Perspectives in protein-fold recognition

Andrew E Torda

Fold recognition force fields based on statistics from native structures have become commonplace. New, nonphysical force fields based on optimizing parameters rather than reflecting Boltzmann statistics may offer improvement in force-field performance for threading and other applications. Improvements in sequence-to-structure alignments will also be essential for improved fold recognition.

Addresses
Research School of Chemistry, Australian National University, Canberra, ACT 0200, Australia;
e-mail: Andrew.Torda@anu.edu.au

Current Opinion in Structural Biology 1997, 7:200–205

Electronic identifier: 0959-440X-007-00200

Abbreviations
2D two-dimensional
3D three-dimensional

Introduction

One may well argue that the field of protein-fold recognition is suffering from an excess of scoring functions and an overabundance of review articles [1–5]. An enthusiast, however, may note there are some new ideas in force-field construction and some new approaches to the sequence–structure alignment problem.

This review is concerned with protein-structure prediction for cases in which a target sequence does not have unambiguous sequence homology to any known structure. It is also centred on protein-fold recognition as opposed to *ab initio* protein-structure prediction. Fold recognition (usually by threading [6]) means that one has a library of candidate structures and uses a score function to see which is the most appropriate for a sequence of interest. *Ab initio* probably means one uses any method capable of generating coordinates for a protein sequence without relying on a library. This wording is deliberately vague as search methods include molecular dynamics and Monte Carlo simulations, genetic algorithms [7,8], hierarchical methods [9,10] or procedures similar to those used to generate structures based on NMR data [11,12].

At the heart of almost every method is some kind of score function or force field, which behaves like potential energy or perhaps even free energy. Such force fields often have little fundamental physical basis but come directly from observations of known protein structures, so one might prefer to use terms such as structure-derived, statistical or knowledge-based force fields. More generally, these are low resolution force fields, as single interaction sites may represent more than one atom and some recent constructions do actually have a real physical basis. If one is limited to protein-fold recognition and forgoes all claim to physical energies, a more appropriate term is simply sequence–structure fitness functions.

Such score functions may behave quite differently depending on the method of construction and the intended application. If a score function is only going to be challenged with prefolded, compact structures from a library, then it may not be necessary for it to perform well on less native-like structures. A function may perform perfectly when matching sequences to library structures but produce an absurd answer when applied to a totally wrong, unfolded structure. This kind of function would not be very general, but it might be perfectly suited to fold recognition and threading. If, however, a function will be used for scoring in *ab initio* structure prediction, it may well be challenged with different structures generated by the search algorithm. In an extreme case, consider a molecular dynamics or Monte Carlo simulation of infinite length. These methods are supposed to be ergodic and may visit (infrequently) conformations that are in no way protein-like. It is essential that a score function in this context should penalize these conformations appropriately. This consideration has practical implications; for example, most force fields for *ab initio* structure prediction will produce a very high energy if two interaction sites (maybe atoms) lie on top of each other. A fold-recognition force field may never be confronted with this situation and may be perfectly adequate with absolutely no excluded volume or repulsive term.

None of this reasoning forbids the discovery of a reliable and perfectly general scoring function that works in all situations. It merely suggests that it might be easier to construct a function with limited goals; however, this reasoning does not dissuade many workers and some of the force fields discussed below are geared towards being generally applicable.

If one is going to improve fold recognition, there are several possible areas that could be worked on, although some will be more fruitful than others. Can one build better force fields or better scoring functions or add more information? Are there better methods for dealing with alignments (including gaps and insertions) when fitting sequences to structures? Can one build better fold libraries? These are some of the points that I will address in this review.

New methods for force-field construction

The most common method for building a sequence–structure fitness function is by using statistics from collections

of native structures [6,13,14]. First, one presumes that protein structures from a data bank constitute a kind of ensemble, and that the distances between interaction sites within proteins distribute themselves according to a Boltzmann distribution. Second, one can calculate the potential of mean force responsible for the observed statistics via the Boltzmann equation. Implementing this usually means discretizing distances and producing table-driven force fields or contact terms, rather than the continuous interaction functions familiar to molecular mechanics. Although the framework sounds like a simple prescription, there is still room for variation. For example, a force field based only on Cα coordinates has fold-recognition capability, but adding information from Cβ atoms or sidechain centroids is better [15], and different research groups use different interaction sites. One could certainly add more atom-based interaction sites, possibly to better account for hydrogen bonding, but at some point one may introduce noise or redundant information or possibly double counting. The framework of the Boltzmann relation is not limited to distances. One may well add in angular dependence [16,17]. There is also no consensus as to how the interactions should be divided up. Typically, statistics are collected separately for different topological distances. A force field may have different contributions from residues separated by one intervening residue (*i*,*i*+2), two intervening residues (*i*,*i*+3) and so on. These classes can be grouped and, beyond a certain topological distance, treated as long range. There is little reason to believe that any two groups in the world have used the same categorization.

