

Citation for published version:

Andres C. Burgos and Daniel Polani, 'An Informational Study of the Evolution of Codes and of Emerging Concepts in Populations of Agents', *Artificial Life*, Vol. 22 (2): 196-210, Spring 2016.

DOI:

https://doi.org/10.1162/ARTL_a_00199

Document Version:

This is the Accepted Manuscript version.

The version in the University of Hertfordshire Research Archive may differ from the final published version. **Users should always cite the published version of record.**

Copyright and Reuse:

This Manuscript version is distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution, and reproduction in any medium, provided the original work is properly cited.

Enquiries

If you believe this document infringes copyright, please contact the Research & Scholarly Communications Team at rsc@herts.ac.uk

An informational study of the evolution of codes and of emerging concepts in populations of agents

Andrés C. Burgos¹

Daniel Polani²

^{1,2}Adaptive Systems Research Group, University of Hertfordshire, Hatfield, UK

¹Email address: a.c.burgos@herts.ac.uk, Tel: +44 1707 28 4490

January 19, 2015

Abstract

We consider the problem of the evolution of a code within a structured population of agents. The agents try to maximise their information about their environment by acquiring information from the outputs of other agents in the population. A naive use of information-theoretic methods would assume that every agent knows how to “interpret” the information offered by other agents. However, this assumes that one “knows” which other agents one observes, and thus which code they use. In our model, however, we wish to preclude that: it is not clear which other agents an agent is observing, and the resulting usable information is therefore influenced by the universality of the code used and by which agents an agent is “listening” to. We further investigate whether an agent who does not directly perceive the environment can distinguish states by observing other agents’ outputs. For this purpose, we consider a population of different types of agents “talking” about different concepts, and try to extract new ones by considering their outputs only.

Keywords: information theory, code evolution, semantics, emerging concepts

1 Introduction

If we consider organisms capable of processing information, then we can argue that they must be able to internally assign meaning to the symbols they perceive in a code-based manner [10]. For instance, bacteria perceives chemical molecules in their environment and interprets them in order to better estimate environmental conditions and (stochastically) decide their phenotype [24, 1, 23, 27]. Plants detect airborne signals released by other plants, being able to interpret them as attacks of pathogens or herbivores [13, 29]. Therefore, a correspondence between environmental conditions and chemical molecules must be established. It is in this way that Barbieri characterises codes, and he proposes three fundamental characteristics for them: they connect two independent worlds; they add meaning to information; and they are community rules [2].

Codes connect two independent worlds by establishing a correspondence or mapping between them. These worlds are independent and thus there are no material constraints for establishing arbitrary mappings. The meaning of information comes exclusively from the mapping: symbols by themselves are meaningless. Finally, the third property requires that the correspondence between the two worlds constitutes an integrated system.

For instance, human languages establish a correspondence between words and objects [2]; in bacteria it is between chemical molecules and environmental and social conditions [35, 36]. Words (or chemical molecules) by themselves do not have any meaning, and each individual of a population can define, arbitrarily to some extent, their own set with its mapping. However, populations of individuals sharing the same code are ubiquitous in nature. How is it that codes come to be shared by many individuals when their constitution involve arbitrary choices for each individual? This question is what we are investigating in the present paper.

For this work, we assume a simple scenario where organisms live in a fluctuating environment. If they can perfectly predict the future environmental conditions, they can prepare themselves by adopting a proper phenotype, and, therefore, survive. However, when uncertainty about the environment remains, organisms will follow a bet-hedging strategy [31, 28], where they try to maximise their long-term growth rate by adopting the phenotype that matches the environment in proportions based on the information they have about it. For example, seeds of annual plants germinate stochastically in different periods in relation to the probability of rainfalls, and their chances of survival are maximised when they match this probability [5].

The relation between information and long-term growth rate can be expressed elegantly in information theoretic terms, where an increase in the environmental information of an organism

is translated into an increase in its long-term growth rate [30, 17, 18, 8, 26]. Such models achieve the maximisation of the long-term growth rate by maximising an organism’s information about the environment. If we assume this behaviour in organisms, then those obtaining additional environmental information (other than that from their sensors, which we assume it does not completely eliminate environmental uncertainty) from other individuals will have an advantage over those that do not, since they would be able to better predict the future conditions. However, for individuals to be able to communicate with each other, they must be able to translate symbols into environmental conditions, where the output of these symbols results from an individual’s code. We consider the code of an individual as a stochastic mapping from its sensors states to a set of outputs.

For this study, we consider outputs (or messages) of individuals (or agents) as conventional signs. In semiotics, the science of all processes in which signs are originated, stored, communicated, and being effective [10], two types of signs are traditionally recognised: *conventional signs* and *natural signs* [7]. In conventional signs there is no physical constraint on the possible mappings, they are established by conventions. Although in physical systems there can be limitations to the possible mappings that can be implemented, in this work we assume complete freedom of choice. On the other hand, in natural signs, there is always a physical link between the signifier and signified, such as smoke as a sign of fire, odours as signs of food, etc. [3].

In this work, we are not interested in the particular detailed mechanisms by which an agent implements its code, nor how the agent decodes the outputs of other agents. Instead, we focus on the theoretical limits on the amount of environmental information an agent can possibly acquire resulting from different scenarios of population structure and codes distribution. The natural framework to analyse such quantities is information theory [30]. However, it does not take semantic aspects into account, it only deals with frequencies of symbols instead of what they symbolise. Codes, on the other hand, add meaning to information, which makes the integration of sciences such as semiotics with information theory non-trivial [9, 4]. In the following section, we present an information-theoretic model which incorporates the necessity of conventions by dropping from the model the usual implicit assumption of knowing the identity of the communicating units.

2 Model

To introduce the model in a progressive manner, let us first consider three agents, θ_1 , θ_2 and θ_3 . Each of these agents depend on the same environmental conditions for survival, which are

modelled by a random variable μ . Agents acquire information about the environment through their sensors, which are modelled by random variables Y_{θ_1} , Y_{θ_2} and Y_{θ_3} , all three conditioned on μ , for agents θ_1 , θ_2 and θ_3 , respectively. We assume each agent acquires the same amount and aspects of environmental information from μ , *i.e.* $p(Y_{\theta_1}|\mu) = p(Y_{\theta_2}|\mu) = p(Y_{\theta_3}|\mu)$. Let us further assume that the information each agent acquires about the environment does not eliminate its uncertainty, *i.e.* $H(\mu|Y_{\theta_i}) > 0$ for $1 \leq i \leq 3$. The code of an agent is a stochastic mapping from its sensor states into a set of outputs, and is represented by the conditional probabilities $p(X_{\theta_1}|Y_{\theta_1})$, $p(X_{\theta_2}|Y_{\theta_2})$ and $p(X_{\theta_3}|Y_{\theta_3})$ for agents θ_1 , θ_2 and θ_3 , respectively (see Fig. 1).

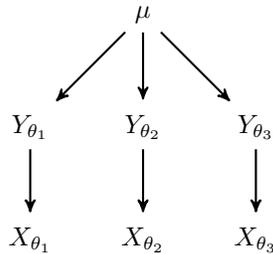


Figure 1: Bayesian network representing the relationship between the sensor and output variables of three agents.

Let us assume that agent θ_1 perceives only the outputs of agents θ_2 and θ_3 . One possible way of computing the information about the environment agent θ_1 has is to consider the mutual information between μ and the joint distribution of the sensor of θ_1 and the outputs of θ_2 and θ_3 : $I(\mu; Y_{\theta_1}, X_{\theta_2}, X_{\theta_3})$. However, by writing down this quantity, we are implicitly assuming that agent θ_1 “knows” which output corresponds to θ_2 and which output corresponds to θ_3 . Therefore, in this consideration, an agent can theoretically do the translations of the outputs according to some internal model of other agents and infer the mentioned amount of information about its environment.

2.1 Indistinguishable sources of messages

For this study, on the contrary, we consider an agent observing other agents’ messages, but under the assumption that the originator of a message cannot be identified. In this way, the total amount of information an agent can infer from the outputs of other agents will depend on to which extent it either can identify who the other agents are or can rely on them using a coding scheme that does not depend too much on their particular identity. For instance, if agents θ_2 and θ_3 both agree on the output for each of the environmental conditions, then agent θ_1 should be able to infer

more environmental information than if they disagree on the output for each of the environmental conditions, given that agent θ_1 does not know which of the agents it is observing.

