STAGE DE MASTER 2 RECHERCHE EN INFORMATIQUE

Towards a Statistical Analysis of Emergent Communication in Evolved Virtual Agents

Auteur :	*Encadrant* :
Andrew Szabados	Stefano Nolfi

Organisme d'accueil :

Laboratory of Autonomous Robotics and Artificial Life
Institute of Cognitive Sciences and Technologies
National Research Council
Rome, Italy

Table of Contents

0 Abstract
1 Introduction
- Context of this work
- Evolutionary Robotics
  - Evolutionary Computation
  - Embodied cognition
  - Agent/Robot-based simulation
  - Minimal cognition
  - Minimal bias
- Studying Language with Evolutionary Robotics
  - Studying Language through Communication
  - A question to be answered

2 Previous Work
- Set-up/Methodology
- Results
- Analysis
  - Progressive complexification
  - Behavioral Analysis of an exemplar solution
  - Productive and reactive abilities
- Commentary

3 Goals/Methodology
- Goals/Motivation
- Methodology
  - Cataloging the solutions
- A Pilot Study
  - Preliminary Behavioral Analysis
  - Towards a statistical analysis
  - Entropy
  - Mutual Information between Signal and Motor Output
  - Mutual Information between Signal and System State
  - Input Signal
  - Summary of Measures
- Generalization

4 Results
- By Measure
  - Entropy
  - Mutual Information between Signal and Motor Output
  - Mutual Information between Signal and System State
- By Goal
  - Measuring onset of communication
  - Measuring production/meaningfulness
  - Measuring reactivity
  - Evidence for Producer Bias or Receiver Bias
  - Other
- Generalization

5 Conclusions and Perspective
- Conclusion
  - Measuring the onset of communication
  - Regarding producer bias vs receiver bias
  - From the preliminary generalization
  - Lack of objectivity
- Perspective
  - Statistical test of significance
  - Framework for Generalization

6 Remarks
7 Bibliography

0 Summary/Abstract

Summary/Abstract

The field of Evolutionary Robotics has given Cognitive Scientists a new way to study Cognition. Using simulated, embodied agents one can seek to understand the role of basic mechanisms in cognitive processes and how neural control structures give rise to these processes. One such process is language whose minimal cognitive correlate, communication, has been studied by de Greeff and Nolfi in an experiment in which they successfully evolved communicative abilities in virtual agents. While demonstrating that communication can be evolved, questions remain as to exactly how this occurs given that a communicative act requires both the production of a meaningful signal and the appropriate reaction to this signal. Since it is unlikely that these two abilities emerge at the same point, it has been proposed that either a relevant but purposeless production or receptive sensitivity arises first and is then either exploited or exapted to give rise to functional communication. With this work I wish to show that a statistical approach, unlike the original behavioral analysis, can provide certain answers or at least clues regarding the emergence of these abilities. Taking a specific evolved agent and developing information-theoretic measures I demonstrate the utility of these measures in providing evidence for the presence of a producer bias. I then try to generalize these results to other evolved agents. I conclude that while this approach appears promising there is still work to do to create a more generic analysis.

keywords: Evolutionary Robotics, Communication, Producer Bias, Receiver Bias, Information Theory, Entropy, Mutual Information

1 Introduction

In this report I will relate the project I have been working on while visiting the LARAL-ISTC-CNR in Rome. After introducing this work, including its context, the field of Evolutionary Robotics, and its use in the study of language, I will detail some previous work, specifically the research of which this present project is a continuation. I will then detail my goals and methodology, followed by a presentation and discussion of the results.

Context

I set out with the motivation to understand language. More specifically, as a cognitive scientist seeing language as a complex high-level cognitive process, I was hoping to be able to develop a model which could be used to demonstrate that categories form the basis of communication. First though, I would have to become more familiar with the existing tools and models in the field of Evolutionary Robotics. In doing so, I became familiar with the work of my colleagues at the CNR and even tried developing an idea for a research project based on their work. For the first portion of my stay, I was interested in building on the notion of communication as a sensorimotor action the control of which could be evolved for. After encountering limitations which I believe to be as much ideological as they are technical, I decided to change focus and work instead on an experimental paradigm more in line with the potential of Evolutionary Robotics, the so-called Target-switching experiment on the emergence of communication which I will present below.

Evolutionary Robotics

Evolutionary robotics is a technique for the design and study of neural control mechanisms of either robotic or simulated agents based on evolutionary computational methods. As Turing already envisioned, it is the "designing of brainlike networks through genetic search" having as its goal though the "understanding [of] cognition" (Harvey et al. 2005). It is an approach that makes use of several different lines of research, which are perhaps best introduced individually.

Evolutionary Computation

One of the most significant elements of Evolutionary Robotics is the evolutionary aspect. Here it is meant to implicate the use of Evolutionary Strategies and Genetic Algorithms to perform a search within the space of agent neural controllers. The search is thus parameterized in the form of a *genetic code* and the objective function takes the form of a fitness function measuring the agent's performance on a defined task. The basic idea is that of its Darwinian inspiration, that over numerous generations advanced, that is “robust adaptive” (Harvey et al. 2005), behavioral strategies will emerge through a process of gradual complexification.

Embodied cognition

The other element of Evolutionary Robotics, that is the robotic aspect, has its roots in the notion of embedded cognition, the idea that cognitive processes are not abstract computational processes which occur offline and in isolation from the environment, but instead depend on the constant interaction and feedback with the external world provided by being situated in an interactive, potentially dynamic environment. The concept of active perception, i.e. that even perceptual processes such as vision actually require constant interaction with the environment (Varela et al. 1992), exemplifies the justification of seeing cognition as an embodied process.

Agent/Robot-based simulation

While nominally associated with robotics, the reality is that most often the agents are entirely virtual simulations inside a virtual world that may or may not be based on real robots. The reasoning is that evolutionary computation requires faster than realtime simulations to be able to run numerous trials and generations in a reasonable amount of time, and in practice real robots are only used to verify the results obtained in silico.

Minimal cognition

Perhaps the most important idea associated with this approach is that of minimal cognition, that is that Evolutionary Robotics is a tool to explore cognition by seeking the phylogenetically more simplistic building blocks of cognition. Harvey et al. prefer to define cognition not as an advanced process akin to human cognition but rather as “the capability of an agent interacting with its environment so as to maintain some viability constraint”.

Minimal bias

While the constraint of studying these simplistic building blocks of cognition may seem limiting, this approach presents one major advantage; it permits the reduction of prejudice and bias to a minimum (Harvey et al. 2005). Specifically, this is a consequence of the only constraints on the behavior of the agent being the “viability constraints” that is their tasks and any selective pressures. Any evolved behaviors are not the direct result of programming and are thus less biased towards solutions envisioned by the programmer.

Studying Language with Evolutionary Robotics

Evolutionary robotics can be used as an instrument to study myriad cognitive phenomena, including high-level cognitive functions by means of their minimal-cognitive correlates. Thus, even a complex high-level cognitive function such as language can be studied when looking at its reduced, or phylogenetically antecedent version, *communication*.

Studying Language through Communication

While language is a complex phenomenon, its abstraction, its minimally-cognitive correlate, communication is simply “the execution of a behavior that alters the behavior of another individual (or individuals) that has evolved because it is beneficial to either one of both individuals” (Nolfi 2010).

Numerous Evolutionary Robotics experiments have thus been performed within this view on communication. Two major lines of research have been the studies on its emergence (e.g. Marocco et al 2003), or its characteristics, such as the development of referential communication (Beer 2008), or the development of categorization (Beer 2003, Nolfi 2005, Hanard 2005).

A question to be answered

In the line of research on emergence of communication, Nolfi has identified one major question that can hopefully be answered with these techniques:

> “Whether and how communication can emerge and evolve despite the need to concurrently develop two skills at the same time (an ability to produce signals encoding useful information and an ability to react to those signals appropriately)” (Nolfi 2010)

There are two elements to this question, whether and how. That is, the first question pertains to the possibility of the emergence of such a system and the second pertains to the details regarding the nature of this emergence. The first question has largely been answered in work by Nolfi et al. which I will detail in the section on previous work. One of the goals of my work is to address the second question within the following framework.

As mentioned in Nolfi's question, in order to have a functional communicative act two conditions must be satisfied: there must exist the production of a signal encoding useful information and there must exist an ability to react to those signals appropriately. The question is then in which order these abilities occur in the evolutionary timeline. Several scenarios have been proposed.

It can be imagined that the ability to produced a signal containing information pertaining to the current state of the agent is developed simply through the tendency of information to be propagated throughout a neural controller. This is akin to saying that agents have the tendency to report on their current state, either sensory or motor. This scenario is known as a producer bias. (Mirolli and Parisi 2008, 2010)

Alternatively, it could be imagined that first a sensitivity to external signals is developed which is subsequently exploited by the evolutionary process to increase performance, a scenario known as a receiver bias. (Mirolli and Parisi 2008, 2010)

These two scenarios should not be taken to be mutually exclusive or to form a closed set. One could imagine the perhaps less probable scenario in which a mutation engenders both the production and sensitivity at the same time, a case in which the interest would then be identifying the evolutionary and genetic dynamics that give rise to such a fortuitous event.

