Journal Pre-proof

An Agent-Based Lattice Model for the Emergence of Anti-Microbial Resistance

R.C. Alamino

PII: S0022-5193(19)30449-7
DOI: https://doi.org/10.1016/j.jtbi.2019.110080
Reference: YJTBI 110080

To appear in: *Journal of Theoretical Biology*

Received date: 5 April 2019
Revised date: 8 November 2019
Accepted date: 11 November 2019

Please cite this article as: R.C. Alamino, An Agent-Based Lattice Model for the Emergence of Anti-Microbial Resistance, *Journal of Theoretical Biology* (2019), doi: https://doi.org/10.1016/j.jtbi.2019.110080

Highlights

- Probabilistic agent-based model of interacting cells.
- Emergence of a unique genotype-phenotype map.
- Reversal of resistance is statistically characterised.
- New microscopic proxy for measuring resistance.
- Machine learning techniques can allow inclusion of real molecular data.

Noname manuscript No.
(will be inserted by the editor)

An Agent-Based Lattice Model for the Emergence of Anti-Microbial Resistance

Alamino, R.C.

the date of receipt and acceptance should be inserted later

Abstract This work introduces a new probabilistic agent-based lattice model for studying the emergence of anti-microbial resistance (AMR) and proposes a new proxy to measure it: the average death probability of the population under the action of the AMD. Both analytical studies and computer simulations of the microscopic behaviour of a bacterial culture interacting with anti-microbial drugs on a discrete lattice are carried out by focusing on the dynamics of this quantity. A unique genotype-phenotype map and classes of AMDs follow as emergent properties and their effects on the possible reversal of resistance are analysed. We also discuss briefly the possibility of using machine learning techniques to learn the model parameters.

Keywords probabilistic model; perceptron; resistance reversal; single-drug protocol

Aston University, Birmingham, B4 7ET, UK
E-mail: alaminrc@aston.ac.uk

1 Introduction

Anti-microbial resistance (AMR), the resistance of pathogens to anti-microbial drugs (AMDs), has been dramatically increasing during the last decades, quickly becoming a threat as dangerous as climate change (Woodford and Livermore 2009; Andersson and Hughes 2010). With levels of resistance dangerously close to the pre-antibiotic era and decreasing rates of AMD discovery (Charles and Grayson 2004), we risk becoming once again defenceless against infections. The issue is so pressing that the World Health Organization (WHO) has suggested a global action plan to address this problem (WHO 2015).

Although not restricted to them, current studies focus mostly on bacteria as they are responsible for a large number of serious diseases, can develop AMR in a plethora of ways and their evolution can be quickly analysed *in vitro*. In addition to mutations in their chromosomal DNA, bacteria can also benefit from horizontal gene transfer (HGT) through the intermediate action of phages (bacteria-infecting viruses) or by acquiring and exchanging circular DNA fragments called *plasmids*, which frequently encode resistance genes (Ng et al 2010). Plasmids are independent from the chromosomal DNA, can also mutate and largely contribute to HGT (San Millan et al 2015; Baltrus 2013; Andersson and Hughes 2010) by being exchanged between bacteria of different species or freed into the environment upon death as the cell's membrane breaks down.

There are several attempted methods to deal with resistance, although none of them has been enough to alleviate the problem satisfactorily. One common practice is simply to avoid (for a potentially long amount of time) using a certain AMD for which resistance has emerged. Although resistance seems to decrease in general with this protocol, its efficiency has been disputed (Barbosa and Levy 2000). Evidence shows that the reversal rate can be slow (Austin

et al 1999; Tan et al 2011) and, even if reversal is observed, AMR may not go back to original levels. Although it might be possible that resistance disappears if one waits long enough, the time scales might be so large that they can be considered unattainable for all practical purposes. Other common treatment protocols make simultaneous use of more than one AMD (Bonhoeffer et al 1997) – a practice that, however, can open the door for the emergence of multiple resistance. As a rule of thumb, the less AMDs are used, the less problems we have with resistance.

There is a large literature on modelling the dynamics of AMR using differential equations which have been used to predict and analyse different aspects of the problem (Bonhoeffer et al 1997; Alexander et al 2007; D'Agata et al 2008; Obolski and Hadany 2012; Ternent et al 2015). These are effective models in which the relevant dynamical variables are obtained by averaging local quantities over large populations. The dynamical equations contain a possibly high number of macroscopic adjustable parameters that have to be included *ad hoc* to allow for a better fitting of observed features – the coefficients of the terms of the differential equations. In biology, they are known simply as *continuous models*.

In *microscopic models*, one seeks the behaviour of the same quantities, but the aim is to derive their dynamics from the interactions of the systems components (e.g., bacterial cells). By finding the rules that set the fundamental processes of life and death of an individual cell and the way it interacts with its environment (including other cells), one seeks to derive from them the same equations as before, but connecting the macroscopic parameters with the microscopic quantities. This strategy provides a better framework for microbiology experiments and also allows for finer modelling, taking into account the statistical variability of the results which is smoothed out by the effective equations. This variability is the result of several sources of stochasticity in

the process as the variance in genotypes, the variations in initial and environmental conditions in each realisation of the process and uncontrollable sources of noise which can appear in the physical and chemical processes involved, which includes the death, reproduction and response to the AMDs from each particular cell.

Effective and microscopic models are complementary rather than competing techniques, with their own range of applicability, advantages and drawbacks. The former are usually of lower computational cost, while the latter help in understanding basic principles behind the phenomenon, leading possibly to a better control and to refinements and/or corrections. From a more technical point of view, differential equations might suffer from serious convergence problems, while microscopic models are better controlled and are only limited by their computational running time when the number of microscopic components is very large (although it must be reckoned that this is typically the case).

The use of microscopic models is not uncommon in several areas of biology (Anderson et al 2007), in particular systems biology. There has been, for instance, a growing literature on using them to model cancer growth (Gerlee and Anderson 2009; Rejniak and Anderson 2011). As the strategies used by cancer cells to avoid the immune system can be similar to AMR strategies, those models can shed some light on the mechanisms of resistance. Because cells are autonomous units, the models are usually said to be *agent-based*. Their popularity resulted in the online availability of many general purpose open-source programs to simulate those models with varied degree of detail (Tisue and Wilensky 2004; Holcombe et al 2012; Gorochowski et al 2012).

Here we introduce a new microscopic agent-based model for the study of AMR emergence and show how it can be used to improve our understanding of it. By modelling bacterial cells as agents on a lattice, we have a better resolu-

tion of the involved quantities, allowing the statistical analysis of the problem and highlighting the probabilistic processes involved. We try to make the model as simple as possible, but still sharing some key features with real situations. The resulting framework is easily adaptable to include other mechanisms (like HGT) and provides a scenario for direct comparison with experimental *in vitro* tests.

In the following, we show how the model is based on reasonable biological assumptions and present results of the simulations which are in qualitative agreement with observations. We do not claim that such a simplified model can become a sharp decision-making tool for treatments in its present form, but we believe that more sophisticated versions of it, obtained with further inputs from experiments, will increasingly contribute in the assessment and development of new strategies against AMR.

In section 2 we introduce the lattice model representing an artificial bacterial culture in a Petri torus (a mathematical idealization of a Petri dish). The cell's response to a certain AMD depends only on two functions taking as arguments the cell DNA and the AMD to which it is being exposed. The resulting fitness landscape is analysed in section 3. The introduced model is used to analyse the effects of a single drug protocol in section 4. Conclusions and further discussions are presented in section 5.

An open-source C code for the simulations generating figs. 6, 7 and 8 is available online at https://github.com/robertoalamino/AMR.

2 Artificial Bacteria

The model here introduced is inspired by a typical laboratory setting. A bacterial culture is grown on a Petri dish containing some pre-defined and fixed concentration of an AMD and its population is recorded as a function of time.

The Petri dish is modelled by a square lattice with bacteria living on its sites. A certain initial configuration of cells allocated at random in the dish (spread randomly on the sites of the lattice) is evolved in time according to pre-defined rules. The temporal data on the appropriate proxies provide information about how different treatments will influence the rates of AMR emergence in the culture.