To model this idea, let us assume a random variable Θ' denoting the selected agent. This agent depends on the same environmental conditions for survival as θ_1 , which are modelled, as above, by a random variable μ . Agents acquire information about the environment through their sensors, which are modelled by a random variable $Y_{\Theta'}$ conditioned on the index variable denoting the agent under consideration, Θ' , and μ . The amount of acquired sensory information of a specific agent θ' about μ is given by $I(\mu; Y_{\theta'})$. As above, the code of an agent is a stochastic mapping from its sensor states into a set of messages, and is represented by the conditional probability $p(X_{\theta'}|Y_{\theta'})$ for an agent θ' (see Fig. 2).

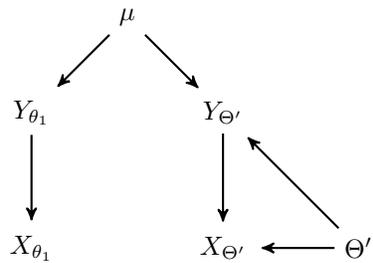


Figure 2: Bayesian network representing the relationships as described above (see text).

However, now we want to model the fact that we do not know which agent is observed. In the case with maximum uncertainty, Θ is uniformly distributed, and then this parametrisation of the codes considers the outputs of all agents in Θ' altogether, such that if we are not observing Θ' , we cannot identify whose agent's output we are observing. In Eq. 3 and Eq. 4 we show two examples of codes for agents θ_2 and θ_3 , while their sensor states are define by the Eq. 2 (Eq. 1 defines the sensors states of agent θ_1). We compute how much information about the environment there is when the outputs of both agents (θ_2 and θ_3) are considered together by agent θ_1 .

$$Pr(Y_{\theta_1}|\mu) := \begin{matrix} & y_1 & y_2 \\ \begin{matrix} \mu_1 \\ \mu_2 \end{matrix} & \begin{pmatrix} 1-\epsilon & \epsilon \\ \epsilon & 1-\epsilon \end{pmatrix} \end{matrix} \quad (1) \quad Pr(Y_{\Theta'}|\mu, \Theta') := \begin{matrix} & y_1 & y_2 \\ \begin{matrix} \theta_2, \mu_1 \\ \theta_2, \mu_2 \\ \theta_3, \mu_1 \\ \theta_3, \mu_2 \end{matrix} & \begin{pmatrix} 1-\epsilon & \epsilon \\ \epsilon & 1-\epsilon \\ 1-\epsilon & \epsilon \\ \epsilon & 1-\epsilon \end{pmatrix} \end{matrix} \quad (2)$$

If we assume $p(\theta_2) = p(\theta_3) = 1/2$, and $p(\mu_1) = p(\mu_2) = 1/2$ and $\epsilon = 0.01$, then if we consider the codes shown in Eq. 3, we have that $I(\mu; Y_{\theta_1}, X_{\Theta'}) = 0.97872$ bits, where Θ' consists of agents θ_2 and θ_3 . However, had θ_2 and θ_3 “opposite” codes as shown in Eq. 4, then $I(\mu; Y_{\theta_1}, X_{\Theta'}) = 0.9192$

$$Pr(X_{\Theta'}|Y_{\Theta'}, \Theta') := \begin{matrix} & x_1 & x_2 \\ \theta_2, y_1 & \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 0 \\ 0 & 1 \end{pmatrix} \\ \theta_2, y_2 & \\ \theta_3, y_1 & \\ \theta_3, y_2 & \end{matrix} \quad (3) \quad Pr(X_{\Theta'}|Y_{\Theta'}, \Theta') := \begin{matrix} & x_1 & x_2 \\ \theta_2, y_1 & \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 1 \\ 1 & 0 \end{pmatrix} \\ \theta_2, y_2 & \\ \theta_3, y_1 & \\ \theta_3, y_2 & \end{matrix} \quad (4)$$

bits, which is exactly $I(\mu; Y_{\theta_1})$, that is, $I(\mu; X_{\Theta'}|Y_{\theta_1}) = 0$ bits (agent θ_1 cannot acquire any side information from the outputs of agents θ_2 and θ_3). We should note here that the sensor states y_1 and y_2 of agents θ_2 and θ_3 in the conditional probability shown in Eq. 1 and 2 refer almost deterministically to the same environmental condition, and therefore the loss of side information is thus entirely due to the incompatible codes. The conditional probabilities of sensor states given the environmental conditions further defined throughout the paper are also assumed to be almost deterministically.

2.2 Environmental information of a population

The model shown in Fig. 2 considers the environmental information of agent θ_1 , ignoring its own output X_{θ_1} . Nevertheless, agents ignoring their outputs is contrary to our assumption over the incapability of agents to identify the sources of the outputs. On the other hand, we are assuming a specific type of communication, one which could be classified as persistent within the different classifications of stigmergy ([37, 33, 22], see [14] for a summary). To incorporate this option in the model shown in Fig. 2, we could consider the state space of Θ' as the set $\{\theta_1, \theta_2, \theta_3\}$. Then, to express not only the environmental information of agent θ_1 , but the average environmental information of the whole population, we can parametrise the agent by a random variable Θ (defined over the same state space, representing the same set of agents as Θ'), such that $p(Y_{\Theta}|\mu, \Theta) = p(Y_{\Theta'}|\mu, \Theta')$ (i.e., $Y_{\Theta'}$ is *i.i.d.* to Y_{Θ} , and vice versa).

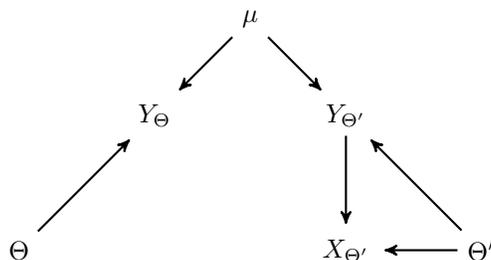


Figure 3: Bayesian network representing the sensor variables of a set of agents indexed by the random variable Θ , and the sensor and output variables of a copy of the set of agents indexed by Θ named Θ' .

In this way, the average environmental information of a population of the agents selected by Θ is given by $I(\mu; Y_\Theta, X_{\Theta'})$ (see Fig. 3). This measure can be considered as the objective function to maximise in our model. However, we would be making two important assumptions: first, this objective function assumes agents have access to the environmental conditions μ , which they indirectly do but only through their sensors; and second, every agent would perceive the output of every other agent, including itself. In this work, instead, we propose that agents follow a behaviour such that it maximises the similarity of their outputs (via their codes) with those of which the agent perceives. A consequence of this behaviour is that the average information about μ is also maximised. In addition, we will introduce a potentially flexible “population structure”, so that we can specify which agents interact with which.

2.3 Code similarity

First, we introduce a copy of the codes of the agents, such that when we instantiate the variables X_Θ and $X_{\Theta'}$, the probabilities are the same. The structure of the population is then given by $p(\Theta, \Theta') = p(\Theta)p(\Theta')$. However, the conditional independence of Θ and Θ' restricts significantly the diversity of the structures that can be represented. In such cases, the agents selected by Θ perceive the outputs of *all* the agents selected by Θ' and vice versa. In order to model a general interaction structure between agents, we consider $p(\Theta, \Theta')$ not independent, as shown in the Bayesian network in Fig. 4, where we introduce a helper variable Ξ . This allows different agents selected by Θ to perceive outputs from exclusive agents selected by Θ' .

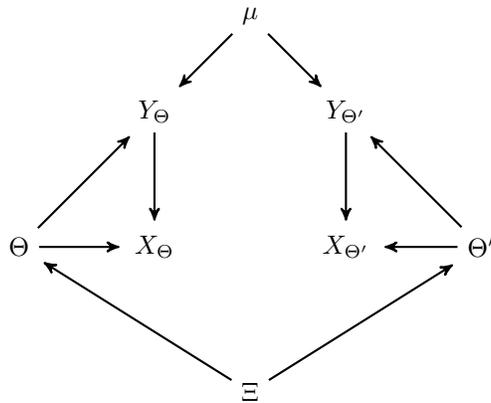


Figure 4: Bayesian network representing the relationship of the variables in the model of code evolution. $Y_{\Theta'}$ is an *i.i.d* copy of Y_Θ and $X_{\Theta'}$ is an *i.i.d* copy of X_Θ . Θ' covers the same set of agents as Θ , but its probability distribution is not necessary the same.

