



Evolutionary dynamics of a lattice dimer: a toy model for stability vs. affinity trade-offs in proteins

E Loffredo^{1,*} , E Vesconi^{1,7}, R Razban², O Peleg^{3,4,5} ,
E Shakhnovich⁶, S Cocco¹ and R Monasson¹

¹ Laboratory of Physics of the Ecole Normale Supérieure, CNRS UMR 8023 and PSL Research, Sorbonne Université, 24 rue Lhomond, 75231 Paris Cedex 05, France

² Laufer Center for Physical and Quantitative Biology, Stony Brook University, Stony Brook, NY, United States of America

³ Computer Science Dept, University of Colorado Boulder, Boulder, CO 80309, United States of America

⁴ Santa Fe Institute, Santa Fe, NM 87501, United States of America

⁵ Ecology & Evolutionary Biology, Physics, Applied Math Dept, University of Colorado Boulder, Boulder, CO 80309, United States of America

⁶ Department of Chemistry and Chemical Biology, Harvard University, Cambridge, MA 02138, United States of America

E-mail: emanuele.loffredo@phys.ens.fr

Received 20 March 2023; revised 6 September 2023

Accepted for publication 27 September 2023

Published 13 October 2023



CrossMark

Abstract

Understanding how a stressor applied on a biological system shapes its evolution is key to achieving targeted evolutionary control. Here we present a toy model of two interacting lattice proteins to quantify the response to the selective pressure defined by the binding energy. We generate sequence data of proteins and study how the sequence and structural properties of dimers are affected by the applied selective pressure, both during the evolutionary process and in the stationary regime. In particular we show that internal contacts of native structures lose strength, while inter-structure contacts are strengthened due to the folding-binding competition. We discuss how dimerization is achieved through enhanced mutability on the interacting faces, and how the designability of each native structure changes upon introduction of the stressor.

Keywords: protein interaction, evolutionary trade-offs, lattice protein

(Some figures may appear in colour only in the online journal)

⁷ Currently at Planven Entrepreneur Ventures, Glockengasse 9, 8001, Zurich.

* Author to whom any correspondence should be addressed.

1. Introduction

To be functional proteins must satisfy a variety of constraints, related to their stability, activity, specificity etc. . . . Most proteins reach an adequately folded stable conformation to efficiently perform their tasks involving binding with other biomolecules, such as DNA, peptidic ligand or other proteins [1]. How these various evolutionary constraints combine to shape the evolutionary landscape of proteins is of great interest. This question is particularly crucial when additional stressors, e.g. an antibiotic on a bacterial population, are dynamically applied during evolution. To what extent organisms can accommodate new constraints, while still fulfilling constitutive ones, is an important issue.

In this work we introduce a minimal setting to address this question from a theoretical point of view. We consider lattice proteins (LPs) [2–5], an exactly solvable model of amino acids chains constrained to fold on the sites of a cube. Despite their simplicity, LPs share many features with real proteins and have been proved to be a useful tool for studying protein folding and designability. In addition to require that proteins acquire their native folds with high probability, we impose that they bind each other and form a stable dimer [6]. The study of the conflict between these two requirements, and its consequence on the distribution of adequate sequences is the goal of the present work [6–9].

In a first part of the work we study the stationary regime in which sequences evolve through a mutational dynamics under a two-fold selection pressure requiring the native folds and the dimer conformation to be achieved. We make use of the so-called direct coupling approximation (DCA) [10–13]—a graphical model based approach—to model the distribution of sequences subject to selection constraints. This inverse modeling approach consists in finding effective energetic parameters (fields and couplings in a Potts Hamiltonian) describing the empirical distribution of sequences. Applying DCA to the exactly solvable model of LPs allows us to describe in great detail our dimeric system. We show that the inferred couplings are excellent predictors of the intra- and inter-protein structure; moreover, they indirectly keep track of the effect of selection due to the mentioned constraints.

In the second part we focus on the full evolutionary history, from the initial state with two non-interacting LPs in their native structures to the final state where a bounded dimer is formed. We look at the competitive dynamics between folding into the native structures vs. realizing the functional protein–protein interaction to characterize the evolutionary trade-offs due to the binding constraints.

2. The model

2.1. Native folds

We focus on an exactly solvable model, namely LPs, to study the formation of protein dimer, that is a macromolecular complex formed by two protein monomers. Each model protein consists of a chain of $L = 27$ amino acids that occupy the sites of a $3 \times 3 \times 3$ cubic lattice [4, 5]. A valid conformation, hereafter called structure or fold, is a non-interacting chain that visits each site once, and there are 103 406 of such possible structures (excluding global symmetries) [5]. Two examples of folds, called S_A and S_C , are shown in figure 1. For computational efficiency, we restrict ourselves to a representative subset of $\mathcal{N} = 10000$ structures [14].




Figure 1. Representation of the protein dimer studied in this work, composed of a structure S_A (right) attached to a structure S_C (left) through a specific face and orientation, which we refer to as functional binding mode (red dotted lines in figure). For each structure the backbone of protein is highlighted with solid black lines. Dashed lines highlight binding contacts between the two structures.

Two amino acids are said to be in contact if they are nearest neighbors on the cubic lattice (but not on the backbone). The contact matrix \mathbf{c}^S of structure S is the 27×27 adjacency matrix such that

$$c_{ij}^S = \begin{cases} 1 & i, j \text{ in contact,} \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

This matrix fully defines the structure S . Given a sequence $\mathbf{A} = (a_1, \dots, a_{27})$ of amino acids folded into structure S , we can assign the energy

$$\mathcal{E}(\mathbf{A}|S) = \sum_{i < j} c_{ij}^S E(a_i, a_j), \quad (2)$$

where residues in contact interact via the Miyazawa–Jernigan potential E [15], an empirical, symmetric 20×20 -dimensional matrix containing effective interaction energies for each couple of amino acids. The probability that sequence \mathbf{A} folds into structure S , hereafter called native probability, is given by the Gibbs–Boltzmann distribution at unit temperature, i.e.

$$P_{\text{nat}}(S|\mathbf{A}) = \frac{e^{-\mathcal{E}(\mathbf{A}|S)}}{\sum_{S'=1}^{\mathcal{N}} e^{-\mathcal{E}(\mathbf{A}|S')}}. \quad (3)$$

Stable structures S for the sequence \mathbf{A} are the ones that maximize the gap between their energy $\mathcal{E}(\mathbf{A}|S)$ and the one of competing structures S' .

2.2. Protein–protein interaction

We now move the attention to dimer LPs (hereafter only referred to as dimer), i.e. aggregate of two LPs bounded together through one of their faces, whose chain is made of a sequence of $L = 54$ amino acids. To model protein–protein interaction we decide to look at only one among the $6 \times 6 \times 4 = 144$ possible binding modes, so the two LPs interact functionally via two specific faces with a specific orientation. Indeed, here one interface is considered *functional* while the 143 remaining might exist but are not deemed as functional in this model. An example of

dimer—which corresponds to the one we studied in this work—is shown in figure 1. The interaction between faces of structure S_1 and S_2 brings into the arena a new energy contribution in the form

$$\mathcal{E}_{\text{int}}(\mathbf{A}_1, \mathbf{A}_2 | S_1 + S_2, k) = \sum_{i \in S_1, j \in S_2} b_{ij}^k E(a_i, a_j), \quad (4)$$

where $\mathbf{A}_1, \mathbf{A}_2$ are the amino acid sequences folded in S_1, S_2 respectively, and b_{ij}^k is the contact map of the k th binding mode, namely

$$b_{ij}^k = \begin{cases} 1 & i \in S_1, j \in S_2 \text{ in interaction,} \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

Hence, the interaction probability associated to the functional binding mode (hereafter always labeled with $k = 1$) reads

$$P_{\text{int}}(S_1 + S_2 | \mathbf{A}_1, \mathbf{A}_2) = \frac{e^{-\mathcal{E}_{\text{int}}(\mathbf{A}_1, \mathbf{A}_2 | S_1 + S_2, 1)}}{\sum_{m=1}^{144} e^{-\mathcal{E}_{\text{int}}(\mathbf{A}_1, \mathbf{A}_2 | S_1 + S_2, m)}}. \quad (6)$$

2.3. Dimerization

We can now write the full probability for a dimer (up to a normalization constant):

$$P(\mathbf{A}_1, \mathbf{A}_2) \propto P_{\text{nat}}(S_1 | \mathbf{A}_1)^{\hat{\beta}_{\text{nat}}} \times P_{\text{nat}}(S_2 | \mathbf{A}_2)^{\hat{\beta}_{\text{nat}}} \\ \times P_{\text{int}}(S_1 + S_2 | \mathbf{A}_1, \mathbf{A}_2)^{\hat{\beta}_{\text{int}}} \quad (7)$$

The exponents $\hat{\beta}_{\text{nat}}$ and $\hat{\beta}_{\text{int}}$ acts as stressors that control the stringency of evolutionary selection to fold into native conformations and to bind functionally. In practice, we want P_{nat} and P_{int} to reach values very close to 1, e.g. 0.99 or 0.999, implying that $\hat{\beta}_{\text{nat}}$ and $\hat{\beta}_{\text{int}}$ must be very large, of the order of 1000. To avoid manipulating these large numbers we introduce the rescaled stressors

$$\beta_{\text{nat}} = \frac{\hat{\beta}_{\text{nat}}}{1000}, \quad \beta_{\text{int}} = \frac{\hat{\beta}_{\text{int}}}{1000}, \quad (8)$$

where β_{nat} and β_{int} are of the order of 1.

