BIOINFORMATICS

Protein secondary structure: entropy, correlations and prediction

Gavin E. Crooks* and Steven E. Brenner

Department of Plant and Microbial Biology, University of California, 111 Koshland Hall No. 3102, Berkeley, CA 94720-3102, USA

ABSTRACT

Motivation: Is protein secondary structure primarily determined by local interactions between residues closely spaced along the amino acid backbone or by non-local tertiary interactions? To answer this question, we measure the entropy densities of primary and secondary structure sequences, and the local inter-sequence mutual information density.

Results: We find that the important inter-sequence interactions are short ranged, that correlations between neighboring amino acids are essentially uninformative and that only one-fourth of the total information needed to determine the secondary structure is available from local inter-sequence correlations. These observations support the view that the majority of most proteins fold via a cooperative process where secondary and tertiary structure form concurrently. Moreover, existing single-sequence secondary structure prediction algorithms are almost optimal, and we should not expect a dramatic improvement in prediction accuracy.

Availability: Both the data sets and analysis code are freely available from our Web site at http://compbio.berkeley.edu/

Contact: gec@compbio.berkeley.edu

INTRODUCTION

The secondary structure of a protein is a summary of the general conformation and hydrogen bonding pattern of the amino acid backbone (Frishman and Argos, 1995). This structure is frequently simplified to a sequence (one element per residue) of helixes (H), extended strands (E) and unstructured loops (L). It has long been recognized that each residue's secondary structure is appreciably correlated with the local amino acid sequence (Szent-Gyorgyi and Cohen, 1957) and that these correlations may be used to predict the secondary structure (Rost, 2001; Przybylski and Rost, 2002) or as a contribution to threading potentials (Alexandrov *et al.*, 1996; McGuffin and Jones, 2003) and other tertiary structure prediction algorithms (Bowie *et al.*, 1991). The effectiveness of local secondary structure prediction, and the utility of secondary structure potentials, depends upon the extent to which

```
Primary:    YDPEEHHKLSHEAESLPSVVISSQAAGNAVMMGAGYFSP
Secondary:  LLHHHHHHHHHHHHLLLEEEELLHHHHHHHHHHHHLLLLL
```
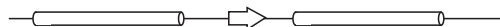
Fig. 1. A protein's amino acid sequence is correlated with the corresponding secondary structure sequence, represented here by a sequence of helixes (H), extended strands (E) and unstructured loops (L). For example, alanines (A) are typically associated with helixes, while glycines (G) are often located near helix breaks. Also note that secondary structure is strongly persistent. Helixes, e.g. are on average about 10 residues long (Schmidler *et al.*, 2000).

a protein's structure, particularly the secondary structure, is determined by local, short-ranged interactions between residues closely spaced along the backbone, as opposed to non-local or long-ranged tertiary interactions.

The strength, organization and relative importance of local sequence–structure interactions can be determined with a statistical analysis of the corpus of known protein structures. We treat the primary and secondary structures of a protein as random sequences comprising either the 20 letter amino acid or the three letter extended strand/helix/other (EHL) structure alphabets, as shown in Figure 1. These sequences contain substantial local sequence and inter-sequence correlations that can be quantified using entropic measures. To ensure accurate results, we employ a large, carefully curated collection of protein structures derived from the structural classification of proteins (SCOP) (Murzin *et al.*, 1995; Lo Conte *et al.*, 2002) database, which contains 2853 sequences.

Sequence information

Entropy is a measure of the information needed to describe a random variable. Specifically, the entropy, $H(X)$, of a discrete random variable X, measured in bits, is defined as

$$H(X) = -\mathrm{E}\big[\log_2 P(X)\big] = -\sum_{x \in \mathcal{X}} P(x)\log_2 P(x), \quad (1)$$

where $\mathcal{X}$ is the alphabet, the set of allowed states, x is an element of $\mathcal{X}$, $\mathrm{E}(X)$ is the expectation and $P(x)$ is the probability of state x. The maximum achievable entropy is $\log_2 |\mathcal{X}|$ bits, where $|\mathcal{X}|$ is the size of the alphabet. For a modern

*To whom correspondence should be addressed.

introduction to information theory, see the book by Cover and Thomas (1991).

When considering the entropy of a collection of variables, it is important to take into account inter-variable correlations. For a statistically homogeneous random sequence with local correlations, the appropriate information measure is the entropy density, h_μ, the rate at which the entropy of the sequence increases with length:

$$h_\mu = \lim_{L \to \infty} \frac{H(X^L) - E_h}{L}. \tag{2}$$

Here, $H(X^L)$ is the entropy of sequence fragments, X^L, of length L. The non-extensive excess entropy, E_h, is the quantity of information explained away by taking account of inter-site correlations. The entropy density is also referred to as the entropy rate or metric entropy (Cover and Thomas, 1991).

A convenient measure of correlation between two discrete random variables X and Y is the mutual information, $I(X; Y)$, defined as

$$I(X; Y) = H(X) + H(Y) - H(X, Y), \tag{3}$$

$$= \sum_{x \in \mathcal{X}, y \in \mathcal{Y}} P(x, y) \log_2 \frac{P(x, y)}{P(x)P(y)}, \tag{4}$$

where $P(x, y)$ is the joint probability of observing states x and y. If the random variables are independent [$P(x, y) = P(x)P(y)$], then the mutual information achieves its lower bound of zero. Mutual information cannot exceed the entropy of either variable, and this upper bound is reached when the variables are perfectly correlated [$P(x, y) = P(x) = P(y)$].

