

# The Emergence of Communication by Evolving Dynamical Systems

Steffen Wischmann<sup>1</sup> and Frank Pasemann<sup>2</sup>

<sup>1</sup> Bernstein Center for Computational Neuroscience,  
Bunsenstr. 10, D-37073 Göttingen, Germany  
[steffen\\_wischmann@web.de](mailto:steffen_wischmann@web.de)

<http://www.chaos.gwdg.de/~steffen>

<sup>2</sup> Fraunhofer Institute for Autonomous Intelligent Systems,  
Schloss Birlinghoven, D-53754 Sankt Augustin, Germany  
[frank.pasemann@ais.fraunhofer.de](mailto:frank.pasemann@ais.fraunhofer.de)

**Abstract.** In the context of minimally cognitive behavior, we used multi-robotic systems to investigate the emergence of communication and cooperation during the evolution of recurrent neural networks. The networks are systematically analyzed to identify their relevant dynamical properties. Evolution efficiently adapts these properties through small structural changes within the networks when specific environmental conditions are altered, such as the number of interacting robots. The findings signify the importance of reducing the predefined knowledge about resulting behaviors, dynamical properties of control, and the topology of neural networks in order to utilize the strength of the Evolutionary Robotics approach to Artificial Life.

## 1 Introduction

The dynamical systems approach to cognition [1,2,3] aims at the study of natural cognitive systems as dynamical systems. While concrete dynamical models of cognitive phenomena are still under construction, “one powerful way to improve our intuitions, clarify the key issues and sharpen the debate is through a careful study of simpler idealized models of *minimally cognitive behavior*, the simplest behavior that raises issues of genuine cognitive interest” [4]. We consider minimal cognition as metabolism-independent sensorimotor behavior [5] and presuppose that cognitive behavior generally results from perception-action couplings [6].

The Evolutionary Robotics [7,8] approach to Artificial Life [9] aims at the emergence of such perception-action couplings during the evolution of *complete brain-body-environment systems* [4,10].

As artificial brain structures we utilize recurrent neural networks (RNNs) which can be described as parameterized dynamical systems [11]. We distinguish two types of parameters. The first type concerns parameters of single neurons (bias terms), the structural coupling (the topology), and the strength of these couplings (the synaptic weights). These parameters are shaped by an evolutionary algorithm. The second type of parameters are characterized by the sensor

states of a robot, which are represented by the activation of input neurons. They dynamically change during the interaction of a robot with the environment.

Here, we will investigate the dynamical properties of evolved RNNs and their relation to observable collective behavior in groups of robots, especially by systematically exploring the sensor input activations provoked by robot-environment interactions. The detailed analysis of communication underlying dynamical properties of recurrent neural networks and their relation to structural changes during evolutionary processes distinguishes our work from pioneer studies on the evolution of emergent communication among artificial agents [12,13] as well as from more recent studies [14,15]. These studies describe significant results but they mainly focus on a detailed analysis at the behavioral level of communicating artificial agents.

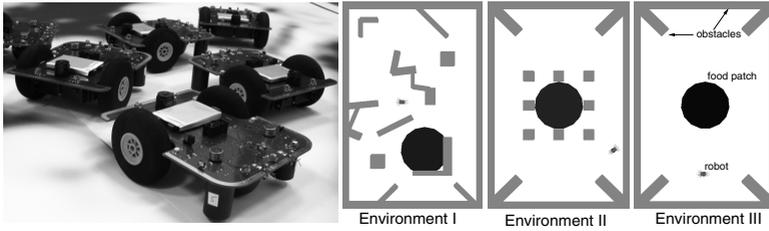
In earlier studies we utilized RNNs to coordinate conflicting behaviors in very large robot groups [16]. There, we *manually designed* a local communication system between several robots in order to synchronize individual internal neural rhythms which determine the behavior of each robot. For the following experiments, we used the same robot platform, but implemented a much simpler task in order to investigate how communication can *emerge* as the basis of cooperation by reducing the predefined knowledge assigned to the evolutionary process. Furthermore, we investigate how evolution shapes certain parameters of behavior underlying dynamical systems and how it adapts these parameters to specific changes of the environment.