The effective Hamiltonian of the two-sequence system therefore reads

$$\mathcal{H}(\mathbf{A}_1, \mathbf{A}_2) = -1000 \beta_{\text{nat}} \log P_{\text{nat}}(S_1 | \mathbf{A}_1) \\ - 1000 \beta_{\text{nat}} \log P_{\text{nat}}(S_2 | \mathbf{A}_2) \\ - 1000 \beta_{\text{int}} \log P_{\text{int}}(S_1 + S_2 | \mathbf{A}_1, \mathbf{A}_2). \quad (9)$$

3. Sampling dimer space

3.1. Monte Carlo (MC) metropolis sampling

We generate a multiple sequence alignment (MSA) for the dimer formed by two LPs folded in two specific structure, labeled with S_A and S_C as in [16], through MC simulations via the Metropolis rule. For each dimer in the MSA, the simulation starts with two sequences folded in structure $S_A (= S_1)$ and $S_C (= S_2)$ randomly taken from the two MSAs built in [16]. At each time step we perform the following routine to update the sequences:

1. **Mutation:** a move consisting of mutating at random two amino acids (one on each protein) is attempted—say, $a_{1i} \rightarrow a'_{1i}$ and $a_{2i} \rightarrow a'_{2i}$; let $\mathbf{A}'_1, \mathbf{A}'_2$ denote the mutated dimer, while $\mathbf{A}_1, \mathbf{A}_2$ refers to the old dimer.
2. **Selection:** if $P(\mathbf{A}'_1, \mathbf{A}'_2) > P(\mathbf{A}_1, \mathbf{A}_2)$ the move is accepted, otherwise it is accepted with probability $P(\mathbf{A}'_1, \mathbf{A}'_2)/P(\mathbf{A}_1, \mathbf{A}_2)$.

For each couple of value $(\beta_{\text{nat}}, \beta_{\text{int}})$ we construct a MSA of $N = 23\,000$ dimer sequences, separated by $N_{it} = 5000$ MC steps to avoid correlations. The time taken by the dimer to relax to equilibrium is independent from the two initial sequences and slightly varies on the values $(\beta_{\text{nat}}, \beta_{\text{int}})$, but remains smaller than N_{it} steps. Moreover, in order to study the evolutionary dynamics of dimer we also built MSA at intermediate steps, i.e. before reaching thermalization. In practice, during MC simulations, we collected MSA of evolving dimer at steps 0; 50; 200; 500; 1000; 3000; 5000—with step = 0 corresponding to the starting situation, while step = N_{it} corresponds to the final dimer once thermalization has occurred.

3.2. Evolutionary model sampling

An alternative way to construct a MSA dimer is to simulate the evolutionary dynamics of a population of individuals (sequences), hereafter referred to as \mathbf{X} . Having an additional MSA generated with a completely different approach allows us to perform the analysis on both of them independently and check whether we obtain the same results.

The evolutionary model starts with a population of size N , where each individual has a genome made by the two LP chains. The initial sequences have high folding probabilities but do not generally interact, i.e. they were generated through the previous MC dynamics with $\beta_{\text{nat}} = 1, \beta_{\text{int}} = 0$. For the sake of simplicity we then randomly pick up one of the sequences \mathbf{A}_2 for S_C and freeze it. The evolutionary process will focus on sequences \mathbf{A}_1 for S_A only.

At each generation of the evolutionary dynamics, two steps are carried out to make evolve the parent population \mathbf{X} into a new population, \mathbf{X}^* :

1. **Mutation:** for each individual and for each site of the sequence we draw a binomial random variable $m_i = 0, 1$ with mean μ . If $m_i = 1$ the amino acid on site i is replaced with a new one, drawn from a background distribution of frequencies of amino acids in the MSA associated to structure S_A . We end up with a new population \mathbf{X}' .
2. **Selection:** to each individual $X'_i = (\mathbf{A}_1, \mathbf{A}_2)$ in the mutated population \mathbf{X}' we associate a fitness value $P(X'_i)$ given by equation (7). Then we draw a multinomial random variable $\kappa = 1, \dots, N$ according to the fitnesses of all individuals in the population, i.e. the probability of drawing κ is

$$\mathcal{P}_\kappa = \frac{P(X'_\kappa)}{\sum_i P(X'_i)}. \quad (10)$$

We thus generate an off-spring identical to κ th mutated individual in \mathbf{X}' . This random extraction process is repeated N times (with replacement), and we end up with the population of off-springs \mathbf{X}^* , where the numbers of copies of individuals in \mathbf{X}' are, on average, proportional to their fitnesses.

These two steps are repeated T times, and results are averaged over multiple sample populations.

4. Equilibrium properties for various stressor strengths

In this section we investigate how the value of the stressor β_{int} affects the sequence distributions after evolution has reached an equilibrium state. The dynamical transient preceding this equilibrium regime will be studied in the next section.

The study of equilibrium properties of the dimer can be carried out from sequence data generated either with MC sampling or population dynamics, with similar results from a qualitative point of view. Hereafter, we use MC generated data to characterize the statistics of residues in the dimer and its connection with structural properties (inter- and intra-protein contacts). We then focus on sequences produced by the population dynamics model to understand the interplay between the population size, the number of mutations per individual, and the response to the stressor.

Unless otherwise said, we perform simulations sharing the same parameter $\beta_{\text{nat}} = 1$ for the two models. We also set for population dynamic model $\mu = 1/L$, so that, on average, we have one mutation per individual as in the MC evolution. The population dynamics model is characterized by an additional parameter to tune, namely the population size.

4.1. Foldability–dimerization trade-offs, and their effects on sequence statistics

We start by tuning β_{int} to see how changing the selection pressure applied along the P_{int} direction affects the capability of the individual proteins to reach their native folds. The behaviors of $P_{\text{nat}}(S_A|\mathbf{A}_1)$, $P_{\text{nat}}(S_C|\mathbf{A}_2)$, and $P_{\text{int}}(S_A + S_C|\mathbf{A}_1, \mathbf{A}_2)$ are shown in figure 2. As a general trend, P_{int} is an increasing function of β_{int} as expected from the explicit dependence of P on β_{int} in equation (7). We also observe that the two P_{nat} decrease with β_{int} , an effect of the evolutionary trade-offs between two competing fitness components. As β_{int} grows to large values, the probabilities that the two proteins adopt their native folds become small: in this regime the selective pressure favoring protein–protein binding is too strong to cope with the folding constraints. For intermediate values of β_{int} we are able to obtain both good foldings and high interaction between the two structures.

The presence of trade-offs observed at the phenotypic level in figure 2 can be studied at the sequence level. To obtain a fine characterization of the sequence statistics resulting from the evolutionary constraints, we infer a pairwise Potts model with $q = 20$ -state variables (corresponding to the 20 possible amino acids), and $N = 54$ variables (corresponding to the full dimer length) [17]. The energy of this effective model is the sum of two contributions, featuring local and interacting terms:

$$\mathcal{H}^{\text{Potts}}(\mathbf{A}_1, \mathbf{A}_2) = - \sum_{i < j} J_{ij}(u, v) \delta_{a_i, u} \delta_{a_j, v} - \sum_i h_i(v) \delta_{a_i, v}, \quad (11)$$

where the indices i, j run along all the dimer sequence, i.e. $i, j = 1, \dots, L$. Here, $J_{ij}(u, v)$ represents the coupling between amino acid u at position i and amino acid v at position j along the sequence; given two positions i, j in contact in the 3D dimer structure, they reproduce the interacting MJ potential $E(u, v)$. The parameters $h_i(u)$ are local fields acting on position i that depends on specific amino acid u . To infer the values of these parameters we follow a Boltzmann machine (BM) learning procedure [18], consisting in maximizing the average log-likelihood $\langle \log \mathcal{P} \rangle_{\text{data}}$ of the model over the sequence data. In practice, we assume the data to be Gibbs-distributed as

$$\mathcal{P}^{\text{Potts}} = \frac{e^{-\mathcal{H}^{\text{Potts}}}}{Z}, \quad (12)$$




Figure 2. Values of P_{int} , $P_{\text{nat}}(S_A)$, $P_{\text{nat}}(S_C)$ as functions of β_{int} at equilibrium. Inset: zoom in the region $\beta_{\text{int}} \sim 1$, where all probabilities have high values. Averages are carried out over 10^4 realizations.

where $\mathcal{H}^{\text{Potts}}$ is given by equation (11). We compute $\langle \log \mathcal{P} \rangle_{\text{data}}$ over the MSA dataset for an initial guess of the parameters; we then proceed to its maximization by numerically ascending the gradient $\nabla \langle \log \mathcal{P} \rangle_{\text{data}}$, until the log-likelihood is maximized (see appendix A for further details on BM learning). The meaning of the inferred inter- and intra-protein couplings is studied below in sections 4.2 and 4.3.

