

Cognitive conditions to the emergence of sign interpretation in artificial creatures

Angelo Loula^{1,2}, Ricardo Gudwin² and João Queiroz^{3*}

¹ Informatics Area, Department of Exact Sciences, State University of Feira de Santana (UEFS), Brazil

² Department of Computer Engineering and Industrial Automation, School of Electrical and Computer Engineering, State University of Campinas (UNICAMP), Brazil

³ Institute of Arts and Design, Federal University of Juiz de Fora (UFJF), Brazil
*queirozj@pq.cnpq.br

Abstract

Although the emergence of communication has been the topic of various Artificial Life experiments, the study of underlying representational processes finds little discussion. We have previously differentiated between symbolic and indexical interpretation and proposed that symbolic interpretation may act as a shortcut to cognitive traits already acquired. Here we evaluate conditions of this acquired cognitive trait for the emergence of different modalities of sign interpretation. Results show that symbolic processes may act as a cognitive shortcut to a previous acquired cognitive competence even if minimally functional or initially not available.

Introduction

Computational simulation approaches, such as Artificial Life experiments, are considered to have an important role in the study and modeling of general semiotic processes (see Christiansen and Kirby, 2003, Noble et al., 2010, Cangelosi and Parisi, 2001, Steels, 2003). Communication, vocabulary, grammar are among the processes that have been studied by this synthetic approach (for a review, see Nolfi and Mirolli, 2010, Wagner et al. 2003). In these experiments, semiotic processes are simulated in a social context, involving multiple agents. The process in focus is not pre-defined, but it rather emerges during and by means of agents' interactions. As the main form of interaction between agents, in most of these synthetic experiments, communication has, particularly, been a significant research subject. It depends on the production of representations (by an utterer) and the interpretation of them (by an interpreter). Nevertheless, we find little discussion around representation processes underlying communication such as the types of representations involved and how they can represent something to the agents. If agents communicate, the underlying representational processes are an essential issue to be addressed.

We have previously modeled the emergence of two different types of representational processes (symbols and indexes) and how they emerge in a community of simulated creatures (Loula et al., 2010a). We proposed that a symbolic interpretation process can act as a cognitive shortcut to a cognitive competence that is hard to acquire. Here we propose to assess further this hypothesis and evaluate

cognitive conditions to the emergence of interpretation processes, varying availability and reliability. We apply the same scenario previously used, which involves empirical constraints from studies of animal communication and also theoretical constraints from Peircean pragmatic theory of signs.

In the next section, we review related work in the context of the emergence of communication in Artificial Life research. Next, we briefly describe the theoretical principles and biological motivations that guided our experimental design. We then describe the experiment involving the emergence of different interpretational processes in communication events. Results are presented next, summarizing previous results and exhibiting news ones on the conditions for the emergence of sign interpretation. We discuss achievements and draw conclusions and future directions, in the end.

Related work

The simulation of the emergence of communication is the topic of various works, but discussions on the underlying semiotic processes finds little space in such literature. Therefore, we will review two representative works that deals with the emergence of communication that are relevant in the context of this work.

Robots were evolved by de Greeff and Nolfi (2010) to execute a navigation task in which two robots had to exchange places in two target areas. The robots could use wireless sensors for an 'explicit signal' communication or they could use their spatial position as an 'implicit signal'. At the end of an evolution process of neural networks that control the robots, de Greeff and Nolfi (2010) described that the robots were able to use 2 or 3 explicit signals to execute the proposed task, but also used one implicit signal to achieve that. They state that explicit signals codify certain conditions in which the emitter robot finds itself and that the implicit signal is a visual perception of the position of the other robot, and that each signal produces a different reaction. Signals are said to be deictic, dependent of spatial-temporal context, but there was no further discussion on what and how robots representationally interpret such signals.

In an experiment with artificial creatures in a grid world, Cangelosi (2001) simulated the emergence of communication systems to name edible and poisonous mushrooms. He had relied on biological motivations to define a food forage goal for the creatures. He proposed the emergence of different modalities of representations in this experiment on the evolution of communication. To classify communication systems, Cangelosi (2001) differentiated signals, with direct relation with world entities, from symbols, also with relation with world entities but also related to other symbols. In his experiments, neural networks were both evolved and trained in various tasks, and, at the end, a shared communication system emerged, involving signals and symbols, according to Cangelosi. But he did not describe how these signals and symbols were interpreted by the creatures and what they actually represented.