The development of computational models of the evolution of language, or at least communicative behavior, presents a unique opportunity to explore these questions empirically. While the results will only be generalizable to human language to the same, limited extent that the communication evolved is an approximation of human language,

To summarize, Evolutionary Robotics is a tool for studying models of the base elements of cognition. The base element of corresponding to the high-level cognitive process of language is that of communication, which has been studied using these techniques both in the context of its emergence and its composition. The project I have undertaken is in the line of research on its emergence, specifically in response to the question posed regarding the mechanisms by which a functional communicative act requiring both meaningful production and appropriate reaction can evolve.

2 Previous Work

Of the work on Emergence of Communication and thus the study of language with Evolutionary Robotics, one of the most significant is that of Joachim de Greff and Stefano Nolfi (detailed in de Greeff and Nolfi 2010). Their work will serve as the basis for this current project, so I will first present their setup and summarize their results.

Set-up/Methodology

As mentioned above, the work of Nolfi and de Greff seeks to address the possibility of whether communication can emerge in an Evolutionary Robotics set-up. Essentially, teams of simulated mobile robots are rewarded for accomplishing a task for which communication is advantageous.

Agents are composed of wheels, infrared, ground, and rudimentary vision sensors in addition to a fully duplex communication channel. Their neural controller is a multi-layer recurrent neural network having the values of the infrared, vision, ground and communication channel sensors, plus the agent's own communication output and ground sensor at the previous timestep as input connected both to the 4 neurons of the hidden layer and directly to the output layer. The output layer consists of two motor neurons controlling the wheels and an outgoing signal as seen in Figures 2.0.

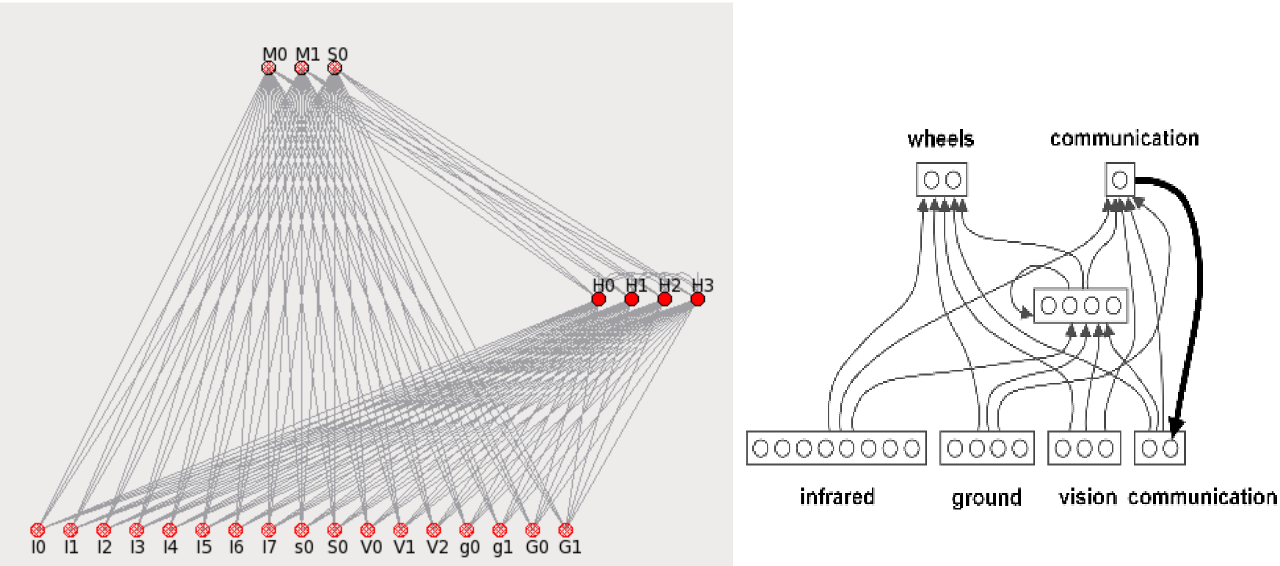

Figures 2.0a and 2.0b: Connectivity diagram of the neural controller (left) and explanatory schematic representation (right).

Regarding the neural equations (de Greeff and Nolfi 2010):

> The output of the motor neurons at time t is computed as the weighted sum of all inputs units and bias, filtered through a sigmoid functions:
>
> $$O_j(t) = \sigma\left(\sum_i w_{ij} I_i(t) + \beta_j\right), \qquad \sigma(z) = \frac{1}{1+e^{-z}},$$
>
> where I_i (t) corresponds to the activation of the i^{tn} neuron at time t, w_{ij} is the weight of the synaptic connection between the input neuron i^{th} neuron and the current neuron j, and B_j is a bias term.
>
> The output of the internal neurons at time t is computed by the following equation:
>
> $$O_i\,(t) = \tau \cdot Oi\,(t-1) + (1-\tau) \cdot Oi\,(t)$$
>
> Where Oi (t – 1) represents the output of the neuron at time t-1, and Oi (t) represents the weighted sum of all input units and bias filtered through a sigmoid function (see above), τ represents a time constant ranging between [0.0, 1.0].

The environment consists of a rectangular arena with two randomly placed circular areas, one gray, the other black (see Figure 2.1). Agents are rewarded for alternating their position between the two circular areas in a coordinated way. That is, they receive fitness points when they successfully switch areas, staying in opposite areas at the same time. They are not rewarded directly for any other abilities, neither communication nor such tactics as obstacle avoidance. Their task is known as the T(arget)SWITCH task.

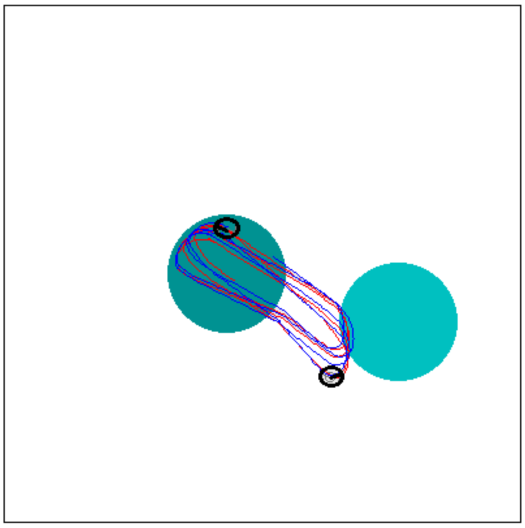

Figure 2.1: Arena during execution of successful strategy. The agents are represented by black circles and their trails are the red and blue lines. Here they are successfully alternating between the two areas, shown in dark green and light green.

The parameters of their neural networks, including synaptic connection and bias values comprise the genetic code which is evolved on using a 10020+1 Evolutionary Strategy for 2000 generations. During each generation the agents are allowed to live for 4000 timesteps and their fitness is calculated as the average of 20 trials.

Results

The results were that agents evolved to "exploit the possibility to communicate through explicit signals in most of the replications".

They observed two families of strategies. The first strategy which they called the symmetrical strategy was characterized by "a synchronized target-switching behavior in which the two robots, located in the two different target areas, simultaneously leave their current target area and move directly toward the other target area". The second strategy which they called the asymmetrical strategy in which there is "a switching behaviour organized in two phases in which first a robot exits from its target area and travels toward the other target area containing the second robot and then the latter robot exits from its target area and travels directly toward the target area previously occupied by the former robot."

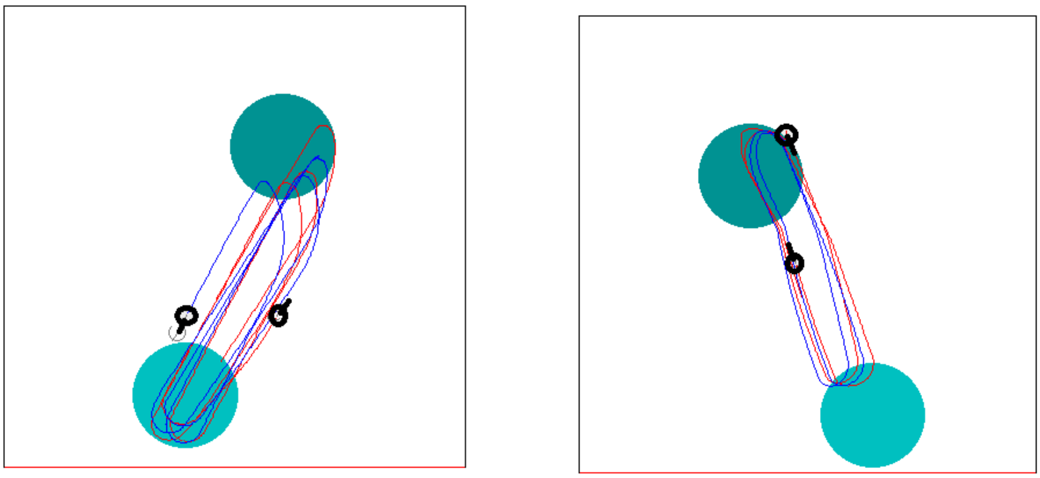

Figures 2.2 and 2.3: Snapshots of the two observed "families" of strategies evolved. On the left asymmetrical strategy in which both agents have exited the areas contemporaneously, and on the right the asymmetrical strategy in which one agent awaits the arrival of the other before leaving.