The appropriate entropic correlation measure for a pair of statistically homogeneous random sequences is the mutual information density, i_μ,

$$i_\mu = \lim_{L \to \infty} \frac{I(X^L; Y^L) - E_i}{L}. \tag{5}$$

Here, E_i is the excess mutual information.

When we consider the correlations between three random variables, it is often useful to consider $I(X; Y|Z)$, the conditional mutual information (Cover and Thomas, 1991) of X and Y, given a third variable Z. This quantity can be conveniently defined in terms of mutual information:

$$I(X; Y|Z) = I(X; Y) + I(X, Y; Z) - I(X; Z) - I(Y; Z). \tag{6}$$

Conditioning on a third random variable may increase or decrease the mutual information.

MATERIALS AND METHODS

Secondary structure library

Ideally, a secondary structure library should be based upon a representative, high-quality and non-redundant subset of available protein structures. The Protein Data Bank (PDB)

(Berman *et al.*, 2000) currently contains over 20 000 publicly accessible structures, but many of these are very similar, and many are of relatively low quality. The SCOP (Murzin *et al.*, 1995; Lo Conte *et al.*, 2002) database provides a convenient decomposition of PDB structures into domains, and the ASTRAL (Brenner *et al.*, 2000; Chandonia *et al.*, 2002) compendium provides representative subsets of SCOP domains, filtered so that no two domains share more than a given percentage level of sequence identity. This filtering preferentially retains higher quality structures, as judged by AEROSPACI scores (Chandonia *et al.*, 2002), an agglomeration of several structure quality measures. We selected the ASTRAL 40% sequence identity subset of SCOP release 1.61, which was further filtered to remove multi-sequence domains, SCOP classes f (membrane and cell surface proteins) and g (small proteins) and retain only those structures determined by X-ray diffraction at better than 2.5 Å resolution. The protein sequences were taken from the ASTRAL rapid access format (RAF) sequence mappings (Chandonia *et al.*, 2002), which provide a more reliable and convenient representation of the true sequence than the PDB ATOM or SEQRES records. The secondary structure sequences were determined by the program STRIDE (Frishman and Argos, 1995) using each protein's hydrogen bonding pattern and backbone torsional angles. STRIDE was unable to process a small fraction of SCOP domains that were consequentially removed from further consideration. The resulting library contains 2853 protein domains and 553 373 residues.

For comparative purposes, we also studied the secondary structure library of Cuff and Barton (1999, 2000), which consists of 513 proteins and 84 091 residues. Secondary structure assignments are provided by both STRIDE and the program DSSP (Kabsch and Sander, 1983). This data set also includes, for each structure, a multiple alignment of homologous sequences. These multiple sequence alignments were converted to amino acid probability profiles (Durbin *et al.*, 1998) using the program `hmmbuild` from HMMER (v2.3) (Eddy, 2001, http://hmmer.wustl.edu/).

Both DSSP and STRIDE assign each residue's secondary structure to one of eight classes: α-helix (H), 3_{10} helix (G), π-helix (I), β-strand (E), β-bridge (B or b), Coil (C, L or space), Turn (T) and Bend (S). Unstructured or poorly resolved regions of the protein are unassigned (X). These eight classes were reduced to the three letter alphabet, E (extended strand), H (helix) and L (loop/Other), using the common CK mapping (Chandonia and Karplus, 1995; Frishman and Argos, 1997; Cuff and Barton, 1999) E → E; H → H; all others → L. We also considered another common reduction, the 'EHL' mapping (Moult *et al.*, 2001; Rost and Eyrich, 2001) E, B → E; H, G, I → H; all others → L.

Entropy estimation and bias correction

The entropy of a discrete probability can be estimated by sampling from the distribution and then replacing the true

probabilities, $P(x)$, by the observed frequencies, $f(x) = n_x/N$. Here, N is the total number of samples and n_x is the number of observations of state x. A useful alternative approach is to construct an approximation of the true probabilities, $g(x) \approx P(x)$ [e.g. Equation (11)], and then estimate the entropy by the mean log-likelihood of the data (Moddemeijer, 2000).

$$
\begin{aligned}
H(X) &= E\big[\log_2 P(X)\big] \\
&\geq E\big[\log_2 g(X)\big] \\
&\approx \sum_{i=1}^{N} \frac{1}{N} \log_2 g(x_i).
\end{aligned} \tag{7}
$$

Similarly, the mutual information can be related to the mean log odds since

$$
I(X;Y) = E\left[\log_2 \frac{P(X|Y)}{P(X)}\right]. \tag{8}
$$

A serious problem with either approach is that entropies estimated from limited amounts of data tend to be significantly biased ~~significantly~~ (Miller, 1955), resulting in a systematic underestimation of the true entropy or overestimation of the mutual information. We used non-parametric bootstrap resampling (Efron and Tibshirani, 1993) to correct for this bias and to estimate standard statistical errors. A total of 50 replicas of the original data are generated by sampling, with replacement, from the available sequences. This resampling has associated systematic and random errors that are approximately the same as the errors introduced by the original finite sampling of sequences from the true random distribution. These error estimates were not significantly improved when the number of replicas was increased from 50 to 500. The requisite pseudo-random numbers were drawn from the Mersenne Twister generator (Matsumoto and Nishimura, 1998; Gough, 2003).

