Side-chain and Backbone Flexibility in Protein Core Design

John R. Desjarlais and Tracy M. Handel*

Department of Molecular and
Cell Biology, University of
California, Berkeley
CA, 94720, USA

We have developed a computational approach for the design and prediction of hydrophobic cores that includes explicit backbone flexibility. The program consists of a two-stage combination of a genetic algorithm and monte carlo sampling using a torsional model of the protein. Backbone structures are evaluated either by a canonical force-field or a constraining potential that emphasizes the preservation of local geometry. The utility of the method for protein design and engineering is explored by designing three novel hydrophobic core variants of the protein 434 cro. We use the new method to evaluate these and previously designed 434 cro variants, as well as a series of phage T4 lysozyme variants. In order to properly evaluate the influence of backbone flexibility, we have also analyzed the effects of varying amounts of side-chain flexibility on the performance of fixed backbone methods. Comparison of results using a fixed *versus* flexible backbone reveals that, surprisingly, the two methods are almost equivalent in their abilities to predict relative experimental stabilities, but only when full side-chain flexibility is allowed. The prediction of core side-chain structure can vary dramatically between methods. In some, but not all, cases the flexible backbone method is a better predictor of structure. The development of a flexible backbone approach to core design is particularly important for attempts at *de novo* protein design, where there is no prior knowledge of a precise backbone structure.

© 1999 Academic Press

Keywords: protein design; backbone flexibility; hydrophobic core design; genetic algorithm; monte carlo

*Corresponding author

Introduction

A number of computational approaches have been developed for the optimization of side-chain sequence and/or structure for use in either protein design or comparative modeling. Almost all of these methods involve the use of a fixed backbone structure. Much of this work has focused on evaluation of the thermodynamic consequences of core mutations (Dahiyat & Mayo, 1996; Desjarlais & Handel, 1995; Hellinga & Richards, 1994; Lazar *et al.*, 1997; Lee & Levitt, 1991), prediction of hydrophobic core sequences (Dahiyat & Mayo,

1996; Desjarlais & Handel, 1995; Kono & Doi, 1994; Lazar *et al.*, 1997; Ponder & Richards, 1987), or prediction of core structure (Dahiyat & Mayo, 1996; Desjarlais & Handel, 1995; Harbury *et al.*, 1995; Hellinga & Richards, 1994; Kono & Doi, 1994; Lee & Subbiah, 1991; Tuffery *et al.*, 1991; Wilson *et al.*, 1993). A small number of these methods have now been confirmed by experimental characterization of core variants designed with the programs. In our own previous work, we demonstrated the successful design of several hydrophobic core variants of the proteins 434 cro and ubiquitin using a genetic algorithm called ROC (repacking of cores) which assumes a fixed backbone structure (Desjarlais & Handel, 1995; Lazar *et al.*, 1997; Johnson *et al.*, 1999). Mayo and co-workers have demonstrated similar success in the design of alternative core sequences for the proteins GCN4 and the B1 domain of protein G (Dahiyat & Mayo, 1996, 1997). Matthews and co-workers have used the original Ponder & Richards (1987) program to redesign a portion of the hydrophobic core of

Present address: J. R. Desjarlais, 408 Chandlee Laboratory, Department of Chemistry, Pennsylvania State University, University Park, PA 16802, USA.

Abbreviations used: ROC, repacking of cores; GA, genetic algorithm; MC, Monte Carlo; PLG, preservation of local geometry.

E-mail address of the corresponding author:
handel@paradise1.berkeley.edu

phage T4 lysozyme (Hurley *et al.*, 1992). The conceptual similarity of most of these programs implies that they would all, if tested, be at least partially successful in designing novel cores for known proteins.

The assumption of a fixed backbone is extremely powerful because it dramatically reduces complexity and computation time. An additional advantage is that it avoids reliance on existing force-field potentials to discriminate between favorable and unfavorable backbone structures. Despite these benefits, it has limitations. First, significant backbone relaxation has been clearly demonstrated in a number of crystallographic studies of hydrophobic core variants (Baldwin *et al.*, 1993; Eriksson *et al.*, 1992; Hurley *et al.*, 1992; Jackson *et al.*, 1993; Lim *et al.*, 1994). Thus, in core replacement studies, sequences found to be experimentally stable are sometimes predicted to be unallowed when using algorithms that assume a fixed backbone structure. Investigators have concluded that this is because the backbone cannot relax to accommodate the novel hydrophobic core residues, although this notion has been called into question by other results (Lee, 1996). A further limitation of fixed backbone approaches is in their proposed use for *de novo* protein design, since starting structures are, by definition, not precisely defined. A hydrophobic core optimization procedure that can begin with an approximate backbone structure and explore the surrounding structural space would clearly be advantageous. Harbury and co-workers have reported a flexible backbone approach to the design of hydrophobic cores for coiled coil motifs (Harbury *et al.*, 1995, 1998). The method employs parametrization of coiled-coil geometries to explore potential backbone structures. Making use of symmetry relations intrinsic to coiled-coil systems, the authors are able to reduce the problem to one which permits exhaustive searches of side-chain rotamer combinations. Impressive success is demonstrated with this method, both in terms of prediction and design. However, since most proteins do not contain such a high degree of symmetry, the general applicability of the method remains to be determined. Su & Mayo (1997) have examined the influence of backbone movement on core design by systematically manipulating the relative orientations of super secondary structural elements in the B1 domain of protein G (GB1). They created a series of fixed backbone structures which were then fed into a core design algorithm (Dahiyat & Mayo, 1996). Using this approach, they demonstrated some tolerance of core sequence prediction to small but significant perturbations in structure, and successfully designed a series of novel GB1 variants which have thermal stabilities and other properties similar to those of the wild-type protein.

Here, we describe the development and characterization of a first generation program for the simultaneous optimization of hydrophobic core sequence/structure and backbone structure. It

differs from the methods mentioned above in that it allows explicit backbone flexibility at every position in the structure, and its use is not limited to systems with a high degree of symmetry. Since this is a modification of the original program ROC that adds capability for backbone flexibility, we named it SoftROC. We have used one version of SoftROC to design three new core variants of the bacteriophage protein 434 cro. These variants and several previously designed variants of 434 cro (Desjarlais & Handel, 1995) are then evaluated using either fixed or flexible backbone models and an Amber-based potential for evaluation of backbone energies. Based on these results, a second flexible backbone version was developed which uses the preservation of local backbone structure as a criterion for evaluating alternative backbone models. Use of this constraining potential results in a modest correlation between calculated and experimental stabilities, and appears to be a promising approach for protein engineering efforts.

The generality of the method was tested by its application to a series of hydrophobic core variants of phage T4 lysozyme, some of which are known to have significant backbone displacements (Baldwin *et al.*, 1993). For the 434 cro and T4 lysozyme variants, we compare the prediction of relative stabilities, side-chain rotamer structures, and backbone structures using fixed backbone *versus* constrained, flexible backbone models. The results establish that use of a fixed backbone model can lead to a good correlation of calculated *versus* experimental stabilities, provided that side-chains are given a wide range of flexibility. We find that when backbone flexibility is incorporated, this correlation remains but is slightly diminished. However, the core side-chain structures predicted using either approach are often significantly different.