Analysis

Progressive complexification

Beyond the result of the specific families of strategies evolved, evidence was found for a process of progressive complexification by which simple behaviors and signals acquired new meaning or become more diversified or otherwise more complex as they “adapt to their task/environment”.

Behavioral Analysis of an exemplar solution

Within this perspective, they performed a detailed analysis looking at “how the behavioral and communication skills exhibited by robots of succeeding generations vary over the course of evolution”. They did so by cataloging “elementary motor and communicative behaviors” trying to find those which are functionally productive and how they are related to each other across generations. The result of this is a fairly coherent and highly detailed cross-generational behavioral analysis which brings light on the mechanisms the agents use, at a behavioral level, to perform their task, explaining the role of communication signals and their content as shown in Figure 2.4.

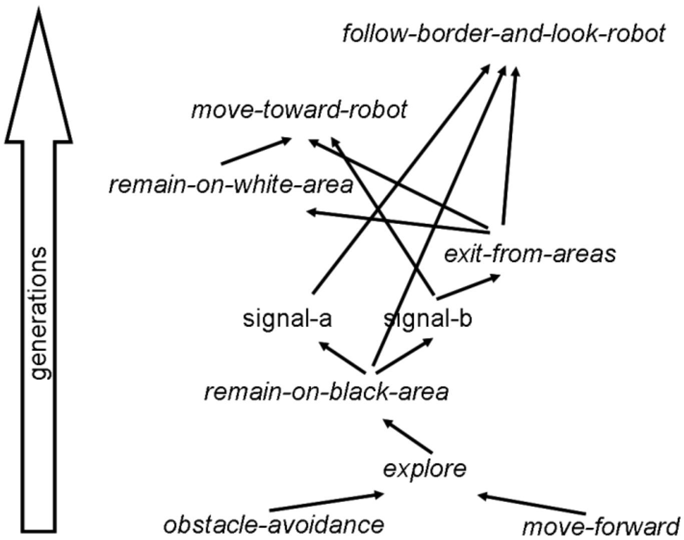

Figure 2.4: (From original article by de Greeff and Nolfi 2010). Displays the “relations between different behavioral and communication skills”. This analysis shows that more advanced communicative behaviors appearing in later generations are based on behaviors developed previously.

Productive and reactive abilities

Lastly, they discuss their results in the context of the question presented above, that is whether and how communication can emerge even though both productive and reactive abilities are necessary for communication to be adaptively influential. They provide examples of both cases in which a preexisting signal acquires communicative functionality through changes in the agents' reaction and in which a preexisting ability to react to signals in a specific way acquires functionality through variation of the signal produced but not the agents' reaction.

Commentary

In the end they were successful in responding to the first portion of the question, on whether or not communication can be evolved. With respect to the second part they have only presented evidence for the presence of both such mechanisms at work in behavioral and communicative functionality. There are two main problems with this. Firstly, the response is not definitive as to which bias is responsible, or more responsible for the evolved skills. While both mechanisms are at play over time, both cases they indicate deal with the evolution of one behavior into another. Secondly, even these results depend entirely on the act of cataloging the behaviors and only make sense in this context.

I wish to avoid entirely behavioral characterization of evolved skills, instead looking only towards statistical methods which can be much more easily generalized to multiple replications. Their analysis was almost entirely ad-hoc in that it was not much more than a detailed inspection of the agents at different points along the evolutionary history.

3 Goals/Methodology

Here I explain what I hoped to accomplish with this project and the methodology that I used.

Goals/Motivation

While the work of Nolfi and de Greeff has demonstrated the emergence of communication in this experimental set-up, thus responding to the first question, that is the possibility of the emergence of such communication, their behavioral analysis does not provide any clear answers regarding the question of either a producer or receiver bias. My first motivation is thus to attempt a statistical and information-theoretic, instead of behavioral, analysis of the evolved communication with the hope that this will be able to provide information relevant to the questions of interest. Specifically, the questions of interest are: can statistical methods provide information regarding the onset of functional communication, evidence for either a producer bias, or a receiver bias?

Methodology

In order to accomplish this I was given 80 seeds from the study of Nolfi and de Greeff of which approximately 10 evolved successful strategies, presumably all involving communications tactics, to solve the TSWITCH task. My first task was then to identify exactly which of these seeds resulted in interesting solutions. After cleaning up the data and obtaining fitness histories for each seed, it was apparent that only 8 achieved a best fitness consistently above 7, as can be seen in Figure 3.0.

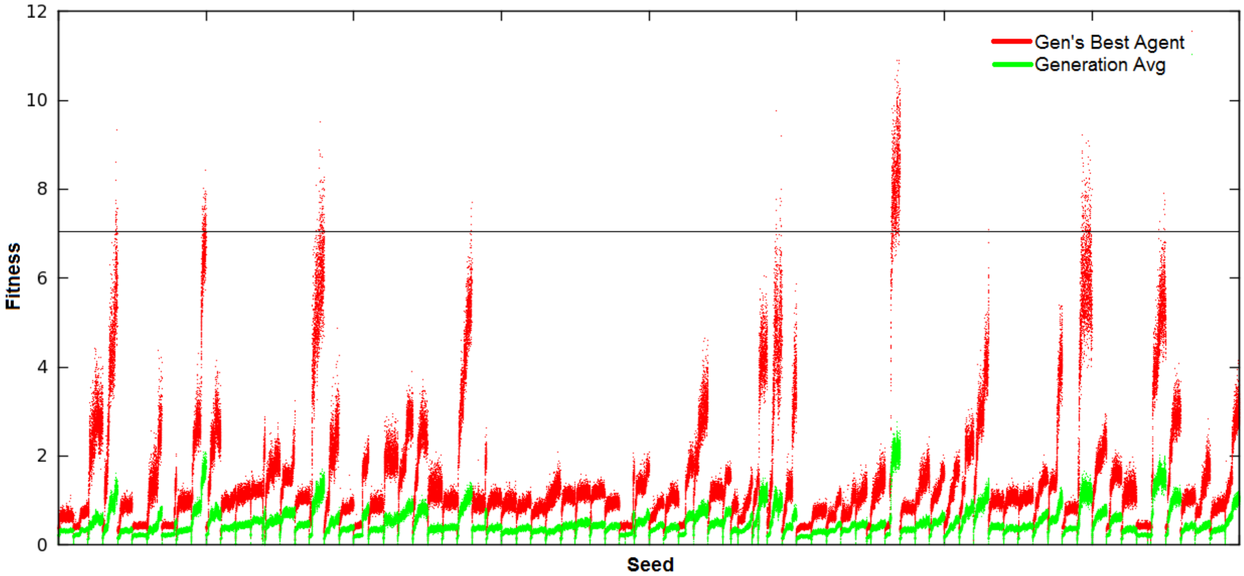

Figure 3.0: Fitness of each generation separated by seed and (arbitrary) threshold line shown at 7. I consider only those 8 seeds for which the best agent's fitness exceeded this threshold. It is also visible that Seed 71 attained a fitness considerably higher than the others.

Cataloging the solutions

The first task was then to cataloging these solutions. This was done looking both at the evolved strategies and the fitness/evolutionary profile. Of note was that the two types of strategies, as previously identified by Nolfi and de Greeff were present, symmetrical and asymmetrical. 5 of the 8 observed used the asymmetrical strategy, including the seed for which the fitness was notably higher than in any other run – Seed 71. Also of note was that in 6 of the 8 fitness histories, there were considerable, sudden jumps in fitness. Again, Seed 71 presented itself as the run with the most drastic and clear-cut jumps. It was thus selected for a pilot study.

A pilot study

Preliminary Behavioral Analysis

The first step in analyzing this agent was to understand how the evolution proceeded. From observation of the fitness history, it is clear that there are 3 distinct phases (marked in Figure 3.1). The agents were then tested in each phase of development and an attempt to understand and categorize their behavior was made.