A stringent test of the accuracy of the inferred Potts model is its ability to generate new dimer sequences that have both the right target native structures and high, specific face–face interaction. In practice, we generate sequences with MC from the inferred distribution

$$\mathcal{P}^{\text{Potts}}(\mathbf{A}_1, \mathbf{A}_2) \propto e^{-\mathcal{H}^{\text{Potts}}(\mathbf{A}_1, \mathbf{A}_2)/T}, \quad (13)$$

where T is a fictitious sampling temperature used to control the broadness of the sampled region, hence the diversity of the generated MSA. By choosing T lower than unity, which is the implicit value of the inference temperature, we are able to sample sequences with low energies. For further details about the generation procedure we refer to appendix A.

We show in figure 3 the histograms of the ground-truth probabilities—equations (3) and (6)—which show these sequences are good dimers. We see that these probabilities are quite high, proving that most of the necessary information needed to model a good dimer can be captured by pairwise interactions and local biases, in agreement with some recent works [19, 20]. In addition, the generated sequences are far from the old ones and have high diversity between each other, as it can be seen in figure 3 (bottom right panel) where we plot the distribution of the Hamming distance for each pair of sequence; hence we are sampling a different subset of the dimer space.

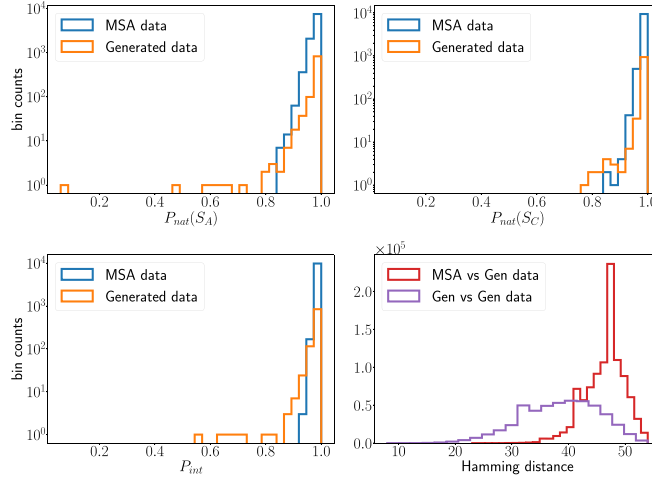


Figure 3. Distribution of $P_{\text{nat}}(S_A)$, $P_{\text{nat}}(S_C)$, P_{int} for the MC MSA and generated MSA (blue and orange, respectively). Bottom right: distribution of Hamming distance between the MC MSA and generated MSA and inside the generated MSA (red and purple, respectively). We sample new sequences with $T = 0.5$.

4.2. Structural significance of the statistical couplings

Following the standard direct-coupling approach [11], we rank the inferred couplings in descending order of their Frobenius scores (L_2 norms),

$$F_{ij} = \sqrt{\sum_{a_i, a_j} J_{ij}(a_i, a_j)^2}, \quad (14)$$

and use this ranking as a predictor of the contacts in the dimer. Informally speaking, we expect the strongest inferred couplings to be those corresponding to real contacts among amino acids in the dimer (see appendix A). We compute the positive predictive value (PPV) at rank k as the fraction of the top- k scores whose pairs of positions along the sequence are effectively in contact in the 3D dimer.

We show in figure 4 the PPV used for contacts prediction in structure S_A and S_C (left and right panels, respectively), for different values of the interaction strength β_{int} . Between the two structures, we can see that—whatever the value of β_{int} is—the PPV performs better on predicting contacts in structure S_A rather than in S_C . Indeed, for S_A we always predict the first 19 contacts out of the total 28, while for S_C we miss more contacts. The missed contacts always include the central site and the central site of the binding mode, see below. Both structures S_A and S_C can host an extremely large number of sequences: the higher this number, the more designable the structure is said to be. However, it appears that having high designability is harmful for inferring couplings: between S_A and S_C , the former can host less sequences, which means it is more specific and allows for better inference and contacts prediction. As for the different interaction strengths, we cannot identify a particular trend for increasing values of β_{int} both for structure S_A and structure S_C .

We also compute the PPV for the binding mode for several values of β_{int} in order to assess the quality of contact prediction between the two structures (see figure 4, middle panel). Regardless of the interaction strength, we only miss one contact out of the nine present on




Figure 4. PPV for structure S_A (left panel, circles) and structure S_C (right panel, squares) at different values of β_{int} . Middle panel: PPV for contacts belonging to the functional binding mode (stars). Data points at different β_{int} overlap, as we are able to predict all contacts but the central one. The MSA is made of 23 000 dimer sequences evolved for 5000 MC steps.

the interacting layer. Interestingly, the missed contact is always the same for all β_{int} , and corresponds to the central contact between sites 1–19 in figure 1. In fact, contacts involving central sites are generally the ones predicted worse, as central sites have more neighbors and are often present in several competing structures [16].

4.3. Characterization of dimer interactions

Additionally, we characterize the binding modes between the two structures using the quantity λ_{ij} —which is another score computed again from the inferred couplings J_{ij} that we used to assess the quality of the functional binding mode against the remaining 143 modes. Indeed, in [16] the authors observe a linear dependency between the couplings and the MJ energy matrix, with a slope given by λ_{ij} for each pair (i, j) . Such slope can be computed as

$$\lambda_{ij} = -\frac{\sum_{a,b} J_{ij}(a,b) E(a,b)}{\sum_{a,b} E(a,b)^2}, \quad (15)$$

where the sum runs over all the possible amino acids. The quantities λ_{ij} can then be seen as a measure of the coevolutionary constraints imposed by the design of the two structures.

The projection scores $\langle \lambda \rangle$, averaged over the nine binding contacts for each binding mode, are shown in figure 5 (left panel). Among them, we have identified some relevant subsets depending on the number of maintained/mismatched contacts involving the interacting faces wrt the functional binding mode contacts (see figure 5, right panel for a visual representation with the associated color code):

- single red peak corresponds to the functional binding and it is the highest one;
- four orange peaks corresponding to binding modes where three functional contacts are still maintained. An example of such configuration is in figure 5, right panel (a);
- black subset corresponds to binding modes where one out of nine binding contacts is still present. An example of such configuration is in figure 5, right panel (b);
- pink subset contains all binding modes where there is at least one mismatched contact. An example of such configuration is in figure 5, right panel (c);
- three green binding modes that have eight out of nine mismatched contacts. An example of such configuration is in figure 5, right panel (d).

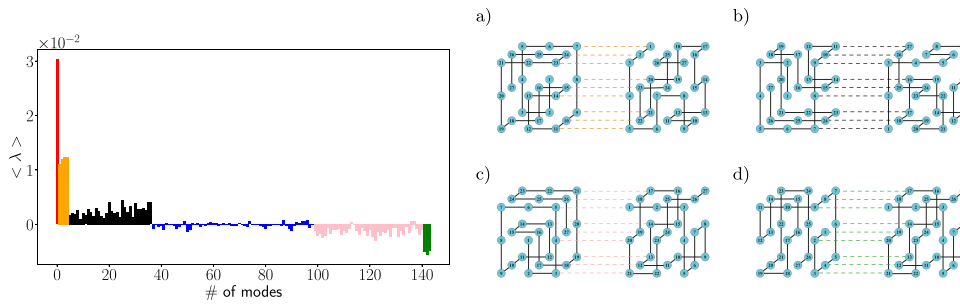


Figure 5. Left: scores λ_{ij} averaged over all the nine binding contacts for each mode, computed at equilibrium according to equation (15). Except for the scores associated to binding modes described in the main text (see section 4.3 for orange, black, pink and green bars and figure 1 for the red bar related to functional binding), the remaining binding modes—the 62 blue ones—exhibit a rather flat landscape. They are associated to binding modes that do not involve sites belonging to the interacting faces of the functional binding mode (e.g. binding between the two back faces), hence they are just slightly favored or disfavored depending on the specific case. Right: schematic of representative binding modes belonging to each of the subsets described in the main text: (a) orange mode, (b) black mode, (c) pink mode, (d) green mode.