We define the objective function as $I(X_\Theta; X_{\Theta'})$, that is the average code similarity of a population of agents according to the population structure $p(\Theta, \Theta')$. For instance, if the interaction probability of two agents is zero, then the similarity of the codes of these two agents is irrelevant for the objective function. On the other hand, they interact with probability bigger than zero ($p(\theta, \theta') > 0$, for some agents θ and θ'), then how similar their codes are will influence $I(X_\Theta; X_{\Theta'})$.

If we consider our system as a process in time, then at each time-step two agents are chosen according to $p(\Theta, \Theta')$. Agent Θ reads the output of agent Θ' (generated via its code, which is *i.i.d* over time), and let us assume that it stores the pair $(Y_\Theta, X_{\Theta'})$, *i.e.* its current sensor state together with the perceived output. If this is repeated a large number of times, then the total amount of environmental information that can be inferred from the collected statistics by the population is bounded by $I(\mu; Y_\Theta, X_{\Theta'})$. This is the theoretical limit to which we refer in the introduction, and for this study we are not interested in how the inference is computed. However, we implicitly assume that agents decode the perceived outputs according to their codes.

2.4 Distance between two codes

In order to visualise the evolution of codes, we define the distance between the codes of two agents θ_i and θ_j as the square root of the *Jensen-Shannon divergence* [40, 19] between them. This measure has the property that $0 \leq JSD(\theta_i, \theta_j) \leq 1$ when \log_2 is used, and the square root yields a metric. Let us note that this distance requires the sensor states Y to be named identically (for the corresponding states of μ) among agents in order to be meaningful. As we stated above, this is (closely) the case in all our experiments. This requirement over the sensor states discards the possibility of using other measures such as mutual information.

$$\begin{aligned} dist(\theta_i, \theta_j) &= \sqrt{JSD(p(X_{\theta_i}|Y_{\theta_i})||p(X_{\theta_j}|Y_{\theta_j}))} & (5) \\ &= \sqrt{\frac{1}{2}D(p(X_{\theta_i}|Y_{\theta_i})||p(X_{\theta_k}|Y_{\theta_k})) + \frac{1}{2}D(p(X_{\theta_j}|Y_{\theta_j})||p(X_{\theta_k}|Y_{\theta_k}))} \end{aligned}$$

where $p(X_{\theta_k}|Y_{\theta_k}) = \frac{1}{2}(p(X_{\theta_i}|Y_{\theta_i}) + p(X_{\theta_j}|Y_{\theta_j}))$.

3 Methods

To illustrate the behaviour of our model, we consider four different scenarios, which are described in Sec. 4. The common parameters for the first two experiments are the following: the population consists of 25 agents; the amount and quality of the acquired sensory information is the same for every agent, that is $p(Y_{\Theta}|\mu) = p(Y_{\Theta'}|\mu)$. For the third scenario, the only difference is that we consider only 15 agents, since the dimensions to consider with a flexible structure grows quadratically with the number of agents.

The optimisation algorithm used in the following experiments is CMA-ES (Covariance Matrix Adaptation Evolution Strategy), which is a stochastic derivative-free method for non-linear optimisation problems [12]. We utilised the implementation provided by the Shark library v3.0.0 [15] with its default parameters, which implements the CMA-ES algorithm described in [11]. The evolutionary algorithm used for optimisation does not intend to represent the actual evolution of the codes. Instead, we are interested in the solutions of this optimisation process, which are representative of the possible outcomes of evolution.

To visualise the evolution of the codes of the agents, we use the method of multidimensional scaling provided by R version 2.14.1 (2011-12-22). This method takes as input the distance matrix between codes, and plots them in a two-dimensional space preserving the distances as well as possible. To visualise, not only the distances between the resulting codes, but also how they relate to the distances between initial codes, we provide a distance matrix of both initial and resulting codes. The initial codes are randomly set by the evolutionary algorithm.

4 Results

In this section, we analyse the outcome of the four different scenarios where code similarity is maximised. While the outcomes are particular for one simulation, they are illustrative of the richness that the model is able to capture, which is described for each scenario. The outcomes are typical solutions, and we cannot perform statistics over simulations since the many solutions are qualitatively different. However, the outcome of each scenario is presented together with a description of alternative outcomes, giving indicators of achievement of local/global optimum.

4.1 Well-mixed population

In the first scenario, each agent θ_i perceives the output of every other possible agent θ_j with the same probability, that is $p(\theta_i, \theta_j) = 1/25^2$ for every $i, j \in [1, 25]$. The maximum average code similarity is bounded by $I(Y_{\Theta}; Y_{\Theta'}) = 1.71908$ bits, which is achieved under two conditions: first, every code must be a one-to-one mapping; second, the code must be universal. This is indeed the outcome of the performed optimisation, as we show in Fig. 5: the optimised codes (blue points) converged into a universal code (the distance between any of them is zero). Each red (diamond) point correspond to an initial code.

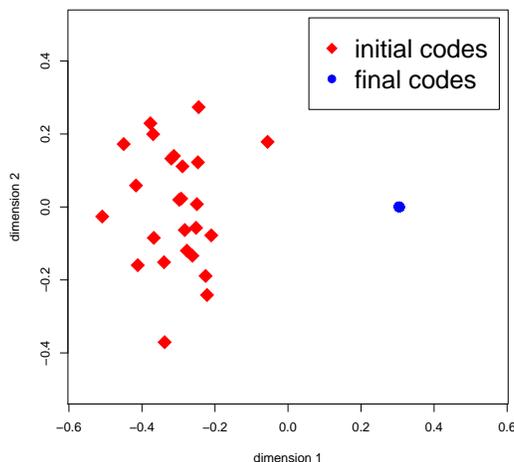


Figure 5: 2-dimensional plot of code distance: red points are codes at the beginning of the optimisation process; blue points are codes at the end of the optimisation process (where the distance between every pair of codes is zero).

The resulting code adopted by the population is a one-to-one mapping between sensor states and outputs, and any of the 24 possible one-to-one mappings is a global maximum (there are 4 sensor states and 4 possible outputs). However, it is still interesting to briefly analyse the possible paths towards a universal and optimal code. In Fig. 6, we show the distribution of the adopted codes by the agents of the population in an iteration of the optimisation process where the average code similarity is $I(X_{\Theta}; X_{\Theta'}) = 1.18276$ bits. Here, the most popular code is the suboptimal code shown in Fig. 6 (a). This results from the particular initialised codes, driving the agents temporarily towards a suboptimal code. However, once any of the many-to-one codes becomes (nearly) universally adopted, then any code's deviation improving the code similarity will eventually drive the convention towards optimality. The fact that it does not need simultaneous changes in the code increases the likeliness of improving the code similarity.

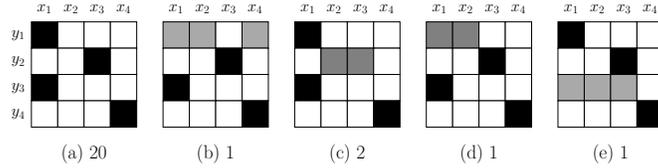


Figure 6: Representation of the codes $p(x|y)$ by a heat-map using inverse grayscale. For each evolved code, we output the number of agents adopting it. This code distribution was achieved with 25 agents in a well-mixed population.

4.2 Spatially-structured population

In another set-up, we assume the agents are structured in a 5×5 grid, where $p(\theta, \theta') = 1/105$ if θ and θ' are neighbours or when $\theta = \theta'$ (see Fig. 8 for a representation of the structure). After randomly initialising the codes, the performed optimisation plateaued on an average code similarity of $I(X_{\Theta}; X_{\Theta'}) = 1.13536$ bits. As in the former scenario, here the optimal solution is also a universal code with a one-to-one mapping. However, in this case, the result is not a universal code, as can be appreciated in Fig. 7. Spatially structured populations are sensitive to the initial codes and how codes are updated.

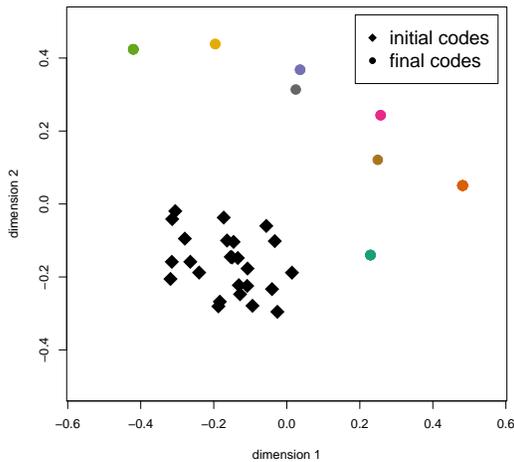


Figure 7: 2-dimensional plot of code distance: points in diamond shape represent codes at the beginning of the optimisation process; rounded points represent codes at the end of the optimisation process. The points are coloured in order to be able to relate this plot with the figure beside it.

