

Informational parasites in code evolution

Andrés C. Burgos¹
Daniel Polani²

^{1,2}Adaptive Systems Research Group, University of Hertfordshire, Hatfield, UK
¹a.c.burgos@herts.ac.uk

May 22, 2015

Abstract

In a previous study, we considered an information-theoretic model of code evolution. In this model, agents bet on (common) environmental conditions using their sensory information as well as that obtained from messages of other agents, which is determined by an interaction probability (the structure of the population). For an agent to understand another agent’s messages, the former must either know the identity of the latter, or the code producing the messages must be universally interpretable.

A universal code, however, introduces a vulnerability: a parasitic entity can take advantage of it. Here, we investigate this problem. In our specific setting, we consider a parasite to be an agent that tries to inflict as much damage as possible in the mutual understanding of the population (*i.e.* the parasite acts as a disinformation agent). We show that, after introducing a parasite in the population, the former adopts a code such that it captures the information about the environment that is missing in the population. Such an agent would be of great value, but only if the rest of the population can understand its messages. However, it is of little use here, since the parasite utilises the most common messages in the population to express different concepts.

Now we let the population respond by updating their codes such that, in this arms race, they again maximise their mutual understanding. As a result, there is a code drift in the population where the utilisation of the messages of the parasite is avoided. A consequence of this is that the information that the parasite possesses but which the agents lack becomes understandable and readily available.

1 Introduction

Codes shared among entities are ubiquitous in nature, not only present in biological systems, but also, at the least, in technological ones (Doyle, 2010). We define a code as a probabilistic mapping from an “input” random variable (*e.g.* environmental variable) to a set of outputs (*e.g.* messages). A code, then, implies a representation of the input variable. When representations are shared among entities, they become conventions which are used for communication (Burgos and Polani, 2014, 2015). The correct use of these conventions for communicating can be interpreted as “honest signalling”. For instance, the TCP/IP protocol allows the interaction of hardware and software in a code-based, “plug-and-play” fashion, as long as they obey the protocol (Doyle, 2010). In biology, the genetic code acts as an innovation-sharing protocol, one that allows the exchange of innovations through horizontal gene transfer (HGT) (Woese, 2004).

However, communication protocols introduce vulnerabilities: parasitic agents can take advantage of them (Ackley and Littman, 1994; Doyle, 2010). For instance, the chemical cues that ant colonies use to recognise nest-mates can be mimicked by slave-making workers for social integration (D’Ettorre et al., 2002). On the Internet, one can take advantage of machine communication

protocols (TCP/IP) to force target computers to perform computations on behalf of a remote node, in what is called “parasitic computing” (Barabási et al., 2001).

Parasites benefit from their interaction with other agents, while reducing the fitness of the attacked hosts. Nevertheless, parasites can be a positive force in evolution. For instance, they can be generators of biodiversity, achieving more resistance to future attacks (Brockhurst et al., 2004). In an artificial setting, the presence of parasites was shown to attain more efficient communication between agents of a population, increasing their reproductive success (Robbins, 1994). Furthermore, some authors suggest that a healthy ecosystem is one rich in parasites (Hudson et al., 2006). In this work, we study this apparent contradiction from an information-theoretic perspective.

We look at some aspects of the co-evolutionary arms race between host and parasite. Particularly, we would like to characterise informationally the behaviour of parasites and the consequences for the host. For this purpose, we assume a simple scenario where organisms seek to maximise their long-term growth rate by following a bet-hedging strategy (Seger and Brockmann, 1987). We know that maximising their information about the environment achieves this (Shannon, 1948; Kelly, 1956). Then, individuals obtaining extra environmental information from other individuals will have an advantage over those that do not, since they would be able to better predict the future environmental conditions (Donaldson-Matasci et al., 2010). However, as we showed in a previous work, for simple agents which do not have the ability to identify who they are listening to, a shared code among the population is necessary to interpret the transmitted information and therefore improve predictions (Burgos and Polani, 2014, 2015). Here, we keep this assumption with respect to the agents, and we study the effects of introducing a parasite in a population that previously evolved its codes as well as its structure.

2 Model

In our previous model of code evolution (Burgos and Polani, 2014, 2015), the outputs or messages of an agent were produced according to a code, which was a conditional probability from sensor states to messages. The probability of each sensor state of an agent conditioned on the environmental variable μ was given. The information about the environment of each agent was obtained by considering the mutual information between the environmental variable and its sensor variable, together with the outputs of other agents. These outputs would be perceived or not, according to the structure of the population. The codes, as well as the population structure, were optimised in order to maximise what was called the similarity of the codes (we will introduce a more suitable term below) among the interacting agents of the population.

Here, instead, we consider a simplified model where the sensor states of an agent are the agent’s messages, which are represented by a random variable X_Θ . That is, $p(X_\Theta | \mu, \Theta = \theta)$ gives the probability distribution of the sensor states (and, simultaneously, the messages) of an agent θ given the environmental conditions μ .

