

2019

How communication evolved: reflections on the evolution process of an in silico embodied robotic system

Xiaoqing Xu
Vassar College

Follow this and additional works at: https://digitalwindow.vassar.edu/senior_capstone

Recommended Citation

Xu, Xiaoqing, "How communication evolved: reflections on the evolution process of an in silico embodied robotic system" (2019). *Senior Capstone Projects*. 849.
https://digitalwindow.vassar.edu/senior_capstone/849

This Open Access is brought to you for free and open access by Digital Window @ Vassar. It has been accepted for inclusion in Senior Capstone Projects by an authorized administrator of Digital Window @ Vassar. For more information, please contact library_thesis@vassar.edu.

How Communication Evolved: Reflections on the Evolution Process of an *In Silico* Embodied
Robotic System

Xiaoqing Xu

Vassar College

April 24, 2019

Senior thesis submitted in partial fulfillment of the
requirements for major in Cognitive Science

First reader: Josh de Leeuw

Second reader: John Long

Research in biological communication has often analyzed evolved communication systems, but fewer have investigated how evolutionary history of a species influences the evolution of communication-related traits. The current study speculates that agents must develop cognitive behaviors related to conditions of interest before they can meaningfully communicate about these conditions. In an evolutionary robotics model known to exhibit communication behaviors after selection for performance in a coordinating foraging task, genetic and behavioral factors contributing to fitness were analyzed at both the initial and the final stages of evolution. Results showed that the selection for exploring the environment and staying in target area drove the beginning of the evolution process while navigation and coordination using IR and communication sensors were selected later on. Surprisingly, during initial stages, communication channels contributed to fitness by facilitating exploration rather than coordination. These findings demonstrated that the evolution of non-communication behaviors might have served as precursors and constraints for the evolution of communication.

Table of Contents

I.	Introduction (p. 4):
	A. The Computational Model of Communication (p. 5)
	B. The Biological Evolution of Communication (p. 9)
	C. Studying Communication in Evolutionary Robotics (p. 14)
	D. Objectives for the Current Project (p. 18)
II.	Methods (p. 22)
	A. The Evolutionary Robotics System (p.22)
	B. Experiment Setup (p.31)
	C. Data Analysis (p. 33)
III.	Results (p. 44)
	A. Main Effects on Fitness of Access to Communication (p. 44)
	B. The trend of Fitness Throughout Evolution (p.45)
	C. Selection Gradient Analyses Comparisons (p. 47)
	D. Selection of Different Behaviors At Different Evolutionary Times (p. 50)
	E. Correspondence Between Traits and Fitness Across Evolution (p. 56)
	F. Qualitative Trial Analysis of Selective Groups (p. 60)
IV.	Discussion (p. 64)
	A. Replication of Marocco and Nolfi (2007) (p.64)
	B. Functional Behavior Level (p. 65)
	C. Motor Behaviors Serves as Precursors (p.67)
	D. Evolution of Signal Avoidance Behavior (p. 69)
	E. Multiple Realizability Across Levels (p.71)
V.	Reference (p. 73)

Introduction

There is little doubt that communication is something that we, as members of the human species, engage in every day and almost constantly. You are doing it right now, as every sentence in this thesis is me trying to *communicate* some messages to you, and by reading them, you engage in this act of communication as the receiver of these messages. But what is communication? How has it come about, and why?

This thesis project is driven by the motivation to better understand these questions. Building on the basis of a previous study by Marocco and Nolfi (2007) that has demonstrated that functional communication systems evolve after adaptive evolution in a specific evolutionary robotics model, the current project takes a closer look at the process of how communication systems evolve in this model. In Marocco and Nolfi (2007), after 100 generations of selection through a foraging tasking that may be facilitated via the ability to coordinate, robotic teams that had access to sensors dedicated for communication exhibited significantly higher fitness scores than other agents. This result was replicated for ten different initial randomly generated populations in their study, which is strong evidence that the evolution of the resulting communication system was not accidental. Marocco and Nolfi (2007) also analyzed the communication behaviors of the best performing group of agents in the final generation, categorizing specific signals-receiver coupling underlying the success of this group. However, the plot thickens if we take a closer look into the specifics of the robotic agent's morphological and functional embodiment, properties of the environment, which Marocco and Nolfi (2007) did not address. How do these factors interact with each other through evolutionary to lead to the emergence of communication systems? How does the evolution process unfold across time?

With a goal to address these questions, the current project first successfully replicated the results from Marocco and Nolfi (2007), and then took a closer look at the process of the evolution of communication in this evolutionary robotics model. Top traits that contribute to fitness across different time points during the evolution are identified using a multi-level selection gradient analysis on three levels: genetic, network behavior and functional behavior. The results showed that **natural selection for motor behaviors precedes selection for communication behaviors and may serve to bootstrap the resulting communication system at the end of evolution.** The behaviors of the best performing group from each of the ten different experimental runs were also analyzed qualitatively; while there were a number of traits that occurred in multiple groups, the evolved solution across groups are highly heterogeneous.

In the following sections, I will first discuss the computational and biological theories of communication and its evolution, introduce the benefits of using evolutionary robotics as a research method for this topic, and review some previous work. I will specifically address the work of Marocco and Nolfi (2007), which serves as the foundation of the current project. Then, I will discuss the experimental design, implementation, results, and implications of the current project.

A Computational model of the Communication System

When discussing communication, the instinct of many computational researchers, including myself, is to turn to Shannon's information theory. The information theory is the first formal computational model of communication (defined as the transmission of a meaningful message from a signaler to a receiver; Shannon, 1949) and has informed the underlying mechanism of most of the information technology available today. At the beginning of canonical

writing of information theory, *The Mathematical Theory of Communication*, a broad definition was offered: “all of the procedures by which one mind may affect another (Weaver, 1949).”

Only a few lines later, a yet broader definition came up: “the procedures by means of which one mechanism... affects another mechanism.” From a biological perspective, a few extra requirements seem to be lacking, but we will get to that later. For now, let’s first go with this abstract, computational framework and consider the definition of a *communication system* provided by Shannon (1949) in this landmark text. It consists of five parts (Figure 1):

- (1) an information source;
- (2) a transmitter;
- (3) a transmission channel;
- (4) a receiver;
- (5) a destination.

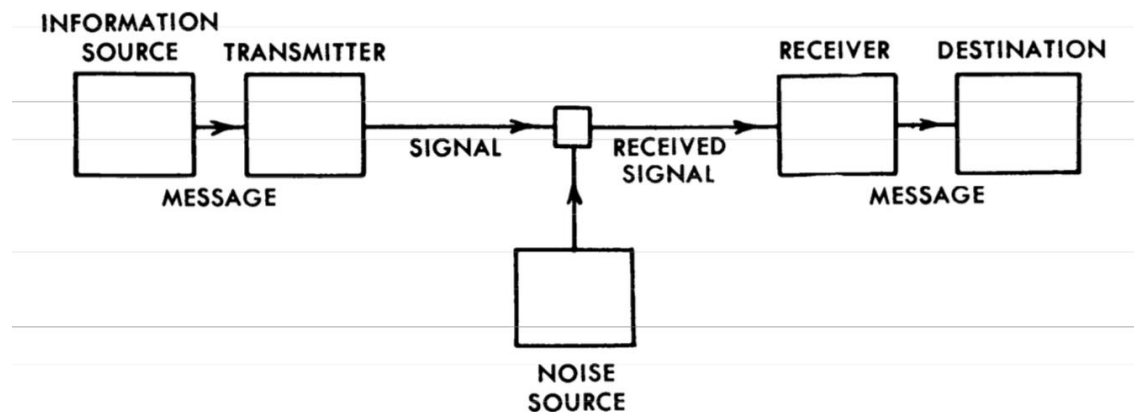


Figure 1. Schematic diagram of a general communication system, from Shannon (1949).

In this system, the information source refers to the component of the system that produces the messages to be communicated to the receiving destination. This “source information” may take any form and be of any nature; for example, let’s say I discovered a great

eatery just now and would like to communicate this to a friend. The transmitter operates on the message produced by the information source and transforms it into a signal that is suitable for transmission through the channel. The channel refers to the medium used to transmit the signal from the transmitter to the receiver. It may also take any form and be of any nature, such as an electric cable, a beam of light, a band or radio frequencies. For example, if I - the transmitter in this example - were to talk to my friend through the channel of Facebook messenger, I might write up my signal in the form of a text: “omg just discovered the best [restaurant emoji] ever [drooling emoji] [drooling emoji] [drooling emoji] must try this weekend” There are noise sources at the end of the transmitter, during the transmission, and at the end of the receiver; I might make a typo, the emoji might show up as something completely different because we were on different phones, and my friend might misread the text due to some distraction. The receiver then decodes the message from the signal, and sends it to the destination; in the current example, my friend serves as both the receiver and the destination, would read the text, understand it, and decide what to do with this piece of information.

Weaver (1949) proposed that the communication system is required to solve problems on three levels: the technical problem, the semantic problem, and the effectiveness problem. The *technical problem* concerns how accurately a signal is transmitted from the sender to the receiver. If the emojis did not show up as the same one I sent because our phones were different, and the restaurant emoji turned into hospital emoji and the drooling emoji turned into angry emoji, my friend would be getting a very different text than the one I sent - an inaccurate signal due to technical difficulties in the transmission. The *semantic problem*, on the other hand, concerns whether the receiver interprets the signal as representing the same meaning as intended by the message. For example, if my intention was to tell my friend to actually go to the restaurant next

weekend, but my friend, despite getting the accurate message, interpreted it as me simply expressing my positive emotions about the restaurant and didn't mean actually going. The semantic problem manifests here as it is the semantic meaning that got lost in translation. Finally, the *effectiveness problem* concerns whether the signal affects the receiver's conduct as intended after its meaning has been successfully conveyed. Let's say my friend knew very well what I meant by the text, yet failed to make any plans to go to the restaurant for the weekend regardless of my communicative effort. Other than that my friend is a terrible friend, this would also mean this communication had not been effective.

While Weaver (1949) suggested that the effectiveness and the semantic problems are closely related and contain most of the philosophical "meaning" of the communication, he emphasized that they rely heavily on the accurate operations on the technical level. It's not hard to see that issues on the technical aspects of transmission would likely lead to problems on the other two levels. If an erroneous text was delivered to my friend, it would only be fair that my friend might fail to accurately interpret its semantic meaning, and the signal would also likely fail to drive my friend to go to the restaurant as I intended. Focusing on the importance of the technical problem, Shannon's (1949) mathematical theory of communication dismissed the semantic aspects of communication as "irrelevant to the engineering problem." He proposed that the key is that there are a number of possible messages, the goal is to choose the actual intended message from them. Rather than being the message itself, the signal provides information and reduces uncertainty as to what is the actual message.

Although my example has featured two human agents communicating via language and text messages, Shannon's communication system seems most fit for describing a mechanical system designed for signal transmission between sources and destinations of messages where the

possible messages are known to begin with. This mathematical model of communication became an important advance in information technology with a wide range of applications; however, it doesn't address how a communication system comes about. What determines what possible messages are out there? What determines what is the best message to send? What determines what is the most effective behavior for the receiver? What signals are encoded and transmitted, and why those specific signals?

In the organic world, these questions seem to be more thorny than the technical problem, which may be avoided by choosing different channels. For example, knowing that the Facebook messenger can be unreliable sometimes, I could opt for a different communication channel, such as hollering at my friend across a playground, or leaving a print out of the message and tap it to my friend's door. But to decide what signal to represent requires more nuanced knowledge than simply the set of all possible messages; the communication system requires knowledge of what signals I know to represent, what signals my friend knows to interpret, and how a given signal may affect my friend's behavior. The questions that I find interesting concerns the process through which a communication system come to know these things, which Shannon and Weaver (1949) have categorized as concerning semantic or effectiveness, rather than the technicality of the signal transmission. To look for better answers, let's then turned to evolutionary biology and artificial life models.

The evolution of biological communication

The father of sociobiology, E. O. Wilson (1975) famously defined communication as “the action on the part of one organism that alters the probability pattern of behavior in another organism in a fashion adaptive to either one or both of the participants (Wilson, 1975).”

Communication, he pointed out, is “neither the signal by itself nor the response; it is instead the relation between the two.” Wilson (1975) especially emphasized two key components: a communication signal must change the behavior pattern in the receiver in some way, and that the trait of communication must be adaptive, as in it must have gone through some kind of natural selection on the level of the genes. In other words, in the process of communication, a signaler transfers a piece of information, a receiver receives this information, which then leads to *a change in behavior of the receiver* in a way that leads to a *net increase in their expected fitness*, which then leads to an increase in relative ratio of the genes underlying the traits involved in such act of communication (Greenfield, 2014).

The signaler-receiver coupling, as well as the communication goal as eliciting behavioral changes in the receiver, echoes the specification in the general model of a communication system proposed by Shannon and Weaver (1949). However, the emphasis on evolutionary adaptation to an agent’s environment was missing from the general computational model. While not explicitly discussed, the environment is not completely absent in Shannon and Weaver (1949). For instance, the message source, the channel, and the message destination may each be viewed as an aspect about the environment, respectively as a source of some message worth communicating, a source of noise, and a state (possibly connecting with the message source) that could be affected by the receiver’s changed behaviors. However, even so, these are all relatively passive roles assumed to be designed as part of the system. In the evolutionary biology perspective, the environment takes on a much more active role in shaping a communication system via determining the fitness function for the communicating agents (Greenfield, 2014).

The evolution of biological communication is a complex process. On the one hand, there are many strict requirements before a communication system may evolve. Every step in the

communication (e.g., sending messages, receiving messages, and interpreting messages) may consist of distinct traits, and each of these may be subject to very different selection pressures. Furthermore, even if a trait is selected, it is not guaranteed that it would go through adaptive evolution because it may not be genetically inheritable. Further complicating the story, each disparate traits may evolve independently and does not guarantee the evolution of other traits needed for meaningful communication, but because communication necessitates all of these traits and have implications to an agent's fitness, the evolution of some traits may depend on other traits in a temporal fashion (Greenfield, 2014). For example, a receiver meerkat's response to an alarm call may necessitate a sender meerkat making that alarm call upon seeing a predatory eagle (Greenfield, 2014). We can perhaps confidently call the alarm call a signal *after* the receiver's response has evolved. But is the alarm call a signal *before* that?

There is considerable controversy surrounding the question of what counts as a *signal* and what is simply a *cue* (Greenfield, 2014). A cue is a piece of sensory information at least partly correlated with a condition of interest that an animal may perceive and assess; it is typically the byproduct of other biological processes that an animal may learn to take advantage of inadvertently (Bradbury & Vehrencamp, 2011). A signal, on the other hand, are stimuli produced by as sender and monitored by a receiver to the average net benefit of both parties; unlike cues, the main function of its production is for the sender to communicate with the receiver (Bradbury & Vehrencamp, 2011). Animals nearly always emit various inadvertent cues that may be perceived by another animal, and in many cases, such perception may benefit one or both individuals, but this alone does not suffice the biological definition of communication signaling (Wilson, 1975). Communication signaling often requires precursors such as cues and takes advantages of sensory biases that developed for other functions such as foraging or

predator avoidance, but then goes through specific adaptive evolution for the traits that support the production and reception of the signals (Greenfield, 2014). Some of the criteria for specific evolution may include extra energy expenditure beyond the necessary amount for the primary physiological function of the signal modality, specialized structures that improve the release of the signal to the outside environment, or a temporally specific routine for the release of the signal that coincides with the periodicity of a behavior that benefits from communication, such as mating behavior and male receptivity (Greenfield, 2014). If none of these indicators are present, it is hard to claim that a stimulus is a signal rather than a cue. This poses some interesting questions for studying the evolution of communication in artificial biological systems, which often build in sensorimotor modalities dedicated for communication only and do not incorporate the factor of energy expenditure (Quinn, 2001).

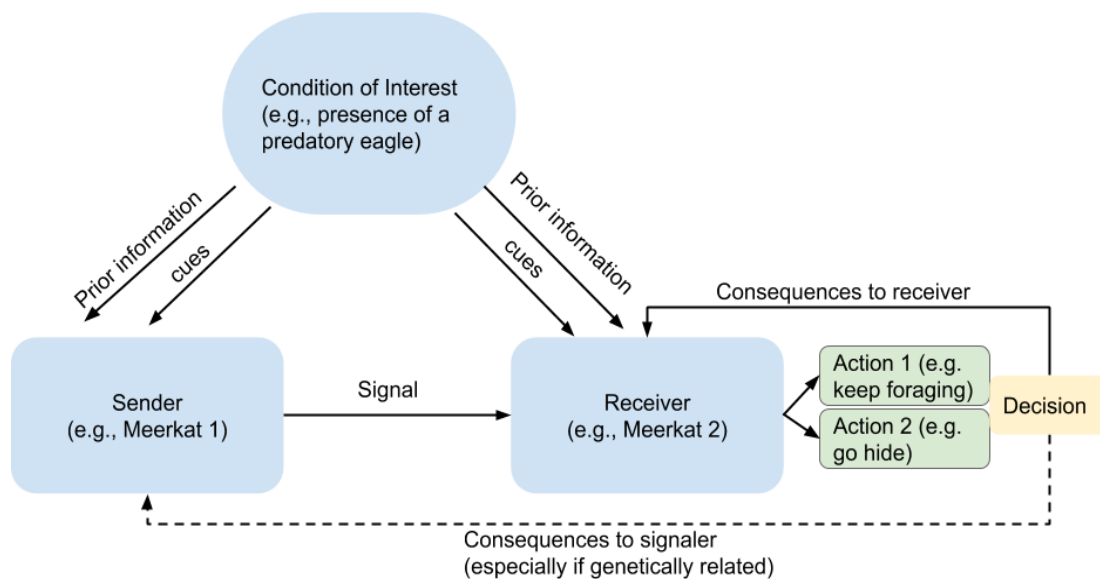


Figure 2. The process of communication in a biological system (adapted from Bradbury & Vehrencamp, 2011). The sender sends a signal to the receiver containing information about the environment (predator activity) that then affects the process of the receiver, which then has consequences to both the receiver (behavioral outcome) and to the sender (if it is genetically closely related to the receiver).