## 2 The Ingredients for the Emergence of Communication

To keep the analysis of evolved RNNs, concerning the dynamical properties and their relation to behavior, still tractable, we use a neuron model with only two parameters, a bias term and a synaptic self-weight [17]. A network consisting of  $n$  units is then defined as a parameterized discrete-time dynamical system:

$$a_i(t+1) = \theta_i + \sum_{j=1}^n w_{ij} f(a_j(t)), \quad i = 1, \dots, n, \quad (1)$$

where  $a_i \in \mathbf{R}$  denotes the activity of neuron  $i$ ,  $w_{ij}$  the synaptic strength of the connection from neuron  $j$  to neuron  $i$ , and  $\theta_i$  its fixed bias term. The output  $o_i = f(a_i)$  of a unit  $i$  is given by a sigmoid transfer function, here by  $f := \tanh$  (i.e.,  $o_i \in (-1, 1)$ ). Although this neuron model is rather simple, already small recurrent networks of this type can generate complex dynamics, such as periodic, quasi-periodic, or even chaotic attractors [11]. For the evolution of these dynamical systems we used an implementation (see [18] for details) of the evolutionary algorithm *ENS*<sup>3</sup> [19]. This algorithm optimizes the parameters intrinsic to the RNN, such as synaptic weights, bias terms, *and the topology of the network*. For evolution and analysis a physical simulation environment was created. There, we implemented important properties, such as noise of sensors and motors, in accordance with results of measurements done with the physical robot.



**Fig. 1.** The physical robot *Do:Little* (left) and three simulation environments of decreasing complexity (I  $\rightarrow$  III)

We utilized the *Do:Little* robot (illustrated in Fig. 1) as an artificial creature. The advantages of this robot are its simple but very reliable sensor and motor capabilities. Besides infrared sensors for obstacle detection and floor sensors for measuring the gray scale of the ground, we especially made use of its robust communication system. It consists of a stereo microphone which can detect the direction of sound signals emitted by nearby robots. In every interval of the robot's update cycle (100 ms) the robot can produce several sound signals of the same frequency. Signals are differentiated by a unique sequence of pulses within one update cycle. The advantage of not coding different signals with different frequencies is that the robots are able to detect their sound signals very reliable even in rather noisy environments. Thus only acoustic signal peaks can be detected but no continuous sound signals. We will see later how this constraint will influence the evolution of communication behavior if we change the population size.

Neurons  $I_1$ ,  $I_2$ , and  $I_3$  represent the left, right, and back infrared sensors, respectively. The sensor inputs are linearly mapped onto  $[-1, 1]$ , where  $-1$  means no obstacle detection and  $+1$  indicates very close obstacles. Neuron  $I_4$  represents the floor sensor. The inputs are also linearly mapped onto  $[-1, 1]$ , where  $-1$  indicates white colored and  $+1$  black colored ground. For communication we only used one acoustic signal. The angle  $\alpha$  of a perceived sound signal to the heading direction of the robot is represented by  $o(I_5) = 0.5 \cdot (1 + \sin(\alpha))$  and  $o(I_6) = 0.5 \cdot (1 - \sin(\alpha))$ . The speed of the left and right wheel are calculated by  $c \cdot (o(O_1) - o(O_2))$  and  $c \cdot (o(O_3) - o(O_4))$ , respectively, where  $c$  is a speed factor. Important for the understanding of the described communication systems is that the robot emits a single sound signal when  $o(O_5)$  switches from a negative to a positive value.

In the following experiments we wanted to know more about the minimal requirements necessary to provoke the emergence of communication within a population of robots during the evolution of their control architectures. Hence, we defined the following simple task: A single robot can increase its fitness by exploring the environment and finding patches of food while avoiding collisions with obstacles and other robots. Thereby an individual can benefit from the behavior of other robots if they cooperate. Such robots, sharing a common environment, are conspecific because they are identical with respect to their morphology and control, and the selection process during evolution is group based because the

mean fitness of all robots in a group is taken<sup>1</sup>. However, this does not necessarily mean that only cooperative behavior can be successful or communication will inevitably emerge. For instance, even solitary behavior can be efficient if each individual is able to locate food patches reliably without running into obstacles or other robots.