Agents perceive the sensor states (messages) of other agents according to the structure of the population, which is given by $p(\Theta, \Theta')$. This joint probability induces a weighted graph, where agents represent the nodes of the graph and there is an edge from agent θ to an agent θ' if $p(\theta, \theta') > 0$ (which is the weight of the edge). We interpret $p(\theta, \theta')$ as the probability of interaction between these two agents, and thus we require that $p(\theta, \theta') = p(\theta', \theta)$ (interactions are symmetrical) and $p(\theta, \theta) = 0$ for every agent θ (self-interaction is excluded).

We now consider a population of agents where interacting agents maximise their *mutual understanding*. This is formally defined by $I(X_\Theta ; X_{\Theta'})$, and, when this value achieves its maximum, the mapping that results from the agents’ codes, $p(X_\Theta | X_{\Theta'})$, is deterministic. It is important to note here that this model allows the agents to cluster into different sub-populations due to differences in their codes. Therefore, each sub-population could have its own convention for representing different aspects about the environment, and the conventions used can be as varied as possible, as long as the mapping $p(X_\Theta | X_{\Theta'})$ is universal among all sub-populations.

For cases where the mutual understanding is locally optimal, the codes of the agents are related to each other in one of two possible manners: (a) within each sub-population, all agents have the

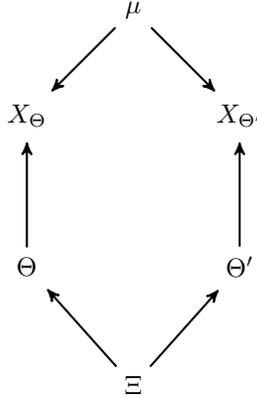


Figure 1: Bayesian network representing the relation of the variables in the simplified model of code evolution. $X_{\Theta'}$ is an *i.i.d* copy of X_Θ . Θ and Θ' selects agents from the same set, but their probability distributions are not necessary the same. These two variables depend on a common variable Ξ to model more general interaction structures.

same code, and the mapping $p(X_\Theta | X_{\Theta'})$ is the identity matrix; or (b) within each sub-population, there are two types of agents (where the type is given by the agent’s code), and the interaction is only between agents of different type. In this case, the graph induced by the interaction probability is bipartite between the types.

For cases where an agent interacts with more than one type of agent, then $p(X_\Theta | X_{\Theta'})$ will necessarily be probabilistic, and thus the mutual understanding among the population will decrease. We can measure the total amount of information about the environment of an agent θ by $I(\mu; X_\Theta, X_{\Theta'} | \Theta = \theta)$. And, since the interaction probability is symmetric, the proposed measure for agent θ is equal to $I(\mu; X_\Theta, X_{\Theta'} | \Theta' = \theta)$. Let us note that, whenever the mutual understanding of a population is optimal, then the individuals that interact necessarily capture the same aspects of the environment. Then, at the optimum of mutual understanding in a population, agents do not increase their information by reading other agent’s messages, although this indeed plays an important role in the evolution of codes. Nevertheless, the ties that an agent establishes are relevant for other purposes, which we study in the following sections.

2.1 Informational parasitism

There are different ways to define an informational parasite. Here, we adopt the model that characterises an informational parasite as an agent π that tries to minimise the mutual understanding between the agents with whom it interacts. An informationally antagonistic parasite is not typical for biology, as the parasite is concentrating at abusing the host system for its own interest, but does not care about the host except for avoiding detection. However, in the context of social networks or news sources, such a parasite can be considered a “troll” or a “disinformation” (FUD) agent who has direct interest in damaging the mutual understanding of the other agents of the population and/or their confidence in their knowledge of the true state of the environment.

In our case, the parasite will choose its code in such a way that the value $I(X_\Theta; X_{\Theta'})$ is minimised. This is an extreme case of parasitism, where the parasite may kill its host as a result of maximising damage. Usually, the known parasites manipulate their hosts in order to benefit from it, decreasing their fitness such that it would not kill them (Schmid-Hempel, 1998). Although the defined type of parasite is not common in biology, it is still a possibility in the range of behaviours that decrease fitness of the host while increasing the attacker’s fitness.