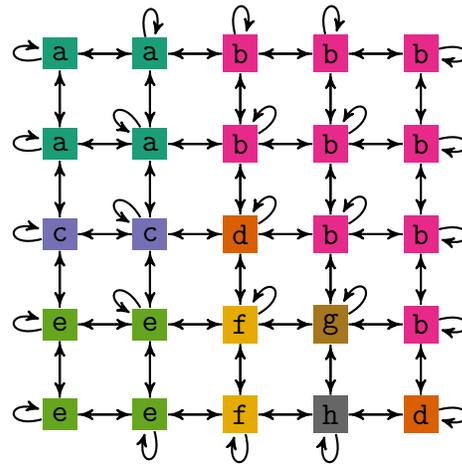


Figure 8: Representation of the spatial structure utilised for the experiment. Agents are assumed to be distributed in a grid: an edge from one agent to another means that one agent perceives the output of the other. Agents are labelled (see Fig. 9) and coloured according to their adopted code.

The resulting code distribution among the population is shown in Fig. 9, with 8 different codes

in the population. Where well-mixed populations evolved the use of common codes, agreement on codes only occurred among neighbours in spatially structured populations. As a consequence, many local conventions are established within neighbourhoods, and, once this situation is reached, the improvement of the total code similarity requires simultaneous changes to the agent’s codes. For instance, the code shown in Fig. 9 (e) could increase the average similarity of the population if $p(x_2|y_1) = 1$, as it is in the rest of the codes. However, for this to happen (in this particular case), at least two agents need to change their code simultaneously (otherwise the average similarity decreases), which makes the deviation from the resulting code distribution unlikely.

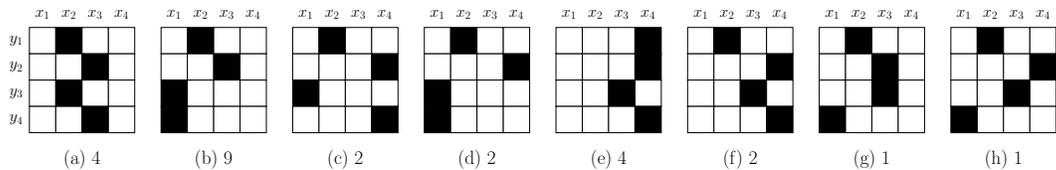


Figure 9: Representation of the codes $p(x|y)$ by a heat-map using inverse grey scale. For each evolved code, we output the number of agents adopting it. This code distribution was achieved with 25 agents in a grid structure.

4.3 Flexible population structure

For the third scenario, we let the structure co-evolve with the codes without any constraint (the probability distribution of the interaction between agents, $p(\Xi)$, is optimised together with the codes). In this case, the resulting average code similarity is nearly optimal, but the code is not necessarily universal. This is because, when the structure is not fixed, agents form roughly disconnected clusters of related codes. In this process, the interaction probability of agents with unrelated codes will vanish. However, once the clusters are formed, if it is not a single isolated agent (such that no other agent perceives its output), then codes of agents are universal within each cluster. This is exemplified by the code distribution and population structure we obtained (see Fig. 10). Here, we have two clusters with universal codes, one optimal (in red) and the other suboptimal (in yellow). Agents with dissimilar codes from every other agent they interact with will become isolated in the optimisation process, as the example shows for two agents (light and dark blue).

To summarise, the optimal code similarity equals $I(Y_{\Theta}; Y_{\Theta'})$, and is achieved, for instance, when all agents adopt the same one-to-one mapping. Nevertheless, the interaction probability allows agents to form disconnected clusters of related codes, where several one-to-one mappings could

result while still achieving optimality. Theoretically, we could have as many one-to-one mappings as the minimum between the amount of agents and the total one-to-one mapping combinations (24 in this case).

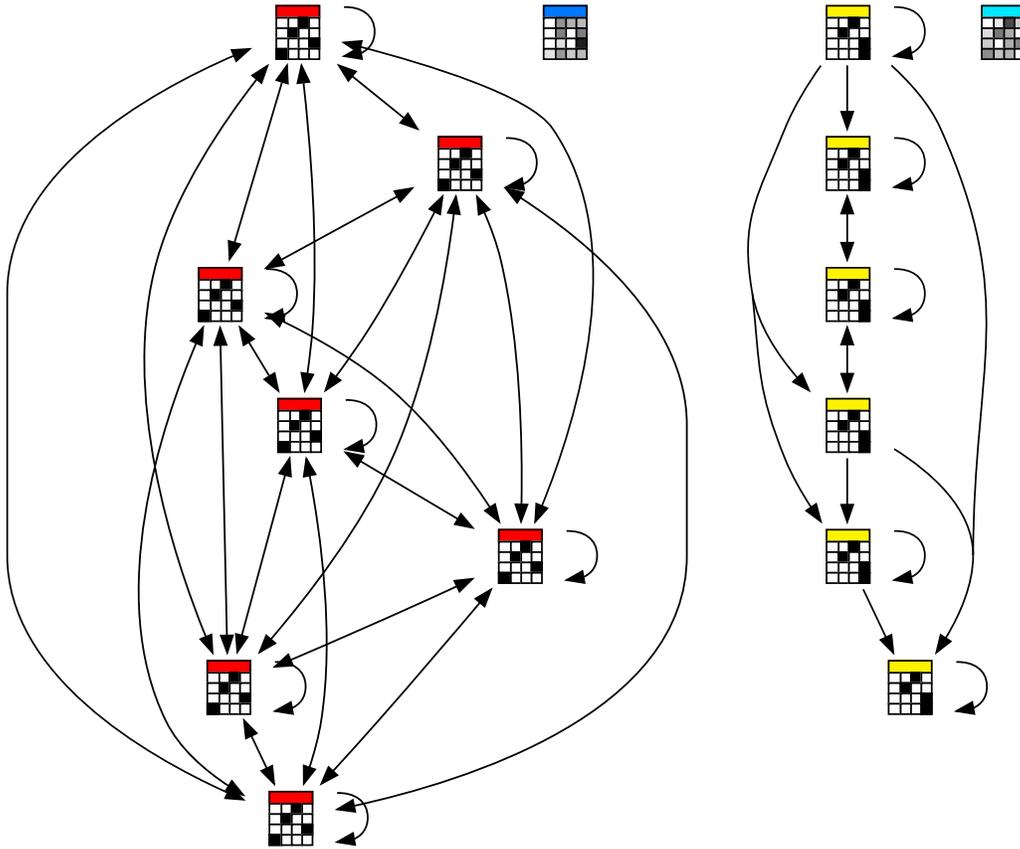


Figure 10: Each node in the graph corresponds to the code of an agent. There is a weighted edge between agent θ_i and θ_j if $p(\theta_i, \theta_j) > 0$ (which is the weight). We omit weights of edges in the graph since they all are roughly of similar value. The temperature colours on top of the nodes indicate the amount of environmental information they would contribute to any agent perceiving only that agents output.

4.4 Emerging concepts in a well-mixed heterogeneous population

So far, we have only considered populations of agents that acquired the same aspects of information from μ (*i.e.*, $p(Y_{\theta_i}|\mu) = p(Y_{\theta_j}|\mu)$ for any pair of agents θ_i, θ_j). The assumption was that the information that was relevant for the survival of the agents was the same among the agents of the population, and this was represented by μ . Now, we consider a more general scenario, where different types of agents acquire different aspects from the environmental conditions μ . We investigate whether it is possible for an agent that does not directly perceive the environment at

all (we call this type of agent “blind”) to predict conditions based solely on the outputs of other agents. We consider a well-mixed population, such that different types of agents are forced to talk to each other. Considerations with a flexible population structure are not interesting for our purposes, since in these cases, each type of agent forms a cluster disconnected from clusters of other types. This was confirmed by simulations which are not shown here.