In the four orange configurations, the scores λ_{ij} associated to the three binding contacts of the functional mode are very high and thus are responsible for a large value of $\langle \lambda \rangle$. The same holds for black configurations, where this time only one λ_{ij} is high. Surprisingly the orange peaks are almost 1/3 of the red one and the black ones are on average almost 1/9 of the red one.

Conversely, the pink and green configurations have, respectively, one and eight negative scores thus resulting in binding modes that we strongly avoid.

5. Transient responses to stressor

Up to now we have only discussed the stable state of the dimer sequence distribution. We now consider the effects of a rapid change of the stressor value β_{int} .

5.1. Dynamical recovery of structural fitness

5.1.1. Response to a step-like change. In order to achieve a clear view of what happens during MC evolution, we compute the native probability P_{nat} and the interaction probability P_{int} at each time step for the whole length of the MC simulation and for $N_{\text{seq}} = 1000$ dimer sequences. We then average $P_{\text{nat}}(t)$ and $P_{\text{int}}(t)$ over these N_{seq} dimer sequences.

Figure 6 (top) shows the evolution of $P_{\text{nat}}(t)$ and $P_{\text{int}}(t)$ in time for two values of β_{int} . For large β_{int} we observe a huge drop in the folding probabilities P_{nat} associated to both structures S_A and S_C , while P_{int} monotonically increases to reach a β_{int} -dependent plateau value. This dynamical evolution is a direct consequence of the evolutionary trade-offs between P_{nat} and P_{int} , as the amino acids on the interacting faces enter both in equations (3) and (6). The evolutionary trajectories in figure 6 show that, in order to increase the interaction probability associated to the binding mode, the two sequences are forced to go through sub-optimal—even very bad—states from a structural point of view (P_{nat} values down to 0.2). After this strong drop and the amino acids controlling the dimer interaction have been optimized enough both




Figure 6. Top: evolutionary trajectories of P_{nat} for S_A (blue) and S_C (orange) and of P_{int} (green) for weak (left, solid line with $\beta_{int} = 0.5$) and strong (right, dotted line with $\beta_{int} = 10.0$) protein-protein interactions. Curves are averaged over 3000 dimer sequences. Bottom: evolutionary trajectory $P_{nat}(S_A)$ for increasing values of τ . The smaller τ , i.e. the sharper is the $\beta_{int}(t)$ dependence, the more we experience the out of equilibrium effect. Let us note that here we are plotting the short-time evolution, since we are only interested to see how the sharpness of the stressor β_{int} affects the evolution. Here exceptionally $\beta_{nat} = 0.2$ and $\beta_{int}(0) = 0.0$, $\beta_{int}^F = 2.0$.

P_{nat} start increasing again but reach values lower than the ones at the beginning of the evolutionary trajectory, i.e. in the absence of interaction. The constraint arising from the binding mode does not allow them to optimally maximize their single structure folding; in addition, as expected, the larger β_{int} , the bigger the drop in P_{nat} is for both structures, and the lower the final value of P_{nat} [6, 21].

Interestingly, for any β_{int} value, the drop in P_{nat} for structure S_A is always larger than the one for S_C . Even if both structures are undergoing roughly the same number of mutations, structure S_C remains more stable compared to S_A , a fact related to its larger designability [16]. We will study this point in more details in section 5.2.

5.1.2. Case of smooth increases of the stressor strength. The drop in P_{nat} seen in figure 6 is an out-of-equilibrium effect, resulting from the abrupt change of selection pressure from 0 to β_{int} . To better study this effect, we consider a smooth, time-dependent stressor during the MC evolution

$$\beta_{int}(t) = \beta_{int}^F \tanh\left(\frac{t}{\tau}\right), \quad (16)$$

where τ sets the time scale of the stressor ($\tau \rightarrow 0$ gives back the step-like function studied so far). In figure 6 (bottom panel) we plot the transient dynamics of P_{nat} for several values




Figure 7. Time course of the diversity $D(t)$ for different selection pressures β_{int} . Inset: fitnesses $P_{\text{nat}}(S_A)$, P_{int} at equilibrium after population dynamic evolution for several stressor values β_{int} . We recall that P_{int} cannot increase further as we are keeping S_C fixed compared to MC evolution. The parameters used are $\beta_{\text{nat}} = 10^{-3}$, $N = 1500$ and $\mu = 1/L$, with β_{int} kept fixed to its value during the evolution. Note that the values of β_{int} are much smaller than in MC dynamics, due to the large population size, see section 6.

of τ with $\beta_{\text{int}}(t=0) = 0$ and $\beta_{\text{int}}(t \rightarrow \infty) = \beta_{\text{int}}^F = 2.0$. While for long enough simulations, i.e. when $t \gg \tau$, all curves reach the same plateau, the drop in P_{nat} decays with the time scale τ in equation (16). For large enough τ , i.e. when the evolution can be considered adiabatic, P_{nat} decreases monotonically over time.

5.1.3. Evolution of a non-clonal population. While we have so far studied the evolution of a single sequence, we now consider the case of population of constant size N . Our aim is to characterize how the selection pressure β_{int} affects the substitution rate, the diversity of population, and also how the effects of the stressor relates to the population size.

We first focus on the diversity $D(t)$, as the fraction of diverse individuals (unique sequences) present in the population at time t . The time behavior of $D(t)$ is shown for various values of β_{int} (at fixed size and mutation rates) in figure 7. Increasing β_{int} makes the population less diverse, as fewer sequences satisfy the selection constraints and give rise to off-springs. The loss of diversity is maximal at the drop in P_{nat} . The inset of figure 7 shows the stationary values of P_{nat} and P_{int} ; we recall that the latter cannot reach as high a value as with MC evolution since the second sequence in the dimer (associated to protein S_C is not allowed to evolve).

We then study the substitution rate $m(t)$, defined as the average Hamming distances between the sequences at step $t+1$ and their parents at time t ; without selection this rate would be on average equal to the number of mutations proposed per individual, i.e. μL . The substitution rate per individual is plotted as a function of the evolutionary time in figure 8. We observe, for sizes $N > 1$, a peak in m at the beginning of dynamics, decaying to a plateau. Conversely, for $N = 1$, the substitution rate fluctuates around the average value $\langle m \rangle = \mu L$. The initial peak in $m(t)$ is therefore mostly due to selection, rather than to mutations (we recall that, in MC, $N = 1$ and mutations are always proposed, an analogous quantity to m being the acceptance rate).

The time behaviors of the diversity and of the substitution rate can be qualitatively understood in a simplified scenario, in which the maximum of m is achieved in only one evolutionary time point (cf figure 9 that supports our argument):




Figure 8. Substitution rate $m(t)$ for different population sizes N . The parameters used are $\beta_{\text{int}} = 10$ and $\mu = 1/L$.




Figure 9. Distribution of log-fitnesses during the first three steps of evolution. The parameter values are $\beta_{\text{nat}} = 10^{-3}$, $\beta_{\text{int}} = 10^{-2}$, $\mu = 1/L$ and $N = 1000$.

- At $t = 0$, all sequences are distinct, hence $D = 1$, and the fitness distribution is very broad. This broad profile is maintained after random mutations are introduced.
- Hence, at $t = 1$, selection will only keep the few strains that are fitter in the population, resulting in a poorly diversity D . The fitness distribution is now strongly concentrated. Under mutations, the distribution widens (and shifts to the left, since on average, there are more deleterious than beneficial mutations).
- At $t = 2$ selection step picks just such mutated sequences that increased fitness, and as a result the substitution rate $m_{1 \rightarrow 2}$ is high. The diversity $D(2) > D(1)$ and the fitness profile is less peaked than before.
- The mutation step does not impact the distribution of fitness, and we expect $m(3) < m(2)$. At later times, the substitution rate will decreases and the diversity increases until both reach their plateau values. Fluctuations in $m(t)$ follow inverted fluctuations in D at step $t - 1$.




Figure 10. Time dependence of fitnesses $P_{\text{nat}}(S_A)$ (left panel), $P_{\text{nat}}(S_C)$ (middle panel) and P_{int} (right panel) for the three possible protocols *evolve A* (blue), *evolve C* (orange) and *evolve AC* (green). Same values of stressor, $\beta_{\text{int}} = 20$. MC evolution.