RESULTS

Entropy and correlations

In Figure 2, we plot the entropies for secondary structure sequence blocks of length up to 9 ($3^9 = 19\,683$ states). Of the half million residues in our data set, about 23% are assigned to strand, 39% to helix and 38% to other, resulting in a relatively large single-site secondary structure entropy of 1.53 bits. (The maximum entropy for three states is $\log_2 3 \approx 1.59$ bits.) However, neighboring secondary structure elements are strongly correlated, resulting in a relatively large nearest neighbor mutual information value, $I(S_i; S_{i+1}) \approx 0.89$ bits. A linear regression to the asymptotic functional form, $H(S^L) \sim Lh_\mu + E_h$ ($L \geq 3$), gives an excess entropy of $E_h = 0.997 \pm 0.004$ bits and an entropy density of $h_\mu = 0.598 \pm 0.001$ bits per residue. This entropy density, the amount of information needed to describe the secondary structure sequence, is considerably less than the single-site entropy
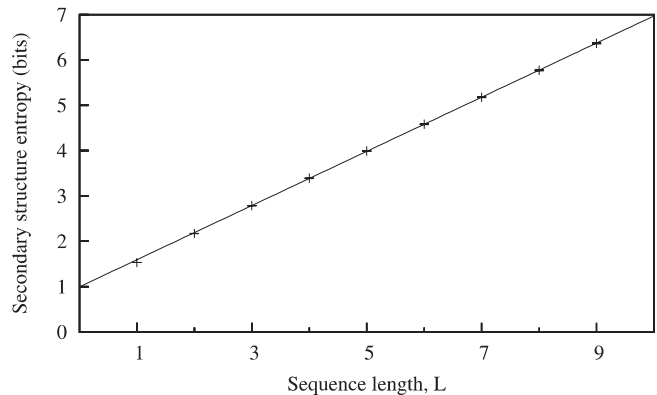
Fig. 2. Secondary structure sequences are strongly correlated, but the correlations have a simple structure. In this figure we plot the entropy of secondary structure blocks, $H(S^L)$, as a function of block length, L (points). Bootstrapped confidence intervals are smaller than the data point symbols. The linear increase in block entropies is indicative of a simple sequence, one that can, to a good first order approximation, be modeled as a low order Markov chain. A linear regression to the data (solid line) gives an excess entropy of $E_h \approx 1.0$ bits (zero intercept) and a true secondary structure entropy density of $h_\mu \approx 0.60$ bits per residue. Over half of the single-site entropy is explained away when we look beyond single-site statistics.

(1.53 bits) due to the strong inter-site correlations that may be observed in Figure 1 and Table 1.

It is notable that the entropies for short blocks are almost identical to the asymptotic linear extrapolation used to estimate entropy density and excess entropy (Fig. 2). This property is indicative of a sequence with a simple structure and suggests that many of the important statistical features of secondary structure sequences can be modeled successfully using a low order Markov chain, as discussed by Crutchfield and Feldman (2003).

In contrast to secondary structure, neighboring amino acids are only weakly correlated. The nearest neighbor mutual information, $I(R_i; R_{i+1}) \approx 0.006$ bits, is small relative to the single-site entropy of $H(R_i) \approx 4.18$ bits, which, consequentially, is almost identical to the primary sequence entropy density. Moreover, the mutual information between neighboring amino acids, conditioned upon the corresponding secondary structure [Equation (6)], is also relatively insignificant: $I(R_i; R_{i+1}|S_i S_{i+1}) \approx 0.016$ bits. Neighboring amino acids are approximately independent irrespective of the local structure, a result previously noted by Weiss *et al.* (2000). The correlations between more distantly separated residues are also very small.

The strength of the primary to secondary structure sequence correlations is quantified by the inter-sequence mutual information density. However, the mutual information can only be directly calculated for short sequence blocks due to the large effective alphabet of 60 ($= 3 \times 20$) symbols. The observed
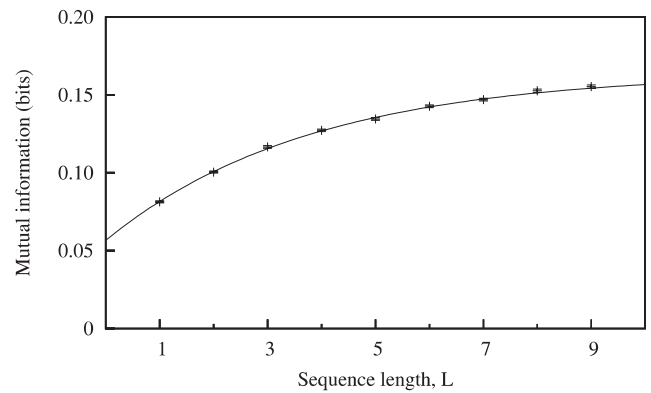
Table 1. Summary of primary structure (R) and secondary structure (S) sequence and inter-sequence information measures

|		Bits
Primary		
Residue entropy	$H(R_i)$	4.179 ± 0.001
Neighbor mutual info.	$I(R_i; R_{i+1})$	0.006 ± 0.002
Conditional neighbor MI	$I(R_i; R_{i+1}|S_i S_{i+1})$	0.0159 ± 0.0004
Entropy density	$h_\mu(R)$	4.173 ± 0.003
Secondary		
Residue entropy	$H(S_i)$	1.533 ± 0.002
Neighbor mutual info.	$I(S_i; S_{i+1})$	0.893 ± 0.003
Entropy density	$h_\mu(S)$	0.598 ± 0.001
Excess entropy	$E_h(S)$	0.997 ± 0.005
Inter-sequence		
Monomer mutual info.	$I(R_i; S_i)$	0.0813 ± 0.0007
Dipeptide mutual info.	$I(R_i R_{i+1}; S_i S_{i+1})$	0.208 ± 0.002
Mutual info. density	$i_\mu(R; S)$	0.164 ± 0.003