Results

Program description

The current version of the flexible backbone program was developed and optimized for redesigning the hydrophobic cores of existing proteins. Therefore, an inherent assumption in developing the program was that backbone structures for novel core sequences would be similar to the backbone structure of the parent protein. In general, knowledge of the structure of the original backbone is incorporated directly into the algorithm, either by appropriately defining the dihedral angle regions to be sampled, including an explicit constraint toward the original structure, or both. Sampling of structures is accomplished by varying dihedral angles alone, starting from a torsional model derived from the starting structure (see Methods).

The first major stage of the program consists of a genetic algorithm (GA) optimization of a population of structures (typically 200), all of which

have different backbone and side-chain structures. The initial population of structures is generated by random sampling of side-chain sequence and/or conformation, and sampling of backbone dihedral angles within a predefined range about the dihedral angles of the template structure. We find that adding random continuous adjustments of up to $\pm 3^\circ$ to each backbone dihedral angle yields a sufficiently diverse starting population of backbone structures which, on average, deviate by approximately 1.4 Å r.m.s.d. from the crystal structure backbone. This is larger than the range of deviations typically observed when proteins relax to accommodate core mutations (Baldwin *et al.*, 1993; Lim *et al.*, 1994).

Starting with this initial population of structures, evolution of the population toward lower energy structures is encouraged using a combination of selective recombination and random mutation. The energies of each model in a given population are evaluated according to the defined potential. In all current versions of SoftROC the Amber force-field (Weiner *et al.*, 1984) with the OPLS non-bonded parameters (Jorgensen & Tirado-Rives, 1988) is used to evaluate the energies of side-chain to main-chain and side-chain to side-chain interactions. In one version, this same force-field is used to evaluate the energies of the main-chain as well. A second version, described in detail below, utilizes a novel pseudo-potential which constrains the main-chain structure to retain similarity to the template structure. According to the calculated energies and a temperature factor, each model is assigned a Boltzmann probability which governs the number of times that the model will recombine with randomly chosen partners from the population. Thus, models with lower energy will recombine more frequently and pass their information to the next generation. The models are encoded as linear arrays of side-chain and main-chain dihedral angles, so recombination between models is accomplished by exchanging segments of the arrays of two models. We use uniform recombination, in which each array element of an offspring model is randomly chosen from either of the parent models. The use of a Boltzmann probability based on the energy is useful in that a variable temperature factor can be used to influence the convergence rate of the population. After all recombinations are carried out for a cycle of evolution, the resulting population is subject to a small amount of random mutation of side-chain identity (for core design), large or small changes in side-chain dihedral angles, or small changes in backbone dihedral angles.

Performance tests of SoftROC revealed that although a GA is superior during early stages of a simulation, a Monte Carlo (MC) procedure is preferable in the refinement stages. Therefore, after the GA stage of the algorithm is complete, the lowest energy structure of the final GA population is subject to a simulated annealing procedure with Metropolis sampling and a decreasing temperature

factor. Different types of dihedral angle sampling, recombination temperatures, etc. are found to be optimal at different stages of the simulation. These, and a more detailed description of the SoftROC algorithm are included in the Methods.

Designing novel hydrophobic cores for 434 cro

We used SoftROC to design new hydrophobic core variants of 434 cro. These variants add to the set which were previously created using an early version of ROC (Desjarlais & Handel, 1995). The earlier studies demonstrated the success of our fixed backbone approach for the design of novel core sequences which have significant differences from the native core. Nevertheless, fixed backbone approaches have been criticized for their tendency to predict false negatives: some hydrophobic core variants which are predicted to be unacceptable by fixed backbone criteria are actually found experimentally to have near wild-type stability. To address this issue, the first designs were deliberate attempts to generate core sequences that were experimentally stable and calculated to have low energies using SoftROC, yet were predicted to be unallowed by the fixed backbone version of the program. This was encouraged by forcing a disruptive mutation at position 59 (L59F), and designing the remainder of the core sequence to accommodate the mutation using SoftROC. Designs were created with a version of SoftROC that uses the Amber potential to calculate energies of all parts of a model. Several runs were performed and the two most promising sequences from over 20 runs were chosen for design. The sequences each have six substitutions relative to the native core (including the forced L59F mutation), and differ from each other at two positions. As noted in Table 1, we will refer to these proteins as SD-F59A and SD-F59B. Another protein with this constraint, F59C, was designed with an earlier version of the program which contained an error in the energy calculations. This protein was characterized as well. A third designed sequence, SD-7, was generated with no restrictions on the amino acid residue in position 59: in this experiment all of the core residues are free to vary between Val, Ile, Leu, Phe, and Ala. Finally, we constructed a 434 cro variant called cro-rep which replaces the natural core residues with those found in 434 repressor, a close homologue of 434 cro. The new sequences are shown in Table 1.

Genes were constructed by standard methods and the proteins were expressed and purified as described (Desjarlais & Handel, 1995). Thermal denaturation profiles of these designs are shown in Figure 1, along with representative designs and controls from earlier work (Desjarlais & Handel, 1995). The results indicate that the new designs are stable, well-packed proteins with cooperative folding transitions. As with the fixed backbone designs, the new proteins have stabilities comparable to that of the positive control cro* (wild-type

Table 1. Hydrophobic core variants of 434 cro

	2	6	13	20	24	26	31	34	45	48	52	54	58	59	T_m (°C)
Wild-type	L	L	L	L	A	V	I	I	L	I	L	C	W	L	
cro*	L	L	L	L	A	V	I	I	L	I	L	V	W	L	56
R-6	I	L	I	L	A	V	L	I	I	L	L	V	W	L	-
SD-F59A	I	I	L	L	A	V	I	L	L	L	L	V	W	F	48
SD-F59B	I	I	L	L	A	I	I	L	L	I	L	V	W	F	47
F59C	I	L	L	L	A	I	I	L	I	I	L	V	W	F	20
SD-7	F	I	L	L	I	I	L	L	L	I	L	V	W	L	48
cro-rep	I	V	L	L	V	V	I	L	L	L	L	V	W	L	55

All positions which were subject to change are listed. Substitutions relative to the wild-type sequence are shown in bold. Variants with an SD prefix were designed with an algorithm that incorporates backbone flexibility. The proteins cro* and R-6 were designed and characterized in previous work (Desjarlais & Handel, 1995). Variants SD-F59A and B were designed with the constraint that position 59 contain a phenylalanine. Variant F59C was designed with an earlier version of the algorithm that contained an error in the energy calculations. Variant SD-7 was designed with no constraints at any position, other than hydrophobicity (Val, Ile, Leu, Phe, or Ala). The cro-rep variant replaces the native core of 434 cro with core residues from 434 repressor.

cro with a C54V mutation, Table 1), and are substantially more stable than a protein with a randomly designed hydrophobic core of composition similar to that of the native protein (R-6, Table 1). This protein is completely unfolded. The results demonstrate the success of our flexible backbone approach in designing novel core sequences which are compatible with a given fold. The newly designed sequences are not included in a large set of designed sequences derived from the original ROC algorithm, suggesting that the addition of backbone flexibility has expanded the range of predicted sequences.

The cro-rep protein has a thermal unfolding transition and melting temperature almost identical to that of cro* (Figure 1). The core replacement constitutes a total of six substitutions relative to wild-type cro (five relative to cro*). This result suggests, not surprisingly, that hydrophobic cores of proteins with similar structure are generally transferable. For these two proteins, the backbone

structures superimpose with an r.m.s.d. of 0.72 Å (Mondragon *et al.*, 1989a,b).