5.2. Short- and long-term consequences of protein designability

LP structures S differ in how much they are designable, i.e. in how many sequences A have high folding probabilities, say, $p_{\text{nat}}(S|A) > 0.99$. Designability has been studied in the literature [4, 22], and it is known that S_C is more designable than S_A [16]. In practice, to assess the designability of a given structure, we can either evaluate the entropy σ of the Potts model inferred on the MSA (equation (11)) [23] or compute the mean identity (MI_d) of the MSA. To do so, we compute the consensus sequence (made of the most frequent amino acids, site after site), and define MI_d as the average number of sites carrying consensus amino acids. We note that the entropy of the (single-sequence) Potts model is bounded from above by $27 \log 20 \simeq 80.9$, corresponding to a totally unconstrained LP where each amino acid is randomly chosen.

The designability of these structures, and how they are affected by the introduction of the binding interaction constraint may help understand the evolutionary trajectories discussed above. We hereafter consider three different classes of binding constraints:

- introducing mutations only on S_A , keeping fixed S_C (labeled *evolve A*),
- introducing mutations only on S_C , keeping fixed S_A (labeled *evolve C*),
- introducing mutations on both S_A and S_C (labeled *evolve AC*).

We report the time dependence of P_{nat} and P_{int} for these three protocols for the same values of $\beta_{\text{nat}}, \beta_{\text{int}}$ in figure 10;. Protocol *evolve AC* is the most advantageous one, as it produces better configurations in terms of P_{nat} and P_{int} . Allowing both sequences to mutate gives rise to a larger number of possibilities to satisfy the constraints; this statement is also confirmed by estimation of the entropy $\sigma^{\text{fixed}}(S_C) \simeq 38.16$ computed on the MSA of *evolve C* protocol, and the conditional entropy $\sigma^{\text{cond}}(S_C) = \sigma(S_A, S_C) - \sigma(S_A) \simeq 43.55$ computed on the MSA of *evolve AC* protocol.

We report in figure 11 (left panel) the entropy σ together with the MI_d for several values of β_{int} and for S_A and S_C , evolved with, respectively, protocols *evolve A* and *evolve C*. The reported values are relative to end-point evolution, i.e. at equilibrium. As the stressor intensity β_{int} increases, the designability of both structures decrease, as it is harder for sequences to cope with the constraint on P_{int} . For all tested $\beta_{\text{int}} > 0$, the structure S_A is realized by more sequences than S_C , contrary to what happens for non-interacting structure ($\beta_{\text{int}} = 0$).

We again resort to the Potts model inferred from sequence data at different time points during evolution to characterize this phenomenon. In figure 11 (top right panel) we observe that the P_{nat} cross each other, signaling an inversion in the designability of the two structures, see figure 11 (bottom right panel). In other words, evolving a dimer surface through binding to

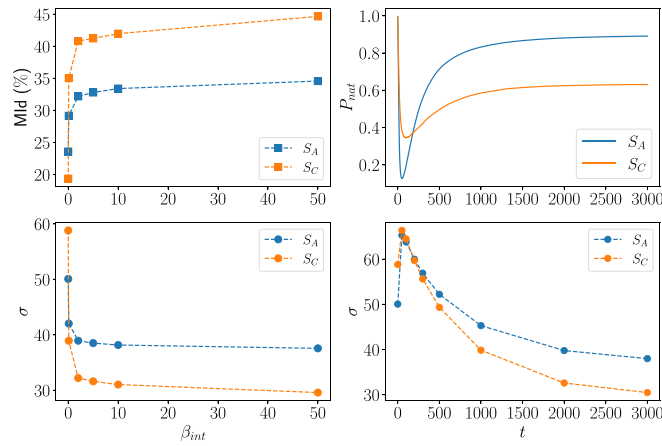


Figure 11. Left: entropy σ (bottom) and mean sequence identity Mid (top) for structures S_A and S_C (blue and orange, respectively) at equilibrium. Right: time dependence of $P_{\text{nat}}(S_A)$, $P_{\text{nat}}(S_C)$ (top) and of the entropy $\sigma(S_A)$, $\sigma(S_C)$ (bottom) under protocol *evolve A* and *evolve C*, respectively. Here σ at intermediate time is computed from the Potts model inferred on the MSA at that time. Same values of stressor, $\beta_{\text{int}} = 20$. MC evolution.

a fixed protein is more harmful in terms of fitness when the fixed protein has low designability. It is worth mentioning that this result is valid in general and does not depend on the particular choice of the two structures, as we validate in appendix C using different structures.

5.3. Microscopic mechanisms

5.3.1. Propagation of constraints on the interacting face. The consequences of having a binding selection pressure β_{int} during the evolution can be seen from the point of view of single protein structure and folding stability. At the beginning of the evolution ($t = 0$), each fold is very stable, which corresponds to large energy gap between the native structures and the competing structures (cf figure 12 and appendix D). As the evolution starts both proteins become less stable for some time, as the energy levels of competing structures get closer to the native folds. Eventually, as equilibrium is approached and the dimer is formed, a large energy gap is restored (albeit smaller than at $t = 0$). Figure 12 shows energy levels for the first competing structures of S_A during a typical MC evolution. The minimal energy gap is reached at similar evolutionary time points, regardless of the stressor strength β_{int} . This time corresponds to sequences where the amino acids on the interacting face have been mutated to favor the binding mode, thus destabilizing the native fold and making alternative competing folds more likely. Statistically, the time needed to update—at least once—all the nine amino acids on the interacting face is $\tau_{\text{typical}} \sim 9/p$, where p is the probability of proposing and accepting a mutation on a site of the interacting face. We estimate it from the MC simulation as the acceptance rate of mutations on interacting sites. It results $p \sim 0.12$ [24], giving τ_{typical} in agreement with the one observed in figure 6. The sequence logos [25], displayed in figure 13 for S_A , visually show enhanced mutability on the interacting face compared to other faces during the highly unstable transient. For example, positive (blue) and negative (red) charged residues on binding sites 1, 18 are highly conserved in the native structure, as they form an electrostatic bridge with amino acids of opposite charge. Non-binding sites 2, 25 also display an electrostatic mode in




Figure 12. Time evolution of the mean energy for protein sequences folded in S_A and in its first three competing structures (labeled as S_1 , S_2 and S_3 , see appendix D), namely the three structures with energies closest to S_A . All other structures have energy levels higher than S_3 . We note that the three competing structures share with the native folding S_A more than ten contacts, whereas an average random structure share just five contacts (see appendix D). Interestingly, there is an inversion in the order of competing structures going from $t = 0$ until the equilibrium state at $t = 3000$.




Figure 13. Sequence logos showing amino acids conservation on each site for protein sequences folded in S_A , averaged over the full MSA. On each site, the total height corresponds to the local conservation (see [25]) and the letter sizes correspond to their site-frequencies: the bigger a letter is, the more frequent that amino acid is in the MSA. We compare sites involved in binding (left) and non-binding sites (right) at $t = 0$ (top) and $t = 60$ (bottom) to show larger mutability on sites in interaction with the other protein. Sites 1–18 (or, equivalently, 2–25) form a crucial electrostatic mode in the native structure (top row), which is lost in correspondence to the minimum of P_{nat} (bottom row). Color code: blue for basic, red for acidic, green for Cysteine, black for hydrophobic and gold for aromatic ones.

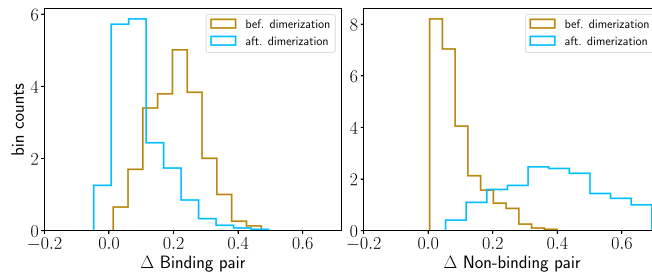


Figure 14. Distribution of scores Δ for two contact pairs, one involved in binding (left, pair 1–18 in figure 1) and the other not involved (right, pair 3–26 in figure 1), before and after dimerization takes place (gold and sky blue histograms, respectively). The scores are computed over the MSA of structure S_A at $\beta_{\text{int}} = 5.0$. Similar results hold for other contact pairs.

the native structure. At $t = 60$, the electrostatic contact 1–18 is depleted due to the ongoing binding with the other structure, while contact 2–25 remains less impacted. Enhanced mutability allowing for dimerization leads to residual frustration at single monomer level, which can be quantitatively evaluated along [26, 27]. Given a contact pair $i - j$, for each sequence \mathbf{A} in the MSA we compute $P_{\text{nat}}(S|\mathbf{A})$ and $P_{\text{nat}}(S|\mathbf{A}')$, where \mathbf{A}' differs from \mathbf{A} by replacing amino acids a_i, a_j with two other amino acids chosen uniformly at random. We then define the following score

$$\Delta_{i,j}(\mathbf{A}) = P_{\text{nat}}(S|\mathbf{A}) - \langle P_{\text{nat}}(S|\mathbf{A}') \rangle, \quad (17)$$

where the average is carried out over all possible pairs of amino acids on i, j excluding the actual one. The distribution of scores across the MSA of structure S_A is shown in figure 14 for a contact pair involved (i.e. binding pair) or not (i.e. non-binding pair) in the binding of the dimer. Upon dimerization, the scores of binding pairs are shifted to lower values: those pairs of sites now carry amino acids mostly optimized for binding, and almost any other amino-acid pair result in the same P_{nat} value. Therefore, such pairs are not crucial to maintain structural stability and we deem them as highly frustrated. Residual frustration at single monomer level is again the trademark of stability-affinity competition.

