

Bayesian Protein Structure Alignment

Abel Rodriguez

Department of Applied Mathematics and Statistics
University of California Santa Cruz

Scott C. Schmidler

Department of Statistical Science
Duke University

March 5, 2009

Abstract

The analysis of the three dimensional structure of proteins is an important topic in molecular biochemistry. Structure plays a critical role in defining the function of proteins, and is more strongly conserved than amino acid sequence over evolutionary timescales. A key challenge is the identification and evaluation of structural similarity between proteins; such analysis can aid in understanding the role of newly discovered proteins, and help elucidate evolutionary relationships between organisms. Computational biologists have developed many clever algorithmic techniques for comparing protein structures; however, all are based on heuristic optimization criteria making statistical interpretation somewhat difficult. Here we present a fully probabilistic framework for pairwise structural alignment of proteins. Our approach has several advantages, including the ability to capture alignment uncertainty, and to estimate key 'gap' parameters which critically affect the quality of the alignment. We show that several existing alignment methods arise as *maximum a posteriori* estimates under specific choices of prior distributions and error models. Our probabilistic framework is also easily extended to incorporate additional information - we demonstrate this by inclusion of primary sequence information to generate simultaneous sequence-structure alignments that can resolve ambiguities obtained using structure alone. This combined model also provides a natural approach for the difficult task of estimating evolutionary distance based on structural alignments. The model is illustrated by comparison with well-established methods on several challenging protein alignment examples.

1 Introduction

Protein alignment is among the most powerful and widely used tools available for inferring homology and function of gene products, as well as determining evolutionary relationships between organisms. In particular, protein sequence alignment uses information about the identity of amino acids to establish regions of similarity, and has a long history of providing valuable insights. For example, the alignment₁ of a putative human colon cancer gene with

a yeast mismatch repair gene played a crucial role in its identification and characterization (Bronner, 1994; Papadopoulos, 1994; Zhu et al., 1998).

Sequence alignment is most useful for shorter evolutionary distances, when amino acid composition has not drifted dramatically from a common ancestor. However, when comparing proteins that are distantly related, sequence conservation may be too dilute to establish meaningful relationships. Because a protein's function is largely determined by its three dimensional structure, and significant sequence mutation can occur while maintaining this structure, it is widely recognized that structural similarity is conserved over much longer evolutionary timescales than sequence similarity. In addition, sequence alignment cannot detect convergent evolution, when proteins with similar 3D structure and carrying out similar functions have evolved from unrelated genes.

Aligning 3D structures requires choosing which amino acids to match as in sequence alignment, but has the added complexity of handling coordinate frames arising from arbitrary rotation and translation. Early work in structural alignment (Rao and Rossmann, 1973; Rossmann and Argos, 1975, 1976) developed techniques that iterate between a rigid body registration and an alignment step, and Satow et al. (1986) introduced the use of Dynamic Programming (applied to sequence alignment by Needleman and Wunsch (1970)) as an efficient way to construct the alignment given a registration. Similar methods have been adopted by many authors Cohen (1997); Gerstein and Levitt (1998); Wu et al. (1998). Most work uses penalized root mean squared deviation (RMSD) between corresponding backbone α -carbon (C_α) atoms to measure quality of the alignments, but several other measures have been proposed, including soap-bubble surface metrics (Falicov and Cohen, 1996), differential geometry (Kotlovski et al., 2003), and heuristic rules like the SSAP method of (Taylor and Orengo, 1989).

An alternative to iterative methods is the use of distance geometry avoid the registration problem, representing each protein by a pairwise distance matrix between all C_α atoms; the popular DALI (Holm and Sander, 1993) method is an example of this approach. Other techniques are specially tailored for the large-scale computational demands of rapid searching of large protein databases, sometimes employing highly redundant representations of the data; these include geometric hashing (Altschul et al., 1990; Fischer et al., 1994; Wallace et al., 1996), graph algorithms (Taylor, 2002), and clustering methods like VAST (Gibrat et al., 1996). And some authors combine these ideas with additional heuristics produce faster or more accurate algorithms, including CE (Shindyalov and Bourne, 1998) and PROSUP (Lackner et al., 2000). Detailed reviews on pairwise structural alignment methods can be found in Brown et al. (1996), Eidhammer et al. (2000) and Lemmen and Lengauer (2000).

The profusion of methods shows the difficulties involved in performing structural alignments: in defining how to measure alignment quality, and in computing 'best' alignments efficiently. It has been well documented in the literature that different algorithms can produce alignments sharing very few amino acid pairings, and are sensitive to both the initial alignment and the specific choice of algorithm parameters Godzik (1996); Zu-Kang and Sippl (1996); Gerstein and Levitt (1998). Additional complications arise when trying to determine the significance of different alignments. Although significant effort has been devoted to this point and important progress made (Lipman and Pearson, 1985; Mizuguchi and Go, 1995; Levitt and Gerstein, 1998; Gerstein and Levitt, 1998), the solutions remain based on heuristics and upper bounds that are difficult to interpret. Finally, all the methods

described above approach the structural alignment as an optimization problem, finding a single best alignment. However, structural comparisons obtain substantial uncertainty arising from evolutionary divergence, experimental measurement error and protein conformational variability, and parameter choices for comparison metrics and optimization algorithms. To address these sources of variability, approaches based on explicit statistical modeling are desirable, and the results of structural comparisons require careful analysis to understand the impact of uncertainty.

In this paper, we develop a full Bayesian statistical approach to pairwise protein structure alignment, combining techniques from statistical shape analysis (Dryden and Mardia, 1998a; Small, 1996; Kendall et al., 1999) and Bayesian sequence alignment (Zhu et al., 1998; Webb et al., 2002; Liu and Lawrence, 1999). This represents one aspect of a general Bayesian framework developed here and elsewhere (Schmidler, 2003, 2004, 2006, 2007; Wang and Schmidler, 2008). Green and Mardia (2006) have independently developed a related approach for hierarchical Bayesian alignment of protein active sites. However, our approach differs in a number of important points: we introduce hierarchical priors on the space of alignments that are equivalent to the standard affine gap penalty of classical alignment approaches, but allow us to estimate the parameters controlling the complexity of the alignment. We also introduce an efficient computational approach that allows rapid computation and both enables identification of alternative alignments and provides measures of their uncertainty. A significant advantage of our formulation is the unification of many existing alternative methods for structural alignment, which can be seen as special cases of MAP alignment under different specific choices of error models or alignment priors. This provides valuable insight into the relationships and properties of existing algorithms.

Another powerful advantage of a fully probabilistic framework is the ability to incorporate disparate sources of information in a natural and coherent fashion. Using our Bayesian structural alignment model as a platform, we also develop a fully probabilistic approach for *simultaneous sequence-and-structure* alignment, which combines information from both primary sequence *and* 3D structure. In the absence of unambiguous geometric matching for highly-divergent proteins or low-resolution structural data, amino acid identities or preferred substitutions can significantly alleviate the remaining uncertainty. We demonstrate this approach on some difficult structural alignment problems from the literature. Finally, we show that our simultaneous alignment approach provides a natural method for estimating evolutionary distances directly from structure comparison, a notoriously difficult task.

2 Bayesian protein structure alignment

Let $X_{n \times 3}$ and $Y_{m \times 3}$ be coordinate matrices for two proteins, with rows x_i (y_i) containing coordinates of the C_α of the i^{th} amino acid. Denote an alignment between X and Y by the match matrix M such that $[M]_{ij} = 1$ if residues X_i and Y_j are matched, and 0 otherwise. Also, denote by X_M and Y_M the $|M| = \sum_{ij} [M]_{ij}$ non-zero rows of $M'X$ and $M'MY$, giving the coordinates of the matched residues. Since each position in X can be matched to at most one position in Y , each row and column of M contains at most one non-zero entry; M is the adjacency matrix for a bipartite graph.

We adopt a Bayesian approach to structure alignment which defines a prior distribution

on alignments $P(M)$ and, given a probability model for the coordinates matrices X and Y conditional on M , obtains the posterior distribution:

$$P(M | X, Y) = \frac{P(X, Y | M)P(M)}{\sum_M P(X, Y | M)P(M)}$$

where the marginal likelihood $P(X, Y) = \sum_M P(X, Y | M)P(M)$ involves a sum over all possible alignments. Although the number of matchings is exponential in n and m , inferences may be obtained by Monte Carlo sampling from the posterior $P(M | X, Y)$ to approximate posterior summaries such as the posterior mode

$$\hat{M} = \arg \max_M P(M | X, Y) \quad (1)$$

or marginal alignment matrix (p_{ij}) for $p_{ij} = \sum_M m_{ij}P(M | X, Y)$.