If one does not feel that pairwise distances are adequate, one may introduce other terms. The early, successful, threading force fields used an additional term to account for solvation [6,18], whereas Matsuo and Nishikawa [16,17] tried to account for solvation, hydrogen bonding and local conformational preferences or tripeptide preferences. This suggests an implicit weighting of terms; for example, hydrogen bonding must indirectly influence the general atom–atom pair functions. Adding an additional term based on the appropriate N–O distances means that hydrogen bonding is counted once on N–O distances and once via effects on other atoms. More recently, Miyazawa and Jernigan [19] noted that statistically derived potentials of mean force suffer from packing and higher order effects (more than from pairwise effects) and introduced an additional repulsive term to account for these.

Force-field construction is really so arbitrary that if one feels that Boltzmann-based statistics are inadequate, one can simply add in any term that appears useful. For example, one might rank a set of candidate structures for a sequence using a statistical force field. Probable guesses could then be filtered using a conventional electrostatic term [20]. At least two implications of this approach exist: first, it must be the case that the form of the first force field does not quite capture the electrostatic contribution of the second force field, otherwise the second calculation would not be useful; second, using one function for ranking, followed by a second for filtering is similar to multiplying the two force-field terms. This is unusual in a world in which we are used to the additivity of molecular-mechanics-type interaction functions.

Most force fields have been tested by checking that some sequences are most suited to their native structures hidden within some threading library. Park and Levitt [21••] applied a more difficult test. They used an off-lattice, discrete model to generate vast numbers of misfolded but native-like structures. Several force field terms were then checked for their performance on fold recognition. It may not be surprising that none of the contact term, the Lennard–Jones-like term or the surface-area-based term was alone sufficient. What is more interesting is that the authors tried to generate a force field that would work in a broad problem domain. As Park and Levitt [21••] suggest, their work may be of use in folding simulations, rather than just in fold recognition.

In contrast to the enthusiasm for statistically based pseudoenergy functions, several groups have attempted to identify weaknesses in this kind of approach. Thomas and Dill [22••] used simple interaction functions and a 2D lattice to show how residue pairs may not be independent within a protein. This could be seen as damning for a methodology in which one assumes that the statistics of a pair AB are independent of a pair CD. Thomas and Dill [22••] went on to give a geometric explanation of how this lack of independence will depend on protein size. They further challenged the statistical force fields by suggesting that a great number of observations can be explained simply on the basis of hydrophobic/polar considerations. This last point has also clearly been made by Huang *et al.* [23], who performed well on a large number of fold-recognition tests, using such a simple hydrophobic/polar model.

If one wants to build a set of functions for fold recognition, there is no reason to be a slave to Herrn Boltzmann's statistics. One may construct a fold-recognition force field by searching for a formulation that favours a group of native sequence–structure pairs over a group of incorrect nonnative sequence–structure pairs [24]. This is quite a different philosophy to the methods discussed above. In the statistical methods, one looks at native sequence–structure pairs, formulates a potential energy and expects it to be able to recognize nonnative pairs as having high energy. Some approaches now exist that build native/nonnative recognition into the parametrization from the start. One particularly cumbersome and unwieldy approach has the entertaining feature of using quasi-Newtonian dynamics in parameter space as a method for optimizing the force field's performance [25]. This approach also uses an objective method for clustering interaction types so as to reduce the number of adjustable

parameters, trying to gracefully handle combinatorial degrees of freedom on top of the more tangible degrees of parameter freedom.

