



Cite this article: Greenbury SF, Johnston IG, Louis AA, Ahnert SE. 2014 A tractable genotype–phenotype map modelling the self-assembly of protein quaternary structure. *J. R. Soc. Interface* **11**: 20140249. <http://dx.doi.org/10.1098/rsif.2014.0249>

Received: 8 March 2014

Accepted: 18 March 2014

Subject Areas:

computational biology

Keywords:

genotype–phenotype map, self-assembly, robustness, evolvability, polyomino, protein quaternary structure

Author for correspondence:

Sam F. Greenbury

e-mail: sfg31@cam.ac.uk

A tractable genotype–phenotype map modelling the self-assembly of protein quaternary structure

Sam F. Greenbury¹, Iain G. Johnston², Ard A. Louis³ and Sebastian E. Ahnert¹

¹Theory of Condensed Matter Group, Cavendish Laboratory, University of Cambridge, Cambridge, UK

²Department of Mathematics, Imperial College London, London, UK

³Rudolf Peierls Centre for Theoretical Physics, University of Oxford, Oxford, UK

The mapping between biological genotypes and phenotypes is central to the study of biological evolution. Here, we introduce a rich, intuitive and biologically realistic genotype–phenotype (GP) map that serves as a model of self-assembling biological structures, such as protein complexes, and remains computationally and analytically tractable. Our GP map arises naturally from the self-assembly of polyomino structures on a two-dimensional lattice and exhibits a number of properties: *redundancy* (genotypes vastly outnumber phenotypes), *phenotype bias* (genotypic redundancy varies greatly between phenotypes), *genotype component disconnectivity* (phenotypes consist of disconnected mutational networks) and *shape space covering* (most phenotypes can be reached in a small number of mutations). We also show that the mutational robustness of phenotypes scales very roughly logarithmically with phenotype redundancy and is positively correlated with phenotypic evolvability. Although our GP map describes the assembly of disconnected objects, it shares many properties with other popular GP maps for connected units, such as models for RNA secondary structure or the hydrophobic-polar (HP) lattice model for protein tertiary structure. The remarkable fact that these important properties similarly emerge from such different models suggests the possibility that universal features underlie a much wider class of biologically realistic GP maps.

1. Introduction

Evolution is one of the most fundamental principles in biology. While organismal genotypes are becoming accessible owing to rapid advances in sequencing technology, further understanding of the complicated mapping from sequence to phenotype is necessary for a richer understanding of evolutionary dynamics [1–4]. While the terms genotype and phenotype can be flexibly assigned in a biological system, genotypes are generally defined as the genetic material upon which mutations act and phenotypes capture the properties of the organism on which selection can differentiate. As such, the mapping from genotype to phenotype—the *GP map*—links mutations to potentially selectable variation and is therefore of critical importance in understanding evolutionary systems. GP maps also provide a basis for understanding important biological concepts such as mutational robustness and evolvability, which may profoundly affect evolutionary dynamics, and help determine the fundamental topologies of the landscapes upon which evolutionary processes occur [5].

In general, it is intractable to directly model the details of even small parts of a whole organismal GP map, owing to both the very large numbers of genotypes, and a lack of knowledge of all possible phenotypes. Advances have been made in recent years, however, with the use of simplified models. Three particular systems have been modelled with notable success. Firstly, genetic regulatory networks have been approximated using a variety of abstract models, including Boolean networks [6,7]. Despite their simplicity, Boolean networks have demonstrated a remarkable ability to produce biologically realistic results. For example,

they have been shown to reproduce key aspects of the yeast cell cycle [7]. Secondly, RNA secondary structure can be routinely and accurately predicted via a host of different methods [8], as a result of which it has become one of the best-known GP maps [4,9], particularly for the study of evolutionary dynamics [9,10]. Finally, the complex problem of a protein folding into a well-defined tertiary structure has been investigated using various models including the highly simplified HP lattice model, where folds are represented as self-avoiding walks on a lattice, and the full sequence is reduced to binary alphabet (H stands for hydrophobic and P for polar amino acids) [11]. Despite its heavily coarse-grained nature, this model has produced important biological insights [12] and has been shown to accurately model known features of protein tertiary structure [13]. While these models have been successful for specific biological examples, another very important advantage of their tractability is the potential for extracting more general underlying principles of GP maps, thus increasing our understanding of how evolution shapes the natural world [2,4].

In this paper, we extend this family of coarse-grained models for biological structure, exploring a very recently introduced model for tile self-assembly [14,15] which can be viewed as a highly coarse-grained GP map for protein quaternary structure (protein complexes). Understanding the formation of protein complexes is important as demonstrated by the fact that most proteins form complexes in the cell (around 80% in yeast [16]) and the function of these complexes is often strongly linked to their physical form. Protein complexes are formed by the interaction of multiple individual protein chains to form larger structures. The interaction between two chains is predominantly mediated by hydrophobic forces acting to pack together non-polar amino acids to provide fewer energetically unfavourable interactions with water [17]. Invaluable resources for the study of protein complexes are the Protein Data Bank, providing a database containing experimentally known protein quaternary structures [18], and the three-dimensional complex database [16] categorizing PDB structures with a graph representation of the interactions between different subunits and a characterization of the symmetry in a complex.

The relationship between the topology of a protein complex and the individual amino acids that make up its protein chains is highly complex owing to the multiple functionalities encoded in the protein sequence. For example, correct tertiary structure, folding pathways and other inter-protein interactions are all potential requirements for a single protein chain. Given these complexities, a direct and complete GP map of protein quaternary structure is intractable as including all required functionality would be unfeasible. Instead, in the spirit of the highly simplified HP model for protein tertiary structure, we represent the proteins as square tiles on a lattice [14,15]. Interactions between tiles model the protein–protein interactions that lead to self-assembly. In proteins, these patches can be made up of a number of different amino acids; therefore, our model provides a coarse-grained representation of the key interactions between proteins. In the model, a genotype is a sequence of characters describing interactions on the edges of the tiles, which, when combined with a self-assembly process on a two-dimensional square lattice, leads to the formation of phenotypes comprising different square tile building blocks conjoined along interacting edges. These square tile structures are known as polyominoes and thus, we refer to the model as the Polyomino model and the resulting GP map as the Polyomino GP map. They are closely related to a

wider class of lattice tiling models that have a long and important history in mathematics and computer science (which we discuss further in §2).

Despite these great simplifications, and in analogy with the highly schematic HP model, RNA secondary structure models and Boolean network models of GRNs, we expect the Polyomino model to provide insight into the general structure of the full GP map for the formation of protein complexes from folded proteins. In fact, our earlier work [15] has already discussed the evolutionary dynamics of the Polyomino model, demonstrating that it can be used to rationalize the preference of dihedral over cyclic symmetries in homomeric protein tetramers [19,20].

In figure 1, we compare coarse-grained representations for RNA secondary structure, protein tertiary structure and protein complexes. On the left-hand side, the biological representation is shown and on the right, a given model representation. Through this figure, we wish to highlight two points. Firstly, that each of the model systems is a dramatically simplified version of the corresponding biological system and that the Polyomino model is of a similar order of coarse-graining to previous models. And secondly, that the Polyomino model can provide a concise representation of real protein complexes by capturing features such as the symmetry of the subunit arrangements (C_4 in this case).

In contrast to RNA secondary structure or protein tertiary structure, where structure forms through the folding of a connected string of individual entities (nucleotides or amino acids), protein complexes are built by joining separate individual entities (protein chains). Furthermore, while for string-like self-assembly the final structure size is constrained, proteins can form unbounded structures that can be ordered or disordered. In our work here, we focus on the subset of deterministic, finite-sized structures to model protein quaternary structure. But, in principle, the Polyomino model can also capture the more general phenomenology of unbounded and non-deterministic assembly [14,15].

As an example of the link between protein quaternary structure and the structures supported by our model, in figure 2 (top) we show the haemoglobin protein complex (Hb A, [25]) and the mutant that causes sickle cell anaemia (Hb S, [26]), alongside a polyomino representation of the wild-type and mutant phenotype (figure 2, bottom). In both cases, we show that a single mutation to the respective wild-types results in the production of extended unbound structures: Hb S forms long fibre aggregates, while the polyomino case forms unbound tiling of the plane (further details in figure). While the exact geometrical results in this analogy are different, the qualitative behaviour that results from a single patch changing is the same, highlighting the utility of our model as a GP map modelling protein quaternary structure. In this paper, we do not study extended structures, but rather structures that reversibly self-assemble to the same finite structure each time.

Finally, tiling models have been used to study synthetic self-assembling systems [28], and a better understanding of the design space of polyominoes may aid in the design of these artificial systems.