Predicting stabilities: flexible versus fixed backbones

In developing a flexible backbone approach to hydrophobic core evaluation, we sought to alleviate several limitations of a fixed backbone approach, one being the overly restrictive prediction of allowable core sequences. However, the incorporation of backbone flexibility may introduce a new problem of alternative sequence predictions being too permissive: core substitutions calculated to be acceptable may actually be found to be experimentally destabilized. In order to investigate these issues, we compare here the calculated stabilities of the previously designed variants of 434 cro (Desjarlais & Handel, 1995) and the new variants, using both the fixed and flexible backbone programs. Three versions of the fixed backbone algorithm are used to illustrate some important differences between fixed backbone approaches. Some fixed backbone approaches use a small number of statistically derived side-chain rotamers with rigid dihedral angle values (Ponder & Richards, 1987). A more common approach includes additional discrete rotamers at defined increments, for example at $\pm 15^\circ$ or $\pm$ one standard deviation from the statistically derived rotamer value (Dahiyat & Mayo, 1996; Tuffery *et al.*, 1991). For our analysis, we have also used a version of our fixed backbone algorithm that incorporates a wide range of flexibility into the side-chains, with fine sampling of side-chain dihedral angles (5°) at almost all possible angles (Lazar *et al.*, 1997). An analysis of side-chain flexibility effects is crucial at this stage, since the flexible backbone method allows complete side-chain flexibility as well. We compare the ability of each version to predict the relative stabilities of the set of 434 cro variants. Consistency in the comparison is ensured by using the same force-field parameters for all versions. The nomenclature is as follows: fixed backbone with rigid rotamers (fxd + rgd), fixed backbone

Figure 1. Thermal denaturation of 434 cro hydrophobic core variants. Unfolding of each variant was monitored by circular dichroism at 222 nm. The data is represented as the apparent fraction of unfolded protein F_u , after curve fitting of the raw data.

Figure 2. Energetic predictions using a fixed backbone and varying amounts of side-chain flexibility. Each set of calculated energies is plotted against the melting temperatures (T_m , $^{\circ}\text{C}$) of a series of 434 cro hydrophobic core variants. (a) Stabilities calculated using a fixed backbone (fxd) with a rigid (rgd) library of side-chain rotamers. (b) Stabilities calculated using a fixed backbone with a typical degree of side-chain flexibility (rgd $\pm 15^{\circ}$). (c) Stabilities calculated with a fixed back-

bone and a nearly continuous (cont) set of side-chain orientations. The negative T_m representing the stability of the R-6 variant is derived from an extrapolation of a series of T_m values determined in water and 40% (v/v) glycerol for several of the variants (R-6 is completely unfolded in water). Error bars in this Figure and figures 3 and 6 represent the standard deviation among multiple independent simulations.

with semi-flexible rotamers (fxd + rgd $\pm 15^{\circ}$), fixed backbone with continuous rotamers (fxd + cont), and flexible backbone with continuous rotamers, using the Amber potential for backbone evaluation (flx + AMB).

Figure 2 shows plots of calculated *versus* experimental stabilities using the three fixed backbone implementations of ROC, for all hydrophobic core variants of 434 cro. We use melting temperature as an experimental measure of stability. Since the calculated energies, are purely conformational energies and melting temperatures are not equivalent to free energies, a perfect correlation is not expected. However, some correlation is anticipated if core design programs are to be considered useful. As seen in Figure 2(a), there is almost no correlation between experimental melting temperatures (T_m) and the energies calculated using a fixed backbone with rigid rotamers. The extreme outliers are due to both of the new core variants of 434 cro which contain a Phe in position 59. When a more typical kind of rotamer set is used (rigid rotamers $\pm 15^{\circ}$), the correlation is somewhat improved, but remains modest (Figure 2(b)). When full flexibility is allowed in the side-chains, the correlation is dramatically improved, as shown in Figure 2(c).

Unfortunately, the incorporation of backbone flexibility, using the Amber/OPLS potential to evaluate backbone energies, leads to a complete absence of correlation between calculated and observed stabilities (data not shown). A particularly disconcerting result is that the randomly designed 434 cro variant (R-6), which is experimentally found to be severely destabilized, is calculated to have a stability close to that of the wild-type, and higher than several experimentally stable variants. Clearly this version of the flexible backbone program is unable to sufficiently discriminate between acceptable and unacceptable core sequences. While the reasons for this are not clear, the backbone structures of the flx + AMB models generally have more local structural deviation than is experimentally observed for proteins with core substitutions (Baldwin *et al.*, 1993; Lim *et al.*, 1994). This suggests that better parametrization or more information is required for accurate prediction of low energy backbone structures, and it calls into question the use of a standard force-field in this context. Consequently, we attempted to create a novel backbone potential that uses information

directly from the parent structure, yet still allows a significant degree of flexibility, as described below.

Constructing a backbone potential

Analysis of the crystal structures of hydrophobic core variants of T4 lysozyme (Baldwin *et al.*, 1993) and lambda repressor (Lim *et al.*, 1994) reveals an interesting phenomenon: the response of the backbone structure to core replacement rarely involves large local deviations of structure. Instead, backbone relaxation is achieved by combinations of very minor local adjustments throughout the chain. The accumulation of these minor local effects over large segments of the chain can result in significant displacement of atoms that belong to residues distant in sequence. In fact, the extent of local deviation of structure for most variants is generally no greater than that seen in comparing structures of the same protein solved in different space groups (Zhang *et al.*, 1995).

We have created a virtual energy term that is designed to preserve local structure in a manner similar to that seen for the T4 lysozyme and lambda repressor hydrophobic core variants. The energy term consists of small windows of backbone r.m.s.d. between the model and the wild-type crystal structure, summed over the whole structure. Instead of a computationally costly canonical r.m.s.d. calculation, which compares inter-structure atomic distances after optimal superposition, we use a similar criterion which compares the differences in intra-structure atomic distances (see Methods). A window size of five residues works well, and is a natural choice for helical proteins. This energy term replaces the energy of the backbone structure that was calculated in the original version according to an Amber-derived potential (side-chain to side-chain and side-chain to backbone energies are still calculated with Amber/OPLS). We will refer to this energy term as the PLG (preservation of local geometry) potential.

The PLG potential was incorporated into Soft-ROC and energies were calculated for all of the 434 cro variants (flx-PLG). The PLG energy is included in the final total calculated energies, with a weight of 2.0 relative to the other energy terms. The calculated energies correlate with the experimentally determined stabilities, as shown in Figure 3. The data are displayed as mean energy values taken over a large number of independent simulations, with error bars indicating the standard variance of the calculated energy. The presence of a correlation suggests that the knowledge-based PLG energy is more useful than an Amber-based energy for evaluating alternative backbone structures for core replacement studies.

While inclusion of the PLG energy term results in a dramatic improvement of the flexible backbone algorithm, several aspects require further development. The quality of the correlation is decreased relative to that of the fixed backbone calculations using continuous rotamers. The corre-

Figure 3. Energetic predictions using a flexible backbone and a backbone potential that emphasizes the preservation of local geometry (PLG).

lation is clearly dominated by the R-6 data, and is almost absent without it, demonstrating that this method is at present only appropriate for coarse evaluation of protein cores. Moreover, the variance of the calculated energies is significant, and greater than that seen for fixed backbone methods, diminishing the confidence in the evaluations. These problems can hopefully be addressed by a systematic exploration of the PLG approach, with varying window sizes and weights, along with additional efforts to improve the convergence properties of the algorithm.