We work with a square $N \times N$ lattice with periodic boundary conditions (PBC) in both directions, which we call a *Petri torus* due to the resulting topology. The choice of periodic instead of open boundary conditions is not expected to have a significant impact on the results of the simulations for larger lattices and was a question of convenience. To each lattice site (i, j), $i, j = 1, ..., N$, we associate a binary variable σ_{ij}, which is 1 if the site is occupied by a cell and 0 if it is empty, similarly to a lattice gas model (Baxter 1982). PBC imply $\sigma_{i+N,j} = \sigma_{i,j+N} = \sigma_{ij}$. The occupation state of the Petri torus at each instant can then be represented by the *occupation matrix* $\boldsymbol{\sigma}(t) = (\sigma_{ij}(t))_{N\times N}$, which can be understood as a function of the time t and whose entries are functions of both time t and space coordinates (i, j). The configuration shown in fig. 1 for instance has $N = 4$ and its occupation matrix is shown above the lattice in the same figure.

Bacteria will not be allowed to move from one site of the lattice to another. Therefore, their life cycle is equivalent to a probabilistic cellular automaton (Wolfram 2002). The "natural" bacterial life cycle, which excludes the action of the AMDs, depends on two probabilities. At each time step t, every cell has a *reproducing probability* r of dividing in two. The position of the new cell is chosen with equal probability from the empty neighbouring sites to the parent cell. If there is no empty neighbour, the cell does not reproduce. A *natural death probability* d for each living cell at time t includes all *non-AMD* related

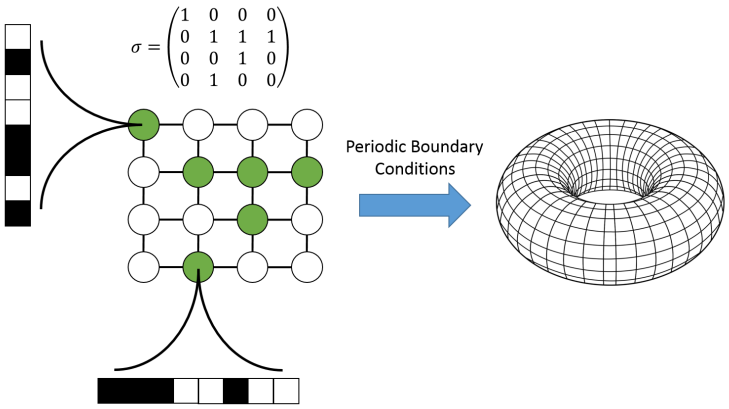

Fig. 1 Artificial bacteria on the Petri torus. The picture shows an example of a small Petri torus with artificial bacteria living on it. The torus is a square grid with periodic boundary conditions. Full circles represent sites occupied by cells. Each cell has its corresponding DNA of dimension $D = 8$, which in the picture is represented by the row/column of boxes (black boxes correspond to the +1 value and white boxes to the -1 value for easier visualization). The matrix just above the lattice is its *occupation matrix*, where occupied sites have the value 1 and empty ones have the value 0.

processes like other adverse environmental conditions, ageing and the patient's immune response.

It remains to model the response of the bacteria to the applied treatment, the latter characterised by the concentrations of the AMD at each site of the lattice. Pharmaceutical companies usually measure the efficiency of AMDs by their *Minimum Inhibitory Concentration (MIC)*, the lowest drug concentration which prevents bacterial growth after a defined incubation period (Davey et al 2015). The MIC is convenient in clinical trials as it avoids the difficulties in isolating the effects of the patients' immune system, but it contains no information about pharmacodynamical properties of the drug (how bacterial growth changes with variations in drug concentrations). It has been proposed that a better proxy is given by the *Minimum Bactericidal Concentration (MBC)*, the concentration that kills at least 99.9% of the bacteria within 24 hours (French 2006). Here we introduce a more convenient quantity from the microscopic

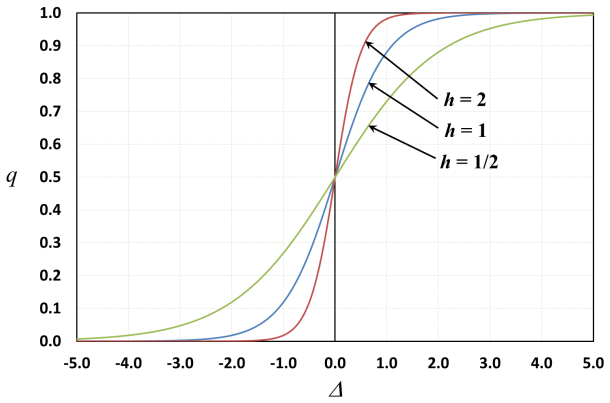

Fig. 2 Death probability. The plot shows the death probability q as a function of the difference Δ between the actual AMD concentration and $\tilde{c}$ for three different values of h. The greater the value of h, the steeper is the curve at zero. In the limit $h \to \infty$, the function has a discontinuous jump from zero to one.

point of view – the concentration $\tilde{c}$ of AMD below which the probability of a cell to die from its action is less than $1/2$. As a concentration, it can vary in the continuous interval $[0, \infty)$. This definition is simpler to implement in a probabilistic microscopic model and, because the studied scenario comprises *in vitro* cultures, it can be actually measured with controlled experiments.

For simplicity, it is assumed that each cell has one single DNA strand encoding its AMD response, which is given by the *total AMD death probability* q_{ij}, the probability that the cell occupying the site (i, j) dies if exposed to the *local* concentration c_{ij} of the AMD. Here we do not consider any HGT, which will be left for future work.

The total AMD death probability is modelled by the heuristic formula

$$q_{ij} = \frac{1 + \tanh\left(h\Delta_{ij}\right)}{2}, \tag{1}$$

where $\Delta_{ij} \equiv c_{ij} - \tilde{c}_{ij}$ and whose plot is shown in fig. 2.

Notice that q_{ij} *is not* the actual fraction of cells dying at a certain moment. It rather measures the response of the organism to a potential use of the AMD. Therefore, even in the absence of any concentration of AMD, random genetic mutations can change the probability of death.

The local value of c_{ij} is fixed during the whole process. We are not considering any diffusion of AMD, although it clearly has an important effect in any bacterial culture on actual agar. Although such a variation with time would be desirable in a more realistic scenario and could be implemented in principle, we will refrain from doing that in the present study for simplicity.

There are two other local (cell/site-dependent) quantities in the above expression. First, the concentration $\tilde{c}_{ij} \in [0, \infty)$, whose objective is to account for physical and chemical mechanisms having threshold behaviours as, for instance, chemical pumps which can become saturated or membranes whose thickness up to a certain point can prevent the AMDs from entering the cell interior. Second, the *sensitivity* $h_{ij} \in [0, \infty)$ regulates the increase/decrease in cell death with variations of AMD concentration and is related to the actual *toxicity* of the applied substance. Both quantities can, in principle, be obtained from actual designed experiments by measuring changes in bacterial populations.

The genotype of each cell will be encoded by a binary chain $\boldsymbol{\pi}_{ij} \in \{\pm 1\}^D$, where D is its integer dimension, i.e., the total number of coordinates in it (see fig. 1) representing abstractly its biological information content. The DNA is responsible for storing the organism's information about how to survive to its environment. The way this information is translated is convoluted as it depends on a series of hierarchical processes. Still, such a mapping is necessary to allow *learning* and *forgetting* (in other words, *evolution*) from the environment. Both are essential informational requirements of adaptation through natural selection and integral elements of mathematical evolutionary models

and genetic algorithms (Mitchell 1998; Schecter and Gintis 2016). Learning in the present model is required for adaptation in the form of AMR acquisition, while forgetting allows acquired adaptations to fade away, as would be the case for reversal of AMR. One of the most studied statistical physics models of biological phenomena with these characteristics is the *Ising perceptron* (Rosenblatt 1958; Engel and van den Broeck 2001), designed to model neuronal responses, with neuron's synapses and stimuli both encoded by binary chains.