To overcome the well known bootstrap problem of the evolutionary approach to the development of behavior [7], we applied a so called semi-restrictive incremental method [18]. Therefore, in the first evolutionary step the task of a single robot was to explore its environment as good as possible without running into obstacles. For this task robots were equipped with infrared sensors for avoiding obstacles. The topology of the neural network was not determined, only input and output neurons were defined. Structural elements, such as synapses and hidden neurons, could freely emerge in between. In the second evolutionary step, robots could additionally access a floor sensor for detecting black food patches on the ground. We selected several different RNNs which were successful in solving the exploration task as a basis for evolving RNNs which are now supposed to force the robots to stay on a food patch as soon as they find one. During this second evolutionary process, already existing structural elements were not allowed to be removed (whereas their parameters could change), but new structural elements could emerge within the *whole* network which now also has new sensor inputs. The same technique was applied for the last step, where robots, in addition, could access a speaker and a stereo microphone for emitting and sensing sound signals. RNNs resulting from the preceding step provided the basis for this evolutionary run.

Consequently, after the first evolutionary step we always put a certain predefined knowledge in each subsequent step. However, this was only done to provide basic behaviors for the evolution of more complex behaviors, for which we never defined how a network eventually should look like. Therefore, we argue that the emergence of communication during evolution was neither explicitly forced by a given network structure nor by the fitness function. Hence, the remaining constraint was the design of the environment.

At first, we thought that a complex environment, such as Environment I, shown in Fig. 1, would enforce the emergence of cooperation. In this environment it is rather complicated for a solitary individual to quickly find the food patch. Once an individual find it perchance, it could use its communication system to guide the others. Surprisingly, even after many repetitions of the evolutionary process no cooperation emerged. In our opinion this is because of the bootstrap problem [7]. It may take too much time until an individual finds the food patch. And consequently, even when it then starts to call other robots, this would not significantly increase the performance compared to robust solitary behavior. Therefore, a step-wise refinement of the communication system (note, the robots had to learn signal-

---

<sup>1</sup> The fitness of a single robot  $i$  is  $F_i = 600 \frac{k_i}{T}$ , where  $T$  is the number of evaluation time steps and  $k_i$  is defined by how often the robot is able to find a food patch in  $T$  (whenever the robot finds a food patch, it recharges its virtual battery and is then replaced randomly within the environment).

ing *and* the appropriate responses to other signals) may become very improbable during evolution (this first assumption has to be verified in future work). Thus, we decreased the environmental complexity by removing obstacles and placing the food patch in the center (Environment II), but even there no cooperation emerged during evolution. Only further removal of obstacles (Environment III) enabled the emergence of cooperation, as it will be discussed in the following section. Note, that although all RNNs were evolved in the rather simple Environment III, the resulting cooperative behavior was robust enough that, in the end, we could also observe better performance in the more difficult Environment I and II compared to solitary behavior without any additional optimization.

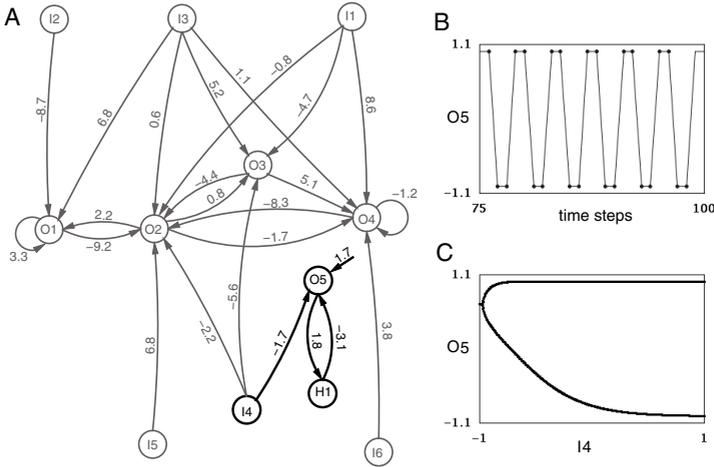
### 3 Dynamics of Evolved Communication Systems