We now analyse the consequences of introducing a parasite in a population for a few very simple, but illustrative, scenarios. First, let us define the following types of codes:

Let us analyse a few simple scenarios where a parasite attacks a population. Let us assume that two (non-parasitic) agents share the same code, for instance, agents θ_1 and θ_2 are of type ϕ_1 ,

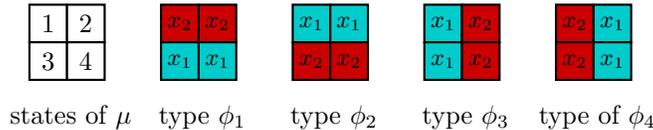


Figure 2: The left-most grid shows an illustration of the environment, although it does not denote its real structure. Then, each type shows a partition of the environmental states induced by an agent’s code (codes here are deterministic). The types ϕ_1 and ϕ_2 capture the first bit of μ . Types ϕ_3 and ϕ_4 capture the second bit of μ .

and that these agents interact only with each other (their mutual understanding is 1 bit). Now, if we minimise their mutual understanding by introducing a parasite π , the optimal configuration is the following: the parasite interacts with one agent only and the parasite’s code is of type ϕ_2 (the “opposite” of type ϕ_1). The mutual understanding between all three agents now is zero. Let us note that, in this case, the environmental information of each agent, before and after introducing the parasite, is 1 bit, which is the amount of information each of them acquire through their corresponding sensors.

Let us consider now a more interesting case: two sub-populations of two agents each, where agents only interact with agents within the sub-population and where the codes are the following: in the first sub-population, we have agents θ_1 and θ_2 of code type ϕ_1 (they capture the first bit of μ); and in the second sub-population, we have agents θ_3 and θ_4 of code type ϕ_3 (they capture the second bit of μ). Now, when we introduce a parasite, two configurations achieve zero mutual understanding: in both, the parasite interacts with every agent, but in the first one, it adopts the “opposite” code of agents θ_1 and θ_2 , which is code type ϕ_2 , while in the second one, it adopts the “opposite” code of agents θ_3 and θ_4 , which is code type ϕ_4 . Here, in the first case, the environmental information of agents θ_3 and θ_4 increases, since the parasite conveys information captured by agents θ_1 and θ_2 (captured by the parasite through its “opposite” code) that the former two agents do not possess. In the same way, for the second configuration, agents θ_1 and θ_2 are benefited by their interactions with the parasite, since here also, the parasite conveys information they lack.

3 Methods

All optimisations were performed using a genetic algorithm (GA). We utilised the C++ library GALib v2.4.7 (Wall, 1996). The GA searches for a particular (possibly local) optimum, and this optimum corresponds to an evolutionary process that has converged. In order to accelerate computations, we assume the following: the probability distribution over μ is uniform, the codes of all agents are deterministic, and the interaction probability between any two agents is given by one over the amount of interactions. 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.

4 Results

We study the introduction of a parasite in a population where the mutual understanding was previously maximised. We consider a population of 256 agents, in an environment with 16 equally likely states, where agents can encode the environment using 16 different symbols. In this way, agents can potentially capture by themselves all the information about μ . As a result of the optimisation process, we obtained 5 sub-populations, where, in each of them, the induced interaction graph is bipartite. Therefore, there are 10 different codes globally, two per sub-population. In Fig. 3 we show the distance between the resulting codes, with point’s size proportional to the number of

agents that adopted each code. The distance used is the Jensen-Shannon divergence (see (Burgos and Polani, 2014)). The average mutual understanding in the population is $I(X_\Theta; X_{\Theta'}) = 3.93$ bits, which, coincidentally, is also the mutual understanding within each sub-population (this does not need to hold necessarily).

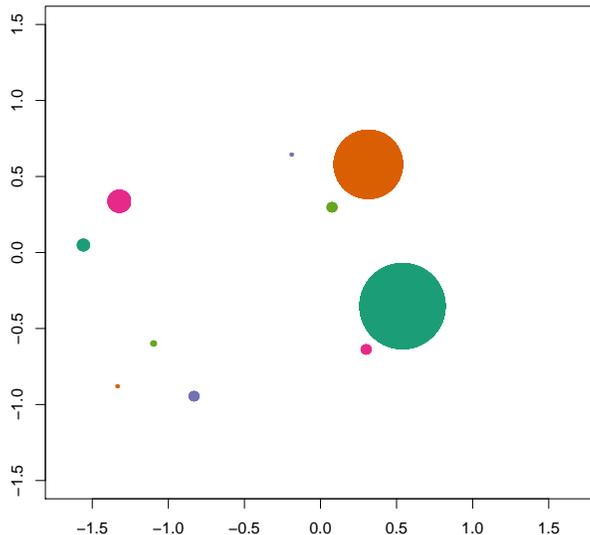


Figure 3: 2-dimensional plot of the distance between the codes: each point represents a particular code, and its size is relative to the number of agents adopting that particular code. The colour of the points denotes the sub-population to which the codes belong.

4.1 Blending in with the crowd

We introduce now a parasite π in the population, and we let it freely choose with whom it interacts, as well as its code (the parasite is introduced before the optimisation process begins, at generation 0). However, the parasite is allowed to use 32 symbols to encode μ , instead of 16, as we did for the rest of the agents. The reason for allowing a larger set of symbols to the parasite is that, otherwise, we will be forcing the parasite to use the symbols used by the population. We allow the double amount of symbols to enable the parasite to perfectly encode the environment while avoiding all symbols already in use.