The perceptron is one of the simplest known machine learning models and it is characterised by a function taking a multidimensional vector into a number, called the *activation function* for biological reasons. This is usually a general function of the scalar product of its *synaptic vector*, which is a multidimensional parameter encoding the information learned by the perceptron, and the *input vector*, a vector with the same dimensions as the synaptic vector and which encodes the stimuli provided by the environment to which the perceptron reacts. A more detailed description of the perceptron and how it works is provided in appendix A.

In this work, input vectors correspond to the binary encodings of the AMD into a binary chain of dimension 2D, $\boldsymbol{A} = (\boldsymbol{\alpha}, \boldsymbol{\beta}) \in \{\pm 1\}^{2D}$. The two model parameters then become functions of both the environmental conditions (applied AMDs) and the cell genotype (the DNA)

$$c_{ij} = \frac{1 + \langle \boldsymbol{\alpha}, \boldsymbol{\pi}_{ij} \rangle}{1 - \langle \boldsymbol{\alpha}, \boldsymbol{\pi}_{ij} \rangle}, \quad h_{ij} = \frac{1 + \langle \boldsymbol{\beta}, \boldsymbol{\pi}_{ij} \rangle}{1 - \langle \boldsymbol{\beta}, \boldsymbol{\pi}_{ij} \rangle}, \tag{2}$$

where we defined $\langle \boldsymbol{x}, \boldsymbol{y} \rangle \equiv \boldsymbol{x} \cdot \boldsymbol{y}/D$, i.e., the normalised cross product between the two vector arguments. These maps have been chosen as the simplest mappings of the given inputs into the relevant intervals.

Although simple perceptrons cannot approximate general functions, it has been shown that adding one extra *layer*, corresponding to another set of perceptrons doing an intermediate processing of information between the stimuli and the final response, turns them into universal approximators (Cybenko 1989). Variations with several layers, known as *deep neural networks*, have been successfully used in machine learning applications and recently provided a solution for the long sought problem of creating a computer algorithm capable of playing Go on a level comparable to human masters (Silver et al 2016). For our purposes, one layer is sufficient, but the model is flexible enough to allow an extension to a more complex neural network structure.

The dimension of the AMD vector was chosen to be twice as large as the DNA's dimension to allow mutations in the latter to simultaneously affect both model parameters, the well-known phenomenon of *pleiotropy* (Stearns 2010). In nature, each protein usually participates in more than one metabolic process simultaneously. As a consequence, each single mutation might affect more than one of them. Although the standard definition of pleiotropy concerns genes and not bases, we are using here a generalised version of it. One could think about working directly with genes, as it is done in a similar model called MQT model (Taylor and Higgs 2000) which considers random associations and a linear fitness function, in which case the phenomenon would be the standard one. From here on, whenever we use the term pleiotropy, we mean this general definition.

One should notice that $\tilde{c}$ and h are not constants of the model, they are functions of the genotype ($\boldsymbol{\pi}_{ij}$) and AMD ($\boldsymbol{A}$) only. They can vary according to evolving conditions, but only through the variations of $\boldsymbol{\pi}_{ij}$ and $\boldsymbol{A}$, nothing else. More precisely, the freedom comes from the choice of the activation functions leading to them, leaving open the possibility of choosing more

complicated genotype-phenotype maps which might lead to better agreement between observation and theory if needed.

3 Fitness Landscape

According to our model, bacteria acquire resistance by either increasing $\tilde{c}$, decreasing their sensitivity h or both simultaneously. The fitness landscape generated by a given AMD, i.e., the value of q as a function of π might, in general, be a complicated function. We now show that the present model naturally leads to a convenient simplification once we take the limit of large DNA sequences $D \gg 1$.

The simulations to follow will analyse AMDs generated by random distributions. We can then show that the scalar product inside the activation functions, in the limit where the number of bases D is very large, depends only on the means of these distributions. In the case where the coordinates of the vectors $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are i.i.d., with means represented by overbars on the variables, we can calculate the distributions of $\tilde{c}$ and h while D is still finite (but large).

The activation functions have the general form

$$x = \frac{1+y}{1-y}, \qquad y = \frac{1}{D}\sum_i z_i \pi_i, \tag{3}$$

where the $z_i's$ are the appropriate sections of the AMD encoding and the $\pi_i's$ are the coordinates of the DNA. We can prove that the resulting distributions are (see appendix B)

$$\mathcal{P}(x) = \frac{1}{(1+x)^2}\sqrt{\frac{2}{\pi\sigma_y^2}}\exp\left[-\frac{1}{2\sigma_y^2}\left(\frac{x-1}{x+1}-\bar{y}\right)^2\right], \tag{4}$$

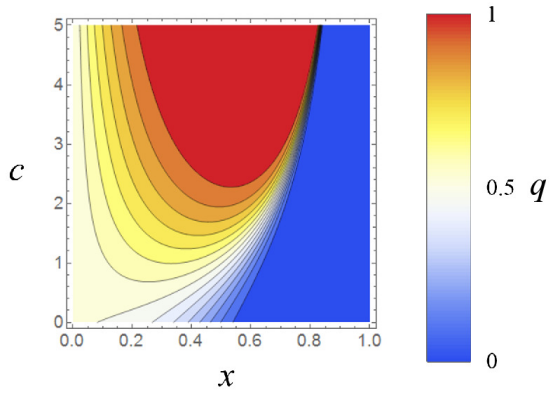

Fig. 3 Fitness landscape. The plot shows the fitness landscape for an AMD with $\bar{\alpha} = \bar{\beta} = -1$ as a contour plot representing the total death probability q as a function of the AMD concentration c and the fraction x of -1 coordinates in the DNA. The landscape shows a global maximum of q at the centre extending to the top of the graph surrounded by decreasing profiles to both side. While higher concentrations will eventually kill all cells leading to $q = 1$, one can see that by changing x in any direction decreases q and, therefore, increases resistance.

with

$$\bar{y} = \frac{\bar{z}}{D}\sum_i \pi_i = \bar{z}\bar{\pi}, \qquad \sigma_y^2 = \frac{\sigma_z^2}{D^2}\sum_i \pi_i^2 = \frac{\sigma_z^2}{D}, \tag{5}$$

where σ_z^2 is the variance of z. For $D \to \infty$, this becomes a delta distribution. Then we can readily calculate the mean of each parameter in this limit as

$$\tilde{c} \to \langle \tilde{c} \rangle = \frac{1 + \bar{\alpha}\bar{\pi}}{1 - \bar{\alpha}\bar{\pi}}, \qquad h \to \langle h \rangle = \frac{1 + \bar{\beta}\bar{\pi}}{1 - \bar{\beta}\bar{\pi}}. \tag{6}$$

As an example, consider the case $\bar{\alpha} = \bar{\beta} = -1$, which gives

$$\tilde{c} = h = \frac{x}{1 - x}, \tag{7}$$

where $x = 1 - \bar{\pi}$ is the fraction of coordinates in the DNA vector whose value is -1. In this case it becomes simpler to plot the fitness landscape. Because it depends not only on the type of AMD, but also on its concentration, it is then convenient to present it as a contour plot of q as a function of both x and the AMD concentration c as in fig. 3.

The particular encodings above were explicitly chosen such that $\tilde{c}$ and the sensitivity h are equal by design, which will not usually be the case in practice. Qualitatively though, this is an important situation where both parameters play an antagonistic role in the evolution of the genotype, contributing in opposite ways for the emergence of AMR. While a higher $\tilde{c}$ improves resistance, a higher sensitivity decreases it. This phenomenon in which the same genes codes for two phenotypes, one of which is beneficial and the other is detrimental to the organism's survival, is known by the name of *antagonistic pleiotropy* (Williams 1957).

Antagonistic pleiotropy is here reflected in the fitness landscape by the maximum of q surrounded by descending profiles both to the left and right. Because this landscape has only one global maximum, we can more clearly see that resistance will eventually emerge as we move away from it in the x direction, which is always observed in simulations.