2.1 Likelihood

We adopt a probabilistic model for matched regions of the proteins given by

$$[Y_M] = [X_M] + \epsilon \quad \epsilon \sim \mathbf{N}(\mathbf{0}, \sigma^2 I) \quad (2)$$

where $[X]$ denotes the *size-and-shape* of X , formally defined (Dryden and Mardia, 1998b; Kendall et al., 1999) as an equivalence class of invariant matrices under the group of Euclidean transformations:

$$[X] = \{XR + \mu : R \in \mathbf{SO}(3) \ \mu \in \mathbf{R}^3\}.$$

Here $\mathbf{SO}(3)$ is the special orthogonal group of 3×3 rotation matrices. Alignment using non-rigid transformations is described elsewhere Schmidler (2007).

The likelihood is then given by the joint density:

$$\begin{aligned} P(X, Y | M) &= P(Y_M | X_M)P(X)P(Y_{\bar{M}}) \\ &= (2\pi\sigma^2)^{-\frac{3|M|}{2}} \exp -\frac{1}{2\sigma^2} \|Y_M - (X_M \hat{R} + \mathbf{1}\hat{\mu}')\|_F^2 P(X) \prod_{y_i \in Y_{\bar{M}}} f(y_i | \lambda) \end{aligned} \quad (3)$$

where $\|X\|_F = \text{tr}(X'X)$ is the Frobenious norm, $f(\cdot; \lambda)$ is a one-parameter density for inserted/deleted positions, $P(X)$ the probability distribution describing the shape of the reference protein X . Here $(\hat{R}_M, \hat{\mu}_M)$ are the optimal rotation and translation placing X and Y on a common coordinate system, obtained by least-squares calculations $\hat{R} = UV^T$ and $\hat{\mu} = \bar{Y}_M - \bar{X}_M \hat{R}$ with $V, U \in \mathbf{SO}(3)$ obtained from singular value decomposition $Y_M^T C X_M = UDV^T$ where $C = [I - \frac{1}{L}11^T]$ is a centering matrix.

The likelihood (3) may be interpreted as the *profile* likelihood for M , maximizing over nuisance parameters (μ, R) conditional on M . For most structural alignments, uncertainty on the (μ, R) given M is minimal. Alternatively, parameters (μ, R) may be assigned prior distributions and integrated out. Green and Mardia (2006) and Wang and Schmidler (2008) adopt this approach, and Schmidler (2006, 2007) discuss both approaches.

We may also interpret (3) as a sampling density defined directly on (a local tangent space approximation to) the underlying shape space of the configurations, replacing $3|M|$ with

$3|M| - 6$, the dimension of the shape manifold, in the normalizing constant. The exponent $\|Y_M - (X_M \hat{R}_M + \mathbf{1} \hat{\mu}'_M)\|_F^2 = d_P^2(X, Y)$ is known as the (squared) *partial Procrustes distance*, and serves as the Riemannian metric on this (size-and) shape space (Dryden and Mardia, 1998a; Kendall et al., 1999). Therefore (3) effectively *defines* a rotation/translation pair associated with every matching M , allowing inference to proceed over the space of possible M only rather than (M, R, μ) .

In what follows, we take $f(\cdot|\lambda) = \lambda = 1/|\Omega|$ uniform over a bounded region Ω ; then λ can be interpreted as a baseline gap penalty as discussed in Section 2.4. Note that the factorization $P(X, Y) = P(Y | X)P(X)$ means the marginal distribution $P(X)$ cancels in the posterior distribution, and may be left unspecified so long as it is assumed independent of parameters M and σ^2 . This is similar to a proportional hazards model where the baseline risk is left unspecified to obtain a semiparametric survival model. In addition, the isotropic error model ensures the model is symmetric in X and Y if we take

$$P(X) = \prod_{x_i \in X} f(x_i|\lambda).$$

2.2 Prior on the alignment matrix

Prior distributions on matchings $P(M)$ may be specified in a variety of ways; here we adopt an *gap-penalty* formulation familiar in the sequence and structure alignment literature, where unmatched stretches of amino acids are penalized by the affine function:

$$u(M; g, h) = g s(M) + h \sum_{i=1}^{s(M)} l_i(M)$$

with gap-opening penalty g and gap-extension penalty h , where $s(M)$ is the number of gaps in alignment M and $l_i(M)$ is the length of the i^{th} gap. Exponentiating and normalizing this function provides a prior on M (Liu and Lawrence, 1999), essentially a Markov chain parametrized as a 'Boltzmann chain' Gibbs random field. (Saul and Jordan, 1995; Schmidler et al., 2007).

$$P(M | g, h) = Z(g, h) e^{-u(M; g, h)} \tag{4}$$

with normalizing constant Z . This prior encourages grouping of matches together along the protein backbone. It allows for explicit control over the number of gaps, compared to e.g. the prior of Green and Mardia (2006) which only controls the expected total length.

Under the affine-gap-penalty prior, sampling may be done efficiently using stochastic recursions analogous to those of standard sequence alignment algorithms (Liu and Lawrence, 1999), along with additional Monte Carlo steps, as described below. Note that this prior requires the alignment to be preserve the sequential order along the polypeptide backbone, requiring topological equivalence of the two proteins. More general priors applicable for comparing proteins of potentially different topologies (convergent evolution) are easily accommodated with the introduction of additional Monte Carlo steps, but will be described elsewhere. Although the prior allows for simultaneous gaps on both proteins, for identifiability purposes we do not allow gaps in X to be followed by gaps in Y (see Webb et al. (2002) for details).

2.3 Hyperpriors

In standard sequence and structure alignment algorithms, the gap parameters g and h are assigned fixed values. However, they have a critical effect on the resulting alignment, with large opening gap penalties g tending to produce alignments with few gaps and vice versa. In the context of sequence alignment, Liu and Lawrence (1999) treat (g, h) as nuisance parameters and assign hyperpriors, integrating them out to obtain a marginal posterior distribution over alignments. We similarly assign g and h Gamma hyperpriors

$$g \sim \text{Ga}(a_g, b_g) \quad h \sim \text{Ga}(a_h, b_h) \quad (5)$$

with hyperparameters (a_g, b_g, a_h, b_h) chosen to be diffuse. An alternative is to utilize manually-obtained reference alignments (for example, BAliBASE, see Thompson et al. (1999) and Thompson et al. (2005)) to obtain informative priors for g and h . The model is completed by specifying inverse-gamma prior $\sigma^2 \sim \text{IGa}(a_\sigma, b_\sigma)$ on variance parameter σ^2 .

2.4 Many existing structure alignment algorithms are special cases

Rather than summarize the posterior $P(M | X, Y)$ by Monte Carlo sampling, we may instead obtain a single MAP alignment (1). by maximizing the (log-) posterior. Conditioning on parameters $\theta = (\sigma^2, g, h, \lambda)$, we obtain

$$\begin{aligned} \log(P(M | X, Y, \theta)) = & -\frac{3}{2}|M| \log(2\pi\sigma) - \frac{1}{2\sigma^2} d_P^2(X_M, Y_M) + (n + m - |M|) \log(\lambda) \\ & + \log(Z(g, h)) - u(M; g, h) \quad (6) \end{aligned}$$

and noting that $\sum_{i=1}^{s(M)} l_i(M) = (n + m) - 2|M|$ this is equivalent to minimizing

$$d_P^2(X_M, Y_M) + u(M; g^*, h^*) + C(\sigma^2, \lambda, g, h)$$

where $g^* = \sigma^2(g + \frac{3}{2} \log(2\pi\sigma) + \log \lambda)$ and $h^* = \sigma^2 h$, and $C(\sigma^2, \lambda, g, h)$ is independent of M . Therefore the MAP estimate for M with $(g, h, \lambda, \sigma^2)$ fixed corresponds to a global alignment obtained via dynamic programming (Needleman and Wunsch, 1970), using RMSD under optimal least-squares rotation/translation as the dissimilarity metric, with gap opening and extension penalties given by g^* and h^* . Here $(\frac{3}{2} \log(2\pi\sigma) + \log \lambda)$ serves as a lower bound on the gap extension penalty.

Note that relative posterior probabilities of two alignments differing by an unmatched pair (x_i, y_j) is greater than one if and only if

$$|y_j - (x_i R + \mu')| < g^*(1 - \xi_{ij}) + h^* \xi_{ij}$$

where ξ_{ij} is an indicator taking value 1 if removing the pair (x_i, y_j) creates a new gap in the alignment and 0 otherwise. Thus the model favors inclusion of pairs below a dynamically estimated threshold given by g^* and h^* . Since these threshold parameters are estimated from the data (compared to standard optimization-based alignment algorithms where they are fixed *a priori*), our approach automatically controls for the error rates associated with multiple comparisons.