This article proceeds as follows. We first describe in detail the Polyomino model and some of its fundamental properties (§2). We then analyse a wide range of properties of the resultant GP map and compare these properties to those described, for example in [29], for the HP model and the RNA secondary structure map. In particular, we show in §3.1–3.3 that the mapping from sequences to phenotypes

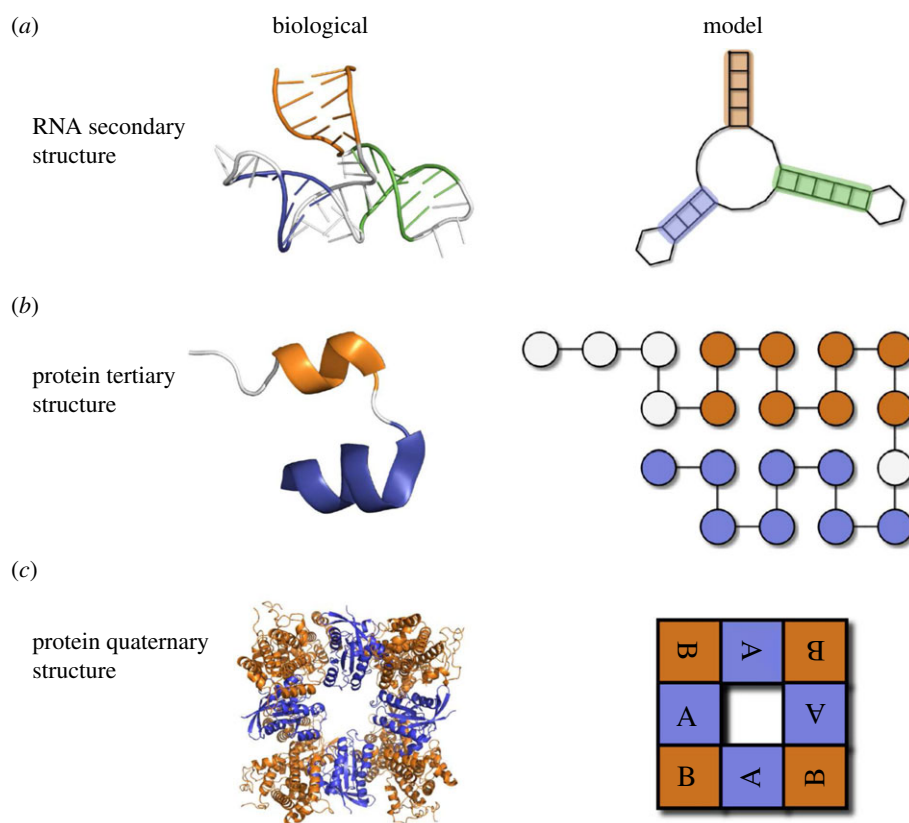


Figure 1. A comparison of how three different biological structures may be represented in a corresponding model system. Row (a) compares a version of the iconic hammerhead ribozyme (PDB reference 1RMN [21]) with its secondary structure representation (showing the bonding pattern of nucleotides) produced by the Vienna RNA package [22] shown alongside. The orange, blue and green colours in each part represent the bonded stems in the structure. Row (b) depicts a cartoon of the tertiary structure of a single chain of length 21 (chain A) from an insulin protein (PDB reference 1APH [23]) which is compared to a schematic HP lattice protein interpretation on the right. The orange and blue colours are used to demonstrate the structural feature of α -helices in the pictures. Finally in row (c), we show a protein complex (PDB reference 1BKD [24]) alongside a polyomino representation. The orange and blue colours represent the different subunits involved in the protein. The ability of the polyomino representation to capture the C_4 symmetry of 1BKD is apparent from the rotation of the labelling on the subunits.

for all three models share the following general properties: *redundancy* (there are many more genotypes than phenotypes) leading to large *neutral sets* (the collection of all genotypes that map to a given phenotype) and *phenotype bias* (some phenotypes are associated with many more genotypes than others). A more fine-grained analysis shows that the neutral sets also exhibit *component disconnectivity* (not all genomes in a neutral set can be linked with single mutational steps). We proceed with a more detailed comparison of the Polyomino and RNA systems, through considering *shape space covering* (most phenotypes can be reached from any other phenotype with just a small number of mutations), before showing the mean mutational robustness of a phenotype (the *phenotypic robustness*) scales very roughly logarithmically with the redundancy of a phenotype, and finally that it is positively correlated with the *evolvability* (defined here as the number of other phenotypes *potentially* accessible from a phenotype), as postulated to hold more generally by Wagner [30]. Finally, in §4 we discuss some implications of the remarkable agreement we find between the structure of our Polyomino GP map and those of the better studied RNA secondary structure and HP maps.

2. The Polyomino self-assembly model and its associated genotype–phenotype map

The process of tiling and its connection with computer science was first developed by Wang [31]. Since then, tiling models

have been shown to be capable of computation and, in particular, Turing-universal computation under the condition that cooperative binding is allowed between tiles, demonstrating the ability of two-dimensional tiling systems to model computational as well as structure-forming processes [32]. Rothmund & Winfree [33] studied the program size complexity necessary to build a structure of a given size. More general considerations of the complexity of tiled structures have since been discussed in [34], with a more biological slant given by Ahnert *et al.* [14] and applications to artificial biological systems discussed by Rothmund *et al.* [28].

Here, rather than focusing on these tiles as potential computing devices or as models for complexity, we explore how they can be used to understand the GP map of a specific biological system, namely the self-assembly of finite-sized protein structures. Nevertheless, we are aware that some of our conclusions may have applications for a wider class of systems.

We now proceed with a more detailed description of the Polyomino model as a GP map.

2.1. Summary of the Polyomino genotype–phenotype map

The genotype is modelled as a character string representation of a set of N_t tiles which make up an *assembly kit*. The edges of each tile in the assembly kit are given a number which represents the *interface type*. Interactions between interface types are defined via an *interface interaction matrix* A_{ij} . In our work

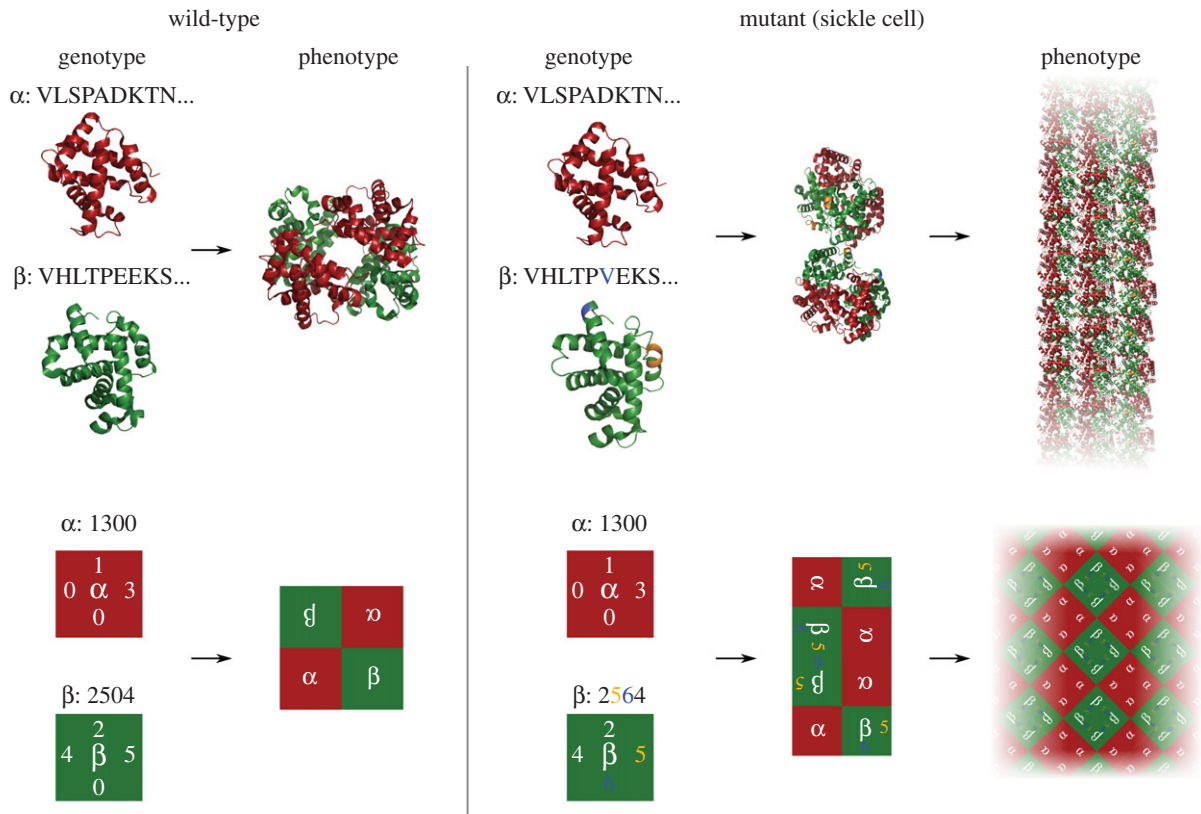


Figure 2. The haemoglobin phenotype polymorphism with a representation in a Polyomino GP map. A Polyomino genotype (bottom left) is constructed to mimic the transition from the haemoglobin wild-type (top left, Hb A, PDB: 1GZX [25]) to the sickle cell mutant (top right, Hb S, PDB: 2HBS [26]). In Hb S, residue 6 of the β chain is mutated from glutamic acid (E) to valine (V), shown in blue. The mutation leads to a hydrophobic patch being created at this position which allows binding to a pocket between residues 85 and 88 (shaded orange) on the second β chain. A single binding is depicted in the middle stage of the top right. The binding allows for large-scale polymerization of Hb S in a double strand form, which further aggregates to form extended fibres [27] (a cartoon depiction of this in the top right). In the polyomino example (bottom), we assume that odd-even numbered interfaces may interact (full details of the model are explained in S2). In the bottom right, a single change (from 0 to a 6 at the third residue of the β tile) has occurred resulting in a new interaction with the second position also on the β tile ($5 \leftrightarrow 6$). The new interaction leads the self-assembly process to produce an unbound aggregate, tiling the two-dimensional plane. This provides an example of the potential for the Polyomino GP map to capture basic features of protein assembly, including extended structures.

here, we only consider one such interaction matrix type, with a total of c interface types. There are no self-interacting interface types and interface types interact in unique pairs ($1 \leftrightarrow 2, 3 \leftrightarrow 4, \dots$), with the only fully neutral types being 0 and $c - 1$. Defining interface interactions in this way allows the single parameter c to control the number of potential unique bond types. As such, the two parameters that describe a given parametrization of the Polyomino GP map are the tile number N_t and the total number of interface types c , allowing a particular Polyomino GP map to be labelled as $S_{N_t, c}$.

