Support Vector Training of Protein Alignment Models

CHUN-NAM JOHN YU,¹ THORSTEN JOACHIMS,¹ RON ELBER,¹
and JAROSLAW PILLARDY²

ABSTRACT

Sequence to structure alignment is an important step in homology modeling of protein structures. Incorporation of features such as secondary structure, solvent accessibility, or evolutionary information improve sequence to structure alignment accuracy, but conventional generative estimation techniques for alignment models impose independence assumptions that make these features difficult to include in a principled way. In this paper, we overcome this problem using a Support Vector Machine (SVM) method that provides a well-founded way of estimating complex alignment models with hundred of thousands of parameters. Furthermore, we show that the method can be trained using a variety of loss functions. In a rigorous empirical evaluation, the SVM algorithm outperforms the generative alignment method SSALN, a highly accurate generative alignment model that incorporates structural information. The alignment model learned by the SVM aligns 50% of the residues correctly and aligns over 70% of the residues within a shift of four positions.

Key words: machine learning, pairwise sequence alignment, protein structure prediction.

1. INTRODUCTION

SEQUENCE TO STRUCTURE ALIGNMENT is a crucial step in building accurate three-dimensional protein models in homology modeling. Most alignment methods are based on dynamic programming on a linear cost model. In the simplest cases, such as alignment with the BLOSUM matrices, the linear model specifies costs for substituting one amino acid with another and for inserting a gap. The choice of these costs greatly determines the quality of alignments. While it is well understood how to estimate the substitution costs for such simple models, sequence to structure alignment requires more complex cost models to take advantage of the structural information available. For these complex cost models, conventional generative estimation techniques such as hidden Markov models (HMM) are difficult to use due to the independence assumptions they make, for example, in assuming that the amino acid sequence and the secondary structure labels of a protein are generated by independent processes.

¹Department of Computer Science and ²Cornell Theory Center, Cornell University, Ithaca, New York.

This paper explores a Support Vector Machine (SVM) method for learning an application-specific cost model from training data for sequence to structure alignment. The advantages of this SVM method over conventional generative techniques are threefold. First, unlike conventional generative estimation techniques, the SVM method does not require independence assumptions among features and therefore provides a well-founded way to learn cost models where each aligned position is described not only by its amino acid identity, but by a potentially high-dimensional feature vector. This feature vector may describe additional properties of the aligned position (e.g., predicted secondary structure) as well as properties of surrounding aligned positions (e.g., the previous amino acid is hydrophobic). This provides great flexibility in building expressive models. Second, the SVM method inherits the benefits of conventional classification SVMs, in particular its robustness to overfitting for high-dimensional and sparse data. Third, the SVM method allows optimizing for different loss functions. This allows accounting for uncertainty in the training data, as well as specifying which types of alignment errors are more costly than others.

The work reported in the following shows that the SVM algorithm can be used to learn highly accurate alignment models for sequence to structure alignments. It also provides the first large-scale implementation and empirical validation of this SVM alignment algorithm, extending the basic algorithm first proposed in Joachims (2003) and Joachims et al. (2005) to include loss functions. We show that this SVM algorithm can effectively learn alignment models with hundreds of thousands of features, outperforming the accuracy of state-of-the-art generative estimation techniques (Qiu and Elber, 2006). Finally, we show that loss functions can be incorporated into the SVM training problem while maintaining polynomial runtime guarantees. We find that the use of application-dependent loss functions during training, in particular by only counting alignment errors if the shift from the correct residue is more than 4, is effective in modeling the uncertainty in the training data. This is an extension of the RECOMB version of the paper by including new results on the use of profile information in the alignment models and also a more accurate estimation of performance with a fivefold cross validation methodology. The alignment program and data of this work is available for download at www.cs.cornell.edu/~cnyu/svm_alignment/.

2. RELATED WORK

Conventional estimation techniques for alignment models (Bucher and Hofmann, 1996; Durbin et al., 1998; Henikoff and Henikoff, 1992; Dayhoff et al., 1978; Ristad and Yianilos, 1998) take the view of a generative probabilistic model. A generative alignment model (e.g., HMM) aims to model the process that generates the data as the joint probability distribution $P(S, T, Y)$, where $\mathbf{s} = (s^1, \dots, s^{|\mathbf{s}|})$ and $\mathbf{t} = (t^1, \dots, t^{|\mathbf{t}|})$ are two sequences and $\mathbf{y}$ is the alignment. We denote the length of a sequence with $|\cdot|$. If $P(S, T, Y)$ (or a good estimate thereof) is known,