We found that after optimisation, the parasite interacts with every agent of the population, and its code distance to every other agent is maximal (the distance is 4, the maximum achievable with 16 states). The resulting average mutual understanding now is $I(X_\Theta; X_{\Theta'}) = 2.55$ bits (before the attack, it was 3.93 bits), and the code of the parasite is shown in Fig. 4.

To understand the choice of code by the parasite, we analyse the joint probability $p(X_\Theta, X_{\Theta'})$ before introducing the parasite in the population. Our results show that the parasite encodes the environment through the messages that are most commonly used among the population. In this case, the parasite chose 4 messages (x_3 , x_{12} , x_{13} and x_{16}), all of them among the most popular in the population (see Fig. 5).

As a consequence of this antagonistic behaviour, the parasite blends in the population. This suggests that the parasite would try to avoid being identified by its messages. Our model allows us to measure how “identifiable” agents are by comparing the average joint messages. For instance, this can be measured by the mutual information between the agent selector and the joint messages, $I(\Theta | X_\Theta, X_{\Theta'})$. For a population with a universal code, this measure is zero, that is, we cannot identify any agent. Here, we want to know particularly how much we can identify the parasite by its messages. Then, we can consider the following measure:

$$D_{KL}(p(X_\Theta, X_{\Theta'} | \Theta = \pi) || p(X_\Theta, X_{\Theta'})) \quad (1)$$

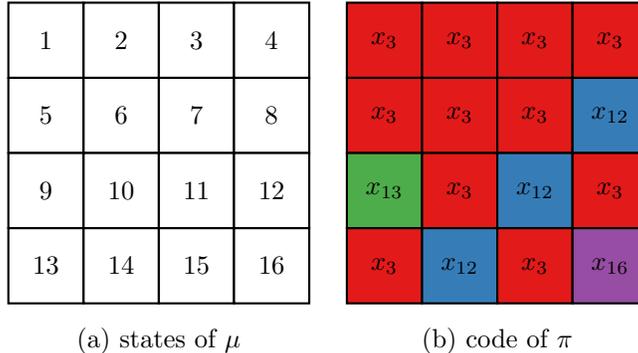


Figure 4: (a) Illustration of the environment μ , although the grid does not denote its real structure. (b) Partition of the environmental states induced by the code of the parasite π .

In Fig. 6, we show the values of Eq. 1 during the optimisation process, which shows that the divergence diminishes as the parasite minimises the mutual understanding of the population.

4.2 Missing environmental aspects

However, it is not only a question of choosing common messages: they must not coincide with other agent’s codes given the environmental conditions. In other words, these messages that are popular among the population will be used by the parasite to express different aspects of μ . Otherwise, if the parasite expresses overlapping aspects, then there might be coincidences with one or more sub-population’s adopted conventions. Consequently, the parasite will capture missing aspects in the population. This can be measured by how much information the code of the parasite adds about μ if we look at the average messages. Formally,

$$I(\mu ; X_{\Theta'} | X_{\Theta}, \Theta' = \pi) \tag{2}$$

The value of Eq. 2 is plotted in Fig. 7 during the optimisation process. After convergence, we have that $I(\mu ; X_{\Theta'} | X_{\Theta}, \Theta' = \pi) = 1.30$ bits, while what the parasite acquires, through its sensors only, is $I(\mu ; X_{\Theta} | \Theta = \pi) = 1.32$ bits. That is, almost all the information it captures is missing in the population. If we consider the perceived information from μ together with the environmental information provided by the population, the parasite captures $I(\mu ; X_{\Theta}, X_{\Theta'} | \Theta = \pi) = 4$ bits, which is the maximum possible, and this means that the parasite always correctly predicts the environment.

4.3 Robustness against parasites

After the parasitic attack, each sub-population has diminished its mutual understanding by different quantities. Although the (former) sub-populations now share a common agent (the parasite, and thus are not strictly speaking different sub-populations), we maintain the colours used in Fig. 3 to identify them. In Table 8 we show a summary of the outcome of the parasitic attack.

As Table 8 shows, in general, larger sub-populations are less damaged by a parasitic intrusion. This phenomenon is due to the large number of interactions among friendly agents, which diminishes the influence of any single agent by considering the average of the perceived messages. The exception in the example is the second largest sub-population, which becomes more damaged than the third largest sub-population. The reason why we see this is that the former sub-population is highly unbalanced, having a small number of agents of one type. Then, agents of the most numerous type interact only with a small number of agents, and therefore are more vulnerable to malicious agents.