Below we discuss the genotype, phenotype and map used in the Polyomino GP map.

2.2. Genotype

The genotype for the set of building blocks in the assembly kit is written as a character string of length L in base K . The base we use in this paper is the total number of interface types c used in a given GP map. This allows each base to mutate to any other base at each site. The procedure of converting a genotype string into the assembly kit is part of the *map*.

2.3. Phenotype

In the context of the self-assembly mapping, there are several ways of classifying a polyomino structure. These may include

criteria based on its overall shape and the individual tile types making up the final polyomino structure, as well as individual tile orientations. These different possibilities are discussed in [15].

Here, we will classify the phenotype according to the overall shape of the polyomino independent of origin translation and C_4 (90°) rotations. Note that chiral counterparts of polyominoes represent distinct phenotypes.

2.4. Map

The map has two stages: conversion of the genotype into the assembly kit, followed by assembly of a polyomino from an assembly process involving the assembly kit and the interface interaction matrix. A diagram representing this process is shown in figure 3.

2.4.1. Genotype to assembly kit

The characters of the genotype are read from left to right along the string. Characters are assigned to the next blank edge of a square tile in the assembly kit. The edges are taken in clockwise order (from the top side) with all edges being written before moving on to a new block. The total number of tiles, in terms of the genotype string length, can be expressed as $N_t = L/4$.

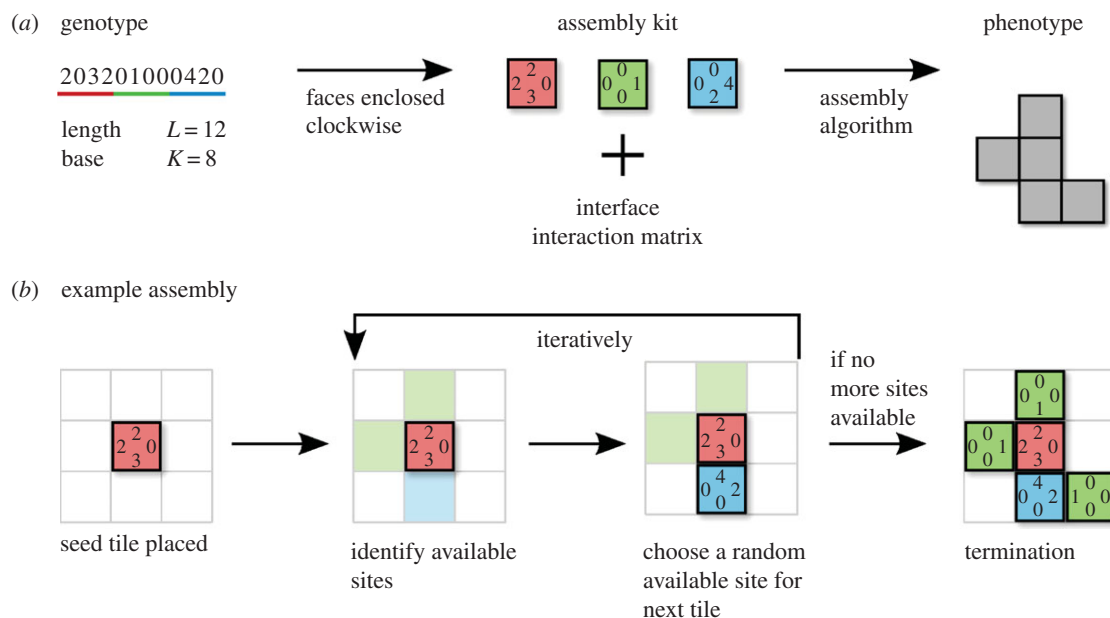


Figure 3. The *Polyomino GP map*. (a) The translation from an example genotype to phenotype is depicted, with the intermediate conversion into an assembly kit shown. In the genotype, the characters are underlined with the colour of the block they are assigned to. The interface interaction matrix is not shown explicitly, but throughout our work we assume the convention of interface types interacting in pairs ($1 \leftrightarrow 2$, $3 \leftrightarrow 4 \dots$), with 0 and $c - 1$ being neutral. The interaction matrix together with the assembly kit are passed to the assembly algorithm, which is used to produce the phenotype. (b) An example of the assembly process is shown, which proceeds in discrete time steps. The first tile in the assembly kit is used to seed the assembly. Further time steps follow, with the identification of *available sites* (depicted in light green and blue in the second picture from left), followed by the random choice of an available site and placement of the corresponding tile (a blue tile is placed on the blue available site in this example). Assembly is terminated when there are either no more available sites (as above), or if the structure is larger than D_{\max} (number of blocks) in width or height after which we assume the structure will no longer terminate but grow indefinitely.

2.4.2. Assembly kit to phenotype

The assembly of the two-dimensional polyomino takes place on a square lattice where individual tiles from the assembly kit are placed. The interface types on the edges of tiles can form an attractive interaction as determined by the binary interface interaction matrix A_{ij} , with 1 denoting attraction and 0 neutrality. A bond may form between tiles if two adjacent (interacting) edges have interface types which attract, as defined by the interaction matrix.

The assembly process is initialized by placing (seeding) a single tile on the lattice. We will consider only GP maps, in which the seed tile corresponds to the first tile described in the genotype. A different protocol where any tile may be used to seed the assembly is also possible and does not significantly affect the results presented here (see appendix A). The assembly then proceeds as follows:

- (1) *Available sites* on the lattice are identified. These are places on the lattice where a tile may be placed in a particular orientation such that it will form a bond to an adjacent tile that has already been placed. In the assembly algorithm, a list is kept of the position, tile type and orientation for each possible tile placement. A potentially infinite number of tiles, in equal proportions, are available.
- (2) A random available site on the lattice is chosen.
- (3) The chosen site is filled with the associated tile and with the corresponding orientation.

These steps are repeated until either:

- There are no available sites for bonding.
- The structure grows beyond a certain width or height D_{\max} , which is taken as a proxy for unbounded assembly, so that the resulting phenotype is described as *unbound*.

We set $D_{\max} = 16$ here, but our results are not sensitive to this cut-off as, for the polyomino systems we study here, there are no bounded structures larger than this.

At this point, the assembly process is terminated and the structure produced is recorded. No re-seeding takes place as we only consider the assembly of a single connected structure. To test whether the structure is *deterministic*, the assembly process is repeated k times, with each C_4 rotation of the final structure checked against the recorded structure. If there are any differences between any of the k assemblies, the phenotype is classified as *non-deterministic*. Phenotypes that are classified as unbound or non-deterministic structures are represented by a single phenotype, which we refer to as the undetermined (*UND*) phenotype and is assumed not to be biologically relevant in this context. A more detailed discussion of the classification of polyomino structures is given in [15].

2.5. Methods for RNA and hydrophobic-polar lattice protein models

RNA secondary structure was modelled with the Vienna package [22] with the default parameters. Single isolated base pairs were permitted (see [35] for further discussion of this choice). The HP lattice model data were taken from the enumeration performed by Irbäck & Troein [36] over the set of non-compact folds for $L = 25$ and reproduced results discussed in their original work as well as that of Ferrada & Wagner [29].

3. Properties of the Polyomino genotype–phenotype map

In this section, we analyse the Polyomino GP map by making measurements of redundancy, phenotype bias and component

Table 1. Redundancy, phenotype bias and components in Polyomino, RNA and HP GP maps. Comparing the number of phenotypes (N_P) to the number of genotypes (N_G) for each GP map highlights large-scale redundancy present. Phenotype bias is demonstrated in each map with the measure of the fraction of phenotypes that covers 95% of the genotypes (N_{cov} is the number of phenotypes that covers the 95% of genotypes). In all cases the fraction of phenotypes is significantly smaller than the fraction of genotypes being covered, indicating the presence of a strong phenotype bias. The final column is the total number of genotype components (N_C) in each GP map. In all cases (non-computable values left out), the number of components is larger than the number of phenotypes, indicating phenotypes tend to be spread out over multiple disconnected components. RNA data for $L = 12$ were computed from the Vienna package [22] and taken from [41] for $L = 15, 20$. The value of N_P for the Polyomino $S_{4,16}$ GP map is approximate as it is found from large-scale sampling of the GP map over multiple runs of the algorithm presented in [42]. All other Polyomino results were found by exhaustive enumeration. The HP results were calculated from the data made available by Irbäck & Troein [36]. Non-deterministic phenotypes and the trivial structure in RNA are excluded from the statistics.

GP map	N_G	N_P	N_{cov}/N_P (%)	N_C
polyomino $S_{2,8}$	1.7×10^7	13	54	1347
polyomino $S_{3,8}$	6.9×10^{10}	147	16	—
polyomino $S_{4,16}$	1.8×10^{19}	~ 2237	3	—
RNA, $L = 12$	1.7×10^7	57	47	645
RNA, $L = 15$	1.1×10^9	431	23	12 526
RNA, $L = 20$	1.1×10^{12}	11 218	10	—
HP, $L = 25$	3.4×10^7	107 336	68	148 254

numbers, before moving on to analyse the properties of shape space covering, robustness and the relationship between robustness and evolvability. For each measurement, comparisons are made with the RNA secondary structure model and, for the first three of these properties, a HP lattice protein model.