Other works have also studied the emergence of communication traits and the acquisition of vocabulary or language among artificial agents (see Nolfi and Mirolli, 2010, Wagner et al. 2003). Nevertheless, we have not found works that have studied the emergence of different types of interpretations processes and differentiated the interpretation processes that emerged.

Theoretical and Empirical Constraints

Synthetic experiments such as Artificial Life ones are heavily influenced by theoretical principles and biological motivations, and that such background should be an essential part of any synthetic experiment (Parisi, 2001, see also Noble 1997, Loula et al., 2010b). Theoretical principles and biological motivations act as requirements and constraints during the design of the experiments, and influences modeling on different degrees depending on how it constrains the model being built and what decisions it leaves to the experimenter.

To model the emergence of communication processes based on different types of representation, it is certainly important to look at theoretical models and principles, and also look for biological motivations, and avoid arbitrary or naïve assumptions about the underlying processes.

Sign-mediated processes, such as the interpretation of representations in communicative contexts, show a remarkable variety. A basic typology (and the most fundamental one), proposed by Peirce (1958; see Short 2007), differentiates between iconic, indexical, and symbolic processes. Icons are signs that stand for their objects by a similarity or resemblance, no matter if they show any spatio-temporal physical correlation with an existent object. In this case, a sign refers to an object in virtue of a certain quality which is shared between them. Indexes are signs which refer to their objects due to a direct physical connection between them. Since (in this case) the sign should be determined by the object (e.g. by means of a causal relationship) both must exist as actual events. Spatio-temporal co-variation is the most characteristic property of indexical processes. Symbols are signs that are related to their object through a determinative

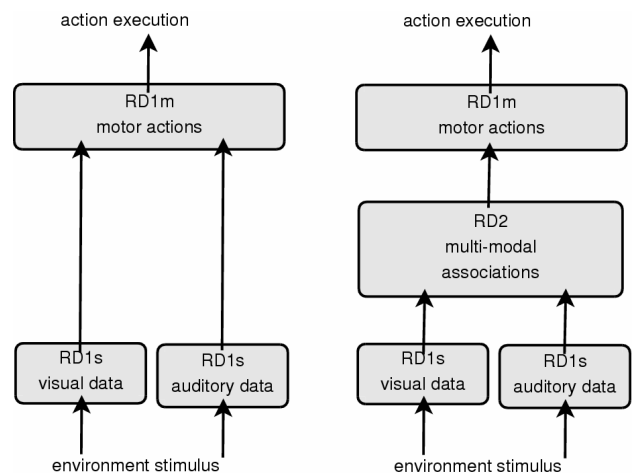


Figure 1: Cognitive architectures for representations interpretations. Left: Type 1 architecture, RD1s are connected directly to RD1m. Right: Type 2 architecture, data from visual RD1s and auditory RD1s can be associated in RD2 before connecting to RD1m.

relation of law, rule or convention¹. A symbol becomes a sign of some object merely or mainly by the fact that it is used and understood as such by the interpreter, who establishes this connection.

Communication is a process that occurs among natural systems and as such we can employ empirical evidences on building our synthetic experiment. Animals communicate in various situations, from courtship and dominance to predator warning and food calls (see Hauser, 1997). And following Peirce's definition of symbols, many animals can actually be capable of communicating by means of symbols (Ribeiro et al., 2007).

To further explore the mechanisms behind communication, a minimum brain model can be useful to understand what cognitive resources might be available and process underlining certain behaviors. Queiroz and Ribeiro (2002) described a minimum vertebrate brain for vervet monkeys predator warning vocalization behavior (Seyfarth et al 1980). It was modeled as being composed by three major representational relays or domains: the sensory, the associative and the motor. According to such minimalist design, different first-order sensory representational domains (RD1s) receive unimodal stimuli, which are then associated in a second-order multi-modal representation domain (RD2) so as to elicit symbolic responses to alarm-calls by means of a first-order motor representation domain (RD1m).