Another way in which a population can defend itself against parasitic attacks is through diversification of their codes. Particularly, agents can reduce damage by using synonyms to express

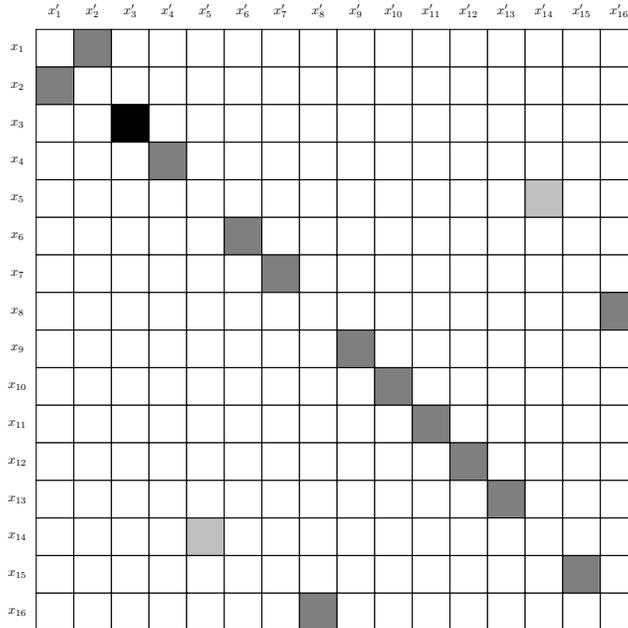


Figure 5: Joint probability of messages $p(X_\Theta, X_{\Theta'})$ of the population before the parasitic attack. Values are normalised to the maximum of all values, in this case $p(x_3, x'_3) = 0.125$. White squares have probability zero.

the same conditions. The presence of synonyms presents an obstacle for the parasite: when trying to confuse agents by expressing different conditions with a chosen symbol, the meaning of the correspondent synonym is not obfuscated.

We study this by comparing populations with different amounts of code types, while maintaining the same population structure. The setup is the following: the population is well-mixed (every agent interacts with every other agent), and first we randomly sample a code for every one of the 64 agents with symbols in the range $[1, 16]$. This population has one type of code only, and the used sample has a mutual understanding of 3.5 bits. Then, we produce a new population by modifying the code of half of the agents, such that $p(x + 16 | \mu) := p(x | \mu)$ and then we set $p(x | \mu) := 0$ (here, $x + 16$ is a synonym of x). In this way, the mutual understanding remains the same for the modified population, which now has two types. Each new type is introduced in a similar fashion, always mapping the original code to a set of (16) unused symbols.

We perform the minimisation of the mutual understanding on each population by introducing a parasite until convergence. We show in Fig. 9 the values of $I(X_\Theta ; X_{\Theta'})$ during the optimisation process for each population. Our results confirm our expectations: more diverse populations are more resistant to parasitic attacks.

4.4 Code diversification

Now we let the population respond to the parasitic attack. If we let the rest of the agents respond to the parasite by freely changing the structure of the population, then our simulations shows that the parasite becomes isolated from the population, which is the expected outcome. However, we consider here a scenario where the structure of the population is maintained, and agents can only respond to the attack by updating their codes. In the same way as we did with the parasite, we allow the agents to choose from a larger set of messages (we consider 32 symbols to give the option to agents of changing completely their encoding of μ).

After convergence, the population's mutual understanding recovered to a value of $I(X_\Theta ; X_{\Theta'}) = 3.50$ bits (see Table 8 for a summary). In response to the parasitic attack, the agents of the popula-

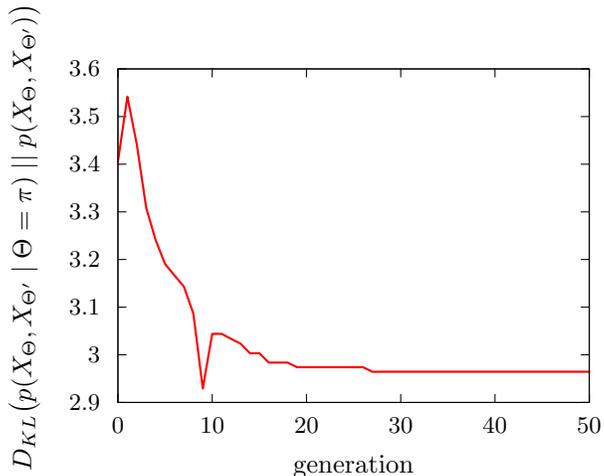


Figure 6: We measure how much we can identify the parasite from its messages. As a consequence of maximising damage to the population, the parasite blends in. The parasite is introduced at generation 0.