It is worth emphasizing that while an individual's fitness is judged on the basis of functional behavior through natural selection, evolutionary adaptation necessitates the heritability of these behaviors on the level of genotypes, and therefore takes into consideration the fitness of both the signaler and the receiver (Lande & Arnold, 1983; Wilson 1975; Greenfield, 2014)). As a result, even if only either the receiver or the signaler benefit from the signal, if the two agents are genetically related, a stipulation of mutual benefit still holds (Greenfield, 2014). For example, by producing an alarm call, a meerkat may put itself at higher risk of being exposed to the predator, thus lowering its individual fitness; however, if the call helps other meerkat receivers that are genetically related to the sender and increase their chances of survival enough to lead to a net increase in fitness, this trait may still be selected (Figure 2; Bradbury & Vehrencamp, 2011). Similarly, even if the signal may not work all of the time and may not work for every individual within a population, leading to decreased fitness in some individuals, as long as it works for most individuals in that population and leads to a net increase in fitness, it may still be considered adaptive trait with the potential to be passed on. For example, the sexual selection in some species involves high-fitness males leaving chemical signals to attract females; some low-fitness males ("satellites") may hang around the same areas so the females would mistake them for high-fitness males to mate with them (Greenfield, 2014). While this may lowering the fitness of the female's offsprings, as long as it happens infrequently enough, the mutual benefit stipulation still holds, and this communication system would likely still be preserved in these species (Greenfield, 2014).

Unfortunately, as with most behavioral traits, the evolution of animal communication is difficult to study because it seldom leaves direct fossil evidence. In trying to reconstruct this

ephemeral process, researchers have often resorted to indirect, phylogenetic comparisons and/or logical inferences, such as computational modeling (Greenfield, 2014).

Communication in Evolutionary robotics

Artificial life systems, especially evolutionary robotics, turn out to be a good candidate for overcoming some of these caveats. Evolutionary robotics is the subfield of robotics that uses the selection, variation, and heredity principles of natural evolution to aid the development and design of embodied robotic systems (Figure 3; Doucieux et al., 2015).

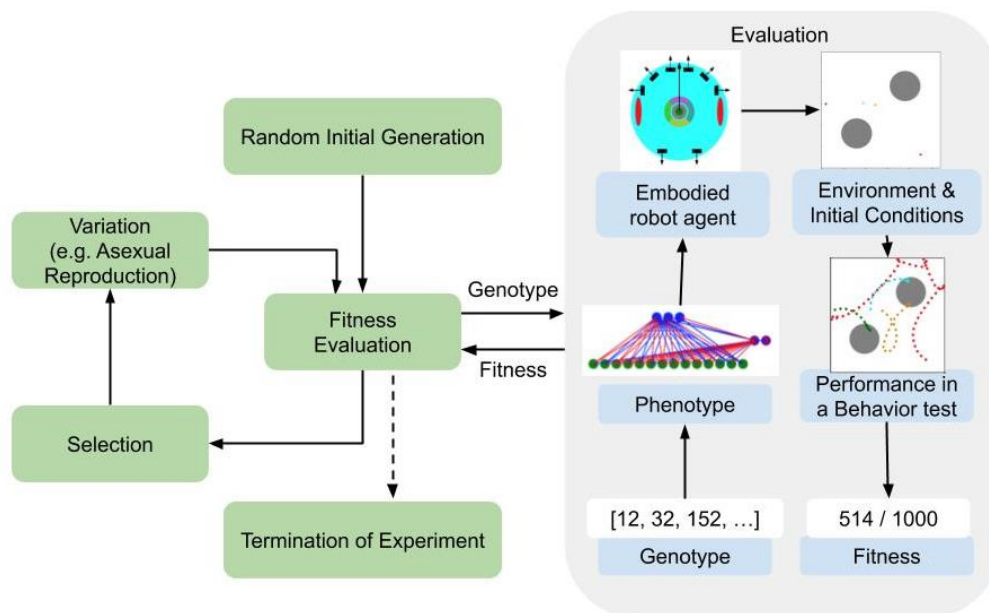


Figure 3. The principle workflow of an evolutionary robotic system (adapted from Doucieux et al., 2015 to fit the current system). The evolution begins with a random generation of genotypes. The genotypes are evaluated through behavior in an embodied, situated agents. A genotype is first converted into a phenotype, typically parameters of a neural controller and/or of a robotic agent's morphology. The robotic agent then goes through some kind of behavioral trial within a task environment, and its performance is evaluated through a fitness function. The fitness determines whether a genotype is selected for reproduction; random variations are then introduced to the selected genotypes through mutation and/or sexual reproduction to produce the new generation of genotypes. The new generation is evaluated again until the end of the experiment (e.g., reaching the number of generations specified before the experiment, or exceeding the required performance).

Although this field is often considered as an alternative method for generating better designs of robotic applications that solve specific problems, the use of embodied robotic systems in studying biological phenomena including communication is also prevalent (Doucieux et al., 2015). In particular, embodied robots enable the simulation of co-evolving, interacting organisms at the level of changes in behavior and perception (Quinn, 2001). Studies using evolutionary robotics typically expose simulated agents to challenges where the use of communication may enhance the agents' performance, such as synergistic collaborative tasks (Szathmáry et al., 2007). These studies ask the questions of in what context does a communication system develop within a population of embodied agents capable of communicating with each other, and aims to identify pre-requisite initial conditions (Marocco & Nolfi, 2007).

Many studies have shown that agents evolve to use communication channels when the task performance requires or may be significantly enhanced by coordination between agents (Marocco & Nolfi, 2007; Quinn et al., 2003; Quinn, 2000; Quinn, 2001; Szathmáry et al., 2007; Floreano, Mitri, Magnenat & Keller, 2007). However, we lack a consensus on specific methods from genome encoding to behavioral tasks across different studies (Doucieux et al., 2015); each of the studies I have come across has used a different robotic system and a different behavioral task. Because the evolution of the robotic agents' behavior may be specific to the environment and the task conditions they have evolved under and are sensitive to changes in these conditions (Lipson, 2007), the heterogeneity in research paradigms makes it hard to generalize findings beyond the specific setups used in the study.

For this reason, the current thesis project chooses to follow an existing paradigm rather than designing a new one. The project follows the system used in Marocco and Nolfi (2007), which is itself an adaptation of the robotic system used in Quinn (Quinn 2000; Quinn 2001).

Both Quinn and Marocco and Nolfi were interested in the evolution of communication in a homogenous robot system, where all agents within the population share the same genotype and phenotype (neural network controller). Marocco and Nolfi (2007) investigate the effect of communication sensor access in a four-robot system evolved in a coordinated foraging task, where the fitness increased for each robot located within one of the two target area (emulating a food patch), but decreased at a higher rate for each robot beyond two robots (emulating the consequences of exceeding carrying capacity of a food patch). They found that when the robots had access to dedicated communication channels, they achieved significantly higher fitness (measured by performance in a coordinated foraging task) after 100 generations of evolution. Additionally, after removing access to the communication sensor, the fitness of the communicating systems dropped significantly, confirming that the communication and coordination behavior of these agents underlies a portion of the increase in their fitness. They also conducted an ethological analysis of the best performing networks, documented five distinct signals and the context of their production, the signaler and receiver of each signal, and inferred the semantic meaning of these signals (Table 1; Fig. 4), including:

- (1) Signal A: medium intensity signal produced by robots outside of target area not interacting with other robots (receivers of this signal produces signal E);
- (2) Signal B: high-intensity signal produced by robots located alone inside a target area (receivers modify motor behaviors to approach signal source and produce signal D);
- (3) Signal C: an oscillating signal produced by robots inside a target area that also contains another robot (receivers of this signal modify their motor behavior to move away from the signal source);

- (4) Signal D: a very low-intensity signal produced by robots outside target areas approaching a target area containing another robot located alone inside (produced in the reaction of receiving signal A; no specified receiver);
- (5) Signal E: an oscillatory signal produced by robots located outside target areas interacting with other robots that are outside target areas (receivers modify their motor behavior to better explore the environment).

Table 1. The five signals categorized by Marocco and Nolfi (2007) in the best performing group.

	Intensity	Sender Status	Receiver Response
Signal A	medium	Alone outside target area	Produce signal E
Signal B	high	Alone inside target area	Approach source; produce signal D
Signal C	oscillating	Inside target with another robot	Avoid signal source
Signal D	low	Approaching signal A	No receiver
Signal E	oscillating	Outside target receiving signal A	Explore the environment

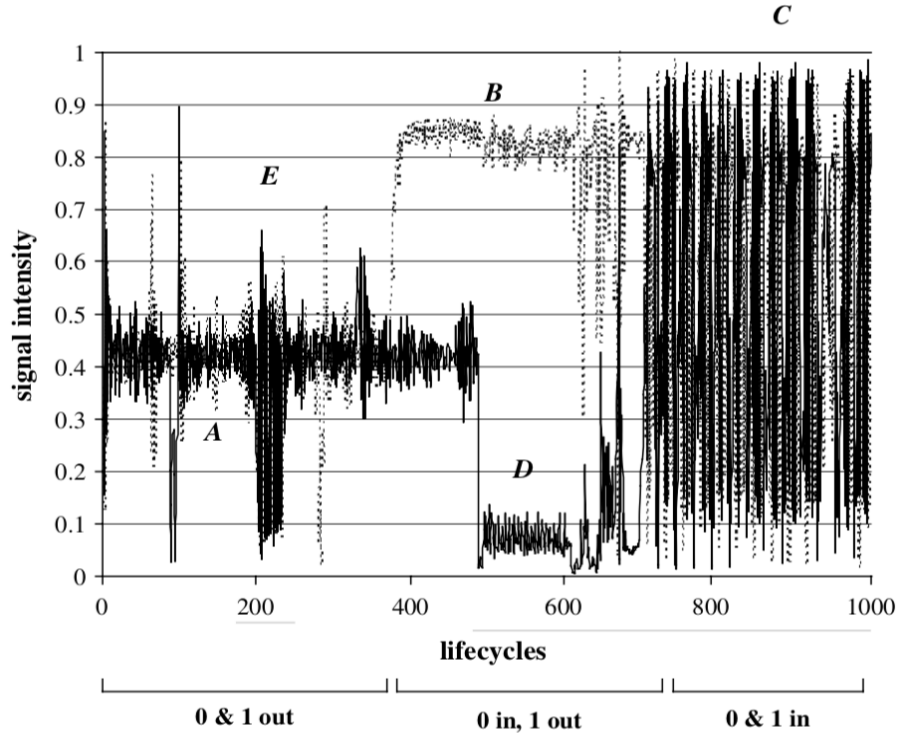


Figure 4. The five signals classified by Marocco and Nolfi (2007). The signals are from a test trial with only two agents and one target area; the agent shared the network of the best-performing evolved group.

While this finding is interesting, it is unclear how prevalent or replicable this particular communication system is beyond the current system or even within the current system, given that it's the only solution Marocco and Nolfi reported. Additionally, according to the biological definitions of a communication signal, it is dubious if signal D, while certainly still a functional behavior, should be considered as a communication signal. While the other signals may have gone through adaptive evolution, no evidence for such evolution was provided. Marocco and Nolfi also did not provide detailed analyses of the embodied and environmental aspects of the current system. For example, they did not discuss the effect of the directionality of the communication sensors, or any behavior related to the IR sensors.

The current project aims to better understand some of these new and remaining questions stemming from the Marocco and Nolfi (2007) study by taking a closer look at the same system. The specific aims, as well as an overview of the methods used to achieve these aims, are outlined below.

Objectives for the Current Project

As previously mentioned, the current project aims to build on the basis of Marocco and Nolfi (2007) who has demonstrated some of the end products evolution of communication in their system, and more importantly to expand the analysis with a focus on the process of how these communication systems evolved.

This first aim was **to examine whether the network controller structure of the system had confounded the results in Marocco and Nolfi (2007)**. In Marocco and Nolfi (2007), the robotic agents also had sensory access to its own communication signal emitted at the previous time step, allowing them to functionally “hear themselves.” In the current study, the effect of this

“hearing oneself” trait from that of “hearing others” trait was examined by adding two conditions on top of the original (comm vs. no comm) condition in Marocco and Nolfi (2007) where this “hearing oneself” sensory access was removed. As the results showed that removing this trait did not lead to any statistically significant difference in the final fitness of the comm and no comm groups, the rest of the analyses were conducted using only the original conditions in Marocco and Nolfi in order to say comparable to their study.

Secondly, the current project investigated the contribution of traits to fitness on different levels, with an aim **to identify the best level of analysis for understanding how traits relate to fitness**. Traits were defined on three levels of analyses: from bottom to top, the genes, the network, and the functional behavior. Each level is related to the level below it and was in some manner calculated bottom-up from the level below. The traits on each level were used to predict fitness for selection gradient analyses using linear multiple regression models (see Moorad & Wade, 2013; Lande & Arnold, 1983). The performance of these statistical models was compared with each other; this revealed that the models on the functional behavioral level lead to the best prediction and explained most of the variation in the fitness data. Consequent analyses then focused on this level.

The third aim of the project was **to investigate whether different behaviors are preferred by selection during different stages of evolution, as well as whether some behaviors that evolved early on served as precursors to the later evolution of communicative behaviors**. Marocco and Nolfi (2007) only examined the resulting network and did not address anything about how the process of evolution unfolded in terms of what selections and adaptive evolution the evolve communication traits went through. However, many of the most important signals within the communication system they described and hypothesized to

contribute to fitness, such as signal B (sent by robots in target areas alone to attract robots outside) and signal C (sent by robots in target areas with another robot to drive away extra robots that may otherwise come within the target area), seem to require precursor traits related to the condition of interest (the target area). For example, for the signal A to be considered a meaningful signal, the agents may need to first evolve the behavior of stopping upon detection of a target area. This observation leads to a new question: are there required precursor traits before communication behaviors could evolve?

Quinn (2001) reported that artificial evolution, much like biological evolution, goes through stages where some general motor behaviors, such as exploring the environment, are selected first before other task-related traits. Following this idea, the current project identified traits that contribute the most to fitness on each of the three levels of analyses at the beginning and the end of the evolutionary trials across the 10 different experimental runs. The main finding confirmed Quinn's (2010) discovery that behaviors that promote exploration of the environment, including a high baseline speed, a slight turn, and a reduction in speed upon detecting target area were favored at the beginning of the evolution, while more complex traits such as IR avoidance, and communication-related traits emerged later during the evolution. It is worth noting that while the selection of behaviors at the beginning of the evolution was highly similar across different runs, while the communication strategies evolved at the end of the evolution were much more heterogeneous.

The final aim of the current project is to investigate **whether the same functional behaviors in this system are multiply realizable via different genotypes and network phenotypes**. While Marocco and Nolfi only focused on the evolved communication system in one group of agents, I find it hard to believe that there would only be one possible solution. It's

possible that there are many genotypes that are considerably different from each other, but share the same level of high fitness due to similarities in their functionally categorized behavior - which is the level selection operates on (Lande & Arnold, 1983). This also relates to the multiple realizability in the functionalist perspective of the theory of mind and cognition: mental states (or cognitive behavior) exist (or evolved) to serve specific functions, and each function may be achieved in a number of ways - i.e., they are multiply realizable (Levin, 2018). This line of reasoning leads to the argument that mental states should be defined based on their properties at the functional behavioral level rather than their physical properties (e.g., the neurological mechanism underlying pain, or the specific pattern of weights in a neural network). How much similarity is there among the genotypes, and how closely do they relate to the phenotypes of the agents? What is permitted to change, and what are the underlying constraints? If we have a good theory about what the functional purposes the agent's behavior serve, does that lead to a better way to categorize their behaviors? The current project examines these questions via comparisons between groups of agents that have different phenotypes and genotypes but share similar motor, IR obstacle avoidance, and communication behaviors, as well as between groups of agents. A surprising finding is that it seemed that functional behaviors were selected without regard to the dedicated sensor modality. Certain genotypes underlying dedicated communication sensors and motor traits were preferred at the beginning of the evolution because motor behaviors that they drive rather than communication behavior. On the other hand, communication-like coordination behaviors occasionally evolved under conditions with no access to communication channels through between-agent interactions via non-communication channels (IR sensors) by the end of the evolution. This gets to the question of the signal vs. cue divide and what should be qualified as communication.

Methods

The Overall Evolutionary Robotics System

The evolutionary robotic system used in this project is implemented in Python (version 3.6.1) exactly as specified in Marocco and Nolfi (2007), Quinn (2001) and Nolfi (2003). All figures and videos of this system, including the environment, the agents, the trial plot and animation, are generated using the Python package Matplotlib (version 2.2.2; Hunter, 2007). The code files are available at <https://osf.io/zpve7/>.

Environment

As specified in Marocco and Nolfi (2007), the simulated environment is a 270 cm x 270 cm square arena. The arena is represented by a two-dimensional Cartesian coordinate system, with the x range = [0, 270] and the y range = [0, 270] (Fig. 5). There are two target areas in the arena, each with a radius of 35 cm, with a center at (80, 80) and (190, 190), respectively.




Figure 5. Diagrams of the environment, with the agents. Left: diagram from Marocco and Nolfi (2007). Right: diagram generated via the current implementation of the system. The target areas are represented with the color grey. The agents are represented by small circles; the four agents (agent 0, 1, 2, 3) are represented with colored red, orange, blue, and green, in that order. The diagram of the current implementation is generated to scale (x- and y-axes represent centimeters).

Agent

The robotic agent is an adaptation of the Khepera mini-robot system that had been used by Quinn and colleagues in a number of projects (Quinn, 2000; Quinn, 2001). Each agent has a circular body with a radius of 2.9 cm diameter, or 5.8 cm in diameter (Fig. 5; Quinn, 2001). Each agent is equipped with 8 IR proximity sensors, 1 ground sensor, 4 communication sensors, 1 communication unit, and 2 wheels. Each agent is controlled by an artificial neural network with a fixed structure including 14 input nodes and 3 output nodes. Specifications for sensors, controller and actuators are detailed in the following sections. Different agents are represented via different colors; agent_0 is always represented in red, agent_1 in orange, agent_3 in blue, and agent_4 in green.




Figure 6. Diagrams of an individual agent, pointing towards 90°. Left: diagram from Quinn (2001; 2000). The IR sensors are labeled with respective ID numbers. The center of the robot is marked with a cross. Right: diagram generated via the current implementation of the system. The positions of the IR sensors, represented in black rectangles, correspond with those in Quinn (2001; 2000). The four communication sensors point to the four cardinal directions as specified in Marocco and Nolfi (2007) and are filled with the color purple (front, comm_0), light green (left, comm_1), yellow (back, comm_2), and blue (right, comm_3). The red ovals represent the left and right wheels. The dark green dot at the center of the robot represents the comm_unit, and the gray dot represents the ground sensor. The diagram of the current implementation is generated to scale (x- and y-axes represent centimeters).