#### 3.1 Communication in Small Groups of Robots

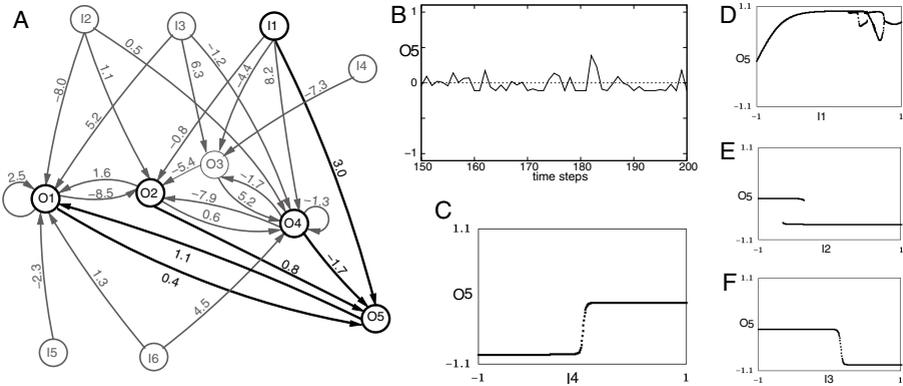
One small sized network resulting from the evolution of robot groups containing 10 individuals is drawn in Fig. 2A (we call individuals with this RNN as control architecture individuals of type *A1*). For completeness, the whole network and its parameters are given, but in the following we will concentrate only on the communication system intrinsic to the RNN. By means of an odd loop with over-critical synaptic weights<sup>2</sup>, the sound generating output neuron *O5* is connected with a hidden neuron (*H1*). This loop acts as a switchable oscillator [11] depending on the value of *I4*, the floor sensor input. *I4* is equal to -1.0 as long as the robot is moving on white ground. As we can see in the bifurcation diagram (Fig. 2C) the oscillation, caused by a period-4 attractor, is switched on by an increased activation of *I4*. The bifurcation point is very close to  $I4 = -1.0$ , and therefore, it can already be crossed by noise of the floor sensor. However, in order to emit a sound signal at least two points of the periodic orbit have to be in the negative and in the positive domain. This is only the case for  $I4 > -0.7$  (never reached by sensor noise only). As a constraint of the environment, food patches always provoke sensor signals of *I4* within  $[0.8, 1.0]$ . For these values the output of *O5* oscillates as shown in Fig. 2B. Thus, communication is context-sensitive: whenever a robot detects a food patch, it emits a sound signal every 4 time steps. This signal triggers a positive taxis in nearby perceiving robots through the input neurons *I5* and *I6*. Consequently, these robots will approach the signaling robots until they reach the food patch where they then also immediately start signaling. Therefore, communication is unidirectional, because the signaling of one robot alters the behavior of another robot, but this behavioral change does not influence the behavior of the signaler.

Another RNN with a completely different solution for context-sensitive communication is shown in Fig. 3A. We call individuals with this control to be of type *A2*. There, communication is realized by utilizing sensor noise. If no obstacle is close to the left side of the robot ( $I1 = -1.0$ ), the robot will stay on

<sup>2</sup> Here, due to the use of *tanh* as activation function, over-critical means a synaptic weight  $|w_{ij}| > 1.0$ . See [3] for details.



**Fig. 2.** A: RNN of *A1* individuals (group size:  $n = 10$  individuals during evolution) with inputs from left (*I1*), right (*I2*), and back (*I3*) distance sensors, left (*I5*) and right (*I6*) sound sensors, and the floor sensor (*I4*). Output neurons *O1* & *O2* and *O3* & *O4* steer the left and right wheel, respectively. *O5* controls the signaling. *H1* is a hidden neuron. B: Signals of *O5* when a robot stays on a food patch ( $I1 = I2 = I3 = -1.0; I4 = 1.0; I5 = I6 = 0.0$ ). C: Bifurcation diagram for *O5* by varying *I4* ( $I1 = I2 = I3 = -1.0; I5 = I6 = 0.0$ ).



**Fig. 3.** A: RNN of *A2* individuals ( $n = 10$  during evolution, cf. Fig. 2A). B: Signals of *O5* when the robots stays on a food patch (cf. Fig. 2B). C: Bifurcation diagram for *O5* by varying *I4* (cf. Fig. 2C). D-F: Bifurcation diagram for *O5* by varying each distance sensor (*I1*, *I2*, *I3*) input, where  $I4 = 1.0; I5 = I6 = 0.0$  (not varied distance sensor inputs are set to -1.0).