Fig. 3. The direct local interactions between primary and secondary structure are short-ranged. Here, the mutual information, $I_c(R^1, S^L)$, between a block of secondary structure of length L and the single amino acid located at the center (odd L) or immediately left of center (even L) of that block is plotted against block length (points). Bootstrapped confidence intervals are smaller than the data point symbols. A non-linear regression to an empirical exponential functional form gives a characteristic length scale of about four residues, and a limiting value of $I_c(R^1, S^\infty) \approx 0.164$, which is a reasonable approximation to the total mutual information density, $i_\mu(R; S)$.

single-site mutual information is $I(S_i; R_i) \approx 0.081$ bits, and the dipeptide mutual information is $I(R_i R_{i+1}; S_i S_{i+1}) \approx 0.208$ bits, or 0.104 bits per residue. Fortunately, to a good approximation we can neglect the correlations between amino acids since neighboring residues are almost (conditionally) independent. For example, the dipeptide mutual information, $I(R_i R_{i+1}; S_i S_{i+1}) \approx 0.208$ bits, can be approximated by $I(R_i; S_i S_{i+1}) + I(R_{i+1}; S_i S_{i+1}) \approx 0.198$, an expression that explicitly ignores amino acid correlations. The relatively small error of 0.010 bits ($<5\%$ of the dipeptide mutual information) is directly related to the mutual information between neighboring amino acids since [by Equation (6)]

$$I(R_i R_{i+1}; S_i S_{i+1}) - I(R_i; S_i S_{i+1}) - I(R_{i+1}; S_i S_{i+1})$$
$$= I(R_i; R_{i+1}|S_i S_{i+1}) - I(R_i; R_{i+1}).$$

It follows that the inter-sequence mutual information density can be estimated by examining $I_c(R^1; S^L)$, the mutual information between a block of secondary structure and the single amino acid located at the center of that block (Fig. 3). Empirically, we expected these entropies to decay exponentially toward their limiting value as block lengths increase (Crutchfield and Feldman, 2003). A non-linear regression to the functional form $a - b \exp(-L/c)$ (using data from odd block lengths only) gives $c = 3.8 \pm 0.3$ residues for the characteristic length scale, $b = 0.108 \pm 0.002$ for the scaling prefactor and $a = 0.164 \pm 0.003$ bits for the central amino acid to secondary structure mutual information in the infinity block length limit. This last value is a good approximation to the inter-sequence mutual information density, $i_\mu(R; S)$, with a bias, due to neglecting amino acid correlations, that is, probably $<10\%$.

In summary, the direct local interactions are short-ranged, neighboring amino acids are almost independent, secondary structure sequences are correlated, but essentially Markovian, and the important inter-sequence correlations are local, with a characteristic length scale of about 4. The inherent information content of secondary structure sequences is 0.60 bits per residue, about four times greater than the 0.16 bits per residue of local mutual information between primary and secondary structure. These measurements place severe constraints on any single-sequence prediction algorithm that purports to extract secondary structure information from local sequence correlations. In particular, no analysis can extract additional information from the signal [the data processing inequality (Cover and Thomas, 1991)], and therefore any sequence local prediction of secondary structure can contain no more information than that contained in the local primary–secondary sequence correlations.

Secondary structure hidden Markov model

Many different algorithms have been proposed for predicting secondary structure from local inter-sequence correlations. Interestingly, the underlying organization of the majority of these algorithms does not reflect the underlying organization of the intra- and inter-sequence interactions elucidated in the preceding section. Typically, these methods use a large primary structure window of around 15–27 residues to predict the single secondary structure element at the center of that window and often assume that inter-amino acid correlations are informative. However, even nearest neighboring
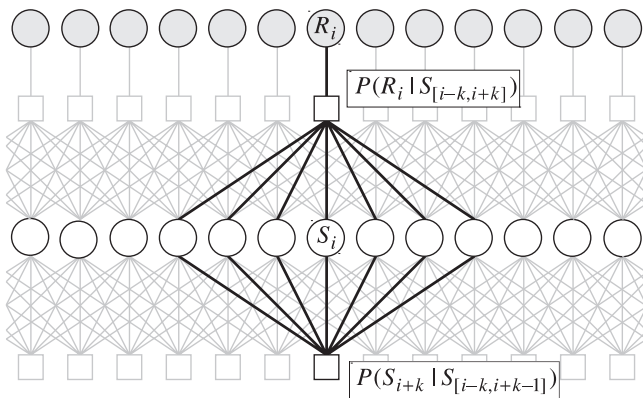
Fig. 4. A factor graph (Kschischang *et al.*, 2001) for $P(S|R)$, representing the decomposition of this complex, many-variable function into simpler parts [Equations (9)–(11)]. Circles represent variables and squares represent factors, local functions of relatively few variables. The upper and lower rows of circles represent the primary and secondary structure sequences, respectively. In this diagram, $k = 2m = 6$, for a window of size $k + 1 = 7$. One set of factors, centered on sequence position i, has been highlighted. The bottom factor connects $k + 1$ neighboring secondary structure elements and represents the approximation of the secondary structure sequence probability by a k-th order Markov chain [Equation (11)]. The factor between the chains represents the inter-sequence dependance [Equation (10)]. Thus, each residue is directly dependent upon a window of secondary structure (length $2m + 1$), and is conditionally independent of neighboring residues.

amino acids on the chain are only weakly correlated, and these correlations provide negligible information about the local structure.