3.1. Redundancy

It is now well established that many different sequences can generate similar protein or RNA phenotypes [37]. This large-scale redundancy is the basis, for example, of the molecular clock hypothesis, which is widely used for inferring phylogenetic relationships [38]. It is therefore not surprising that such redundancy (also known in the literature on GP maps as degeneracy/designability) has been observed in model RNA [9], HP lattice protein [13] and genetic regulatory [39] GP maps, as well as in more general model systems such as signalling networks [40]. As such, redundancy is expected to be a typical feature of GP maps.

In table 1, we show the number of genotypes (N_G) and phenotypes (N_P) for different polyomino, RNA and HP GP

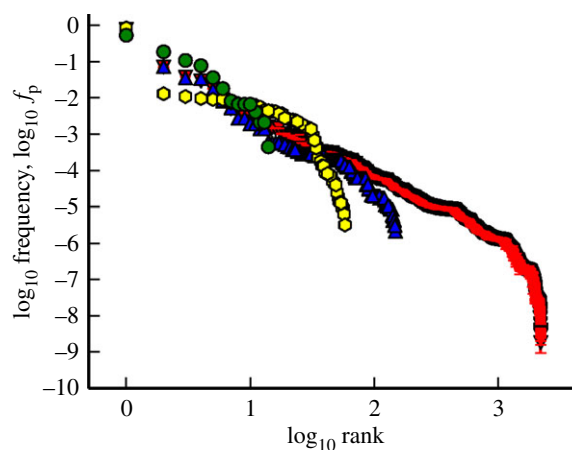


Figure 4. Phenotype bias in GP maps. The plot displays three Polyomino GP maps with increasing tile numbers $S_{2,8}$ (green), $S_{3,8}$ (blue) and $S_{4,8}$ (red) alongside the RNA $L = 12$ GP map (yellow). Note the log–log scale. The frequency of a given phenotype (structure) is plotted against the rank of its frequency within the map. In all three Polyomino maps, we see an approximately linear trend between \log -frequency and \log -rank. The error bars for the Polyomino $S_{4,16}$ map are the associated standard error on the mean for the estimates of frequencies calculated using the estimation algorithm for large maps developed by Jörg *et al.* [42].

maps. As expected, all three models show significant redundancy. The Polyomino model displays large-scale redundancy shown by the vastly fewer phenotypes in comparison to genotypes for each of the GP maps presented. As can be seen in table 1, this is of a similar order to the RNA GP maps. For example, for RNA $L = 12$ and Polyomino $S_{2,8}$ (which both have 1.7×10^7 genotypes), there are 57 phenotypes for the former and 13 for the latter.

3.2. Phenotype bias

On top of the measure of redundancy in GP maps, it is also possible to consider *phenotype bias*, the idea that some phenotypes are much more redundant than others. Examples of bias can again be seen in RNA [9,35] and the HP model [12,13,29]. We demonstrate bias in the GP maps in two ways. In table 1, we present the fraction of phenotypes (N_{cov}/N_P , where N_{cov} is the number of phenotypes) that cover 95% of the genotype space. This statistic shows that in all three GP maps (RNA, HP and Polyomino), of the deterministic/non-trivial space, a smaller fraction of phenotypes is required to cover a given amount of the genotype space, indicating a bias in the number of genotypes assigned to each phenotype. In figure 4, we plot the frequency (fractional coverage of genotype space) of each phenotype against its frequency rank in three different Polyomino GP maps, as well as for the RNA $L = 12$ GP map. In each of the Polyomino GP maps, we see an approximately linear trend between \log -frequency and \log -rank, although there are not enough data for this to be considered a power law. The data from the RNA GP map also show a distinct bias, although the relationship is more complicated than the more apparent linear form seen in the Polyomino GP maps. Taken together these statistics show a significant amount of phenotype bias in the Polyomino GP map along with the RNA and HP GP maps.

3.3. Component disconnectivity

Further to the consideration of redundancy and bias of phenotypes in genotype space, we can ask how many components a

given phenotype forms in a particular GP map. In graph theory, a component is a set of nodes that are connected, and thus when applied to genotype spaces, is a set of connected genotypes with the same phenotype. Here, we consider only point mutations. In RNA, simple biophysical considerations show that two point mutations are needed to turn a purine–pyrimidine bond into a pyrimidine–purine bond. This neutral reciprocal sign epistasis leads to many individual components [43,44]. We expect similar behaviour in polyominoes, where there may be several ways of encoding a given phenotype that are not connected through neutral point mutations owing to there being multiple interface types.

In the final column of table 1, we show the number of genotype components (N_C) for four different GP maps: Polyomino $S_{2,8}$, RNA $L = 12$, RNA $L = 15$ and HP $L = 25$ (Polyomino $S_{3,8}$ and $S_{4,16}$ and RNA $L = 20$ GP maps were unobtainable owing to computational expense). In all four, we see a substantially larger number of components than phenotypes, indicating that phenotypes tend to form disconnected components. Although many components may be small, it is still the case that whole-phenotype properties need to be considered carefully in the context of evolutionary dynamics, as it is not typically possible to access every genotype of a given phenotype through neutral point mutations.

3.4. Shape space covering

GP maps typically exist in high-dimensional genotype (sequence) spaces, a property with counterintuitive consequences and a topic of current biological research for GP maps [45,46]. For example, for an alphabet of K letters (e.g. $K = 4$ for RNA or $K = 20$ for proteins) the number of genotypes scales exponentially as K^L , while the number of mutations needed to reach any two genotypes is at most L , and so scales linearly. In practice, considerably fewer than L mutations are needed to connect any two phenotypes. This property has been studied most extensively for RNA GP maps, and in that context Schuster *et al.* [9], borrowed the term shape space covering from its original use in immunology [47]. The HP model has also been considered with respect to shape space covering [13,29] with seemingly less rapid space covering.

We explore shape space covering in a similar way to Ferrada & Wagner [29], who compared RNA and HP models, building on work of earlier investigators. Shape space covering is measured by counting the average fraction of phenotypes found when applying a given number of independent mutations (the radius, n) to a genotype in the GP map. This process involves picking a sample of genotypes and then measuring the number of phenotypes found in the ball of genotypes n -mutations around each of them. Ferrada & Wagner [29] found that in both the RNA secondary structure and HP lattice protein GP maps, the fraction of phenotypes discovered increases in a roughly sigmoidal fashion with the increasing number of independent mutations applied to the source genotype and at a greater rate for RNA than the HP GP map.

In our work, we measure shape space covering in a similar manner: we picked a sample of 1000 genotypes uniformly across all phenotypes and measured the phenotype of genotypes at each given radius (n). In figure 5, we plot the average fraction of phenotypes found for the Polyomino $S_{2,8}$, $S_{3,8}$ and RNA $L = 12$ GP maps. The general behaviour of both the

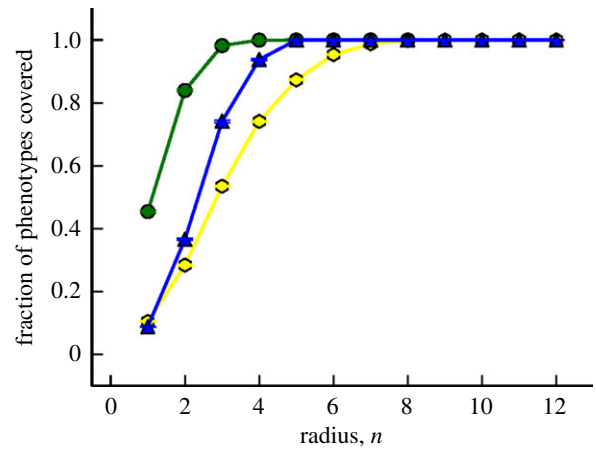


Figure 5. Shape space covering in Polyomino $S_{2,8}$ (green), $S_{3,8}$ (blue) and RNA $L = 12$ (yellow) GP maps. A random sample of 100 genotypes for each phenotype is taken, and the fraction of phenotypes discovered in a ball of radius n is measured. This is then averaged over all sample phenotypes to give an approximate phenotype space covering for a given GP map. Both Polyomino and RNA GP maps exhibit similar behaviour—shape space is almost completely covered in only a few mutations highlighting that in highly redundant spaces, most phenotypes can be reached in only a few mutations. Error bars are the standard error on the mean. The expected asymptotic value of unity is used for the Polyomino $S_{3,8}$ GP map for $n > 4$ owing to computational unfeasibility.

Polyomino and RNA GP maps is similar. Phenotypes are almost completely covered within half the sequence length, and at a slightly higher rate in the Polyomino GP maps. This provides clear evidence to suggest that the Polyomino GP map has shape space covering—most phenotypes can be found within a ball containing many fewer genotypes than the entire space.

Ferrada & Wagner [29] found that shape space covering was dependent on the exact group of phenotypes under consideration—exploring around genotype networks of the most frequent phenotypes, resulted in a more rapid shape space covering than from random genotypes. In appendix A, we provide additional discussion of phenotype parametrizations in the Polyomino GP map. From the general lack of variation in Polyomino GP map properties such as redundancy and phenotype bias discussed there, we would expect shape space covering to exhibit similar behaviour for other conceivable deterministic, bound phenotype parametrizations.

