Robustness and Efficiency in Inverse Protein Folding

Thomas M. Fink[†] and Robin C. Ball[§]
Theory of Condensed Matter, Cavendish Laboratory, Cambridge CB3 0HE, UK
[†]tmf20@cam.ac.uk [§]rcb1@phy.cam.ac.uk

Successful protein design is characterized by two criteria: thermodynamic robustness, the probability of occupation of the target conformation, and kinetic efficiency, the ability of the protein to quickly fold to its target state. We observe a conflict between robustness and efficiency upon variation of both the pair potential matrix and the designed sequence. We argue that marginal reduction in thermodynamic robustness can provide significant increase in kinetic efficiency, thereby allowing improved protein design.

A. Introduction

Protein folding consists of determining the ground state conformation of a given a sequence of amino acids. Inverse protein folding, or protein design, asks what sequence of amino acids possesses a given conformation as its ground state. The interest in the former problem lies not in the end, that is, the ground state conformation, but rather the means, the ability of the protein to achieve its native conformation in a time much less than that required to explore the entire conformational space. Unlike protein folding, the process of protein design occurs in nature on a time scale too long for our observation. Moreover, practical applications of protein design suggest that we do not necessarily wish to mimic nature in determining appropriate sequences; here we *are* interested primarily in the end.

While the protein folding problem possesses a unique solution, the solution to the inverse problem may be manifold. We thus characterize successful inverse protein folding as the design of sequences which robustly and efficiently fold to desired target conformations. Robustness necessitates a pronounced energy gap between the native and set of non-native states such that the target conformation is thermodynamically stable. We characterize efficiency by the mean first passage time from a denatured state to the target conformation. Both criteria may be expressed in terms of the corresponding energy landscape: stability coincides with a deep energy well while accessibility requires a landscape topography characterized by a folding funnel sloping toward the target.

It has been suggested that the former condition implies the latter, *i.e.*, that thermodynamically oriented selection of sequences solves the problem of kinetic accessibility as well [1]. We claim that while selecting for a pronounced minimum of the native state energy makes protein design possible, optimizing stability does not provide adequate accessibility. Moreover, we argue that in the region of sufficient thermodynamic stability, these two conditions are in conflict. In particular, marginal reduction in robustness can provide significant increase in efficiency.

B. Details of Model

We study the design and folding of model proteins N monomers in length, each of A possible species, constrained to a cubic lattice with nearest neighbor interactions. Proteins are designed by minimizing the target state energy with respect to sequence; we observe their folding kinetics via Monte Carlo simulations from self-avoiding walk (SAW) conformations to their respective target structures.

A protein conformation is designated $\Gamma = \{\mathbf{r_i}\}$, where $\mathbf{r}_i$ represents the lattice coordinates of monomer i, $i \in [1, N]$; the protein sequence is given by $S = \{\sigma_i\}$, where σ_i is the species of monomer i. Nearest neighbor monomers interact according to their species by the pair potential $U(\sigma_i, \sigma_j)$; we used the 20×20 one derived in [2] from the distribution of contact energies in native proteins.

The folding Hamiltonian may thus be expressed

$$E(S,\Gamma) = \frac{1}{2}\sum_{i,j=1}^{N} U(\sigma_i,\sigma_j)\Delta(\mathbf{r}_i - \mathbf{r}_j) - \sum_{i=1}^{N-1} U(\sigma_i,\sigma_{i+1}), \tag{1}$$

where $\Delta(\mathbf{r}_i - \mathbf{r}_j) = 1$ if $\mathbf{r}_i$ and $\mathbf{r}_j$ are nearest neighbors and 0 otherwise. The last term removes the energy of interaction along the backbone; since they cannot be avoided by the conformation, backbone interactions cannot influence the folding dynamics.

Given a desired target conformation Γ_{target}, we design a sequence S_{design} by annealing the sequence variables σ_i with respect to $E(S,\Gamma)$ while the conformation coordinates $\mathbf{r}_i$ remain quenched at Γ_{target} [1]. Folding of S_{design} is simulated on a cubic lattice at constant temperature by the Metropolis application of the moveset containing end bends, corner flips, and crank-shaft motions, where multiple occupation of lattice sites is forbidden. Such a move set is ergodic and generates SAW statistics at infinite temperature.

Starting from denatured (SAW) initial conformations, simulation continues until the energy $E(\Gamma_{\text{target}}|S_{\text{design}})$ of the target conformation is reached. The number of

attempted moves required for this to occur is the first passage time. We find that the energy of the designed sequence $S_{\rm design}$ embedded in the target conformation $\Gamma_{\rm target}$ is achieved only by $\Gamma_{\rm target}$ for compact (maximally bonded) structures, in accordance with [1], as well as for more open ones.