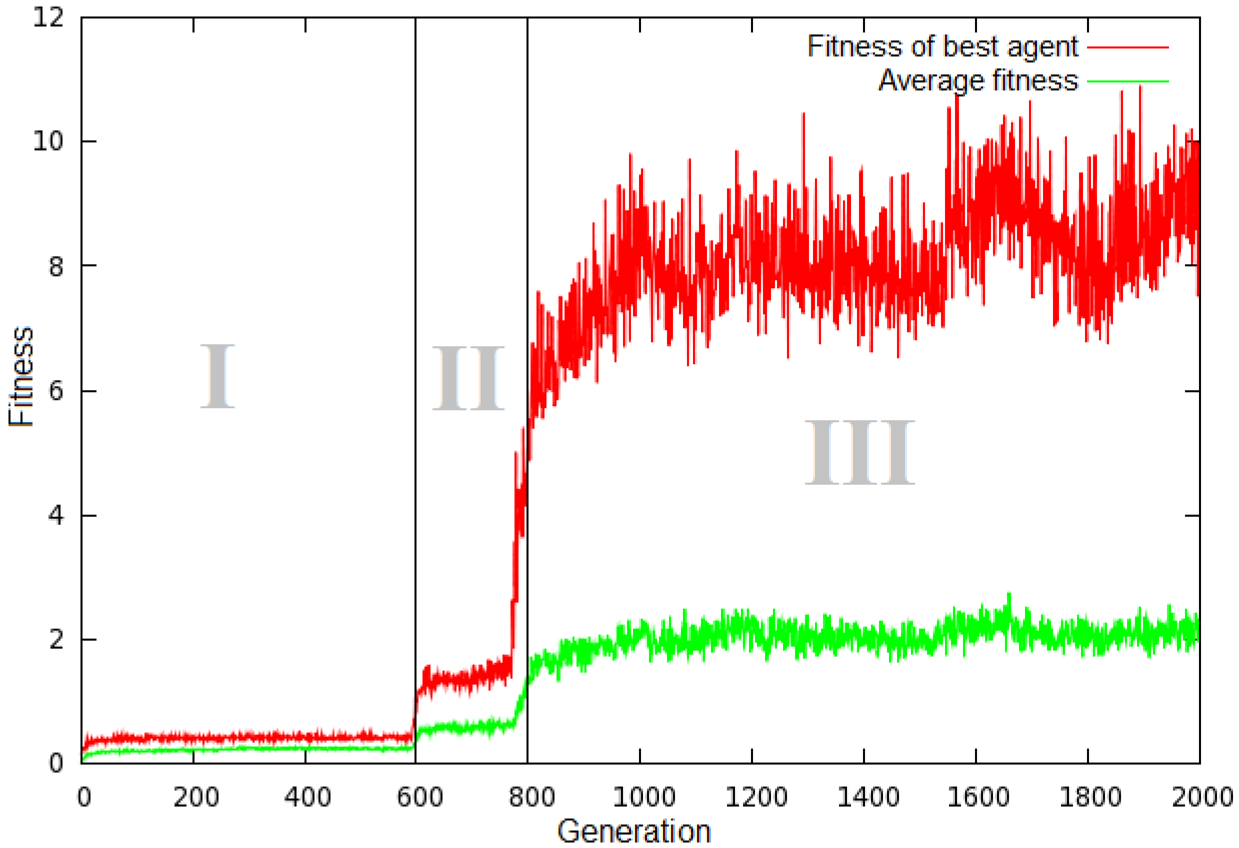

Figure 3.1: Plot of the fitness values across generations for Seed 71. The sudden jumps in fitness are used to separate the history into 3 phases, marked in gray.

What follows is a brief summary of the observations.

Phase I: The agents learn basic skills such as avoiding the walls and each other (they are penalized for collisions) and the ability remain in the target area (whichever one they encounter first) since they are given a small reward for finding the two areas (i.e. when each is in a different area). It does not appear that they use their communication channel, if not for some occasional noise relating to the behavior they have once in the black target area (cycling).

Phase II: By phase II the agents have acquired the ability to solve the task with relative success, and do so using rudimentary communication. It appears that whichever agent is in the black area emits a fairly stable signal. When the other arrives the first one leaves. So essentially the first waits for the second in the black area. I presume that they use the signal to negotiate the fact that only one can remain there (and thus emit the signal).

Phase III: By the third phase, the agents have developed an additional sensitivity to the presence of one another through their visual system, a fact which they exploit upon entering the black area to recover the direction from which the other agent has arrived. The signaling is similar in that a signal is still emitted while in the black area, but a new signal occurs at the moment when it enters. The old signal may now be indicating that the agent is on the edge of the area (where it waits for the other one to arrive) and the new signal indicates that it has not yet reached a stable position on the border. There may also be a similar signal upon exiting.

While intentionally preliminary and incomplete, it is clear that a more detailed picture of the communicative and behavioral skills evolved could be made as was done by de Greeff and Nolfi. This is where our analyses separate.

Towards a statistical analysis

A first step towards a statistical vision of the communicative abilities evolved in these agents, and specifically its evolutionary history, is a histogram of the communication channel throughout the evolutionary run. As can been seen in Figure 3.2, this view shows the variability of the signal, and the formation of stable signals. This picture confirms that during the first phase there is no variability in the signal, with its value remaining almost constantly 0. During Phase II, there already appears to be a stable signal, and during Phase III, there is the appearance of a third signal and corresponding shift in the first signal (which through observation was confirmed to indeed be behaviorally related in both phases).

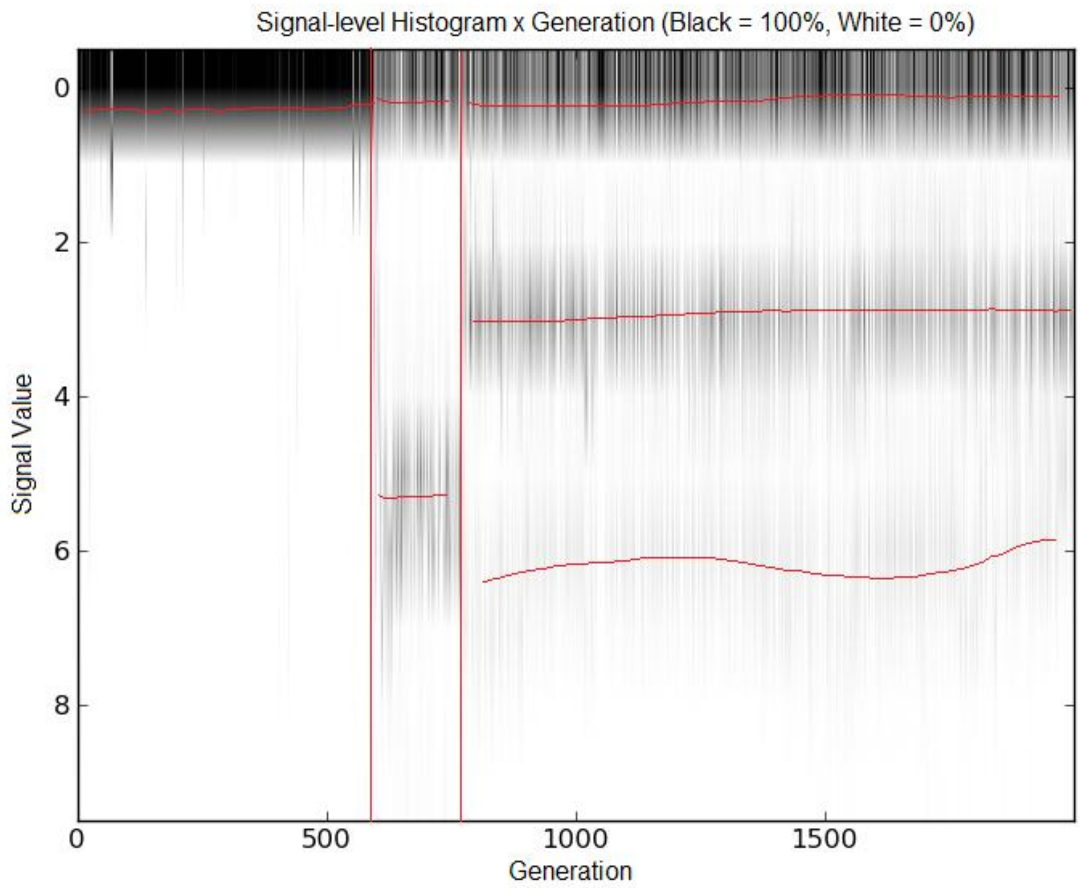

Figure 3.2: Signal histogram across generations with phases and apparent signals marked.

Entropy

The first statistical characterization of the communicative abilities of the evolved agents, specifically the variability of their signal, was thus the entropy of the produced signal. Formally entropy is defined as the average uncertainty in a random variable, intuitively, its information content, (Cover and Thomas 2006) and is given by the following formula:

$$H[X] = -\sum_{x \in \mathcal{X}} p(x) \cdot \log_2 \, p(x).$$

where X is the random variable and p(x) is its probability distribution function. In this work the random variable is the value of the signal or neuron activity level and its distribution function must be estimated numerically. To do so the simplest model is used, according to which a count is maintained of the times that a variable obtains a certain value.

$$p(x) = \frac{\sum_{x=s_t} 1}{T}$$

That is, the probability that random variable X takes on value x is the number of times the signal s was observed with that value divided by the total number of timesteps T. In practice the signal was first discretized into 10 bins.