Sensors

IR sensor. As specified in Quinn (2001), each agent is equipped with 8 short range infrared (IR) proximity sensors (Fig. 6). All IR sensors are placed along the edge of the agent's body; one sensor points to the 90 degrees left of the robot, one to 45 degrees left, two directly in front, one to 45 degrees right, one to 90 degrees right, and two directly towards the back. Each IR sensor has a detection range of 5 cm.

Both wall and other agents may induce IR sensor readings. The sensor will produce a reading of 0 if no wall or another agent is present directly in front of it within its detection range. If a wall or another agent is detected, the sensor will produce a reading between 0 - 1, depending on the distance between the detected wall or agent closest to the sensor. The relationship between the distance, d , and reading, r , is expressed as:

$$r = 1 - (\text{range_max} - d / \text{range_max})$$

The detection process for each IR sensor is implemented by the *get_ir_reading()* function. This function takes an IR sensor's position, angle, and detection range and the position and radius of other agents as inputs. The function first checks whether the wall is within the detection range by checking whether the coordinates of the point of minimum signal detection goes beyond the range of the environment (i.e., $x < 0$, $x > 270$, $y < 0$, or $y > 270$). If so, this means the wall is within the detection range; the distance between the IR sensor and the wall is then calculated using principles of trigonometry. Then, for each other agent, the function checks whether the agent is within the detection range. If so, the distance between the IR sensor and each agent is recorded. Finally, if the function detected both the wall and other agents, the obstacle with the shortest distance will be selected to produce the final reading of the IR sensor.

Ground sensor. Each agent is equipped with a ground sensor, placed at the center of the body. The ground sensor determines whether the agent is in or outside of a target area. The detection process for the ground sensor is implemented through the *get_ground_reading()* function. For each target area, the function checks whether the distance between the agent's current position and the center of the target area is less than the radius of the target area. If this is true of one of the target areas, the ground sensor will read 1; otherwise, the ground sensor will read 0.

Communication sensors. Each agent is equipped with 4 communication sensors. Each sensor has a detection range of 100 cm. The communication sensors are placed along the edge of the body and in a way that each is responsible for detecting communication signals from one of the four orthogonal directions, covering an angular range of 90 degrees; i.e., if the front is 0 degree, the angular coverage of the sensors are respectively: frontal [315°-44°], rear [45°-134°], left[135°-224°], and right[225°-314°]) (Fig. 6). This means that as long as an agent A is within 100 cm of a second agent B, the communication signal emitted by A at the current timestep will be received by B through one of its communication sensors. The sensor reading depends on the intensity of the signal emitted by A in a linear fashion. The detection process for communication sensors is implemented through a single function, *get_comm_reading()*. The function first checks whether each of the other agents is within the detection range (100 cm) of the current agent. If so, the coordinates and signal intensity of that agent is recorded. Then, the function classifies the direction of each detected agent into one of the four orthogonal directions, identifying its corresponding communication sensor. If more than one agent is detected by the same sensor, only the signal from the agent closest to the current agent will determine the reading.

Controller

The controller is implemented exactly as specified in Marocco and Nolfi (2007). It consists of 14 sensory neurons, 2 internal neurons, and 3 motor neurons (Fig. 7). The 14 sensory neurons correspond with 8 IR sensors, 1 ground sensor, 4 communication units, and an extra node for the state of communication unit at the previous timestep (`comm_self`), respectively. The 3 motor units correspond with the left motor output, the right motor output, and the communication unit (`comm_unit`) output, respectively.

Each of the 14 sensory neurons is connected directly to each of the two motor unit neurons. Additionally, except for the `comm_self` neuron, each of the 13 other sensory neurons is connected to internal neuron 1, which is then connected to each of the 3 motor neurons. The `comm_self` neuron is also the only sensory neuron directly connected to the communication motor unit. The 2 internal neurons are connected with each other in a reciprocal fashion; internal neuron 2 only forms connections with internal neuron 1. The `comm_unit` neuron is also connected to the `comm_self` sensory neuron, with its output serving the input for the `comm_self` at the next timestep. Marocco and Nolfi (2007) did not give an elaborate reason for the special treatment of the `comm_self` neuron, the only node not connected to the internal neuron, but a possible reason for this design is to avoid the formation of a positive feedback loop where the value of the `comm_unit` becomes overtaken by its own “echoes” due to the multiple sources of recurrence. In the current project, in two of the four conditions, the connections involving the `comm_self` neuron is removed from the genome in order to investigate the effect of an agent’s access to its own communication signal (Fig. 7).

Each of the sensory neurons and internal neurons has a time constant taking values within the range of [0.0, 1.0]. Each of the internal neurons and the output neurons has a bias, normalized

to the range of $[-5.0, +5.0]$. There is also a connection weight for each of the possible connections between neurons, normalized to the range of $[-5.0, +5.0]$. This leads to a total of 16 times constants (14 for sensory neurons, 2 for internal neurons), 5 biases (2 for internal neurons, 3 for motor neurons), and 48 connection weights (28 for sensory neurons to motor unit neurons, 1 for comm_self to comm unit neuron, 13 for sensory neurons (excluding comm_self) to internal neurons, 2 between the internal neurons, 3 for internal neuron 1 to the 3 motor neurons, and 1 for comm_unit motor neuron to the comm_self sensory neuron). This leads to a total of 69 gene loci. In the no_self condition where the four connections involving comm_self are removed from the genome, the size of each genotype becomes a total of 65 gene loci (Fig. 7). Each genotype is represented by an array of the respective length for the genome size (either 69 or 65), which contains integers between the inclusive range of 0 to 255, simulating the 8-byte representation used in Marocco and Nolfi (2007). Each position in the array is associated with a corresponding network attribute (time constant, bias, or connection weight); the correspondence relationship is stored as a lookup table and used for genotype-phenotype mapping. The activation of the nodes are calculated exactly as specified in Marocco and Nolfi (2007) (Fig. 8).

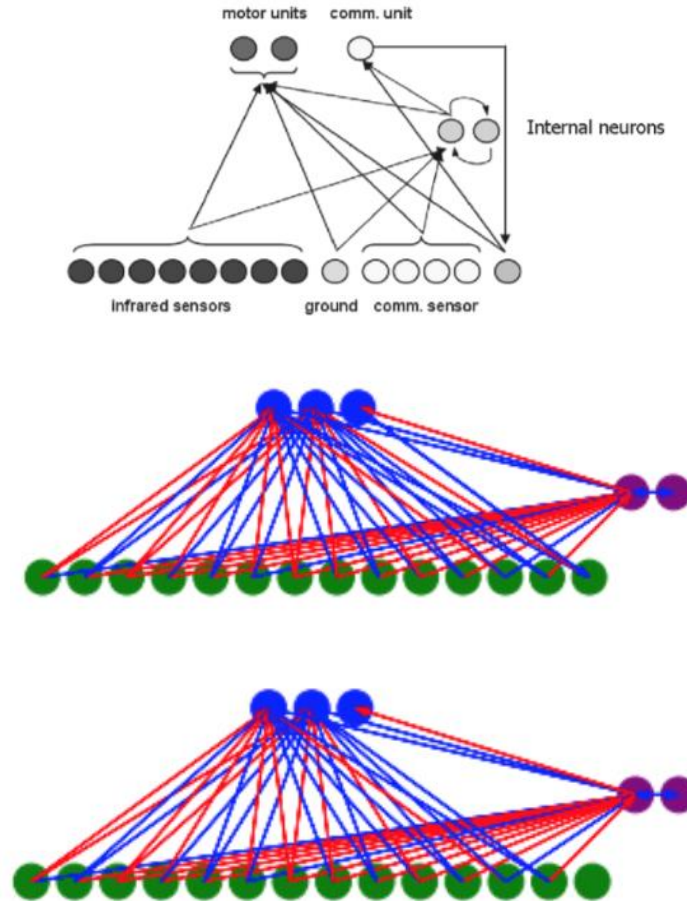


Figure 7. Diagrams of the network controller. Top: controller diagram from Marocco and Nolfi (2007). Middle: diagram of a controller generated in the current system in conditions with `comm_self` neuron connected. Bottom: a controller in the current system without connections to or from the `comm_self` neuron. This network is chosen from a random individual. The positions of the neurons correspond with those in Marocco and Nolfi diagram. The input neurons are IR 0-7, ground sensor, comm 0-3, `comm_self`, respectively from left to right. The internal neurons are internal 1, internal 2 respectively. The output neurons are motor left, motor right, and `comm_unit`, respectively. Red represents positive connection weights, blue represents negative weights.

$$A_j = t_j + \sum_i w_{ij} O_i \quad (1)$$

$$O_j = \frac{1}{1 + e^{-A_j}} \quad (2)$$

$$O_j = O_j^{(t-1)} \tau + I_j (1 - \tau_j) \quad (3)$$

$$O_j = O_j^{(t-1)} \tau + (1 + e^{-A_j})^{-1} (1 - \tau_j) \quad (4)$$

Figure 8. The update rules of the network as specified in Marocco and Nolfi (2007). Equation (1) noted the calculation of raw activation of the internal and the output neurons. Equations (2) noted the activation function of the output neurons, (3) for sensory neurons, and (4) for internal neurons. Notations: A_j : activity of the j th neuron; t_j : bias of the j th neuron; w_{ij} : weight of the incoming connections from the i th to the j th neuron; O_i : the output of the i th neuron; $O_j^{(t-1)}$: the output of the j th neuron at the previous time step; τ_j : time constant of the j th neuron; I_j : intensity reading of the j th sensor.

Actuators

Motor wheels. Each agent has two wheels, one on the left side of the body and one on the right, each placed approximately 2.6 cm away from the center of the body (Fig. 6). In the simulation, the relationship between wheel movement and the agent's displacement and angular change is approximated as follows, as specified in Quinn (2001).

The network output values for the motors are scaled to provide the actual motor velocity. Each motor has a maximum speed of 8 cm/s (while the minimum speed is 0 cm/s). Noise is then applied by multiplying each motor velocity by a random value chosen from a uniform distribution with the range [0.9-1.1]. The relationship between the velocities for the left motor (v_l) and the right motor (v_r) and the change in agent's orientation ($\Delta\theta$) and position (ΔP) is as follows:

$$\Delta\theta = (v_r - v_l) / (w * u)$$

$$\Delta P = (v_r + v_l) / (2 * u)$$

Here, u and w are both constants, u represents the number of simulation updates per second (set to 10 in this experiment), and w represents the distance separating the two wheels in Khepera robots (5.2 cm). The new position of the robot is updated by moving the robot ΔP cm in the new direction $(\theta + \Delta\theta)$. If the updated position is beyond the walls of the arena after updating, it is correct by placing the agent in the given direction as far as the wall goes. Agent collision is not currently implemented in the system.

Communication Unit. Each agent is equipped with 1 communication unit, placed at the center of the body (Fig. 6). The communication unit may emit a communication signal at each time step. The intensity of the signal may vary between 0-1.

Coordination Foraging Task

The behavior task and fitness function used in the experiment simulates a foraging task, with additional requirements of coordination across agents. Following the protocol in Marocco and Nolfi (2007), at the beginning of each trial, four agents that share the same genotype are each assigned a random starting position in the arena outside of the target areas. Then, the group of agents is allowed to propagate for 1000 timesteps within the environment, each time step set to represent 100 ms in physical reality. The fitness of the system at every timestep is calculated based on the relative positions of the four agents, according to the following rules:

- (1) While the total number of agents in a target area does not exceed two, for every agent within this target area, the group gains 0.25 point in fitness score;
- (2) For every agent more than two agents within the same target area, the group receives a deduction of 1.0 in fitness score.

For example, if each target area contains two agents at a time step, the fitness of this time step would be 1.0 (Fig. 8). If one target area contains four agents at a time step, the fitness of this time step would be -1.5 (Fig. 9). The total fitness of a trial is calculated as the sum of all time steps in the trial, with the maximum fitness capped at 1000 points (in Marocco and Nolfi (2007) this is calculated as the mean across all times steps, with a cap of 1 instead of 1000).




Figure 9. Examples illustrating fitness calculation of every time step. Left: each target area contains two agents, each agent gains 0.25 points in fitness score, leading to a total score of 1.0. Right: the top target area contains four agents. For two of these four agents, each gains 0.25 points in fitness score, adding to 0.5; however, each of the other two leads to a 1.0 point deduction, leading to a total deduction of 2.0 points. The total fitness score for this group at the current time step is $0.5 - 2 = -1.5$.

Experimental Design

The first experimental manipulation concerns the access of an agent to the signal produced by other agents within the environment. This is the same manipulation carried out by Marocco and Nolfi (2007) and achieved in the same way by always assigning an input of 0 to the four communication sensors before the sensory inputs enter the controller computation.

An additional experimental manipulation is the access of an agent to the signal produced by itself at the previous time step. This is achieved by removing all of the four genes that encode connection weights involving the comm_self sensory neuron (Fig. 7).

The resulting four conditions are, respectively:

- Condition 1: access to both comm_others and comm_self;
- Condition 2: access to comm_others, but not comm_self;
- Condition 3: no access to comm_others, but access to comm_self;
- Condition 4: no access to either comm_others or comm_self.

Each experiment run starts with a set of randomly generated first generations, containing 100 genotypes. This first starting generation is separately evolved under each of the four conditions. During the evaluation phase, each genotype is converted into four copies of identical network controllers, one for each agent. The fitness score of a group of agents is determined by taking the average fitness score from 20 trials, each lasting 1000 times steps. The genotypes of the 20 top-performing groups ($\frac{1}{5}$) of the generation are selected for reproduction; 5 copies of each genotype are generated with a mutation rate of 0.02 at each locus. The experiment terminates after the 100th generation. The experiment is replicated 10 times, each with a different set of randomly generated first generations.

In Marocco and Nolfi (2007), an additional “testing phase” followed the termination of every experimental run, and the genotypes in the newly generated 101st generation each goes through an extensive evaluation of 1000 trials; the average fitness across these additional trials are used as dependent measures. In the current project, this step is omitted due to the limitation of time and computational resources, and the dependent measures are simply the fitness scores of the 100th generation. Nonetheless, the same as Marocco and Nolfi did, to confirm any increase in fitness of the 100th generation evolved under condition 1 is caused by the access to communication sensors, these fitness scores of these agents are additionally tested in a “deprived” condition for 20 task trials where the agents did not have access to comm_others sensors just as

in condition 3 (Marocco and Nolfi did this for 1000 trials). If the results from Marocco and Nolfi is robust, it should be safe to expect this measures can still reproduce the results in Marocco and Nolfi and lead to meaningful analysis despite the decreased amount of samples and the therefore increased variance.

Data Analysis

Fitness comparisons between conditions. To investigate the effect of comm_others and comm_self analysis, a two-way repeated measures ANOVA was conducted using the statsmodels package (Seabold, Skipper, & Perktold, 2010). In order to test for successful replication of Marocco and Nolfi (2007), additional paired t-tests were carried out between the normal agents (final generation evolved with access to comm_others sensors) and the deprived condition 1 (same genotypes but tested without access to comm_others sensors), and between the normal gents and the no comm agents (final generation evolved without access to comm_others sensors).

Analyses of Selection gradient. Multiple regressions and follow up correlation tests were conducted for each condition at the first generation and the last generation to assess the contributions and correlations of each of agents' traits to fitness at different time points of the evolution. These analyses were carried out on three different levels: genes, network behavior, and functional behavior. Separate analyses were conducted for the beginning (initial generation) and the end (last generation) of the evolution.

For the selection gradients on the level of the genes, the factors considered included all genetic loci (69 factors for comm_self conditions, 65 for conditions with no comm_self) for all networks in the condition within the given generation.

For the selection gradients on the level of the network behavior, three types of expected baseline motor behaviors were measured: average comm intensity (between 0-1), average speed (between 0-1), and average turn in radian (between 0-1). The baseline behavior measures were taken for when the sensory inputs simulated what the agent would perceive when outside (ground sensor input = 0) vs. inside (ground sensor input = 1) the target areas without perceiving any other sensory inputs. This is achieved by feeding the network controller the same expected input for 100 timesteps and taking the average values across these timesteps. Then, sensor-dependent changes in baseline motor behaviors (when the agent is outside of the target areas and ground sensor reading = 0) were taken for each of the 13 sensors (8 IR sensors, 4 comm sensors; comm_self was not included in this analysis) by the same method. Baseline motor behaviors for when the agent is inside of the target area (ground sensor reading = 1). This resulted in 45 network level traits (effects on each of three motor behaviors for each of the 13 sensors and 2 baselines).

For the selection gradients on the level of the functional behavior, the fitness-related consequences of some of an agent's various behaviors within the given environment are first examined. This analysis focused on 4 functional behaviors, including both communicative and non-communicative behaviors, which are driven by one or more network behavior level traits; the traits underlying these functional behaviors were extracted from the basic, network level motor behavioral traits and treated as predicting factors in the multiple regression tests on this level:

- (1) Exploration (Fig. 10; Fig. 11): explore the environment to find the target area; primarily driven by `baseline_speed` and the absolute value of `baseline_turn`. The sign of the `baseline_turn` determines the “handedness” of the agent: if the value is negative, the agent naturally turns right; if the value is positive, the agent naturally turns left.
- (2) Stop at target (Figure 12): driven by the reduction in speed of one or both of the motor units via ground sensor activation (`ground_speed_delta`).
- (3) Wall avoidance (Figure 13): driven by the net changes in turning resulting from directional IR sensors on either or both sides; there are a number of ways to avoid being stuck along the wall, which is discussed below.
- (4) Communication signal from agents inside the target area to attract agents outside: same as signal B described in Marocco and Nolfi (2007). This requires two behavioral traits. First, the sender must be able to stop or slow down when a target area is reached (driven by `ground_speed_delta`), and the ground reading must send a signal strong enough to affect the receiver’s behavior (driven by `ground_comm_delta`). Secondly, the receiver must respond to the signal by approaching it, which requires its directional comm sensors to produce a change in turning in the correct direction, which depends on the agent’s handedness (driven by `comm_0_turn_delta`, `comm_1_turn_delta`, `comm_2_turn_delta`, and `comm_3_turn_delta`).

It is worth noting that the definition of the fourth trait (attraction to communication signal) is relatively arbitrary. Because for the agents that receive this signal, avoiding other agents may be equally if not more beneficial to their fitness. However, in Marocco and Nolfi (2007) this

avoidance behavior (signal C) is described as involving three agents, and the current set up did not include inter-agent level analysis. Therefore, we chose to operationally define the approach behavior similar to signal B rather than the avoidance behavior.