As an alternative prediction algorithm, we have constructed a relatively simple hidden Markov model (HMM) [Equations (9)–(11) and Fig. 4] that embodies three key approximations: that protein sequences are statistically homogeneous, that direct secondary structure to primary structure interactions are local along the chain and that amino acids at neighboring sites are independent. Instead of a large primary structure window, we use short, overlapping secondary structure windows. Similar models, with similar assumptions, can be found in the work of Thompson and Goldstein (1997) and Schmidler *et al.* (2000).

The probability, $P(S|R)$, of a secondary structure sequence S, given the primary sequence, R, can be rewritten using Bayes' rule as

$$P(S|R) = \frac{P(R|S)P(S)}{P(R)}. \tag{9}$$

Since the probability of an amino acid residue depends on the local secondary structure and is almost independent of the identity of neighboring residues, to a good approximation the probability of each residue can be estimated from a short

window of local secondary structure,

$$P(R|S) \approx \prod_{i=1}^{L} P\big(R_i|S_{[i-m,i+m]}\big). \tag{10}$$

Here, X_i is the element at position i and $X_{[i,j]}$ is a subsequence of length $i - j + 1$ starting at position i and ending at j. Residues beyond the termini of the actual sequence ($i < 1$, $i > L$) are treated as undetermined. The window size, $2m + 1$, is an adjustable parameter of the model and need not be particularly long since the inter-sequence correlations have a characteristic length scale of only about four residues (Fig. 3). We approximate the prior probability of the secondary structure sequence by a k-th order Markov chain,

$$P(S) \approx P\big(S_{[1,k]}\big) \prod_{i=1}^{L-k} P\big(S_{i+k}|S_{[i,i+k-1]}\big). \tag{11}$$

The primary structure sequence probabilities, $P(R)$, can be determined from normalization. Combining the preceding approximations, [Equations (10) and (11)], using $k = 2m$ for consistency, generates a HMM (summarized in Fig. 4) that emits the primary structure sequence on transitions between blocks of secondary structure of length k.

The probabilistic model of Equations (9)–(11) can be generalized so that the prediction is based upon a multiple sequence alignment (MSA) of homologous protein sequences. First, we convert the multiple sequence alignment into an amino acid profile, $\theta = \{\theta_1(r), \theta_2(r), \ldots, \theta_L(r)\}$, that represents the probabilities of each amino acid at each position of the protein of interest (Durbin *et al.*, 1998). The secondary structure probability, given this profile, may then be approximated as

$$P(S|\theta) = \frac{P(\theta|S)P(S)}{P(\theta)}, \tag{12}$$

$$P(\theta|S) \approx \prod_{i=1}^{L} P(\theta_i|S_{[i-m,i+m]}). \tag{13}$$

We expect that each residue's observed homology profile, $\theta_i(r)$, will vary from the structure profile, $P(r|S_{[i-m,i+m]})$, due to sampling errors, random site-to-site variation and inter-protein structural variation and because each residue is under different structural variation and evolutionary constraints. As a simple approximation, we use the large deviation distribution (Cover and Thomas, 1991) to model the variation of the observed profile from the expected profile:

$$P(\theta_i|\beta, S_{[i-m,i+m]})$$
$$\approx \exp\left\{-\beta D\big(\theta_i(r) \,\|\, P(r|S_{[i-m,i+m]})\big)\right\}. \tag{14}$$

Here, $D(p\|q) = \sum_i p_i \ln(p_i/q_i)$ is the relative entropy. We treat β as an empirical dispersion parameter that is

independent of the secondary structure or primary structure profile.

Computationally, the conditional secondary structure probabilities can be derived from the amino acid sequence using the standard forward–backward dynamic programming algorithm (Rabiner, 1989). The time and memory complexities for a naive implementation are $O(L3^k)$, which, despite the exponential scaling, is feasible for moderate k. For example, with $k = 7$, training on one half of our library (2853 sequences) required 4 s from a modest contemporary PC (667 MHz PowerPC G4), and prediction of the other half required ~5 min, or about five sequences per second. In principle, a more efficient implementation is possible since, although the total number of secondary structure sequences scales as 3^L, the number of typical sequences with non-negligible probability scales as $2^{H(S^L)} \approx 1.5^L$, by the asymptotic equipartition principle (Cover and Thomas, 1991). The optimal prediction at a particular site is the secondary structure element with the greatest posterior probability.

The available sequence data were partitioned every other sequence into disjoint test and training sets of approximately equal size. The training set was used to estimate secondary structure block probabilities, $P\left(S_{[i,i+k]}\right)$ (regularized with a Laplace pseudocount of 1) and corresponding amino acid profiles, $P\left(R_i | S_{[i-m,i+m]}\right)$ (regularized with a pseudocount 20 times the amino acid background probability). Statistical errors were estimated from a full bootstrap resampling of both the test and training sequences.

We estimate the amount of information that the HMM successfully extracts by measuring the mean log odds of the observed secondary structure fragments [Equation (8)], and then extrapolating across different length scales to estimate the model mutual information density, i_μ^{HMM} [Equation (5)]. Since the maximum amount of information that can be extracted is the previously estimated inter-sequence mutual information density, $i_\mu(R; S)$ (Fig. 3), we may consider profitably the efficiency ratio, $R = i_\mu^{\text{HMM}}/i_\mu$, which is plotted in Figure 5. This model is able to extract over 90% of the available information with a modest secondary structure window size of only $L = 7$. In other words, the prediction algorithm is almost optimal.