While the entropy seems to provide a reliable picture of the onset of the signal, for reasons to be discussed below, it does not provide a complete picture. It can be characterized as providing evidence for the presence of a produced signal, but not whether that signal is meaningful to either the agent that produces it or receives it, or whether there is any functional communication. The next measures seek to do this.

Mutual Information between Signal and Motor Output

The first idea was to look at the relationship between the signal produced and the motor output, hoping that there would be a meaningful relationship between these two quantities. This was done calculating their mutual information. Mutual information is formally defined as the amount of information one random variable contains about another (Cover and Thomas 2006) and is given by the following formula:

$$MI[X;Y] = -\sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} p(xy) \cdot \log_2 \frac{p(x) \cdot p(y)}{p(xy)}.$$

Again, the picture described by these measures is incomplete. Specifically, the question is whether the motor output is a relevant value. The motor output, that is, the level of the motor neurons, is used to control the left and right wheels. Thus the motor output can be mapped to the movement of the agent, but not simply, as the direction of movement is related to the difference in output between the motor neurons. Thus instead of measuring the degree to which the movement of the of the agent is related to its output signal it is measuring this relation between the output signal and the motor profile of the agent, which depends on things such as the distance between the target areas, among other things. An attempt to counter for this fact was made by also calculating the mutual information against the difference between the value of the two motor neurons.

Mutual Information between Signal and System State

With the hope of understanding more about the controlling mechanism, the idea was developed to use, instead of the motor output, the sensory state, at least inasmuch as it pertains to the task devised, as the basis of a measure. One of the sensory inputs is the ground state, that is whether the agent is on either the black or gray zone, or if it is outside them. This sensor thus has three levels. In fact, the current system state, that is the location of the two agents in these terms, is enough to calculate the fitness function and it thus seems reasonable to think that this value may be related to the communicated information. Thus two versions of a system state variable were created, one of 3 states detailing an individual agent's current state, as it would be aware of through its sensory input, and one of the entire system state (3x3). Essentially, the system state was devised as a more abstract measure of the agent's behavior, not susceptible to the variability and complexities of the motor output.

Input Signal

Since the goal was not only to be to identify the moment in which a useful, pertinent, signal carrying relevant information began being produced, but also when a sensitivity to such a signal is developed, the above-described measures were repeated against the received signal.

Summary of Measures:

Entropy:

- produced signal – goal of showing presence or absence of signal by measuring amount of information contained in signal

Mutual Information:

- between the distribution of the signal received and the motor output with the goal of capturing degree to which received signal influences motor output.
- between the distribution of the signal produced and the motor output with the goal of capturing evidence of *meaningfulness* of the produced signal.
- between the signal produced and the state of the agent producing it
- between the signal produced or received and the state of both agents
- between the received signal and the state of the agent with the hope of measuring the degree to which the incoming signal influences the agent's behavior

Generalization

The last component is an attempt to generalize these methods and results to the entire dataset of evolved agents. While the measures were developed with the S71 dataset in mind, the size of the datasets and the time required for computation prevented the testing of the validity of these measures on the other seeds. Afterward however, these procedures have been run on the other datasets.

It should be noted that this present work is simply a pilot study the end goal of which is the verification of the validity of these measures. A longer term goal however is to be able to apply these tests generally so this attempt at generalization is truly preliminary and but a step towards the longer-term goal of being able to increase the scope of the answer obtained.

In this section I will present my results. This section is divided into first a general presentation of the results by each measure and then a specific presentation organized by the question of interest.

By measure

Entropy

Entropy, or the amount of information contained in the signal is shown here to correspond rather nicely to the phases as defined by the fitness measure (Figure 4.0). In the first phase there is minimal entropy, with occasional spikes that amount to noise as occasionally agents produced signals randomly. At the moment of the first spike in fitness, there is a spike in entropy. The entropy remains at a fairly constant level for the remainder of the evolutionary run. Of interest, and perhaps importance, is that there is a variation of entropy that begins approximately 50 generations before the first jump in fitness (Figure 4.1).

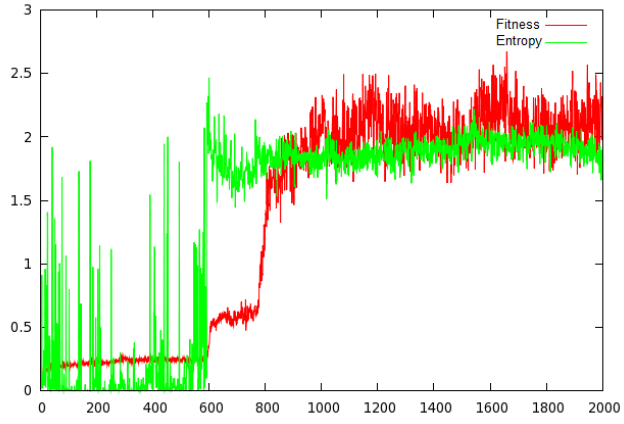

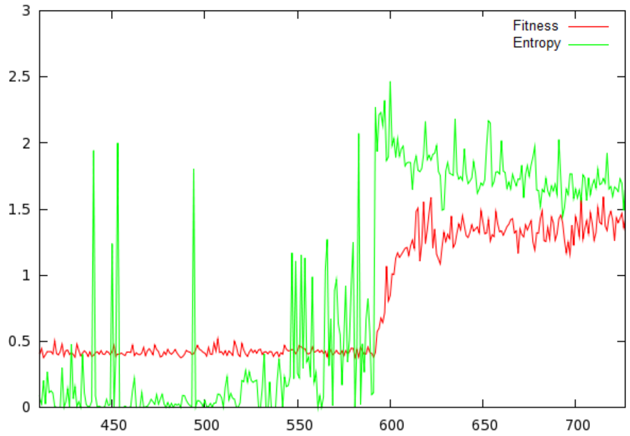

Figures 4.0 and 4.1: Entropy and fitness at each generation, for all 2000 generations, or area of interest.

Mutual Information between Signal and Motor Output

Unlike the entropy measure shown above, the various measures of mutual information between the signals and the motor outputs are less easily interpretable. Depending on the specific measure, there appear to be three patterns. The first is the mutual information between the signals and motor neuron 1 which rises with the fitness at its first jump and then remain high. The second is the mutual information between the signals and motor neuron 2 which rises only at the second fitness jump and then remain high. The third is the mutual information between the signals and the difference between the motor neurons which rises at the first jump and then decrease significantly (although not to 0) at the moment of the second jump. The difficulty in the interpretation of these data lies in the fact that the pattern depends not on whether or not the received or produced signal is used, but which motor neuron.

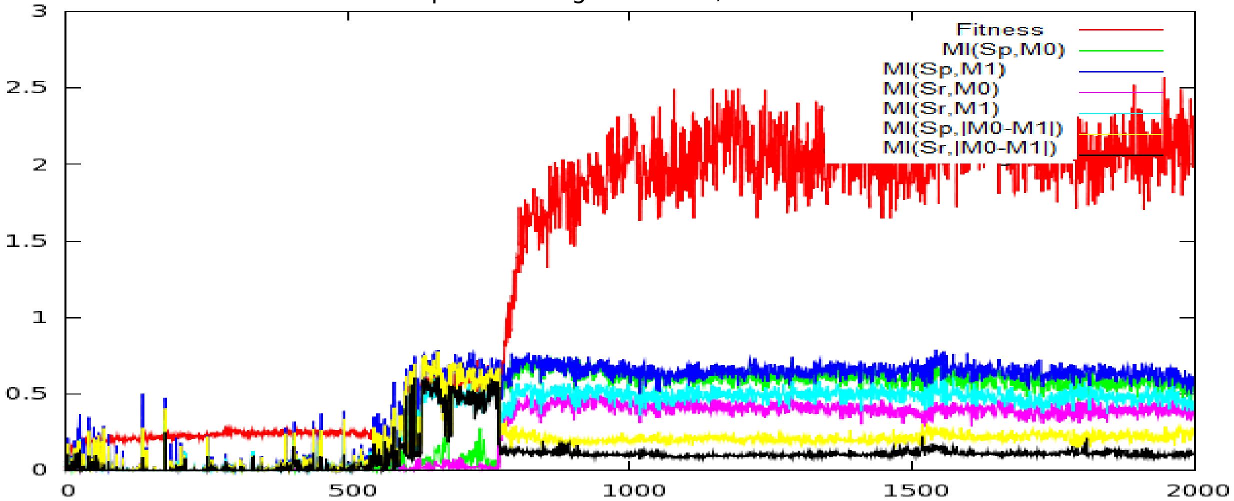

Figure 4.2: Mutual information (MI) between {Signal Produced (Sp),Signal Received (Sr)} and {Motor neurons (M0, M1), Difference between M0 and M1 (|M0-M1|)} for each generation

Mutual Information between Signal and System State

As was the intended goal of the System State-based measures, a much cleaner picture is presented here (Figure 4.3). Both values appear to rise at the same moment in time, in conjunction with the first jump in fitness. While there appears to be a momentary perturbation at the moment of the second jump in fitness, these values remain for the most part stable after the first jump.