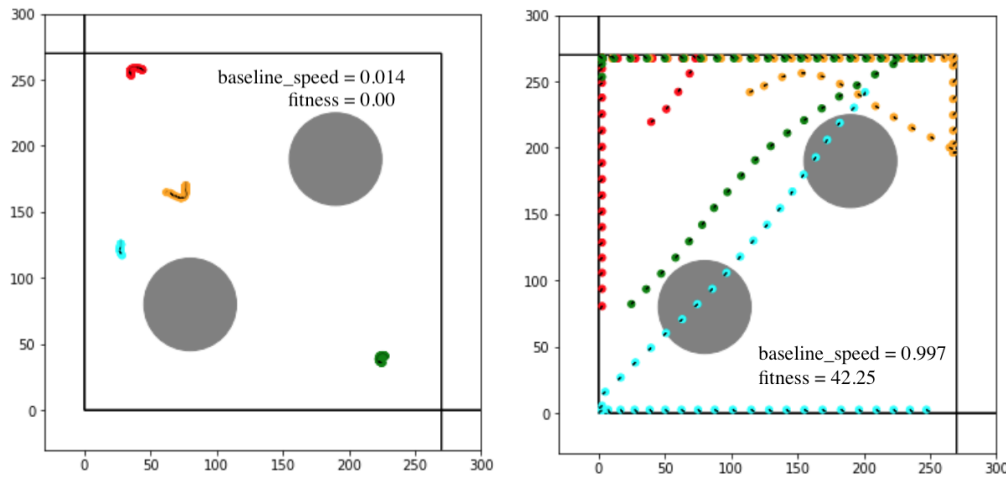


Figure 10. Comparisons between random networks with low and high values for baseline_speed. The four agents are represented by circles colored red, orange, blue, and green, respectively; the black arrow on each agent represents the direction that the agent is pointing towards at a given time step. Agents with higher baseline_speeds are hypothesized to have higher fitness due to a better exploration of the environment and a higher chance of encountering the target areas. Agents with low mobility can rarely achieve a non-zero fitness value.

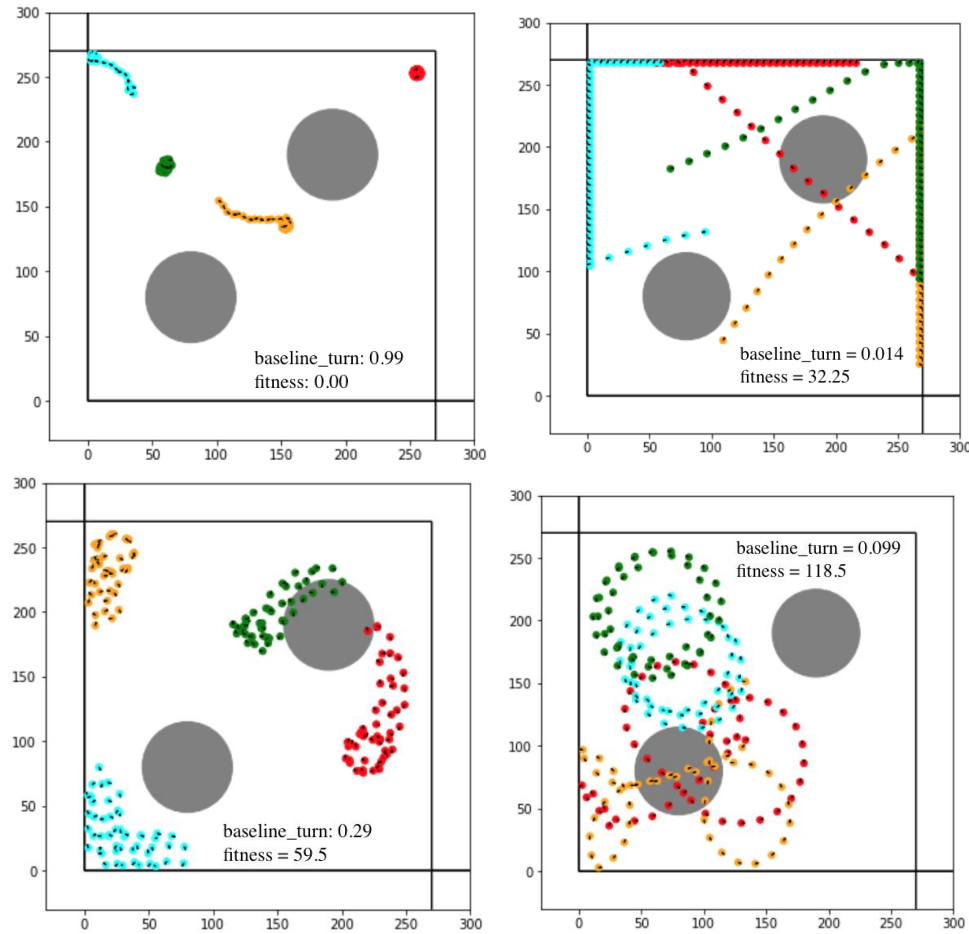


Figure 11. Comparisons between random networks with varying values for `baseline_turn`. The amplitude of the turn determines how tight the turn is, while the sign determines the direction (negative value: right turn; positive value: left turn). We may call an agent “right-handed” if the baseline turn is a right turn (all agents shown above are right-handed). Neither a very tight turn (top left) nor a very small turn (top right) is good for exploring the environment; the former leads to low mobility, and the latter leads to a higher chance of hitting the wall and getting stuck. There is an ideal value around $0.04 \sim 0.1$ that leads to a curved trial that allows the robots to explore the environment efficiently without hitting the wall too often (similar to bottom right).

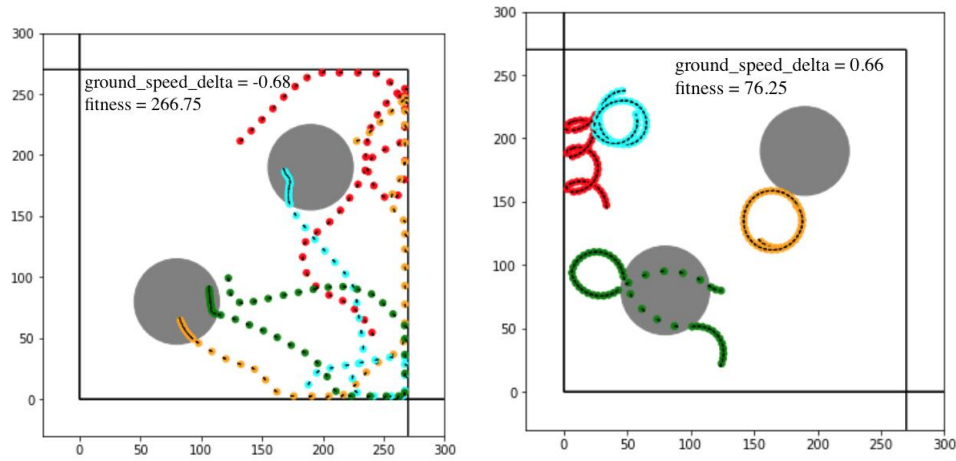


Figure 12. Comparisons between random networks that slow down upon perceiving a target area (top left) and those that do not stop and/or move faster (top right). There are a few possible behavioral changes due to the perception of the target area, among which being able to stop one or both motor units tends to lead to a higher fitness by prolonging the total number of timestep a given agent is within a target area.

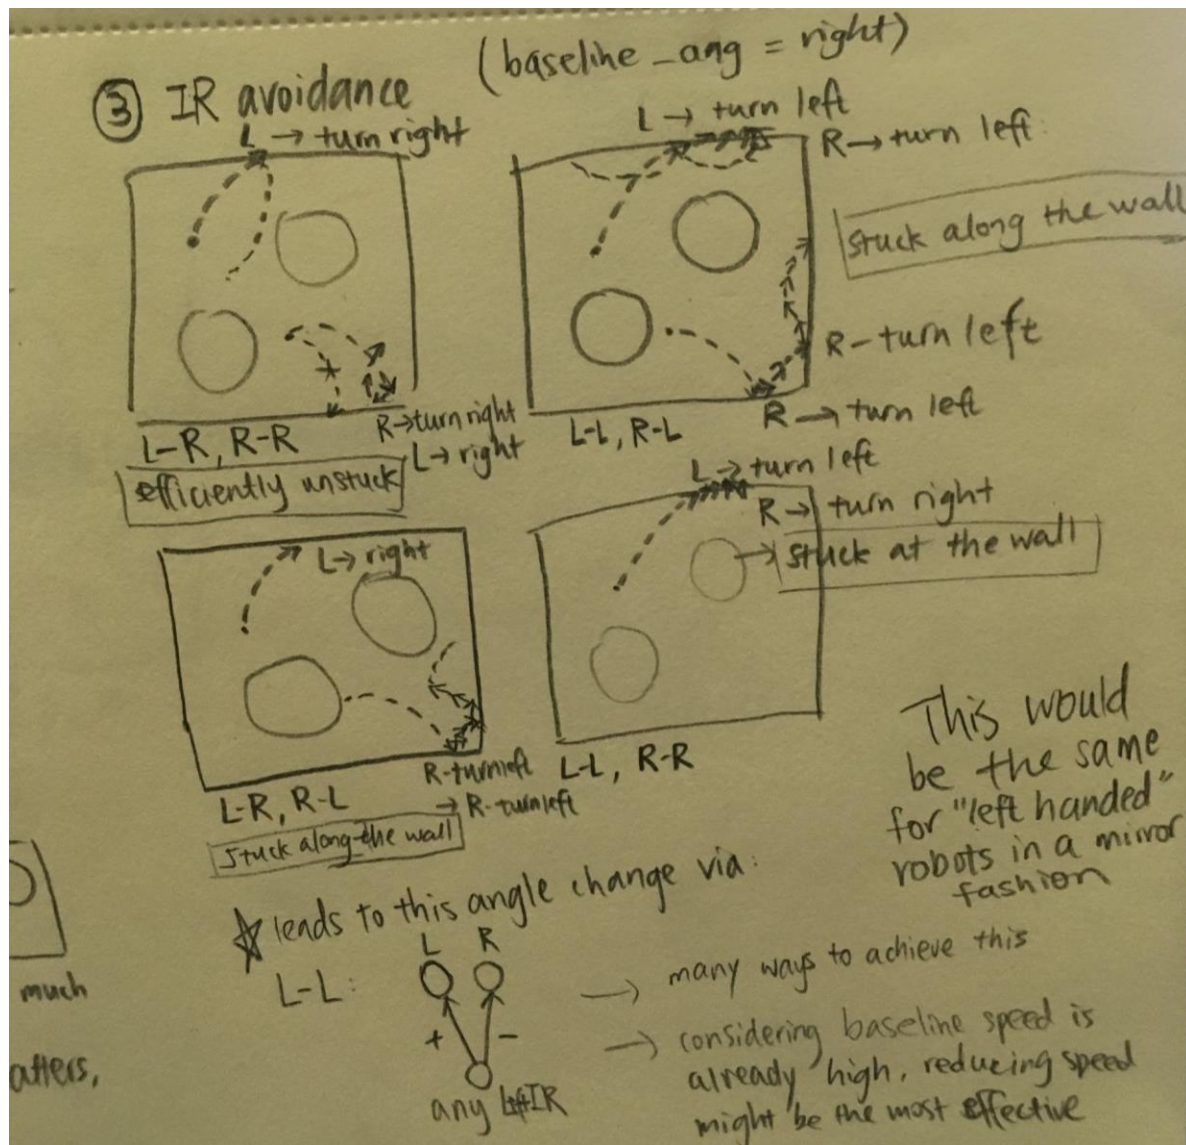


Figure 13. Possible IR-driven wall avoidance behaviors in the system. Assuming that an agent is right-handed, the most efficient way to get unstuck from the wall is by turning right upon the perception of the wall from IR sensors on the left side (IR 0, IR 1) or IR sensors on the right side of the body (IR 4, IR 5). If the left IR sensors lead to a net left turn, and/or the right IR sensors lead to a net right turn, the agent is at a higher risk of getting stuck along the wall.

Considering the traits implicated in these four behaviors, the functional behavior traits consists of the following:

- Baseline_comm: same as baseline_comm intensity from the network level
- Baseline_speed: same as the baseline_speed from the network level
- Baseline_turn: the absolute value of the baseline_turn from the network level
(ignores handedness)
- Ground_baseline_comm: same as the ground_baseline_comm from the network level
- Ground_baseline_speed: same as the ground_baseline_speed from the network level
- Ground_baseline_turn: absolute value of ground_baseline_turn from the network level
- Ground_speed_delta: same as the ground_speed_delta from the network level
- Ground_comm_delta: same as the ground_speed_delta from the network level
- IR_comm_delta: the sum of the change in comm unit output intensity from all IR sensors
- IR_left_turn: the net turning in the right direction of IR_0 and IR_1 sensors
(towards the right for right-handed agents and left for left-handed agents)
- IR_right_turn: the net turning in the right direction of IR_4 and IR_5 sensors
(towards the right for right handed agents and left for left handed agents)
- IR_left_speed: the net change in speed IR_0 and IR_1 sensors
- IR_right_speed: the net change in speed IR_4 and IR_5 sensors
- IR_back_speed: the net change in speed IR_6 and IR_7 sensors

- Comm_comm_delta: the average change in comm signaling from all the comm sensors
- Comm_speed_delta: the average change in motor speed across all the comm sensors
- comm_0_turn_delta, turning (in radian) upon receiving a signal in the right direction (in right for right handed agents, in left for left handed agents)
- comm_1_turn_delta, turning (in radian) upon receiving a signal in the right direction (in right for right handed agents, in left for left handed agents)
- comm_2_turn_delta, absolute value of turning (in radian) upon receiving a signal
- comm_3_turn_delta: turning (in radian) upon receiving a signal in the right direction (in right for right handed agents, in left for left handed agents)

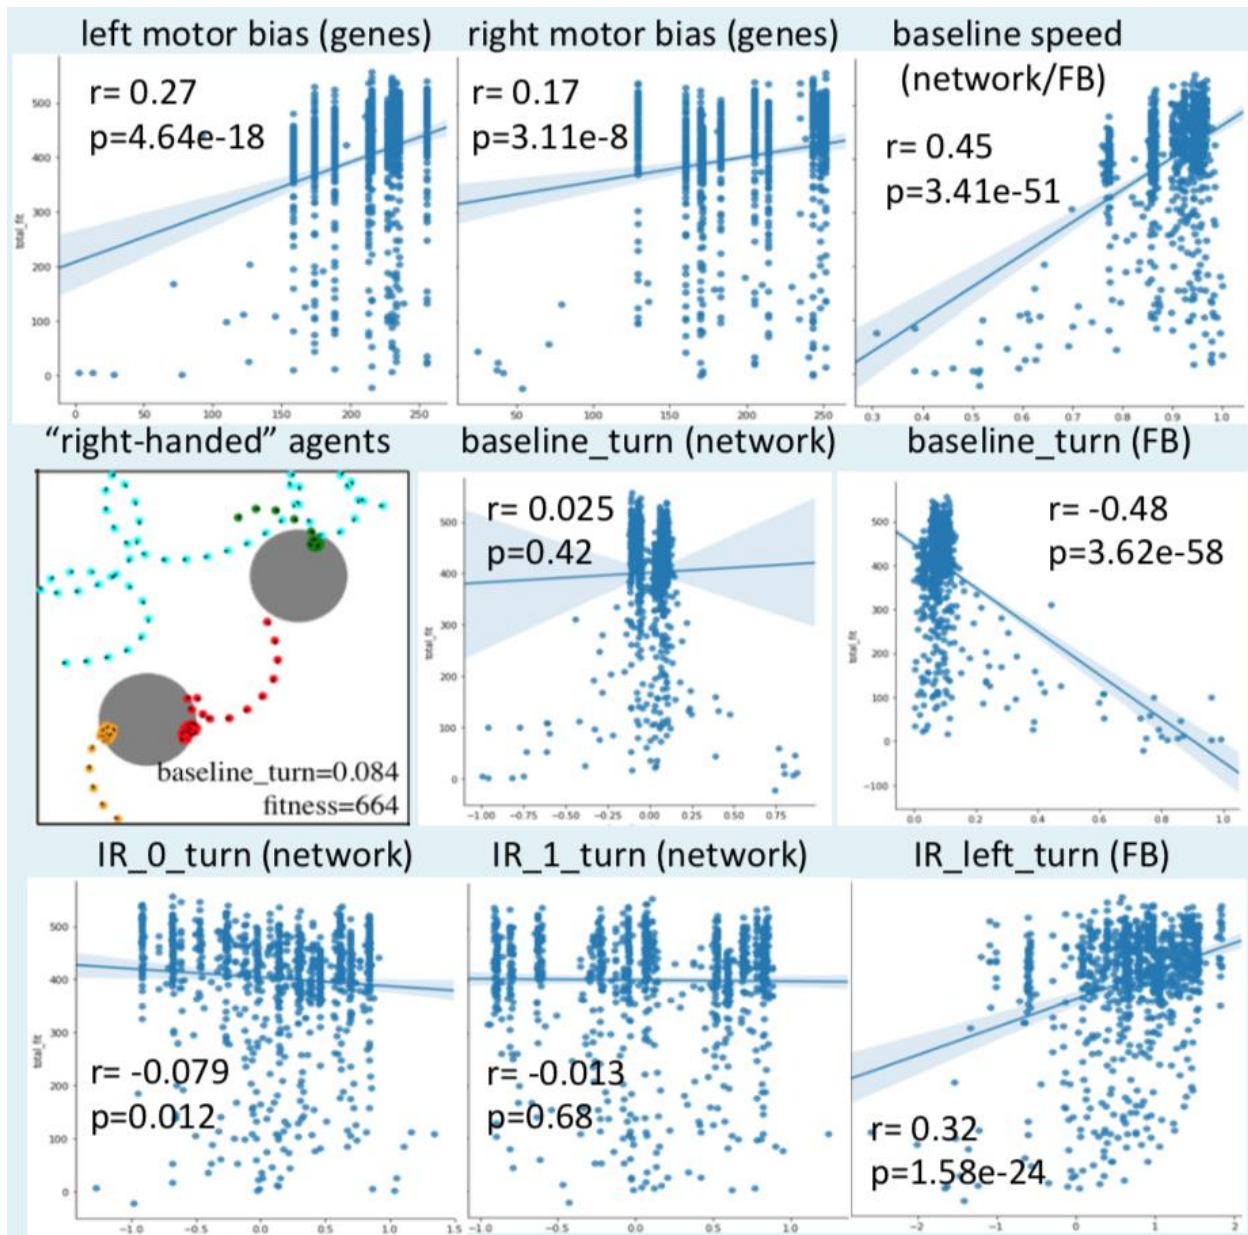


Figure 14. Pearson's R correlation tests show that functional behavior level traits capture linear relation between traits and fitness that are not capture on the genes level or the network level. Data are from the last generation of condition 1, all runs included. The x-axes represent the relative value of the trait, and the y-axes represent average fitness.