It is also worth noting that the assumption of normally distributed errors in (2) may be replaced with an alternative error model, altering the d_P term in (3) and the corresponding posterior distribution. In particular, robust error models with heavy tails may be considered (e.g. Student-t or double exponential distributions) to account for possible outliers. In this way, our probabilistic formulation provides statistical insight into various existing optimization-based algorithms.

For example, Gerstein and Levitt (1998) define the similarity between residues x_i and y_j by

$$S_{ij} = \frac{c}{1 + \left(\frac{d_{ij}}{d_0}\right)^2}$$

where d_{ij} denotes the distance between i and j under the current optimal registration and c and d_0 are arbitrarily chosen constants. Then dynamic programming is employed to obtain the alignment M maximizing the similarity between proteins, defined by $\sum_{(i,j) \in M} S_{ij}$. This is equivalent to obtaining the MAP estimate under our Bayesian model when the distribution of the error ϵ is given by $f(x) \propto \exp\left\{-M\left(1 + (x/d_0^2)\right)^{-1}\right\}$, which is an exponentiated Cauchy density. (Note that $f(x)$ is indeed a proper density as $\int_{-\infty}^{\infty} \exp\left\{-\frac{M}{1+(x/d_0)^2}\right\} dx < \infty$). Thus our unified probabilistic framework allows us to interpret such heuristics in terms of their underlying assumptions about the data generating process.

3 Computational algorithms

Combining (3), (4) and (5) we obtain the posterior distribution:

$$P(M, g, h, \sigma^2 | X, Y) \propto P(X, Y | M, \sigma^2) P(M | g, h) P(\sigma^2) P(g) P(h)$$

This posterior can be explored using a Markov chain Monte Carlo algorithm that iterates between sampling from the conditional distributions, $P(M | g, h, \sigma^2, X, Y)$, $P(g, h | M, X, Y)$ and $P(\sigma^2 | M, X, Y)$. The variance σ^2 is obtained by standard conjugate update:

$$\sigma^2 | M, X, Y \sim \text{IGa} \left(a_\sigma + \frac{3}{2}|M|, b_\sigma + \frac{1}{2}d_P^2(X_M, Y_M) \right).$$

The gap penalty parameters (g, h) are updated jointly by a two-dimensional geometric random walk proposal with Metropolis-Hastings acceptance probability $\gamma((g, h), (g', h'))$ equal to

$$1 \wedge \frac{Z(g', h') e^{-u(M; g', h')}}{Z(g, h) e^{-u(M; g, h)}} \frac{g' h'}{g h} (g'/g)^{a_g - 1} (h'/h)^{a_h - 1} e^{-(b_g(g' - g) + b_h(h' - h))} \quad (7)$$

The normalizing constants $Z(g, h)$ required for (7). can be calculated efficiently via the recursions provided in Appendix A.

As shown by Schmidler (2003), if we *condition* on registration parameters (R, μ) , the alignment matrix M may be sampled from its full conditional distribution using a forward-backward algorithm similar to that of sequence alignment (Zhu et al., 1998; Liu and

Lawrence, 1999) and described in Appendix A. Wang and Schmidler (2008) use this approach for structural alignment. However, here we have instead defined the likelihood (3) directly on shape space using maximal values $(\hat{R}, \hat{\mu})$ associated with each distinct M , so this is no longer the case. But we may still use this efficient block Gibbs step to generate efficient Metropolis-Hastings proposals $P(M \rightarrow M')$ with distribution $q(M'; R_M, \mu_M)$ where (R_M, μ_M) is the registration associated with the *current* state M , and

$$q(M'; R, \mu) \propto P(M') (2\pi\sigma^2)^{-\frac{3|M'|}{2}} e^{-\frac{1}{2\sigma^2} \|Y_{M'} - \hat{X}_{M'}\|_F^2} \prod_{y_i \in Y_{M'}} f(y_i|\lambda)$$

where $\hat{X}_{M'} = X_{M'}\hat{R} + \mathbf{1}\hat{\mu}'$. This q can be sampled efficiently using the recursions of Appendix A. The proposed M' is then accepted according to the Metropolis-Hastings criteria

$$\gamma = \min \left\{ 1, \frac{\pi(M')q(M; R_{M'}, \mu_{M'})}{\pi(M)q(M'; R_M, \mu_M)} \right\}.$$

with the necessary normalizing constants of q obtained from the sampling recursions.

These dynamic programming proposals are highly efficient for local sampling, and sufficient for closely matched proteins. However, when multiple alternative alignments with distinct rotation/translations exist, mixing between them will be slow. We therefore add an additional Metropolized independence step where global moves are proposed without conditioning on (R, μ) . To construct the independence proposal distribution, we generate a library of viable registrations using the following procedure:

1. Compute the least-squares registration for each pair of consecutive 6-residue subsequences on protein X to each such subsequence on Y .
2. If the subsequence RMSD is less than threshold δ , include the corresponding registration in the library.

The library is computed only once when the algorithm is initialized, and stored for use throughout the simulation. At each independence proposal a registration (R', μ') is drawn uniformly at random from the library, and a new alignment M' proposed from $q(M'; R', \mu')$, again using the forward-backward algorithm, and accepted according to the Metropolis-Hastings criteria

$$\gamma = \min \left\{ 1, \frac{\pi(M')q(M; R', \mu')}{\pi(M)q(M'; R', \mu')} \right\}.$$

4 Bayesian synthesis of sequence and structure information

Another advantage of the Bayesian probabilistic framework given above is the ability to seamlessly incorporate additional information when available. For example, our approach leads to a natural algorithm for performing alignments based on primary sequence and tertiary structure simultaneously. This allows alignments which synthesize two types of information - geometric conservation of the protein architecture and physico-chemical properties and evolutionary information provided by sidechain identities. As an important consequence, our approach enables the estimation evolutionary distances from *structure comparison*, which has previously been quite difficult (Chothia and Lesk, 1986; Johnson et al.,

1990; Grishin, 1997; Levitt and Gerstein, 1998; Wood and Pearson, 1999; Koehl and Levitt, 2002). This has important implications because structure is much more strongly conserved than sequence, enabling comparisons across much longer evolutionary timescales.

The model given by (3) for structural observations is easily extended to account simultaneously for both sequence *and* structure information by assuming the structure and sequence to be conditionally independent given the alignment M , i.e. $P(A^x, A^y, X, Y | M, \theta) = P(A^x, A^y | M, \theta)P(X, Y | M, \theta)$. We take the conditional likelihood of the sequences given the alignment to be:

$$P(A^x, A^y | M, \Theta) = \prod_{(i,j) \in M} \Theta(A_i^x, A_j^y) \prod_{i \notin M} \Theta(A_i^x, \cdot) \prod_{j \notin M} \Theta(\cdot, A_j^y) \quad (8)$$

where A_i^y is the i -th amino acid in protein x , $\Theta(a, b)$ gives the probability of residues a and b being matched on related sequences and $\Theta(a, \cdot) = \Theta(\cdot, a)$ gives the marginal probability for residue a . This is the standard likelihood form of sequence alignment (Bishop and Thompson, 1986), and these joint and marginal distributions are the basis of standard sequence alignment substitution matrices such as PAM and BLOSSUM (Dayhoff and Eck, 1968; Henikoff and Henikoff, 1992), where the distributions are estimated from alignments of closely related proteins. For example the PAM- k substitution can be written as $\Psi_k = (\Psi_k(a, b))$ where

$$\Psi_k(a, b) = 10 \log_{10} \left(\frac{\Theta_k(a, b)}{\Theta(a, \cdot)\Theta(\cdot, b)} \right)$$

Sequence alignment may then be formulated as a maximum-likelihood or Bayesian inference problem (Durbin et al., 1998; Zhu et al., 1998; Liu and Lawrence, 1999; Bishop and Thompson, 1986), including inferences on k which estimate the evolutionary distance between the two protein sequences.

Multiplication of equations (3) and (8) directly yields a model for simultaneous inference on M which combines both sequence and structure information. However as pointed out structure is generally much more strongly conserved than sequence; thus we would like to weight the contribution of structure information in determining the alignment more heavily than that of sequence. In this way the sequence alignment information will server primarily to provide supplementary information in portions of the alignment where structural information leaves uncertainty; as we will see it also permits the estimation of evolutionary distance from the largely structure-based alignment.

