental joint attention, we measure the information flow to show the actual information transfer due to communication. In addition to information flow, we also show that communicative behavior is separate from the agent's usual movements, as in the predefined channel model, as its properties are invariant with changes in motion gain.

4.1 Mutuality

The primary characteristic of communication is mutuality. Treverthen et al. and Nadel et al. have shown that mutual-




Fig. 8 Ratios of the final roles that were assigned to each agent. See the main text for further details.




Fig. 9 Agent behaviors in the recording experiment, where one of the three agents just reproduced recorded movements. The agents could not stay close and did not demonstrate coordinated behavior.

ity is required to maintain communication [23,24]. In this section, we investigate whether the communicative behavior that we have seen also demonstrates mutuality, using recordings to test this. Here, the movements of one agent were replaced with movements that had been recorded during a different successful trial. This means that the recorded agent appeared to be behaving properly, but it could not respond to the other two agents in real time. If trials including the recorded agent fail, this means that the original interactions among the agents require mutuality.

Figure 9 shows an example of the agents' behavior when one agent was replaced with a recording that just reproduced its previous movements, showing that the agents failed to establish coordination. The recorded agent separated from the other reactive agents, and the number of collisions increased compared with normal trials because they used mutual interaction to avoid collisions. This shows that the agents were unable to coordinate their behavior without the real-time responses. The performance score reduced to 0.56, compared with 0.95 (averaged over 100 trials) for the normal trials.

Table 2 Fitness scores for each role for the recorded trial.

	Role		
	Front	Center	Rear
Fitness Score	0.95	0.95	0.42

Therefore, the recorded agent clearly caused communication to fail.

Since the recorded experiments started from the beginning of the trial, communication could not be properly established. Next, we investigated what happens when the recording begins during the goal-directed phase via the following experiments. We first recorded the movement of all of the agents in a trial and then reproduced the first 100 steps (so that the communication process was complete). After that, one of the three agents reproduced its recorded movements from the first trial, and the other agents used their neural controllers to behave reactively. The noise sequences were newly generated each time. We performed this experiment for each role and compared the fitness scores. Since the communicative behaviors were complete after 100 time steps, we were able to investigate mutuality during the goal-directed phase using the recording for the last 200 steps.

Table 2 shows a comparison of the resulting fitness scores when each role was replaced by a recorded movement during the goal-directed phase. Even when the movements of the front and center agents were given by recordings, the agents succeeded just as well in the task, which means that the front and center agents did not need to be responsive to the other agents' movements. However, when the rear agent was replaced with a recording, they failed at the task, and the fitness decreased significantly. This must be because the rear agent had to react to the movements of the front and center agents, i.e., it had to follow them. This shows that the front and center agents were not reactive, and the rear agent just followed the other two, which means that their interaction were one-directional.




Fig. 10 Transfer entropy between pairs of agents with different roles.




Fig. 11 Relative amount of information carried by each role.

4.2 Information flow

In order to evaluate quantitatively how much one agent's behavior affects those of the others and how information flows, we measured the transfer entropy (TE), which calculates the amount of information flow between two time series [25]. If their behaviors are independent, the TE should be zero. Here, we calculated the TE s for time series from pairs of agents with different roles. The TE from role J to role I was calculated as follows:

$$TE_{J \rightarrow I} = \sum p(i_{t+1}, i_t^{(k)}, j_t^{(l)}) \log_2 \left(\frac{p(i_{t+1} | i_t^{(k)}, j_t^{(l)})}{p(i_{t+1} | i_t^{(k)})} \right). \quad (10)$$

Here, t is the time index, and k and l represent the numbers of time steps required to define discrete states for each role. In this paper, k and l are both set to 1. In addition, $i_t^{(k)}$ and $j_t^{(l)}$ are the states (from the set of all possible states) of the two agents at time t .

The discretized states of the agents at time t were defined by their movements. Since there are three output patterns for each wheel, the total number of possible states was nine. In

this paper, the TE was calculated for sliding time windows in order to see the way the information flow changed for communicative and goal-directed behaviors. The width of the time window was 50.

Figures 10(a), (b), and (c) show the TE results for all possible pairs of roles. Because the width of the time window was 50, the graph starts at time step 50. The TE during the communicative phase was larger than that during the goal-directed phase, showing that more information was transferred among the agents during this. Since the roles had yet to be decided, we do not see directional information flow. After the goal-directed phase began, the total amount of information flow decreased, but it became asymmetric between the agents. For example, more information flowed from the front to the rear agent than went the other way. To estimate which agents produced the information and which received it, we calculated the TE difference between the outward and inward information flow for each role as follows:

$$TE_I = (TE_{I \rightarrow J} + TE_{I \rightarrow K}) - (TE_{J \rightarrow I} + TE_{K \rightarrow I}). \quad (11)$$

Here, I , J , and K indicate the agents' roles.