This list of functional behaviors is far from exhaustive, and makes many arbitrary assumptions. It should be viewed as a preliminary investigation into a complex problem rather than an well-formed hypothesis, but it's a start, and it does capture some of the more straightforward relationships between traits and fitness that the bottom levels failed to capture (Fig. 14). Lande & Arnold (1983) pointed out that natural selection acts on phenotypes regardless of their genetic basis, but evolutionary response to selection is dependent on genetic variations and shifts that occurs from one generation to the next. The complex question here is how do these levels connect with each other. When studying communication or any other complex behaviors in biological animals, we often do not have access to data across many generations, nor are we usually able to obtain comprehensive information about both the physiological/mechanistic and the environmental factors that contributes to their fitness (Lande & Arnold, 1983; Greenfield, 2014). These barriers are less problematic in a simple artificial system granting access to all genetic and phenotypic materials and capable of generating an abundance of data; the hope here is to see whether such functional level behavior serves as a good way to bridge fitness and genotypes.

Qualitative behavioral analysis. To better understand the communication and non-communication behavior exhibited by the agents, trials are performed and recorded for the best performing genotype of each run, and the sensor and motor outputs are visualized to guide the analyses.

Results

Effect of comm_self neurons access on fitness in the final generation. A two-way repeated measures ANOVA revealed a significant main effect of treatment for comm_others sensors (comm_other: $M = 401.96$, $SD = 66.79$; no comm_other: $M = 297.11$, $SD = 5.89$), $F(1, 9) = 479.92$, $p < .0001$ (Fig. 15). No significant main effect was found for comm_self neurons or the interaction.




Figure 15. A significant effect of access to comm_others sensors on fitness at the final generation was found via a two-way repeated measures ANOVA. The agents evolved under condition 1 (comm_self, comm_others) and 2 (no comm_self, comm_others) achieved significantly higher fitness scores than those evolved under condition 3 (comm_self, no comm_others) and 4 (no comm_self, no comm_others). ***: $p < .0001$

Successful replication of the effect of comm_others sensor on fitness. Two paired-samples t-tests were conducted to confirm that general findings in Marocco and Nolfi (2007) that agents that evolved in condition 1 (the “normal” condition) had significantly higher fitness after 100 generations of adaptive evolution ($t(9) = 14.07$, $p < .0001$) (Fig. 16), and that when these same agents are tested again without access to comm_others sensors, the performance decreased significantly ($t(9) = 11.41$, $p < .0001$) (Fig. 16).




Figure 16. Agents evolved under the condition 1 (labeled “normal”; comm_self, comm_others) performed significantly better than both agents evolved under condition 3 (labeled “No Comm”, comm_self, no comm_others) and the same agents evolved under condition 1 but tested without comm_others access (labeled “Deprived”). ***: $p < .0001$

Fitness increases throughout evolution. In the overall trend (average across 10 experimental runs), the increase of fitness across generation appears stable, with a higher rate of increase in condition 1 and 2 (due to the access to comm_others) (Figure 6). However, when examining fitness trend across generations within a single run, there are much more variations between each generation and the next idiosyncratic to each specific run. Additionally, it seems that the rates of increase in fitness in condition 1, 2 start out similar as condition 3, 4, but then increase with a steep slope for 5-10 generations (staring at gen 18 for cond 2 and gen 30 for cond 1), after which they return to a plateau. This may imply that a small number of traits that may be achieved via a few mutations underlies the increased fitness in conditions with comm sensor access (Fig. 17; Fig. 18).




Figure 17. The trend of fitness across 100 generations of evolution. X-axis: generation number; Y-axis: fitness score (maximum possible score is 1000). Top: the trend of average fitness trend across all 10 runs. The two comm_others conditions (1, 2) seems to evolve at a faster rate than the two no comm_others conditions (3, 4) in the first 25 or so generations; they also reaches a higher fitness at the plateau. Bottom: the trend of average fitness for a single run, run 0. When examining the fitness trajectory of an individual run, it seems that the initial rate of increase did not differ between the four conditions, but groups in conditions 1 and 2 seems to acquire new traits through mutation that allowed fitness to spur up rapidly, reaching a higher plateau. Red: condition 1 (comm_self + comm_other); orange: condition 2 (no comm_self + comm_others); blue: condition 3 (comm_self + no comm_others), green: condition 4 (no comm_self + no comm_others).

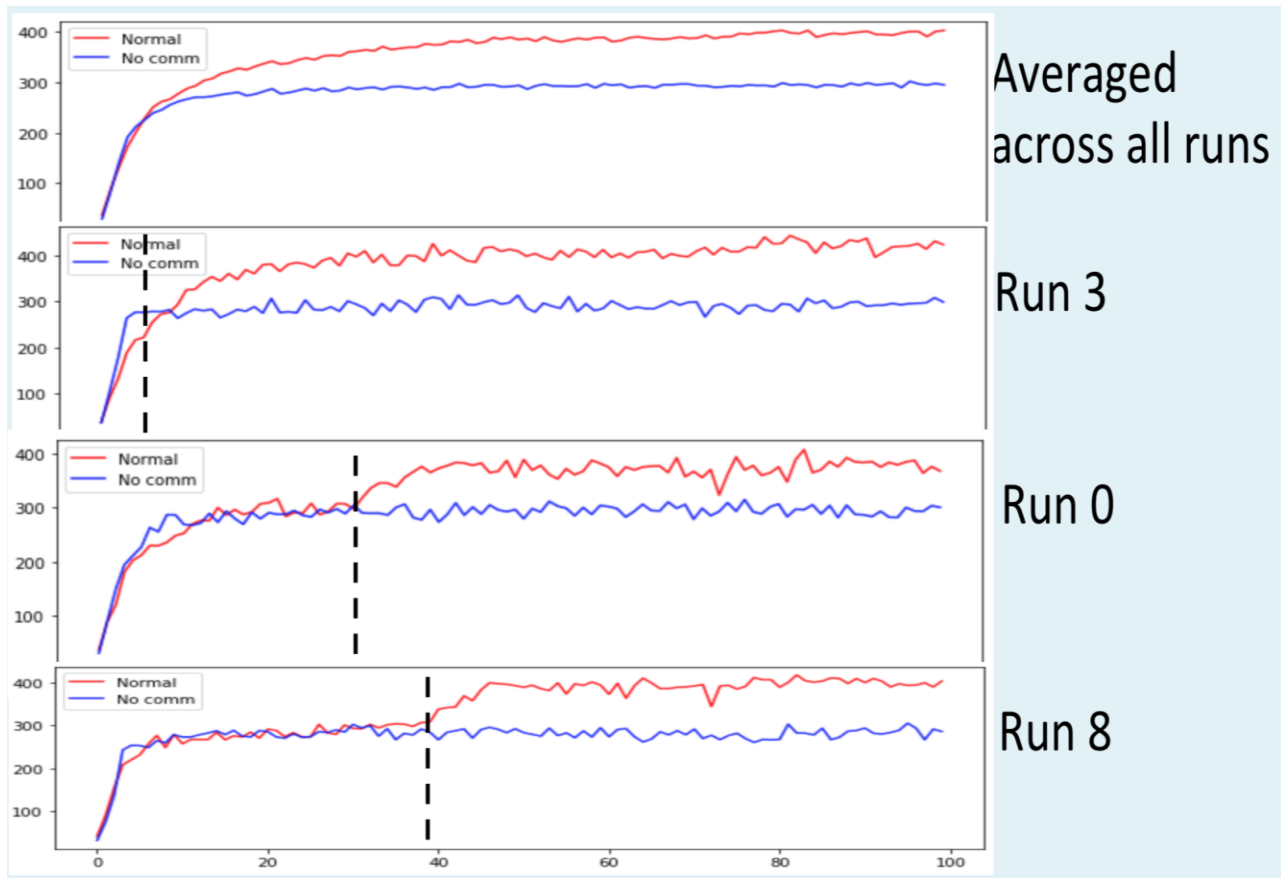


Figure 18. Fitness trends overall and in representative runs for condition 1 (normal) and condition 3 (no comm). While the fitness advantage in the normal condition seemed to emerge smoothly and is present in the beginning, this was not always the case for each individual run. In some of the runs, the fitness trend increased at a similar rate as in the no comm condition in the beginning of the evolution; the advantage only emerged later and evolved quickly before reaching a plateau again. The Black dash lines mark the initial time point when this advantage seemed to have evolved.

Upper level functional analyses led to better models for fitness. Model statistics showed that for every condition, at both evolutionary time points, the functional behavior level analysis leads to the best model, with the largest R^2 and adjusted R^2 (indicating percentation of variance within data explained by the model), as well as the smallest Akaike's information criterion (AIC) and Bayesian information criterion (BIC) (penalizes model complexity; smaller values indicate a model is a better fit; Kuha, 2004). The network level analysis is a relatively better model for than the genotype level analysis.

Additionally, models at all three levels appeared to explain more of the variance within the data for the evolved generation (Table 2) rather than the initial generation (Table 3), as indicated by larger R^2 and adjusted R^2 values. For example, in condition 1, the adjusted R^2 for the genes, network and functional behavior levels are 0.193, 0.266, and 0.356, respectively, for the initial random generation. For the evolved final generation, the adjusted R^2 were 0.199, 0.494, and 0.503, respectively. However, the complexity of the models also increased after evolution. For example, in condition 1, BIC for the genes, network, and functional behavior analyses are 10370, 10120, and 9846, respectively. For the same condition, the BIC of the three levels are 12400, 11790, and 11630, respectively. This pattern is consistent across all four conditions.

Table 2. Selection gradient analysis of the random initial generation across conditions.

Cond.	Level	Random Initial Generation						
		R^2	Adj. R^2	F-statistics	p-value	Log-likelihood	AIC	BIC
1	Genes	0.248	0.193	4.456	2.21e-26	-4943.7	10030	10370
	Network	0.297	0.266	9.617	5.24e-49	-4910.5	9907	10120
	fBehav.	0.368	0.356	31.71	6.26e-85	-4857.2	9752	9846
2	Genes	0.227	0.176	4.436	3.40e-24	-5002.9	10130	10440
	Network	0.259	0.227	7.969	2.50e-39	-4981.6	10050	10260
	fBehav.	0.337	0.325	27.70	4.44e-75	-4926.1	9890	9984

3	Genes	0.229	0.172	4.002	2.23e-22	-4968.3	10080	10420
	Network	0.297	0.266	9.617	5.24e-49	-4910.5	9907	10120
	fBehav.	0.368	0.356	31.71	6.26e-85	-4857.2	9752	9846
4	Genes	0.227	0.176	4.434	3.56e-24	-4942.4	10010	10320
	Network	0.294	0.263	9.478	3.34e-48	-4897.1	9880	10090
	fBehav.	0.463	0.453	46.99	6.70e-119	-4760.1	9558	9651

Table 3. Selection gradient analysis of the final generation after 100 generations of adaptive evolution across conditions.

		After 100 Generations of Evolution						
Cond.	Level	R ²	Adj. R ²	F-statistics	p-value	Log-likelihood	AIC	BIC
1	Genes	0.254	0.199	4.596	1.28e-27	-5960.1	12060	12400
	Network	0.515	0.494	24.19	2.08e-121	-5745.1	11580	11790
	fBehav.	0.512	0.503	57.16	6.94e-139	-5748.2	11530	11630
2	Genes	0.312	0.267	6.865	6.76e-44	-5847.4	11820	12130
	Network	0.549	0.529	27.72	6.25e-136	-5636.6	11360	11570
	fBehav.	0.602	0.594	82.37	1.10e-181	-5574.2	11190	11280
3	Genes	0.163	0.101	2.633	9.74e-11	-5810.1	11760	12100

	Network	0.515	0.494	24.19	2.08e-121	-5745.1	11580	11790
	fBehav.	0.512	0.503	57.16	6.94e-139	-5748.2	11530	11630
4	Genes	0.217	0.165	4.194	3.21e-22	-5839.6	11810	12110
	Network	0.421	0.395	16.55	2.89e-86	-5689.1	11460	11680
	fBehav.	0.430	0.419	41.10	2.06e-106	-5681.0	11400	11490

Different behaviors are selected during different stages of evolution. In the beginning of the evolution, traits that significantly contribute to the overall fitness of the individual seems highly driven by a high baseline speed (Table 5). At the end of the evolution, more traits involved in tropotaxis seem to be going through selection and starting to underlie fitness (Table 6). Below are examples from condition 1. Traits were sorted into four categories based on their functional contribution: exploratory motor behaviors (color coded blue), stopping at target area (green), IR navigation (red), and communication (yellow) (Table 5, Table 6).

In condition 1, at the beginning of the evolution, fitness seemed to be driven mostly by exploratory motor traits (6 on the gene level, 2 on network, 4 on functional behaviors, while others categories only had 1 or 2 on each level) (Table 4). At the end of the evolution, these explorative motor traits are no longer the only main contributor to the fitness as number of other traits that significantly contributed to the fitness grew. Many communication sensors related traits contributed to fitness on the level of the network comparing to other groups (5 on the gene level, 14 on the network level, 2 on the behavior level) (Table 4), and four of these are modulated by IR sensors. It's worth noting, nonetheless, that the data for the last generation are

likely not parametric as genome of successful agents become overrepresented in the data; this may undermine the efficacy of the significance based statistical tests.

Table 4. A summary table of the count for four categories of traits that significantly contributed to fitness at the beginning and the end of the evolution in condition 1 (normal access to comm).

# of traits	Exploratory	Stopping at Target	IR navigation	Communication
The Beginning of Evolution				
Gene	6	2	2	1
Network	2	1	2	2
F. Behavior	4	2	1	1
The End of Evolution				
Gene	2	3	4	5
Network	2	2	5	14
F. Behavior	1	3	2	2

Table 5. Traits that significantly contribute to fitness in condition 1 at the first generation. The traits are clustered into functional groups and color coded; blue: exploration; green: stopping or slowing at target area; red: wall avoidance via IR sensors; yellow: communication behaviors, especially ones that simulate signal B in Marocco and Nolfi (2007).

level	name	Coefficient	Std. Err	t	p-value	Confidence interval
gene	motor_left bias	0.1199	0.015	7.945	0.000	[0.090, 0.149]
gene	motor_right bias	0.1335	0.016	8.611	0.000	[0.103, 0.164]
gene	IR_0 to motor_left	0.0326	0.015	2.184	0.029	[0.003, 0.062]
gene	IR_2 to motor_right	-0.0305	0.015	-1.972	0.049	[-0.061, -0.000]
gene	comm_3 to motor_right	0.0421	0.016	2.686	0.007	[0.011, 0.073]
gene	ground to motor_left	-0.0755	0.016	-4.751	0.000	[-0.107, -0.044]
gene	ground to motor_right	-0.0656	0.016	-4.173	0.000	[-0.097, -0.035]
gene	comm_self to motor_right	0.0472	0.016	2.967	0.003	[0.016, 0.078]
gene	internal_1 to motor_left	0.0414	0.016	2.584	0.01	[0.010, 0.073]
gene	internal_1 to motor_right	0.0586	0.015	3.842	0.000	[0.029, 0.089]
gene	internal_2 to internal_1	-0.0431	0.016	-2.768	0.006	[-0.074, -0.013]
network	IR_0_turn_delta	7.374	2.75	2.682	0.007	[1.978, 12.770]
network	IR_5_turn_delta	-6.0421	2.671	-2.262	0.024	[-11.284, -0.800]
network	baseline_comm	10.0223	3.799	2.638	0.008	[2.566, 17.478]
network	baseline_speed	51.8329	3.685	14.068	0.000	[44.602, 59.064]
network	comm_1_turn_delta	11.4171	2.74	4.167	0.000	[6.040, 16.794]

network	comm_3_speed_delta	22.3623	5.389	4.15	0.000	[11.787, 32.938]
network	ground_speed_delta	-47.1176	3.76	-12.533	0.000	[-54.495, -39.740]
F. behavior	baseline_comm	8.8156	3.425	2.574	0.01	[2.093, 15.538]
F. behavior	baseline_speed	50.9278	3.429	14.854	0.000	[44.200, 57.656]
F. behavior	baseline_turn	-38.1515	3.822	-9.981	0.000	[-45.652, -30.651]
F. behavior	ground_baseline_turn	11.7191	3.142	3.73	0.000	[5.553, 17.885]
F. behavior	ground_speed_delta	-46.5155	3.45	-13.482	0.000	[-53.286, -39.745]
F. behavior	IR_left_turn	3.3977	1.67	2.035	0.042	[0.121, 6.675]
F. behavior	comm_speed_delta	9.596	2.286	4.197	0.000	[5.110, 14.082]
F. behavior	comm_2_turn_delta	16.1638	3.731	4.333	0.000	[8.843, 23.485]

Table 6. Traits that significantly contribute to fitness in condition 1 after 100 generations. The traits are clustered into functional groups and color coded; blue: exploration; green: stopping or slowing at target area; red: wall avoidance via IR sensors; yellow: communication behaviors, especially ones that simulate signal B in Marocco and Nolfi (2007).