a detected food patch because the outputs of *O1* and *O2* are equal (within the upper saturation domain of the transfer function, i.e., 1.0) as well as the outputs of *O3* and *O4* (within the lower saturation domain, i.e., -1.0). Correlating these

values with the connections projecting to  $O5$  (emphasized in Fig. 3A) leads to an activation of  $O5 = 0$ . In this situation the asymptotically stable fixed point (SFP in the following) is in the *linear domain* of the transfer function. Hence, the neural activation is highly sensitive to the sensor noise of  $I1$ . Therefore the output crosses the zero-line randomly from the negative to the positive domain (see Fig. 3B). Diagram C shows how the SFP of  $O5$  is shifted from the *lower saturation domain* to the *linear domain* when the robot detects a food patch via  $I4$ . Once the robot stays on a food patch, this SFP can be shifted away depending on the activation of  $I1$ ,  $I2$ , and  $I3$  (Fig. 3D-F). That means, if another robot approaches a signaling robot, and consequently activates its infrared sensors accordingly, the signaling will cease.

Using noise for behavior control of autonomous robots is usually not wanted and engineers try to eliminate it from their systems as often as possible. In our example, it is a quite efficient solution for signaling. Infrared sensors are always noisy, and we tested different noise levels in the simulation environment<sup>3</sup> with the result that the behavior does not qualitatively change when we vary the noise level between 2% and 10%. In contrast to the prevention of noise in most technical applications, for biological systems it is well known that noise can significantly enhance sensorimotor patterns by means of a mechanism called Stochastic Resonance [20].

When we compare  $A1$  and  $A2$  with respect to their performance depending on the group size (see Fig. 5), we see that the more individuals are interacting in the same environment, the better  $A2$  performs compared to  $A1$  (if  $n > 7$ ). The reason is the described constraint of the physical communication system, namely the ability for perceiving only sound signal peaks. The more individuals of  $A1$  are signaling at the same time, the higher the probability that their signals will sum up to a continuous signal which cannot be perceived anymore by other robots still searching for food. Already four individuals of  $A1$  can produce a continuous signal when they are all signaling with different phases. Note, this is not simply an artifact of the simulation. Experiments with physical robots have also shown that the maximal frequency, where two subsequent signals can be distinguished, is  $5Hz$ .

In contrast to the constant period-4 signals in  $A1$ , the individuals of  $A2$  signal rather randomly. Consequently, the probability of producing a continuous signal for a longer time period is rather low in larger groups. Additionally, whenever a food patch becomes crowded, signaling robots will perceive nearby robots by their infrared sensors which in turn will stop their signaling, as we have discussed above (Fig. 3D-F). Thus, communication can no longer be described as unidirectional because signaling of food patch locations will attract other robots which in turn influence this signaling behavior as soon as they come close to the signaler.

Nevertheless, both control architectures resulted from the evolution with a group size of 10 individuals. For this size the performance difference between  $A1$  and  $A2$  is not as significant as it becomes with increased group size (Fig. 5).

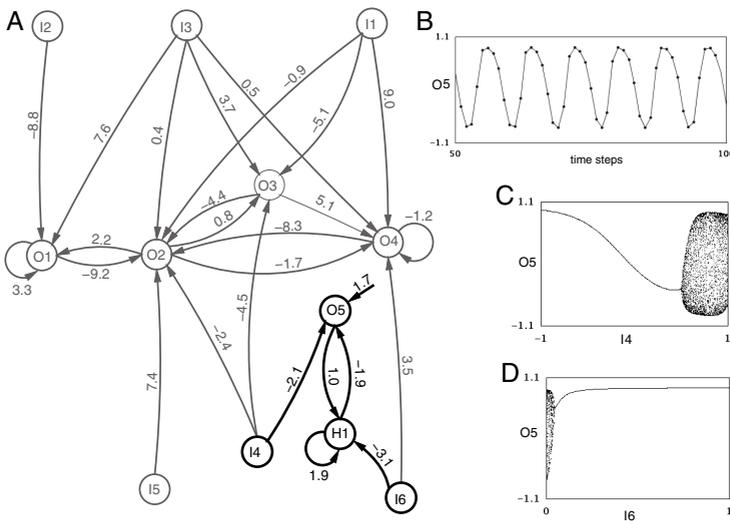
---

<sup>3</sup> The noise level of the physical infrared sensors is between 4% and 6%.