A crucial point in evolutionary biology is to find appropriate genotype-phenotype maps (GPMs) (Stadler and Stadler 2006; Ahnert 2017). In the above case, we have a GPM that is completely determined by the way the genome interacts with the environment which is given by

$$\Pi(\pi) = D\bar{\pi} = \sum_{i=1}^{D} \pi_i. \tag{8}$$

The phenotype is therefore an *emergent property* – it appears as a collective effect of a very large number of DNA (binary) bases.

From the structural properties of GPMs considered in (Ahnert 2017), the map Π possess *redundancy*, as it is a many-to-one map which maps the set $\pm 1^D$ of dimension 2^D into the set $(-D, -D+2, ..., D-2, D)$ of dimension $D+1$ (an exponential dimensional reduction in size) and *bias*, meaning that the number of genotypes for each phenotype is not the same. More precisely,

the number of genotypes of dimension D corresponding to a phenotype Π is given by the binomial coefficient

$$\binom{D}{(D+\Pi)/2}. \tag{9}$$

The other two properties, *robustness* and *evolvability*, need to be discussed better. The simulations indicate that the map allows the artificial cells to adapt, but on the other hand it is not robust. The *neutral networks* (pre-images in genotype space) of the phenotypes in this GPM are not connected. Whenever a single mutation happens, it necessarily changes the phenotype either by +2 or -2. On the other hand, clearly, the network of phenotypes is fully connected by single mutations and has the topology of a line segment, which allows any phenotype or genotype to be reached from any other by an appropriate sequence of single mutations. Although the map is not robust in principle, in practice, for large D, neighbouring phenotypes result in values of $\tilde{c}$ and h that differ typically by very small quantities, which we can see as a *quasi*-robustness. Single mutations will only affect significantly the answer to the phenotypes that lie very close to the value $+D$, but this effect is very limited and does not seem to affect the evolvability of the the cells. In addition, because the number of genotypes for each phenotype is given by a binomial coefficient, for large D it will concentrate around $\Pi = 0$ with lower probability of being found on those extremes.

This GPM is, of course, a crude approximation to real ones, but as our objective is to study qualitative aspects, we will not seek for further sophistication in the present study. The source of this GPM is the linearity of the dot product used in the activation function, which means that if one wishes a more realistic mapping, additional modelling of the response to the environment should be made. It is true that the resulting delta function can be

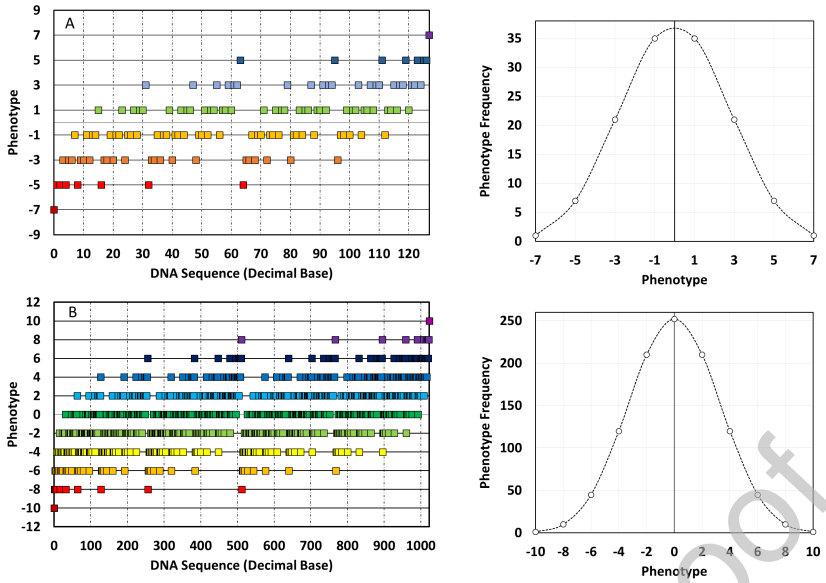

Fig. 4 Genotype-Phenotype Map The diagrams show the genotype-phenotype map induced by the choice of the functional parameters for (A) $D = 7$ and (B) $D = 10$. The squares squares indicate the values of the phenotypes Π given in the vertical axis, while the horizontal axis has the decimal representation of the binary DNA sequences. The plots to the right of each diagram shows the phenotype frequency (vertical axis) as the number of different genotypes corresponding to each phenotype (horizontal axis) out of the respective values of $2^7 = 128$ and $2^{10} = 1024$ possible genotypes.

softened by a different choice of scaling for the scalar product, but that would lead to artificial values for the case we are studying. Fig. 4 illustrates the GPM above for two different values of D, respectively 7 and 10. The value of Π is plotted against the integer representation of the binary sequence of DNA with the convention that -1's are represented by 0's. For instance, $(-1, 1, -1)$ becomes the binary number 010, which corresponds to the integer value 2. The plots to the right of each diagram show the number of genotypes per phenotype for each case.

The above formulas show that, by judiciously choosing the values for the means, one can choose what type of AMD we want to study. Clearly, one would like to approximate the behaviour actual AMDs. Conversely, one can search for

AMDs with certain desired means. For instance, for this very simplified model, there exists one AMD against which resistance cannot evolve, namely the one with $\bar{\alpha} = \bar{\beta} = 0$. This choice makes the activation functions independent of the DNA, completely hindering adaptation. If we do not consider the collective effects, as is the case of biofilm formation which is not allowed in the present model, but only the resistance of each individual, this would be the analogous of bleach or soap for real cells. Both have physical actions that destroy the membrane against which there is no known single cell adaptation.

4 Single Resistance Emergence

Given that most bacterial DNA sequences have between $10^5 - 10^7$ bases, it is reasonable to use equations (6) as approximations in most cases. Throughout our simulations, the values used for D are sufficient for this approximation to be within acceptable precision.

Each initial Petri torus occupation is set randomly by putting a cell in each site with probability 1/2. We will use two different methods to set the initial distribution of genotypes in the lattice. The fastest method computationally is to simply distribute the genotypes uniformly with the same probability. This will generate a binomial distribution with a fixed variance around $\Pi = 0$. We will use this initial configuration throughout the simulations.

To isolate the effect of the AMD, we set the natural death probability to $d = 0$, i.e., cells do not die unless killed by the AMD (we are therefore ignoring cell age, any influence of the immune system or other additional environmental toxicity). The dynamics then follows two steps at each t:

(1) *Reproduction with Mutation:* all living cells are drawn once and only once with the same probability and checked whether or not they will reproduce with

probability r. If reproduction is chosen, then one of the empty neighbours of that cell is chosen with the same probability and the cell generates a child on it. The child cell has a certain probability of *mutation*, which we will later on specify in more details. Basically, one or more coordinates of its DNA can be flipped randomly each time the cell reproduces.

(2) *AMD Death:* one checks whether each cell dies according to the probability q_{ij} for each local AMD. As all cells are checked at this step, the order is irrelevant.

These two steps above are illustrated in the example given in fig. 5. Each run of the simulation consists of T time steps and the information recorded is a double average of the death probability q_{ij} – the average over all living cells and the *quenched average* over initial configurations. To avoid finite size effects, like the stalling of adaptation due to lack of physical space for reproduction, we arbitrarily kill 50% of the cells at random whenever the network becomes fully occupied. For large populations, we expect this to not affect the results significantly. In the simulations we present, this threshold has rarely been reached.

The mutation rate at genotype level can be translated to one at phenotype level, making all equations dependent only on the phenotype. A mutation in the DNA means a flip (i.e., a change of sign) of one of the coordinates of π. We assume that each coordinate has an independent probability m of flipping at each time the cells reproduce, i.e., m is the mutation rate *per base*. Notice that each flip in the DNA means a change of either +2, 0 or -2 in the phenotype. Therefore, it is convenient to write $\Pi_{t+1} = \Pi_t \pm 2\Delta$, such that $\Delta \in \{0, 1, 2, ..., D\}$.