level	name	Coefficient	Std. Err	t	p-value	Confidence interval
genes	ground time const	0.4762	0.116	4.091	0.000	[0.248, 0.705]
genes	motor_left bias	1.7464	0.214	8.143	0.000	[1.326, 2.167]
genes	motor_right bias	1.1748	0.226	5.207	0.000	[0.732, 1.618]
genes	comm_unit bias	-1.1733	0.259	-4.524	0.000	[-1.682, -0.664]
genes	IR_4 to motor_right	-0.4766	0.146	-3.267	0.001	[-0.763, -0.190]
genes	IR_5 to motor_left	-0.2974	0.129	-2.301	0.022	[-0.551, -0.044]

genes	IR_6 to motor_left	0.2672	0.116	2.302	0.022	[0.039, 0.495]
genes	IR_6 to motor_right	0.1765	0.08	2.218	0.027	[0.020, 0.333]
genes	comm_0 to motor_right	-0.3367	0.139	-2.421	0.016	[-0.610, -0.064]
genes	comm_2 to motor_left	-0.3348	0.144	-2.33	0.02	[-0.617, -0.053]
genes	comm_3 to motor_right	0.3299	0.142	2.327	0.02	[0.052, 0.608]
genes	ground to motor_left	-1.3481	0.197	-6.828	0.000	[-1.736, -0.961]
genes	ground to motor_right	-0.5703	0.139	-4.092	0.000	[-0.844, -0.297]
genes	comm_self to comm_unit	-0.2733	0.137	-1.994	0.046	[-0.542, -0.004]
genes	internal_1 to motor_left	-0.2529	0.116	-2.182	0.029	[-0.480, -0.025]
network	IR_0_comm_delta	136.5452	29.905	4.566	0.000	[77.858, 195.232]
network	IR_0_speed_delta	-186.2164	29.066	-6.407	0.000	[-243.256, -129.177]
network	IR_1_speed_delta	-68.4565	22.389	-3.058	0.002	[-112.394, -24.519]
network	IR_1_turn_delta	34.3714	10.396	3.306	0.001	[13.969, 54.774]
network	IR_2_comm_delta	120.2839	35.183	3.419	0.001	[51.239, 189.329]
network	IR_3_comm_delta	-123.6881	32.805	-3.77	0.000	[-188.066, -59.310]
network	IR_3_speed_delta	-47.6052	23.744	-2.005	0.045	[-94.202, -1.009]
network	IR_4_comm_delta	95.1865	40.833	2.331	0.02	[15.054, 175.319]
network	IR_5_speed_delta	-76.8256	24.013	-3.199	0.001	[-123.949, -29.702]
network	IR_5_turn_delta	-47.4201	11.657	-4.068	0.000	[-70.296, -24.544]

network	IR_6_turn_delta	27.7373	12.342	2.247	0.025	[3.516, 51.958]
network	IR_7_comm_delta	-72.3199	36.752	-1.968	0.049	[-144.444, -0.196]
network	baseline_comm	-43.8289	19.404	-2.259	0.024	[-81.909, -5.749]
network	baseline_speed	468.6745	35.723	13.12	0.000	[398.569, 538.780]
network	comm_0_speed_delta	-70.3098	29.468	-2.386	0.017	[-128.140, -12.480]
network	comm_0_turn_delta	37.6307	15.864	2.372	0.018	[6.498, 68.764]
network	comm_1_comm_delta	-117.8412	54.4	-2.166	0.031	[-224.597, -11.085]
network	comm_1_speed_delta	162.2492	40.083	4.048	0.000	[83.589, 240.910]
network	comm_1_turn_delta	74.2255	18.503	4.012	0.000	[37.915, 110.536]
network	comm_3_speed_delta	246.4665	34.055	7.237	0.000	[179.635, 313.298]
network	comm_3_turn_delta	-97.4111	16.932	-5.753	0.000	[-130.639, -64.183]
network	ground_baseline_comm	76.401	20.236	3.776	0.000	[36.689, 116.113]
network	ground_baseline_turn	-31.5209	10.119	-3.115	0.002	[-51.379, -11.663]
network	ground_comm_delta	120.23	29.815	4.032	0.000	[61.719, 178.741]
network	ground_speed_delta	-434.1702	25.581	-16.97	0.000	[-484.372, -383.969]
f. behavior	baseline_turn	-418.3178	46.85	-8.929	0.000	[-510.255, -326.381]
f. behavior	ground_baseline_speed	-167.3341	28.043	-5.967	0.000	[-222.365, -112.303]
f. behavior	ground_baseline_turn	201.8385	15.107	13.361	0.000	[172.193, 231.484]
f. behavior	ground_speed_delta	-186.1495	36.073	-5.16	0.000	[-256.938, -115.361]

f. behavior	IR_left_turn	22.7378	6.327	3.594	0.001 [10.321, 35.154]
f. behavior	IR_right_speed	-48.7123	12.281	-3.967	0.000 [-72.812, -24.613]
f. behavior	comm_0_turn_delta	38.0933	9.403	4.051	0.000 [19.642, 56.545]
f. behavior	comm_2_turn_delta	64.2171	18.939	3.391	0.000 [27.051, 101.383]

Comparisons between fitness trends and behavioral trait selection in selective runs. In order to better understand the advantage that develops later and rapidly, the values for a series of traits at every 5 generations were normalized and plotted along with fitness for Run 0 (Fig. 19).

Communication traits, specifically turning upon comm 0 (front), comm 2 (back) and comm 3 (right) activations and communication signal level when in target area (ground baseline comm), seemed to have changed right before or right after the sudden increase in fitness (Fig. 19).

Additionally, IR right turn also underwent selection around this time (Fig. 19). These changes are not seen in condition 3 (Fig. 19). As elaborated in the next section, the result of this set of behavioral trait is differential avoidance behavior towards signal perceived as coming in front of the agent depending on the presence of signal perceived in the back (Fig. 22).

Before this sudden increase happened, the exploratory and ground sensor-based traits for stopping at target areas had already evolved (Fig. 20). The evolution of these communication traits were not universal across runs; for example, in some other runs, ground baseline communication did not go through adaptive evolution (Fig. 21). For runs where this trait did not go through evolution, it is possible that this trait was already present in the initial random genome at the beginning of the run, or that this trait was not necessary for the evolution of functional behaviors that lead to the increase in fitness in condition 1. As elaborated in the next section, the answer may be inconclusive and a mix of both.

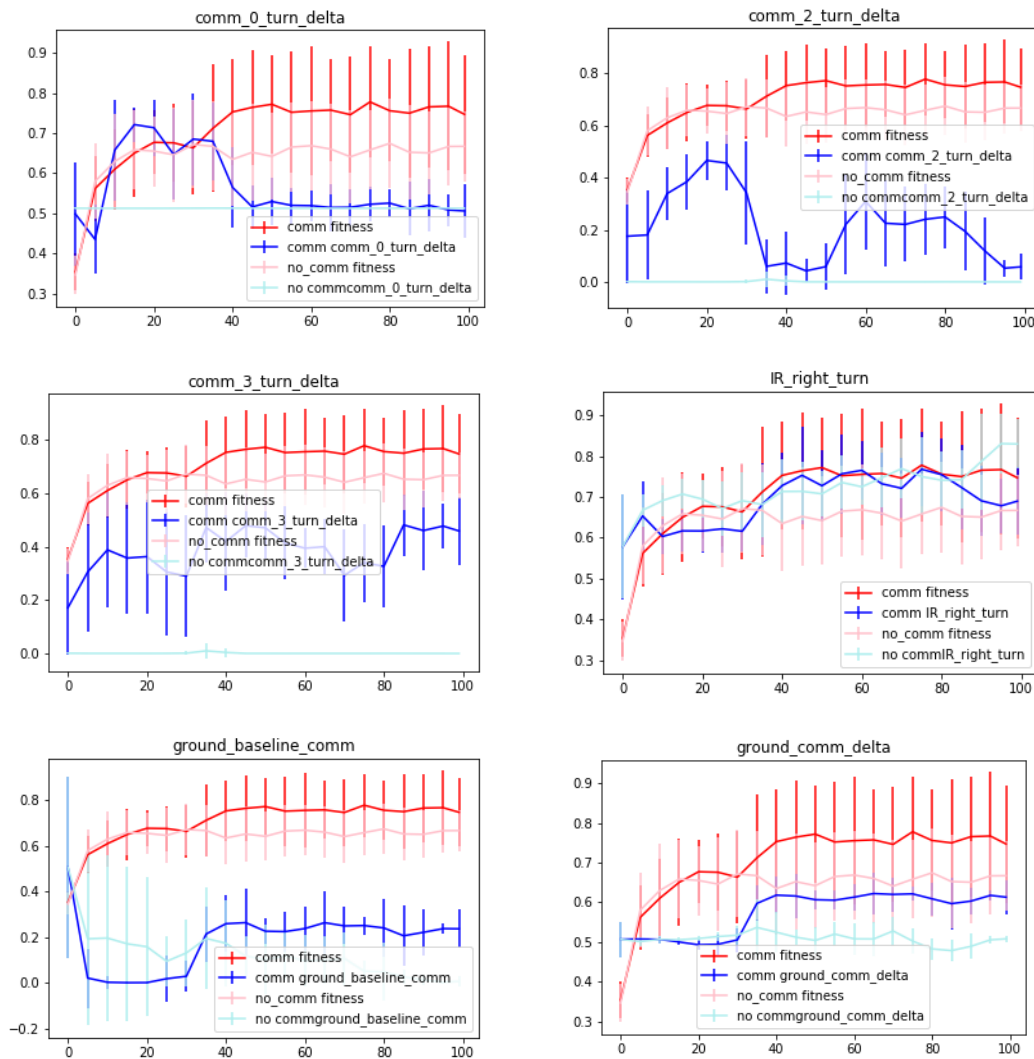


Figure 19. Traits that varied corresponding to the change in fitness at around generation 30 in Run 0. Some of these, such as comm_2_turn_delta (top right), seems to have been selected since before the increase in fitness, while most others concurrently with fitness, and comm_0_turn_delta (top left) shortly after the increase in fitness. Red lines represents average fitness, and blue line the average trait value. Fitness score and trait value for the same generation in condition 3 are presented in pink and cyan for reference. Data are normalized between condition 1 and 3. The vertical error bars represent standard errors.

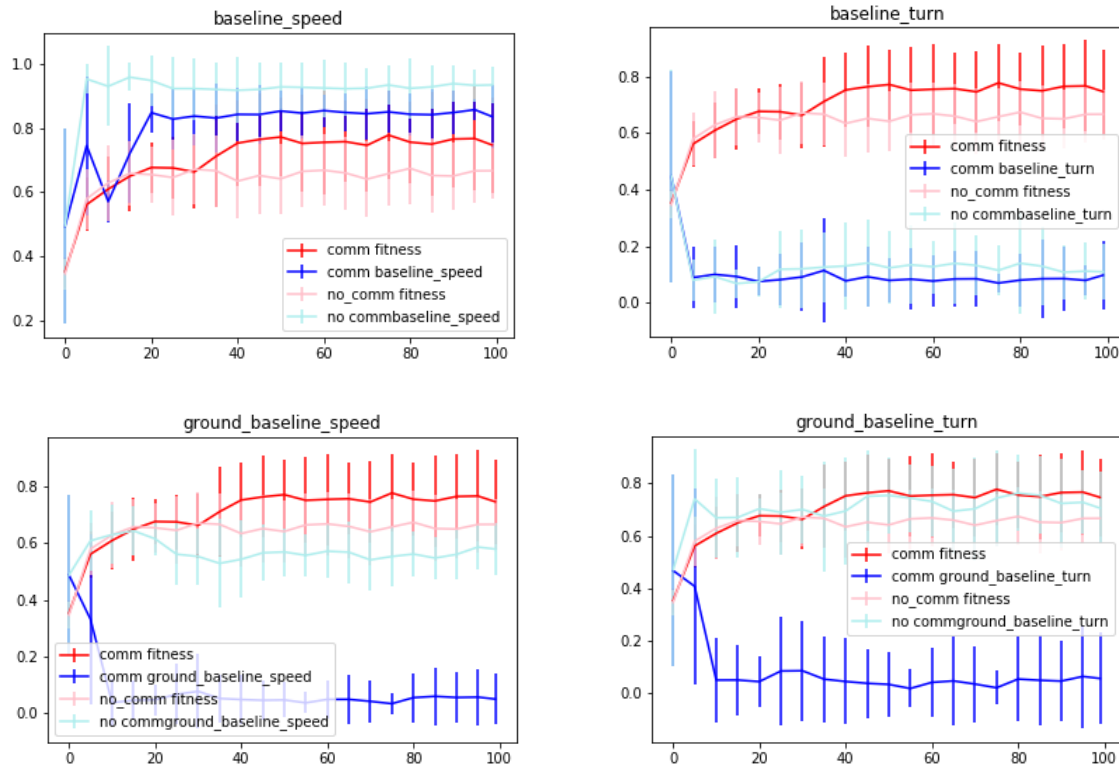


Figure 20. Exploratory and stopping at target motor behaviors during early generations in Run 0. All of these traits have undergone selection and reached their final values in the first 10 or 20 generations, considerably earlier than the increase in fitness occurred around generation 30. Fitness score and trait value for the same generation in condition 3 are presented in pink and cyan for reference. Data are normalized between condition 1 and 3. The vertical error bars represent standard errors.

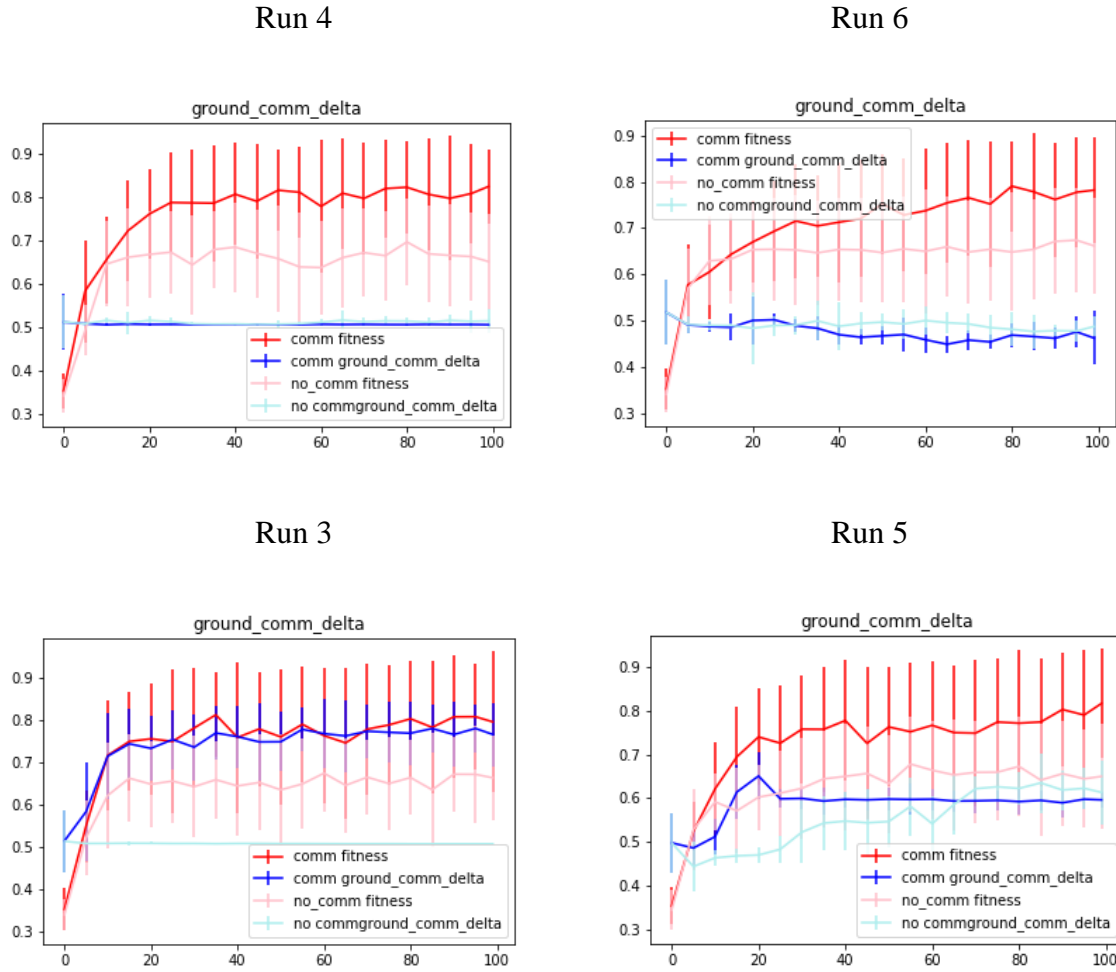


Figure 21. Correspondence between fitness and trait value for `ground_comm_delta` in selective runs.

`Ground_comm_delta` is one of the traits that went through considerable selection around the time of fitness increase in Run 0; it showed varying patterns of selection in other runs. In some runs such as Run 4 (top left) and Run 6 (top right), it did not appear to go through selection at all. In some runs such as Run 3 (bottom left) and Run 5 (bottom right), it corresponded with the increase in fitness.

Qualitative trial analysis for selective runs. Across the ten different runs, a variety of strategies evolved. A commonly occurring strategy is to avoid signals from other agents, especially through increased turning upon perception of a signal through the frontal sensor comm 0 (Fig. 22). Notably, this signal avoidance behavior may depend on the presence of other signals in the environment; for example, in the best performing team in Run 0, an agent would not avoid a signal in front of it if there is also a signal at its back (perceived through the back-facing comm 2 sensor) (Fig. 22). This context-dependent signal avoidance behavior is seen across runs; in most runs this is coupled with the emission of a more intense signal when an agent is inside the target area (Fig. 22), and occasionally an even more intense signal when two agents are inside the same target area (Fig. 22), but this is not seen in all runs (Fig. 23).