The most common measure of secondary structure prediction quality is the average three-state accuracy, Q_3, the average fraction of residues that are correctly classified as helix, strand or other. Prediction accuracy increases monotonically with window length, reaching $65.9 \pm 0.3\%$ at $L = 9$ (Fig. 5). We cannot reliably increase the window size further due to the finite size of the training and test data sets.

The prediction accuracy can vary considerably due to variations in secondary structure assignment and due to variations in the underlying data set itself. Our standard data set consists of 2853 sequences derived from the 40% subset of SCOP release 1.61, with STRIDE (Frishman and Argos, 1995) secondary structure assignments. We also considered prediction
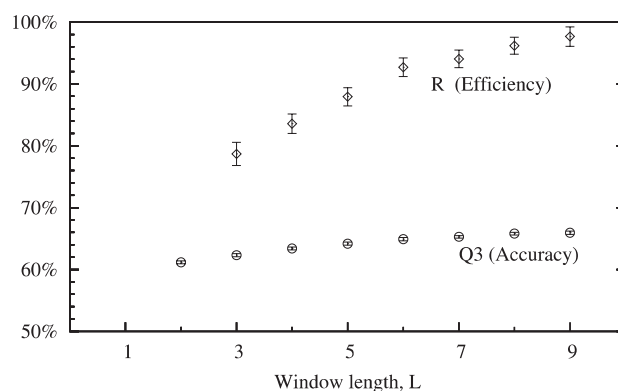
Fig. 5. The HMM defined in [Equations (9)–(11)] is able to extract over 95% of the available inter-sequence information. Here, the efficiency, $R = i_\mu^{\text{HMM}}/i_\mu$ and average three state accuracy (Q_3) are plotted against HMM window size ($L = k + 1$) for single-sequence prediction on the SCOP 1.61 40% STRIDE/CK data set. Results for the Barton data set are similar. Window sizes cannot be extended reliably beyond those shown here due to finite sequence data. The model information density, i_μ^{HMM}, approaches (but cannot exceed) the inter-sequence mutual information density, i_μ, indicating that the model is almost optimal. The prediction accuracy, $Q_3 = 65.9 \pm 0.3\%$ at $L = 9$, is the same (within statistical errors) as the accuracy of a variety of comparable secondary structure prediction algorithms, suggesting that these algorithms are also almost optimal.

accuracies for the Cuff–Barton (Cuff and Barton, 2000) library of 513 sequences, using STRIDE and DSSP (Kabsch and Sander, 1983) secondary structure assignments, and two different reductions of the STRIDE and DSSP alphabets to three states, the CK and EHL mappings (For details, see Materials and Methods). At $L = 7$, the accuracy ranges from 63.6 ± 0.6 to $66.4 \pm 0.7\%$. The maximum accuracy is achieved with the CK mapping, irrespective of the secondary structure assignment program. Essentially, the CK mapping produces more coherent, less random secondary structure sequences than the EHL mapping, which leads to more facile prediction. Using the smaller library of only 513 sequences leads to substantial standard errors of ~0.6% and to a large estimated bias of ~0.7%. (Without a bias correction our maximum reported accuracy would be 67%.) A number of different secondary structure prediction algorithms have been tested upon the Barton data set. However, given these small sample errors and the variation due to changes in secondary structure assignment, we cannot statistically distinguish accuracies separated by less than about two points (Rost and Eyrich, 2001; Przybylski and Rost, 2002). Since the range of reported accuracies is about 65–68% (Cuff and Barton, 1999; Schmidler *et al.*, 2000; Kloczkowski *et al.*, 2002), we are obliged to conclude that many, very different secondary structure prediction algorithms are statistically indistinguishable.

Our HMM model is almost optimal, in the sense that it extracts almost all the available information. Moreover, the

accuracy of our model is approximately the same (within statistical and systematic errors) as the maximum accuracy of a variety of other secondary structure prediction methods that utilize only local sequence–sequence correlations (Chandonia and Karplus, 1996; Cuff and Barton, 1999; Schmidler *et al.*, 2000; Kloczkowski *et al.*, 2002). This suggests that these algorithms are also almost optimal and that the modest prediction accuracy is due to the fundamental lack of local structure information. Conversely, the fact that these diverse, sophisticated prediction algorithms are not able to extract any additional signal from local correlations indicates that we have not overlooked some subtle source of secondary structure information in our analysis of local inter-sequence correlations.

It has been found that the secondary structure prediction accuracy can be enhanced substantially by basing the prediction upon a MSA of homologous protein sequences (Rost and Sander, 1993) rather than just a single sequence. Since protein structure tends to evolve relatively slowly, the MSA essentially represents many semi-independent amino acid sequences, each associated with approximately the same secondary structure sequence. How informative are these additional data? We extended our HMM to handle this evolutionary information [described in Equations (11)–(14)] and tested the model on the MSAs provided with the Barton data set (Cuff and Barton, 2000). This resulted in a three-state accuracy of $72.2 \pm 0.6\%$, an improvement of about six points over the equivalent single-sequence results. Since the information ratio for this data set was $R \approx 1.3$, this modest accuracy increase actually represents a considerable increase in information. This accuracy is similar to reported accuracies of a number of other algorithms tested on this data set (Schmidler *et al.*, 2000; Kloczkowski *et al.*, 2002). It may be that many profile-based secondary structure prediction algorithms are essentially equivalent and that the differing results are due, primarily, to differences in the quality of the input alignments (Rost, 2001; Kloczkowski *et al.*, 2002).