3.5. Robustness

Biological systems need to be robust and consistently produce the same phenotype in response to environmental perturbations or to genetic mutations [2]. Here, we consider only the phenotypic effects of mutations to the genotype. Mutational robustness describes the invariance of a phenotype as a result of mutations to the genotype. We focus here on single point mutations such that the fraction of 1-mutants that have the same phenotype as the original genotype under examination is defined as the *genotypic robustness*. For a genotype g , with phenotype p , the genotype robustness can thus be expressed as follows:

$$\rho_g = \frac{n_{p,g}}{(K-1)L} \quad (3.1)$$

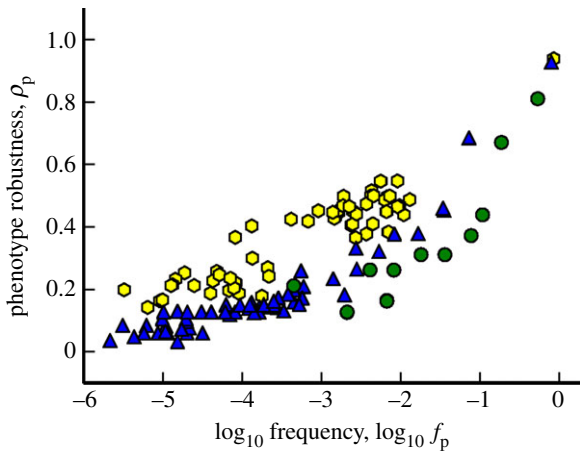


Figure 6. Robustness in GP maps. Phenotypic robustness is plotted against frequency (log scale) for the Polyomino $S_{2,8}$ (green), $S_{3,8}$ (blue) and RNA $L = 12$ (yellow) GP maps. The top three points in the right-hand side of the plot are the UND/trivial structure in the respective Polyomino and RNA systems. In the RNA case, robustness scales linearly with log-frequency, with the exact coefficient of scaling being phenotype dependent [43]. While the trend is very roughly linear for the Polyomino GP maps, there is still a strong positive scaling relationship, demonstrating that polyomino phenotypes exhibit phenotypic robustness in a similar manner to the RNA GP map.

where ρ_g is the genotypic robustness of g , $n_{p,g}$ is the number of 1-mutant neighbours of g with phenotype p , K is the base and L is the sequence length (there are a total of $(K - 1)L$ 1-mutants for any genotype). The robustness of the phenotype can then be considered as the average of this quantity over all genotypes with phenotype p , resulting in

$$\rho_p = \frac{1}{|\mathcal{P}|} \sum_{g \in \mathcal{P}} \rho_g \quad (3.2)$$

where ρ_p is the *phenotypic robustness* and \mathcal{P} is the set of genotypes with phenotype p .

In figure 6, we plot the phenotypic robustness ρ_p against the frequency of each phenotype f_p for the two Polyomino GP maps $S_{2,8}$ and $S_{3,8}$ and the RNA $L = 12$ GP map. Corroborating previous results [41,43], it can be seen that the RNA phenotypic robustness scales logarithmically with phenotype frequency—the exact coefficient of scaling being phenotype dependent [43]. This linear trend with log-frequency is very approximate for the Polyomino GP maps, although the robustness can still be seen to strongly scale with phenotype frequency. Both Polyomino GP maps exhibit a slightly smaller phenotypic robustness at a given frequency when compared with the RNA GP map. Nonetheless, despite these two observations, the clear indication here is that Polyomino GP maps have a phenotypic robustness that scales similar to phenotypes in the RNA GP map.

3.6. Robustness versus evolvability

Robustness and evolvability are two evolutionary properties that have received much recent attention in the literature [2,30,48–51]. As discussed in §3.5, robustness is the ability of an organism to maintain its phenotype if its genotype is mutated. Evolvability on the other hand is considered as the capacity for producing phenotypic variation [1,30]. One might expect that for an individual to be robust it

would have to compromise its ability to produce variation (to be evolvable). Using the RNA GP map, Wagner [30] demonstrated that this expected trade-off exists for individual genotypes, but that the two properties are in fact positively correlated at the level of phenotypes. A possible geometric explanation is that phenotypes can become more mutationally robust as a result of increased redundancy (more neutral neighbours). This in turn leads to connected components spreading out further across genotype space and possessing a greater ‘surface’, thus allowing a greater number of phenotypes to be accessible through neutral mutations.

Wagner’s argument is really about the static structure of genotype space and does not account for any dynamical evolutionary effects. Other authors have considered further static properties, including a different notion of phenotypic evolvability (diversity evolvability) [52] defined as being the probability that two non-neutral mutations lead to different phenotypes. Such a definition of evolvability was found not to be correlated with robustness in the RNA GP map. A dynamical study by Draghi *et al.* [53] demonstrated that evolvability could depend non-monotonically on robustness. In our work here, we simply consider whether the static properties of the Polyomino GP map behave in a similar manner to those of the RNA GP map as demonstrated by Wagner [30] without here commenting on the wider debate of how robustness and evolvability relate.

In the left-hand plots of figure 7, we show fractional evolvability versus fractional robustness, for genotypes and phenotypes in both the RNA $L = 12$ and Polyomino $S_{3,8}$ GP maps. The genotype plot is a binned version of 100 randomly sampled genotypes for each phenotype (apart from the non-deterministic/trivial structure) in each GP map. The fraction of all phenotypes that can be produced in any 1-mutation to each sampled genotype (genotypic evolvability, ϵ_g) is plotted against the fraction of those 1-mutations that produce the same phenotype as that of the genotype being tested (genotypic robustness, equation (3.1)). For both polyominoes and RNA, we see a significant negative correlation. This expected negative correlation simply represents the trade-off between genotypic evolvability and robustness at the level of the individual genotype.

In the right-hand plots of figure 7, we plot phenotypic evolvability against phenotypic robustness. The phenotypic evolvability ϵ_p of phenotype p is defined here as the fraction of all phenotypes that can be reached in a single mutation from any genotype with phenotype p . We also refer to this as the *potential* evolvability because it represents the potential number of phenotypes that could be reached. Whether they can be reached depends, of course, on details of the evolutionary dynamics. For example, if only single mutations are available then ϵ_p should really be defined with respect to the relevant component [44]. Phenotypic robustness is defined in the same way as in §3.5, as the average genotypic robustness over all genotypes with phenotype p (equation (3.1)). In the plots, we see strong positive correlations for both the RNA $L = 12$ and the Polyomino $S_{3,8}$ GP maps, as expected. In other words, for both maps, phenotypes with a greater redundancy can be generated by a larger number of genotypes and are therefore, on average, more likely to be mutationally robust. Furthermore, such phenotypes are also likely to have a larger set of genotypes, and therefore a greater diversity of other phenotypes potentially