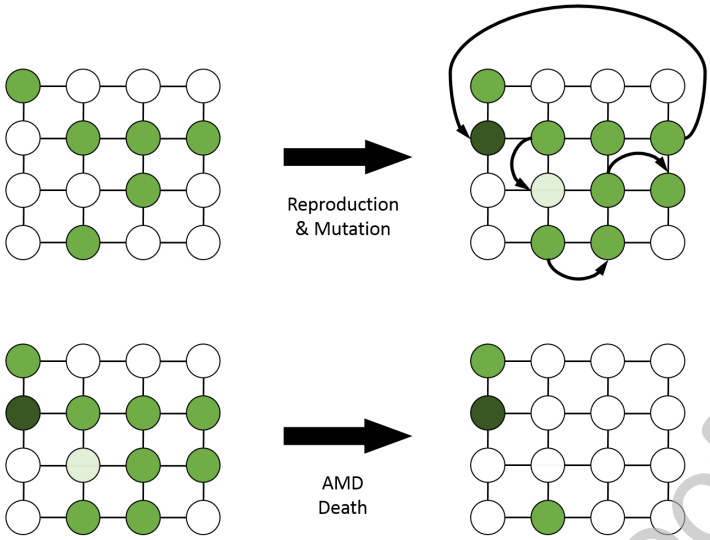

Fig. 5 Simulation steps. The top picture shows an example of a possible change in configuration of the Petri torus after a reproduction with mutation simulation step. Arrows indicate the direction of spreading for each reproducing cell (notice the periodic boundary conditions). Occupied sites with different shades represent mutated cells with respect to their parents. The bottom picture illustrates a possible result of applying an AMD in which one of the mutants from the previous round is more resistant to the AMD (the "darker" mutant) than the other (the "lighter" mutant).

The derivation can be found in appendix C and the final result is given by

$$\mathcal{P}(\Pi_{t+1} = \Pi_t \pm 2\Delta) = \sum_{k=0}^{\lfloor D-\Delta \rfloor/2} \binom{(D \pm \Pi_t)/2}{k} \binom{(D \mp \Pi_t)/2}{\Delta + k} \times m^{\Delta+2k}(1-m)^{D-(\Delta+2k)}, \tag{10}$$

where $\lfloor x \rfloor$ is the floor function, i.e., the greatest integer smaller than x. The above expression needs to be calculated only once for the $(D+1)^2$ possible values of the pair (Π_t, Δ) and can be stored in a file to be accessed each time the simulation is run. Although the above expression is useful for analytical purposes, computationally it is still more convenient to work with genotypes.

In the following simulations, we choose a concentration of AMD (fixed during the whole simulation) such that the initial value of q is 0.3 for the wild-type Π_0 and set $r = 1$ for convenience. This value is chosen to allow time for the cells to adapt to the AMD before the population is wiped out. Higher concentrations will kill all cells before adaptation occurs and will not be useful to obtain information about the dynamics of the model.

Fig. 6 shows the behaviour of the model for large t by plotting the value of q for different AMDs. We used a lattice of linear dimensions $N = 50$ (2500 sites in total), DNA sequences of size $D = 50$ and zero mutation rate, which means that adaptation becomes only a function of the diversity of the initial population. The plot shows the value of q after $T = 3000$ time steps averaged over 100 initial configurations of cell occupations. Although there are no guarantees that at $T = 3000$ adaptation has reached a stationary state, this gives an indication of the relative difficulty to adapt to each AMD. We varied $\bar{\alpha}$ and $\bar{\beta}$ from -1 to 1 in steps of 0.2, giving a total of $11^2 = 121$ data points. The (approximate) symmetry in the plot is clear, with the presence of some antibiotics to which adaptation is relatively easy and those to which it is difficult. Mutation will surely change this picture and we will soon analyse scenarios where it is present.

As the response to each AMD can be very different, we will work with a set of parameters which allows for an easy visualisation of the properties we would like to assess. An analysis of the data presented in fig. 6 indicates that the values $\bar{\alpha} = \bar{\beta} = 0.5$ allows for the emergence of resistance within a reasonable time frame throughout the simulations and will be therefore used for the sake of convenience.

Probably, one of the most popular questions concerning AMR is whether resistance is reversible. As stated before, reversal is a process that is slower than adaptation. On average, the reversal is not complete. The answer, in

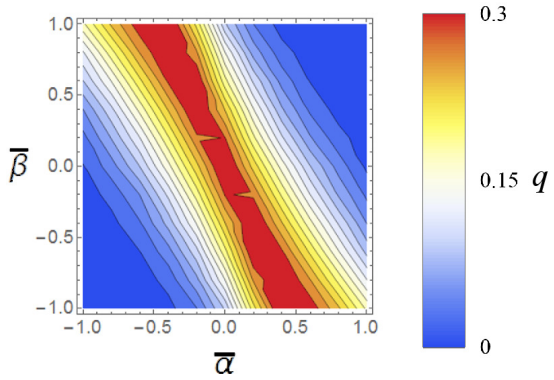

Fig. 6 Comparative adaptation against different AMDs. The plot shows the value of q after 3000 time steps for different values of the pair $(\bar{\alpha}, \bar{\beta})$ for $m = 0$ (no mutations). The values vary from -1 to 1 in steps of 0.2.

fact, depend on a careful consideration of the actual scenario in which AMR appears. For instance, it is a trivial mathematical observation that if the initial *genotype* population is distributed uniformly before the selection pressure, any mutation rate will result in a regression to the original state. If the selection pressure is the introduction of an AMD, by stopping the treatment one would observe a total reversal of AMR. This does not seem to be the actual observed behaviour in real cases.

The difference comes from the fact that the original bacterial population in a patient is already under selection pressure from the environment. By administering an AMD, one creates an additional pressure. When the treatment is stopped, the population has to guarantee that it will remain adapted to the original environment. In order to simulate an analogue situation using our model, we here use an initial population which is uniformly distributed in *genotype* space, which is diverse enough to be able to adapt to a wide range of AMDs. The choice of working with an initial uniform genotype distribution here is that we will take averages over 1000 realisations of each process. In this case, adjusting the initial population using Metropolis-Hastings in-

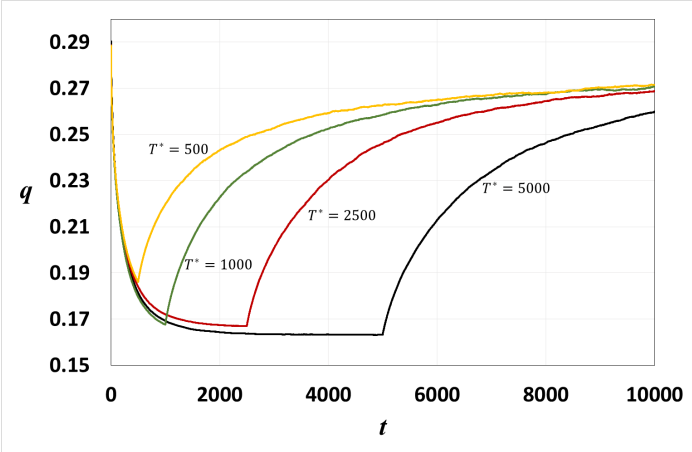

Fig. 7 AMR Reversal. Average death probability q as a function of time t for $m = 0.001$. Each treatment is stopped at a different time, $T^* = 500, 1000, 2500, 5000$.

creases too much the computation time, but because of the long initial adaptive phase before treatment and of the presence of "mutations" (the moves of the Metropolis-Hastings algorithm), does not result in significant differences in the dynamics.

The population is then subjected to a randomly chosen AMD for the first 5000 time steps. This first AMD simulates the action of the original environment. After that, the clock is reset and another randomly chosen AMD is administered for a certain fixed interval of time T^*. Fig. 7 shows the results obtained for different time intervals of treatment.

The curves shown in fig. 7 are averages over 1000 different pairs of AMDs. On average, we see that AMR indeed takes longer to reverse. Although the running times should be extended in order to provide more reliable information, we can see that in practical time-scales reversal is not total.