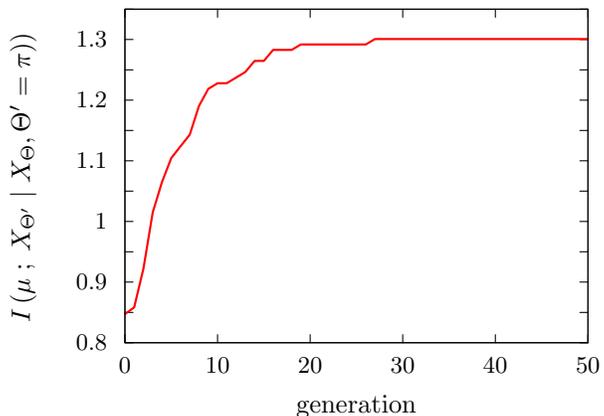


Figure 7: Amount of information the parasite possesses that the population lacks. The parasite is introduced at generation 0.

tion replaced, mostly, the symbols utilised by the parasite with unused ones. In Fig. 10, we can see how the joint probability $p(X_{\Theta}, X_{\Theta'})$ changed after the population’s response. Here, the symbols present in the parasite’s code (x_3 , x_{12} , x_{13} and x_{16}) are mostly removed from the population’s codes.

Three important features follow from the population’s response: first, new code profiles are created in the population. For instance, the orange, purple and green sub-populations shown in Fig. 3 now consist of three types of codes (see Fig. 11). Nevertheless, the bipartite property is kept, but, instead, synonyms were adopted by one type in these sub-populations. This is due to the large amount of symbols available that, in the case they are not in use within the agent’s type, are detached from any meaning and thus would not create confusion.

This can be appreciated in Fig. 12, where we represent the code of an agent before and after the population’s response to the parasitic attack. This agent updated its code such that most symbols used by the parasite are avoided (x_3 and x_{12} are changed for x_{29} and x_{21} , respectively). On the other hand, x_{13} is kept. To check whether this is an optimal solution, we manually updated the code of all agents of the same type, changing x_{13} with every other possible symbol. Indeed, using this particular symbol occupied by the parasite maximises the population’s mutual understanding. The reason for this is that, since all other symbols are occupied by more than one agent, x_{13} is

Colour	Size	Sizes of types	I_1	I_2	I_3
	47	41, 6	3.93	2.85	3.59
	35	33, 2	3.93	2.15	3.05
	16	11, 5	3.93	2.54	3.33
	8	5, 3	3.93	2.16	2.98
	7	5, 2	3.93	1.83	2.65
	113	95, 18	3.93	2.55	3.50

Figure 8: Summary of the parasitic attack for each sub-population. The colours of each sub-population are the same as the ones in Fig. 3. I_1 is the mutual understanding before the parasitic attack, I_2 is the mutual understanding after the parasitic attack, and I_3 is the mutual understanding after the population’s response.

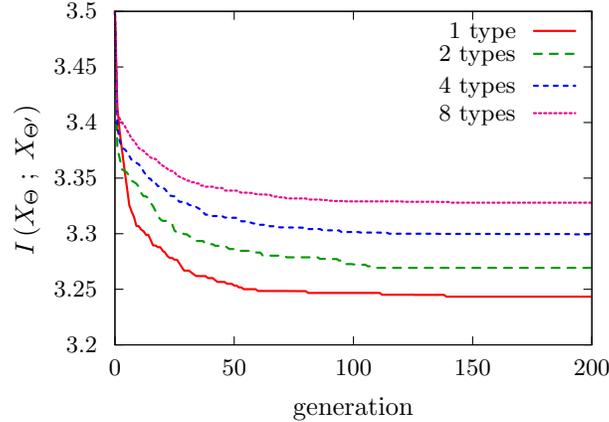


Figure 9: Mutual understanding of populations with varying amounts of types of codes during optimisation. Populations with more types of codes are more resistant to parasitic attacks. In all cases, the parasite was introduced at generation 0.

the one that confuses the population the least.

Second, by drifting from the parasite’s symbols, agents may update their codes in such a way that, after the update, they capture more environmental information. This is the case of the type shown in Fig. 12: before the update, environmental states 9 and 16 (see Fig. 4a to locate these states) were represented by x_3 , while after the update, these states are distinguished from one another.

Third, and most important, the information that the parasite offers can now be understood (although not entirely) by the population: the missing information is mostly expressed using symbols that are not occupied any more by the agents of the population. This cannot be shown in the example, since changes in the agent’s codes after the response to the attack may result in an overlap with the information that the parasite captures. However, we can manipulate the resulting configuration after the parasitic attack to show that agents now consider the information offered by the parasite. For each agent that is not the parasite, we update its code such that $p(x + 16 | \mu) := p(x | \mu)$ and then we set $p(x | \mu) := 0$. In this way, we make sure that all agents capture the same aspects of μ as before the update, without interference (all of the parasite’s symbols are in the range $[1, 16]$).