Predicting structures: flexible versus fixed backbones

The incorporation of backbone flexibility into hydrophobic core packing algorithms not only influences the calculated stabilities, but also the predicted structures. The effect on structure prediction is particularly important for the application of these methods in the field of comparative modeling, where the structure of one protein is predicted using the structure of a close homologue as a template. If the sequences are reasonably homologous, the backbone structures are likely to be very similar, but not identical. For modeling of low homology sequences, the ability to deal with the imperfect relationship between the structures of the target and template proteins remains a challenge (Chung & Subbiah, 1995; Dunbrack *et al.*, 1997; Eisenhaber *et al.*, 1995). The set of core variants of 434 cro can be thought of as a simple case of comparative modeling, where the sequences all have an extremely high degree of homology. Here, we point out significant differences between the predicted structures for the new designs when using either fixed backbone or flexible backbone methods. Since use of the PLG potential leads to a significantly better correlation of calculated and

experimental energies, we restrict our discussion to predictions derived using PLG to calculate backbone energies. Fixed backbone predictions are from the version of ROC which allows full flexibility of side-chains (fxd + cont).

Table 2 shows the number of side-chain rotamers of each 434 cro variant predicted differently by the two methods. When the lowest energy models generated using either method are used for comparison (Table 2, columns 2 and 3), as many as six side-chains are predicted to be in different rotamers. However, the two approaches yield highly correlated estimates of relative stabilities. This apparent contradiction can be rationalized by analysis of the different models. As an example of this, a comparison of model structures of the SD-F59B variant, derived using either method, is shown in Figure 4. The two predicted structures are clearly very different. Analysis of the side-chain dihedral angles predicted by either method reveals that six χ_1 angles differ by at least 40° . The major affect on the different predictions is most likely due to the introduction of a Phe residue at position 59. In the fixed backbone model, Phe59 is predicted to have a χ_1 angle of 233° , compared to 188° in the flexible backbone model. Since our potential includes the torsional angle term of the Amber potential, this represents a calculated energy difference of 3.7 kcal/mol for the Phe side-chain alone. The fixed backbone models of both F59 variants have similar χ_1 angles predicted for Phe59, with both predicting it to be close to a perfectly eclipsed configuration. This indicates that with the fixed backbone constraint, a Phe at position 59 causes a large amount of strain. As described earlier (see Figure 2), when only a small number of rigid side-chain rotamers are allowed for model evaluation (fxd + rgd), both of these models are calculated to have extremely high energies. In the fixed backbone model with flexible side-chain structures (fxd + cont), this high strain is relieved partially,

Figure 4. Model structures of the SD-F59B variant, based on the lowest energy models generated by two methods. The model predicted using a fixed backbone assumption and a continuous set of side-chain orientations is shown in blue. The flexible backbone model is shown in red. Note that although both methods predict a similar amount of relative strain for this variant, the structural rationale for the strain is very different. Six of the side-chains are predicted to be in different orientations and are highlighted with thicker side-chain representations.

but not completely, by forcing Phe59 to adopt an eclipsed rotamer conformation. This conformation may be propagated to the other side-chains, causing the large number of different side-chain orientations predicted throughout the model structures. Finally, in the flexible backbone model, the strain is almost completely relieved by allowing the backbone to adjust to the presence of Phe59. This, in turn, allows Phe59 to adopt a more favorable rotamer configuration. Any strain inherent in the pre-

Table 2. Differences in prediction of side-chain orientations when using fixed (cont) or flexible backbone (PLG) methods

Variant	cont/PLG (lowest E) χ_1	cont/PLG (lowest E) χ_2	cont/cont (all) χ_1	cont/cont (all) χ_2	PLG/PLG (all) χ_1	PLG/PLG (all) χ_2	cont/PLG (all) χ_1	cont/PLG (all) χ_2
cro*	2	1	0.60 ± 1.21	0.44 ± 0.50	1.52 ± 1.03	0.61 ± 0.66	1.24 ± 1.13	0.63 ± 0.71
SD-F59A	5	1	2.96 ± 1.59	0.84 ± 0.88	2.68 ± 1.64	1.08 ± 0.77	3.46 ± 1.64	1.42 ± 0.80
SD-F59B	6	0	2.58 ± 1.69	0.31 ± 0.56	2.52 ± 1.22	0.96 ± 0.72	4.03 ± 1.16	0.46 ± 0.63
F59C	1	2	0.93 ± 0.69	0.60 ± 0.86	1.20 ± 1.00	1.37 ± 1.13	1.31 ± 1.09	1.68 ± 0.95
SD-7	5	0	2.02 ± 1.54	0.18 ± 0.39	3.18 ± 1.99	0.63 ± 0.63	4.28 ± 1.29	0.32 ± 0.57
cro-rep	3	1	0.00 ± 0.00	0.53 ± 0.50	0.91 ± 0.67	0.36 ± 0.53	2.50 ± 0.50	0.52 ± 0.58
D-5	3	1	1.00 ± 0.74	0.00 ± 0.00	2.10 ± 1.29	0.46 ± 0.70	3.06 ± 1.19	1.09 ± 0.43
D-8	2	0	0.00 ± 0.00	0.00 ± 0.00	2.73 ± 1.38	1.07 ± 0.80	4.07 ± 1.34	0.65 ± 0.78
D-7	3	0	3.16 ± 1.83	0.24 ± 0.43	1.25 ± 0.99	0.85 ± 0.71	3.18 ± 1.46	0.75 ± 0.66
M-5	3	0	2.31 ± 1.59	0.33 ± 0.48	2.05 ± 1.40	0.00 ± 0.00	3.29 ± 1.32	0.32 ± 0.61
R-6	6	0	2.80 ± 1.69	1.20 ± 1.03	1.07 ± 0.75	0.79 ± 0.66	4.54 ± 1.16	0.61 ± 0.74

Differences in χ angles of greater than 40° are reported. χ_2 columns correspond to differences only in the χ_2 angle, with the same cutoff. Comparisons between the lowest energy model generated using either method are shown in the first two columns. All other comparisons report the mean and standard deviation of differences between all models of a given set. Comparisons involving a single set (cont/cont, etc.) indicate the variance of side-chain prediction within one method. In general, the variance between models generated using different methods is greater than that between models using the same method.

sence of a Phe at position 59 is now presumably taken up in the energetic cost of allowing the backbone to move from its more favorable wild-type structure. Hence both methods appear to be capable of assessing the amount of energetic strain caused by core substitution, but the structural explanation of how the strain is relieved is very different.

Statistics collected for comparisons between all models generated by either method reveal that the tendency to predict different side-chain structures is general. As shown in Table 2, the number of χ_1 and χ_2 angles predicted differently between models generated using a single method is on average less than that predicted differently between models generated using different methods.