At least two groups have presented very elegant approaches to building force fields that maximize the extent to which they favour native sequence–structure pairs over misfolded structures. This could be seen as either increasing the statistical confidence of predictions or the foldability of the force field. Wolynes and coworkers [26••] have first defined a force field quality function, based on the energetic separation of native sequence–structure pairs from misfolded pairs. They have used an iterative approach so that at each step the parameters are optimized, the alignments recalculated and the quality function evaluated (see below). Hao and Scheraga [27••] have used a different gradient-descent approach, but one also aimed towards optimizing the force-field performance. Their work is the first to show how parameters can be optimized for fold recognition against structures generated by threading. The force field has been further improved by optimizing its discrimination of the native conformation from low energy structures generated by a sampling method such as Monte Carlo.

Once one abandons Boltzmann formalism and treats the force field as something to be rationally optimized, there are other properties one might build in. Crippen [28••] has used simple lattice models to first show how one can extract a useful potential energy function from knowledge of native structures alone and avoid the problems seen by Thomas and Dill [22••]. Crippen [28••] has subsequently noted that even if a pseudoenergy function performs well at correctly recognizing folds, only a poor correlation between energy and proximity to the native structure may exist; however, such a correlation is desirable for two reasons. First, if one is building a force field for threading, it is not only necessary to perform well on the native sequence structure, it is also necessary to perform well for closely related structures: one's library of candidate templates will probably not contain a perfectly correct structure. Second, if one wants to use some general search method, the task is easier if the force field provides some guidance, and low energies are more likely to correspond to structures that are at least close to the desired answer. Crippen [28••] explicitly built this property into the force field by casting it in the form of a series of inequalities and solving these for parameters using a linear programming method.

Finkelstein and Reva [29] have used molecular field theory to score protein sequence–structure alignments. Their work has bearings not only on force-field construction but also an alignment methodology (see below), as the molecular field and sequence–structure alignment are iterated to self-consistency. Although it may not be obvious from the title, the method does rely on physically based force-field terms. Finkelstein and Reva [30] present an application based on β-sandwich proteins but presumably have other calculations in store.

A final consideration of force fields is that if one can not produce a better sequence–structure fitness function, can one use a better sequence? Defay and Cohen [31] incorporated information from multiple sequence alignments into a threading scheme. This is an idea which has been mooted informally but not tested (or at least published) formally. Intuitively, information from multiple sequence alignments would be expected to be useful. Frequently, there is no known structure for a protein sequence but there are many highly homologous sequences known that would be expected to fold into the same structure. Clearly, a score based on more than one sequence will average over contributions from residues that can be changed without changing the overall fold.

Alignment methodology

Whichever way one views progress on scoring functions, almost any fold-recognition method has to address the problem of sequence-to-structure alignment. This is one area in which there may be some consensus: current methods do not perform well; and score functions are frequently applied to nonoptimal alignments of a target sequence to a candidate structure [32]. Ultimately, this will never be an easy problem as, for any pairwise interaction function, finding the optimal alignment is NP-complete if one allows gaps or insertions of any length at any position [33]. In other words, one can not bound the running time by a polynomial function of the input size. This does not really mean that one is likely to spend time exponentially proportional to the size of the sequence or structure of interest; it means that some approximations will be made, or perhaps one's probability of finding the correct answer will decrease logarithmically with increasing sequence length.

Not all force fields suffer from the combinatorial explosion. If a residue can be scored in a position on a structure without reference to the identity of its neighbours, a standard dynamic programming algorithm can be used to find the optimal alignment in $O(n^3)$ time. This is the case for some hydrophobic/polar force fields (although it has not always been exploited) and for other force fields that only depend on the type of a residue and the coordinates of the template that it is fit to [34].

If one is using a force field that requires knowledge of both interaction partners, then some approximations or restrictions will usually be used to make the calculations feasible. For example, one could allow gaps and insertions in some parts of the structure but could set gap penalties to be infinitely high within recognized secondary structure [35].

Another approximation is the frozen approximation in which one scores each residue from the sequence in the

field due to the residues present on the protein template. Westhead *et al.* [36] have compared this approach with the two-level dynamic programming method of Jones *et al.* [6]. Neither method is greatly superior, but Westhead *et al.* [36] have produced some useful methodological results. They find that, using the frozen approximation, two stages are useful. First, a family of folds can be recognized using high gap penalties. Second, the alignments can be optimized for the individual structures using lower gap penalties. Westhead *et al.* [36] have also described computational shortcuts to speed up the alignment calculations; for example, they use constant gap penalties in some cases and infinite gap penalties within certain secondary structures. Some energy calculations are avoided entirely by screening on the basis of one component (solvation) of the pseudoenergy term.