The average curves, however, should be taken with a grain of salt. Due to the fact that resistance can vary widely for different AMDs, the variances of the curves are very high. In fact, a better characterisation is given by the analysis

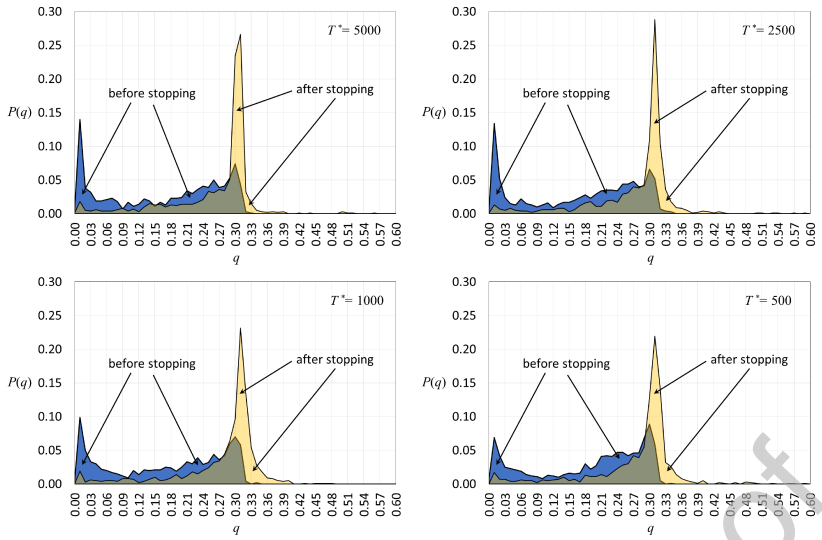

Fig. 8 Reversal frequencies. In each plot, the back shaded histogram (online blue) shows the distribution of q at $t = T^*$, which represents the results *before* the treatment with the relevant AMD is stopped. The front shaded histogram (online beige) is the same distribution at $t = 10000$, which gives the long-term consequences of stopping the treatment at the corresponding T^*. as a function of time t for $m = 0.001$. The different stopping times are given by $T^* = 500, 1000, 2500, 5000$ and indicated on the plots.

of the distribution of values at the relevant time-steps for each situation. This information is provided by the plots of fig. 8

The two shaded plots show the frequencies of the values of q at $t = T^*$ (back shaded are, online blue) and $t = 10000$ (front shaded area, online beige). The bin size is 0.01. The distribution obtained before the treatment is stopped shows that, the longer the AMD is used, the higher the peak near $q = 0$. This means that a *larger* fraction of the population adapts very well. There is a second peak around $q = 0.3$, implying another large number of cells which however cannot adapt. This pushes the average adaptation to lower values, but it also signals that there is a group of cells that will become almost completely resistant. For instance, respectively for $T^* = 500, 1000, 2500, 5000$, the fraction of cells with $q \leq 0.1$ is 25%, 31%, 32% and 34%, all of them very high values.

The long-term distributions obtained after the treatment has stopped tell a more optimistic story. They show that, although on average reversal is not complete, there are very high peaks around the initial value before adaptation $q = 0.3$. In fact, there are even cases in which the levels of susceptibility to the AMD increase above this value, which we will call an *over-reversal*. It seems odd though that the peaks are smaller when adaptation is less efficient, i.e., when the AMD is used for a smaller interval of time. However, this is a result of the fact that there is a larger spread of over-reversals. Respectively, for $T^* = 500, 1000, 2500, 5000$, the fraction of cells that end up with $q \geq 0.27$ is about 73%, 73%, 72% and 70%, which shows that there is more reversal if the AMD is used for a smaller amount of time. The difference does not seem to be too significant, which means that more extensive studies need to be done.

5 Conclusions and Discussion

This work has introduced a new tool for studying the emergence of antimicrobial resistance – an agent-based microscopic model (also know as a single-cell-based model (Anderson et al 2007)) whose agents are perceptrons, the simplest kind of machine learning model. This methodology provides a new point of view from which to study the dynamical mechanisms of resistance spreading by allowing the modelling and analysis of its inherent stochastic aspects.

The use of an agent-based model required the introduction of a new proxy for measuring AMR. We argued that usual ones, MBC or BIC, are not convenient for our simulations and we proposed to use the average probability of death by AMD q. This represents the fact that the reaction to an AMD is not completely deterministic in a population, with several unknown or uncon-

trollable factors contributing to its stochasticity, which is also present in the reproduction with mutation and natural death of the cells.

Unlike deterministic differential equations, the model allows for the consideration of the variance in adaptation, as different contingent paths can lead to different final rates of resistance. This indicates that it is important to analyse the distribution of resistance, which has been overlooked in previous studies. The case of stochastic differential equations would allow for modelling these aspects but, to our knowledge, no model based on them has been proposed so far.

The model presented here is minimal, with few assumptions about the details of biochemical mechanisms in an attempt to be as general as possible. Instead of being implicitly represented by parametrised terms, as is the case in continuous models based on differential equations, the relevant microscopic processes are modelled explicitly. This makes the model flexible enough to be expanded and generalised, including processes that are here not taken into consideration.

For instance, HGT can be incorporated in the model by introducing a probability of exchanging the DNA configuration between adjacent cells. Cell mobility can be achieved by erasing a cell from one site and recreating it in another one. Another important process would be AMD diffusion, which could be simulated by a spread of the AMD to adjacent sites with a corresponding dilution of its concentration.

In the limit of a large number of degrees of freedom, which in the present case means large DNA chains with $D \gg 1$ and grid size $N \gg 1$ (both usually the case for real life scenarios), the model reveals interesting emergent behaviour. In particular, we identified the following emergent properties:

- **AMD Classes:** in the large system limit, the model leads to a fitness landscape that is a function of the average values of the AMD encodings only. This means that AMDs can be classified into large groups with the same typical anti-microbial properties. This is very similar to the real case, where drugs are classified in families like *penicillins* or *cephalosporins*, with some variations inside these groups.

  The above classes of AMDs include some drugs which are impossible to adapt to. Although, this would seem to be an exciting possibility in reality, this kind of AMD already exists, except that they are those substances which are also toxic to the patient. This might contribute to the search for AMDs that can be used efficiently to kill the bacteria without compromising the patient's health. One possible modification of the model to allow for this kind of study could be to introduce a second structurally different agent representing the patient's cells.
- **GPM:** interestingly, this model induces a unique GPM which has many of the most important properties of real GPMs, including the exponential decrease in number from microscopic states (genotype) to macroscopic ones (phenotype). Also, not only the phenotypes, but their distribution emerges in the large system limit too.

The application of this model to the case of single-drug treatments revealed a series of interesting aspects of AMR modelling. For instance, the results of the simulations showed that one must be careful when choosing the initial distribution of bacterial populations. If one uses the simplest choice of a uniform distribution for the genotypes, reversal of resistance *always* happens in this case simply because, no matter what is the mutation rate, they will eventually randomise the DNA chains and reproduce the initial population.

This indicates that it might be more reasonable to use phenotype distributions as selection pressures act on the former instead of the latter. Because every (relatively) stable population will be in equilibrium with respect to some selection pressures defining the environment it is inhabiting, it is more natural to assume that the population that will be treated with some AMD is stable under some selection pressure that will drive the phenotype (or genotype) distribution away from uniformity and favour some particular value. The AMD brings new selection pressures, which forces the population to adapt to the mixed environment. After the AMD is removed, reversal means that the population needs to re-adapt to the initial pressures. In order to simulate this scenario, the initial population was generated by introducing a "dummy" AMD representing the initial environment which acts continuously in the population. Although this would seem to be equivalent to a two-drugs protocol, this is not exactly the same as the first "dummy" AMD (the environment) is never removed.

The simulations then showed that, even if average results for the reversal rate are in qualitative agreement with actual observations, they might hide some crucial information, which we uncovered by looking at frequency plots of the death probability q at key times during the treatment protocol.