To control the relative weighting of sequence and structure information we introduce a discount (or inverse temperature) parameter η , which results in the modified sequence likelihood

$$\begin{aligned} \Pr(A^x, A^y | M, \Theta, \eta) &= \frac{\prod_{(i,j) \in M} \Theta(A_i^x, A_j^y)^\eta \prod_{i \notin M} \Theta(A_i^x, \cdot)^\eta \prod_{j \notin M} \Theta(\cdot, A_j^y)^\eta}{\sum_{A^{x*}, A^{y*}} \prod_{(i,j) \in M} \Theta(A_i^{x*}, A_j^{y*})^\eta \prod_{i \notin M} \Theta(A_i^x, \cdot)^\eta \prod_{j \notin M} \Theta(\cdot, A_j^y)^\eta} \\ &= \prod_{(i,j) \in M} \frac{\Theta(A_i^x, A_j^y)^\eta}{\sum_{A_r, A_s} \Theta(A_r, A_s)^\eta} \prod_{i \notin M} \frac{\Theta(A_i^x, \cdot)^\eta}{\sum_{A_r} \Theta(A_r, \cdot)^\eta} \prod_{j \notin M} \frac{\Theta(\cdot, A_j^y)^\eta}{\sum_{A_s} \Theta(\cdot, A_s)^\eta} \end{aligned}$$

Note that $\eta = 1$ corresponds to simple multiplication of the sequence and structure likelihoods (3) and (8). However, as $\eta \rightarrow 0$, $\Pr(A^x, A^y | M, \Theta, \eta)$ approaches a uniform distribution

for every Θ , with $\eta = 0$ reducing to the structure-only model (8). In fact, we can consider estimating η itself, generally restricting $\hat{\eta} < 1$. Then $\hat{\eta}$ may be interpreted as a measure of agreement between sequence and structure, and may shrink to downweight the importance of sequence matching when it is clearly incompatible with the structural information.

5 Examples

We apply our Bayesian structural alignment algorithm to a number of illustrative examples. Hyperparameter values used are given in Table 1: the prior distribution for σ^2 has mean 1.5\AA and variance 1.0\AA , in line with the results for analogous proteins in Chothia and Lesk (1986), and the prior mean for h is about 40 times larger than the prior mean for g following Gerstein and Levitt (1998). Results were robust to moderate changes in these hyperparameters. All inferences described are based on 100,000 samples obtained after a burn-in period of 20,000 iterations, with convergence verified by visual inspection of the trace plots and using the Gelman-Rubin convergence test (Gelman and Rubin, 1992). Monitored quantities include the length of the alignment, the rotation angles corresponding to rotation matrix \hat{R}_M , the translation vector $\hat{\mu}_M$ and the two gap penalty parameters (g, h). We report MAP alignments unless otherwise noted.

| a_σ | b_σ | a_h | b_h | a_g | b_g |
|------------|------------|-------|-------|-------|-------|
| 2.25 | 1.5 | 2 | 1/2 | 2 | 20 |

Table 1: Hyperparameter values used in the examples.

We first analyze 16 pairs of proteins from Ortiz et al. (2002). This list includes pairs of very different lengths and proteins from various structural classes, including mainly α , mainly β , and $\alpha + \beta$. Table 2 summarizes the results obtained using three different values for λ ranging from a relatively low (7.6) to the relatively high (9.6), and compares the Bayesian alignments against those obtained using the popular CE algorithm (Shindyalov and Bourne, 1998). In most cases, the differences between Bayesian and CE alignments are important; in more than half the examples less than 20% of the matched residues coincide. These differences are mostly due to the way CE handles gaps: to reduce the computational complexity, CE assumes that gaps cannot be introduced simultaneously in both proteins. Similar restrictions can be easily introduced in our model by setting $q_{i,j}(2,3) = 0$, and when this is done the results for both methods tend to agree. Generally speaking, the added flexibility means that the quality of the Bayesian alignments is superior to CE: it tends to produce alignments containing more matched residues but with a lower RMSD.

Some pairs of proteins seem to be somewhat sensitive to the choice of λ (for example, 1ACX-1COB), while the alignment of other pairs seem to be remarkably robust (for example, 1UBQ-FRD). In general, larger values of λ (which imply larger penalties for both opening and extension) tend to generate longer alignments. The sensitivity of the model to λ is not surprising; indeed, setting $\lambda = 0$ immediately implies that the optimal alignment is empty for any pair of proteins. The model is robust to small changes in λ , which can be absorbed by adapting on the value of g and h . However, when λ is increased by a large

amount, we set a new baseline threshold that pairs of residues need to satisfy in order to be included in the alignment, favoring the alignment of sections that were previously not close enough to be aligned. When structural similarity varies dramatically along the alignment (as is the case for some pairs in our list), this “threshold effect” can produce important changes in the resulting alignment. In general we can think of λ as controlling the tightness of the alignment. In our our experience, a conservative value such as $\lambda = 7.6$ works well in most applications, and we use this value in further illustrations.

Table 2 also shows the posterior median of the gap penalties for each alignment. Opening penalties range between 5 and 9, while extension penalties range from 0.01 to nearly 0.9, reflecting the differing levels of similarity across different pairs.

Next, we consider in detail the alignment of two α proteins from the globin family, 5MBN and 2HBG. Figure 1 presents both the marginal alignment matrix (which provides information on the uncertainty associated with the alignment) and the MAP alignment, comparing it against that obtained from CE. The most striking feature about this example is that different alignment methods tend to disagree in regions where the uncertainty in the Bayesian alignment is high (the regions surrounding the gap between residues 47 and 62 of 5MBN:–, the extremes of the helix between residues 81 and 98 of 5MBN:–, and at the very end of the alignment). This is an additional argument for using a probabilistic alignment framework, rather than relying on a single optimum. In Figure 2 we show the prior and posterior distributions associated with both gap penalty parameters in this example, which demonstrate that the model does learn about and adapt these parameters.

Finally we explore the alignment of the $\alpha + \beta$ proteins 1CEW:I and 1OUN:A. Lackner et al. (2000) describe two alternative alignment for these proteins having a comparable number of equivalent residues (70 vs. 68) and RMSD (2.4Å both), which arise by shifts in the alignment of the secondary structures. Figure 3 shows the marginal alignment probabilities for all pairs of residues. Unlike the previous example, uncertainty levels in this alignment are very high, particularly in the α -helix region between residues 10 and 20. The two alternative alignments for this region correspond to the two alignments described in Lackner et al. (2000). However, the alignment of the rest of the proteins corresponds to the 70 residue alignment discussed by authors. This example shows how the global sampling of the full posterior enables the model to automatically weight the relative importance of closely related alternative alignments, and how the estimation of gap penalties can further improve this.

5.1 Combined Sequence-Structure Alignment

To illustrate the performance of our simultaneous sequence-and-structure alignment approach, we consider two pairs of proteins that have been previously analyzed in the literature. For convenience we consider a discrete set of discount factors ranging from 0 to 1 in increments of 0.1, along with 21 PAM matrices ranging from PAM100 to PAM300. Non-informative uniform prior distributions on are used for both discount factors and PAM matrices. All results are based on 130,000 iterations of the Gibbs sampler, after a burn-in period of 30,000 iterations.

In the first example we analyze two kinases studied by Bayesian sequence alignment in Zhu et al. (1998); a guanylate kinase from yeast (1GKY) and an adenylane kinase from the

beef heart mitochondrial matrix(2AK3_A), which are VAST structural neighbors (Gibrat et al., 1996). A structural alignment for these proteins using the combinatorial extension (CE) algorithm (Shindyalov and Bourne, 1998) shows very little sequence similarity (under 13% identity). Figure 4 compares our structural and simultaneous sequence-structure alignments for these two proteins by showing the marginal probability of aligning any pair of residues, integrated over all other parameters in the model (including PAM matrices and discount factors). Both algorithms tend to agree on which regions should be aligned. For example, both avoid aligning the section of the α -helix located between residues 150-162 in 1GKY and residues 175-191 in 2AK3:A (marked III in Figure 4). The axes for these helices are not parallel, producing a big divergence at the C terminus.

In spite of the similarities, some differences are evident among both models. For example, a section of the alignment starting at residue 108 of 1GKY (marked II in Figure 4) is excluded when the sequence information is included in the analysis. Both proteins present a short helix in this region, and they can be structurally aligned reasonably well. However, there are important incompatibilities in the two sequences for these helices, which suggests that this section is not functionally important. Table 3 presents the sequence correspondence associated with the structural alignment of this section, along with the scores for each site. Note that the structural alignment implies no conserved residues in the area and the substitution of various basic and acidic amino acids by either hydrophobic or hydrophilic residues. Indeed, of the eight substitutions, only one happens between members of a common chemical group. This is local discrepancy between sequence and structure which is not seen in other regions of the proteins, and suggests that the region should be dropped from the alignment. Similarly, a couple of short regions in the remote site for mono and triphosphate binding located between residues 35-80 for 1GKY and 38-73 in 2AK3:A (marked I in Figure 4), that show a moderate probability of being aligned under the structural model, are downweighted (but not completely removed) when the sequence information is included. This region, which was probably functionally important in an ancestor, has degraded since both proteins diverged and does not seem functionally active in these proteins. These two minimal changes in the alignment lowers the RMSD from 3.5Å to a median of 1.95Å (with 90% high posterior density interval of (1.84, 2.17)).