$$\operatorname{argmax}_{\mathbf{y}} P(S = \mathbf{s}, T = \mathbf{t}, Y = \mathbf{y}) \quad (1)$$

predicts an alignment $\mathbf{y}$ from two sequences $\mathbf{s}$ and $\mathbf{t}$. To make estimation of $P(S, T, Y)$ tractable, it is decomposed by making independence assumptions on the process that generates $\mathbf{s}$ and $\mathbf{t}$. While this leads to efficient and simple estimation problems, the independence assumptions restrict the interactions within the sequences $\mathbf{s}$ and $\mathbf{t}$ that we could model.

Machine learning research over the last decade has provided substantial evidence that discriminative learning (e.g., SVMs, MaxEnt classifiers) typically produces more accurate rules than generative learning (e.g., naive Bayes classifiers, HMMs) (Joachims, 1998; Tsochantaridis et al., 2005; Lafferty et al., 2001). This can be explained as follows. Since $P(Y|S, T)$ is already sufficient for making an optimal prediction

$$\operatorname{argmax}_{\mathbf{y}} P(Y = \mathbf{y} | S = \mathbf{s}, T = \mathbf{t}), \quad (2)$$

modeling the joint distribution of the input sequences S and T is not necessary, and generative methods might be wasting effort in trying to do so. Discriminative learning applied to the alignment problem would directly estimate $P(Y|S, T)$ or a related discriminant function, thus focusing on the relevant part of the estimation problem.

Only few approaches to discriminative training of alignment models exist to date. While not motivated from this learning theoretical perspective, work on *inverse alignment* is closely related, since our SVM

method can be viewed as solving an inverse alignment problem. Inverse alignment is the task of finding a cost model under which a given alignment algorithm outputs a desired alignment $\mathbf{y}$ for sequences $\mathbf{s}$ and $\mathbf{t}$. This problem was first formulated in Gusfield and Stelling (1996). They discuss inverse alignment in the context of parametric sequence alignment and identify geometric properties of the space of cost model. In more detail, the work in Pachter and Sturmfields (2004) analyzes the space of models and shows that some aspects of its complexity grow only polynomially.

The first concrete algorithm for inverse sequence alignment was proposed in Sun et al. (2002). While they prove that their algorithm finds a consistent cost model in polynomial time, their algorithm is limited to particular cost models with at most three parameters.

The work presented in this paper follows the Structural SVM algorithm first proposed in Joachims (2003) and Joachims et al. (2005) for sequence alignment and later generalized to a wide class of multivariate prediction problems (Tsochantaridis et al., 2004, 2005). In this paper, we present the first large-scale empirical evaluation of this type of algorithm for alignment. We also extend the algorithm to optimize particular loss functions, and show how the resulting optimization problems can be solved in polynomial time.

Related to our SVM approach are Conditional Random Fields (CRFs), which have recently been proposed for sequence alignment as well (Do et al., 2006; McCallum et al., 2005). While CRFs share with SVMs the benefits of discriminative training, they do not allow the use of application-dependent loss functions.

Independent of the work on Structural SVMs in the machine learning community, recently an algorithm for inverse alignment was proposed in Kececiloglu and Kim (2006). Their formulation of the problem is similar to a Structural SVM and the algorithm resembles the cutting-plane method used for training Structural SVMs. However, their approach is based on a linear programming formulation instead of the quadratic programming formulation used in SVMs. Furthermore, their approach to handling infeasibilities in the resulting optimization problem is different and it is unclear how it relates to an intuitively meaningful loss function. While they give empirical results, they are on a scale of 10 training examples and 212 features. In the following, we will explore models trained over thousands of examples and hundred of thousands of features.

3. SEQUENCE ALIGNMENT

We begin by introducing the class of alignment models considered in this paper. Since we focus on the problem of sequence to structure alignment, we will describe the model in these terms. However, the methods can obviously be used for other applications as well. Let $(\mathbf{s}, \mathbf{t})$ be a pair of target and template sequence that we wish to align. For an alignment $\mathbf{y}$ of $(\mathbf{s}, \mathbf{t})$, we write $\mathbf{y}$ as a sequence of alignment operations $(y^1, y^2, \dots, y^{|\mathbf{y}|})$. Each y^k is an alignment operation of the form (i, j) , where i, j are positions of characters in $\mathbf{s}$ and $\mathbf{t}$, respectively, or the special gap symbol “—.”