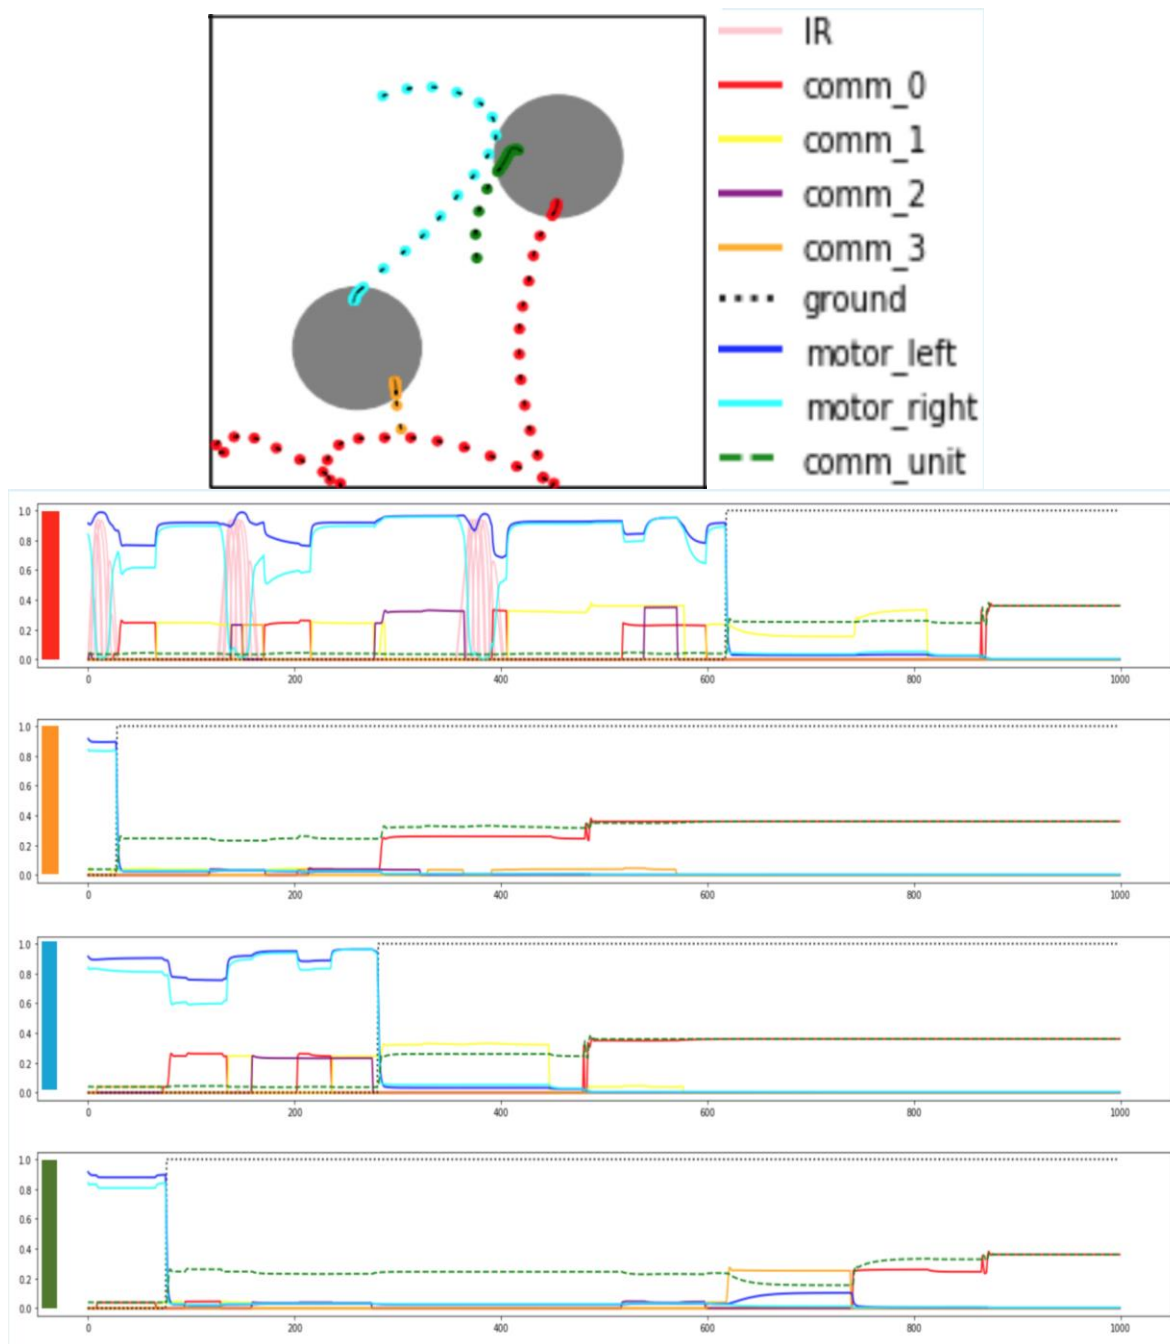


Figure 22. A representative trial from the best performing group in Run 0. Top figure shows the path of the agent across time, and the bottom panels show the state of their sensory inputs and motor outputs (legend on top right). In this trial, agent 1 (yellow) and agent 3 (green) each arrived at a target area shortly after the beginning, while agent 0 (red) and agent 2 (cyan) found their way to a target area later; both went to a target area that was further away than the one they were close to. Agent 0 avoided the target area with agent 1 when it perceived agent 1's signal from its frontal comm 0 sensor, eliciting a right turn caused by a decrease in right wheel speed. However, later when it was approaching the top target area with agent 3 in it, because of a concurrent perception of the signal from agent 0 in the back comm 2 sensor, the perception of agent 3's signal from its frontal comm 0 sensor did not have the same effect. The same perceptual behavioral pattern may be seen in agent 2. The fitness of the current trial is 760, and the total fitness of the team is 524 (out of 1000). Animation is available at <https://osf.io/ku5wy/>.

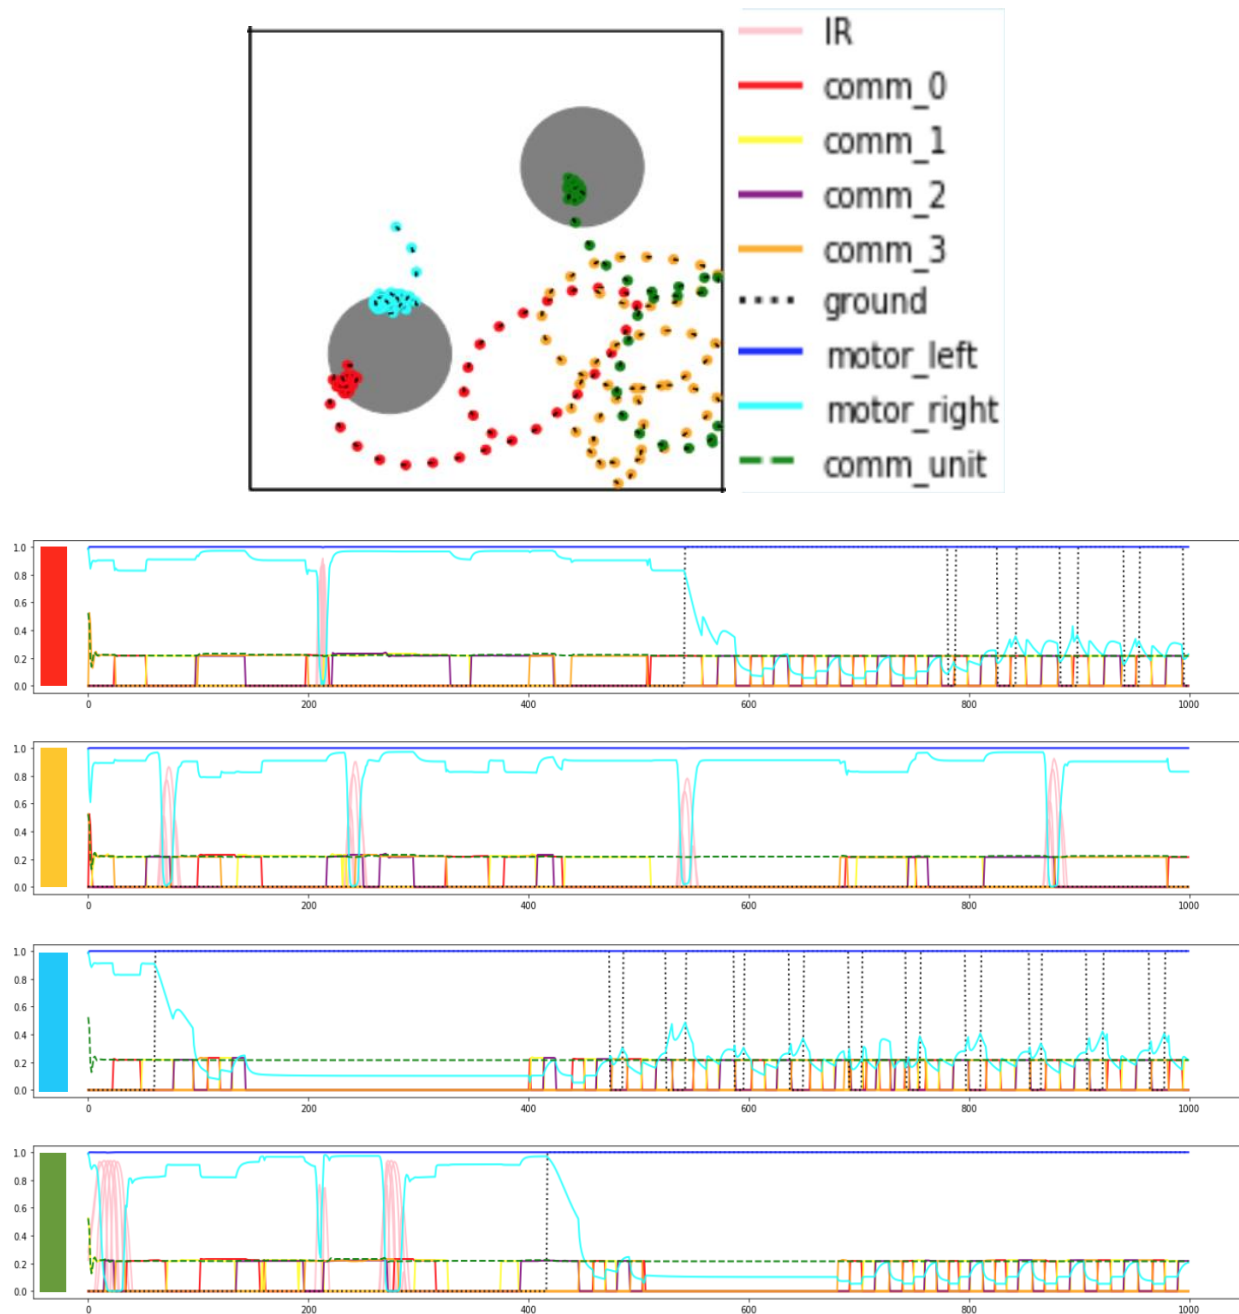


Figure 23. A representative trial from the best performing group in Run 4. Top figure shows the path of the agent across time, and the bottom panels show the state of their sensory inputs and motor outputs (legend on top right). In this trial, agent 2 (cyan) and agent 3 (green) each arrived at a target area shortly after the beginning, while agent 0 (red) and agent 1 (yellow) found their way to a target area later. Very similar to in Run 0, agent 0 avoided the target area already containing agent 2 when it perceived agent 2 via the front comm 0 (red lines) or right comm 3 (orange lines) sensors, until agent 1 happened to be behind it (purple lines), at which point it went straight instead of turning away. It's worth noting that in this group of agents, the signaling behavior of an agent did not change at all upon entering a target area; an agent will just avoid signals from other agents regardless of whether that agent is already in a target area. The fitness of the current trial is 489, and the total fitness of the team is 522 (out of 1000). Animation is available at <https://osf.io/us4ex/>.

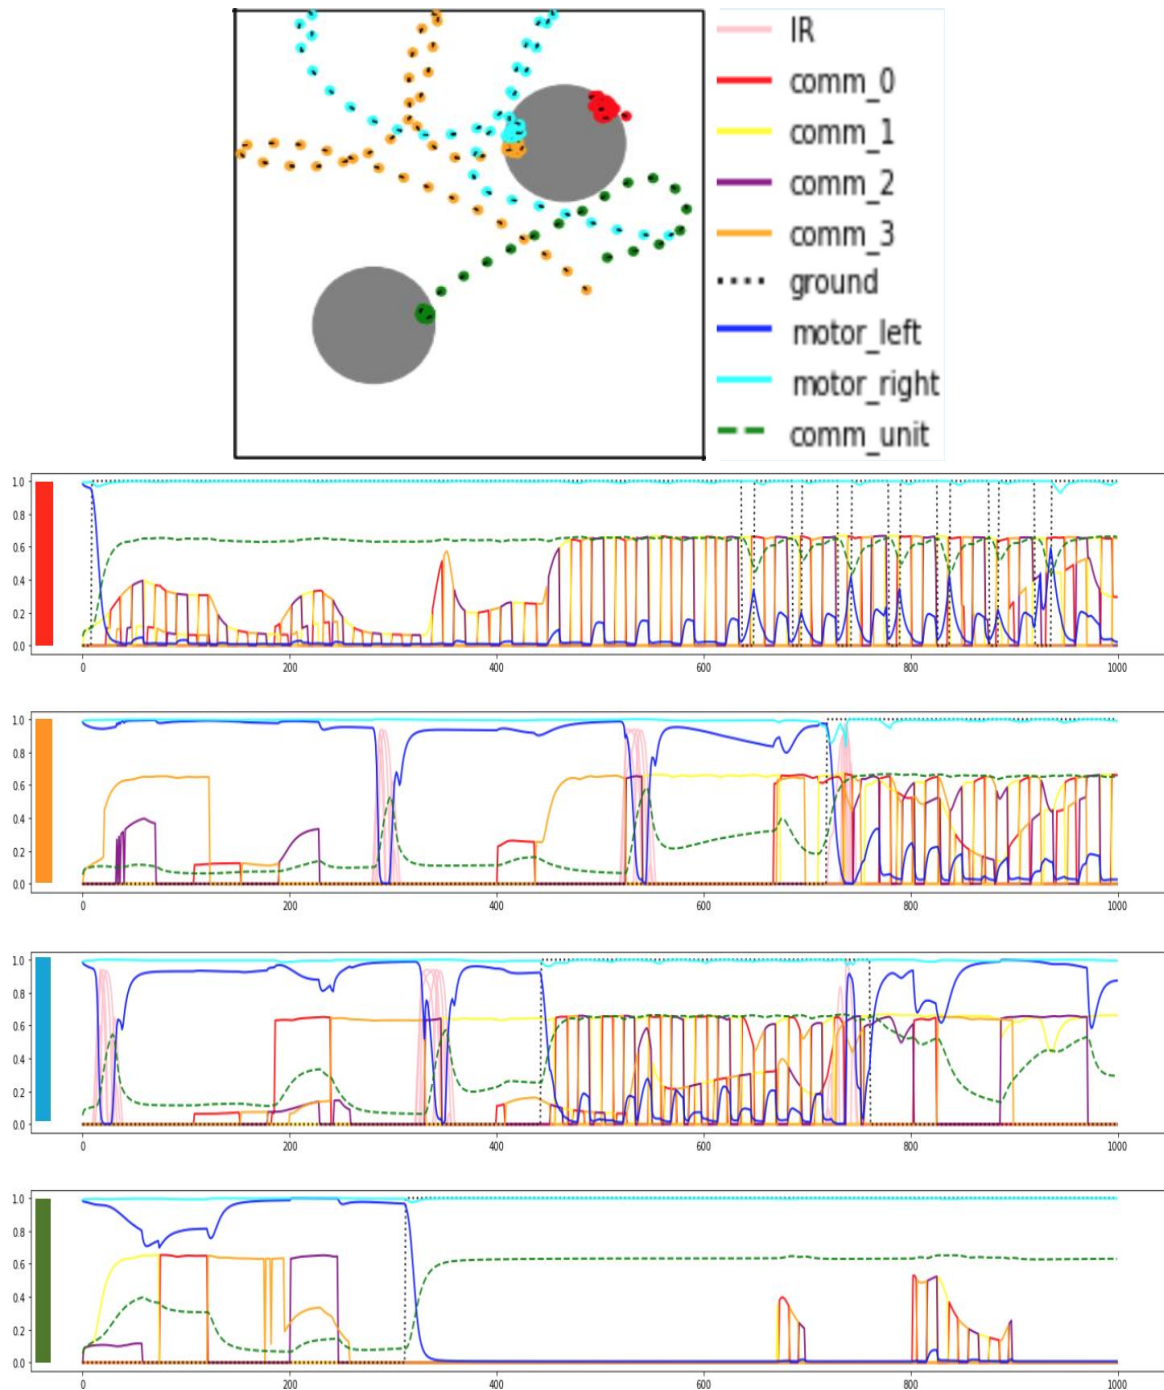


Figure 24. A representative trial from the best performing group in Run 2. Top figure shows the path of the agent across time, and the bottom panels show the state of their sensory inputs and motor outputs (legend on top right). In this trial, agent 0 (red) found the top target area soon after the trial started. Agent 3 (green) first avoided agent 0 via its left comm 1 (yellow) and front comm 0 sensors (red), and found its way to the bottom target area. Agent 2 (cyan) found its way to the top target area; later on, agent 1 (yellow) also entered the top target area and were close to agent 2 (cyan). Upon detecting agent 1 with its IR sensors, agent 2 left the target area and avoided overcrowding. The fitness of the current trial is 511, and the total fitness of the team is 546.7 (out of 1000). Animation is available at <https://osf.io/5kafv/>.

Discussion

The current project successfully reproduced the findings regarding the effect of communication in Marocco and Nolfi (2007) (Fig. 15). Some of the additional aims of the project were addressed through the analyses and are discussed below:

- (1) Does being able to hear one's own communication signal contribute to fitness in a significant way?
- (2) What is the correct level of understanding the relationship between traits and fitness?
- (3) Are different traits being selected across different stages of the evolution?
- (4) Are the same evolved functional behaviors multiply realizable through different genotypes?

No significant effect of comm_self sensor on fitness.

This first aim was to examine whether the network controller structure of the system had confounded the results in Marocco and Nolfi (2007). Results showed that having sensory access to its own communication signal from the previous times step did not affect fitness (Fig. 14). It is possible that it may lead to some minor differences in evolutionary process in some cases by virtue of affecting network behaviors given the same genotype. As we can see in a sample of individual run, the initial trajectory between condition 1, 3 and condition 2, 4 are slightly different in (Fig. 16). However, perhaps this small difference does not have a strong enough impact on the efficacy of the evolved functional behavior of the agents in different groups to cause a statistically meaningful difference. This in a way demonstrates the multiple realizability principle in that similar levels of fitness may be achieved by agents in condition 1 and 2 despite some differences in their network structure. On the other hand, the different network structure

between condition 1 and 3 did lead to a significant difference in fitness outcome. This further validates Marocco and Nolfi (2007)'s conclusion that this fitness increase is caused by the access to communication channels in condition 1. In the following analyses, we will focus mostly on condition 1 (comm) and condition 3 (no comm).

Analysis on Functional Behavior Level Led to Better Models For Fitness

The current study summarized traits on three different levels that related to each other in a bottom-up fashion. Selection gradient analyses showed that each of the upper level is better than the level below it both at the beginning (Table 1) and at the end of the evolution (Table 2). Although by a small margin, the functional behavior level exceeded the network level based on all measures of the model, explaining more proportion of the variance (as indicated by a larger R^2), a lower probability that the results happened by chance (as indicated by a larger F-statistics), and a simpler model (as indicated by AIC and BIC scores). It is worth noting that even the small advantage is already impressive, because the functional behaviors considered here are far from exhaustive. The definition of comm censor traits -- approach signals, following the definition of Marocco and Nolfi (2007) for signal B (Table 1) -- was also problematic as almost all agents in the last generation showed avoidance when perceiving signals from others instead. If better categorization of the functional behaviors may be extracted from more carefully designed ethological analysis, perhaps its advantage would increased from the marginal amount we saw here.