Figure 5 shows the marginal posterior probability distribution over PAM matrices that arises from our joint sequence-structure model, contrasting it with the results in Zhu et al. (1998). Whereas the sequence-based analysis in the original paper led to a multimodal posterior with modes at PAMs 110, 140 and 200, our posterior is smooth and unimodal, with its mode located between PAM200 and PAM210. This demonstrates the strong additional information obtained by aligning based on structure and sequence simultaneously: virtually all sequence alignments which are compatible with structural alignment indicate the larger evolutionary distance (posterior mean 212, median 206). The marginal mode for the temperature is 1 (posterior probability 0.57), indicating that there is very little need to discount sequence with respect to structure information.

Our second example focuses on comparing the single-chain fused Monellin from the Serendipity berry (1MOL:A) and the chicken egg white Cystatin (1CEW:I) analyzed previously in Lackner et al. (2000) and Kotlovyyi et al. (2003). Figure 6 shows the Bayesian alignments obtained with and without inclusion of sequence information. Again the two alignments are quite similar as expected, but the sequence information leads to small re-

finements in the structural alignment. For example, two alternative alignments of the initial strand are supported by structure-only alignment, with the one where 1:MOL:A is shifted towards the C terminus being slightly preferred (this is also the one preferred by CE). However, incorporation of sequence information reverses this to prefer the N-terminus shifted alignment (approximate posterior probabilities of 0.85 vs 0.15), and examination of the sequences strongly supports this choice. Table 5 shows the sequence alignment under both alternatives, with amino acids colored by a simple classification according to physico-chemical properties (Table 4) to demonstrate the improved similarity on top of amino acid identity. The sequence-structure alignment yields six matches in amino acid type, including an additional two identities and a hydrophilic match on top of the three hydrophobics achieved by the structural alignment. The corresponding price paid in structural distance (mean RMSD of 1.91Å versus 1.89Å, with both 90% hpd regions being (1.81Å,2.05Å)) is insignificant. This example clearly shows that incorporation of sequence information can refine structural alignments in areas where the structure alignment is ambiguous.

Figure 7 shows the joint posterior distribution over PAM matrices and discount factors for this example. Relative to the previous example, there is more uncertainty in both the evolutionary distance and the discount factor. The diagonal pattern in the plot suggests an obvious dependence between these two parameters. This is to be expected, as both η and evolutionary distance increase the entropy of the joint amino acid distribution. Nevertheless, the results point towards a relatively large divergence time (recall one is a plant protein and the other is an animal protein), with the mode of the distance at 210.

To avoid confounding of PAM and tempering parameters, one parameter may be chosen in advance and fixed. For example, the substitution matrix may be chosen to reflect prior information about evolutionary distance and inference performed only on the *discount factor*, or viceversa. When 1MOL_A and 1CEW_I are aligned using PAM250 as the fixed substitution matrix, the resulting distribution for *discount factor* is very similar: the mode is located at $\eta = 0.6$ with a posterior probability of 0.32, and most the remaining mass concentrates in $\eta = 0.5$ and $\eta = 0.7$, both with posterior probability of 0.24. Differences in the actual alignments are not obvious from the marginal distribution plot (not shown). However, a more detailed look at the values shows that fixing the PAM matrix further decreases the probability of the CE-like alignment below 10%.

The examples show that simultaneous estimation of PAM distance and discount factor may be difficult. Since larger evolutionary distances increase sequence divergence/decrease sequence conservation, both low discount values and high PAM distances imply more tolerance to substitutions. One way to measure substitution tolerance is via the Shannon entropy of the joint distributions (Figure 8). Although increasing η and k both increase this entropy, they do so in slightly different ways. We observe that PAM100 with a temperature of 0.8 has roughly equivalent entropy to untempered PAM200. However, PAM100/0.8 assign a much larger probability of match than does PAM200/1.0, as seen by the darker diagonal. In addition, temperature increases treat all combinations of amino acids in the same way, and thus low probability regions tend to disappear quickly. This is not so for increases in the evolutionary distances. In the limit of k , the PAM joint distribution will converge to the product of independent marginal distributions given by the stationary distribution of the underlying Markov chain (estimated as overall population frequencies). In contrast, as the discount factor approaches 0, the joint distribution and thus the marginal distributions

converge to uniform. Figure 8a also shows that differences between PAM matrices grow weaker as the discount factor decreases. In the extreme case when $\eta = 0$, all matrices are equivalent. Thus restricting or fixing the discount factor is important to improve inference on evolutionary distances.

Finally, it is important to mention that we have not found alternative methodologies in the literature capable of this type of information synthesis, against which to compare our results. One of the few methods available is an extension of the combinatorial extension (CE) method Shindyalov and Bourne (1998), accessible via <http://cl.sdsc.edu/ce.html>. However, in this implementation there is little control on the choice of substitution matrices and, for the examples we have studied, the sequences seems to have little practical influence in the final results.

6 Conclusions

We have presented a unifying probabilistic framework for protein structure alignment, based on Bayesian hierarchical modeling. Computationally efficient MCMC algorithms for sampling the posterior distribution enable us to directly account for uncertainty over alignments, including identification of alternative alignments and evaluation of their relative importance. Our model provides insights into the relations between and assumptions of standard optimization-based alignment techniques, along with a unifying framework that facilitates comparisons between them. It also naturally incorporates additional information, such as the inclusion of sequence information in structural alignments. As a byproduct of the latter, we obtained a model which can estimate evolutionary distance directly from structural alignment, an otherwise difficult task. The examples shown clearly highlight how these advantages of our model aid in identification of functionally relevant regions and in resolving ambiguities in alignments. By introducing a discount parameter, we are able to control the influence of the sequence information on the final alignment, an important characteristic missing in previous attempts to combine sequence and structure. As noted, PAM distance and discount factor are correlated, and inference on evolutionary distance will therefore be more reliable if additional information is used to determine the discount factor; this is an area for additional study. Finally, we feel that that sequence-structure alignments provide the most insight when used in conjunction with structure-only alignments as done in the examples. Comparisons between the two appear to provide more direct information on conservation than do comparisons between structure-only and sequence-only alignments.

Acknowledgments

This work was partially supported by NSF grant DMS-0605141 (SCS).

A Dynamic programming forward-backward sampling

As shown by (Schmidler, 2003), if we *condition* on registration parameters (R, μ) , the alignment matrix M may be sampled from its full conditional distribution using a forward-backward algorithm similar to that of sequence alignment (Zhu et al., 1998; Liu and

Lawrence, 1999). Let $v_{i,j}(k)$ be the probability of the alignment of the i^{th} prefix of X and the j^{th} prefix of Y ending in type k , with $k = 1$ meaning that both final residues are aligned, $k = 2$ inserts a gap in X and $k = 3$ inserts a gap in Y . Then

$$\begin{aligned} v_{i,j}(1) &= \sum_{k=1}^3 q_{i,j}(k, 1) v_{i-1, j-1}(k) & v_{i,j}(2) &= \sum_{k=1}^3 q_{i,j}(k, 2) v_{i-1, j}(k) \\ v_{i,j}(3) &= \sum_{k=1}^3 q_{i,j}(k, 3) v_{i, j-1}(k) \end{aligned}$$

and letting $d_{ij}^2 = \|y_j - (x_i R + \mathbf{1}\mu')\|^2$ the transition weights are given by

$$q_{i,j}(l, k) = \begin{cases} \frac{\lambda}{(2\pi\sigma^2)^{\frac{3}{2}}} \exp\left\{-\frac{1}{2\sigma^2} d_{ij}^2\right\} & k = 1 \\ \exp\{g + h\} & (l, k) = (1, 2) \text{ or } (1, 3) \\ \exp\{g\} & (l, k) = (2, 2) \text{ or } (3, 3) \text{ or } (2, 3) \\ 0 & (l, k) = (3, 2) \end{cases}$$

In order to ensure identifiability of the alignments, we do not allow a gap in Y to follow a gap in X , hence $q_{3,2} = 0$. The initialization of these recursions are

$$\begin{aligned} v_{1,1}(1) &= \frac{\lambda}{(2\pi\sigma^2)^{3/2}} \exp\left\{-\frac{1}{2\sigma^2} d_{11}^2\right\} \\ v_{i,1}(1) &= \frac{\lambda}{(2\pi\sigma^2)^{3/2}} \exp\left\{-\frac{1}{2\sigma^2} d_{1i}^2 + (i-1)g + h\right\} \\ v_{1,j}(1) &= \frac{\lambda}{(2\pi\sigma^2)^{3/2}} \exp\left\{-\frac{1}{2\sigma^2} d_{j1}^2 + (j-1)g + h\right\} \\ v_{1,j}(2) &= \exp\{(j+1)g + h\} \quad \text{and} \quad v_{i,1}(3) = 0 \end{aligned}$$

Note that $v_{n,m}$ contains the sum over all alignments, and given (g, h) the same algorithm with $q_{i,j}(l, 1) = 1$ can be used to efficiently compute the normalizing constant $Z(g, h)$ in the gap-penalty prior (4), as required for the acceptance probability (7). Once $q_{i,j}(k)$ is available for all (i, j) , the alignment is sampled backwards, starting with

$$u_{n,m}(k) = \frac{v_{n,m}(k)}{\sum_{l=1}^3 v_{n,m}(l)}$$

and then conditionally adding a matched pair or a gap on one of the proteins with probabilities.

$$u_{i,j}(k, l) = \frac{q_{i-1, j-1}(l, k) v_{i-1, j-1}(k)}{\sum_{k=1}^3 q_{i-1, j-1}(l, k) v_{i-1, j-1}(k)}$$

References

Altschul, S., Gish, W., Miller, W., Myers, E., and Lipman, D. (1990). Basic local alignment search tool. *Journal of Molecular Biology*, 215:403–410.