DISCUSSION

Although local inter-sequence information is insufficient to determine secondary structure accurately, such correlations are still useful to statistical tertiary structure prediction algorithms. For example, in protein threading (Jones *et al.*, 1992), a primary sequence is matched to a structural template using an amino acid contact potential and other similar potentials derived from sequence–structure correlations. Recently, the information contained in amino acid contacts was estimated to be about 0.04 bits per contact, or 0.06 bits per residue (Cline *et al.*, 2002), which can be compared with our estimate of 0.16 bits per residue of primary with secondary structure mutual information. Therefore, local structure potentials may be of under-appreciated importance to threading and other similar statistical structure prediction methods. Many such methods do consider secondary structure (McGuffin and

Jones, 2003), but some of these only consider the direct correlation between an amino acid and the secondary structure class at that one residue (Bowie *et al.*, 1991; Alexandrov *et al.*, 1996). By ignoring the correlations between an amino acid and an extended segment of local secondary structure, such methods lose over half of the available local signal and, unlike secondary structure prediction algorithms, are not optimal.

Protein folding is also constrained by the scarcity of local structure information since the mechanism by which information is extracted, using either a computer or physics, is irrelevant. Secondary structure must be predominately determined by non-local interactions, which in turn depend on the overall native fold of the protein. But the native fold cannot be achieved until the native secondary structure has formed. Therefore, protein folding must typically proceed by a cooperative mechanism (Baldwin and Rose, 1999) where secondary and tertiary structures form concurrently. Note, however, that since this conclusion is based upon a statistical analysis, it applies only to proteins on the average and does not preclude particular proteins, or parts of proteins, from folding via a hierarchal mechanism (Baldwin and Rose, 1999) where pre-organized local secondary structure elements collapse successively into ever-larger structures. For example, it has been suggested that the B-domain of staphylococcal protein A (Myers and Oas, 2001), a small single-domain protein, can fold extremely quickly because of its strongly defined native secondary structure, which persists even in the unfolded state. If this is a general property of fast folding proteins, then the widely divergent folding rates of single-domain proteins may be strongly correlated with the accuracy to which a particular protein's secondary structure can be predicted from the primary sequence.

There are at least two approaches to prediction that aim to circumvent the lack of local structure information. One is to utilize evolutionary information. Since protein sequences evolve more rapidly than protein structure, a MSA of a homologous family represents many semi-independent sequence samples of approximately the same protein structure. Local structure prediction quality is then limited by the size of the family, the divergence of structure across the family and the quality of the alignment. This strategy is commonly employed in secondary structure prediction (Rost and Sander, 1993), and improvements in accuracy to about $Q_3 \approx 75\% \pm 3$ are routine (Cuff and Barton, 1999, 2000; Chandonia and Karplus, 1999; Rost, 2001; Eyrich *et al.*, 2003). By modifying our HMM to use evolutionary profiles, we find that even a modest increase in prediction accuracy represents a substantial increase in secondary structure information.

The alternative approach is to incorporate explicitly non-local interactions. This is essentially what threading attempts to do, although the relatively small magnitude of contact potential information suggests that the bulk of non-local information is subtle, and difficult to extract. Of course, in principle we can determine the full three-dimensional

structure of a protein using an atomic detailed molecular simulation. Until this becomes routinely feasible, computational structure determination will have to proceed via less direct, statistical approaches.

ACKNOWLEDGEMENTS

REFERENCES

Alexandrov,N.N., Nussinov,R. and Zimmer,R.M. (1996) Fast protein fold recognition via sequence to structure alignment and contact capacity potentials. *Pac. Symp. Biocomput.*, 53–72.

Baldwin,R.L. and Rose,G.D. (1999) Is protein folding hierarchic? I. Local structure and peptide folding. *Trends Biochem. Sci.*, **24**, 26–33.

Berman,H.M., Westbrook,J., Feng,Z., Gilliland,G., Bhat,T.N., Weissig,H., Shindyalov,I.N. and Bourne,P.E. (2000) The Protein Data Bank. *Nucleic Acids Res.*, **28**, 235–242.

Bowie,J.U., Lüthy,R. and Eisenberg,D. (1991) A method to identify protein sequences that fold into a known three-dimensional structure. *Science*, **253**, 164–170.

Brenner,S.E., Koehl,P. and Levitt,M. (2000) The ASTRAL compendium for protein structure and sequence analysis. *Nucleic Acids Res.*, **28**, 254–256.

Chandonia,J.-M. and Karplus,M. (1995) Neural networks for secondary structure and structural class predictions. *Protein Sci.*, **4**, 275–285.

Chandonia,J.-M. and Karplus,M. (1996) The importance of larger data sets for protein secondary structure prediction with neural networks. *Protein Sci.*, **5**, 768–774.

Chandonia,J.-M. and Karplus,M. (1999) New methods for accurate prediction of protein secondary structure. *Proteins*, **35**, 293–306.

Chandonia,J.-M., Walker,N.S., Lo Conte,L., Koehl,P., Levitt,M. and Brenner,S.E. (2002) ASTRAL compendium enhancements. *Nucleic Acids Res.*, **30**, 260–263.

Cline,M.S., Karplus,K., Lathrop,R.H., Smith,T.F., Rogers,R.G. and Haussler,D. (2002) Information-theoretic dissection of pairwise contact potentials. *Proteins*, **49**, 7–14.