The cro-rep variant, the result of an experiment in which the native core of 434 cro is replaced with that of 434 repressor, represents an interesting situation for structure prediction. Examination of the crystal structures of wild-type 434 cro and repressor proteins reveals a single phenylalanine residue which is completely buried in the repressor structure, yet fully solvent-exposed in cro, as shown in Figure 5(a). Because of the potential of this Phe to be buried in the cro homologue, this residue was included in all calculations of variants, including cro-rep. Figure 5(b) shows the lowest energy fxd + cont and flx-PLG models for cro-rep. While the fixed backbone model fails to predict the burial of this residue, the flexible backbone model predicts a repressor-like orientation of this and several other core side-chains affected by the difference. This result alone makes a case for the inclusion of backbone flexibility in comparative modeling efforts.

T4 Lysozyme hydrophobic core variants

Because structures have not yet been determined for the 434 cro variants, we can only surmise the quality of the predictions from the different versions of ROC. However, an extensive number of studies have been conducted which examine the effect of hydrophobic core substitutions on the stabilities and structures of T4 lysozyme. In order to extend our analysis to include a different model system, we analyzed the differences in predictive ability of ROC and SoftROC using a set of eight variants of T4 lysozyme. The structures and stabilities of these variants have been reported by Matthews and co-workers (Baldwin *et al.*, 1993). Table 3 describes the core substitutions and measured stabilities of these proteins.

In general, comparisons of the different versions of ROC and SoftROC for the T4 lysozyme variants parallel the results obtained for the 434 cro variants. Plots of calculated *versus* experimental energies are shown in Figure 6. Using a fixed backbone and rigid rotamers again results in the calculation of extremely high energies for some of the T4 variants. This problem is alleviated by allowing a wide range of flexibility for core side-chains, and results in a significantly improved correlation between calculated and experimentally determined stabilities. A correlation is also observed when including backbone flexibility with the PLG potential.

As is the case of the 434 cro variants, many of the predicted structures differ depending on whether the fixed or flexible backbone method is used, although to a lesser extent than seen for the cro variants. For all T4 lysozyme variants, however, crystal structures have been determined

Figure 5. Modeling of cro-rep. Structures are displayed as alpha carbon traces with hydrophobic core side-chains. (a) Comparison of the crystal structures of 434 cro (black) and 434 repressor (red). Note the difference in burial of Phe44 (arrow). (b) Comparison of fixed (black) and flexible (red) backbone models of cro-rep. The fixed backbone method predicts a cro-like orientation of Phe44, while the flexible backbone method predicts a core structure more similar to that of repressor.

Figure 6. Correlation of calculated and experimental stabilities of T4 lysozyme variants. Each set of calculated energies is plotted against the reported free energies of T4 lysozyme variants (Baldwin *et al.*, 1993). (a) Stabilities calculated using a fixed backbone (fxd) with a rigid (rgd) library of side-chain rotamers. (b) Stabilities calculated with a fixed backbone and a nearly continuous (cont) set of side-chain orientations. (c) Stabilities calcu-

(Baldwin *et al.*, 1993), enabling a more meaningful comparison of the accuracy of each method. Figure 7 shows superpositions of models of the wild-type and a representative variant derived using either ROC (fxd + cont) or SoftROC (flx-PLG). In general, orientations of only a small number of side-chains are predicted differently for each variant, with some of the incorrect predictions involving large side-chains such as tryptophan. For the wild-type protein, the ROC prediction is more accurate than that of SoftROC, as shown in Figure 7(a) and (b). This is not unexpected because ROC is given the actual backbone structure of the wild-type protein. The major discrepancy in prediction involves a rotamer change of Leu133 from the most statistically favorable leucine rotamer to the second most favorable (McGregor *et al.*, 1987). For variant II, SoftROC gives the more accurate prediction, as shown in Figure 7(c) and (d). The difference in prediction involves a major change in the orientation of tryptophan 153 and a correlated change of Leu129 into a statistically unfavorable rotamer in the ROC model (McGregor *et al.*, 1987). In this case, the incorrect predictions are presumably due to the inability of ROC to accommodate the large tryptophan residue in the context of the wild-type backbone structure. The strain caused by this forces the mispredicted side-chains into suboptimal rotamer structures.

A summary of the success of ROC and SoftROC in predicting side-chain structure is shown in the form of prediction penalties as a function of backbone r.m.s.d. between the crystal structure of each variant and that of wild-type lysozyme (Figure 8). The penalties are defined as the sum of the number of atoms misplaced by each dihedral angle that is incorrectly predicted by at least 40° . This normalizes the effect of incorrect side-chain predictions on residues of different size (e.g. a mispredicted tryptophan residue has a much more serious impact on the structure than a mispredicted valine residue). The results of the comparison are interesting. For those variants whose actual structures are very similar to that of wild-type T4 lysozyme (backbone r.m.s.d. less than $0.3 \text{ \AA}$), the fixed backbone assumption clearly yields superior predictions. However, for variants whose structures are significantly different from the wild-type structure, the flexible backbone method represents an improvement. It is not surprising that fixed backbone ROC performs less well as structures deviate from the wild-type. A more mysterious result is the observation that prediction penalties for SoftROC are highest for structures that are most similar to wild-type. In fact, SoftROC incorrectly predicts a small number of side-chain rotamers of

lated with a flexible backbone (flx) and a constraining potential (PLG) and continuous side-chain orientations. For each plot, the linear correlation coefficient (R) is reported.

Table 3. Hydrophobic core variants and experimentally determined folding free energies of T4 lysozyme (Baldwin *et al.*, 1993)

	121	129	133	149	153	ΔG^0 (kcal/mol)
Wild-type	L	A	L	V	F	0.0
I	A	M	L	V	L	1.1
II	M	L	M	I	W	1.3
III	I	L	M	V	W	1.3
IV	A	M	L	I	W	1.4
V	I	W	M	V	F	1.4
VI	A	V	M	V	L	2.3
VII	M	A	V	V	L	2.5
VIII	A	V	A	V	L	3.5

the wild-type protein. Further analysis of the side-chain structure prediction properties of SoftROC will be required to determine if this is a general phenomenon, although it seems unlikely. The current results suggest that predictions from both methods should be accepted as possibilities, perhaps to be analyzed further using other methods.

Analysis of the SoftROC-predicted backbone structures reveals that at this point, high-resolution prediction of backbone adjustment is not possible. However, there does appear to be some potential for the use of SoftROC with PLG for the prediction of gross backbone adjustment. Variant number V has the largest deviation in structure (0.63 Å r.m.s.d.) from wild-type lysozyme, and therefore is expected to most readily reveal trends in the accuracy of backbone predictions. The SoftROC model of variant V is significantly more similar to the

crystal structure than are models of all other variants, suggesting that there is some specificity in the predictions of SoftROC. The r.m.s.d. for backbone atoms between the variant V SoftROC model and variant V crystal structure is 0.77 Å, whereas the closest of the other models to the same crystal structure is 0.94 Å, and the most dissimilar is 1.17 Å. While this observation suggests some promise for prediction of large changes in backbone structure, it is important to stress that because of the nature of the PLG potential, SoftROC currently has no ability to predict fine structural changes in response to core substitution.

Discussion

We have described the development of a new core optimization algorithm which simultaneously

Figure 7. Comparison of model and crystal structures of T4 lysozyme variants. (a) Superposition of a fixed backbone model and crystal structure of wild-type T4 lysozyme. (b) Superposition of a flexible backbone model and crystal structure of wild-type T4 lysozyme. (c) Superposition of a fixed backbone model and crystal structure of T4 lysozyme variant III. (d) Superposition of a flexible backbone model and crystal structure of T4 lysozyme variant III. Crystal structures are displayed as black and models are displayed as blue. Atoms affected by a misprediction of a side-chain dihedral angle of more than 40° are displayed as red.