The presence of evolutionary trade-offs between P_{nat} and P_{int} also affects contacts' strength, assessed by the Frobenius norm computed as described in appendix A. In figure 15 we show the evolution of F_{ij} during the MC simulation, for both structures. For the LP folded in S_A the contacts belonging to the interacting face are the ones that suffer most from the binding constraint, as the corresponding Frobenius norms undergo large drop; conversely, the Frobenius norms associated to the back face (opposed to the interacting face) reach larger values after equilibration than the initial time point.

The Frobenius norm of structure S_C is smaller at any time point, suggesting that there is less coevolutionary pressure between sites due to the native design in structure S_C compared to S_A (cf figure 15); this is consistent with structure S_C being more designable than S_A .

5.3.2. Learning the binding mode through local fields. In this section we want to better characterize the selection pressure imposed by the binding mode through $\beta_{\text{int}} \neq 0$. We use the *evolve A* protocol and model the binding interactions as external fields h_i (exerted by the protein S_C) on the sites of the structure S_A belonging to the interacting face. To check whether this simple modeling approach, which neglects couplings J_{ij} within the interacting face is




Figure 15. Time evolution of $\langle F_{ij} \rangle$ averaged over contacts on interacting face (red), on back faces (purple) and over all 28 contacts (grey), for structure S_A (circle) and S_C (square). See color code of figure 1 for red and purple faces we are referring to. The norm is computed for a MSA at $\beta_{\text{int}} = 2.0$.

Table 1. Correlation coefficients of BM fields versus the MJ energy $\langle E(a_i, a_j) \rangle_{a_j}$ averaged over columns, for the model inferred on the *modified case* scenario. At $t = 0$ both fields of interacting and non-interacting sites are not (or negatively) correlated to the MJ values; after evolving the dimer surface, a strong positive correlation can be observed for sites on the interacting face.

	Interacting sites	Non-interacting sites
$t = 0$	0.05	−0.08
$t = 100$	0.92	−0.09

plausible we develop a slightly *modified protocol* of MC evolution, where at each Metropolis step, P_{int} in equation (6) is replaced with

$$P_{\text{int}}^M(S_1 + S_2 | \mathbf{A}_1, \mathbf{A}_2) \propto e^{-\mathcal{E}_{\text{int}}(\mathbf{A}_1, \mathbf{A}_2 | S_1 + S_2, 1)}. \quad (18)$$

In practice, discarding the denominator in equation (6) amounts to approximate P_{int} as a product over the sites of the binding interface. The competition between interfaces, usually leading to negative design, is therefore neglected.

This *modified framework* gives a qualitatively similar behavior for the time dependencies of $P_{\text{nat}}(S_A)$ and of P_{int} . We learn a Potts model from the sequence data in the same way as section 4.1, expecting that the inferred fields of sites on the interacting faces are highly correlated with the Myazawa–Jernigan (MJ) energy matrix. This is especially true when the stability of the protein is compromised, i.e. in the low P_{nat} region, because the binding selection pressure (fields) dominates over the internal contacts (couplings). We show in table 1 the average correlation coefficients between inferred fields and the MJ matrix for sites on the interacting face and not, at two different MC time steps.

To assess the validity of this field-based model, we consider a new learning procedure for the Potts model that consist in two steps. We collect in our dataset sequences having S_A as their native conformation ($P_{\text{nat}}(S_A) > 0.99$), and produce a smaller dataset with sequences for S_A all bounded to the same sequence folded in S_C . Inference of the Potts model works as follows:

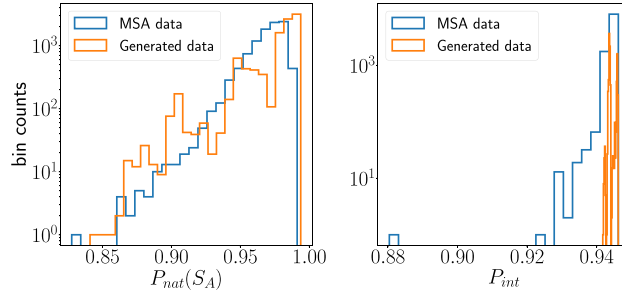


Figure 16. Distribution of $P_{\text{nat}}(S_A)$ (left) and P_{int} (right) for the training dataset and the generated MSA. To ensure high fitnesses on the new sequences, we sample the space with low temperature, $T = 0.2$.

- (i) we learn couplings and fields of the Potts model on the first dataset, having high $P_{\text{nat}}(S_A)$ and $P_{\text{int}} \simeq 0$;
- (ii) we now learn fields of the Potts model on the second dataset, using less data (25% of sequences used for the previous step), and with couplings J frozen to their values obtained in step (i).

In other words, we first have a background Potts model, able to model the distribution of S_A sequences, and then we infer only new fields to capture the features related to the interaction with the other protein sequence. We show in figure 16 that learning local fields is enough to generate good sequences in structure S_A that bind well with the given protein sequence in S_C . Inferred local fields are sufficient to reproduce the P_{nat} distribution (see figure 16, left panel), suggesting that couplings J_{ij} inferred at initial stage with no interaction are still meaningful. However, if one removes couplings at the single structure level attempting to use only fields to reproduce the bounded dimer distribution, the approach completely fails: it means that only inter-protein interactions can be modeled with local fields. Let us note that the generated sequences have mean identity (as defined in section 5.2) $\text{Mid} \sim 80\%$, compared to 60% for the training data, likely due to the presence of strong fields that force interacting sites to be very conserved. The generated sequences have an average Hamming distance of 11.75 amino acids to the ones used in the training.

6. Discussion

In this work we analyze the case of two LPs that evolve a dimer surface under a binding selective pressure, through a Monte Carlo Metropolis or population dynamics approach. In section 4 we study the equilibrium properties for different stressor strengths β_{int} , using the inference of a Potts model over the dimer MSA to reveal the effect of foldability-dimerization trade-offs on inter- and intra-protein couplings. In particular, we reproduce the majority of internal contacts of both structures, with contacts in S_C always less predictable than in S_A due to the higher designability of the former structure. As for the binding contacts, we succeed in reconstructing eight out of nine of them, missing always the central contact. In fact, the latter is shared in the functional and non-functional rotated modes, and our inferred model is not able to capture the *negative design* associated to such contact of competing modes [16]. We also

characterize the dimer interaction, assessing the quality of the function binding mode over the competing ones.

We then study in section 5 the ongoing evolutionary trade-offs during the dimer formation. Putting a selective pressure β_{int} forces the two LPs to explore unstable conformations in order to maximize the binding, at the expense of folding; this results in a temporary loss of stability, which depends on how rapidly the stressor β_{int} is applied during sequence evolution and how strong it is [28]. Indeed, the drop in the native fitnesses P_{nat} , has a strong out-of-equilibrium nature; whereas applying a smooth selection pressure relaxes the constraint over the two interacting faces and allows for a monotonically (decreasing) evolution of P_{nat} . This observation signals the existence of a minimal time for adaptation to new constraints, in agreement with experimental findings for bacterial evolution under stressful conditions [29–31].