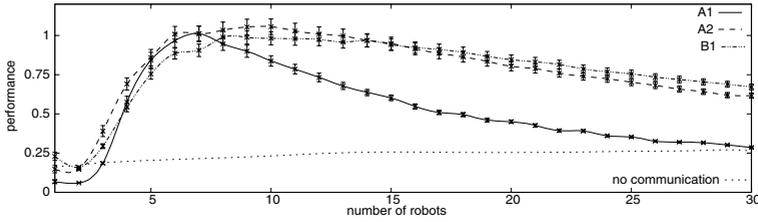
### 3.2 Evolutionary Changes of Communication in Larger Groups

The results of the previous section suggest that performance may improve when we repeat the evolution with a larger group size. More interesting than a simple performance improvement would be to see how the communication system, as a part of the RNN, will change when we increase the number of interacting individuals. Therefore, we started a new evolutionary run with the control architecture of *A1* as initial structure (which was more appealing because of its independence of sensor noise and its lower fitness at larger population size). We increased the group size to 25 individuals and allowed again parameter changes of the initial RNN as well as the emergence of new structural elements.

One resulting RNN is shown in Fig. 4A. When we compare the structural elements responsible for the communication system with the initial RNN of *A1* (Fig. 2A), we notice the same odd loop between *H1* and *O5* with over-critical synaptic weights. In addition, we found an over-critical self-connection at *H1*. With the given weight configuration this module exhibits quasi-periodic oscillations (Fig. 4B) which are switched on by an increased activation of the floor sensor *I4* (see Fig. 4C). We applied a power spectrum analysis to the time series in Fig. 4B and found a mean period length of about 8.7 time steps. The period of time between emitting two subsequent sound signals is now almost twice as long as in *A1*. Although this is presumably a coincidence, the correlation is interesting because the group size used in evolution of *B1* is also almost twice as large as it was used for the evolution of *A1*.



**Fig. 4.** A: RNN of *B1* individuals ( $n = 25$  during evolution, cf. Fig. 2A). B: Signals of *O5* when the robot stays on a food patch (cf. Fig. 2B). C: Bifurcation diagram for *O5* by varying *I4* (cf. Fig. 2C). D: Bifurcation diagram for *O5* by varying *I6* with  $I1 = I2 = I3 = -1.0$ ;  $I4 = 1.0$ ;  $I5 = 0.0$ .



**Fig. 5.** Performance of *A1*, *A2*, and *B1* individuals. For each group size several simulation runs ( $N = 25$ ) were performed, each lasting 18,000 time steps. Performance is the mean number of finding a food patch in 600 time steps for a single robot (mean of all robots in a group is taken). The mean performance of *A1*, *A2*, and *B1* individuals where the communication system was deactivated is drawn as a reference (dotted line).

Another new structural element is the connection between *I6*, the left microphone input, and *H1*. Whenever a robot is staying on a food patch, and therefore  $I4 > 0.8$ , the described quasi-periodic oscillation (see Fig. 4B-C) is responsible for sound emission. As soon as another nearby robot also starts signaling, *I6* will become activated (even when the other robot is to the most right side, which is due to the high noise, approx. 30%, of the sound direction detection). Then, as we can see in the bifurcation diagram of Fig. 4D, the quasi-periodic attractor switches to a SFP, thus the oscillation will cease. Because the sound signal of signaling robots lasts only one time step, these oscillations immediately start again in the next time step ( $I6 = 0$ ). This reset mechanism will lead to a synchronization of the signaling among robots which stay together on the food patch (a mechanism very similar to the synchronization of internal neural rhythms described in [16]). That means, if there are many robots on a food patch, they will not produce a continuous sound signal as it is the case for robots of type *A1*. One can see the improvement of the performance with respect to *A1* in Fig. 5. However, the performance of *B1* is not significantly higher compared to *A2* (although *A2* was only evolved with a population size of 10). The next section will discuss this surprising robustness of *A2* against more complicated environmental conditions which did not occur during its course of evolution.

