Emergent communication for coordination in teams of embodied agents

Kevin Godin-Dubois, Sylvain Cussat-Blanc and Yves Duthen

University of Toulouse, IRIT - CNRS UMR 5505, 2 rue du Doyen Gabriel Marty, 31042 Toulouse, France
{kevin.dubois, sylvain.cussat-blanc, yves.duthen}@irit.fr

Communication is ubiquitous in both the biological (growl, songs, natural language) and the artificial (e.g. wifi and rfid) worlds. Indeed, emergence mechanisms are much studied in the former case e.g. with respect to how they might have first emerged (Suzuki et al., 2021), how to reproduce well-known phenomenon such as the bee's waggle dance (Campos and Froese, 2020) or how agents might learn to segregate between frequencies (Eldridge and Kiefer, 2020) or words (Dubova, 2021). However, a more rational approach is taken with respect to the latter case with little space for emergent approaches.

In this abstract, we expose a refined version of our framework to study emergent neural structures in embodied virtual agents. These can communicate, vocally, with each other through elementary signals thereby allowing arbitrary complex messages. In the context of tag-team physical competition, we observed structured communications despite no *explicit* evolutionary pressure to do so. We further show that this exchange is not random but, instead, is correlated with a joint attack strategy.

Methodology

As in previous work, the use of structured communication to exchange meaningful information is the result of intricate interactions between multiple elementary components. Most have already been described in (Godin-Dubois et al., 2021) and, as such, will only be given a brief overview. However, additional morphological parameters have been introduced and will be presented in more details alongside the experimental protocol.

Succinctly, a creature's genome is composed of two major components: morphological and neural. The former case encompasses the visual parameters as well as the spline-encoding values used to produce the blue structures of figure 1. Vision is based on a (genetically controlled) number of rays perceiving the color at the point of impact thereby mimicking the biological retina. Similarly, audition and vocalization are based on frequencies and motion is performed by either forward/backward (body) or rotational (arms) motors. Finally, the splines are generated by interpolating cubic


Figure 1: Splinoid morphology showing the positions of input neurons (in red) and output neurons (in blue). The environment is perceived by RGB ray-casting (eyes), frequency (ears) and contact (touch) sensors. Actions consists of body motion (motors), arms rotation (articulations) and vocalisation (mouth). The blue structures are produced by gene-controlled cubic splines.

bézier curves which are then dilated, to produce a surface, before being quantized into elementary shapes for use in the underlying physics engine.

The neural genetic component is a Composite Pattern Producing Network (CPPN), an n-dimensional function used by ES-HyperNEAT (Risi and Stanley, 2012) to indirectly encode the position, density and connectivity of hidden neurons. With these algorithms, it is thus possible for a neural topology to emerge as a result of environmental and evolutionary pressures. Indeed, as all of an agent's perceptions and actions are the result of the activity of dedicated neurons, this leads to the CPPN being directly responsible for a creature's intrinsic behavior.

The associated experiment investigated behavioral and neural strategies for physical confrontations in a small enclosed arena. To allow for a sufficiently large fitness landscape, each creature can "grow" up to two symmetrical pairs of static splines, for defensive purposes, and a pair of single-jointed arms, useful both for offensive and defensive actions. The fitness function straightforwardly rewarded inflicting more damage to the opponent(s) than were received.
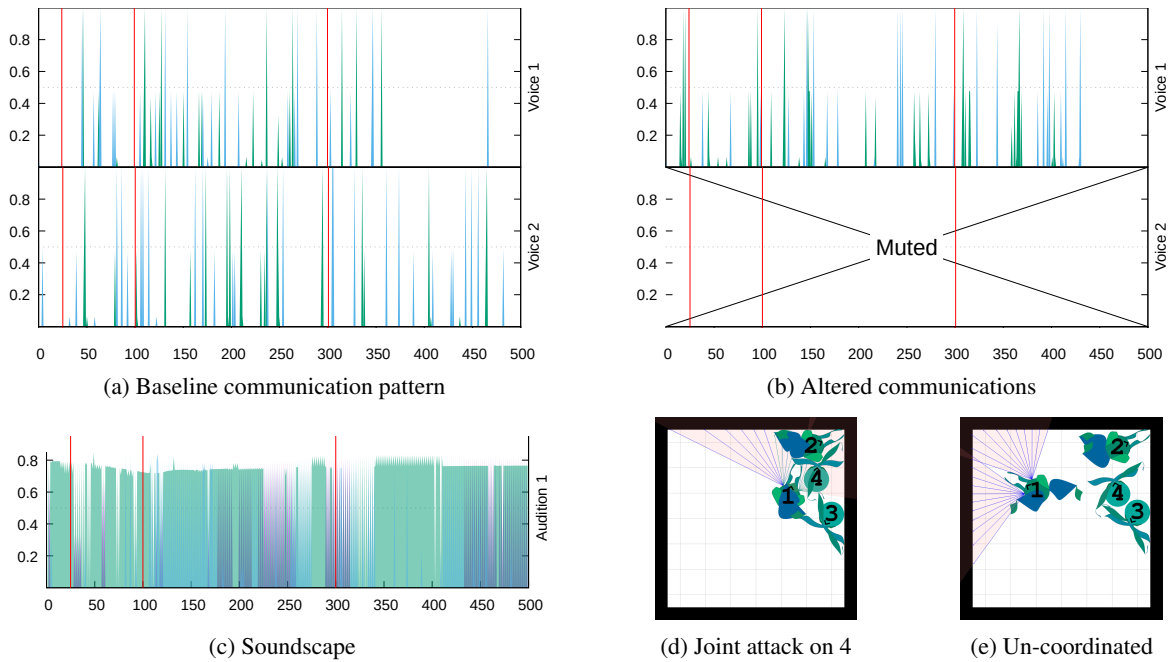
Figure 2: Communication for coordinated action in noisy conditions. (a) In standard conditions, both subjects (1,2) communicate in alternating bursts. (b) When 2 is forcefully muted, the communication pattern of 1 is different. (c) The opponents (3,4) broadcast continuously on all three channels. (d) In normal conditions 1,2 converge on the same target while muting 2 results in a coordination failure (e), drastically reducing the team's efficiency.