Furthermore, we resort to the evolution of a non-clonal population (see section 5.1.3) to understand the interplay between the population size, the mutation rate per individual and the stressor strength. We see that population size here acts as an inverse temperature that sets the stringency of the fitnesses in evolution of stable complexes, which accounts for the lower values of β_{nat} , β_{int} used in the population dynamics compared to MC algorithms. This statement is especially clear when dealing with mono-clonal evolution, where the population size and the fixation probability of a mutation are related through [32]

$$P_{\text{fix}} = \frac{1 - \exp(-2s)}{1 - \exp(-2Ns)}, \quad (19)$$

where s is the selection coefficient and N the population size. Hence, here N is playing the role of an inverse temperature (see e.g. [33, 34]). Furthermore, we discuss how the selection pressure applied on the dimer evolution has consequences on the designability of the structures both at the short- and long-term level. We show that high stressor strengths reduce the diversity of MSA and that binding to a more designable folding (cf *evolve A* protocol) allows to find more optimal sequences.

Eventually, we discuss the microscopic mechanisms underlying the dimer formation showing that the binding interaction can be efficiently encoded in a Potts model on the single structure with local fields that mimic the selective pressure arising from the other structure.

In the future we plan to apply the discussed framework to real data where one can experience evolutionary trade-offs, e.g. bacteria strains evolving under two (or more) competing stressors.

Data availability statement

All data that support the findings of this study are included within the article (and any supplementary files).

Acknowledgment

E L is funded by the CNRS—University of Tokyo ‘80 Prime’ Joint Research Program.

Appendix A. Inferring Potts model

We use a Boltzmann machine (BM) to learn the probability distribution of the data and thus inferring the parameters of the Potts model. A BM is a probabilistic graphical model constituted of a single set of random variables $\mathbf{v} = (v_1, \dots, v_N)$ that interact within each other through

a coupling matrix \mathbf{J} and that are subject to local fields, hereafter called g_i . For our purpose, we need a set of $N = 54$ variables corresponding to the length of the two adjacent amino acids sequences, with each variable assuming $q = 20$ different states. For such BM the probability distribution of the variables set \mathbf{v} is exactly given by the Gibbs–Boltzmann probability at fixed temperature $T_{\text{inf}} = 1$ with energy

$$E(\mathbf{v}) = - \sum_i g_i(v_i) - \sum_{i < j} J_{ij}(v_i, v_j), \quad (\text{A1})$$

where the couplings J_{ij} and fields g_i set the mean and correlation of the variables v_i . Training a BM to infer its parameters consists in fitting numerically the distribution P of the data by maximizing the likelihood $\mathcal{L} = \langle \log P \rangle_{\text{data}}$. Taking the gradient of \mathcal{L} and setting it to zero for its maximization, bring us to solve the following problem (for a generic parameter θ of the model)

$$\nabla_{\theta} \mathcal{L} = - \langle \nabla_{\theta} E(\mathbf{v}) \rangle_{\text{data}} + \langle \nabla_{\theta} E(\mathbf{v}) \rangle_{\text{model}} \quad (\text{A2})$$

where $\langle \cdot \rangle_{\text{data}}$ stands for the expectation value over the data, while $\langle \cdot \rangle_m$ over the model. The gradient update thus consists of decreasing the energy of the data configurations while increasing the ones of the model. In our case, the gradient problem in equation (A2) can be turned into the set of equations

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial g_i} &= \langle v_i \rangle_d - \langle v_i \rangle_m \\ \frac{\partial \mathcal{L}}{\partial J_{ij}} &= \langle v_i v_j \rangle_d - \langle v_i v_j \rangle_m, \end{aligned} \quad (\text{A3})$$

which is a momentum-matching problem. The fitting procedure stops when the two expectation values match. Among the two right hand terms in equation (A3), it is easy to compute the average over the data as it can be done from the dimer MSA just once when the learning procedure starts; on the other hand computing the average over the model is a challenging task. Here, we use standard gradient descent (GD) with a varying learning rate in time and we make use of persistent contrastive divergence (PCD) algorithm with a fixed number of MC steps between each update for sampling data to compute the average over the model. To avoid over-fitting during training, we use a L_1^2 regularization term with strength λ_1^2 .

All in all, we tune the hyper-parameters for learning as follows

- Number of epochs = 150
- Learning rate = 0.005 (it decays with a 0.5 rate after half iterations)
- MC steps between each update = 5
- $\lambda_1^2 = 0.025$

Given the BM parameters, we can also compute the entropy σ of the model as the opposite of the log-likelihood averaged over the full MSA, i.e.

$$\sigma = - \langle \log P(\mathbf{v}) \rangle_{\mathbf{v}}, \quad (\text{A4})$$

which involves the numerical computation of the partition function \mathcal{Z} . The latter is intractable as it consists in summing over all configurations \mathbf{v} ; therefore, we estimate it with the annealed importance sampling (AIS) algorithm [35].




Figure 17. Couplings matrix $J_{ij}(a,b)$ for two specific sites in contact, i.e. $i = 3, j = 25$ of structure S_A , versus the MJ energy matrix $\mathcal{E}(a,b)$.

Appendix B. Dimer sequence generation

We generate new dimer sequences using the BM log-likelihood as the energy in the Gibbs sampling. For computational reasons after a long thermalization, we run 100 chains in parallel each of length 1000, and sample a dimer sequence each 500 steps. Sampling with the inferred Potts model is hard because we must have reconstructed the J_{ij} matrix very precisely; however the BM has learnt what are the sites in contact but it is struggling to understand which couple of amino acids is present on a given pair of sites. Effectively, if we look at the matrix J_{ij} for a given pair of sites in contact, most of its entry are zeros (i.e. the BM has never seen such pair of amino acids in the MSA), some are slightly different from zero and few of them have large values (i.e. those associated to polar amino acids). We show this trend in figure 17. Thus with so few peaks in the matrix, it is hard to generate and to ensure a great diversity of the sampling space, because the BM is exploring just one deep minima in the landscape. A solution for that would be to enlarge the effective depth of the MSA with more diverse dimer sequences, i.e. sampling with MC at higher temperature.

Appendix C. Additional dimeric assemblies

To further corroborate our results we perform an equivalent analysis using a different LP monomer, namely S_B as labeled in [16], that binds S_C . Binding of new different native structures yield the same scenario, suggesting that our results presented in the main text go beyond the specific conformation of S_A , S_C . Here, to speed up computation we restrict ourselves to 15000 MSA. We show results for the dimer $S_B - S_C$ in figure 18, where we plot the PPV for intra-structure contacts as we do in figure 4 for the dimer $S_A - S_C$. Since S_B is less designable than S_C , the picture here confirms that high designability negatively affects contact predictions of the native structure. As in figure 4, the worst predicted couplings involve the central site of the native structure and/or the central site of the binding face.

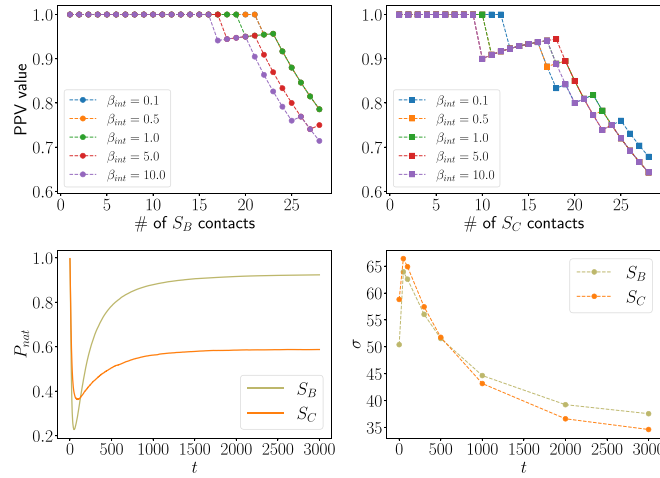


Figure 18. Top: PPV for structure S_B (top left panel, circles) and structure S_C (top right panel, squares) at different values of β_{int} . The MSA is made of 15 000 dimer sequences evolved for 3000 MC steps. Bottom: time dependence of $P_{nat}(S_B)$, $P_{nat}(S_C)$ (bottom left panel) and entropy $\sigma(S_B)$, $\sigma(S_C)$ (bottom right panel) under protocol *evolve B* and *evolve C*, respectively. Same values of stressor, $\beta_{int} = 20$. MC evolution.

Furthermore, we use structure S_B in interaction with S_C to validate our claim that evolution of a dimer keeping one of the two interface fixed is more or less harmful in terms of fitness depending on the designability of the fixed interface (see section 5.2). We thus designed, as we did for the dimer $S_A - S_C$, the two protocols *evolve B*, *evolve C* and computed the entropy at different time steps for both protocols. As in figure 10, we can see in figure 18 that the fitness P_{nat} is lower along protocol *evolve C*, where we keep fixed protein S_B that is less designable than S_C .