We consider alignment algorithms (Smith and Waterman, 1981) that optimize a linear scoring function $D_{\mathbf{w}}(\mathbf{y}, \mathbf{s}, \mathbf{t}) = \mathbf{w} \cdot \Psi(\mathbf{y}, \mathbf{s}, \mathbf{t})$, where Ψ is a function that maps an alignment $\mathbf{y}$ of $\mathbf{s}$ and $\mathbf{t}$ to a feature vector, and $\mathbf{w}$ is a given cost vector. Note that $\mathbf{w}$ contains the parameters of the alignment model that we will learn. We require that $\Psi(\mathbf{y}, \mathbf{s}, \mathbf{t})$ be linear in the individual alignment operations y^k within $\mathbf{y}$, written in terms of equations,

$$\Psi(\mathbf{y}, \mathbf{s}, \mathbf{t}) = \sum_{k=1}^{|\mathbf{y}|} \phi(y^k, \mathbf{s}, \mathbf{t}), \quad (3)$$

where ϕ is a function that maps each individual alignment operation onto a feature vector. Note that this feature vector can be any function that depends on the operation y^k and the *full* target and template sequences, not just the current positions that are aligned by y^k . In general, alignment algorithms compute

$$\operatorname{argmax}_{\mathbf{y} \in \mathcal{Y}} [\mathbf{w} \cdot \Psi(\mathbf{y}, \mathbf{s}, \mathbf{t})] = \operatorname{argmax}_{\mathbf{y} \in \mathcal{Y}} \left[\sum_{k=1}^{|\mathbf{y}|} \mathbf{w} \cdot \phi(y^k, \mathbf{s}, \mathbf{t}) \right] \quad (4)$$

to determine the alignment, where $\mathcal{Y}$ is the set of all possible (local or global, as desired) alignments between $\mathbf{s}$ and $\mathbf{t}$. This is typically computed using dynamic programming (Smith and Waterman, 1981). Note that our setting includes the common scenarios of alignment with substitution matrices such as BLOSUM as a special case, where the function $\phi(y^k, \mathbf{s}, \mathbf{t})$ return a sparse vector with exactly one “1” that corresponds to the particular substitution or gap score in $\mathbf{w}$. However, we will consider richer feature mappings ϕ that go beyond amino acid identity and that include structural information of the template sequence.

4. DISCRIMINATIVE TRAINING OF ALIGNMENT MODELS

In the above section, the vector $\mathbf{w}$ parameterizes the scoring function D and has crucial influence on the quality of alignments between $\mathbf{s}$ and $\mathbf{t}$. In the following, we aim to learn $\mathbf{w}$ from a set of training examples

$$\mathcal{S} = ((\mathbf{s}_1, \mathbf{t}_1, \mathbf{y}_1), (\mathbf{s}_2, \mathbf{t}_2, \mathbf{y}_2), \dots, (\mathbf{s}_n, \mathbf{t}_n, \mathbf{y}_n)) \quad (5)$$

of sequence pairs $(\mathbf{s}_i, \mathbf{t}_i)$ for which the (approximately) correct alignment $\mathbf{y}_i$ is known. This training set is assumed to be generated independently and identically distributed (i.i.d.) according to some unknown distribution $P(S, T, Y)$. Thinking of a sequence alignment algorithm as a function,

$$h_{\mathbf{w}}(\mathbf{s}, \mathbf{t}) = \operatorname{argmax}_{\mathbf{y} \in \mathcal{Y}} [\mathbf{w} \cdot \Psi(\mathbf{y}, \mathbf{s}, \mathbf{t})] \quad (6)$$

maps a given sequence pair $(\mathbf{s}, \mathbf{t})$ to an alignment $\mathbf{y}$. Our goal is to find a parameter vector $\mathbf{w}$ so that the predicted alignment $h_{\mathbf{w}}(\mathbf{s}, \mathbf{t})$ matches the correct alignment on new test data as well as possible. In particular, we want to find a $\mathbf{w}$ that minimizes the expected loss (i.e., risk)

$$R_P(h_{\mathbf{w}}) = \int \Delta(\mathbf{y}, h_{\mathbf{w}}(\mathbf{s}, \mathbf{t})) dP(S, T, Y), \quad (7)$$

where $\Delta(\mathbf{y}, \mathbf{y}')$ is a user defined (non-negative) loss function that quantifies how “bad” it is to predict $\mathbf{y}'$ when $\mathbf{y}$ is the correct alignment. For example, one may choose $\Delta(\mathbf{y}, \mathbf{y}')$ to be 1 minus the Q-score (Q-score is the fraction of match operations from $\mathbf{y}$ that are also contained in $\mathbf{y}'$).