Our objective is to model the emergence of indexical and symbolic interpretation competences, so the first step is to specify the requirements for each and also how to recognize each of them in the experiment. Indexical interpretation is a reactive interpretation of signs, such that the interpreter is directed by the sign to recognize its object as something spatio-temporally connected to it, so for our creatures to have

¹ Differently from Cangelosi's (2001) definition of symbol, based on Deacon's approach (1997), Peirce (1958) did not require symbols to be related to each other to be called symbols.

this competence, they must be able to reactively respond to sensory stimulus with prompt motor answer. In the minimum brain model, this corresponds to an individual capable of connecting RD1s to RD1m without the need for RD2. But a symbolic interpretation undergoes the mediation of the interpreter to connect the sign to its object, in such a way that a habit (either inborn or acquired) must be present to establish this association. Thus, in symbolic interpretation, RD2 must be present once it is the only domain able to establish connections between different representation modes. Thus, our artificial creatures must be able to receive sensory data, both visual and auditory, in its respective RD1s, that can be connected directly to RD1m, defining motor actions (Type 1 architecture), or connected to RD1m indirectly, through the mediation of RD2, that associates auditory stimulus to visual stimulus acting as an associative memory module (Type 2 architecture) (see figure 1). To evaluate what conditions might elicit each response type – indexical or symbolic –, we implemented these two possible cognitive processing paths as mutually exclusive paths: either the creature responds to auditory events indexically and reactively responds with motor actions, or the creatures responds to auditory events symbolically and associates them with a visual stimulus and responds as if that was really seen. For an external observer, who only watches the information available to the creature and its motor responses, it may not be possible to see changes in the interpretation process. But the underlying mechanisms behind each semiotic process are qualitatively different.

The experiment

The scenario to test the conditions for the emergence of semiotic processes is inspired by food foraging behavior of animals. One way animals cooperate in such task is by vocalizing for food quality, recruiting other group members to feed. Inspired by such biological motivation, we simulate a scenario of artificial creatures evolved to collect resources in a virtual environment.

Lower quality resources are scattered throughout the environment and a single location receives highest quality resources. One creature (vocalizer) is placed in this high quality resource position, vocalizing a sign continuously. At the start of simulation, the other creatures (interpreters) do not know how to respond appropriately to sensory inputs and neither recognizes the sign vocalized as a sign. But an evolutionary process of variation and selection is applied, allowing the evolution of individuals to better accomplish the task of resource foraging. During the evolutionary process, for each start-up conditions, we observe the emergence of indexical or symbolic interpretation for the vocalizations.

The environment is a 50 by 50 grid world (figure 2) and there are 20 random positions with only one resource unit each. There is also one position with 500 resource units, where an immovable vocalizer creature is placed. The vocalizer's sole behavior is to produce a single vocal sign, reproduced at every instant. Fifty interpreter creatures are randomly placed in this grid.

Interpreter creatures are capable of visually sensing food up to a distance of 4 positions and sensing vocalizations up to a distance of 25 positions. This sensory range difference models




Figure 2: The grid environment. Creatures are blue circles, low quality resource positions are in green cells, and high quality one in the cyan cell in the center.

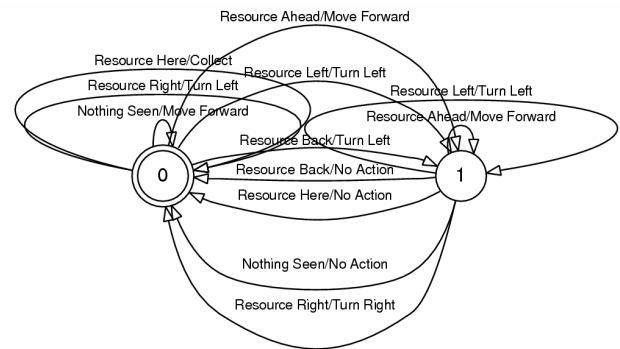


Figure 3: State diagram of a sample FSM that controls the creatures. The circles are states and a double circle marks the start state. Arches represent transitions and are labeled according to the sensory event and the action to take over when that event occurs. This FSM has only visual inputs and 2 states to simplify the diagram, but there can be more arches for auditory inputs (vocalization and its position) and up to 4 states.

an environment where vision is limited by the presence of other elements such as vegetation, restraining far vision such as in an open field. These creatures can either see a resource and its position or hear a vocalization and its position, if any of them is within range.

Interpreters are controlled by finite state machines (FSM), with up to 4 states (see figure 3). An FSM was chosen as the control architecture because the analysis of how it is functioning is quite simple and direct, permitting direct identification of the processes underlying the creatures' cognition. Input events to the FSM include 5 visual events for resource in 5 different positions (ahead, left, right, back, or same position), 4 auditory events for vocalization in 4 positions (ahead, left, right, back), and 1 event for nothing seen or heard. Outputs from FSM can be one of the 5 motor

actions for creatures: move forward, turn left, turn right, collect resource, or do nothing.