The fact that a search space may grow exponentially does not mean that a problem is always intractable. It may be possible to construct a method that is not guaranteed to find the optimal solution but, instead, finds a very good solution with high probability. Lathrop and Smith [37••] described a branch and bound search for finding the optimal sequence–structure alignment. A side effect of having such a good algorithm is that it highlights weaknesses in the scoring functions and Lathrop and Smith outline aspects that they feel need improvement.

A discussion of alignment algorithms assumes that one has an idea of the appropriate gap penalties. In fact, gap penalties are largely determined by trial and error and are nontransferable between force fields. The framework of Boltzmann statistics does not obviously give gap penalties and most methods for optimizing force-field parameters do not treat this problem. In contrast, Koretke *et al.* [26••] have devized a method whereby the penalties are optimized with the rest of the force-field parameters. At each step of the iterative parametrization, alignments are also recalculated using the current set of parameters and these newly aligned structures are used at the next iteration.

Fold libraries

A fold library may refer to the training set of proteins used to parametrize a force field or it may refer to a set of structures used as candidates for some sequence.

At some point, it is very probably that someone will use totally artificial constructs for a fold library rather than rely on the product of nature and structure-determination methods. This possiblity was hinted at by Crippen and Maiorov [38], who have used a geometric measure and estimate all the possible folds for a specific size and difference criterion [38]. Their work has also demonstrated a method for generating smoothed protein structures by applying a discrete cosine transform and setting the smaller coefficients to zero before back-transforming. Given the low resolution nature of most fold-recognition force fields, it may not be unreasonable to remove the bumps and details of real proteins.

Hamprecht *et al.* [39] have also highlighted the arbitrary nature of current fold libraries. They proposed generating folds by reorganizing the topology of existing structures where real proteins would permit this according to geometric criteria. This may not be necessary for threading small sequences onto larger structures, as current alignment methods allow a sequence to align with gaps to a template. There is also no clear route to handling the combinatorial explosion due to all the ways a protein can be reassembled. They do, however, make the very strong argument that folding libraries are short of structures suitable for threading large sequences.

Conclusions

Protein-fold recognition has quickly moved from being purely speculative to a point at which applications can be published [40]; many programmes are currently available and the methods for generating statistical force fields appear to be quite routine. This progress is in despite of the fact that the methods are not reliable and certainly not truly automatic (they still need some amount of interpretation). Furthermore, one is tempted to ignore the current weakness in sequence–structure alignments.

From the point of view of statistical force fields, incremental improvements will certainly take place, but it may be the case that more effort is needed on interpretation [41], rather than more statistics. One approach is to combine information from the length of a sequence, a candidate structure, the pseudoenergy scores and the number of aligned residues from a sequence–structure alignment [42].

The problem of score interpretation is even worse. Most force fields have a sequence or structure bias that seems resistant to any sequence composition correction [43] or to the use of statistical measures such as z-scores [18,44]. More fundamentally, a measure such as the z-score can lose its meaning if one allows deletion of sequence during alignment as the effective protein sequence changes [2].

One apparent direction in the field is the movement away from threading methods. This may mean the construction of force fields that work on a wider variety of misfolded structures [21••,27••] or of force fields for simulation [19], or it may even mean taking force fields that have their roots in fold recognition and using these to extract free energies [45••,46••]. Another trend seems to be the dissatisfaction with the testing of force fields, especially with ungapped alignments [31].

Although many would disagree, it may be that highly potent force fields will be built by forgetting Boltzmann's equation and optimizing the parameters for some functional forms, so as to optimize a measure of

force-field quality [25,26••–28••]. Depending on one's quality measure, this could mean building better force fields in the conventional sense or, perhaps, admitting that in many cases what one wants is a discrimination function (for good or bad sequence–structure matches).

These developments, together with new approaches to alignment problems, suggest that there will be improvement in results from protein-fold recognition, and not merely the stasis that one might see from studying the mass of similar force fields.