Following the principle of (Structural) Empirical Risk Minimization (Vapnik, 1998), finding a $\mathbf{w}$ that predicts well on new data can be achieved by minimizing the empirical loss (i.e., the training error) $R_S(h_{\mathbf{w}}) = \sum_{i=1}^n \Delta(\mathbf{y}_i, h_{\mathbf{w}}(\mathbf{s}_i, \mathbf{t}_i))$ on the training set $\mathcal{S}$. This leads to the computational problem of finding the $\mathbf{w}$ which minimizes $R_S(h_{\mathbf{w}})$ as follows.

5. STRUCTURAL SVMs FOR SEQUENCE ALIGNMENT

In the framework of structural SVMs (Tsochantaridis et al., 2005), we formulate the problem of finding the parameters $\mathbf{w}$ that minimizes the empirical loss $R_S(h_{\mathbf{w}})$ of the sequence alignment algorithm as the following optimization problem:

$$\min_{\mathbf{w}, \xi} \frac{1}{2} \|\mathbf{w}\|^2 + \frac{C}{n} \sum_{i=1}^n \xi_i \quad (8)$$

$$s.t. \quad \forall i \in \{1, \dots, n\} \quad \forall \hat{\mathbf{y}} \in \mathcal{Y}_i \setminus \{\mathbf{y}_i\} : \mathbf{w} \cdot (\Psi(\mathbf{y}_i, \mathbf{s}_i, \mathbf{t}_i) - \Psi(\hat{\mathbf{y}}, \mathbf{s}_i, \mathbf{t}_i)) \geq \Delta(\mathbf{y}_i, \hat{\mathbf{y}}) - \xi_i$$

The objective is the conventional regularized risk used in SVMs. The constraints state that the score $\mathbf{w} \cdot \Psi(\mathbf{y}_i, \mathbf{s}_i, \mathbf{t}_i)$ of the correct alignment $\mathbf{y}_i$ must be greater than the score $\mathbf{w} \cdot \Psi(\hat{\mathbf{y}}, \mathbf{s}_i, \mathbf{t}_i)$ of all alternative alignments $\hat{\mathbf{y}}$. Note that $\mathcal{Y}_i$ (i.e., the set of all possible alignments, for example, i) depends on whether the alignment is local or global. However, the optimization problem is well-formed in either case.

Unlike the formulation in Joachims et al. (2005), our new formulation includes a loss function $\Delta(\mathbf{y}_i, \hat{\mathbf{y}})$ that scales the desired difference in score. Intuitively, the larger the loss of an incorrect alignment $\hat{\mathbf{y}}$,

the further should the score be away from that of the correct alignment $\mathbf{y}_i$. This method for including a loss function is analogous to proposals for other structured prediction problems (Taskar et al., 2003; Tsochantaridis et al., 2005). ξ_i is a slack variable shared among constraints from the same example, since in general the constraint system is not feasible. Following the proof in Tsochantaridis et al. (2005), it is easy to see that the following result holds.

Theorem 1. *If $(\mathbf{w}^*, \xi^*)$ is the solution of the optimization problem in (8), the sum of the slack variables ξ_i^* is an upper bound on the training loss, $R_S(h_{\mathbf{w}}) \leq \sum_{i=1}^n \xi_i^*$.*

This shows that our formulation minimizes training loss, while the SVM-style regularization with the norm of $\mathbf{w}$ in the objective provides protection against overfitting for high-dimensional $\mathbf{w}$. The parameter C allows the user to control the trade-off between training error and regularization.

5.1. Efficient training algorithm

While it is easy to see that the optimization problem in (8) is convex, it unfortunately has an exponential number of constraints. This results from the fact that there are exponentially many “wrong” alignments $\mathcal{Y}_i \setminus \{\mathbf{y}_i\}$ for each given pair of sequences $(\mathbf{s}_i, \mathbf{t}_i)$. Any attempt to solve this type of optimization problem using standard methods that require enumerating all constraints is obviously not tractable for sequences and training sets of interesting size.