Appendix D. Competing structures

In figure 19 we report the folding of the first three competing structure with S_A , labeled in the main text as S_1 , S_2 , S_3 in figure 6. Such structures have been identified among the $\mathcal{N} = 10000$ representative foldings, as the ones having the smallest energy gap with the native structure. The energy for a given folding S has been measured as in equation (2), and averaged over 46000 sequences that fold in S_A . The random structure S_r has been randomly selected and has a large energy gap with the native structure. The more contacts structure S shares with the native folding, the more such structure will be considered as a competing one.

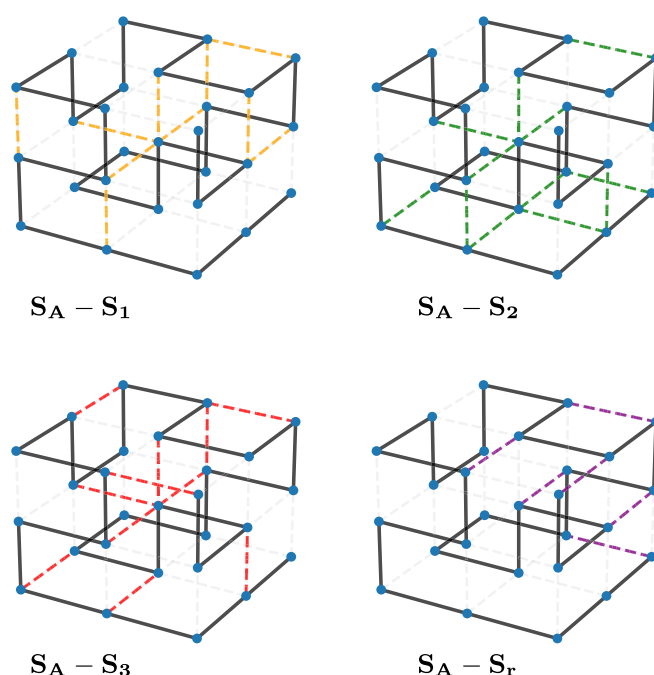


Figure 19. Representation of the competing foldings of S_A , as discussed in figure 6. Solid black line represents the backbone of structure S_A . Dotted grey shaded lines high-light contacts only present in the native folding S_A . Colored dotted lines show contacts of S_A common to its competing structures (orange, green and red refers to S_1, S_2, S_3 , respectively); purple dotted lines show contacts of S_A shared with a random structure S_r .

ORCID iDs

E Loffredo  <https://orcid.org/0009-0004-4882-8250>

O Peleg  <https://orcid.org/0000-0001-9481-7967>

References

- [1] Šali A, Shakhnovich E and Karplus M 1994 How does a protein fold *Nature* **369** 248
- [2] Lau K F and Dill K A 1989 A lattice statistical mechanics model of the conformational and sequence spaces of proteins *Macromolecules* **22** 3986
- [3] Shakhnovich E I and Gutin A M 1993 Engineering of stable and fast-folding sequences of model proteins *Proc. Natl Acad. Sci.* **90** 7195
- [4] Li H, Helling R, Tang C and Wingreen N 1996 Emergence of preferred structures in a simple model of protein folding *Science* **273** 666
- [5] Shakhnovich E and Gutin A 1990 Enumeration of all compact conformations of copolymers with random sequence of links *J. Chem. Phys.* **93** 5967
- [6] Peleg O, Choi J-M and Shakhnovich E I 2014 Evolution of specificity in protein-protein interactions *Biophys. J.* **107** 1686
- [7] Marchetti F, Capelli R, Rizzato F, Laio A and Colombo G 2019 The subtle trade-off between evolutionary and energetic constraints in protein-protein interactions *J. Phys. Chem. Lett.* **10** 1489
- [8] Tiana G and Broglia R A 2002 Design and folding of dimeric proteins *Proteins* **49** 82

- [9] Tiana G, Provasi D and Broglia R 2003 Role of bulk and of interface contacts in the behavior of lattice model dimeric proteins *Phys. Rev. E* **67** 051909
- [10] Weigt M, White R A, Szurmant H, Hoch J A and Hwa T 2009 Identification of direct residue contacts in protein–protein interaction by message passing *Proc. Natl Acad. Sci.* **106** 67
- [11] Cocco S, Feinauer C, Figliuzzi M, Monasson R and Weigt M 2018 Inverse statistical physics of protein sequences: a key issues review *Rep. Prog. Phys.* **81** 032601
- [12] Coucke A, Uguzzoni G, Oteri F, Cocco S, Monasson R and Weigt M 2016 Direct coevolutionary couplings reflect biophysical residue interactions in proteins *J. Chem. Phys.* **145** 174102
- [13] Morcos F, Pagnani A, Lunt B, Bertolino A, Marks D S, Sander C, Zecchina R, Onuchic J N, Hwa T and Weigt M 2011 Direct-coupling analysis of residue coevolution captures native contacts across many protein families *Proc. Natl Acad. Sci.* **108** E1293
- [14] Heo M, Maslov S and Shakhnovich E 2011 Topology of protein interaction network shapes protein abundances and strengths of their functional and nonspecific interactions *Proc. Natl Acad. Sci.* **108** 4258
- [15] Miyazawa S and Jernigan R L 1985 Estimation of effective interresidue contact energies from protein crystal structures: quasi-chemical approximation *Macromolecules* **18** 534
- [16] Jacquin H, Gilson A, Shakhnovich E, Cocco S and Monasson R 2016 Benchmarking inverse statistical approaches for protein structure and design with exactly solvable models *PLoS Comput. Biol.* **12** e1004889
- [17] Wu F-Y 1982 The Potts model *Rev. Mod. Phys.* **54** 235
- [18] Tubiana J, Cocco S and Monasson R 2019 Learning protein constitutive motifs from sequence data *Elife* **8** e39397
- [19] Marmier G, Weigt M and Bitbol A-F 2019 Phylogenetic correlations can suffice to infer protein partners from sequences *PLoS Comput. Biol.* **15** e1007179
- [20] Gandarilla-Pérez C A, Mergny P, Weigt M and Bitbol A-F 2020 Statistical physics of interacting proteins: impact of dataset size and quality assessed in synthetic sequences *Phys. Rev. E* **101** 032413
- [21] Nooren I M and Thornton J M 2003 Diversity of protein–protein interactions *EMBO J.* **22** 3486
- [22] England J L and Shakhnovich E I 2003 Structural determinant of protein designability *Phys. Rev. Lett.* **90** 218101
- [23] Barton J P, Chakraborty A K, Cocco S, Jacquin H and Monasson R 2016 On the entropy of protein families *J. Stat. Phys.* **162** 1267
- [24] This estimate is based on the short-time transient, while the overall acceptance rate is much lower
- [25] Sequence logos consist of a visual representation of conserved amino acids expressed in bits. On the i th site the total height is $C_i = \log_2 Q - \sum_{a=1}^Q \log_2 f_i(a)$, where Q corresponds to the total number of amino acids; the size of each letter is then proportional to its frequency $f_i(a)$
- [26] Ferreira D U, Hegler J A, Komives E A and Wolynes P G 2007 Localizing frustration in native proteins and protein assemblies *Proc. Natl Acad. Sci.* **104** 19819
- [27] Norbiato F, Seno F, Trovato A and Baiesi M 2019 Folding rate optimization promotes frustrated interactions in entangled protein structures *Int. J. Mol. Sci.* **21** 213
- [28] Šali A, Shakhnovich E and Karplus M 1994 Kinetics of protein folding: a lattice model study of the requirements for folding to the native state *J. Mol. Biol.* **235** 1614–36
- [29] Iwasawa J, Maeda T, Shibai A, Kotani H, Kawada M and Furusawa C 2022 Analysis of the evolution of resistance to multiple antibiotics enables prediction of the escherichia coli phenotype-based fitness landscape *PLoS Biol.* **20** e3001920
- [30] Yen P and Papin J A 2017 History of antibiotic adaptation influences microbial evolutionary dynamics during subsequent treatment *PLoS Biol.* **15** e2001586
- [31] Swings T, Van den Bergh B, Wuyts S, Oeyen E, Voordeckers K, Verstrepen K J, Fauvart M, Verstraeten N and Michiels J 2017 Adaptive tuning of mutation rates allows fast response to lethal stress in escherichia coli *Elife* **6** e22939
- [32] Kimura M 1962 On the probability of fixation of mutant genes in a population *Genetics* **47** 713
- [33] Sella G and Hirsh A E 2005 The application of statistical physics to evolutionary biology *Proc. Natl Acad. Sci.* **102** 9541
- [34] Rotem A *et al* 2018 Evolution on the biophysical fitness landscape of an RNA virus *Mol. Biol. Evol.* **35** 2390
- [35] Neal R M 2001 Annealed importance sampling *Stat. Comput.* **11** 125