The creatures can respond to visual inputs with one of the motor actions, and can also respond to auditory input with a direct motor action (a reactive, indexical process) (Type 1 architecture). Alternatively, before an input is sent to the FSM, they can also choose to establish an internal association between the heard stimulus and the visual representation domain (Type 2 architecture). This internal association links what is heard with the view of a collectible resource, i.e. the creature can interpret the sign heard as a resource and act as if the resource was seen. As a result, an auditory input is exchanged by an equivalent visual input and the FSM is executed with that input. Additionally, the creature may also ignore the sign heard, interpreting it as nothing and acting as if no sensory data was received.

At start, creatures are controlled by randomly constructed FSMs, and are all placed at random in the environment. They are allowed to collect resources for 10 trials of 100 iterations each trial. Creatures collect resources by executing the specific action, removing one unit from the resource at each time step. When no more units are available at the resource, it disappears.

Creatures evolution

At the all trials, the 10 best creatures in the foraging task (those that collected the most resource units) are selected to create next generation. These 10 individuals are copied to the next population and the 40 remaining individuals are a product of mutations and crossovers of the FSMs of the best individuals.

The mutations can be of changing an action in transition, changing the next state after a transition, changing the start state, add a state and remove a state. There can also be a mutation of the cognitive architecture type, as described below. The number of mutations is selected from a Poisson probability distribution with an expected value of 3. The crossover has a 50% chance of occurring and it exchanges states and transitions originating from the selected states between two FSM in a uniform way. All FSM undergo a correction process to fix error that might occur during these operations, such as a transition pointing to a non-existing state.

The experiment runs for 500 generations, normally with two distinct moments. In the first 200 generations (cycle 1), the vocalizer creature is not present and interpreters do not have an auditory sensor, but this first cycle will be omitted in one of the simulation scenarios. In the 300 subsequent generations the vocalizer creature is present and interpreters are able to hear (cycle 2).

At the start of cycle 2, all creatures are set to ignore the vocalizations, as if it was not relevant, however, there is also a small mutation probability for changing the type of response to vocalizations. These can be of reacting to them by moving towards the resource, or to linking it with the view of a resource. This corresponds to a change to a Type 1 cognitive architecture (indexical) or to a Type 2 cognitive architecture

(symbolic). The probability of going from Type 1 architecture to Type 2 architecture is lower than the other way around to simulate the fact that such a significant cognitive change is not that easy to happen.

We expect that creatures adapt to the foraging task by responding to the auditory input of vocalizations. Since they can not see the high quality resource position, they must rely on the vocalization to guide their movements in this direction. We are interested in observing the overall adaptation process to the foraging task, and are specially focused on the type of interpretation process, related to the cognitive architecture type, that might result.

Results

In previous work, we have run two initial experiments to evaluate the emergence of either an indexical interpretation or a symbolic interpretation of vocalizations (Loula et al., 2010a). Such experiments involved 2 cycles as described above, varying the way motor actions needed to be coordinated. In the first experiment, creatures just had to have the specified action as output of the FSM to execute that action. In this scenario, we observed that indexical interpretation was the competence acquired by creatures to deal with communication, with direct association between auditory signs and motor actions. But in a second experiment, for motor actions to be executed, the creatures needed to first output a null action before any movement action, that way it would be harder to learn motor coordination. In this alternative scenario, symbolic interpretation was the emerging competence, instead of an indexical one like it happened in the previous case. We made the hypothesis that acquiring symbolic competence would act as a *cognitive shortcut*, by reusing a previously acquired ability in cycle 1: to appropriately respond to visual data with motor actions. We proposed that a symbolic interpretation process can happen if a cognitive trait is hard to be acquired and the symbolic interpretation of a sign will connect it with another sign for which the creature already has an appropriate response.

Single cycle scenario

In face of the fact that there should be a previously acquired competence for symbolic interpretation to benefit from, a subsequent question is to ask what would happen to sign interpretation if such previous competence is not present. If the creature does not respond in a proper manner to visual input, a cognitive shortcut to this uncoordinated competence would not help the foraging success. As cycle 1 acts as a first step in which creatures are dedicated to learn visual-motor coordination, we removed this cycle in a new scenario, in which the simulation begins in cycle 2 with the vocalizer at the center and interpreter creatures able to hear but starting with random FSMs. The need to first output a null action before any movement action remains, so it is hard to learn motor coordination. Figure 4 shows the performance of creatures in foraging and the type of interpretation used.