Despite the exponential size, however, it has been shown that cutting-plane algorithms can be used to efficiently approximate the optimal solution of this type of optimization problem (Tsochantaridis et al., 2005). An adaptation of this algorithm to the problem of sequence alignment is given in Figure 1. The algorithm iteratively constructs a subset of all constraints from (8) until this subset constrains the feasible region enough to ensure an ϵ -accurate solution. The desired precision ϵ is provided by the user. In particular, the algorithm starts with an empty set K of constraints. It then adds the most violated constraint among the exponentially many for each example. If no constraint exists that is violated by more than ϵ , the algorithm terminates. Otherwise, it solves the optimization problem over the current set K and repeats. Adapting the result from Tsochantaridis et al. (2005), it can be proved that only a polynomial number of constraints will be added before the algorithm converges.

Theorem 2. *For any $\epsilon > 0$, $C > 0$, and any training sample $S = ((\mathbf{s}_1, \mathbf{t}_1, \mathbf{y}_1), \dots, (\mathbf{s}_n, \mathbf{t}_n, \mathbf{y}_n))$, the algorithm in Figure 1 converges after adding at most $\max\{2n\bar{\Delta}/\epsilon, 8C\bar{\Delta}R^2/\epsilon^2\}$ constraints to K , where $R = \max_{i,y} \|\Psi(\mathbf{y}, \mathbf{s}_i, \mathbf{t}_i)\|$ and $\bar{\Delta}$ is an upper bound on the loss function $\Delta(\mathbf{y}_i, \hat{\mathbf{y}})$.*

Input: sequence pairs $(\mathbf{s}_1, \mathbf{t}_1), \dots, (\mathbf{s}_n, \mathbf{t}_n)$, correct alignments $\mathbf{y}_1, \dots, \mathbf{y}_n$, tolerated error $\epsilon \geq 0$.

$K = \emptyset$, $\mathbf{w} = 0$, $\xi = 0$

repeat

- $K_{\text{org}} = K$
- for i from 1 to n
 - $\hat{\mathbf{y}} = \operatorname{argmax}_{\mathbf{y} \in \mathcal{Y}_i \setminus \{\mathbf{y}_i\}} [\Delta(\mathbf{y}_i, \hat{\mathbf{y}}) + \mathbf{w} \cdot (\Psi(\hat{\mathbf{y}}, \mathbf{s}, \mathbf{t}) - \Psi(\mathbf{y}_i, \mathbf{s}_i, \mathbf{t}_i))]$
 - if $\mathbf{w} \cdot (\Psi(\mathbf{y}_i, \mathbf{s}_i, \mathbf{t}_i) - \Psi(\hat{\mathbf{y}}, \mathbf{s}, \mathbf{t})) < \Delta(\mathbf{y}_i, \hat{\mathbf{y}}) - \xi_i - \epsilon$
 - * $K = K \cup \{\mathbf{w} \cdot (\Psi(\mathbf{y}_i, \mathbf{s}_i, \mathbf{t}_i) - \Psi(\hat{\mathbf{y}}, \mathbf{s}, \mathbf{t})) \geq \Delta(\mathbf{y}_i, \hat{\mathbf{y}}) - \xi_i - \epsilon\}$
 - * $(\mathbf{w}, \xi) = \operatorname{argmin}_{\mathbf{w}, \xi} \frac{1}{2} \|\mathbf{w}\|^2 + \frac{C}{n} \sum_{i=1}^n \xi_i$ subject to K .

until $(K = K_{\text{org}})$

Output: $\mathbf{w}$

FIG. 1. Cutting-plane algorithm for solving the SVM optimization problem.

One crucial aspect of the algorithm, however, is the use of an oracle (often called a separation oracle in optimization theory) that can find the most violated constraint among the exponentially many in polynomial time. This is equivalent to the argmax problem in the algorithm. The following section shows that this argmax can be computed efficiently for a large class of loss functions.

5.2. Loss functions

We first introduce the Q-loss function, which we denote as $\Delta_Q(\mathbf{y}_i, \hat{\mathbf{y}})$ for the loss between a correct alignment $\mathbf{y}_i$ and a predicted alignment $\hat{\mathbf{y}}$. The Q-loss measures the proportion of incorrect matches in a predicted alignment, which we want to minimize. Writing $\hat{\mathbf{y}}$ as a sequence of alignment operations $\hat{\mathbf{y}} = (\hat{y}^1, \hat{y}^2, \dots, \hat{y}^{|\hat{\mathbf{y}}|})$, we can decompose the Q-loss $\Delta_Q(\mathbf{y}_i, \hat{\mathbf{y}})$ into a sum of losses on individual alignment operations $\Delta_Q(\mathbf{y}_i, \hat{\mathbf{y}}) = 1 - \sum_{k=1}^{|\hat{\mathbf{y}}|} \delta_Q(\mathbf{y}_i, \hat{y}^k)$. The function $\delta_Q(\mathbf{y}_i, \hat{y}^k)$ returns $1/M$ when $\hat{y}^k$ is a match contained in the correct alignment $\mathbf{y}_i$, and 0 otherwise. M is the number of matches in $\mathbf{y}_i$. With this decomposition, we can rewrite the computation of the most violated constraint in the cutting plane algorithm as:

$$\operatorname{argmax}_{\hat{\mathbf{y}} \in \mathcal{Y}_i \setminus \{\mathbf{y}_i\}} \left[\sum_{k=1}^{|\hat{\mathbf{y}}|} (\mathbf{w} \cdot \phi(\hat{y}^k, \mathbf{s}_i, \mathbf{t}_i)) - \delta_Q(\mathbf{y}_i, \hat{\mathbf{y}}) \right] + 1 - \mathbf{w} \cdot \Psi(\mathbf{y}_i, \mathbf{s}_i, \mathbf{t}_i). \quad (9)$$

Since the non-constant term in the argmax decomposes into a sum over individual alignment operations $\hat{y}^k$, we can apply dynamic programming similar to (4) with operation costs modified to $\mathbf{w} \cdot \phi(\hat{y}^k, \mathbf{s}_i, \mathbf{t}_i) - \delta_Q(\mathbf{y}_i, \hat{\mathbf{y}})$ for each $\hat{y}^k$. Note that δ_Q can be computed efficiently using table lookup.

Another interesting loss function that we would like to consider is the Q_4 -loss function, where we count a match as correct even if it is slightly shifted, in particular, shifted by not more than four positions. Suppose we have an alignment operation $\hat{y}^k = (i, j)$ in $\hat{\mathbf{y}}$. The shift of $\hat{y}^k$ is less than 4 if there is some match operation $y^l = (u, j)$ in $\mathbf{y}$ with $|i - u| \leq 4$. Similar to the Q-loss, we can decompose Q_4 linearly in its alignment operations as $\Delta_{Q_4}(\mathbf{y}_i, \hat{\mathbf{y}}) = 1 - \sum_{k=1}^{|\hat{\mathbf{y}}|} \delta_{Q_4}(\mathbf{y}_i, \hat{y}^k)$, where $\delta_{Q_4}(\mathbf{y}_i, \hat{y}^k)$ is $1/M$ if $\hat{y}^k$ is a match operation contained in $\mathbf{y}$ with a shift of 4 or less, and 0 otherwise. The same dynamic programming algorithm applies to the computation of most violated constraint with Q_4 -loss, since Q_4 -loss decomposes in the same way as the Q-loss. Again, note that δ_{Q_4} can again be computed efficiently by table lookup of matched characters in the range $(i - 4, j), \dots, (i + 4, j)$ within the correct alignment.

6. EXPERIMENTS

The following experiments evaluate the SVM alignment algorithm on a sequence to structure alignment task. We evaluate whether the algorithm can effectively and efficiently learn complex alignment models with hundred of thousands of features, how optimizing to different loss functions might help, and how the algorithm compares to conventional methods.

In all our experiments, we train the algorithm on a training set, select any parameters and models based on a validation set, and then report performance on an independent test set. Beyond the choice of model $\phi(y^k, \mathbf{s}_i, \mathbf{t}_i)$ and loss function $\Delta(\mathbf{y}, \mathbf{y}')$, our method has only a single parameter to tune, namely the regularization parameter C . We train alignment models with C ranging from 1 to 2^{15} , in powers of 2, all to precision $\epsilon = 0.01$. We then pick the value of C that gives the best performance on the validation set, and re-train the models using this particular value of C with fivefold cross validation. The Q-scores and standard errors reported below are all estimated using this fivefold cross validation procedure.

The training and validation sets are the same as those used in Qiu and Elber (2006). The data set contains 1379 target sequences, and each target sequence $\mathbf{s}$ has one or more template structures $\mathbf{t}$ associated with it. Structural alignments between target and template are generated using the CE program (Shindyalov and Bourne, 1998), and one example $(\mathbf{s}, \mathbf{t}, \mathbf{y})$ is generated whenever the structural alignment $\mathbf{y}$ between the structure of $\mathbf{s}$ and $\mathbf{t}$ has a CE Z-score of at least 4.5. The data set is randomly split into two sets, namely a training set with examples from 690 targets and a validation set with examples from 689 targets. The

resulting training set contains 4542 examples (i.e., pairwise alignments), while the validation set contains 4587 examples.