References and recommended reading

Papers of particular interest, published within the annual period of review, have been highlighted as:

- • of special interest
- •• of outstanding interest

1. Jones DT, Thornton JM: Potential energy functions for threading. *Curr Opin Struct Biol* 1996, 6:210–216.
2. Sippl MJ, Flöckner H: Threading thrills and threats. *Structure* 1996, 4:15–19.
3. Böhm G: New approaches in molecular structure prediction. *Biophys Chem* 1996, 59:1–32.
4. Jernigan RL, Bahar I: Structure-derived potentials and protein simulations. *Curr Opin Struct Biol* 1996, 6:195–209.
5. Sippl MJ: Knowledge-based potentials for proteins. *Curr Opin Struct Biol* 1995, 5:229–235.
6. Jones DT, Taylor WR, Thornton JM: A new approach to protein fold recognition. *Nature* 1992, 358:86–89.
7. Pedersen JT, Moult J: ***Ab initio*** structure prediction for small polypeptides and protein fragments using genetic algorithms. *Proteins* 1995, 23:454–460.
8. Unger R, Moult J: Genetic algorithms for protein folding simulations. *J Mol Biol* 1993, 231:75–81.
9. Srinivasan R, Rose GD: LINUS: a hierarchic procedure to predict the fold of a protein. *Proteins* 1995, 22:81–99.
10. Gunn JR: Minimizing reduced-model proteins using a generalized hierarchical table-lookup potential function. *J Phys Chem* 1996, 100:3264–3272.
11. Pachter R, Fairchild SB, Lupo JA, Adams WW: Biomolecular structure prediction at a low resolution using a neural network and the double-iterated Kalman filter technique. *Biopolymers* 1996, 39:377–386.
12. Reese MG, Lund O, Bohr J, Bohr H, Hansen JE, Brunak S: Distance distributions in proteins: a six-parameter representation. *Protein Eng* 1996, 9:733–740.
13. Tanaka S, Scheraga HA: Medium- and long-range interaction parameters between amino acids for predicting three-dimensional structures of proteins. *Macromolecules* 1976, 9:945–950.
14. Sippl MJ: Calculation of conformational ensembles from potentials of mean force. *J Mol Biol* 1990, 213:859–883.
15. Kocher J-PA, Rooman MJ, Wodak S: Factors influencing the ability of knowledge-based potentials to identify native sequence–structure matches. *J Mol Biol* 1994, 235:1598–1613.
16. Nishikawa K, Matsuo Y: Development of pseudoenergy potentials for assessing protein 3D–1D compatibility and detecting weak homologies. *Protein Eng* 1993, 6:811–820.
17. Matsuo Y, Nishikawa K: Protein structural similarities predicted by a sequence–structure compatibility method. *Protein Sci* 1994, 3:2055–2063.
18. Sippl MJ: Boltzmann's principle, knowledge-based mean fields and protein folding. An approach to the computational determination of protein structures. *J Comput Aided Mol Des* 1993, 7:473–501.
19. Miyazawa S, Jernigan RL: Residue–residue potentials with a favorable contact pair term and an unfavorable high packing density term, for simulation and threading. *J Mol Biol* 1996, 256:623–644.
20. Matsuo Y, Nakamura H, Nishikawa K: Detection of protein 3D–1D compatibility characterized by the evaluation of side-chain packing and electrostatic interactions. *J Biochem* 1995, 118:137–148.
21. •• Park B, Levitt M: Energy functions that discriminate X-ray and near-native folds from well-constructed decoys. *J Mol Biol* 1996, 258:367–392.
In this paper, rather than just generate force fields suitable for threading, the authors offer a series of terms that, when summed, allow fold recognition from a very large number of native-like structures.
22. •• Thomas PD, Dill KA: Statistical potentials extracted from protein structures: how accurate are they? *J Mol Biol* 1996, 257:457–469.
This paper uses a simple lattice model to highlight weaknesses in the calculation of potentials of mean force from native structures. The authors suggest that fold recognition would be almost as powerful using a much simpler hydrophobic/polar model for interactions.
23. Huang ES, Subbiah S, Tsai J, Levitt M: Using a hydrophobic contact potential to evaluate native and near-native folds generated by molecular dynamics simulations. *J Mol Biol* 1996, 257:716–725.
24. Crippen GM, Snow ME: A 1.8 Å resolution potential function for protein folding. *Biopolymers* 1990, 29:1479–1489.
25. Ulrich P, Scott W, Van Gunsteren WF, Torda AE: Protein structure prediction force fields: parametrisation with quasi-Newtonian dynamics. *Proteins* 1997, in press.
26. •• Koretke KK, Luthey-Schulten Z, Wolynes PG: Self-consistently optimized statistical mechanical energy functions for sequence structure alignment. *Protein Sci* 1996, 5:1043–1059.
First, the authors parametrize a force field by optimizing its ability to distinguish native from nonnative folds for a sequence. Second, the authors give a method to optimize gap penalties along with the rest of the force field.
27. •• Hao M-H, Scheraga HA: How optimization of potential functions affects protein folding. *Proc Natl Acad Sci USA* 1996, 93:4984–4989.
As in [26••], the authors present a force field generation method that optimizes fold-recognition capability. They show that even though a force field performs well at finding a native structure amongst alternatives generated by threading, it does not perform well with misfolded structures generated by other means. The authors then improve their own force field to handle these other decoy structures.
28. •• Crippen GM: Easily searched protein folding potentials. *J Mol Biol* 1996, 260:467–475.
This paper gives a method for extracting the correct potentials of mean force from observations of native structures. The authors then show how to improve the suitability of the force field for searching algorithms by imposing some correlation between pseudoenergy and distance from the correct answer.
29. Finkelstein AV, Reva BA: Search for the most stable folds of protein chains. I. Application of a self-consistent molecular field theory to a problem of protein three-dimensional structure prediction. *Protein Eng* 1996, 9:387–397.
30. Reva BA, Finkelstein AV: Search for the most stable folds of protein chains: II. Computation of stable architectures of β-proteins using a self-consistent molecular field theory. *Protein Eng* 1996, 9:399–411.
31. Defay TR, Cohen FE: Multiple sequence information for threading algorithms. *J Mol Biol* 1996, 262:314–323.
32. Lemer CM-R, Rooman MJ, Wodak SJ: Protein structure prediction by threading methods: evaluation of current techniques. *Proteins* 1995, 23:337–355.
33. Lathrop RH: The protein threading problem with sequence amino acid interaction preferences is NP-complete. *Protein Eng* 1994, 7:1059–1068.
34. Ouzonis C, Sander C, Scharf M, Schneider R: Prediction of protein structure by evaluation of sequence–structure fitness. *J Mol Biol* 1993, 232:805–825.
35. Madej T, Gilbrat J-F, Bryant SH: Threading a database of protein cores. *Proteins* 1995, 23:356–369.
36. Westhead DR, Collura VP, Eldridge MD, Firth MA, Li J, Murray CW: Protein fold recognition by threading: comparison of algorithms and analysis of results. *Protein Eng* 1995, 8:1197–1204.