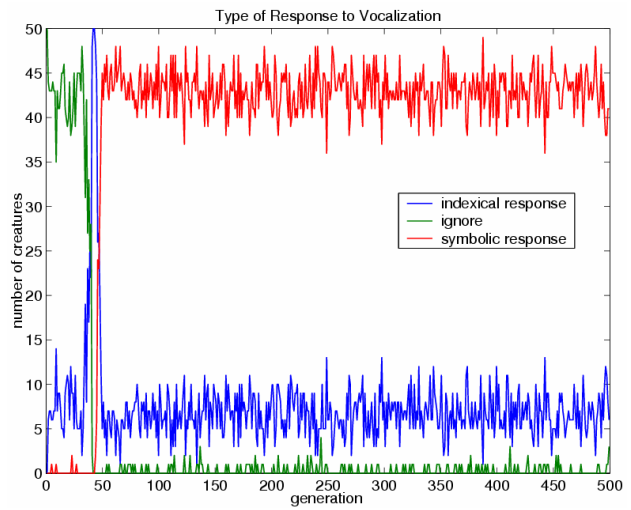
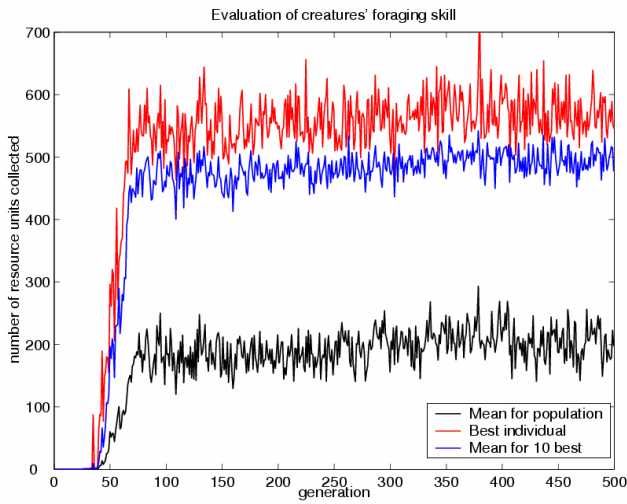


Figure 4: Evaluation of foraging task and type of response to vocalizations along the generations for the one cycle only experiment.


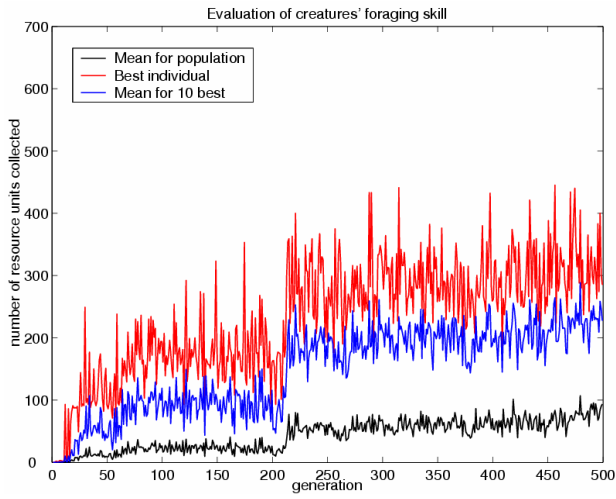


Figure 5: Evaluation of foraging task and type of response to vocalizations along the generations for the 20% failure experiment.