The test set is based on a database of protein structures that is used by the modeling program LOOPP (<http://cbsuapps.tc.cornell.edu/loopp.aspx>). We select 4185 structures from the new PDB structures released between June 2005 and June 2006 via clustering. These structures serve as target sequences in our test set and none of them appear in the training or validation sets since they were developed earlier. Each of these 4185 structures is aligned against all other structures using the structural alignment program TM-align (Zhang and Skolnick, 2005). Pairs that score 0.5 or better are considered homologous and are added to the test set. The selected pairs are then aligned by the structural alignment program CE. Only alignments that have CE Z-score higher than 4.5 are included in the final test case, providing a total of 29,345 alignments to consider.

As described in the following, we use structural annotations as features in our alignment models. The structural annotations of all the target sequences, that is, the secondary structure and the relative exposed surface area, are predicted by the SABLE program (Adamczak et al., 2004). There are three types of secondary structure predicted and the relative exposed surface areas are binned into four types. The structural annotations used in the template structures are computed by the program DSSP (Kabsch and Sander, 1983). The secondary structures are binned into five types, while the exposed surface areas are binned into six types. In the last alignment model we also include evolutionary information in the form of PSI-BLAST profile to investigate its usefulness in alignment.

6.1. Can the SVM algorithm learn complex models effectively?

In our first set of experiments, we evaluate models $\phi(y^k, \mathbf{s}, \mathbf{t})$ of increasing complexity and number of features. Our focus is on exploring how far we can push the complexity of the model and still be able to train them efficiently and effectively.

We use R_s^i, S_s^i, A_s^i to denote the residue, predicted secondary structure, and predicted exposed surface area at the i th position of the target sequence, and R_t^j, S_t^j, A_t^j to denote the residue, actual secondary structure, and actual exposed surface area at the j th position of the template structure. We also use $\mathbf{R}, \mathbf{S}, \mathbf{A}$ to denote the set of possible values for residue, secondary structure, and exposed surface area. In the last alignment model considered, we include PSI-BLAST profile as an extra feature, and we denote them as P_s^i and P_t^j for the profiles at the i th and j th positions in the target and template sequences respectively.

6.1.1. Substitution cost models. For the substitution costs, we consider the following six models for $\phi(y^k, \mathbf{s}, \mathbf{t})$. Since the length of alignments of different examples varies greatly, we normalize each ϕ by dividing with $|\mathbf{s}| + |\mathbf{t}|$.

Simple: In this alignment model, we only consider the substitution cost of single features. Let $y^k = (i, j)$ be a match operation. We define $\phi(y^k, \mathbf{s}, \mathbf{t})$ to be

$$\begin{aligned} \phi_{Simple}(y^k, \mathbf{s}, \mathbf{t}) &= \sum_{r_1, r_2 \in \mathbf{R}} \mathbb{I}[R_s^i = r_1, R_t^j = r_2] + \sum_{r_1 \in \mathbf{R}, s_2 \in \mathbf{S}} \mathbb{I}[R_s^i = r_1, S_t^j = s_2] + \sum_{r_1 \in \mathbf{R}, a_2 \in \mathbf{A}} \mathbb{I}[R_s^i = r_1, A_t^j = a_2] \\ &+ \sum_{s_1 \in \mathbf{S}, r_2 \in \mathbf{R}} \mathbb{I}[S_s^i = s_1, R_t^j = r_2] + \sum_{s_1, s_2 \in \mathbf{S}} \mathbb{I}[S_s^i = s_1, S_t^j = s_2] + \sum_{s_1 \in \mathbf{S}, a_2 \in \mathbf{A}} \mathbb{I}[S_s^i = s_1, A_t^j = a_2] \\ &+ \sum_{a_1 \in \mathbf{A}, r_2 \in \mathbf{R}} \mathbb{I}[A_s^i = a_1, R_t^j = r_2] + \sum_{a_1 \in \mathbf{A}, s_2 \in \mathbf{S}} \mathbb{I}[A_s^i = a_1, S_t^j = s_2] + \sum_{a_1, a_2 \in \mathbf{A}} \mathbb{I}[A_s^i = a_1, A_t^j = a_2], \end{aligned}$$

where $\mathbb{I}[\rho]$ is a function that returns a vector with “1” in the position designated to ρ if the boolean expression ρ is true, and returns “0” otherwise and in all other positions. For example, $\mathbb{I}[R_s^3 = \text{“A,”} S_t^7 = \text{“}\alpha\text{”}]$