Let us illustrate the idea with a relatively simple scenario: we consider five types of agents (we denote the i -th type ϕ_i), where each type can only distinguish whether the current state of the environment belongs to its coloured region or not. The environment consists of 9 states, and the probability of each state is uniformly distributed. We illustrate this environment by a 3×3 grid, as shown in Fig. 11, although the square does not denote the physical structure of the environment. Then, the outputs of each type of agent will be related to the regions they capture. For instance, for agents of type ϕ_2 with the same deterministic code, if $Pr(\mu \in \{1, 2, 4, 5\} | X_\theta = x)$ equals one (for all θ of type ϕ_2), then x will signify that this agent is currently in the region coloured in red in Fig. 11. We say that a population of agents has a joint *concept* of the environment if by considering its representation of the environmental information they capture, we can obtain information about the environment, *i.e.* we require that $I(\mu; X_\Theta) > 0$. For instance, the symbol x in the example above, assuming that it is only utilised by agents of the same type, can be understood as representing the concept “top-left” of the grid.

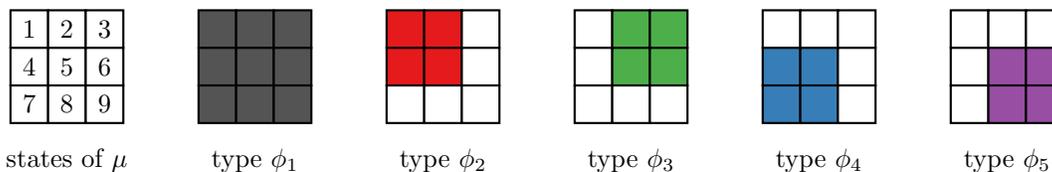


Figure 11: Representation of the conditional probabilities $p(Y_\theta | \mu)$ for an agent θ of each type. These are defined such that each type of agent can only distinguish between the coloured region and the white region. For instance, the sensor of type ϕ_2 is defined as $Pr(Y = y_1 | \mu) = 1$ if $\mu \in \{1, 2, 4, 5\}$, and zero otherwise, and $Pr(Y = y_2 | \mu) = 1$ if $\mu \notin \{1, 2, 4, 5\}$, and zero otherwise. For type ϕ_1 , $Pr(Y = y_1 | \mu) = 0.5$ and $Pr(Y = y_2 | \mu) = 0.5$ ($|Y| = 2$ for all types of agents).

The amount of environmental information that an agent θ of type ϕ_1 (a blind agent) captures is $I(\mu; Y_\theta) = 0$ bits, while all agents θ of the other types capture $I(\mu; Y_\theta) = 0.991076$ bits (note that the total entropy in μ to be resolved is $H(\mu) = 3.16993$ bits). Throughout this study, we considered that agents predict the environment by considering their perceptions together with the outputs of other agents. The blind agent, instead, since it is not able to capture any direct cue

from μ , we consider capable of perceiving the outputs of both of the agents selected by Θ and Θ' . With this relaxed consideration, we say a blind agent *has a concept* of the environment if $I(\mu; X_\Theta, X_{\Theta'}) > 0$, *i.e.* we consider the maximum amount of information an agent can possibly infer from the joint outputs X_Θ and $X_{\Theta'}$.

Let us recall that the structure of the population is well-mixed, and thus the distribution of outputs of all agents is considered, including the blind ones, which are not able to express (via their outputs) any particular concept by themselves (for a blind agent θ , $I(\mu; X_\theta) \leq I(\mu; Y_\theta) = 0$, *i.e.* $I(\mu; X_\theta)$ vanishes). Therefore, whether a blind agent has some concept of the environment will depend, first, on the universality of the codes of each type of agent (agents representing the same information with different symbols may create ambiguities). Second, on the cardinality of the alphabet of X (*i.e.* $|X|$) utilised by the population. A small alphabet will force agents to represent different concepts of the environment with the same symbols, while a large alphabet is likely to result in exclusive representations of concepts for each type of agent.

Taking this into account, we ask, is it possible for a blind agent to identify concepts of the environment? If so, how are these concepts related to the concepts of the individual agents (other than the blind ones)? Is the size of the available alphabet related to the quality of the concepts?

To study these questions, we performed different experiments varying the size of the alphabet $|X|$, where the rest of the parameters remained the same. In these experiments, we optimised the similarity of codes for a population composed of 20 agents, with 4 agents of each of the five types. In Table 1 we show that the cardinality of the alphabet of X affects the limit of the amount of information a blind agent can possibly infer about the environment.

Now, if we measure the uncertainty of the environment for a blind agent for each combination of outputs X_Θ and $X_{\Theta'}$, we find that for some of them, it is zero. For instance, with $|X| = 7$, we found that when $Pr(\mu = 5 | X_\Theta = 1, X_{\Theta'} = 2) = 1.0$ (see Fig. 12, where only combinations with $X_\Theta \leq X_{\Theta'}$ are shown). These distributions are also valid when swapping the values of X_Θ and $X_{\Theta'}$, since in the well-mixed population the structure is symmetric. Looking at the example of the conditional probability in Fig. 12, we can find many other concepts, although none of them—apart from the one already discussed—can uniquely identify a state of the environment. For instance, we have that $Pr(\mu | X_\Theta = 3, X_{\Theta'} = 6) = 0.33$ when $\mu \in \{3, 5, 7\}$, which is a concept for being on a particular diagonal of the environment.

In Fig. 13 we show the resulting codes (which are universal for each type, including the blind one) for this particular experiment. Here, the types ϕ_2 (red) and ϕ_5 (purple) utilise the same

$ X $	$I(\mu; X_\Theta, X_{\Theta'})$
2	0.34621
3	0.56555
4	0.71620
5	0.95467
6	1.08139
7	1.18362
8	1.30919
9	1.30919

Table 1: Results of experiments where the size of the alphabet of a population varies. The maximum amount of environmental information that a blind agent can infer is achieved with $|X| = 8$ and remains equal for bigger alphabets. As the size of the alphabet decreases, this information also decreases.

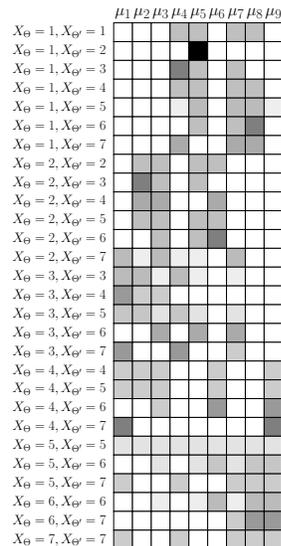


Figure 12: Conditional probability $p(\mu|X_\Theta, X_{\Theta'})$ in inverse grey-scale. Each row represents a combination of values of X_Θ and $X_{\Theta'}$, and each column represents a state of μ .

symbols to represent different environmental conditions. By using a small size of the alphabet for X , we force ambiguities in the population, but these will be chosen (by evolution) such that they are minimal. In this way, we maximise the amount of information we can infer from the outputs (although this can be a local optimum). For instance, the outputs of the blind agents (type ϕ_1) for all the experiments never overlapped that of other types (unless we use $|X| = 2$, where there is no choice). In other words, blind agents always choose one symbol so that they minimise the amount of utilised symbols from the whole population.

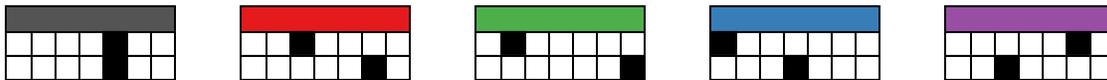


Figure 13: Representation of codes $p(X_\Theta|Y_\Theta, \Theta)$ by a heat-map using inverse grayscale for the experiment with $|X| = 7$. For each node, the rows represent a sensor state y , while the columns represent an output state x . The colours on top of the nodes are used to distinguish the type of agent to whom the code belongs, and colours are related to those shown in Fig. 11.

In all the performed experiments, we found that for values of $|X| \geq 6$, the blind agent can perfectly predict the environmental state $\mu = 5$ for at least one combination of outputs X_Θ and $X_{\Theta'}$. Interestingly, this new concept, which in this particular experiment can be called the “centre”

of the world or environment, cannot be obtained by looking to individual concepts only.

5 Discussion

We considered four different scenarios of code evolution: in the first one, all agents perceived the outputs of all other agents, including itself. We argued that two main stages of evolution can be recognised: in the first stage, a universal code is established, which can be optimal or not. If it is not optimal, then a second stage will achieve optimality. The same result was obtained in [34], in a model of the evolution of the genetic code (represented as a probabilistic mapping between codons and amino acids), although universality and optimality were simultaneously achieved.