One exciting possibility about this model, which will require more involved future research, is to use machine learning algorithms to encode the structure of *actual* AMDs and study them. The genotype can be directly translated to binary code and the macroscopic parameters of the model can be obtained from experiments. Although perceptrons are too simple to approximate general genotype-phenotype maps, it was proven (Cybenko 1989) that more complex networks, as deep networks (Silver et al 2016), are universal approximators and can become powerful tools in the search of real new AMDs and evaluation of resistance scenarios.

There are several other issues which can be tackled by the present model in future versions. Multiple resistance/multi-drug protocols is one of the most important of them. The use of two or more AMDs is based on the hope that it will be difficult for the bacteria to adapt to more than one AMD at the same time. It can be effected in different ways as, for instance, alternating two AMDs or using both at the same time. The efficiency of these protocols is an ongoing object of study with crucial importance to health systems around the world under pressure due to the lack of weapons to fight resistance infections. This kind of protocol is currently being investigated by us and will the object of an upcoming work.

In many situations it might be important to consider the response of the cell to other kinds of external stimuli, like the presence of resources for growth, different physical conditions, competition with other microorganisms and even reaction of the host's body. There are two ways to do that, both equivalent in the mathematical sense. One is to include additional neural networks with relevant parameters contributing to the overall value of q, which would be then more appropriately described as the probability of death given a certain environment. On the other hand, it could be convenient to consider these stimuli separated from the AMD, in which case one could generate a different probability of death by modelling separately additional stimuli and adding another simulation round in which cells are tested against this probability.

Finally, it must be stressed that the major limitation of this approach is the lack of a mechanistic description. This by itself does not prevent the investigation of interesting emergent behaviours, but is a critical hindrance in the use of the model for any actual *in silico* screenings of new anti-microbial drugs as it stands.

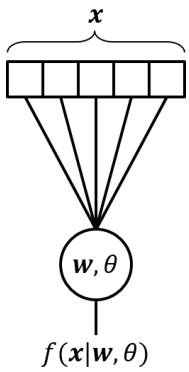

Fig. 9 Perceptron. The perceptron as an elementary processing unit which maps the input vector $\boldsymbol{x}$ into a number through the activation function f given the perceptron's synaptic vector $\boldsymbol{w}$ and activation threshold θ .

Acknowledgements I would like to thank Dr J. Neirotti, Dr M.Stich, Dr M.Chli and Dr A. Cheong for their suggestions and discussions. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

A The Ising Perceptron

The perceptron is the most elementary model of a neural network and is schematically represented in fig. 9. It is intended to do a very basic processing of information by taking a certain input, which is represented by an input vector $\boldsymbol{x}$, and producing an output value y. The function that maps inputs into outputs is in general non-linear (although linear functions can also be used) and parametric, the so called *activation function* f, leading to the equation

$$y = f(\boldsymbol{x}|\boldsymbol{w}, \theta), \tag{11}$$

where the given parameters $\boldsymbol{w}$ and θ are, respectively, the *synaptic vector* $\boldsymbol{w}$ and a real-valued activation threshold. The synaptic vector has the same dimensionality of the input and the activation function is usually written as a function of their scalar product

$$f(\boldsymbol{x}|\boldsymbol{w}, \theta) = g(\boldsymbol{x} \cdot \boldsymbol{w} + \theta). \tag{12}$$

More precisely, the perceptron is intended to simulate the action of a single neuron and more sophisticated neural networks are obtained by connecting them with different network

topologies. They were introduced by Frank Rosenblatt (Rosenblatt 1958) based on the work of McCulloch and Pitts (McCulloch and Pitts 1943) showing that neural networks with units with the features of perceptrons can encode any logical function (although that is not true for one single perceptron).

When the output y is binary, we talk about a *binary perceptron*. The reason for the names *activation function* and *threshold* is that, in this case, g is usually chosen to be either a sign function or a Heaviside theta, depending on the choice of representation for the binary variables. The -/+ or 0/1 results then represent respectively the quiescence or firing of a neuron due to the stimulus $\boldsymbol{x}$ and only occurs if $\boldsymbol{x} \cdot \boldsymbol{w} + \theta \geq 0$.

If the input is binary, we call it an Ising perceptron as, from a statistical physics point of view, each coordinate can be thought as either an up or down spin, represented respectively by the values +1 and -1. It is common to use 0 and 1 for the binary variables, which are the same up to a linear transformation.

The perceptron is capable of learning by adjusting the synaptic weights $\boldsymbol{w}$ such that the correct pairs of input and output $(\boldsymbol{x}, y)$ from a given database of examples are reproduced exactly or within a certain margin of error.

B Distributions of Functional Parameters

Using the general formula for the functional parameters given in the main text, we can write their probability distributions as

$$\mathcal{P}(x) = \int dy \mathcal{P}(y)\mathcal{P}(x|y) = \int dy \mathcal{P}(y)\delta\left(x - \frac{1+y}{a-by}\right), \tag{13}$$

where $\delta(\cdot)$ is the Dirac delta distribution.

If the AMD coordinates are generated independently and equally distributed with mean $\bar{z}$ and variance σ_z^2, then the Central Limit Theorem guarantees that, in the limit $D \gg 1$, we have

$$\mathcal{P}(y) \to \mathcal{N}(y|\bar{y}, \sigma_y^2), \tag{14}$$

i.e., it approaches a Gaussian distribution with mean $\bar{y}$ and variance σ_y^2 given by

$$\bar{y} = \frac{\bar{z}}{D}\sum_i \pi_i, \qquad \sigma_y^2 = \frac{\sigma_z^2}{D^2}\sum_i \pi_i^2 = \frac{\sigma_z^2}{D}, \tag{15}$$

where we used the fact that the DNA sequences are encoded by binary vectors with coordinates in the set $\{\pm 1\}$. Notice that, for $D >> 1$, in the case the π_i's are also i.i.d. with mean $\bar{\pi}$ we can approximate

$$\frac{1}{D}\sum_i \pi_i \approx \bar{\pi}. \tag{16}$$

In order to carry out the y integration, we have to rewrite the Dirac delta using the property

$$\delta(g(y)) = \frac{\delta(y - y^*)}{|g'(y^*)|}, \tag{17}$$

where y^* is the solution of $g(y) = 0$. We can easily show that

$$y^* = \frac{ax - 1}{bx + 1}, \qquad g'(y^*) = -\frac{(1 + bx)^2}{a + b}, \tag{18}$$

and therefore

$$\begin{aligned} \mathcal{P}(x) &= \int dy \mathcal{N}(y|\bar{y}, \sigma_y^2) \frac{a + b}{(1 + bx)^2} \delta(y - y^*) \\ &= \frac{a + b}{(1 + bx)^2} \frac{1}{\sqrt{2\pi\sigma_y^2}} \exp\left[-\frac{1}{2\sigma_y^2}\left(\frac{ax - 1}{bx + 1} - \bar{y}\right)^2\right], \end{aligned} \tag{19}$$

which is equivalent to equation (4). The mean and the variance of x can be calculated using the obtained distribution. In the limit $D \to \infty$ we can find simple expressions if we notice that the Gaussian on y becomes a delta function centred on its mean as the variance goes to zero. Then

$$\langle x \rangle = \left\langle \frac{1 + y}{a - by} \right\rangle = \frac{1 + \bar{y}}{a - b\bar{y}}, \tag{20}$$

and the variance of x also goes to zero, which means that the AMD's cluster around the means. Plugging in the appropriate values of a and b leads to the formulas (6).

C Probability of Phenotype Change

Let us prove equation (10) which gives the probability of changing from phenotype Π_t to $\Pi_{t+1} = \Pi_t \pm 2\Delta$ upon cell reproduction. For simplicity, let us consider the case where D is even. The case of D odd is then easy to obtain.