- Bishop, M. and Thompson, E. (1986). Maximum likelihood alignment of dna sequences. *Journal of Molecular Biology*, 190:159–165.
- Bronner, C. E. *et al.* (1994). Mutation in the DNA mismatch repair gene homologue hMLH1 is associated with hereditary non-polyposis colon cancer. *Nature*, **368**(258-261).
- Brown, N., Orengo, C., and Taylor, W. (1996). A protein structure comparison methodology. *Computational Chemistry*, 20:359–380.
- Chothia, C. and Lesk, A. M. (1986). The relation between the divergence of sequence and structure in proteins. *The EMBO Journal*, **5**:823–826.
- Cohen, G. (1997). ALIGN: A program to superimpose protein coordinates, accounting for insertions and deletions. *Acta Crystallographica*, 30:1160–1161.
- Dayhoff, M. and Eck, R., editors (1968). *Atlas of Protein Sequence and Structure*, volume 3. National Biomedical Research Fundation, Silver Spring, MD.
- Dryden, I. and Mardia, K. (1998a). *Statistical shape analysis*. Wiley.
- Dryden, I. L. and Mardia, K. V. (1998b). *Statistical Shape Analysis*. Wiley.
- Durbin, R., Eddy, S., Krogh, A., and Mitchison, G. (1998). *Biological Sequence Analysis: Probabilistic Models of Proteins and Nucleic Acids*. Cambridge University Press.
- Eidhammer, I., Jonassen, I., and Taylor, W. (2000). Structure comparison and structure patterns. *Journal of Computational Biology*, 7:685–716.
- Falicov, A. and Cohen, F. (1996). A surface of minimal area metric for the structural comparison of proteins. *Journal of Molecular Biology*, 258:871–892.
- Fischer, D., Wolfson, H., Lin, S., and Nussinov, R. (1994). Three-dimensional, sequence order-independent structural comparison of a serine protease against the crystallographic database reveals active site similitates: Potential implications to evolution and to protein folding. *Protein Science*, 3:768–778.
- Gelman, A. and Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science*, **7**:457–472.
- Gerstein, M. and Levitt, M. (1998). Comprehensive assessment of automatic structural alignment against a manual standard, the scop classification of proteins. *Protein Science*, 7:445–456.
- Gibrat, J.-F., Madej, T., and Bryant, S. (1996). Surprising similarities in structure comparison. *Current Opinion in Structural Biology*, 6:377–385.
- Godzik, A. (1996). The structural alignment between two proteins: is there a unique answer? *Protein Engineering*, 5:1325–1338.
- Green, P. J. and Mardia, K. V. (2006). Bayesian alignment using hierarchical models, with applications in protein bioinformatics. *Biometrika*, **93**:235–254.


- Grishin, N. V. (1997). Estimation of evolutionary distances from protein spatial structures. *Journal of Molecular Evolution*, **45**:359–369.
- Henikoff, S. and Henikoff, J. (1992). Amino acid substitution matrices from protein blocks. *Proceedings of National Academy of Science*, **89**:10915–10919.
- Holm, L. and Sander, C. (1993). Protein structure comparison by alignment of distance matrices. *Journal of Molecular Biology*, **233**:123–138.
- Johnson, M. S., Sutcliffe, M. J., and Blundell, T. L. (1990). Molecular anatomy: Phyletic relationships derived from three-dimensional structures of proteins. *Journal of Molecular Evolution*, **30**:43–59.
- Kendall, D. G., Barden, D., Carne, T. K., and Le, H. (1999). *Shape and Shape Theory*. Wiley.
- Koehl, P. and Levitt, M. (2002). Sequence variations within protein families are linearly related to structural variations. *Journal of Molecular Biology*, **323**:551–562.
- Kotlovyy, V., Nichols, W., and Ten Eyck, L. (2003). Protein structural alignment for detection of maximally conserved regions. *Biophysical Chemistry*, **105**:595–608.
- Lackner, P., Koppensteiner, W., Sippl, M., and Domingues, F. (2000). Prosup: a refined tool for protein structure alignment. *Protein Engineering*, **13**:745–752.
- Lemmen, C. and Lengauer, T. (2000). Computational methods for the structural alignment of molecules. *Journal of Computer-Aided Molecular Design*, **14**:215–232.
- Levitt, M. and Gerstein, M. (1998). A unified statistical framework for sequence comparison and structure comparison. *Proceedings of the National Academy of Sciences USA*, **95**:5913–5920.
- Lipman, D. and Pearson, W. (1985). Rapid and sensitive protein similarity searches. *Science*, **227**:1435–1441.
- Liu, J. and Lawrence, C. (1999). Bayesian inference on biopolymer models. *Bioinformatics*, **15**:38–52.
- Mizuguchi, K. and Go, N. (1995). Seeking significance in three-dimensional protein structure comparisons. *Current Opinion in Structural Biology*, **5**:377–382.
- Needleman, S. and Wunsch, C. (1970). A general method applicable to the search for similarities in the amino acid sequence of two proteins. *Journal of Molecular Biology*, **48**:443–453.
- Ortiz, A. R., Strauss, C. E. M., and Olmea, O. (2002). Mammoth (Matching molecular models obtained from theory): An automated method for model comparison. *Protein Science*, **11**:2606–2621.
- Papadopoulos, N. *et al.* (1994). Mutation of a mutL homolog in hereditary colon cancer. *Science*, **263**:1625–1629.

- Rao, S. and Rossmann, M. (1973). Comparison of super-secondary structures in proteins. *Journal of Molecular Biology*, 105:241–256.
- Rossmann, M. and Argos, P. (1975). A comparison of heme binding pocket in globins and cytochrome b5*. *Journal of Biological Chemistry*, 250:7523–7532.
- Rossmann, M. and Argos, P. (1976). Exploring structural homology in proteins. *Journal of Molecular Biology*, 105:75–96.
- Satow, Y., Cohen, G., Padlan, E., and Avies, D. (1986). Phosphocholine binding immunoglobulin fab mcpc603 : An x-ray diffraction study at 2.7 a. *Journal of Molecular Biology*, 190:593–604.
- Saul, L. K. and Jordan, M. I. (1995). Boltzmann chains and hidden Markov models. In Tesauro, G., Touretzky, D. S., and Leen, T., editors, *Advances in Neural Information Processing Systems (NIPS) 7*.
- Schmidler, S. C. (2003). Statistical shape analysis of protein structure families. Presentation at the *LASR workshop on Stochastic Geometry, Biological Structure and Images*, Leeds, UK.
- Schmidler, S. C. (2004). Bayesian shape matching and protein structure alignment. Presentation at the *6th World Congress of the Bernoulli Society and the 67th Annual Meeting of the Institute of Mathematical Statistics*.
- Schmidler, S. C. (2006). Fast Bayesian shape matching using geometric algorithms (with discussion). In Bernardo, J. M., Bayarri, S., Berger, J. O., Dawid, A. P., Heckerman, D., Smith, A. F. M., and West, M., editors, *Bayesian Statistics 8*, pages 471–490, Oxford. Oxford University Press.
- Schmidler, S. C. (2007). Bayesian flexible shape matching with applications to structural bioinformatics. (submitted to *Journal of the American Statistical Association*).
- Schmidler, S. C., Lucas, J., and Oas, T. G. (2007). Statistical estimation in statistical mechanical models: Helix-coil theory and peptide helicity prediction. *Journal of Computational Biology*, 14(10):1287–1310.
- Shindyalov, I. and Bourne, P. (1998). Protein structure alignment by incremental combinatorial extension (ce) of the optimal path. *Protein Engineering*, 11:739–747.
- Small, C. G. (1996). *The Statistical Theory of Shape*. Springer.
- Taylor, W. (2002). Protein structure comparison using bipartite graph matching and its application to protein structure classification. *Molecular and Cellular Proteomics*, 1:334–339.
- Taylor, W. R. and Orengo, C. A. (1989). Protein structure alignment. *Journal of Molecular Biology*, 208:1–22.
- Thompson, J. D., Koehl, P., Ripp, R., and Poch, O. (2005). BALiBASE 3.0: Latest developments of the multiple sequence alignment benchmark. *Proteins*, 61:127–136.