### 3.3 Discussion

The results presented in section 3.1 demonstrate an example of the phenomena called natural drift, which is well known from evolution of biological systems [6]. We started several evolutionary runs for the rather simple task of exploration and obstacle avoidance. The initial conditions were always the same (empty RNN structure, fitness function, environment). Evolution came up with solutions being very different concerning the network's topology, but very similar concerning the observed behavior. Secondly we started again several evolutionary runs, with different RNNs as initial structure which resulted from the preceding step. The task was slightly more complex (i.e., individuals had to find food patches and should stay there). Again, several RNNs resulting from this evolution were selected as

initial structures for the final evolutionary step. This time individuals could develop a communication system in order to cooperate. During this runs we still allowed the emergence of new structural changes within the initial RNNs. And we presented two completely different solutions of emergent communication (*A1* and *A2* individuals) to the same task as a result of different structural changes. Both perform well with respect to the given fitness function and group size with which evolution took place. However, individuals utilizing noise for communication (*A2*) do also perform well in conditions they were not confronted with during their course of evolution. They possess the intrinsic property to be robust against increased population size not only because of the lowered probability of producing unrecognizable continuous sound signals, but also because of the described indirect bidirectional communication behavior. We argue, such solutions can hardly be found when too much predefined knowledge about the fitness function and topology of RNNs is assigned to the evolutionary process.

In section 3.2 we demonstrated another striking result of the described experiments and analysis: the evolutionary adaptivity of RNNs, as dynamical systems for behavior control, to varying environmental conditions, such as the number of interacting robots. We observed how small structural changes within such networks lead to an adaptation of the communication mechanism. During the evolution with small group sizes a context-sensitive communication system developed which is based on a simple two neuron loop that provided period-4 oscillations (*A1* individuals). In this case robots directly communicate the discovery of a food patch unidirectionally to other robots. The behavior was sufficient to improve the performance of the robot group, as we have defined it there. However, by changing the environmental conditions, that is, by increasing the number of robots, it turned out that this strategy was not sufficient enough anymore. In our experiments this was especially due to the physical implementation of the sound perception system. Confronted with this constraint and the larger group size, small structural changes refined this solution. These changes lead to quasi-periodic oscillations of longer periods which then are also synchronized among interacting agents (*B1* individuals). Communication is not longer just a simple stimulus-response action. It is direct bidirectional: the act of signaling also directly influences the signaling behavior of other robots.

## 4 Conclusions

In this paper we showed how communication among interacting autonomous robots emerges by evolving dynamical systems, like recurrent neural networks, in the context of *complete brain-body-environment systems* [4,10]. We have seen that only small structural changes are necessary to alter previously solitary behavior to cooperative behavior among communicating robots.

The presented communication system and collective behavior are indeed rather simple. And often the answer to what is necessary in the Artificial Life approach to place autonomous robots into the same category as animals is to keep up climbing the complexity ladder. However, we agree with [21] in that this is “not the most practical answer .. [because].. seeking such complexity blindly,

by typically restricting the search to achieving more complex behaviors, does not accomplish much". Therefore, our approach is to build simple, and therefore still tractable, models of *minimally cognitive behavior* [22] and to increase their complexity as soon as our understanding improves [23].

There are also some relations to biological systems, we can draw from our observed processes of *evolutionary adaptation*. We setup a communication system in the robot where the perception is limited to signal peaks. This was thought to be a disadvantage for the development of behavior, although it is of great advantage for the interaction with highly noisy real-world environments. In the end, by evolving dynamical systems, we have to argue, that such physical constraints are not necessarily a disadvantage for the development of behavior. Evolution finds solutions which integrate the properties of such physical system very well. On a more abstract phenomenological level, we can compare our artificial system with biological systems, for instance with the synchronized flashing of male fireflies during mating [24]. This astonishing collective process is in general also based on pulse coupled oscillators [25]. Although the process which leads to this synchronization is now well understood, the evolutionary reason why thousand of fireflies synchronize their flashing can only be assumed. One possible explanation is that females are stronger attracted by sudden bright pulses than by a clutter of single flashes. It is also well known that humans, or animals in general, react to sudden changes in their environment stronger than to sustained sensory inputs (it is also known that persistent stimuli can attenuate sensation [26]). Here, we unintentionally put this property into our system. There was no other choice than utilizing only acoustic changes instead of continuous signals. And evolution found solutions able to handle this handicap and adapt to changes in the environment in a very efficient way considering the size of the resulting networks. We argue that such solutions are hardly found when too much predefined knowledge about the topology and the dynamics are assigned to such systems, however compelling this may seem in order to speed up the evolutionary process.