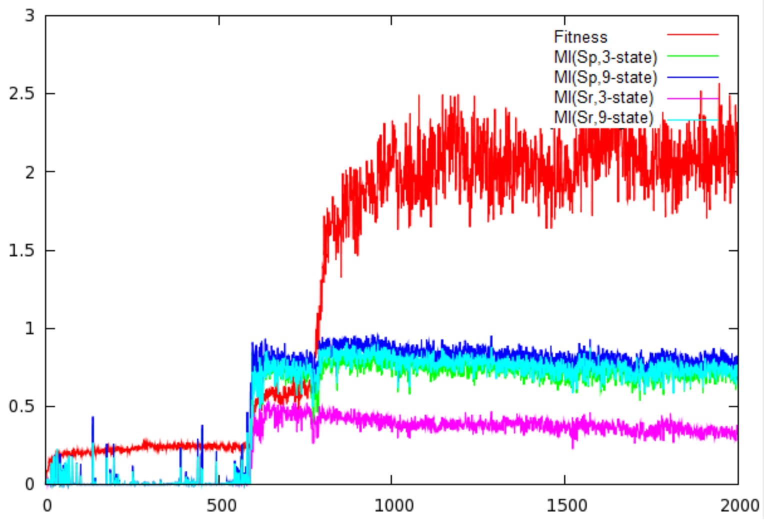

Figure 4.3: Mutual information (MI) between {Signal Produced (Sp),Signal Received (Sr)} and {3-state. 9-state} for each generation

By goal

Measuring onset of communication

From Figure 4.0 it is apparent that the increase in entropy corresponds to the first jump in fitness. This is consistent with the vision of things obtained through the behavioral analysis. Furthermore, the entropy value, which provides an objective measure of the presence of a signal shows that while the entropy increases drastically at the moment of the first jump in fitness, there is a period of approximately 50 generations previous to this jump in which there is activity on this channel. Thus, the entropy measure indicates that a signal is being produced before the development of any functional communicative act (which would be evidenced by an increase in performance) but does not provide any clues as to whether that signal is being modulated with any relevant/useful information, that is, the meaningfulness of the signal.

Measuring production/meaningfulness

Instead, in order to measure its meaningfulness, one must look towards the mutual information. The most pertinent measure of meaningfulness is the mutual information between the produced signal and the state of the agent (3-state). This measure, as seen in blue in Figure 4.4, shows that overall the relation between the signal and the current state of the agent becomes stronger exactly at the moment in which the fitness jumps and successful communication is developed. Additionally, it can be seen that even previous to this, in the 50-generation phase of non-functional communication from g550 to g600, while remaining small, the relation follows almost precisely that of the entropy, suggesting that for every occurrence of a signal, there is at least a minimal correlation between that signal and the entropy.

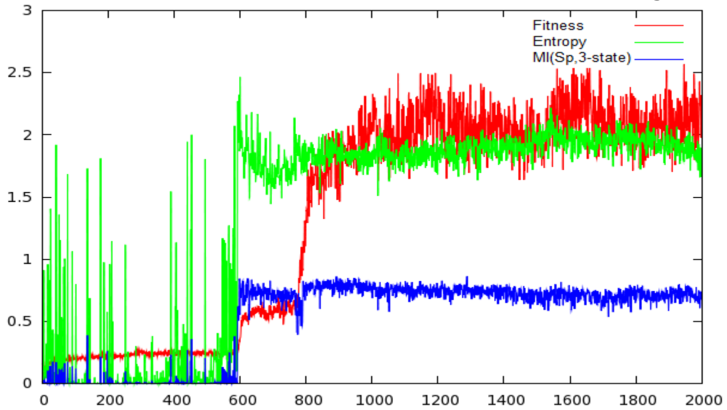

Figure 4.4: Fitness, Entropy, and Mutual Information (MI) between the Signal produced (Sp) and the state of the individual agent (3-state) for each generation.

Measuring reactivity

Now looking at the mutual information between the received signal and the state of the single agent, one should have a measure of the degree to which that agent is reactive to the signal. In fact, looking at this value (purple in Figure 4.5) one sees a similar trend as before in that the value jumps at the same moment as the fitness, entropy, and the mutual information between the signal produced and the state fo the agent, confirming that at that moment the agents develop the ability to react appropriately, or at least coherently, to the signal they receive. It is also the last piece necessary to show that that jump in fitness corresponds to the onset of functional communication in that there is a produced signal, it contains relevant information about the agent who sends it, and the receiving agent has the ability to react appropriately to it.

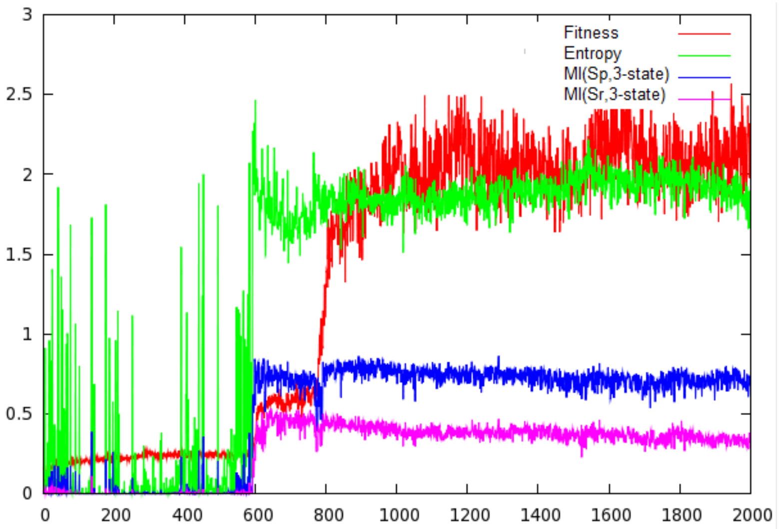

Figure 4.5: Fitness, Entropy, and Mutual Information (MI) between {Signal Produced (Sp), Signal Received (Sr)} and the state of the agent for each generation.

Evidence for Producer Bias or Receiver Bias

Lastly, the question is whether or not anything conclusive can be said about the events just previous to the onset of functional communication. Since there is a measure of reactivity and meaningfulness, which at least macroscopically follow the same trend and undergo a major change at the same moment, it must be seen whether or not they follow the same trend before that moment as well. Looking at the moment before the jump (Figure 4.6) one sees that the reactivity does indeed fail to appear before that moment.

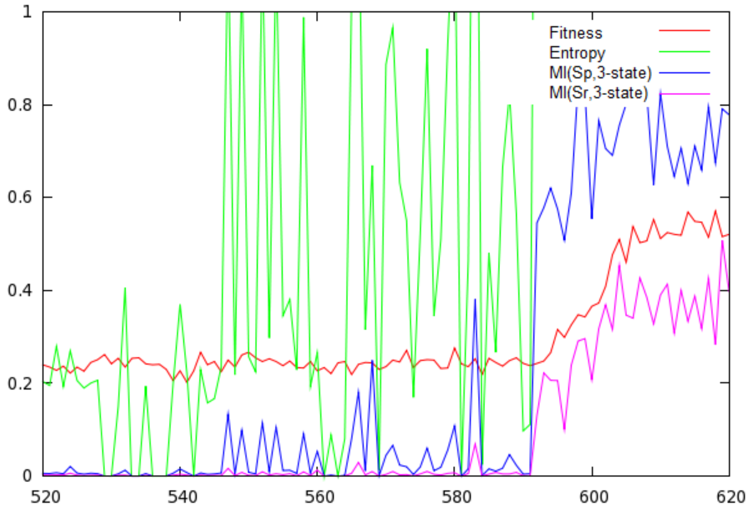

Figure 4.6: Close-up of generations 520-620 of Fitness, Entropy, and Mutual Information (MI) between {Signal Produced (Sp), Signal Received (Sr)} and the state of the agent.

Other

Otherwise not much can be concluded except that perhaps the mutual information measures based on the motor levels may be able to help in the understanding of the second fitness jump. That is, they were the only measures that changed significantly at that point, but without further analysis, such as an analysis of the motor activity itself, nothing more can be said.

Preliminary Generalization

There were several key phenomena observed in the case study that one would like to see in other instances too in order to generalize the findings:

First, there was an increase in entropy associated with an increase in fitness.

Next, the onset of functional communication was identified to be the moment when entropy and the mutual information between the signal (both produced and receive) and the state of the individual agent had all three increased above their base levels.

Lastly, a period of variations in entropy and “meaningfulness” precede “reactivity” that is the onset of communication.