C. Variation of Pair Potential

For a given sequence, the probability that the ground state conformation Γ_0 is occupied is

$$P_{\Gamma_0} = \frac{e^{-\beta E_{\Gamma_0}}}{\sum_\Gamma e^{-\beta E_\Gamma}}, \tag{2}$$

where E_Γ is the energy of the sequence embedded in conformation Γ. Consider shifting the pair potential matrix U by a constant α,

$$U_\alpha(\sigma_i, \sigma_j) = U(\sigma_i, \sigma_j) + \alpha, \quad \alpha \in (-\infty, \infty), \tag{3}$$

where U, taken from [2], has near zero mean value,

$$\sum_{ij=1}^{N} U(\sigma_i, \sigma_j) = 0.018 \simeq 0; \tag{4}$$

note that varying α has no influence on sequence optimization for a fixed conformation. For a compact native state, the probability of occupation of the ground state then appears as

$$P_{\Gamma_0}(\alpha) = \frac{e^{-\beta\alpha B} e^{-\beta E_{\Gamma_0}}}{\sum_\Gamma e^{-\beta\alpha b} e^{-\beta E_\Gamma}}, \tag{5}$$

where B is the number of bonds in the compact target structure Γ_0 and b is the number of bonds in the conformation Γ. Since $b \leq B$ for all Γ, it follows that

$$\frac{\partial P_{\Gamma_0}}{\partial \alpha} = \beta(\langle b \rangle_T - B) P_{\Gamma_0} < 0, \tag{6}$$

where $\langle b \rangle_T$ is the average number of bonds over the thermal ensemble of conformations at temperature T.

From (6), the probability P_{Γ_0} that the native state Γ_0 is occupied is decreasing with α. Accordingly, negatively shifting the potential increases the energy gap between the native and the set of non-native states, thus increasing the thermodynamic stability of the desired conformation.

The extent to which this imposed increase in robustness affects folding efficiency can be investigated via simulation. We define efficiency as the mean first passage time $\langle \tau_{fp} \rangle$ of a given target conformation, that is, the average number of time steps necessary for a protein in a denatured state to fold to its target. We are interested in the mean first passage time as a function of the shift in the pair potential, α. However, the optimal folding temperature $T_{\rm opt}$ itself depends on α, so we must consider $\langle \tau_{fp}(\alpha, T) \rangle$.

As shown in Fig. 1, the mean first passage rate exhibits a peak in the α, T plane near $\alpha = 0$. Making the mean interaction α more negative promotes folding at higher temperatures, but decreases the attainable folding efficiency. As $\alpha \to -\infty$, the protein effectively quenches to a suboptimal compact state; to explore conformational phase-space, it must successively pass large energy barriers before expanding and quenching to other local minima. Thus, the gain in thermodynamic stability is made at the loss of folding efficiency. When $\alpha > 1$, the pair potential is purely repulsive and the protein ceases to fold; when $\alpha > 0.3$, the ground state ceases to be the target state.

D. Variation of Sequence

Improved protein design may be achieved by judiciously choosing a sequence, and hence an energy landscape, which more closely resembles a folding funnel about the target conformation than the ground state sequence, while retaining sufficient thermodynamic stability in the ground state conformation. This requires a sequence spectrum in which many sequences robustly fold to a single target conformation. Assuming the (conformational) ground state of the large majority of sequences is maximally compact, this implies the average number of sequences per compact conformation must be much greater than 1, that is

$$\left\langle \frac{N_{\rm S}}{N_\Gamma} \right\rangle \simeq \frac{A^N}{\kappa^N} \gg 1, \tag{7}$$

where the number of compact conformations of an N-mer follows $\Omega_c(N) \sim \kappa^N$ and $\kappa \simeq 1.9$ [3] on a cubic lattice. Accordingly, we require the number of amino acid species $A \geq A_c \simeq 3$ (which suggests that binary HP folding models are not sufficient for robust, efficient folding).

Generally (*e.g.*, [1]) and in the present work, the Hamiltonians used to optimize protein sequence and structure are equivalent. That is, sequence design consists of minimizing the energy, thereby maximizing thermodynamic stability, with the tacit assumption that stable sequences (deep minima) fold quickly (are funnel shaped). While empirical observations suggest a correlation between the two, it does not imply that such sequences fold *most* efficiently to their target conformations. Roughly put, the deepest wells need not be the most funnel oriented.

This conjecture may be tested as follows. Consider the set of all sequences which fold to a given compact target conformation; such a set is very large since in our case $A = 20 > A_c$. Plot each sequence, and hence landscape, on the efficiency-robustness phase space according to its

mean first passage time and its target (ground) state energy. In practice, it is not feasible to examine all $(\frac{A}{\kappa})^N$ sequences which map to a specific conformation; the set must be suitably sampled.

We prepared an ensemble of independently annealed 27-mer sequences trained to fold to a single compact $3 \times 3 \times 3$ target conformation. The thermodynamic stability of each sequence is approximated by the energy of the sequence in the target conformation, the ground state energy of the sequence; the actual Boltzmann occupation probability depends on the energy gaps between the native and non-native states. The mean first passage time is taken as the mean of N_f folding times from a denatured to the target conformation, where N_f is the number of times the sequence folds to its target in 2×10^7 Monte Carlo steps. Note that inversion of a rate gives better statistics to those sequences which fold more efficiently, which are the ones in which we are particularly interested.