Figure 11 shows the average amount of information for each role, averaged over 100 trials. During the communicative phase, there was no difference among the three roles. However, during the goal-directed phase, the front agent was the largest information source, and the other two roles received more information. The ANOVA results showed no significant difference for any of the role pairs during the communicative phase, but showed significant differences during the goal-directed phase ($p < .05$ after Bonferroni correction). This is because the front agent played a leadership role, deciding which direction they should move in and requiring the others to follow to maintain the formation. The reason why the rear agent needs mutuality during goal-directed behaviors is because the rear has to receive more information by reacting to the front's or center's behaviors. From these results, we can conclude that the agents achieved the task effectively by changing the way the information flowed between the communicative and goal-directed phases.




Fig. 12 Ratios of the final roles assigned to each agent in the speed-limited experiments. The speed of agent 0 was limited to 6 cm/s.




Fig. 13 Communication time for each speed of agent 0.




Fig. 14 Partial fitness scores under normal and speed-limited conditions.

4.3 Invariance of communication against the motion gain

If the established communication is an abstract system, the content and meaning should be consistent regardless of motion gain (e.g., louder or quieter voice). The information does not depend on the volume because the content of our communication is somehow detached from its physical properties. Following this idea, we tested the communicative behaviors that we had obtained in terms of their invariance properties. In this experiment, the speed of one of the three agents was limited to see if the communicative behaviors changed. Our hypothesis was that the communicative behaviors would not be affected by the speed limitation if it was used for communication, similar to the way our communication is independent of the volume of our voices.

The same (i.e., best evolved) agents were used for this experiment. They had evolved under normal conditions, where all of the agents traveled at a speed of 8 cm/s. However, in this experiment, the speed of one of the agents (agent 0) was limited to 6 cm/s, and we investigated how they would behave in such a situation. We analyzed their behavior in the same way as before (see Fig. 8), and we tested whether this affected the time required for the communicative phase.

The experimental results are shown in Figs. 12(a), (b), and (c). In the trial, agents 0, 1, and 2 were allocated the rear,

front, and center roles, respectively. The results show that the communication phase took about 75 time steps, almost the same as under normal conditions (Fig. 8), indicating that motion gain did not affect communication time. To investigate this further, we calculated the communication times for speeds between 6 and 8 cm/s. The agents frequently failed the task when the speed was less than 6.0 cm/s, which means that we could not calculate when the communication phase ended. We judged when communication ended by finding the time step when the ratios for all roles (as shown in Fig. 12) exceeded 95%. After this point, the role allocation did not change. Figure 13 shows the communication times, averaged over 100 trials, for each speed, and the error bars indicate the standard error. This indicates that the communication time was around 75 time steps regardless of speed. This means that the speed limitation did not affect the communication time because communication was established by exchanging signals that did not depend on motion gain. However, since the goal-directed behaviors depended directly on such embodied constraints, the task completion time was affected. This is further evidence that communicative behavior is different from goal-directed behavior. It should be noted, however, that the role allocation ratios were different from those in the previous experiment in the early stages, as shown in Figs. 12(a), (b), and (c). This was because it was

difficult for the speed-limited agent to play the front role because it moved slowly.

Figure 14 shows the changes in the cumulative partial fitness scores over time. The distance d^t that the team traveled and its density s^t are also shown, averaged over five different simulations. These results compare the normal and speed-limited conditions, showing that the distance fitness values were not affected during communication, but that the speed limitation caused the values to decrease after the goal-directed phase began.

5 Conclusion

In this paper, we have shown that communication can evolve purely through behavior in computer simulations of three identical agents and have evaluated the respects in which the evolved behaviors were really communicative. After evolution, the agents' behavior differentiated into communicative behavior during the early stages, followed by goal-directed behavior in the latter stages. We have analyzed these behaviors quantitatively in terms of participatory and instrumental joint attentions, i.e., mutuality, information flow, and invariance against the motion gain. A transfer entropy analysis clearly showed that the communicative behaviors involved more information flow than the other behaviors. Our results show that the agents evolved behaviors featuring these three important communication properties, and hence, that communication can evolve from embodied interactions.