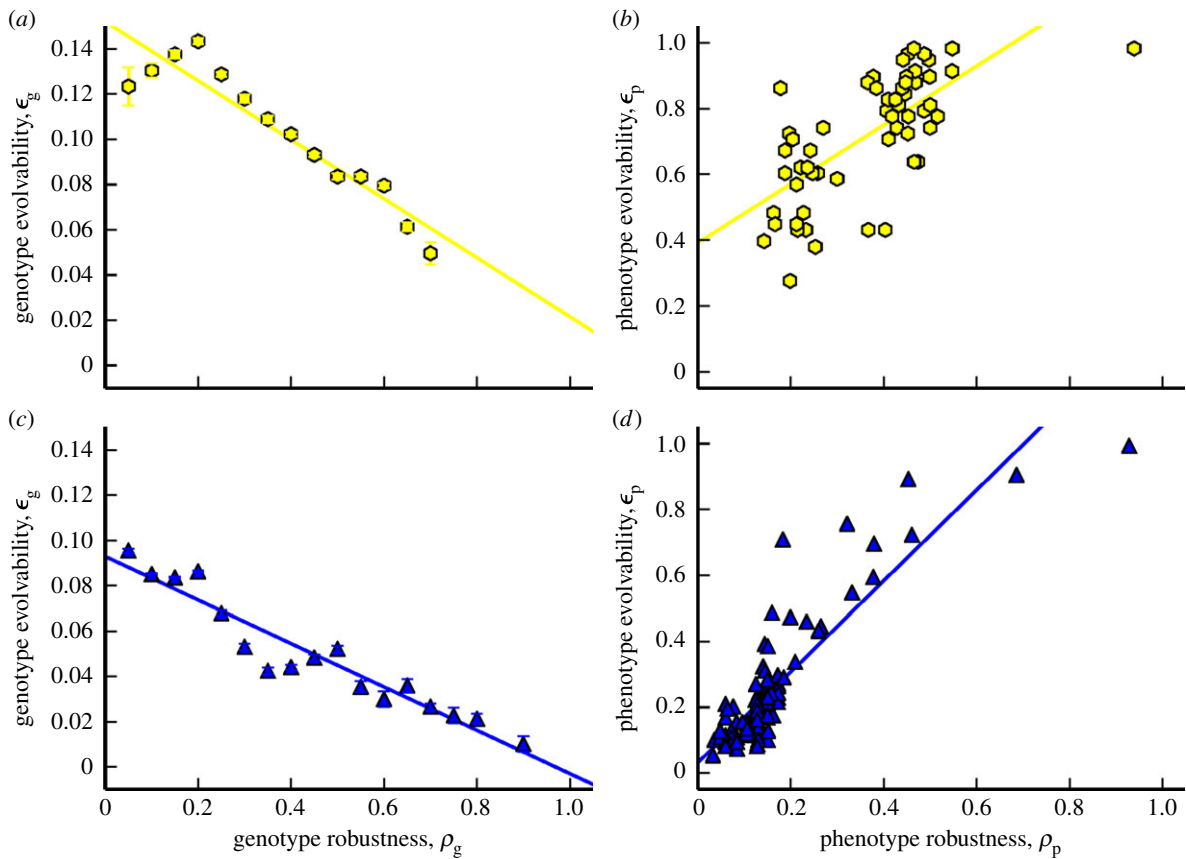


Figure 7. Robustness and evolvability in the RNA and Polyomino GP map. In columns (a,c), we show a significant negative correlation between genotypic robustness and genotypic evolvability, ϵ_g (RNA (yellow, top): p -value = 4.9×10^{-7} , $r^2 = 0.89$ and Polyomino (blue, bottom): p -value = 2.4×10^{-9} , $r^2 = 0.91$). The error bars are the standard error on the mean genotypic evolvability for each genotypic robustness bin. The more robust a genotype is, the fewer phenotypes lie in its 1-mutant neighbourhood. In columns (b,d), there is a significant positive correlation between phenotypic robustness and phenotypic evolvability, ϵ_p (RNA: p -value = 1.5×10^{-9} , $r^2 = 0.48$ and Polyomino: p -value < 2.2×10^{-16} , $r^2 = 0.74$). The far right-hand phenotypes are the trivial structure and UND phenotype in RNA and Polyomino systems, respectively. Both these results are in correspondence with [30].

accessible from this set. Thus, phenotypic robustness and potential evolvability are positively correlated.

4. Discussion

In this paper, we have explored the properties of a GP map for biological self-assembly of the kind exhibited by protein quaternary structure based on a recently introduced Polyomino model for tile assembly [14,15]. We compared its properties to models of RNA secondary structure and the HP model for protein tertiary structure. As is the case for these two well-studied GP maps, we argue that even though our Polyomino model is highly schematic and thus misses many details of protein quaternary structure, it may nevertheless provide important biological insight into the structure of the design space for protein complexes.

Despite the great complexity in potential phenotypes, the polyomino model remains tractable as demonstrated in this paper by the ability to perform a wide variety of useful measurements on the GP map. Our main results are: firstly, the Polyomino model exhibits large-scale redundancy, a strong phenotype bias and the presence of disconnected genotype components across several parametrizations of the Polyomino GP map (§3.1–3.3). Secondly, that shape space may be covered in only a fraction of mutations—that is, all phenotypes are a significantly smaller number of mutations away from each other than the total sequence length

(§3.4). Thirdly, phenotypic robustness scales very roughly with the logarithm of the phenotype frequency (§3.5). And finally, genotypic robustness and genotypic evolvability are negatively correlated, while phenotypic robustness and phenotypic (potential) evolvability are positively correlated (§3.6).

The Polyomino model describes the self-assembly of disconnected units (proteins) into finite-sized structures (protein clusters) that can vary in size. By contrast, for RNA secondary structure and the HP model for protein tertiary structure, strings of connected units (nucleotides or amino acids) assemble into shapes of a fixed size. Given the substantially different class of the phenotypes in our model, it is remarkable that the measured properties of these GP maps turn out to be so similar. This begs the question of whether what we observe is in fact a more general property of self-assembling systems, or even broader, whether a wider class of GP maps will share these properties. This question can be sharpened by looking at the different properties separately. Redundancy should be widely shared across GP maps. Phenotype bias has been observed in models for gene-regulatory networks [54] and developmental networks [55]. Could it be a more general property of GP maps? Disconnected components have also been observed in a Boolean threshold model for gene regulation [56]. To the best of our knowledge, shape space covering has not been studied for other GP maps, but general considerations based on the high-dimensionality of genetic space suggest that something like this may be

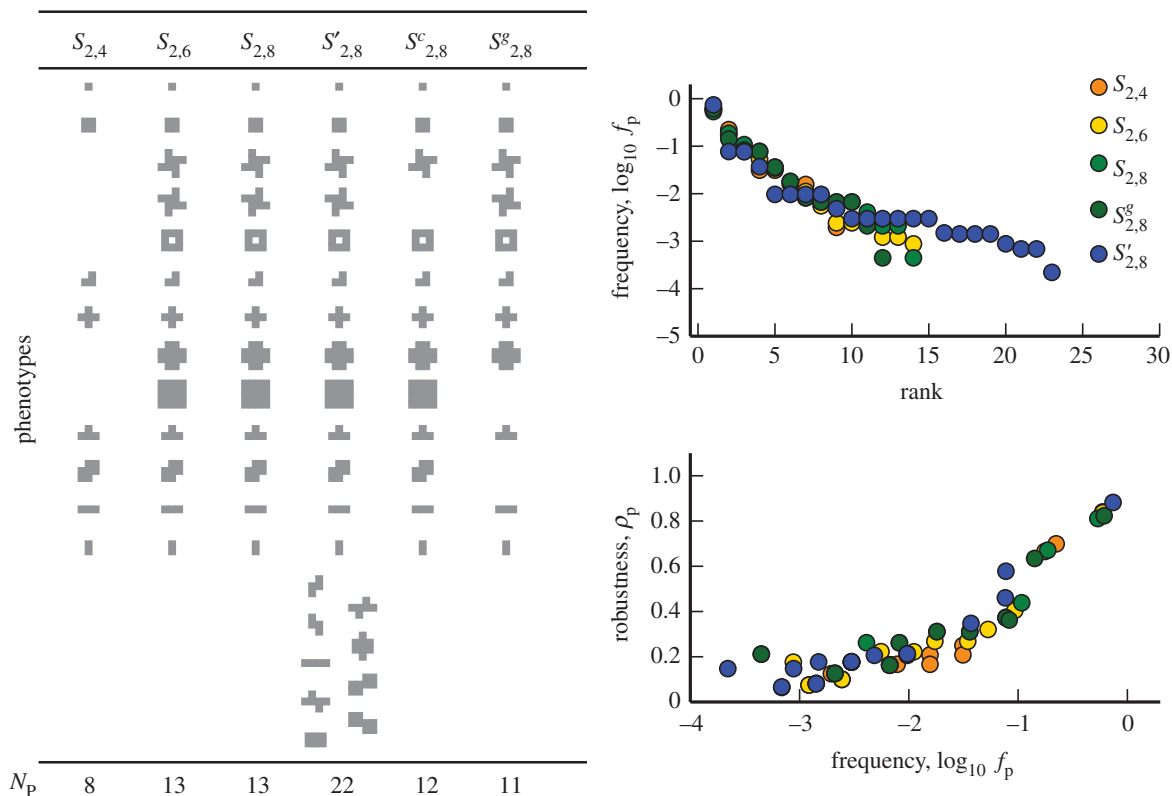


Figure 8. The variation in possible structures and quantitative properties for different Polyomino GP map parametrizations. We tested $S_{2,4}$, $S_{2,6}$, $S_{2,8}$, $S'_{2,8}$ (same as $S_{2,8}$ except the interface type 7 may interact with itself instead of being neutral, $7 \leftrightarrow 7$), $S^c_{2,8}$ (distinct chiral counterparts are treated as a single phenotype) and $S^g_{2,8}$ (any tile may seed the assembly process). In particular, as number of interface types is increased, N_p increases up to saturation where no more new phenotypes may be constructed. The addition of a self-interacting interface type ($S'_{2,8}$) allows a further increase in the number of phenotypes. Defining phenotype identity for chiral counterparts only removes a single phenotype in $S^c_{2,8}$. Allowing the assembly process to be seeded by any tile in the assembly kit in $S^g_{2,8}$ reduces N_p by two phenotypes as further specificity is required in the assembly kit to ensure deterministic assembly. The overall properties of phenotype bias and robustness scaling with log-frequency are qualitatively invariant.

more widely relevant [57]. Finally, the correlation between phenotypic robustness ρ_p and potential evolvability ϵ_p is deeply connected to the geometry of neutral sets and so is likely to be a much more general property of GP maps. How this correlation plays out for realistic evolutionary dynamics is, of course, a much more complicated question.

We are hopeful that more complete answers will be derived through further analysis and comparisons of different model GP maps in a similar manner to the work here or, for example, in [29] or [2]. An important related question is whether model GP maps for parts of systems can be combined to achieve a more complete understanding of the evolution of phenotypic traits in the full organismal GP maps [37].

The Polyomino model can easily be adapted to study unbounded assembly or the assembly of synthetically produced objects like DNA tiles or patchy colloids [28]. Thus, the perspective gained from viewing polyominoes as a GP map may also shed light on the artificial design process for these systems.

Finally, although we introduce the Polyomino model as a coarse-grained model for protein quaternary structure, it is clear that the model is not capable of modelling particular intricacies of some individual proteins. For example, F1-ATPase is a protein whose subunit structure could not be accurately represented with a square tile model. We argue that the Polyomino model nevertheless provides biological insight into questions about the global nature of the GP map that would not be computationally accessible using more complex models with greater biological detail.

However, further work is needed to assess how well this new GP map performs in this respect.