In the mentioned work, which developed further the ideas of [38, 39], the authors argue that the universality of the genetic code is a consequence of early communal evolution, mediated by horizontal gene transfer (HGT) between primitive cells. In this evolutionary process, they argue, larger communities will have access (through the exchange of genetic material) to more innovations, leading to faster evolution than smaller ones. Then, “it is not better genetic codes that give an advantage but more common ones” [34]. Although their model does not explicitly show this property, it is captured in our model. We show that a more common, but not optimal code is widely adopted within a population (see Fig. 6). However, in our model, a code imposes itself as universal not because it provides access to more innovations (in our model there is no “code exchange”, only the outputs are shared), but because the population structure forces the adoption of the most popular code. After this stage, further changes in the code of the agents eventually lead to optimality.

In another related work, [21] explored the origins of language in a scenario consisting of artificial agents with a coupled perception and production of speech sounds. Although this work is focused on plausible mechanisms for the origin of language, it assumes the same similarity principle as we do (hearing a vocalisation increases the probability of producing similar vocalisations), arriving to the same outcome (a universal language, or code). Other works have considered similar principles in the evolution of languages: for instance, the naming game [32] and the imitation game [6]. However, these models assume some common conventions in order to evolve new ones. In this study, our main assumption was that the population of agents depended on common environmental conditions.

Our second scenario, where the structure of the population is a grid, showed how establishing local conventions in early stages of evolution constrains the outcome of the code distribution,

since to reconcile different conventions, several simultaneous changes are needed. On the other hand, in our third scenario, where we let the structure of the population change simultaneously with the codes themselves, such situations are avoided by “disconnecting” clusters with dissimilar conventions. This property enhances evolution, and can potentially lead to the adoption of several different conventions within an increasingly fragmenting, or “speciating” population.

Our last scenario assumed perceptual constraints on the environmental information of each agent, and we looked at emerging concepts within a well-mixed population. This scenario was studied in [20], where, as well as in our study, new conceptualisations of the world emerged as a result of considering together the concepts of every agent. In both studies, the new concept was not representable individually by any agent. Differently from the mentioned study, the new concepts obtained in our study were the result of a simple similarity maximisation principle, while in the work of [20], concepts were obtained through the modelling of an explicit fitness function.

The evolution of conventional codes could be interpreted, in the widest sense, as a form of cultural evolution. For instance, considering the definition of culture given by [25]: “*Culture is information capable of affecting individuals’ behavior that they acquire from other members of their species through teaching, imitation, and other forms of social transmission.*”, it could be argued that a form of cultural information is present in organisms, such as bacteria or plants. Although there is a dependence among the different dimensions on which information is transmitted in organisms (if we assume the dimensions to be, for instance, genetic, epigenetic, behavioural and symbol-based, as proposed by [16]), our model assumes freedom of choice in one dimension, without direct influence on the others.

Finally, communication between individuals of a population opens up the possibility of “signal cheaters”, which could be either individuals that do not produce signals themselves but still perceive those of the others, or individuals who exploit other individual’s learned responses to symbols to their advantage. However, our model does not allow such behaviour, since the code producing the outputs functions, implicitly, as the interpreter of the perceived signals.

6 Conclusion

In the proposed model, we introduced a key assumption which allowed us to evolve, for some structures, universal and optimal codes. This assumption states that an agent cannot distinguish the sources of the outputs it perceives from other agents. Following from this, a universal code will necessarily introduce semantics by relating symbols to environmental conditions (via the internal

states of the agent). Our model proposes an information-theoretic way of measuring the similarity within a population of codes.

In this work, we proposed, as an evolutionary principle, that agents try to maximise their side information about the environment indirectly by maximising their mutual code similarity. This behaviour produces several interesting outcomes in the code distribution of a structured population. Depending on the population structure, it captures the evolution of a universal and optimal code (well-mixed population structure), while also the evolution of different codes organised in clusters (in a freely evolving structure), which allows the establishment of optimal as well as suboptimal conventions.

Finally, we considered a well-mixed heterogeneous population with perceptual constraints on the agents about the environment, and showed how, just by looking at the outputs of agents, it is possible to extract concepts that relate to the environment, concepts that none of the agents of the population could individually represent.

References

- [1] Balázsi, G., Van Oudenaarden, A., & Collins, J. J. (2011). Cellular decision making and biological noise: from microbes to mammals. *Cell*, 144(6), 910–25.
- [2] Barbieri, M. (2003). *The organic codes: An introduction to semantic biology*. Cambridge university press.
- [3] Barbieri, M. (2008). Biosemiotics: a new understanding of life. *Die Naturwissenschaften*, 95(7), 577–99.
- [4] Battail, G. (2009). Applying semiotics and information theory to biology: A critical comparison. *Biosemiotics*, 2(3), 303–320.
- [5] Cohen, D. (1966). Optimizing reproduction in a randomly varying environment. *Journal of Theoretical Biology*, 12(1), 119–129.
- [6] De Boer, B. (2000). Self-organization in vowel systems. *Journal of Phonetics*, 28(4), 441–465.
- [7] Deely, J. (2006). On semiotics as naming the doctrine of signs. *Semiotica*, 2006(158), 1–33.
- [8] Donaldson-Matasci, M. C., Bergstrom, C. T., & Lachmann, M. (2010). The fitness value of information. *Oikos*, 119(2), 219–230.

- [9] Favareau, D. (2007). The evolutionary history of biosemiotics. In *Introduction to biosemiotics* (pp. 1–68). Springer.
- [10] Görlich, D., Artmann, S., & Dittrich, P. (2011). Cells as semantic systems. *Biochimica et biophysica acta*, 1810(10), 914–23.
- [11] Hansen, N., & Kern, S. (2004). Evaluating the CMA evolution strategy on multimodal test functions. In *Parallel Problem Solving from Nature-PPSN VIII*, (pp. 282–291). Springer.
- [12] Hansen, N., & Ostermeier, A. (2001). Completely derandomized self-adaptation in evolution strategies. *Evolutionary computation*, 9(2), 159–195.
- [13] Heil, M., & Karban, R. (2010). Explaining evolution of plant communication by airborne signals. *Trends in ecology & evolution*, 25(3), 137–44.
- [14] Heylighen, F. (2011). *Stigmergy as a generic mechanism for coordination: definition, varieties and aspects* (working paper). Vrije Universiteit Brussel.
- [15] Igel, C., Heidrich-meisner, V., & Glasmachers, T. (2008). Shark. *Journal of Machine Learning Research*, 9, 993–996.
- [16] Jablonka, E., & Lamb, M. J. (2005). *Evolution in four dimensions, revised edition: Genetic, epigenetic, behavioral, and symbolic variation in the history of life*. MIT press.
- [17] Kelly, J. (1956). A new interpretation of information rate. *IEEE Transactions on Information Theory*, 2(3), 185–189.
- [18] Kussell, E., & Leibler, S. (2005). Phenotypic diversity, population growth, and information in fluctuating environments. *Science (New York, N. Y.)*, 309(5743), 2075–8.
- [19] Lin, J. (1991). Divergence measures based on the Shannon entropy. *IEEE Transactions on Information Theory*, 37(1), 145–151.
- [20] Möller, M., & Polani, D. (2008). Common concepts in agent groups, symmetries, and conformity in a simple environment. *Artificial Life*, 11, 420.
- [21] Oudeyer, P. (2005). The self-organization of speech sounds. *Journal of Theoretical Biology*, 233(3), 435–449.
- [22] Parunak, H. V. D. (2006). A survey of environments and mechanisms for human-human stigmergy. In *Environments for Multi-Agent Systems II*, (pp. 163–186). Springer.

- [23] Perkins, T. J., & Swain, P. S. (2009). Strategies for cellular decision-making. *Molecular systems biology*, 5(1).
- [24] Platt, T. G., & Fuqua, C. (2010). What's in a name? The semantics of quorum sensing. *Trends in microbiology*, 18(9), 383–7.
- [25] Richerson, P. J., & Boyd, R. (2005). *Not by genes alone: How culture transformed human evolution*. University of Chicago Press.
- [26] Rivoire, O., & Leibler, S.. The value of information for populations in varying environments. *Journal of Statistical Physics*, 142(6), 1124–1166.
- [27] Schuster, M., Sexton, D. J., Diggle, S. P., & Greenberg, E. P. (2013). Acyl-homoserine lactone quorum sensing: From evolution to application. *Annual review of microbiology*, 67, 43–63.
- [28] Seger, J., & Brockmann, H. J. (1987). What is bet-hedging? *Oxford surveys in evolutionary biology*, 4, 182–211.
- [29] Shah, J. (2009). Plants under attack: Systemic signals in defence. *Current opinion in plant biology*, 12(4), 459–64.
- [30] Shannon, C.E. (1948). A mathematical theory of communication. *Bell Systems Technical Journal*, 27, 379–423.
- [31] Slatkin, M. (1974). Hedging one's evolutionary bets. *Nature*, 250, 704–705.
- [32] Steels, L. (1995). A self-organizing spatial vocabulary. *Artificial life*, 2(3), 319–332.
- [33] Theraulaz, G., & Bonabeau, E. (1999). A brief history of stigmergy. *Artificial life*, 5(2), 97–116.
- [34] Vetsigian, K., Woese, C. R., & Goldenfeld, N. (2006). Collective evolution and the genetic code. *Proceedings of the National Academy of Sciences of the United States of America*, 103(28), 10696–10701.
- [35] Waters, C. M., & Bassler, B. L. (2005). Quorum sensing: cell-to-cell communication in bacteria. *Annual review of cell and developmental biology*, 21, 319–46.
- [36] West, S. A., Griffin, A. S., Gardner, A., & Diggle, S. P. (2006). Social evolution theory for microorganisms. *Nature Reviews Microbiology*, 4(8), 597–607.