References

1. Stefano Nolfi and Dario Floreano. Evolutionary robotics, 2000.
2. Rolf Pfeifer and Christian Scheier. *Understanding intelligence*. MIT press, 2001.
3. Inman Harvey, Ezequiel Di Paolo, Rachel Wood, Matt Quinn, and Elio Tuci. Evolutionary robotics: A new scientific tool for studying cognition. *Artificial life*, Vol. 11, No. 1-2, pp. 79–98, 2005.
4. Hidenori Kawamura, Masahito Yamamoto, Keiji Suzuki, and Azuma Ohuchi. Ants war with evolutive pheromone style communication. In *European Conference on Artificial Life*, pp. 639–643. Springer, 1999.
5. Davide Marocco, Angelo Cangelosi, and Stefano Nolfi. The role of social and cognitive abilities in the emergence of communication: Experiments in evolutionary robotics. In *EPSRC/BBSRC international workshop biologically-inspired robotics*, pp. 174–181, 2002.
6. Gregory M Saunders and Jordan B Pollack. The evolution of communication schemes over continuous channels. *From Animals to Animats*, Vol. 4, pp. 580–589, 1996.
7. Davide Marocco and Stefano Nolfi. Emergence of communication in embodied agents: Co-adapting communicative and non-communicative behaviours. In *Connection Science*. Citeseer, 2005.
8. Angelo Cangelosi. Evolution of communication and language using signals, symbols, and words. *IEEE Transactions on Evolutionary Computation*, Vol. 5, No. 2, pp. 93–101, 2001.
9. Bruno Galantucci. An experimental study of the emergence of human communication systems. *Cognitive science*, Vol. 29, No. 5, pp. 737–767, 2005.
10. Simon Kirby, Hannah Cornish, and Kenny Smith. Cumulative cultural evolution in the laboratory: An experimental approach to the origins of structure in human language. *Proceedings of the National Academy of Sciences*, Vol. 105, No. 31, pp. 10681–10686, 2008.
11. Takashi Ikegami and Hiroyuki Iizuka. Joint attention and dynamics repertoire in coupled dynamical recognizers. *The Proceedings of Imitation in Animals and Artifacts II*, pp. 125–130, 2003.
12. Thomas C Scott-Phillips, Simon Kirby, and Graham RS Ritchie. Signalling signalhood and the emergence of communication. *Cognition*, Vol. 113, No. 2, pp. 226–233, 2009.
13. Hiroyuki Iizuka, Davide Marocco, Hideyuki Ando, and Taro Maeda. Experimental study on co-evolution of categorical perception and communication systems in humans. *Psychological research*, Vol. 77, No. 1, pp. 53–63, 2013.
14. Matt Quinn. Evolving communication without dedicated communication channels. In *European Conference on Artificial Life*, pp. 357–366. Springer, 2001.
15. Matt Quinn. Evolving co-operative homogeneous multi-robot teams. In *Intelligent Robots and Systems, 2000.(IROS 2000). Proceedings. 2000 IEEE/RSJ International Conference on*, Vol. 3, pp. 1798–1803. IEEE, 2000.
16. John H Holland. *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence*. U Michigan Press, 1975.
17. Hiroyuki Iizuka and Takashi Ikegami. Adaptability and diversity in simulated turn-taking behavior. *Artificial Life*, Vol. 10, No. 4, pp. 361–378, 2004.
18. Michael Tomasello. *The cultural origins of human cognition*. Harvard university press, 2009.
19. Ryoko Uno and Takashi Ikegami. Joint attention/prediction and language: A mechanism to align intentionalities. *Papers in Cognitive Linguistics*, Vol. 2, pp. 231–274, 2003.
20. Lynne Murray. Emotional regulations of interactions between two-month-olds and their mothers. *Social perception in infants*, pp. 177–197, 1985.
21. Davide Marocco and Stefano Nolfi. Emergence of communication in embodied agents evolved for the ability to solve a collective navigation problem. *Connection Science*, Vol. 19, No. 1, pp. 53–74, 2007.
22. Ryoko Uno, Davide Marocco, Stefano Nolfi, and Takashi Ikegami. Emergence of protosentences in artificial communicating systems. *IEEE Transactions on Autonomous Mental Development*, Vol. 3, No. 2, pp. 146–153, 2011.
23. Colwyn Treverthen and Lynne Murray. Emotional regulation of interaction between two-month-olds and their mothers. *Social perception in infancy*, 1985.
24. Jacqueline Nadel, Isabelle Carchon, Claude Kervella, Daniel Marcelli, and Denis Réserbat-Plantey. Expectancies for social contingency in 2-month-olds. *Developmental science*, Vol. 2, No. 2, pp. 164–173, 1999.
25. Thomas Schreiber. Measuring information transfer. *Physical review letters*, Vol. 85, No. 2, p. 461, 2000.