## References

1. Port, R., van Gelder, T.: *Mind as Motion*. MIT Press (1995)
2. Thelen, E., Smith, L.: *A Dynamic Systems Approach to the Development of Cognition and Action*. MIT Press (1994)
3. Pasemann, F.: *Neuromodules: A Dynamical Systems Approach to Brain Modeling*. In: *Supercomputing in Brain Research: From Tomography to Neural Networks*. World Scientific (1995) 331–347
4. Beer, R.D.: The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior* **11**(4) (2003) 209–243
5. van Duijn, M., Keijzer, F., Franken, D.: *Principles of minimal cognition*. *Adaptive Behavior* **14**(2) (2006) in press
6. Maturana, H.R., Varela, F.J.: *The Tree of Knowledge: The Biological Roots of Human Understanding*. rev. edn. Shambhala (1992)
7. Nolfi, S., Floreano, D.: *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. MIT Press (2000)

8. Harvey, I., Di Paolo, E.A., Wood, R., Quinn, M., Tuci, E.: Evolutionary robotics: A new scientific tool for studying cognition. *Artificial Life* **11**(1–2) (2005) 79–98
9. Langton, C.: *Artificial Life: An Overview*. MIT Press (1995)
10. Pfeifer, R., Scheier, C.: *Understanding Intelligence*. MIT Press (1999)
11. Pasemann, F.: Complex dynamics and the structure of small neural networks. *Network : Computation in Neural Systems* **13**(2) (2002) 195–216
12. MacLennan: Synthetic ecology: an approach to the study of communication. In Langton, C., Taylor, C., Farmer, J., Rasmussen, S., eds.: *Proc. Artificial Life II*, Addison-Wesley (1991)
13. Werner, G., Dyer, D.: Evolution of communication in artificial organisms. In Langton, C., Taylor, C., Farmer, J., Rasmussen, S., eds.: *Proc. Artificial Life II*, Addison-Wesley (1991)
14. Quinn, M., Smith, L., Mayley, G., Husbands, P.: Evolving controllers for a homogeneous system of physical robots: Structured cooperation with minimal sensors. *Philosophical Transactions of the Royal Society of London, Series A: Mathematical, Physical and Engineering Sciences* **361**(1811) (2003) 2321–2344
15. Marocco, D., Nolfi, S.: Emergence of communication in teams of embodied and situated agents. In: *Proc. of the 6th Int. Conference on the Evolution of Language*. (2006)
16. Wischmann, S., Hülse, M., Knabe, J., Pasemann, F.: Synchronization of internal neural rhythms in multi-robotic systems. *Adaptive Behavior* **14**(2) (2006) in press
17. Pasemann, F.: Dynamics of a single model neuron. *International Journal of Bifurcation and Chaos* **2** (1993) 271–278
18. Hülse, M., Wischmann, S., Pasemann, F.: Structure and function of evolved neuro-controllers for autonomous robots. *Connection Science* **16**(4) (2004) 249–266
19. Dieckmann, U.: Coevolution as an autonomous learning strategy for neuromodules. In: *Supercomputing in Brain Research: From Tomography to Neural Networks*. World Scientific (1995) 427–432
20. Gammaitoni, L., Hänggi, P., Jung, P., Marchesoni, F.: Stochastic resonance. *Reviews of Modern Physics* **70**(1) (1998) 223–287
21. Di Paolo, E.A.: Organismically-Inspired Robotics: Homeostatic Adaptation and Teleology Beyond the Closed Sensorimotor Loop. In: *Dynamical Systems Approach to Embodiment and Sociality*. Advanced Knowledge International (2003) 19–42
22. Beer, R.D.: Toward the evolution of dynamical neural networks for minimally cognitive behavior. In: *From Animals to Animats 4: Proc. of the 4th Int. Conference on Simulation of Adaptive Behavior*. (1996) 421–429
23. Beer, R.D.: Arches and stones in cognitive architecture. *Adaptive Behavior* **11**(4) (2003) 299–305
24. Camazine, S., Deneubourg, J.L., Franks, N.R., Sneyd, J., Theraulaz, G., Bonabeau, E.: *Self-Organization in Biological Systems*. Princeton University Press (2001)
25. Strogatz, S.H., Stewart, I.: Coupled oscillators and biological synchronization. *Scientific American* **269**(6) (1993) 102–109
26. Bays, P.M., Flanagan, J.R., Wolpert, D.M.: Attenuation of self-generated tactile sensations is predictive, not postdictive. *PLoS Biology* **4**(2) (2006)