A typical landscape scatter plot is shown in Fig. 2. The bottom colinear points represent the degenerate (sequential) ground state energy, due to rationality of the pair potential (of two significant figures) and certain sequence rearrangements. The spread of these mean first passage times is significant, notwithstanding the accompanying uncertainties. A sequence annealed to minimal energy is effectively sampled from this range. Of greater interest is the large fraction of sequences with higher target energies and lower mean first passage times. Choosing an arbitrary (sequential) ground state, approximately half, on average, of the sequences fold more efficiently, up to an order of magnitude more so. Disregarding small variations about the target conformation due to entropic considerations, all sequences shown spend the majority of their time in the (conformational) ground state.

E. Conclusion

We have shown, both by variation of the pair potential and the choice of trained sequence, that robustness of the target state and efficiency of folding are in conflict. In particular, maximal stability of the ground state conformation does not correspond to maximal accessibility. A marginal reduction in the robustness of the target conformation allows significant increase in the efficiency of folding.

While we have demonstrated that improved folding efficiency, while retaining sufficient stability, is achievable, we have not addressed how such an optimal sequence should be selected.

We present elsewhere [4] a novel method of kinetically favored sequence selection on the assumption that the widest possible funnel is that which least constrains the dynamics, which we propose is given by the conformations sampled during *unfolding* of the target conformation. Moreover, we provide arguments that lowering the energy of successive conformations, whether correlated, as in the case of a funnel, or independent, such as training to multiple targets, reduces the depth to which such conformations can be trained.

Alternatively, sequences can be annealed explicitly with respect to mean first passage time. Such a technique, impractical in itself due to computational constraints, may be made viable by the use of thermodynamic stability as a guide in choosing acceptable variations of the designed sequence; work to this end is ongoing and will be reported elsewhere.

The authors acknowledge useful communication with E.I. Shakhnovich, University of Harvard, and stimulating discussion with M.M.L. Khimasia, Cavendish Laboratory. T.M.F. is supported by a National Science Foundation Graduate Research Fellowship.

[1] E.I. Shakhnovich, Phys. Rev. Lett. **72**, 3907 (1994).
[2] S. Miyazawa and R. Jernigan, Macromolecules **18**, 534 (1985).
[3] W.J.C. Orr, Trans. Faraday Soc. **43**, 12 (1947).
[4] Thomas M. Fink and Robin C. Ball, submitted to Phys. Rev. Lett. (1996).

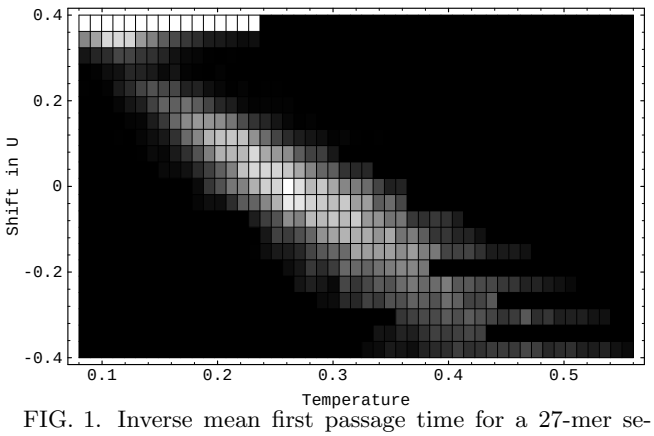

FIG. 1. Inverse mean first passage time for a 27-mer sequence embedded in a cubic conformation as a function of T and α. Each square represents the number of times the native state energy is reached ($N_{\rm ns}$) in 20×10^6 mcsteps. $N_{\rm ns}^{\rm max} = 164$ corresponds to white, 0 to black.

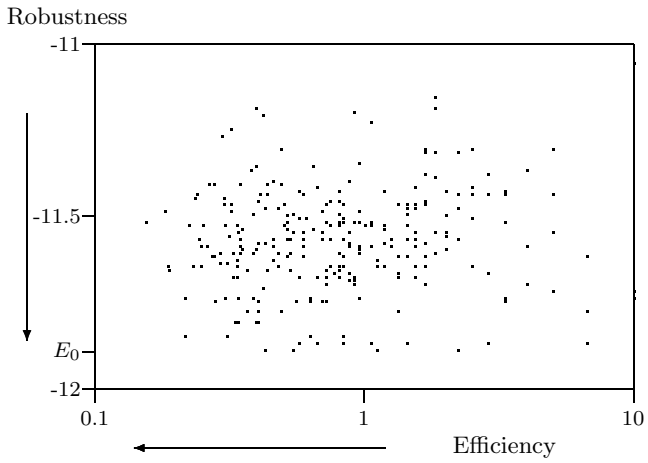

FIG. 2. Ensemble of 27-mer sequences independently trained to fold to a single compact target plotted in efficiency, robustness phase space. Efficiency is measured in millions of attempted Monte Carlo steps. Note that the most robust sequences do not correspond to the most efficient.