Figure 8. Side-chain prediction penalties for T4 lysozyme and variants. Penalties for each model are plotted against the structural similarity between each variant crystal structure and that of wild-type T4 lysozyme. Prediction penalties are defined as the sum of all atoms affected by misprediction of side-chain dihedral angles of more than 40° (red atoms in Figure 7). A penalty of zero implies a very good prediction, as shown in Figure 7(a). Open circles denote penalties for flexible backbone models, and open squares denote penalties for fixed backbone models.

samples side-chain sequence/structure and backbone structure. The development of such an algorithm is important for many reasons. First, backbone adjustment in response to amino acid substitution is well documented (Baldwin *et al.*, 1993; Eriksson *et al.*, 1992, 1993; Hurley *et al.*, 1992; Jackson *et al.*, 1993; Karpusas *et al.*, 1989; Lim *et al.*, 1994). Any algorithm developed for the prediction of core sequence or structure should aim to predict all of the consequences of the core substitutions, including backbone adjustment, especially since this can have a significant effect on function (Lim *et al.*, 1994; Lim & Sauer, 1989). Second, comparative modeling efforts rely on the use of a template backbone structure from one protein to predict the structure of a homologous protein. Proteins with significant homology have similar, but not identical structures, so the ability to adjust a backbone structure to accommodate a different sequence should contribute to these efforts. Third, the useful assumption of a fixed backbone structure has been widely criticized, yet a full analysis of the effects of including backbone flexibility has not been performed. Finally, if core optimization algorithms are intended for use in *de novo* design experiments, the problem remains that target design structures are by nature not precisely defined. The concurrent selection of backbone topology and side-chain identity is one of the major goals of computational protein design.

The assumption of a fixed backbone structure is often called into question, because it is known that backbones do relax to accommodate amino acid

substitutions, and it has been suggested that the lack of backbone flexibility in existing algorithms is the source of their limited success in predicting allowed core sequences (Baldwin *et al.*, 1993; Hurley *et al.*, 1992; Lim *et al.*, 1994). Our analysis suggests that the inclusion of proper amounts of flexibility at any level can improve the predictive performance of core optimization algorithms. Even with the assumption of a fixed backbone structure, the level of side-chain flexibility can have dramatic effects on the prediction of structure or stability. Two of the 434 cro variants described in this work (SD-F59A and SD-F59B) were designed specifically to address this issue. When these variants are evaluated using a set of rigid side-chain rotamers, a correlation between calculated and experimental energies is non-existent, and their sequences are strongly predicted to be unallowed. Experimentally, however, both proteins have stabilities close to that of the wild-type protein, and are significantly more stable than a variant with a randomly chosen hydrophobic core. When a full range of orientations is allowed for each side-chain, this picture changes remarkably, and a very good correlation between calculated and experimental stabilities emerges. Importantly, this is still in the context of a fixed backbone. These results suggest that if a fixed backbone algorithm using a rigid rotamer library performs poorly for predicting allowed sequences, criticism of the fixed backbone assumption itself may be partially misdirected. This view is similar to conclusions published in a separate study on predicting the energetics of single alanine to valine substitutions in T4 lysozyme (Lee, 1996).

The further addition of backbone flexibility is not without effect. Although energetic correlations derived with a version of SoftROC which uses the PLG constraining potential are similar to those derived with ROC and flexible side-chains, the predicted structures can be significantly different. For the 434 cro variants studied herein, up to six side-chain rotamers are predicted differently when using fixed or flexible backbone methods. However, without determining the structures of these variants, assessing which method yields superior predictions is purely speculative. Differences are smaller for the set of T4 lysozyme variants studied. We demonstrate that in some cases, especially for variants which contain more extreme substitutions, the inclusion of explicit backbone flexibility leads to an improvement in structure prediction, and therefore may be a useful tool for comparative modeling efforts. However, since neither method has been shown to be superior for all cases, both should be used conservatively for analysis of protein variants. The prediction of side-chain structure may not be crucial for core replacement studies, but for success in future attempts to design catalytic sites, accurate prediction of side-chain orientations will be key.

A disappointing finding of our study is that when an Amber/OPLS potential is used in SoftROC to evaluate backbone structures there is no

correlation between calculated and experimentally determined stabilities. Although several possibilities exist to explain this, we believe the limitation is in the accuracy of the Amber/OPLS potential for evaluating the local structure of a protein. For core variants whose structures have been determined, there is a remarkable preservation of local structure. However, the Amber/OPLS potential leads to a significantly larger adjustment of local structure than is typically observed, suggesting that the backbone structures are underconstrained. The substitution of the Amber potential with the constraining potential, PLG, improves the predictive ability of SoftROC and results in a more realistic deviation of local structure.

There is considerable room for improvement of the methodology we have begun to develop. At this point, a clear improvement of predictive ability as backbone flexibility is included has not been demonstrated. However, several aspects of the results indicate that backbone flexibility should be included, and that it can potentially be addressed with current sampling methods. Further exploration of different potentials and run-time parameters will prove useful for improving the methods reported here. For instance, the weight of the PLG potential term used in the present calculations appears to be approximately correct, but has not been systematically optimized. The size of the r.m.s.d. window used for the PLG calculation can also be explored further. Finally, the convergence properties of the method require improvement, as indicated by the error bars in Figures 3 and 6(c).

The use of constraint-based potentials for evaluating backbone structures may be the most conservative approach for the incorporation of backbone flexibility into core optimization or comparative modeling algorithms (Sali *et al.*, 1995). In our PLG potential, we placed particular emphasis on the preservation of local geometry since it is strongly conserved in sets of variants whose structures have been analyzed. However, experimentation with other constraints is warranted in order to increase the accuracy of predicted backbone adjustments. The addition of hydrogen bonding constraints may improve performance, and will be particularly important for extending these modeling techniques to β -sheet containing proteins.

Methods

Fixed backbone algorithms

The ROC programs used in this report have been described (Desjarlais & Handel, 1995; Lazar *et al.*, 1997). Briefly, a fixed backbone structure and a list of core residues is read into the program, and side-chain identities and/or conformers for each core position are sampled from a rotamer library. A genetic algorithm is used to optimize the core sequence and structure *via* multiple rounds of recombination and mutation. In the current version, the rotamer library is that reported by Tuffery *et al.* (1991). For the ROC fxd + rgd runs, these rotamers

alone are used as input. For the ROC fxd + rgd $\pm 15^\circ$ runs, this library is expanded by including all rotamers with dihedral angles $\pm 15^\circ$ away from these. For the ROC fxd + cont runs, these rotamers are used as input, but all rotamer values with dihedral angles within $\pm 50^\circ$ are also included, at 5° increments. Between 10 and 20 runs were performed for analysis of each individual variant to establish convergence. Each run of ROC fxd + rgd or ROC fxd + rgd $\pm 15^\circ$ takes approximately one minute on a Silicon Graphics with an R4400 processor. Each run of ROC fxd + cont takes approximately one hour.