Now, we measure the average environmental information before and after the change. Before, the value was $I(\mu ; X_\Theta, X_{\Theta'}) = 3.70$ bits, and after, $I(\mu ; X_\Theta, X_{\Theta'}) = 3.72$ bits. The increase

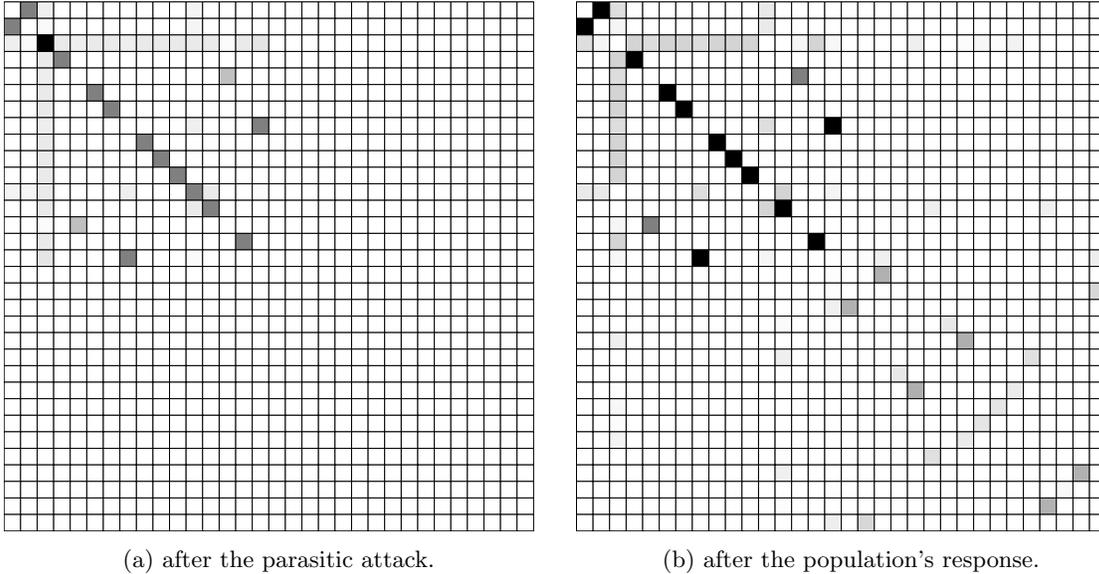


Figure 10: Joint probability of messages $p(X_{\Theta}, X_{\Theta'})$ (a) after the parasitic attack and (b) after the population's response.

is small, but we are considering one parasite only. If we introduce 8 parasites in the population, then the increase in the average environmental information is more significant: from 3.43 bits to 3.73 bits after updating the codes. It is worth noting that, if the parasites interact between themselves, then they would try to capture not only environmental information that is not present in the population, but also that is not captured by the other parasites.

5 Discussion and conclusions

We have considered a scenario where a parasite is introduced in a previously evolved population, and, after convergence, we looked at the response of the population, in one step of many in the co-evolutionary arms race. We considered one step only since, in this setting where the agent's behaviours are not unified, the arms race will cycle continuously.

Our model shows interesting behaviour consistent with empirical observations. For instance, parasites are known to mimic the chemical signatures utilised by the attacked host (D'Etorre et al., 2002; Lorenzi et al., 2014). In this way, identification of the parasite by the population becomes harder. We measured this property during the parasite's attack, showing that as it increased damage in the mutual understanding of the population, it blended in. Additionally, we showed that it becomes parasitically dependent on the population, as most of the environmental information it uses to predict the environment comes from the population.

We have also showed which properties a population may have in order to be robust against parasitic attacks. For instance, large populations are more resilient, since its numerous members provide a solid standard from which perturbations become less significant. Another way in which the population becomes resilient is for the population's agents to utilise synonyms. If the parasite intends to create confusion among the population by using messages that have a different meaning for the rest of the agents, then when synonyms are present, then they do not present any ambiguities.

The presence of parasites in a population can be, in the long term, a positive force (Hudson et al., 2006). For instance, they increase the diversity of the population, which in our scenario was manifested in the creation of new types of codes by using synonyms. As we have seen, this makes the population more robust to subsequent attacks in their co-evolution. Second, parasites

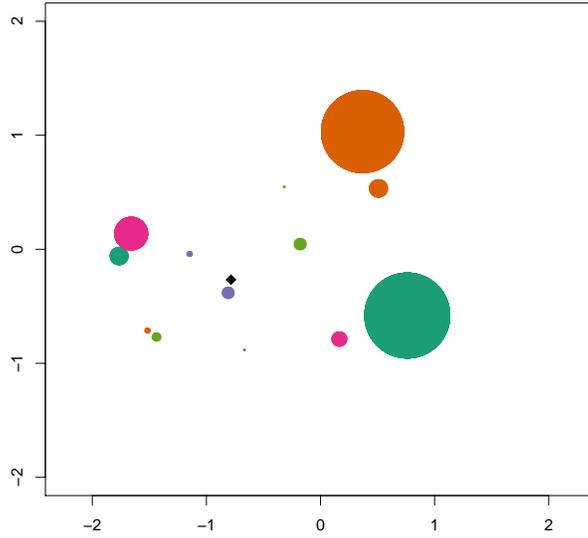


Figure 11: 2-dimensional plot of the distance between the codes after the population’s response to the parasite attack. Each point represents a particular code, and its size is relative to the number of agents adopting that particular code. The colour of the points denotes the sub-population to which the codes belong. The black diamond represents the parasite’s code.