Consider a genotype π with a number n^+ of $+1$'s and a number n^- of -1's. Let us start by considering the case $+2\Delta$. In this case, there is an excess in the flips from -1 to $+1$ of exactly Δ, but any combination of flips satisfying this condition is allowed. Suppose now that $\Delta + k$ negative coordinates flip (k a positive integer), then k positive coordinates

have to flip as well in order to maintain the overall change in the phenotype. Therefore, we have a factor

$$\binom{n^+}{k}\binom{n^-}{\Delta+k}m^{\Delta+2k}(1-m)^{D-(\Delta+2k)}, \tag{21}$$

meaning that we can choose any $\Delta+k$ negative coordinates and k positive coordinates to flip. The probability of this is then just the probability of $\Delta+k+k$ coordinates flipping while the remaining $D-(\Delta+2k)$ don't flip. All that remains is to add these factors for all possible values of k. Now, we need to have at least Δ negative coordinates to flip. Consider the case in which all coordinates are flipped. Clearly we have the constraint $\Delta+2k=D$, which leads to $k=(D-\Delta)/2$. This works if the quantity $D-\Delta$ is even. When it is odd, one needs to keep at least one positive coordinate fixed and, therefore, $k=(D-\Delta-1)/2$, which can be written in the general case as $\lfloor D-\Delta\rfloor/2$.

By noticing that

$$n^+ + n^- = D, \qquad n^+ - n^- = \Pi_t, \tag{22}$$

we can write

$$n^+ = (D+\Pi_t)/2, \qquad n^- = (D-\Pi_t)/2, \tag{23}$$

which gives

$$\begin{aligned}\mathcal{P}(\Pi_{t+1}=\Pi_t+2\Delta) = &\sum_{k=0}^{\lfloor D-\Delta\rfloor/2}\binom{(D+\Pi_t)/2}{k}\binom{(D-\Pi_t)/2}{\Delta+k}\\ &\times m^{\Delta+2k}(1-m)^{D-(\Delta+2k)}.\end{aligned} \tag{24}$$

The case when the change is -2Δ is analogous, simply changing the role of n^+ and n^-. By putting the two expressions together we obtain the required probability.

References

Ahnert SE (2017) Structural properties of genotype–phenotype maps. Journal of The Royal Society Interface 14(132):20170,275

Alexander ME, Bowman CS, Feng Z, Gardam M, Moghadas SM, Röst G, Wu J, Yan P (2007) Emergence of drug resistance: implications for antiviral control of pandemic influenza. Proceedings of the Royal Society of London B: Biological Sciences 274(1619):1675–1684

Anderson ARA, Chaplain MAJ, Rejniak KA (2007) Single-cell-based Models in Biology and Medicine. Springer Science & Business Media, Basel, Switzerland

Andersson DI, Hughes D (2010) Antibiotic resistance and its cost: is it possible to reverse resistance? Nature Reviews Microbiology 8(4):260–271

Austin DJ, Kristinsson KG, Anderson RM (1999) The relationship between the volume of antimicrobial consumption in human communities and the frequency of resistance. Proceedings of the National Academy of Sciences 96(3):1152–1156

Baltrus DA (2013) Exploring the costs of horizontal gene transfer. Trends in Ecology & Evolution 28(8):489–495

Barbosa TM, Levy SB (2000) The impact of antibiotic use on resistance development and persistence. Drug Resistance Updates 3:303–311

Baxter RJ (1982) Exactly Solved Models in Statistical Mechanics. Elsevier, Suffolk

Bonhoeffer S, Lipsitch M, Levin BR (1997) Evaluating treatment protocols to prevent antibiotic resistance. Proceedings of the National Academy of Sciences 94(22):12,106–12,111

Charles PG, Grayson ML (2004) The dearth of new antibiotic development: why we should be worried and what we can do about it. Medical Journal of Australia 181:549–553

Cybenko G (1989) Approximation by superpositions of a sigmoidal function. Mathematics of Control, Signals and Systems 2(4):303–314

D'Agata EM, Dupont-Rouzeyrol M, Magal P, Olivier D, Ruan S (2008) The impact of different antibiotic regimens on the emergence of antimicrobial-resistant bacteria. PLoS One 3(12):e4036

Davey PG, Wilcox MH, Irving WL, Thwaites G, et al (2015) Antimicrobial chemotherapy. Oxford University Press, USA

Engel A, van den Broeck C (2001) Statistical Mechanics of Learning. Cambridge University Press

French G (2006) Bactericidal agents in the treatment of mrsa infectionsthe potential role of daptomycin. Journal of Antimicrobial Chemotherapy 58(6):1107–1117

Gerlee P, Anderson A (2009) Evolution of cell motility in an individual-based model of tumour growth. Journal of theoretical biology 259(1):67–83

Gorochowski TE, Matyjaszkiewicz A, Todd T, Oak N, Kowalska K, Reid S, Tsaneva-Atanasova KT, Savery NJ, Grierson CS, Di Bernardo M (2012) Bsim: an agent-based tool for modeling bacterial populations in systems and synthetic biology. PloS one 7(8):e42,790

Holcombe M, Adra S, Bicak M, Chin S, Coakley S, Graham AI, Green J, Greenough C, Jackson D, Kiran M, et al (2012) Modelling complex biological systems using an agent-based approach. Integrative Biology 4(1):53–64

McCulloch WS, Pitts W (1943) A logical calculus of the ideas immanent in nervous activity. The bulletin of mathematical biophysics 5(4):115–133

730 Mitchell M (1998) An introduction to genetic algorithms. MIT press
731 Ng JW, Chatenay D, Robert J, Poirier MG (2010) Plasmid copy number noise in monoclonal
732 populations of bacteria. Phys Rev E 81(1):011,909
733 Obolski U, Hadany L (2012) Implications of stress-induced genetic variation for minimizing
734 multidrug resistance in bacteria. BMC Medicine 10(1):89
735 Rejniak KA, Anderson AR (2011) Hybrid models of tumor growth. Wiley Interdisciplinary
736 Reviews: Systems Biology and Medicine 3(1):115–125
737 Rosenblatt F (1958) The perceptron: A probabilistic model for information storage and
738 organization in the brain. Psychological Review 65(6):386
739 San Millan A, Toll-Riera M, Qi Q, MacLean RC (2015) Interactions between horizontally
740 acquired genes create a fitness cost in pseudomonas aeruginosa. Nature Communications
741 6
742 Schecter S, Gintis H (2016) Game theory in action: An introduction to classical and evolu-
743 tionary models. Princeton University Press, New Jersey
744 Silver D, Huang A, Maddison CJ, Guez A, Sifre L, Van Den Driessche G, Schrittwieser J,
745 Antonoglou I, Panneershelvam V, Lanctot M, et al (2016) Mastering the game of go with
746 deep neural networks and tree search. Nature 529(7587):484–489
747 Stadler PF, Stadler BM (2006) Genotype-phenotype maps. Biological Theory 1(3):268–279
748 Stearns FW (2010) One hundred years of pleiotropy: a retrospective. Genetics 186(3):767–
749 773
750 Tan L, Serene S, Chao HX, Gore J (2011) Hidden randomness between fitness landscapes
751 limits reverse evolution. Phys Rev Lett 106(19):198,102
752 Taylor CF, Higgs PG (2000) A population genetics model for multiple quantitative traits
753 exhibiting pleiotropy and epistasis. Journal of theoretical Biology 203(4):419–437
754 Ternent L, Dyson RJ, Krachler AM, Jabbari S (2015) Bacterial fitness shapes the population
755 dynamics of antibiotic-resistant and-susceptible bacteria in a model of combined antibiotic
756 and anti-virulence treatment. Journal of Theoretical Biology 372:1–11
757 Tisue S, Wilensky U (2004) Netlogo: A simple environment for modeling complexity. In:
758 International conference on complex systems, vol 21, pp 16–21
759 WHO (2015) Global action plan on antibiotic resistance. Tech. rep., World Health Organi-
760 sation/UNO
761 Williams GC (1957) Pleiotropy, natural selection, and the evolution of senescence. Evolution
762 11(4):398–411
763 Wolfram S (2002) A New Kind of Science. Wolfram Media, USA

764 Woodford N, Livermore DM (2009) Infections caused by gram-positive bacteria: a review of
765 the global challenge. Journal of Infection 59:S4–S16