Flexible backbone algorithms

SoftROC was developed to use torsional models of protein backbones and side-chains. Alternative structures are sampled by modifying dihedral angles instead of directly changing cartesian coordinates. All bond lengths and angles are fixed at the equilibrium values defined by the Amber force-field (Weiner *et al.*, 1984). In order to incorporate knowledge of the parent backbone structure into a simulation, a torsion space model must first be created for the parent structure. Because crystal structures do not contain fixed bond lengths and angles, building a torsion space model using dihedral angles derived directly from the crystal structure yields a model which is significantly different from the crystal structure, due to propagation of small differences along the chain. Therefore, an initial Monte Carlo optimization of the backbone torsional angles is performed to yield a torsional model which is nearly equivalent to the parent crystal structure. This model serves as the template for generation of the initial population of structures to be used in the genetic algorithm. The backbone structure of each member of the initial GA population is generated by randomly choosing ϕ and ψ angles within 3° of those from the parent template. Although this seems minor, these differences result in significantly different backbone structures when propagated throughout the whole chain. The ω values are also randomly sampled, but within a range of 1° from the template values. Initial side-chain identities and/or structures are generated by choosing a rotamer from the rotamer library (Weiner *et al.*, 1984), and adding random amounts up to $\pm 30^\circ$ to each dihedral angle.

In a typical run of SoftROC, a GA population of 200 is subject to 300 rounds of evolution consisting of uniform recombination and random mutation. Uniform recombination is used to generate an offspring model from two parent models by randomly choosing information for the structure of each backbone and side-chain position from either parent model. This is straightforward when using torsional models, because of a simple transfer of dihedral angle information. Recombination is selective, with weight given to lower energy models. This is accomplished by defining a normalized Boltzmann probability for each model according to its calculated energy, temperature, and the population size. Each model is assigned a number of recombinations, defined as the product of its probability and the population size. It then recombines with that number of other models picked randomly from the population. The use of a Boltzmann probability factor allows temperature control of the convergence of the population. To prevent premature convergence, a decaying temperature factor is included in the Boltzmann probability. The decay rate is defined by:

$$T = (T_0 - T_f) \left(1 - \frac{c}{t_c} \right) + T_f$$

where T_0 and T_f are the initial and final temperatures, respectively, c is the current GA cycle, and t_c is the total number of GA cycles (typically 300). For all runs described in the current work, T_0 was 10,000 K and T_f was 100 K. After recombination, models are subject to random mutations of three types: (1) rotamer classes and/or side-chain identities are mutated at an average frequency of 3%. New rotamers are chosen, and random amounts of $\pm 10^\circ$ are added to each dihedral angle. (2) small changes in side-chain dihedral angles within 5° of the previous dihedral angle are made at an average frequency of 10%. (3) Small changes in backbone dihedral angles up to 3° are allowed at an average frequency of 0.5%. At 280 cycles, rotamers and side-chain identities are frozen by turning off the rotamer class mutations (type 1 above). At this point, the maximum backbone dihedral angle change is reduced to 1° .

SoftROC was developed over time to include the genetic algorithm for early phases of the run, and a Monte Carlo (MC) with metropolis sampling and simulated annealing procedure for refinement. The MC phase of the algorithm begins with the lowest energy model from the final GA population. No rotamer class changes are allowed in the MC phase. Side-chain dihedral angle changes of up to 3° are allowed at an average frequency of 10%. Backbone dihedral changes of up to 1° are allowed at an average frequency of 2%. Metropolis sampling with simulated annealing is used to generate and propagate alternative models, with temperature decaying linearly from 200 K to 0 K throughout the run.

Runs were done in three sets, using a trimming procedure to eliminate higher energy side-chain orientations. After a first set of ten runs, side-chain rotamers not found in any of the models are omitted from the input rotamer library. The same trimming is performed after a second set of five runs. A final set of five runs is then performed with the remaining rotamers. The lowest energy model from all 20 runs was chosen as the prediction model for that variant. The lowest energy model typically comes from the second or third set of runs. Each single run of SoftROC takes approximately six hours on a Silicon Graphics with an R4400 processor.

Models

Model representations for the variants studied included all core side-chains, and a small number of non-core side-chains. All core side-chains were allowed equivalent ranges of motion, defined by the algorithm used. Non-core side-chains were fixed into their original dihedral angle states as defined by the template crystal structure. Core side-chains for 434 cro include positions 2, 6, 13, 20, 24, 26, 31, 34, 44, 45, 48, 52, 54, 58 and 59, using the numbering as defined in the crystal structure (2CRO) (Mondragon *et al.*, 1989b). Non-core side-chains included in the 434 cro modeling are R5, R9, R10, M15, Q17, E37, and W60. For all other side-chains, only the α and β carbon atoms were maintained.

For T4 lysozyme modeling, only residues 81 through 157 of the carboxy-terminal domain were considered. The crystal structure of the wild-type protein (2LZM) was used as a template (Weaver & Matthews, 1987). Core residues for T4 lysozyme are 87, 91, 99, 102, 106, 111, 114, 118, 121, 126, 129, 133, 138, 149, 150, and 153. The only non-core side-chain included in the modeling of T4 lysozyme variants was that of Arg95. For all other

side-chains, only the α and β carbon atoms were included.

Energy calculations

For ROC or SoftROC, energies of all side-chain to backbone, side-chain to side-chain, and intra-side-chain interactions are evaluated using the Amber potential (Weiner *et al.*, 1984) with the OPLS non-bonded parameters (Jorgensen & Tirado-Rives, 1988). Bond stretching and angle bending terms were not calculated, since the equilibrium values are used to generate the torsional model. In one version of SoftROC, this same potential is used for evaluating backbone energies. In a second version of SoftROC, backbone energies are estimated using the PLG potential, which is calculated as follows. The total PLG virtual energy consists of a sum of values from a five residue window scanned across the whole model structure. For each five residue window, the distances between atoms in the first residue of the window and atoms within all other residues in the window are evaluated and compared to the equivalent atom pair distances within the original template structure. The square of each difference is calculated and summed for all pairs. The square roots of this value for each five residue window are summed over the whole model structure to yield the PLG virtual energy. This measure is analogous to summing up a series of canonical r.m.s.d. values, but is independent of optimal superposition of structures and is more computationally expedient. Because of the CPU-intensive nature of the simulations, the weight of the PLG energy relative to the other energetic terms has not been systematically optimized. A value of 2.0 was used for the calculations presented here after determination that a value of 1.0 was not sufficiently constraining.

Protein expression and characterization

Proteins were expressed and purified as described (Desjarlais & Handel, 1995). Conditions for circular dichroism and analysis of thermal denaturation data have also been described (Desjarlais & Handel, 1995).

Acknowledgments

T.M.H. is an NSF Young investigator and J.R.D. was a fellow of the Jane Coffin Childs Memorial Fund for Medical Research, fellowship number 61-948, during the course of this study. We gratefully acknowledge these institutions for support.

References

- Baldwin, E. P., Hajiseyedi, O., Baase, W. A. & Matthews, B. W. (1993). The role of backbone flexibility in the accommodation of variants that repack the core of T4 lysozyme. *Science*, **262**, 1715-1718.
- Chung, S. Y. & Subbiah, S. (1995). The use of side-chain packing methods in modeling bacteriophage repressor and cro proteins. *Protein Sci.* **4**, 2300-2309.
- Dahiyat, B. I. & Mayo, S. L. (1996). Protein design automation. *Protein Sci.* **5**, 895-903.
- Dahiyat, B. I. & Mayo, S. L. (1997). Probing the role of packing specificity in protein design. *Proc. Natl Acad. Sci. USA*, **94**, 10172-10177.