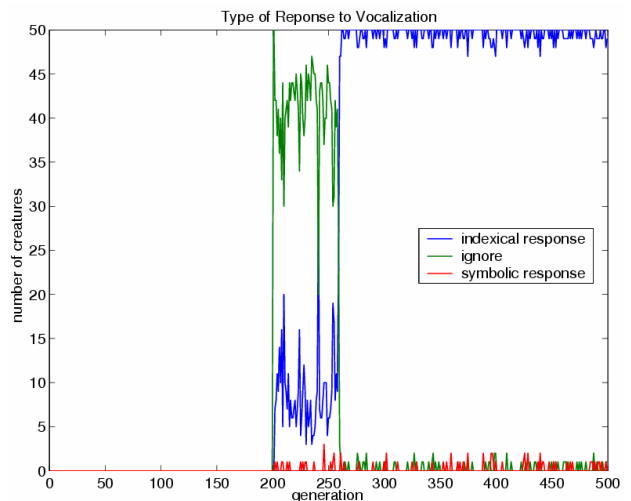
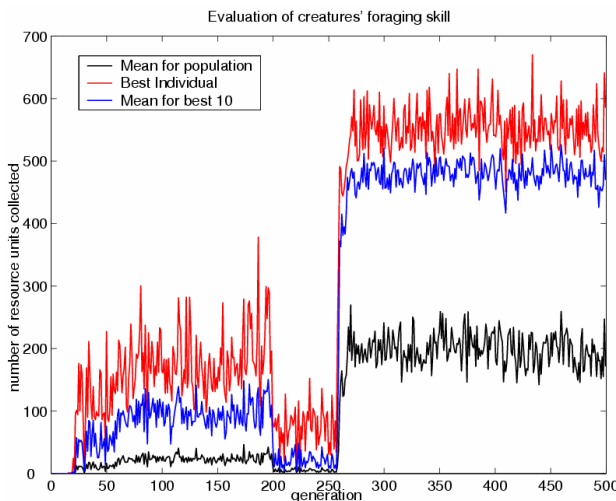


Figure 6: Evaluation of foraging task and type of response to vocalizations along the generations for the 50% failure experiment.

As we can see from the graphs, the experiment had three phases. At first, no resources were collected and creatures opted to ignore signs produced by the vocalizer. Then there was a transition phase, where the amounts of resources increased rapidly along generations and creatures gave up ignoring signs and started an indexical interpretation of them. Then creatures turn to a symbolic interpretation of sign and the amount of resources collected further increases and then stabilizes. To better understand what happened in such transitions, the FSMs of creatures have to be further detailed.

From the first generation until generation 25, creatures did not demonstrate any motor coordination and were not able to collect resources, and *most creatures just ignored signs*. In generation 26, one creature was able to move forward and collect when a resource was in front of it, but it still ignored signs.

This remained the same until generation 39, where one creature was able to turn right when a resource was seen at right side, and this creature was also *responding indexically to signs*, by going towards the vocalizer when a sign was heard in front of it. By generation 40, half of the population is interpreting signs indexically and the other half is ignoring it. Most of the creatures could move towards a seen resource, but there were still some useless outputs from the FSM, state/transition combinations that would make a creature stop responding effectively, and they still would not move when nothing is seen.

At generation 44, one creature starts *interpreting signs as symbols*, relating the sign heard with the view of a resource, able to collect 67 resource units, while the best performing creature collected 77, but interpreting signs as index. Nevertheless, creatures still had problems in motor coordination. By generation 46, half of the creatures were symbolically interpreting signs, and by generation 50, almost all of them did so. From there on, all 10 best performing creatures used symbolic interpretations (and most of the others too), and the number of collected resources increased rapidly as creatures acquired a best performing FSM, that would always respond effectively to the inputs received.

So even though, there was no cycle dedicated to acquire a previous competence that could be re-used by a symbolic interpretation, the evolution process allowed first for visual-motor coordination to appear before sign interpretation (either as index or symbol) started. Thus, there was at least a little visual-motor competence to be re-used by symbolic interpretation.

Cognitive module malfunction scenario

To further evaluate the way symbolic interpretation acts as a cognitive shortcut, we set up one more scenario. Since there is re-use of a previous acquired cognitive competence, we tested how reliable should this competence be for this new symbolic process to connect to it. The scenario is similar to the one above, but we brought back cycle 1 before cycle 2, so the creatures had time to acquire visual-motor coordination. However, in this reliability test, we introduced a failure chance in the visual-motor coordination after cycle 1, simulating a malfunctioning cognitive module. Given an output from the FSM in response to a visual input, this output (an action) would have a chance of changing to a different

one. If the input is 'Resource Left' and the output from the FSM is 'Turn Left', for example, it could be changed to 'Turn Right'. Outputs that are responses to auditory inputs are not subject to such changes. This way visual-motor coordination would be defective and processes relying on it would be jeopardized.

The first simulation of this malfunctioning in visual-motor coordination applied a 20% chance of output change. The results are presented in figure 5. Compared to a previous experiment with 2 cycles but no malfunction (Loula et al., 2010a), it is possible to notice that the number of collected resources during cycle 1 is similar in both experiments, but in the second cycle it is quite different: while in the previous experiment the best creature collected between 500 and 600 units, in this unreliable module experiment, the best creature collect only around 300 units. This shows that the foraging efficiency has dropped down with the imposed malfunctioning. Looking at the type of response, signs ended up having a *symbolic interpretation*, thus the unreliable visual-motor connections were in fact reused, despite the fact that it was not an efficient module. Comparing with the cited previous experiment, the interpretation type graph is quite similar.