Appendix A. Methods: interface types, assembly and phenotype

In this section, we discuss how variation in number of interface types, phenotype definition and assembly process can affect properties of the GP map produced. We performed GP map analysis of $S_{2,4}$, $S_{2,6}$, $S_{2,8}$, $S'_{2,8}$ (same as $S_{2,8}$ except the interface type 7 may interact with itself), $S^c_{2,8}$ (chiral counterparts are regarded as the same phenotype) and $S^g_{2,8}$ (any tile may seed the assembly process) and in figure 8 show how redundancy, phenotype bias and robustness vary.

A.1. Interface types

First, we discuss the impact of varying interface type numbers, c . As the number of interface types is increased ($c = 4, 6, 8$), we see that the number of phenotypes (N_p) increases from 8 before saturating at 13 ($S_{2,c \geq 6}$). For the adjacent integer interface interactions, we have defined for these GP maps, an increase in N_p before saturating for a certain value of c is expected to be a general behaviour due to the following rationale: there are $4N_t$ edges in the assembly kit and thus for $c \geq 4N_t$, every edge may be given a unique interface, resulting in no further new interactions being able to be introduced to the assembly kit by further increasing c . In general,

we expect N_p to saturate for smaller values of c due to the deterministic assembly criterion.

We also consider the effect of introducing a self-interacting interface type ($7 \leftrightarrow 7$) to the standard $S_{2,8}$ GP map. In figure 8, we see what nine more phenotypes are now produced, although the overall number of phenotypes remains of the same order of magnitude.

A.2. Phenotype definition

As discussed in §2, there are several ways of defining a polyomino phenotype and, as well as being biologically more or less relevant, can possibly affect the GP map properties discussed in this paper. The approach taken in our work has focused on treating the self-assembly process in two dimensions, leading to tiles with distinct chirality being used in the assembly process, as well as distinct chirality being applied to the overall phenotypes. In this section, we discuss one particular alternative GP map where the chiral counterparts of a given polyomino are treated as the same phenotype with all other aspects of the GP map definition remaining the same ($S_{N_i,c}^c$).

In figure 8, we depict the effect on phenotype numbers in the column $S_{2,8}^c$ with N_p only reduced from 13 to 12 (a reduction of N_p by 7.14%) as there is only a single phenotype with a distinct chiral counterpart. The change in N_p becomes more pronounced for larger systems—in $S_{3,8}^c$ a 30.41% reduction and in $S_{4,16}^c$ 43.03% reduction. We note that the change in N_p can at most be 50% when every phenotype has a chiral counterpart, and owing to the symmetries of the way assembly kits are constructed in the Polyomino GP map, chiral counterparts would have precisely the same frequencies, component numbers and robustness as each other. With these considerations, we can see that chirality is a property that will not dramatically affect any of the overall GP map properties discussed here, but may be important when considering particular biological situations.

In some GP maps, changes in phenotype definition can have a more dramatic effect. For example, in the HP lattice model, considering phenotypes only from the subset of maximally compact structures greatly affects the number of potential phenotypes. For those of length $L = 25$ the size of the non-compact set (that may be produced by at least a single sequence from the 2^{25}) is 107 336 [36], in comparison to the compact set of 549 [58], a difference of several orders of magnitude. The strong effect here is due to genotypes that have multiple minimum energy structures in the non-compact case being reduced to a unique structure in the compact case.

An analogous definition for the Polyomino GP maps would be a phenotype reclassification that resulted in members from the set of non-deterministic or unbound phenotypes giving rise to large numbers of new meaningful phenotypes. Such ideas are certainly conceivable—unbound behaviour may be deterministic and thus may be categorized as individual phenotypes in their own right by a description of a repeating unit in the structure. However, detailed categorization of such formulations may be regarded as less biologically relevant—most biological functions performed by protein complexes rely on deterministic bound assemblies. However, we note that the Polyomino GP map may be parametrized to consider such cases, as shown in the haemoglobin example in figure 2.

A.3. General tile seeding

A final variation of the GP map considered is the seeding of the assembly process itself. In the $S_{N_i,c}$ GP maps tested here, it is always the first tile in the assembly kit which is used to begin the assembly process. In figure 8, we demonstrate the effect on N_p , phenotype bias and robustness, in allowing any randomly chosen tile to seed the assembly (*general tile seeding*, first discussed by Ahnert *et al.* [14]). We denote this assembly protocol as $S_{N_i,c}^g$. As can be seen in figure 8, $S_{2,8}^g$ has $N_p = 11$, in comparison to $S_{2,8}$ with $N_p = 13$. $S_{3,8}^g$ possesses $N_p = 119$ phenotypes in comparison to $N_p = 148$ for $S_{3,8}$. The reduction in number of phenotypes is due to the extra information required to ensure the deterministic assembly. While the biological relevance of the single and general tile seeding is open to discussion, the GP map properties can be seen to be qualitatively similar.

Appendix B. Computational details for Polyomino genotype – phenotype maps

Three main Polyomino GP maps are tested in this work: $S_{2,8}$, $S_{3,8}$ and $S_{4,16}$. In enumerating the GP maps, $S_{2,8}$ and $S_{3,8}$ were approached exhaustively, while the statistics of $S_{4,16}$ could only be estimated through sampling because of the greater number of genotypes.

For GP maps enumerated exhaustively, the enumeration is performed in two steps: first the space is sampled randomly with large values for k (repeat assemblies) and D_{\max} to check for the largest structures that may be produced deterministically. As mentioned in §2, $D_{\max} = 16$ was found to be sufficient for all GP maps studied here. With regard to determinism checking, for low values of k some genotypes will be mistakenly classified as deterministic, while if they were assembled a greater number of times, a different polyomino would eventually be produced. To reduce the chance of error to essentially 0 (it can never be fully 0 as the process is stochastic), in the $S_{2,8}$ GP map we used values of $k = 1000$. For $S_{3,8}$, computationally more challenging owing to the much greater number of genotypes, we were only able to set $k = 200$ for the full enumeration, although from sampling at $k = 5000$, we found the exact number of phenotypes, as well as finding the $D_{\max} = 16$ value sufficient. Furthermore, from analyses of genotype sampling for values of $k > 200$, we saw only rare misclassification, leading to negligible error in this regard.

Owing to the much greater size of the $S_{4,16}$ GP, this could only be enumerated through a sampling approach and we made use of the algorithm developed by Jörg *et al.* [42]. We used smaller values of $k = 20$ to aid the speed of the algorithm, with any error due to misclassified deterministic assembly being absorbed into the error predicted by the algorithm, which is shown to be small in our work. The algorithm of Jörg *et al.* [42] relies upon a distance measure between phenotypes A and B . We constructed such a measure between polyominoes based upon the number of blocks that would need to be moved to produce B from A . If there is a difference in sizes of A and B , the number of blocks required to be added or removed is also included in the distance. This simple measure allowed an effective implementation of the algorithm.