Moorad and Wade (2013) have discussed that R^2 measures the extent to which the fitness of individuals in a population can be accounted for by the values of some phenotypic traits; the higher R^2 measures in upper level analyses may indicate that these upper level analyses indeed

better explain the sources of variation in the fitness and captures something about the underlying mechanisms of individual fitness. One of the reasons for this is that, as we go up levels, the behavioral traits incorporate more information about the system. By quantifying motor behavior changes induced by isolated sensors, we are not only considering the time constants, the biases and the connection weights of the network encoded by the genotype, but also taking into account the mechanistic relationships between them by virtue of the fixed structure of the network controller. Then, by studying the fitness-related consequences of functional behaviors of the agents with this network, we are not only examining the exact motor behaviors of the network, but rather also the fitness-related consequences of the agent's behavior, which can not be analyzed without consideration of the size of the environment, the placements of the target areas, and the presence of other agents. Such analysis, again, allows us to including the mechanistic relationship between the agent and its environment, thus captures more information that constitutes an agent's fitness.

This resonates with a co-constitutive interactionist perspective of biological analysis, which proposes that for a given organism, the genotype and the environment at each level of analysis *construct* the organism together rather than simply interact and affect each other (assuming the organism = the genotype) (Oyama, 1985). Indeed, one may imagine that some changes about the environment (e.g., turns into an oval rather than a square, or has an additional target area), the efficacy of some of the desired behaviors illustrated in our functional behavior analysis (Fig. 9-13) may quickly dissolve, and the selected genome might look very differently. Currently, however, it seems that despite only taken into consideration an incomprehensive subset of the possibly desired behaviors, the linear functional behavior level model already

serves as a better mechanistic level explanation of individual fitness than a linear gene level model that incorporates all genetic loci.

Nonetheless, as the evolution progresses, certain high fitness genotypes may become over-represented in the population, leading to a skewed dataset that violates the assumptions of parametric analyses. Additionally, multiple regression tests for this dataset tended to have a large condition number, which may suggest strong multicollinearity.

Selection For Motor Behaviors Drive Fitness In the Beginning of the Evolution

In the beginning of the evolution, fitness seems to be driven primarily by the selection for motor behaviors, specifically exploring the environment and stopping at the target area (Table 4, Table 5). The former is purely motor and does not require sensory perception, and the latter requires relatively simple perception-action coordination as the ground sensor input is either 0 or 1. Evolution of traits that depend on comm sensors, IR sensors and comm units are typically selected later in the evolution (Table 4; Table 6).

To understand how the traits evolved and which traits underlied the comm-related fitness advantage, the correspondence between the trends of the normalized values of each trait and the fitness was examined in a selective run where the fitness advantage emerged later. The results reveal some traits that went through considerable selection right before or right after the time increase in fitness emerged for each run; the results from one representative run (Run 0) are discussed as an example (Fig. 19). In Run 0, these traits included turning upon comm 0, comm 2 and comm 3 sensor activation, ground_comm, and turning in the right direction (right for right-handed agents, left for left-handed agents) upon activation of IR 4 and IR 5. The pattern of the

evolution of these traits in terms of both which traits were selected and when they were selected is highly heterogeneous across different runs (Fig. 21).

One possible explanation for this pattern is that a trait that went through selection right before the fitness increase served as a precursor for some other traits to become beneficial to fitness and start to undergo natural selection. For example, in Run 0, the trait where an agent emits an increased signal intensity when it has stopped at a target area may allow other agents to avoid it more efficiently through the frontal comm 0 sensor (Fig. 22). Alternatively, some of these other traits might also have already gone through some kind of selection, but because of the evolution of this new trait, begin to experience another type of selection (Fig. 19).

The simpler motor behaviors, on the other hand, are selected early on during evolution; the resulted values selected during this early period were also conserved in the genome throughout the rest of the evolution runs (Fig. 20). Even some communication sensor related traits that contributed to fitness are mostly just performing non-navigation motor functions and helping agents move faster by taking advantage of a baseline communication signal, similar to signal A and E in Marocco and Nolfi (2007). This pattern of selection was highly similar across all runs and in both conditions with or without comm sensor access. It is possible that the selection for these traits served as constraints for the initial subset of genome within with comm-sensor and IR-sensor navigations were able to evolve. This leads to a possible explanation about why in some runs, the fitness advantage emerges from the beginning, and in some runs it only emerges later (Fig. 18): for the runs where fitness advantage were there in the beginning and where the navigation traits did not go through selection, perhaps the traits existed in the genotypes that also had the desired motor behavior traits (Fig. 21). For the runs where fitness emerged only later, perhaps even if the these traits were present in the random initial generation,

they were eliminated during selection for motor behaviors and had to make their way back into the genome through mutation.

The Evolution of Signal Avoidance

Marocco and Nolfi (2007) reported 5 functional signals (Table 1; Fig. 4). Two of these promote exploratory behaviors (signal A, E), one promotes avoiding areas that already contain two agents (signal C), and two are involved in approaching an area that contains only one agent (signal B, D). However, at least based on initial qualitative trial analysis, the current project did not find these signals to have evolved in many groups, let alone universally across runs. Instead, the most common evolved strategy in the best performing groups across runs was simply always avoiding signals from other robots (Fig. 22). In most runs this is coupled with a higher signal intensity when an agent is in a target area, but this is not required for good performance in the task (Fig. 23). Additionally, many of the groups evolve to respond to a signal in context of the presence of other signals (Fig. 22; Fig. 23). While the signaling system identified in Marocco and Nolfi (2007) is certainly a good solution to the problem, it is possible that it was captured by chance and did not reflect the majority of the system. Nonetheless, it is not unreasonable to hypothesize that given more time, some of the evolved solutions found in the current project may evolve towards the direction of distinguishing one vs. two agents within target areas and act accordingly.

Most of the best performing groups did not evolve solutions where all the agents always end up in a target area. Rather, most groups would have two or three agents by the end of the trial, leaving one or two wandering outside (Fig. 23). Considering the system's harsh penalty for overcrowding a target area, this conservative solution makes sense because the cost of having

one less agent in the target area is lower than the cost of having an extra agent. On an individual level, this signal avoidance behavior seem detrimental to the fitness of the receivers that were kept outside the target area; however, on the group level, this ensured the fitness of the agents that have already found the target area. This result is in support of the mutual benefit stipulation in the biological theory of communication, that a signal would evolve if there is a net benefit to the sender and receiver in terms of genotypes (Greenfield, 2014). Because the signalers and the receivers share the same genotype, this behavior is passed down when it is inheritable (Greenfield, 2014).

In Run 0, the resulting functional behavior of the traits that seemed to underlie the fitness advantage is an alternative avoiding behavior (Fig. 22). An agent will avoid another agent located alone in a target area when its signal is perceived via the frontal comm 0, but not avoid it when the back comm 2 sensor receives the signal from an agent located in the other target area (Fig. 22). This often leads to fairly good performance in trials; logically it is a good strategy given that an agent perceives some agents in front of it and some agents at its back, comparing to when it only perceives an agent in front of it, knows more information about the environment and can be more confident that there might be a vacancy in the target area in front of it. In

This avoidance behavior is generalizable across different runs (Fig. 22, Fig. 23, Fig. 24). Run 4 showed a very similar behavior to Run 0, but did not evolve the behavior of changing signaling intensity when an agent in within the target area (Fig. 23). Surprisingly, in a small number of groups, IR sensor readings also affected the output of comm signal intensity and seem to help facilitate the avoidance behavior for agents in target areas (Table 6; Fig. 24). In these cases, instead of the perception of a signal emitted by another agent through its comm sensors, an agent exhibits avoidance behavior upon the perception of another agent's body through its IR

sensors. Reminiscent of the comm sensor/motor function in increasing motor speed for better exploratory behavior, this is yet another demonstration of the leaky boundaries between dedicated sensorimotor modalities when it comes to the evolution of functional behavior. Both the prevalence of this signal avoidance behavior across different runs and the phenomena of sensors contributing to behaviors outside of its “intended” modality draw attention to the multiple realizability issue and ask deeper questions about the relationship between selection of the genotypes and selection of the phenotypes.

Functional Behaviors Are Multiple Realizable Across Levels

The functionalist concept of multiple realizability is built on the insight that mental states or cognitive behaviors exist to serve specific functions, and each function may be achieved in a number of ways; that they are multiply realizable (Levin, 2018). This coincides with our knowledge from evolutionary biology that natural selection operates on the basis of functional behavior rather than directly on the genotype (Lande & Arnold, 1983).

The current project demonstrates this principle in a number of ways. First, the functional behavioral level is found to be a better level on which to predict fitness than the levels of the network or of the genes, and the network level was also better than that of the genes (Table 2; Table 3). Correlation tests showed that some of the traits defined on a functional behavior level lead to stronger correlation with fitness than the network level, and the network level leads to stronger correlation with fitness than the genes level (Fig. 14). It is worth reiterating that these upper level functional behavior and network traits were not defined in a top-down manner through observation of the actual behaviors of the agent. Rather, they were defined through a bottom-up manner based on the behavior of network structures encoded by the genotype or on

logical inferences about the functional consequences of a network behavior within the task environment. This way, each level serves as a mechanistic translation of the level below it, maintaining the connection with the level of the genes which adaptive evolution operates on (Lande & Arnold, 1983). With this in mind, it follows that the superior models on each of the upper levels serve as supporting evidence that the same network behaviors are multiply realizable by different genotypes, and the same functional behaviors are multiply realizable by different sets of network behaviors. Perhaps, when arriving at a better model by measuring network behavior that emerge from network structures encoded by different genotypes, or by measuring functional behavior that emerge from the consequences of network behaviors within a specific task environment, the variance we removed are the variance between the different genotypes that leads to the same network behavior, and the difference network behaviors that lead to the the same functional behavior.

As the task environment selects specific functional behaviors (the current project found it to be signal avoidance), even though each of the ten runs started with different random initial genotypes and each went through a different and independent evolution process, the resulting behavior converges to a great extent at least judging from preliminary qualitative analysis (Fig. 22; Fig. 23; Fig. 24). The signal avoidance behavior resides a level above the functional behavior level of analysis we have included in the current project, as we only take into account an agent's behavior in regard to input of individual clusters of sensory neurons, but during an actual trial agents typically receive a combination of many inputs at the same time. As a result, we were unable to quantitatively measure an agent's behavior in multi-agent environment. Nonetheless, there are intriguing insights from the observation that sensorimotor mechanisms of signal avoidance behavior is not limited to the dedicated communication channels (Fig. 24). Similarly,

the use of dedicated communication channels were not specific to communication signaling and may have contributed to speed during the early stage of evolution (Table 4; Table 5).

It seems that during evolution, functional behavior evolves without regard to what a sensorimotor channel was *designed for*, and only concerns *how it may be used* in a way to increase fitness. Quinn (2001) has raised this observation by reporting the evolution of coordination behaviors in a groups of robots without dedicated communication channels, and IR sensors only. In the current project, this is manifest in some runs where the IR modality also facilitated the behavior of avoiding other agents. These results again reinforced the idea that, at least when the traits are categorized in a manner that relate to their genetic heritability (as the current project has done), the functional behavior level is a much better level on which to understand fitness.

Nonetheless, it seems dubious whether some of the coordination behavior in Quinn (2001) or in the current project may qualify as a signal under the biological definition, which requires both the signaling modality and the receiving modality to undergo adaptive evolution and/or to account for energy expenditure beyond the necessary baseline for the previous function (Wilson, 1975; Greenfield, 2014). In Quinn (2001), the robotic agents did not evolve their morphology to make it more perceivable to other agents, it's unclear whether the agents evolved more acute perception through IR sensors, and energy expenditure was not considered. These problem also apply to the current project, which also does not take into account the changes in morphology or energy expenditure. It may be easier to make a case for the increase signaling intensity when an agent is inside a target area as a valid signal (Fig. 22), but less so for IR-based avoidance of other agent's morphology; the latter seem more so a reaction to a cue rather than a signal (Fig. 24). Since the agent also does not seem to distinguish an agent with a wall, it may just be a

convenient byproduct of IR avoidance of the wall. Nonetheless, because it has a chance to contribute to fitness and at least does not seem to harm it, this trait was preserved in Run 2 (Fig. 24). On the other hand, for the signal in Run 4 that remained a constant value regardless of the state of the sender, should it be viewed as a signal or almost a part of its “morphology” and therefore a cue? Furthermore, is this really fundamentally different than a signal that is slightly raised when an agent enters a target area to justify treating one as a communicative signal and the other as non-communicative cue, if the fitness outcome is so similar? To answer these questions, more research is needed. One hypothesis is many of these alternative methods for agent avoidance behavior are conserved because energy is not of concern. Future study may test whether their presence is attenuated when the behaviors of the agents have energy costs. Another hypothesis is that comparing to biological environments, the robotic model is much more stable and noise-free, which allowed these alternative methods to evolve. Future study may test whether adding noise in the environment would lead to a different evolution process.

The current study is by and large a preliminary investigation into the complex problem of multi-level analysis of the process of evolution of communication. While some of the main questions (e.g., the functional behaviors space of the system or multiple realizability across levels) were only investigated through qualitative analyses, each objective was addressed and led to some new hypotheses. At the end of the day, these findings demonstrate the fussy mapping problem between mechanisms and function arising from multiple realizability and calls for reflections and new hypotheses on our conceptual constructs and categorizations regarding the evolution of communication.

References

- Bradbury, J. W., & Vehrencamp, S. L. (2011). Principles of animal communication. Sunderland, MA: Sinauer.
- Doncieux, S., Bredeche, N., Mouret, J., & Eiben, A. E. (2015). Evolutionary Robotics: What, Why, and Where to. *Frontiers in Robotics and AI*, 2. doi:10.3389/frobt.2015.00004
- Floreano, D., Mitri, S., Magnenat, S., & Keller, L. (2007). Evolutionary conditions for the emergence of communication. *Curr. Biol.*, 17:514-519
- Greenfield, M. D. (2014). Evolution of Communication. In *The Princeton Guide to Evolution*. Princeton, NJ: Princeton University Press.
- Hunter, J. D. (2007). Matplotlib: A 2D Graphics Environment, Computing in Science & Engineering, 9, 90-95, DOI:10.1109/MCSE.2007.55
- Lande, R., & Arnold, S. J. (1983). The Measurement of Selection on Correlated Characters. *Evolution*, 37(6), 1210. doi:10.2307/2408842
- Levin, J. (2018, July 20). Functionalism. Retrieved from <https://plato.stanford.edu/entries/functionalism/#Bib>
- Lipson, H. (2007). Evolutionary Robotics: Emergence of Communication. *Current Biology*, 17(9). doi:10.1016/j.cub.2007.03.003
- Kuha, J. (2004). AIC and BIC. *Sociological Methods & Research*, 33(2), 188-229. doi:10.1177/0049124103262065

- Marocco, D., & Nolfi, S. (2007). Communication in Natural and Artificial Organisms: Experiments in Evolutionary Robotics. *Emergence of Communication and Language*, 189-205. doi:10.1007/978-1-84628-779-4_9
- Moorad, J. A., & Wade, M. J. (2013). Selection Gradients, the Opportunity for Selection, and the Coefficient of Determination. *The American Naturalist*, 181(3), 291-300. doi:10.1086/669158
- Nolfi, S. (2002). Evolving robots able to self-localize in the environment: The importance of viewing cognition as the result of processes occurring at different time-scales. *Connection Science*, 14(3), 231-244. doi:10.1080/09540090208559329
- Oyama, S. (1985). The ontogeny of information: Developmental systems and evolution. Cambridge: Cambridge University Press.
- Quinn, M. (2000). Evolving co-operative homogeneous multi-robot teams. *Proceedings. 2000 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2000) (Cat. No.00CH37113)*. doi:10.1109/iros.2000.895232
- Quinn, M. (2001a). A comparison of approaches to the evolution of homogeneous multi-robot teams. *Proceedings of the 2001 Congress on Evolutionary Computation (IEEE Cat. No.01TH8546)*. doi:10.1109/cec.2001.934381
- Quinn, M. (2001b). Evolving Communication without Dedicated Communication Channels. *Advances in Artificial Life Lecture Notes in Computer Science*, 357-366. doi:10.1007/3-540-44811-x_38

- Quinn, M., Smith, L., Mayley, G., & Husbands, P. (2003). 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), 2321-2343. doi:10.1098/rsta.2003.1258
- Rowe, C., & Skelhorn, J. (2004). Avian psychology and communication. *Proc. R. Soc. Lond. B*, 271, 1435-1442. doi:10.1098/rspb.2004.2753
- Seabold, Skipper, and Perktold, J. (2010). Statsmodels: Econometric and statistical modeling with python. *Proceedings of the 9th Python in Science Conference*.
- Shannon, C. E. (1948). A Mathematical Theory of Communication. *Bell System Technical Journal*, 27(3), 379-423. doi:10.1002/j.1538-7305.1948.tb01338.x
- Smith, W. J. (1969). Messages of Vertebrate Communication. *Science*, 165(3889), 145-150. doi:10.1126/science.165.3889.145
- Szathmáry, E., Szathmáry, Z., Ittész, P., Orbaán, G., Zachár, I., Huszár, F., . . . Számadó, S. (2007). In silico Evolutionary Developmental Neurobiology and the Origin of Natural Language. *Emergence of Communication and Language*, 151-187. doi:10.1007/978-1-84628-779-4_8
- Waskom, M., Botvinnik, O., Hobson, P., Cole, J. B., Halchenko, Y., Hoyer, S., . . . Allan, D. (2014). seaborn: v0.5.0 (November 2014) (Version v0.5.0). Zenodo. <http://doi.org/10.5281/zenodo.12710>

Weaver, W. (1999). Recent contributions to the mathematical theory of communication. In *The mathematical theory of communication*. Urbana: University of Illinois Press.

Wilson, E. O. (1975) *Sociobiology: The New Synthesis*. Harvard: Belknap Press, Harvard.

McKinney, W. (2010). Data Structures for Statistical Computing in Python, *Proceedings of the 9th Python in Science Conference*, 51-56.