Taking a closer look at simulation outcome, results show that from generation 200 to 210 the foraging performance did not improve. Initially creatures ignored signs, but by generation 202, a few creatures start trying to respond to signs in an indexical manner. These creatures with type 1 architecture, nevertheless, are not able to move towards the vocalizer and still rely in the defective visual-motor coordination. In generation 208, almost all creatures are ignoring signs again.

By generation 210, a symbolic interpreter appeared and it was able to collect more than 200 resource units. Even though visual-motor coordination was degraded, it still performed better than the random actions of a creature trying an indexical response. From this generation on, the number of creatures *interpreting signs symbolically* increased and, by generation 218, almost all creatures followed this type of interpretation.

To further test the effects of a malfunctioning of a cognitive module, the chance of changing actions was increased to 50% in a new simulation, with the expected effect of turning the visual-motor coordination so unreliable that its reuse would be not be possible. Results of this simulation are shown in figure 6.

In this new simulation run, we observe that after cycle 1 the number of collected resources dropped considerably more, to about half of the amount at the end of cycle 1. This was expected since the creatures are using a quite defective control model that is not able to cope with the task of foraging resources efficiently anymore.

Until around generation 250, creatures had this bad performance, but in the meanwhile sign interpretation was varying from ignoring signs to indexical response. The best performing creatures were most ignoring sign, though, indicating that indexical interpreters were not able to successfully respond to signs. One or two creatures with symbolic response were created but disappear right after as its performance was not consistent, due to the dependence on visual-motor coordination.

At generation 258, there appeared a creature with *indexical interpretation*, able to collect 210 resource units. This creature was able to effectively respond to signs by going towards the vocalizer when it was located ahead or to the left. Therefore, this indexical interpreter was able to rely on a direct connection between auditory input and motor actions, and avoided using the faulty cognitive module. The number of collected resource units along generations increased fast, and the best creature (an indexical interpreter) on generation 270 was collecting almost 600 resource units and this performance was consistently kept until the end of simulation. Notice that if we compare the efficiency of creatures in the 20% failure chance simulation with this 50% failure chance simulation, it is clear that even though the second simulation had a worst damage to the visual-motor module, it was able to achieve better performance at the end.

Discussion

In this paper, we continue investigating conditions for qualitative different interpretation processes to emerge in a communicational context. Previous results showed that symbolic interpretation can emerge when the appropriate motor coordination is a hard skill to acquire, and therefore symbolic processes can act as a cognitive shortcut, mapping auditory signs to visual input and reusing visual module mapping to motor actions. Here we test other conditions for this cognitive shortcut to be established.

First we removed the first cycle, when creatures were allowed to acquire visual-motor coordination, which could be reused through a cognitive shortcut. Consequently, adequate auditory and visual responses needed to be acquired at the same time. From this single cycle experiment, it is possible to observe that even though the vocalizer and the hearing sensor were available from start, creatures did not use signs at all in a first moment. It was necessary to first have minimum visual-motor coordination for signs to start being interpreted by creatures. Indexical interpretation was the first attempt as a response to signs. As trying to acquire visual-motor coordination and also a sign-motor coordination is a tough route, the symbolic interpretation diminished this effort and became the dominant strategy.

To further evaluate the cognitive shortcut stability, we imposed a variable malfunctioning to the visual-motor connections. At first, a 20% of changing actions specified by the visual module still conducted to the establishment of symbolic processes, with reuse of a degraded module, but that still allowed a relative increase in foraging efficiency. A higher failure of 50% proved to worsen the performance of the visual control module considerably more, and allowed indexical interpretation of sign to be established, as a way to avoid reusing it. And, even though symbolic processes were established in the 20% failure scenario, it seems that creatures got trapped in a 'local maximum' performance, as the foraging efficacy of creatures in the 50% failure scenario was much better.

Conclusion

Communication necessarily involves an utterer, who produces a sign, conveyed to an interpreter, in whom the sign produces its effect. And signs can be of different types according to the way it is connected during interpretation process to its referent. We proposed that, for two types of signs – indexes and symbols – to be interpreted, different cognitive paths had to be followed, one with direct mapping of signs to motor actions (indexical interpretation) and another one with a mapping of signs into another representation form (symbolic interpretation) and then to motor actions.