Cover,T.M. and Thomas,J.A. (1991) *Elements of Information Theory*. Wiley.

Crutchfield,J.P. and Feldman,D.P. (2003) Regularities unseen, randomness observed: levels of entropy convergence. *Chaos*, **13**, 25–54.

Cuff,J.A. and Barton,G.J. (1999) Evaluation and improvement of multiple sequence methods for protein secondary structure prediction. *Proteins*, **34**, 508–519.

Cuff,J.A. and Barton,G.J. (2000) Application of multiple sequence alignment profiles to improve protein secondary structure prediction. *Proteins*, **40**, 502–511.

Durbin,R., Eddy,S., Krogh,A. and Mitchison,G. (1998) *Biological Sequence Analysis*. Cambridge University Press, Cambridge, UK.

Eddy,S.R. (2001) *HMMER: Profile hidden Markov Models for Biological Sequence Analysis*.

Efron,B. and Tibshirani,R.J. (1993) *An Introduction to the Bootstrap*. CRC Press.

Eyrich,V.A., Przybylski,D., Koh,I.Y., Grana,O., Pazos,F., Valencia,A. and Rost,B. (2003) CAFASP3 in the spotlight of EVA. *Proteins*, **53**(Suppl. 6), 548–560.

Frishman,D. and Argos,P. (1995) Knowledge-based protein secondary structure assignment. *Proteins*, **23**, 566–579.

Frishman,D. and Argos,P. (1997) Seventy-five percent accuracy in protein secondary structure prediction. *Proteins*, **27**, 329–335.

Gough,B. (ed.) (2003) *GNU Scientific Library Reference Manual*, 2nd edn. Network Theory Ltd.

Jones,D.T., Taylor,W.R. and Thornton,J.M. (1992) A new approach to protein fold recognition. *Nature*, **358**, 86–89.

Kabsch,W. and Sander,C. (1983) A dictionary of protein secondary structure. *Biopolymers*, **22**, 2577–2637.

Kloczkowski,A., Ting,K.L., Jernigan,R.L. and Garnier,J. (2002) Combining the GOR V algorithm with evolutionary information for protein secondary structure prediction from amino acid sequence. *Proteins*, **49**, 154–166.

Kschischang,F.R., Frey,B.J. and Loeliger,H.-A. (2001) Factor graphs and the sum-product algorithm. *IEEE Trans. Info. Theory*, **47**, 498.

Lo Conte,L., Brenner,S.E., Hubbard,T.J., Chothia,C. and Murzin,A.G. (2002) SCOP database in 2002: refinements accommodate structural genomics. *Nucleic Acids Res.*, **30**, 264–267.

Matsumoto,M. and Nishimura,T. (1998) Mersenne twister: a 623-dimensionally equidistributed uniform pseudorandom number generator. *ACM Trans. Model. Comput. Simul.*, **8**, 3–30.

McGuffin,L.J. and Jones,D.T. (2003) Improvement of the GenTHREADER method for genomic fold recognition. *Bioinformatics*, **19**, 874–881.

Miller,G.A. (1955) Note on the bias of information estimates. In Quastleri,H. (ed.), *Information Theory in Psychology; Problems and Methods II-B*. Free Press, pp. 95–100.

Moddemeijer,R. (2000) The distribution of entropy estimators based on maximum mean log-likelihood. In Biemond,J. (ed.), *21st Symposium on Information Theory in the Benelux*. pp. 231–238.

Moult,J., Fidelis,K., Zemla,A. and Hubbard,T. (2001) Critical assessment of methods of protein structure prediction (CASP): round IV. *Proteins*, **45**(Suppl. 5), 2–7.

Murzin,A.G., Brenner,S.E., Hubbard,T. and Chothia,C. (1995) SCOP: a structural classification of proteins database for the investigation of sequences and structures. *J. Mol. Biol.*, **247**, 536–540.

Myers,J.K. and Oas,T.G. (2001) Preorganized secondary structure as an important determinant of fast protein folding. *Nat. Struct. Biol.*, **8**, 552–558.

Przybylski,D. and Rost,B. (2002) Alignments grow, secondary structure prediction improves. *Proteins*, **46**, 197–205.

Rabiner,L.R. (1989) A tutorial on hidden markov models and selected applications in speech recognition. *Proc. IEEE*, **77**, 257–248.

Rost,B. (2001) Review: protein secondary structure prediction continues to rise. *J. Struct. Biol.*, **134**, 204–218.

Rost,B. and Eyrich,V.A. (2001) EVA: large-scale analysis of secondary structure prediction. *Proteins*, **45**(Suppl. 5), 192–199.

Rost,B. and Sander,C. (1993) Prediction of protein secondary structure at better than 70% accuracy. *J. Mol. Biol.*, **232**, 584–599.

Schmidler,S.C., Liu,J.S. and Brutlag,D.L. (2000) Bayesian segmentation of protein secondary structure. *J. Comput. Biol.*, **7**, 233–248.

Szent-Gyorgyi,A.G. and Cohen,C. (1957) Role of proline in polypeptide chain configuration of proteins. *Science*, **126**, 697–698.

Thompson,M.J. and Goldstein,R.A. (1997) Predicting protein secondary structure with probabilistic schemata of evolutionarily derived information. *Protein Sci.*, **6**, 1963–1975.

Weiss,O., Jimenez-Montano,M.A. and Herzel, H. (2000) Information content of protein sequences. *J. Theor. Biol.*, **206**, 379–386.