37. •• Lathrop RH, Smith TF: Global optimum protein threading with gapped alignment and empirical pair score functions. *J Mol Biol* 1996, 255:641–665.
The authors give the first practical method for handling sequence–structure alignments using residue-specific scoring functions and without using approximations.

38. Crippen GM, Maiorov VN: How many protein folding motifs are there? *J Mol Biol* 1995, 252:144–151.

39. Hamprecht FA, Scott W, Van Gunsteren WF: Generation of pseudo-native protein structures for threading. *Proteins* 1997, in press.

40. Madej T, Boguski MS, Bryant SH: Threading analysis suggests that the obese gene product may be a helical cytokine. *FEBS Lett* 1995, 373:13–18.

41. Miller RT, Jones DT, Thornton JM: Protein fold recognition by sequence threading: tools and assessment techniques. *FASEB J* 1996, 10:171–178.

42. Jones DT: Two applications of statistical potentials: protein fold recognition and ***de novo*** protein design. Abstract for Molecular Interactions, 15th International Meeting of Molecular Graphics Society, 1996 April 16–19, York.

43. Bryant SH, Lawrence CE: An empirical energy function for threading protein sequence through the folding motif. *Proteins* 1993, 16:92–112.

44. Jones DT, Miller RT, Thornton JM: Successful protein fold recognition by optimal sequence threading validated by rigorous blind testing. *Proteins* 1995, 23:387–397.

45. •• Sippl MJ: Helmholtz free energy of peptide hydrogen bonds in proteins. *J Mol Biol* 1996, 260:644–648.
See annotation [46••].

46. •• Sippl MJ, Ortner M, Jaritz J, Lackner P, Flöckner H: Helmholtz free energies of atom pair interactions in proteins. *Fold Des* 1996, 1:289–298.
These two papers, [45••,46••] are of outstanding interest because of the controversy they arouse and the boldness of the authors' assertion that protein data bank structures constitute an ensemble that allows free-energy estimations.