- [37] Wilson, E. O. (2000). *Sociobiology: The new synthesis*. *Harvard University Press*.
- [38] Woese, C. R. (2002). On the evolution of cells. *Proceedings of the National Academy of Sciences of the United States of America*, 99(13), 8742–7.
- [39] Woese, C. R. (2004). A new biology for a new century. *Microbiology and Molecular Biology Reviews*, 68(2), 173–186.
- [40] Wong, A. K., & You, M. (1985). Entropy and distance of random graphs with application to structural pattern recognition. *IEEE transactions on pattern analysis and machine intelligence*, 7(5), 599–609.

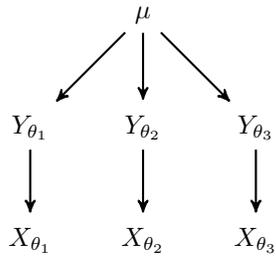


Figure 14: Bayesian network representing the relationship between the sensor and output variables of three agents.

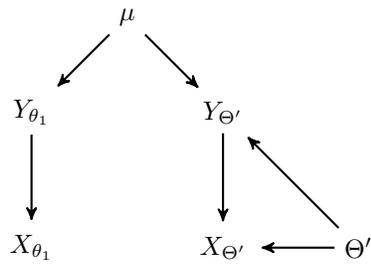


Figure 15: Bayesian network representing the relationships as described above (see text).

$$Pr(Y_{\theta_1}|\mu) := \begin{matrix} & y_1 & y_2 \\ \mu_1 & \begin{pmatrix} 1-\epsilon & \epsilon \end{pmatrix} \\ \mu_2 & \begin{pmatrix} \epsilon & 1-\epsilon \end{pmatrix} \end{matrix} \quad (6) \quad Pr(Y_{\Theta'}|\mu, \Theta') := \begin{matrix} & y_1 & y_2 \\ \theta_2, \mu_1 & \begin{pmatrix} 1-\epsilon & \epsilon \end{pmatrix} \\ \theta_2, \mu_2 & \begin{pmatrix} \epsilon & 1-\epsilon \end{pmatrix} \\ \theta_3, \mu_1 & \begin{pmatrix} 1-\epsilon & \epsilon \end{pmatrix} \\ \theta_3, \mu_2 & \begin{pmatrix} \epsilon & 1-\epsilon \end{pmatrix} \end{matrix} \quad (7)$$

$$Pr(X_{\Theta'}|Y_{\Theta'}, \Theta') := \begin{matrix} & x_1 & x_2 \\ \theta_2, y_1 & \begin{pmatrix} 1 & 0 \end{pmatrix} \\ \theta_2, y_2 & \begin{pmatrix} 0 & 1 \end{pmatrix} \\ \theta_3, y_1 & \begin{pmatrix} 1 & 0 \end{pmatrix} \\ \theta_3, y_2 & \begin{pmatrix} 0 & 1 \end{pmatrix} \end{matrix} \quad (8) \quad Pr(X_{\Theta'}|Y_{\Theta'}, \Theta') := \begin{matrix} & x_1 & x_2 \\ \theta_2, y_1 & \begin{pmatrix} 1 & 0 \end{pmatrix} \\ \theta_2, y_2 & \begin{pmatrix} 0 & 1 \end{pmatrix} \\ \theta_3, y_1 & \begin{pmatrix} 0 & 1 \end{pmatrix} \\ \theta_3, y_2 & \begin{pmatrix} 1 & 0 \end{pmatrix} \end{matrix} \quad (9)$$

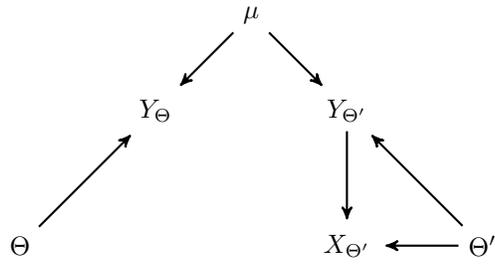


Figure 16: Bayesian network representing the sensor variables of a set of agents indexed by the random variable Θ , and the sensor and output variables of a copy of the set of agents indexed by Θ named Θ' .

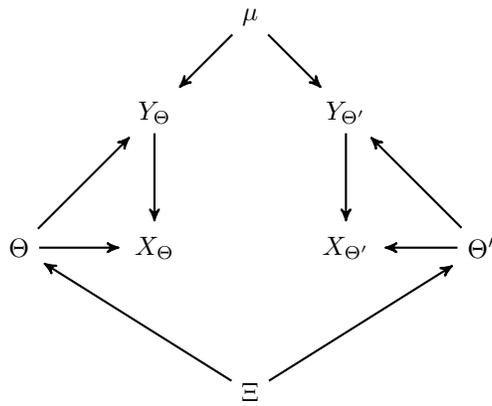


Figure 17: Bayesian network representing the relationship of the variables in the model of code evolution. $Y_{\Theta'}$ is an *i.i.d* copy of Y_Θ and $X_{\Theta'}$ is an *i.i.d* copy of X_Θ . Θ' covers the same set of agents as Θ , but its probability distribution is not necessary the same.

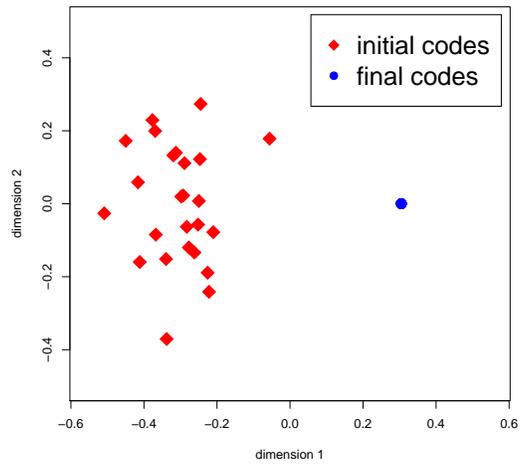


Figure 18: 2-dimensional plot of code distance: red points are codes at the beginning of the optimisation process; blue points are codes at the end of the optimisation process (where the distance between every pair of codes is zero).

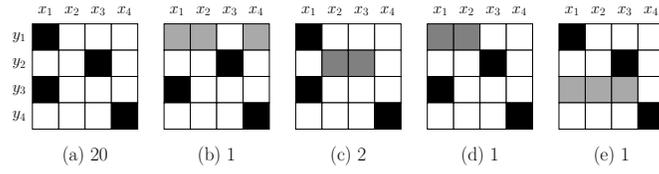


Figure 19: Representation of the codes $p(x|y)$ by a heat-map using inverse grayscale. For each evolved code, we output the number of agents adopting it. This code distribution was achieved with 25 agents in a well-mixed population.

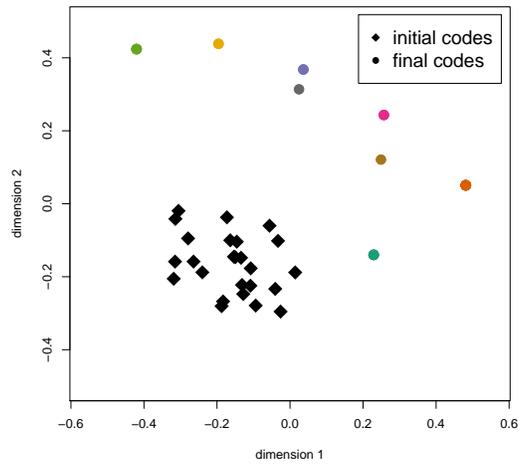


Figure 20: 2-dimensional plot of code distance: points in diamond shape represent codes at the beginning of the optimisation process; rounded points represent codes at the end of the optimisation process. The points are coloured in order to be able to relate this plot with the figure beside it.