We proposed that a cognitive shortcut can be established by symbolic interpretation processes, by establishing bridges to reuse previous acquired competences. We confirmed here that the cognitive module to which the symbolic interpretation is connecting to must be already established, otherwise there is no advantage in such connection. But it does need to be fully functional, as minimal visual-motor coordination is sufficient to begin a symbolic interpretation process, according to the single cycle experiment, and even a moderately damaged module can also be reused as a cognitive shortcut.

Even further investigations on differentiating indexical and symbolic processes have to be done. Other aspects and conditions should be tested to better understand what leads sign interpreters to each of them, for example, how can an agent handle both of them at the same time or how does other cognitive competences influence this process. We expect that the discrimination of these semiotic processes and the cognitive apparatus necessary for each of them will bring forth more discussion on representation process in experiments on the emergence of communication and language.

References

- Cangelosi, A. (2001) Evolution of communication and language using signals, symbols and words. *IEEE Transactions on Evolutionary Computation* 5(2), 93-101
- Cangelosi, A. and Parisi, D. (2001) Computer simulation: A new scientific approach to the study of language evolution. In Cangelosi, A. and Parisi, D., editors, *Simulating the Evolution of Language*, chapter 1, pages 3–28. Springer Verlag.
- Christiansen, M. H. and Kirby, S. (2003) Language evolution: consensus and controversies. *Trends in Cognitive Sciences*, 7(7):300–307.
- de Greeff, J. and Nolfi, S. (2010) Evolution of implicit and explicit communication in mobile robots. In *Evolution of Communication and Language in Embodied Agents*, pages 179–214. Springer Verlag.
- Deacon, T. (1998) *The Symbolic Species: The Co-Evolution of Language and the Brain*. W. W. Norton & Company, New York.
- Hauser, M. D. (1997) *The Evolution of Communication*. MIT Press, Cambridge, MA.
- Loula, A., Gudwin, R., and Queiroz, J. (2010a) On the emergence of indexical and symbolic interpretation in artificial creatures, or What is this I hear? In Fellermann, H., et al., editors, *Artificial Life XII*, pages 862–868. MIT Press.
- Loula, A., Gudwin, R., Ribeiro, S., and Queiroz, J. (2010b) On Building Meaning: a biologically-inspired experiment on symbol-based communication. In Cutsuridis, V., Hussain, A., Barros, A.K., Aleksander, I., Smith, L., and Chrisley, L., editors. *Brain Inspired Cognitive Systems*, pages 77-94. USA: Springer.
- Noble, J. (1997) The scientific status of artificial life. In *Fourth European Conference on Artificial Life (ECAL97)*, pages 28-31. MIT Press/Bradford Books.

- Noble, J., De Ruiter, J.P., and Arnold, K. (2010) From monkey alarm calls to human language: How simulations can fill the gap. *Adaptive Behavior* 18(1):66–82.
- Nolfi, S. and Miralli, M. (2010) *Evolution of Communication and Language in Embodied Agents*. Springer.
- Parisi, D. (2001) *Simulazioni - la realtà rifatta nel computer*. Il Mulino, Bologna.
- Peirce, C.S. (1958) *Collected Papers of Charles Sanders Peirce*. Harvard University Press, Cambridge, Mass.
- Queiroz, J. and Ribeiro, S. (2002) The biological substrate of icons, indexes, and symbols in animal communication: A neurosemiotic analysis of vervet monkey alarm calls. In Shapiro, M., editor, *The Peirce Seminar Papers 5*, pages 69–78. Berghahn Books, New York.
- Ribeiro S., Loula, A., Araújo I., Gudwin, R., and Queiroz, J. (2007) Symbols are not uniquely human. *Biosystems* 90:263-272.
- Seyfarth, R., Cheney, D., and Marler, P. (1980). Monkey responses to three different alarm calls: Evidence of predator classification and semantic communication. *Science*, 210, 801–803.
- Short, T.L. (2007) *Peirce's Theory of Signs*. Cambridge University Press.
- Steels, L. (2003) Evolving grounded communication for robots. *TRENDS in Cognitive Sciences*, 7(7): 308-312.
- Wagner, K., Reggia, J. A., Uriagereka, J., and Wilkinson, G. S. (2003) Progress in the simulation of emergent communication and language. *Adaptive Behavior*, 11(1):37-69.