- Desjarlais, J. R. & Handel, T. M. (1995). *De novo* design of the hydrophobic cores of proteins. *Protein Sci.* **4**, 2006-2018.
- Dunbrack, R. L., Jr, Gerloff, D. L., Bower, M., Chen, X., Lichtarge, O. & Cohen, F. E. (1997). Meeting review: the second meeting on the critical assessment of techniques for protein structure prediction (CASP2), Asilomar, California, December 13-16, 1996. *Fold. Design*, **2**, R27-42.
- Eisenhaber, F., Persson, B. & Argos, P. (1995). Protein structure prediction: recognition of primary, secondary, and tertiary structural features from amino acid sequence. *Crit. Rev. Biochem. Mol. Biol.* **30**, 1-94.
- Eriksson, A. E., Baase, W. A., Zhang, X. J., Heinz, D. W., Blaber, M., Baldwin, E. P. & Matthews, B. W. (1992). Response of a protein structure to cavity-creating mutations and its relation to the hydrophobic effect. *Science*, **255**, 178-183.
- Eriksson, A. E., Baase, W. A. & Matthews, B. W. (1993). Similar hydrophobic replacements of Leu99 and Phe153 within the core of T4 lysozyme have different structural and thermodynamic consequences. *J. Mol. Biol.* **229**, 747-769.
- Harbury, P. B., Tidor, B. & Kim, P. S. (1995). Repacking protein cores with backbone freedom: structure prediction for coiled coils. *Proc. Natl Acad. Sci. USA*, **92**, 8408-8412.
- Harbury, P. B., Plecs, J. J., Tidor, B., Alber, T. & Kim, P. S. (1998). High-resolution protein design with backbone freedom. *Science*, **282**, 1462-1467.
- Hellinga, H. W. & Richards, F. M. (1994). Optimal sequence selection in proteins of known structure by simulated evolution. *Proc. Natl Acad. Sci. USA*, **91**, 5803-5807.
- Hurley, J. H., Baase, W. A. & Matthews, B. W. (1992). Design and structural analysis of alternative hydrophobic core packing arrangements in bacteriophage T4 lysozyme. *J. Mol. Biol.* **224**, 1143-1159.
- Jackson, S. E., Moracci, M., elMasry, N., Johnson, C. M. & Fersht, A. R. (1993). Effect of cavity-creating mutations in the hydrophobic core of chymotrypsin inhibitor 2. *Biochemistry*, **32**, 11259-11269.
- Johnson, E. J., Lazar, G. A., Desjarlais, J. R. & Handel, T. M. (1999). Solution structure and dynamics of a hydrophobic core variant of ubiquitin. *Struct. Fold. Design*, in the press.
- Jorgensen, W. L. & Tirado-Rives, J. (1988). The OPLS potential functions for proteins. Energy minimizations for crystals of cyclic peptides and crambin. *J. Am. Chem. Soc.* **110**, 1657-1666.
- Karpusas, M., Baase, W. A., Matsumura, M. & Matthews, B. W. (1989). Hydrophobic packing in T4 lysozyme probed by cavity-filling mutants. *Proc. Natl Acad. Sci. USA*, **86**, 8237-8241.
- Kono, H. & Doi, J. (1994). Energy minimization method using automata network for sequence and side-chain conformation prediction from given backbone geometry. *Proteins: Struct. Funct. Genet.* **19**, 244-255.
- Lazar, G. A., Desjarlais, J. R. & Handel, T. M. (1997). *De novo* design of the hydrophobic core of ubiquitin. *Protein Sci.* **6**, 1167-1178.
- Lee, C. (1996). Testing homology modeling on mutant proteins: predicting structural and thermodynamic effects in the Ala98 → Val mutants of T4 lysozyme. *Fold. Des.* **1**, 1-12.
- Lee, C. & Levitt, M. (1991). Accurate prediction of the stability and activity effects of site-directed mutagenesis on a protein core. *Nature*, **352**, 448-451.
- Lee, C. & Subbiah, S. (1991). Prediction of protein side-chain conformation by packing optimization. *J. Mol. Biol.* **217**, 373-388.
- Lim, W. A. & Sauer, R. T. (1989). Alternative packing arrangements in the hydrophobic core of lambda repressor. *Nature*, **339**, 31-36.
- Lim, W. A., Hodel, A., Sauer, R. T. & Richards, F. M. (1994). The crystal structure of a mutant protein with altered but improved hydrophobic core packing. *Proc. Natl Acad. Sci. USA*, **91**, 423-427.
- McGregor, M. J., Islam, S. A. & Sternberg, M. J. (1987). Analysis of the relationship between side-chain conformation and secondary structure in globular proteins. *J. Mol. Biol.* **198**, 295-310.
- Mondragon, A., Subbiah, S., Almo, S. C., Drott, M. & Harrison, S. C. (1989a). Structure of the amino-terminal domain of phage 434 repressor at 2.0 Å resolution. *J. Mol. Biol.* **205**, 189-200.
- Mondragon, A., Wolberger, C. & Harrison, S. C. (1989b). Structure of phage 434 Cro protein at 2.35 Å resolution. *J. Mol. Biol.* **205**, 179-188.
- Ponder, J. W. & Richards, F. M. (1987). Tertiary templates for proteins. Use of packing criteria in the enumeration of allowed sequences for different structural classes. *J. Mol. Biol.* **193**, 775-791.
- Sali, A., Potterton, L., Yuan, F., van Vlijmen, H. & Karplus, M. (1995). Evaluation of comparative protein modeling by MODELLER. *Proteins: Struct. Funct. Genet.* **23**, 318-326.
- Su, A. & Mayo, S. L. (1997). Coupling backbone flexibility and amino acid sequence selection in protein design. *Protein Sci.* **6**, 1701-1707.
- Tuffery, P., Etchebest, C., Hazout, S. & Lavery, R. (1991). A new approach to the rapid determination of protein side-chain conformations. *J. Biomol. Struct. Dyn.* **8**, 1267-1289.
- Weaver, L. H. & Matthews, B. W. (1987). Structure of bacteriophage T4 lysozyme refined at 1.7 Å resolution. *J. Mol. Biol.* **193**, 189-199.
- Weiner, S. J., Kollman, P. A., Case, D. A., Singh, U. C., Ghio, C., Alagona, G., Profeta, S. & Weiner, P. (1984). A new force-field for molecular mechanical simulation of nucleic acids and proteins. *J. Am. Chem. Soc.* **106**, 765-784.
- Wilson, C., Gregoret, L. M. & Agard, D. A. (1993). Modeling side-chain conformation for homologous proteins using an energy-based rotamer search. *J. Mol. Biol.* **229**, 996-1006.
- Zhang, X. J., Wozniak, J. A. & Matthews, B. W. (1995). Protein flexibility and adaptability seen in 25 crystal forms of T4 lysozyme. *J. Mol. Biol.* **250**, 527-552.

Edited by A. R. Fersht

(Received 6 April 1998; received in revised form 4 May 1999; accepted 13 May 1999)