The data from five additional replications was examined looking for evidence of these phenomena. In 4 out of the 5 cases, the analysis was rather clear, while in one case (Figures 4.11), further tests would need to be performed to interpret the results, specifically, to identify the onset of communication, or even if the format of communication is qualitatively similar.

In all 4 of the first 4 cases (Figures 4.7-4.10), increases or stabilization of entropy were associated with an increase in fitness.

In 3 of the 4 cases, there is a clear onset of communication as indicated by the presence of the three measures (Figures 4.7,4.8,4.10). In the fourth case, this onset is gradual but appears to still be explained by these measures (Figures 4.9).

In only 1 of the 4 cases was there a clear period of meaningfulness that preceded the onset of communication (Figures 4.7). In 2 other cases (Figures 4.9, 4.10), there was a possible period of meaningfulness, but either the gradual nature of the increase or the brevity of the period made it difficult to identified such a period with any certainty. In the last of the 4 cases (Figures 4.8), the onset of communication is preceded by a period of variability in entropy but apparently no meaning. Meaningfulness, reactivity and increased performance appear to occur contemporaneously.

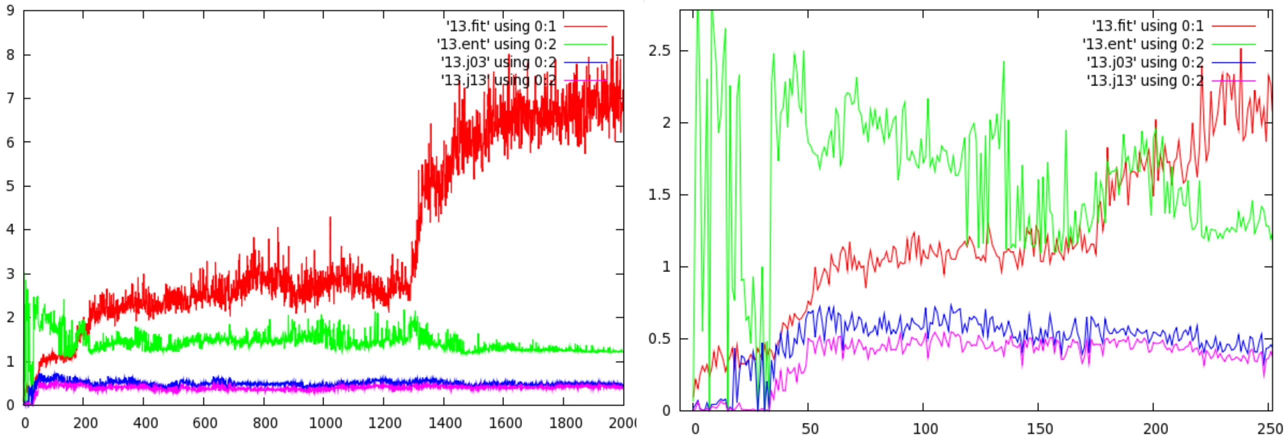

Figures 4.7a and 4.7b: Data from Seed 13, Fitness, Entropy and Mutual Information between {produced, received signal} and agent state. Across all generations on left, zoomed to area of onset on right.

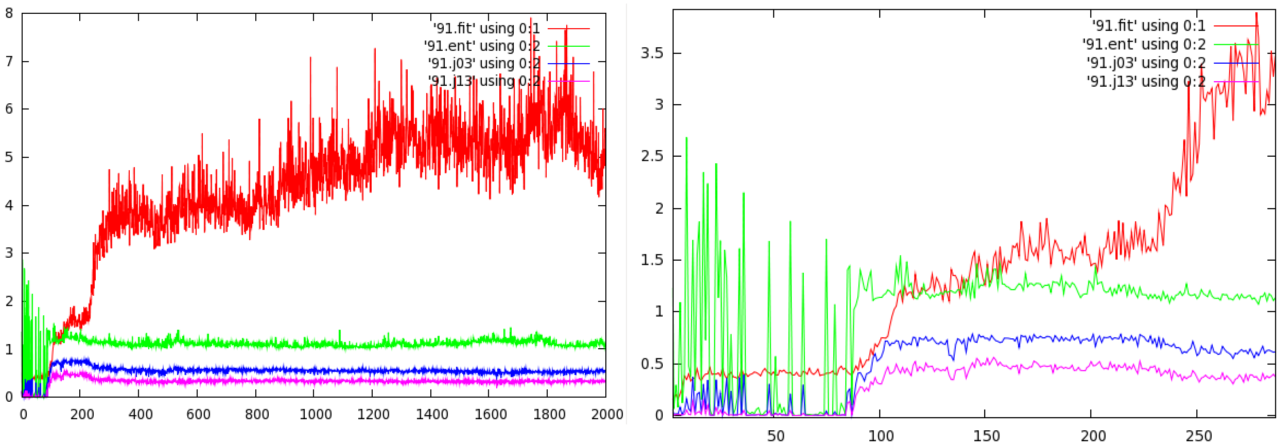

Figures 4.8a and 4.8b: Data from Seed 91,

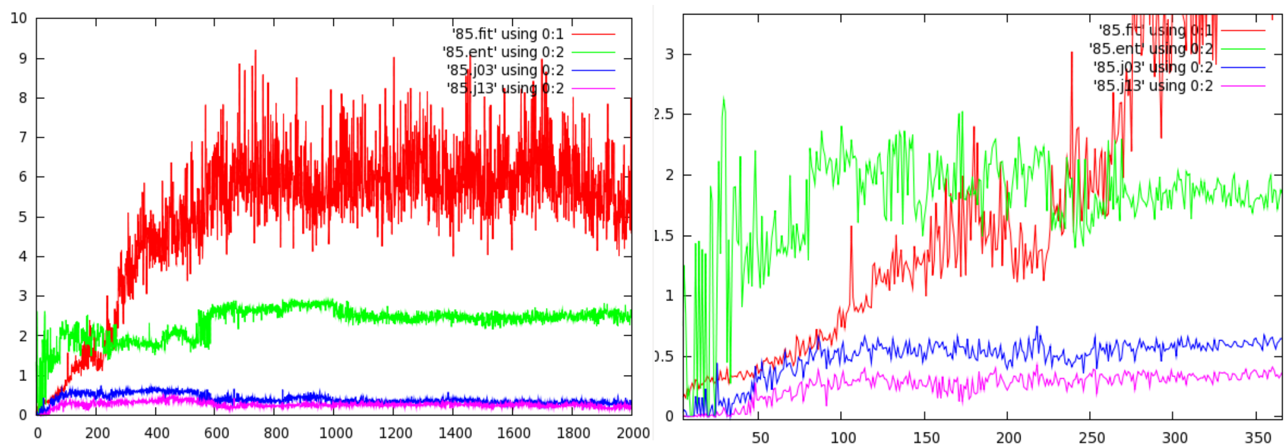

Figures 4.9a and 4.9b: Data from seed 85.

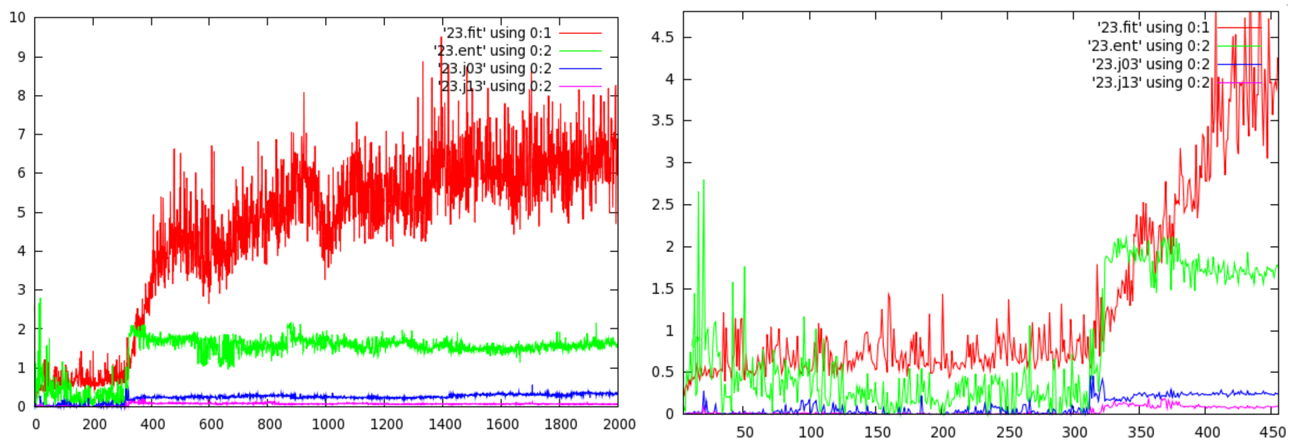

Figures 4.10a and 4.10b: Data from seed 23.