On the use of communication

The resulting creatures displayed numerous emergent abilities of various level of complexity (tracking, evasion, telemetry...) all aiming at maximizing gains in this zero-sum game. Here we single out vocal communication, as illustrated by figure 2, for a co-evolutionary tag-team champion.

This specific pair, which incidentally possesses the densest ANN across all concurrent runs, has a parsimonious approach to communication. This can be opposed to the behavior of their opposing team, with which they share a co-evolutionary history, which consists of a mostly continuous broadcast across the different channels (fig. 2c). Such a strategy is most common amongst the various replicates as it only requires constant activation of the associated output neuron. On the other hand, the pattern displayed in 2a requires much finer control of both the channel and magnitude outputs: both subjects seem engaged in a structured exchange via spikes emitted on alternating channels.

To better investigate whether this form of communication has practical use in the context of tag-team competition, we perform two alternate evaluations with one of the teammates being forcefully muted. This results in drastic modifications of the communication pattern, as illustrated in figure 2b when muting 2. The opposite alternative induces similar responses in terms of vocal activity. Most interesting, however, is the impact of the second teammate's muteness on the performance of the team.

In baseline conditions, the subjects jointly attack one member of the opposing team (fig. 2d). Given that the fitness functions only takes into consideration the health of the worst injured individual in each team, this is an efficient method to maximize gains. When removing the capacity to vocalize, however, we notice that the strategy is drastically impaired: 1 does not join the attack on 4 (fig. 2e). We can track this behavioral divergence to the first sound spike as preventing its emission is sufficient to observe a similar divergence. From this, we can gather that the vocalization pattern emitted by this pair of creatures is not random but, instead, seem to rely on a protocol initiated by a mandatory handshake.

It follows that, while not actively selecting for it, communication emerged as a partial solution to the physical confrontation problem. Moreover, the example detailed here shows two extremely promising characteristics: parsimony and resilience to noise. However, while we showed that both individuals successfully communicated with one-another, there is little ground to give a semantic interpretation. In future work, we will investigate communication in a formal context where the information to convey can be directly interpreted by the human observer. Thus, by gradually increasing the complexity of the data to exchange we will be in position to study the construction of a proto-language, potentially applicable to human-intelligible communication in populations of micro-robots.

References

Campos, J. I. and Froese, T. (2020). From embodied interaction to compositional referential communication: A minimal agent-based model without dedicated communication channels. *Proceedings of the 2019 Conference on Artificial Life: How Can Artificial Life Help Solve Societal Challenges, ALIFE 2019*, pages 79–86.

Dubova, M. (2021). Growing Opportunities to Grow: Toward Open-Ended Multi-Agent Communication Learning.

Eldridge, A. and Kiefer, C. (2020). Toward a synthetic acoustic ecology: Sonically situated, evolutionary agent based models of the acoustic niche hypothesis. *ALIFE 2018 - 2018 Conference on Artificial Life: Beyond AI*, pages 296–303.

Godin-Dubois, K., Cussat-Blanc, S., and Duthen, Y. (2021). Spontaneous modular NeuroEvolution arising from a life/dinner paradox. In *The 2021 Conference on Artificial Life*, page 95, Cambridge, MA. MIT Press.

Risi, S. and Stanley, K. O. (2012). An Enhanced Hypercube-Based Encoding for Evolving the Placement, Density, and Connectivity of Neurons. *Artificial Life*, 18(4):331–363.

Suzuki, R., Seki, R., Banno, T., Kawai, K., and Arita, T. (2021). An Artificial Creature Approach to the Origin of Acoustic Communication. 2(i).

Williams, P. L., Beer, R. D., and Gasser, M. (2008). Evolving referential communication in embodied dynamical agents. *Artificial Life XI: Proceedings of the 11th International Conference on the Simulation and Synthesis of Living Systems, ALIFE 2008*, pages 702–709.