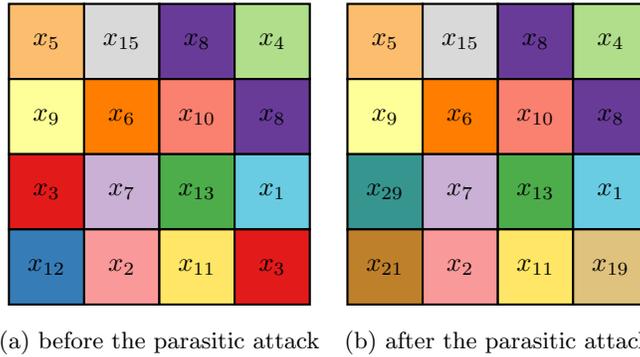


Figure 12: Partition of the environmental states induced by the code of an agent (a) before, and (b) after the parasitic attack.

are able to capture information about the environment which is not captured by any other agent. Most of this information is not understandable by the agents until they respond to the parasitic attack by drifting their codes.

The code drift has two effects: first, it makes the parasite easier to identify, since it is the only agent using a particular set of messages; and second, after the messages used by the parasite are avoided, the parasite’s information becomes understandable for the whole population. Therefore, after the population recovers from the attack, agents can improve their predictions of the environment. The parasite, after the population’s response, can still perfectly predict the state of the environment, but with one major drawback: it becomes easily identifiable, and thus the population have the possibility to take action (for instance, by avoiding interaction) when a future attack begins.

References

Ackley, D. H. and Littman, M. L. (1994). Altruism in the evolution of communication. In *Artificial life IV*, pages 40–48.

- Barabási, A.-L., Freeh, V. W., Jeong, H., and Brockman, J. B. (2001). Parasitic computing. *Nature*, 412(6850):894–897.
- Brockhurst, M. a., Rainey, P. B., and Buckling, A. (2004). The effect of spatial heterogeneity and parasites on the evolution of host diversity. *Proceedings. Biological sciences / The Royal Society*, 271(April 2003):107–111.
- Burgos, A. C. and Polani, D. (2014). An informational study of the evolution of codes in different population structures. In *ALIFE 14: The Fourteenth Conference on the Synthesis and Simulation of Living Systems*, volume 14, pages 352–359.
- Burgos, A. C. and Polani, D. (2015). An informational study of the evolution of codes and of emerging concepts in population of agents. *Accepted to "Artificial Life"*.
- D’Ettorre, P., Mondy, N., Lenoir, A., and Errard, C. (2002). Blending in with the crowd: social parasites integrate into their host colonies using a flexible chemical signature. *Proceedings. Biological sciences / The Royal Society*, 269(1503):1911–8.
- Donaldson-Matasci, M. C., Bergstrom, C. T., and Lachmann, M. (2010). The fitness value of information. *Oikos*, 119(2):219–230.
- Doyle, J. C. (2010). The Architecture of Robust, Evolvable Networks. *The Lee Center*, pages 12–14.
- Hudson, P. J., Dobson, A. P., and Lafferty, K. D. (2006). Is a healthy ecosystem one that is rich in parasites? *Trends in Ecology and Evolution*, 21(7):381–385.
- Kelly, J. (1956). A new interpretation of information rate. *IEEE Transactions on Information Theory*, 2(3):185–189.
- Lorenzi, M. C., Azzani, L., and Bagnères, A. G. (2014). Evolutionary consequences of deception: Complexity and informational content of colony signature are favored by social parasitism. *Current Zoology*, 60:137–148.
- Robbins, P. (1994). The effect of parasitism on the evolution of a communication protocol an artificial life simulation. In *Proceedings of the third international conference on Simulation of adaptive behavior: from animals to animats 3: from animals to animats 3*, pages 431–437. MIT Press.
- Schmid-Hempel, P. (1998). *Parasites in social insects*. Princeton University Press.
- Seger, J. and Brockmann, H. J. (1987). What is bet-hedging? *Oxford surveys in evolutionary biology*.
- Shannon, C. (1948). A mathematical theory of communication. *Bell Systems Technical Journal*, 27:379–423.
- Wall, M. (1996). Galib: A c++ library of genetic algorithm components. *Mechanical Engineering Department, Massachusetts Institute of Technology*, 87:54.
- Woese, C. R. (2004). A new biology for a new century. *Microbiology and Molecular Biology Reviews*, 68(2):173–186.