The last case (Figures 4.11) does not seem to fit with the others for several reasons. Among these are that the first clearly significant increase in fitness corresponds to a significant decrease in entropy and that the "meaningfulness" score remains quite low for the duration of the trial. Despite this, there is the possibility that communication is evolved in a manner consistent with that of the other replications but in the first few generations. In approximately the first 20 generations there appear to be all the components of the onset of communication – increase and stabilization in entropy, increase in meaningfulness and reactivity – except the corresponding drastic increase in fitness. While there is a gradual increase in fitness, the rest of the evolutionary run does not proceed as would be predicted. If these clues are taken to show the presence of communication early on, what happens next is necessarily a significant change in the nature of this communication, with a sudden change in the level (in fact, a decrease) of all three measures of the onset of communication. All that can really be said is that these measures do not provide a complete picture of what happens in this replication.

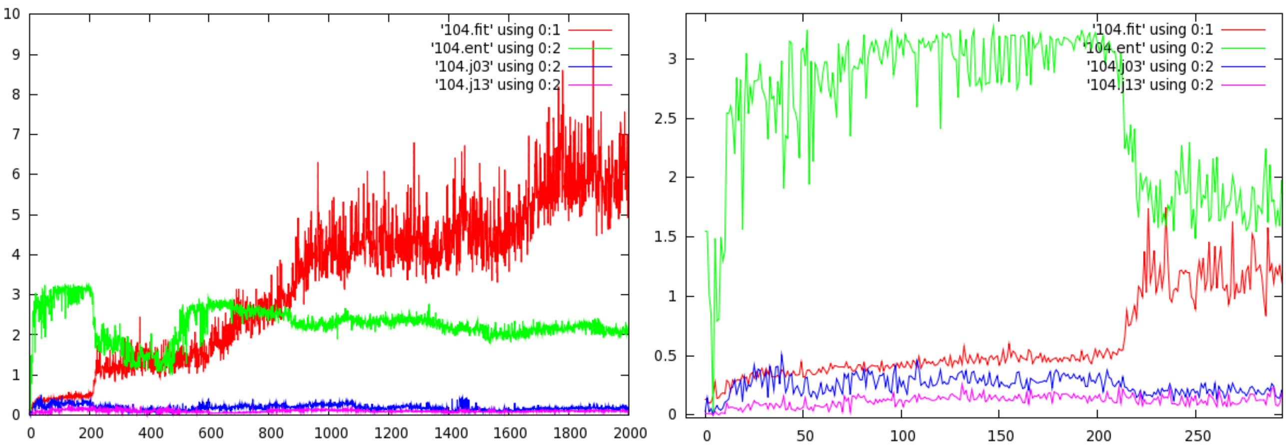

Figures 4.11a and 4.11b: Data from seed 104.

5 Conclusion and Perspective

In this section I will first summarize the conclusions that this work allows for and then I will present some directions for future work on this topic.

Conclusion

This work was undertaken with the goal of moving away from a dependence on the type of analysis done up until this point on the evolved agents in experiments on the emergence of communication, that is a behavior-based analysis which requires a hands-on analysis and observation of the performance of agents and providing a proof of concept for the use of statistical or information theoretical tools to perform similar, but complementary analyses. In this sense, this work was largely successful, in that it was clear that information theoretical tools are relevant and useful to questions regarding communication. Less successful was the demonstration of the power of these tools to replace outright the behavioral analysis. The picture presented in isolation was simply not clear enough and I thus propose that these tools should be used in conjunction with the types of tools used previously.

Measuring the onset of communication

In the case study, the measures developed - the entropy of the signal produced, the mutual information of the signal produced or received and the state of the agent - characterized nicely the onset of functional communication and provided solid evidence for an adaptive advantage of this ability.

Regarding producer bias versus receiver bias

This work was successful in providing evidence in the case study of the development of functional communication arising from a signal produced containing meaningful information but at first going unused. It was thus demonstrated that the replication in question exploited a so-called producer bias. What can be concluded beyond this statement is unclear. Beyond the issues of generalization to other seeds, there is a fundamental issue which needs clarification before any conclusions can be made regarding a generic role of this phenomenon in evolutionary processes. As it stands, the explanatory power of this model is limited to the configuration and parameters used in generating the data.

From the preliminary generalization

Applying of these methods to another 5 seeds, the relevance of these tools was reinforced in that they seem to provide evidence for the onset of functional communication and even its dynamics in the majority of the cases. Equally reinforced however is the need for a framework for generalizing the results.

Lack of objectivity

Another point that can be concluded from this work is that even the application of information-theoretic methods requires a certain degree of subjective interpretation in the format presented. Thus it can be said that one of the weaknesses of the methodology presented here was the lack of a statistical definition of significance, something which is not however unforeseeable. The data presented throughout has been the result of 100 trials at each timestep. While this was done to clean up the data, a confidence measure could also be derived. Likewise, one of the difficulties for generalization was the lack of an objective measure of, for example, the presence of a period of meaningful production without appropriate reaction.

Perspective

I will now address some of the issues raised above presenting ideas for future continuation work.

Statistical test of significance

Addressing the lack of objectivity for the issue of the producer bias, an objective definition for the absence of significant reaction and a statistical test for its presence should be created. An idea for such an objective definition is that, since the reactivity measure usually stabilizes at a value x approximately half of the stabilized value of the meaningfulness measure, it can be said there is a momentary producer bias whenever the value of the reactivity drops below this fraction of the meaningfulness value. A confidence score can then be derived applying this test to each trial (e.g. N=100). A similarly objective method would need to be developed extending the presence of momentary producer biases to a general one, keeping in mind that one also wants to be able to say with confidence that no such bias occurred.

Framework for Generalization

Once such a test and objective definition have been developed, one will be able to generalize the results to as many replications as desired. The results will still be limited in context however to this model (environment, task, neural control structure, with these parameters, etc.). A framework for generalizing these results to evolutionary processes in general and perhaps the origins of language should be a goal.

6 Remarks

Thanks

I wish to thank all those that made this work, and the experience as a whole, possible.

First and foremost I thank Stefano Nolfi for being willing to host me in his laboratory, for all his efforts to introduce me to the field of Evolutionary Robotics and its application to the study of language and communication, for his support and discussion of my various ideas throughout the course of my visit. He was also so kind as to provide me will all the data from the original experiment and to put me in contact will Joachim de Greeff.

I also wish to thank all those at LARAL and the ISTC, especially Vito Trianni, Onofrio Gigliotta, Giuseppe Morlino, Tomassino Ferrauto, Gianluca Massera, Elio Tuci, and others for their helpful insight, advice, and above all for having welcomed me so kindly in their lab.

Lastly, I thank the Master 2 de Recherche en Informatique and the Université de Paris Sud for allowing this project to be done in my laboratory of choice and allowing me to pursue my interests in full liberty as well as the Ecole Normale Supérieure, especially Carole Incana, and the Regione Ile de France for having contributed financially.

7 Bibliography

Works Cited

Beer, R. 2003. The Dynamics of Active Categorical Perception in an Evolved Model Agent. Adaptive Behavior.

Cover, T., Thomas, J. 2006. Elements of Information Theory. Wiley-Interscience.

de Greeff, J., Nolfi, S. 2010. Evolution of Implicit and Explicit Communication in Mobile Robots in Nolfi, S., Mirolli, M. (eds.), Evolution of Communication and Language in Embodied Agents, Springer-Verlag.

Harnad, S. 2005. To Cognize is to Categorize: Cognition is Categorization. In Handbook of Categorization in Cognitive Science. Elsevier.

Harvey, I., Di Paolo, E. Wood, R., Quinn, M., Tuci, E., 2005. Evolutionary Robotics: A New Scientific Tool for Studying Cognition. Artificial Life.

Marocco, D., Cangelosi, A., Nolfi, S. 2003. The Emergence of Communication in Evolutionary Robotics. Philosophical Transactions of the Royal Society.

Mirolli, M., Parisi, D. 2008. How producer biases can favor the evolution of communication: An analysis of evolutionary dynamics. Adaptive Behavior.

Mirolli, M., Parisi, D. 2010. Producer Biases and Kin Selection in the Evolution of Communication in Nolfi, S., Mirolli, M. (eds.), Evolution of Communication and Language in Embodied Agents, Springer-Verlag.

Nolfi, S., Floreano, D. 2000. Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines, MIT Press

Nolfi, S. 2005. Category Formation in Self-Organizing Embodied Agents. In Handbook of Categorization in Cognitive Science. Elsevier.

Nolfi, S. 2010, Artificial life modeling of language evolution. Invited Talk. Summer Institute in Cognitive Sciences on the Origins of Language. Université du Québec à Montréal

Varela, F., Thompson, E., Rosch, E. 1992. The Embodied Mind: Cognitive Science and Human Experience. The MIT Press.

Williams, P., Beer, R., Gasser, M. 2008. Evolving Referential Communication in Embodied Dynamical Agents. Artificial Life.