- Thompson, J. D., Plewniak, F., and Poch, O. (1999). BALiBASE: A benchmark alignment database for the evaluation of multiple alignment programs. *Bioinformatics*, **15**:87–88.
- Wallace, A., Laskowski, R., and Thornton, J. (1996). Tess: A geometric hashing algorithm for deriving 3d coordinate templates for searching structural databases: Applications to enzyme active sites. *Protein Science*, **6**:2308–2323.
- Wang, R. and Schmidler, S. C. (2008). Bayesian multiple protein structure alignment and analysis of protein families. (in preparation).
- Webb, B., Liu, J., and Lawrence, C. (2002). Balsa: Bayesian algorithm for local sequence alignment. *Nucleic Acids Research*, **30**:1268–1277.
- Wood, T. C. and Pearson, W. R. (1999). Evolution of protein sequences and structures. *Journal of Molecular Biology*, **291**:977–995.
- Wu, T. D., Schmidler, S. C., Hastie, T., and Brutlag, D. L. (1998). Regression analysis of multiple protein structures. *Journal of Computational Biology*, **5**(3):597–607.
- Zhu, J., Liu, J., and Lawrence, C. (1998). Bayesian adaptive sequence alignment algorithms. *Bioinformatics*, **14**:25–39.
- Zu-Kang, F. and Sippl, M. (1996). Optimum superimposition of protein structures: Ambiguities and implications. *Folding and design*, **1**:123–132.


| $X - Y$ | n | m | CE | | | $\lambda = 7.6$ | | | | | $\lambda = 8.6$ | | | | | $\lambda = 9.6$ | | | | |
|-----------|-----|-----|-------|------|-----|-----------------|------|----------|------|-----|-----------------|------|----------|------|-----|-----------------|------|----------|------|-----|
| | | | $ M $ | RMSD | h | $ M $ | RMSD | N_{CE} | h | g | $ M $ | RMSD | N_{CE} | h | g | $ M $ | RMSD | N_{CE} | h | g |
| 1ABA-IDSB | 87 | 188 | 56 | 4.5 | 24 | 2.2 | 0 | 9.64 | 0.01 | 57 | 3.7 | 0 | 6.26 | 0.13 | 76 | 4.7 | 14 | 6.10 | 0.42 | |
| 1ABA-ITRS | 87 | 105 | 70 | 2.7 | 65 | 3.0 | 37 | 5.68 | 0.14 | 72 | 3.4 | 38 | 5.68 | 0.22 | 75 | 3.6 | 38 | 5.70 | 0.24 | |
| 1ACX-ICOB | 108 | 151 | 92 | 4.0 | 66 | 2.1 | 49 | 6.89 | 0.06 | 86 | 3.8 | 57 | 6.60 | 0.15 | 93 | 4.1 | 57 | 6.38 | 0.21 | |
| 1ACX-IRBE | 108 | 104 | 56 | 7.3 | 25 | 2.5 | 6 | 9.02 | 0.01 | 31 | 2.8 | 0 | 7.89 | 0.02 | 50 | 4.2 | 15 | 7.45 | 0.03 | |
| 1MJC-5TSS | 69 | 194 | 61 | 2.7 | 52 | 2.3 | 25 | 7.89 | 0.03 | 60 | 3.0 | 29 | 7.24 | 0.36 | 66 | 3.9 | 15 | 7.36 | 0.44 | |
| 1PGB-5TSS | 56 | 194 | 48 | 2.9 | 39 | 2.3 | 19 | 6.52 | 0.56 | 55 | 3.3 | 40 | 6.60 | 0.87 | 55 | 3.1 | 34 | 6.65 | 0.94 | |
| 1PLC-1ACX | 102 | 108 | 80 | 3.3 | 71 | 3.4 | 23 | 5.92 | 0.10 | 84 | 4.0 | 23 | 5.80 | 0.20 | 89 | 4.6 | 23 | 6.30 | 0.22 | |
| 1PTS-1MUP | 119 | 157 | 80 | 4.1 | 76 | 3.0 | 0 | 6.80 | 0.06 | 83 | 3.1 | 0 | 6.60 | 0.09 | 88 | 3.5 | 0 | 6.77 | 0.12 | |
| 1TNF-1BMV | 152 | 185 | 115 | 4.1 | 70 | 2.7 | 3 | 7.86 | 0.02 | 109 | 4.2 | 40 | 7.10 | 0.08 | 113 | 4.3 | 35 | 6.94 | 0.11 | |
| 1UBQ-IFRD | 76 | 98 | 64 | 4.4 | 62 | 3.0 | 28 | 5.41 | 0.15 | 62 | 2.9 | 23 | 5.10 | 0.25 | 65 | 3.1 | 32 | 5.36 | 0.29 | |
| 1UBQ-4FXC | 76 | 98 | 64 | 4.0 | 46 | 2.3 | 22 | 5.46 | 0.13 | 61 | 2.9 | 34 | 5.20 | 0.24 | 66 | 3.4 | 42 | 5.43 | 0.29 | |
| 2GB1-1UBQ | 56 | 76 | 48 | 3.1 | 44 | 2.1 | 0 | 5.38 | 0.18 | 51 | 3.4 | 6 | 5.27 | 0.46 | 51 | 3.3 | 6 | 5.66 | 0.59 | |
| 2GB1-4FXC | 56 | 98 | 48 | 3.6 | 35 | 3.5 | 0 | 9.06 | 0.06 | 53 | 3.9 | 7 | 7.56 | 0.42 | 55 | 4.1 | 0 | 7.55 | 0.62 | |
| 2RSL-3CHY | 119 | 128 | 80 | 4.1 | 43 | 2.6 | 22 | 7.99 | 0.02 | 76 | 3.8 | 31 | 6.76 | 0.08 | 81 | 4.0 | 33 | 6.67 | 0.11 | |
| 2TMV-256B | 154 | 106 | 84 | 3.5 | 65 | 2.3 | 0 | 7.39 | 0.06 | 79 | 2.9 | 0 | 7.13 | 0.10 | 89 | 4.0 | 69 | 6.87 | 0.19 | |
| 3CHY-IRCF | 128 | 169 | 116 | 3.9 | 80 | 3.0 | 38 | 6.89 | 0.06 | 122 | 4.5 | 87 | 5.99 | 0.54 | 126 | 4.7 | 70 | 6.05 | 0.76 | |

Table 2: Bayesian structural alignments of the 16 pairs of proteins from Ortiz et al. (2002), compared against the popular CE algorithm (Shindyalov and Bourne, 1998). $|M|$ is the length of the alignment, N_{CE} denotes the number of matches in common with CE, and RMSD is expressed in Å.



(a)

(b)



(c)

Figure 1: Bayesian structural alignment of 5MBN and 2HBG. (a) Marginal alignment probability matrix for all pairs of residues, showing uncertainty associated with the alignment. (b) Plot of all sampled alignments (c) Comparison of the MAP alignment (red) with the CE alignment (blue); common regions are shown in purple.

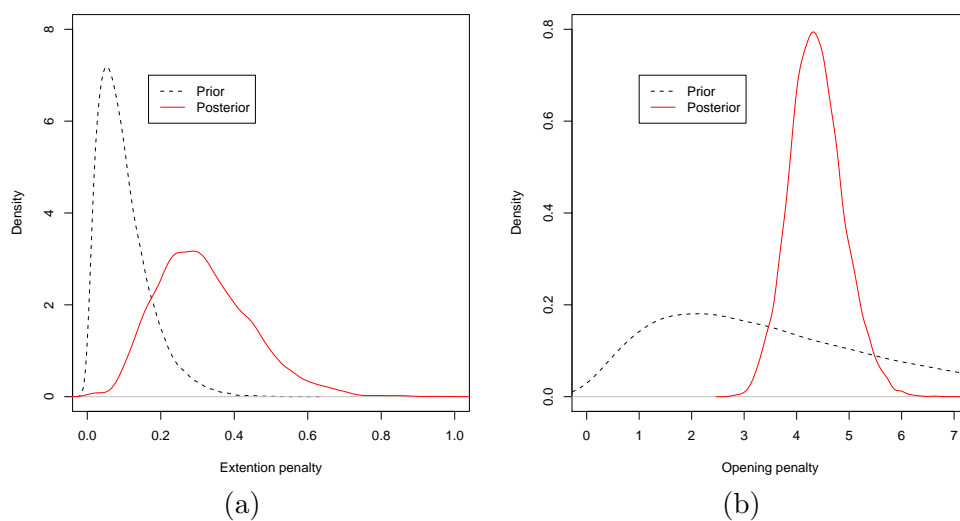


Figure 2: Prior and posterior distributions for gap penalty parameters obtained for the Bayesian alignment of globins 5MBN and 2HBG. The Bayesian approach allows the algorithm to adaptively determine the appropriate gap parameters rather than treating them as fixed.