References

- Alberch P. 1991 From genes to phenotype: dynamical systems and evolvability. *Genetica* **84**, 5–11. (doi:10.1007/BF00123979)
- Wagner A. 2005 *Robustness and evolvability in living systems*. Princeton, NJ: University Press Princeton.
- Noble D. 2006 *The music of life: biology beyond genes*. Oxford, UK: Oxford University Press.
- Pigliucci M. 2010 Genotype–phenotype mapping and the end of the ‘genes as blueprint’ metaphor. *Phil. Trans. R. Soc. B* **365**, 557–566. (doi:10.1098/rstb.2009.0241)
- Svensson E, Calsbeek R. 2012 *The adaptive landscape in evolutionary biology*. Oxford, UK: Oxford University Press.
- Kauffman SA. 1969 Metabolic stability and epigenesis in randomly constructed genetic nets. *J. Theor. Biol.* **22**, 437–467. (doi:10.1016/0022-5193(69)90015-0)
- Li F, Long T, Lu Y, Ouyang Q, Tang C. 2004 The yeast cell-cycle network is robustly designed. *Proc. Natl Acad. Sci. USA* **101**, 4781–4786. (doi:10.1073/pnas.0305937101)
- Mathews DH, Moss WN, Turner DH. 2010 Folding and finding RNA secondary structure. *Cold Spring Harb. Perspect. Biol.* **2**, a003665. (doi:10.1101/cshperspect.a003665)
- Schuster P, Fontana W, Stadler PF, Hofacker IL. 1994 From sequences to shapes and back: a case study in RNA secondary structures. *Proc. R. Soc. Lond. B* **255**, 279–284. (doi:10.1098/rspb.1994.0040)
- Cowperthwaite MC, Meyers LA. 2007 How mutational networks shape evolution: lessons from RNA models. *Annu. Rev. Ecol. Evol. Syst.* **38**, 203–230. (doi:10.1146/annurev.ecolsys.38.091206.095507)
- Dill KA. 1985 Theory for the folding and stability of globular proteins. *Biochemistry* **240**, 1501–1509. (doi:10.1021/bi00327a032)
- Li H, Helling R, Tang C, Wingreen N. 1996 Emergence of preferred structures in a simple model of protein folding. *Science* **273**, 666–669. (doi:10.1126/science.273.5275.666)
- Bornberg-Bauer E. 1997 How are model protein structures distributed in sequence space? *Biophys. J.* **73**, 2393–2403. (doi:10.1016/S0006-3495(97)78268-7)
- Ahnert SE, Johnston IG, Fink TMA, Doye JPK, Louis AA. 2010 Self-assembly, modularity, and physical complexity. *Phys. Rev. E* **82**, 026117. (doi:10.1103/PhysRevE.82.026117)
- Johnston IG, Ahnert SE, Doye JPK, Louis AA. 2011 Evolutionary dynamics in a simple model of self-assembly. *Phys. Rev. E* **83**, 066105. (doi:10.1103/PhysRevE.83.066105)
- Levy ED, Pereira-Leal JB, Chothia C, Teichmann SA. 2006 3D complex: a structural classification of protein complexes. *PLoS Comput. Biol.* **2**, e155. (doi:10.1371/journal.pcbi.0020155)
- Perica T, Marsh J, Sousa F, Natan E, Colwell L, Ahnert S, Teichmann S. 2012 The emergence of protein complexes: quaternary structure, dynamics and allostery. *Biochem. Soc. Trans.* **40**, 475. (doi:10.1042/BST20120056)
- Berman HM, Westbrook J, Feng Z, Gilliland G, Bhat TN, Weissig H, Shindyalov IN, Bourne PE. 2000 The Protein Data Bank. *Nucleic Acids Res.* **28**, 235–242. (doi:10.1093/nar/28.1.235)
- Levy ED, Erba EB, Robinson CV, Teichmann SA. 2008 Assembly reflects evolution of protein complexes. *Nature* **453**, 1262–1265. (doi:10.1038/nature06942)
- Villar G, Wilber AW, Williamson AJ, Thiara P, Doye JPK, Louis AA, Jochum MN, Lewis ACF, Levy ED. 2009 Self-assembly and evolution of homomeric protein complexes. *Phys. Rev. Lett.* **102**, 118106. (doi:10.1103/PhysRevLett.102.118106)
- Tuschl T, Gohlke C, Jovin TM, Westhof E, Eckstein F. 1994 A three-dimensional model for the hammerhead ribozyme based on fluorescence measurements. *Science* **266**, 785–789. (doi:10.1126/science.7973630)
- Hofacker IL. 2003 Vienna RNA secondary structure server. *Nucleic Acids Res.* **31**, 3429–3431. (doi:10.1093/nar/gkg599)
- Gursky O, Badger J, Li Y, Caspar DL. 1992 Conformational changes in cubic insulin crystals in the pH range 7–11. *Biophys. J.* **63**, 1210–1220. (doi:10.1016/S0006-3495(92)81697-1)
- Boriack-Sjodin PA, Margarit SM, Bar-Sagi D, Kuriyan J. 1998 The structural basis of the activation of RAS by SOS. *Nature* **394**, 337–343. (doi:10.1038/28548)
- Paoli M, Liddington R, Tame J, Wilkinson A, Dodson G. 1996 Crystal structure of T state haemoglobin with oxygen bound at all four haems. *J. Mol. Biol.* **256**, 775–792. (doi:10.1006/jmbi.1996.0124)
- Harrington DJ, Adachi K, Royer Jr WE. 1997 The high resolution crystal structure of deoxyhemoglobin S. *J. Mol. Biol.* **272**, 398–407. (doi:10.1006/jmbi.1997.1253)
- Wishner BC, Ward KB, Lattman EE, Love WE. 1975 Crystal structure of sickle-cell deoxyhemoglobin at 5 Å resolution. *J. Mol. Biol.* **98**, 179–194. (doi:10.1016/S0022-2836(75)80108-2)
- Rothmund PWK, Papadakis N, Winfree E. 2004 Algorithmic self-assembly of DNA Sierpinski triangles. *PLoS Biol.* **2**, e424. (doi:10.1371/journal.pbio.0020424)
- Ferrada E, Wagner A. 2012 A comparison of genotype–phenotype maps for RNA and proteins. *Biophys. J.* **102**, 1916–1925. (doi:10.1016/j.bpj.2012.01.047)
- Wagner A. 2008 Robustness and evolvability: a paradox resolved. *Proc. R. Soc. B* **275**, 91–100. (doi:10.1098/rspb.2007.1137)
- Wang H. 1961 Proving theorems by pattern recognition—II. *Bell Syst. Tech. J.* **40**, 1–41. (doi:10.1002/j.1538-7305.1961.tb03975.x)
- Winfree E, Yang X, Seeman NC. 1996 Universal computation via self-assembly of DNA: some theory and experiments. *DNA Based Comp. Two* **44**, 191–213.
- Rothmund PWK, Winfree E. 2000 The program-size complexity of self-assembled squares (extended abstract). In *Proc. Thirty-second Annual ACM Symp. on Theory of Computing, Portland, OR*, pp. 459–468. New York, NY: ACM. (doi:10.1145/335305.335358)
- Soloveichik D, Winfree E. 2007 Complexity of self-assembled shapes. *SIAM J. Comput.* **36**, 1544–1569. (doi:10.1137/S0097539704446712)
- Stich M, Briones C, Manrubia SC. 2008 On the structural repertoire of pools of short, random RNA sequences. *J. Theor. Biol.* **252**, 750–763. (doi:10.1016/j.jtbi.2008.02.018)
- Irbäck A, Troein C. 2002 Enumerating designing sequences in the HP model. *J. Biol. Phys.* **28**, 1–15. (doi:10.1023/A:1016225010659)
- Wagner A. 2011 *The origins of evolutionary innovations: a theory of transformative change in living systems*. Oxford, UK: Oxford University Press.
- Thorpe JP. 1982 The molecular clock hypothesis: biochemical evolution, genetic differentiation and systematics. *Annu. Rev. Ecol. Syst.* **13**, 139–168. (doi:10.1146/annurev.es.13.110182.001035)
- Ciliberti S, Martin OC, Wagner A. 2007 Robustness can evolve gradually in complex regulatory gene networks with varying topology. *PLoS Comput. Biol.* **3**, e15. (doi:10.1371/journal.pcbi.0030015)
- Fernández P, Solé RV. 2007 Neutral fitness landscapes in signalling networks. *J. R. Soc. Interface* **4**, 41–47. (doi:10.1098/rsif.2006.0152)
- Schaper S. 2012 *On the significance of neutral spaces in adaptive evolution*. PhD thesis, University of Oxford, Oxford, UK.
- Jörg T, Martin O, Wagner A. 2008 Neutral network sizes of biological RNA molecules can be computed and are not atypically small. *BMC Bioinform.* **9**, 464. (doi:10.1186/1471-2105-9-464)
- Aguirre J, Buldu JM, Stich M, Manrubia SC. 2011 Topological structure of the space of phenotypes: the case of RNA neutral networks. *PLoS ONE* **6**, e263241. (doi:10.1371/journal.pone.0026324)
- Schaper S, Johnston IG, Louis AA. 2012 Epistasis can lead to fragmented neutral spaces and contingency in evolution. *Proc. R. Soc. B* **279**, 1777–1783. (doi:10.1098/rspb.2011.2183)
- Kauffman S, Levin S. 1987 Towards a general theory of adaptive walks on rugged landscapes. *J. Theor. Biol.* **128**, 11–45. (doi:10.1016/S0022-5193(87)80029-2)
- Gravner J, Pitman D, Gavrillets S. 2007 Percolation on fitness landscapes: effects of correlation, phenotype, and incompatibilities. *J. Theor. Biol.* **248**, 627–645. (doi:10.1016/j.jtbi.2007.07.009)
- Perelson AS, Oster GF. 1979 Theoretical studies of clonal selection: minimal antibody repertoire size and reliability of self-non-self discrimination. *J. Theor. Biol.* **81**, 645–670. (doi:10.1016/0022-5193(79)90275-3)
- Kirschner M, Gerhart J. 1998 Evolvability. *Proc. Natl Acad. Sci. USA* **95**, 8420–8427. (doi:10.1073/pnas.95.15.8420)

49. Bloom J, Lu Z, Chen D, Raval A, Venturelli O, Arnold F. 2007 Evolution favors protein mutational robustness in sufficiently large populations. *BMC Biol.* **5**, 29. (doi:10.1186/1741-7007-5-29)
50. Gallagher A. 2009 *Evolvability: a formal approach*. PhD thesis, University of Oxford, Oxford, UK.
51. Masel J, Trotter MV. 2010 Robustness and evolvability. *Trends Genet.* **26**, 406–414. (doi:10.1016/j.tig.2010.06.002)
52. Cowperthwaite MC, Economo EP, Harcombe WR, Miller EL, Meyers LA. 2008 The ascent of the abundant: how mutational networks constrain evolution. *PLoS Comput. Biol.* **4**, e1000110. (doi:10.1371/journal.pcbi.1000110)
53. Draghi JA, Parsons TL, Wagner GP, Plotkin JB. 2010 Mutational robustness can facilitate adaptation. *Nature* **463**, 353–355. (doi:10.1038/nature08694)
54. Raman K, Wagner A. 2011 The evolvability of programmable hardware. *J. R. Soc. Interface* **8**, 269–281. (doi:10.1098/rsif.2010.0212)
55. Borenstein E, Krakauer DC. 2008 An end to endless forms: epistasis, phenotype distribution bias, and nonuniform evolution. *PLoS Comput. Biol.* **4**, e1000202. (doi:10.1371/journal.pcbi.1000202)
56. Boldhaus G, Klemm K. 2010 Regulatory networks and connected components of the neutral space. *Eur. Phys. J. B* **77**, 233–237. (doi:10.1140/epjb/e2010-00176-4)
57. Gavrilets S. 2004 *Fitness landscapes and the origin of species (MPB-41)*. Princeton, NJ: Princeton University Press.
58. Kloczkowski A, Jernigan RL. 1997 Computer generation and enumeration of compact self-avoiding walks within simple geometries on lattices. *Comput. Theor. Polymer Sci.* **7**, 163–173. (doi:10.1016/S1089-3156(97)00022-6)