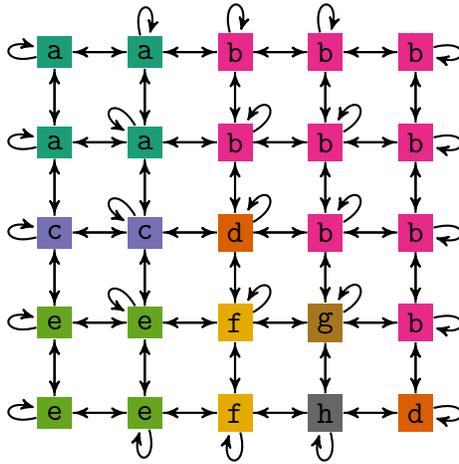


Figure 21: Representation of the spatial structure utilised for the experiment. Agents are assumed to be distributed in a grid: an edge from one agent to another means that one agent perceives the output of the other. Agents are labelled (see Fig. 9) and coloured according to their adopted code.

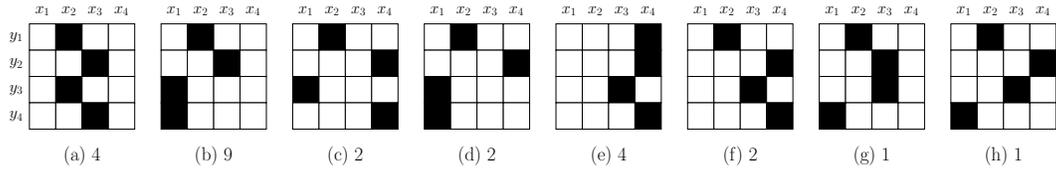


Figure 22: Representation of the codes $p(x|y)$ by a heat-map using inverse grey scale. For each evolved code, we output the number of agents adopting it. This code distribution was achieved with 25 agents in a grid structure.

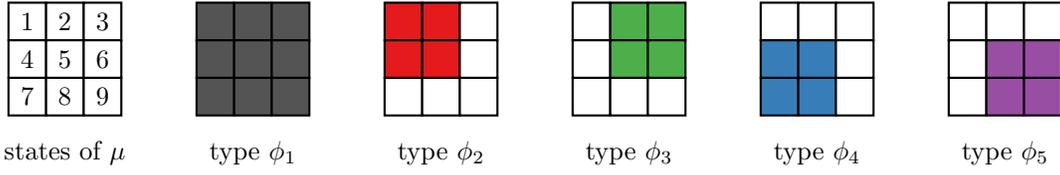


Figure 23: Representation of the conditional probabilities $p(Y_\theta|\mu)$ for an agent θ of each type. These are defined such that each type of agent can only distinguish between the coloured region and the white region. For instance, the sensor of type ϕ_2 is defined as $Pr(Y = y_1|\mu) = 1$ if $\mu \in \{1, 2, 4, 5\}$, and zero otherwise, and $Pr(Y = y_2|\mu) = 1$ if $\mu \notin \{1, 2, 4, 5\}$, and zero otherwise. For type ϕ_1 , $Pr(Y = y_1|\mu) = 0.5$ and $Pr(Y = y_2|\mu) = 0.5$ ($|Y| = 2$ for all types of agents).

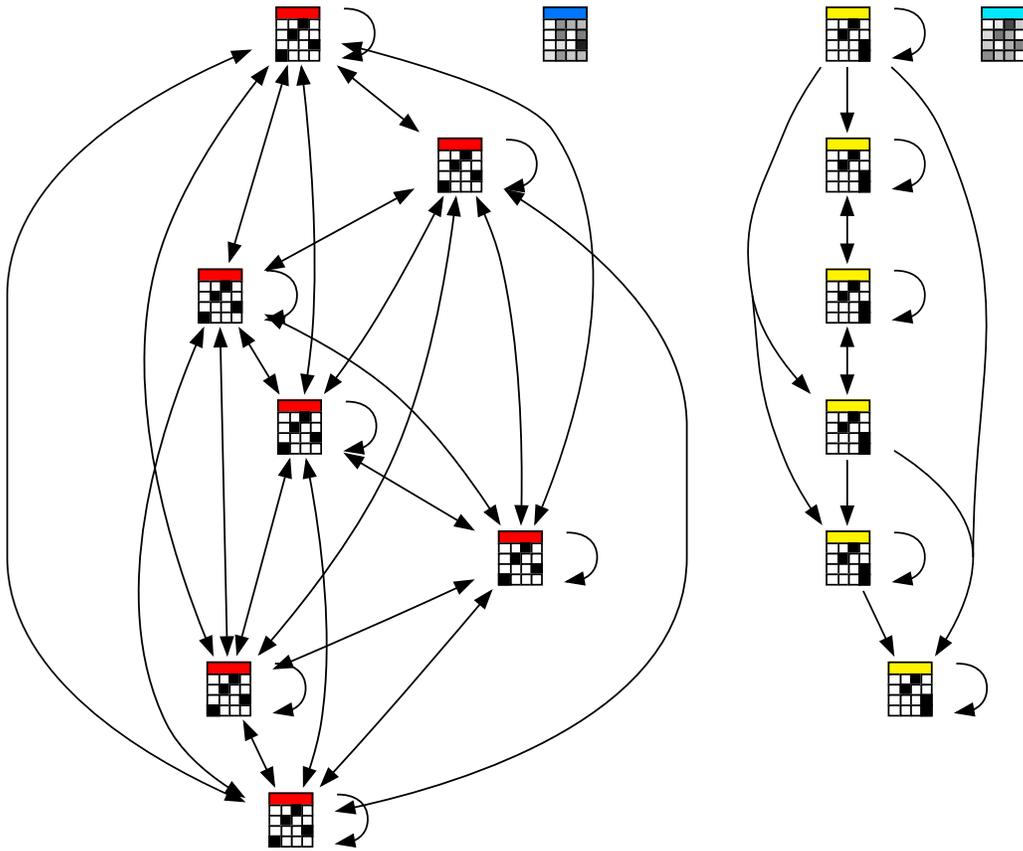


Figure 24: Each node in the graph corresponds to the code of an agent. There is a weighted edge between agent θ_i and θ_j if $p(\theta_i, \theta_j) > 0$ (which is the weight). We omit weights of edges in the graph since they all are roughly of similar value. The temperature colours on top of the nodes indicate the amount of environmental information they would contribute to any agent perceiving only that agents output.

$ X $	$I(\mu; X_{\Theta}, X_{\Theta'})$
2	0.34621
3	0.56555
4	0.71620
5	0.95467
6	1.08139
7	1.18362
8	1.30919
9	1.30919

Figure 25: Results of experiments where the size of the alphabet of a population varies. The maximum amount of environmental information that a blind agent can infer is achieved with $|X| = 8$ and remains equal for bigger alphabets. As the size of the alphabet decreases, this information also decreases.

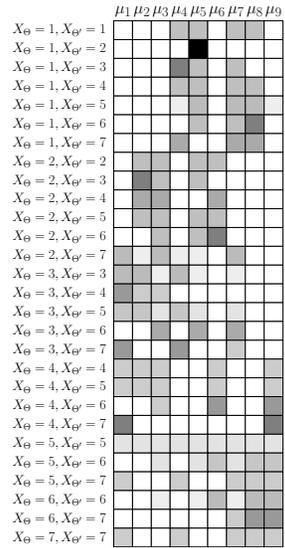


Figure 26: Conditional probability $p(\mu|X_\Theta, X_{\Theta'})$ in inverse grey-scale. Each row represents a combination of values of X_Θ and $X_{\Theta'}$, and each column represents a state of μ .

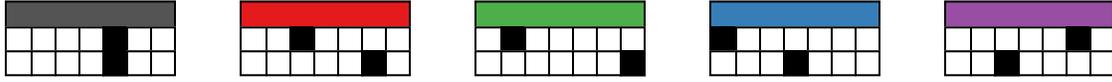


Figure 27: Representation of codes $p(X_\Theta|Y_\Theta, \Theta)$ by a heat-map using inverse grayscale for the experiment with $|X| = 7$. For each node, the rows represent a sensor state y , while the columns represent an output state x . The colours on top of the nodes are used to distinguish the type of agent to whom the code belongs, and colours are related to those shown in Fig. 11.