Figure 3: Marginal alignment matrix for the Bayesian structural alignment of 1OUN:A and 1CEW:I. The posterior uncertainty in the alignment can be seen at the N-terminus, where two possible alignments of the α -helix at positions 10-20 have comparable posterior probabilities.

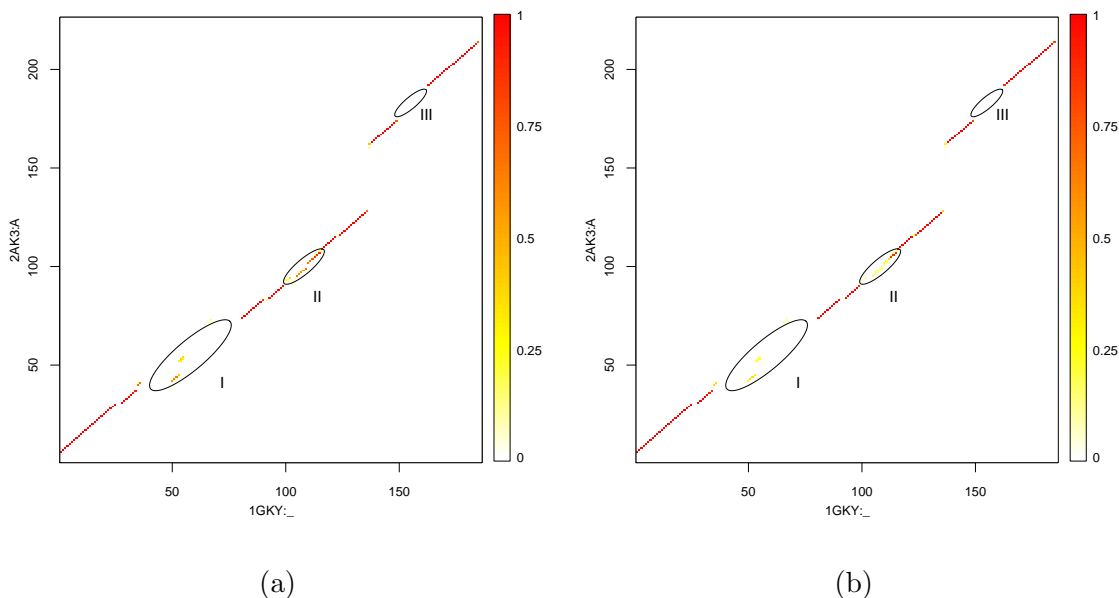


Figure 4: Marginal probabilities over aligned pairs for 1GKY and 2AK3:A. (a) Shows alignments based only on structure, while (b) presents alignments that also incorporate sequence information. Note that although there is some structural similarity in regions I and II, sequence similarity in these areas is low.

| | | | | | | | | |
|--------|----|---|----|---|----|----|---|----|
| 2AK3:A | R | T | L | P | Q | A | E | A |
| 1GKY | G | V | K | S | V | K | A | I |
| | -3 | 0 | -3 | 1 | -2 | -1 | 0 | -1 |

Table 3: Sequence alignment of corresponding to the best structural alignment between region II of 2AK3:A and 1GKY, with residues 93-100 of 2AK3:A matched with residues 103-111 of 1GKY. Numbers correspond to the PAM 250 (log-odds) scores for each matched residue pair, and clearly show that despite the shape similarity, there is little evidence of common ancestry in this region of the protein.

| Group | Type | Amino acids |
|-------|------------------------|------------------------|
| 1 | Non-polar, hydrophobic | A, V, L, I, P, M, F, W |
| 2 | Polar, hydrophilic | G, S, T, C, N, Q, Y |
| 3 | Acidic | D, E |
| 4 | Basic | K, R, H |

Table 4: Simple amino acid classification based on chemical properties




Figure 5: Posterior probabilities of PAM distances based on sequence information alone (a) and based on the Bayesian sequence-structure alignment

| | | | | | | | | | | |
|-----|--------|---|---|---|---|---|---|---|---|---|
| (a) | 1MOL:A | E | W | E | I | I | D | I | G | P |
| | 1CEW:I | G | A | P | V | P | V | D | E | N |
| (b) | 1MOL:A | G | E | W | E | I | I | D | I | G |
| | 1CEW:I | G | A | P | V | P | V | D | E | N |

Table 5: Sequence alignment of the first strand of 1MOL:A and 1CEW:I induced by the two alternative models. (a) Mode using structural information only and (b) Mode under the Bayesian simultaneous sequence-structure alignment. Colors refer to the classification in Table 4; note the improvement in matching of chemical classes.




Figure 6: Marginal probabilities over aligned pairs for 1MOL:A and 1CEW:I. (a) Shows alignments based only on structure, while (b) presents alignments that also incorporate sequence information. Circles show the strand region discussed in Table 5.

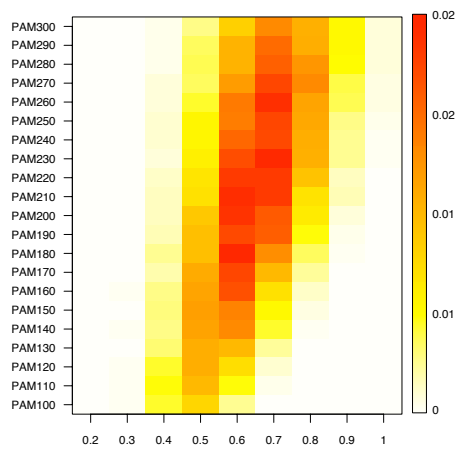


Figure 7: Heat map representation of the joint posterior distribution over discount factors and PAM matrices for 1:MOL:A and 1CEW:I.

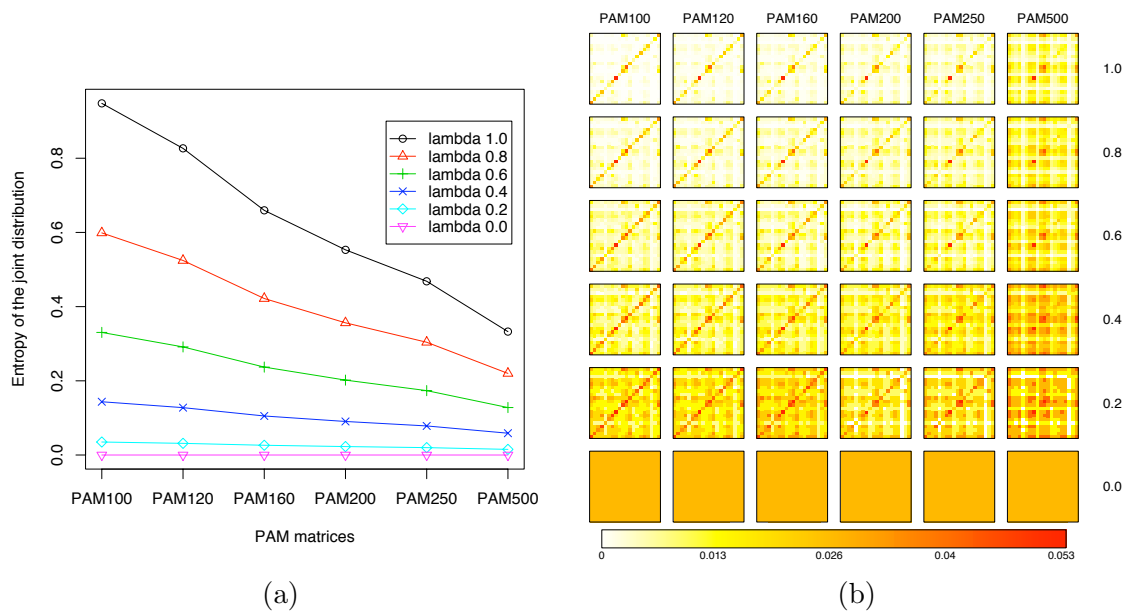


Figure 8: (a) Entropies of the joint from induced by different evolutionary distances and tempering parameters. (b) Heat map plots of the joint distributions. Amino acids are ordered alphabetically order starting Alanine in the lower-left.