returns “1” in the particular dimension corresponding to $\mathbb{I}[R_s^i = \text{“A”}, S_t^j = \text{“}\alpha\text{”}]$, if $y^k = (3, 7)$ aligns the residue alanine “A” in $\mathbf{s}$ with an alpha helix “ α ” in $\mathbf{t}$. Otherwise, it returns “0” in this dimension. For all other dimensions, it always returns “0.” Note that each such dimension corresponds to a particular position in cost vector $\mathbf{w}$. Note also that each feature vector $\phi_{Simple}(y^k, \mathbf{s}, \mathbf{t})$ has exactly nine “1”s corresponding to the nine terms in the sum, and is zero elsewhere.

Anova2: In this more complex feature vector, we take the interactions between pairs of structural annotations at the same position in the sequence into account. We define $\phi(y^k, \mathbf{s}, \mathbf{t})$ to be

$$\begin{aligned}
\phi_{Anova2}(y^k, \mathbf{s}, \mathbf{t}) = & \sum_{r_1, r_2 \in \mathbf{R}, s_1, s_2 \in \mathbf{S}} \mathbb{I}[R_s^i = r_1, S_s^i = s_1, R_t^j = r_2, S_t^j = s_2] \\
& + \sum_{\substack{r_1 \in \mathbf{R}, a_2 \in \mathbf{A}, \\ s_1, s_2 \in \mathbf{S}}} \mathbb{I}[R_s^i = r_1, S_s^i = s_1, S_t^j = s_2, A_t^j = a_2] \\
& + \sum_{\substack{r_1, r_2 \in \mathbf{R}, s_1 \in \mathbf{S}, \\ a_2 \in \mathbf{A}}} \mathbb{I}[R_s^i = r_1, S_s^i = s_1, A_t^j = a_2, R_t^j = r_2] \\
& + \sum_{\substack{r_2 \in \mathbf{R}, a_1 \in \mathbf{A}, \\ s_1, s_2 \in \mathbf{S}}} \mathbb{I}[S_s^i = s_1, A_s^i = a_1, R_t^j = r_2, S_t^j = s_2] \\
& + \sum_{\substack{s_1, s_2 \in \mathbf{S}, \\ a_1, a_2 \in \mathbf{A}}} \mathbb{I}[S_s^i = s_1, A_s^i = a_1, S_t^j = s_2, A_t^j = a_2] \\
& + \sum_{\substack{r_2 \in \mathbf{R}, s_1 \in \mathbf{S}, \\ a_1, a_2 \in \mathbf{A}}} \mathbb{I}[S_s^i = s_1, A_s^i = a_1, A_t^j = a_2, R_t^j = r_2] \\
& + \sum_{\substack{r_1, r_2 \in \mathbf{R}, s_2 \in \mathbf{S}, \\ a_2 \in \mathbf{A}}} \mathbb{I}[A_s^i = a_1, R_s^i = r_1, R_t^j = r_2, S_t^j = s_2] \\
& + \sum_{\substack{r_1 \in \mathbf{R}, s_2 \in \mathbf{S}, \\ a_1, a_2 \in \mathbf{A}}} \mathbb{I}[A_s^i = a_1, R_s^i = r_1, S_t^j = s_2, A_t^j = a_2] \\
& + \sum_{\substack{r_1, r_2 \in \mathbf{R}, \\ a_1, a_2 \in \mathbf{A}}} \mathbb{I}[A_s^i = a_1, R_s^i = r_1, A_t^j = a_2, R_t^j = r_2].
\end{aligned}$$

For example, the term $\mathbb{I}[R_s^i = r_1, S_s^i = s_1, R_t^j = r_2, S_t^j = s_2]$ returns “1” in the appropriate position, if $y^k = (i, j)$ aligns residue of type r_1 in secondary structure s_1 in the target with residue of type r_2 in secondary structure s_2 in the template. These features capture pairwise interaction of structural annotations within the same sequence.

Tensor: In this even more complex alignment model we consider the interaction of all three structural annotations. Note that there is only one non-zero feature in this feature vector.

$$\phi_{Tensor}(y^k, \mathbf{s}, \mathbf{t}) = \sum_{r_1, r_2 \in \mathbf{R}, s_1, s_2 \in \mathbf{S}, a_1, a_2 \in \mathbf{A}} \mathbb{I}[R_s^i = r_1, S_s^i = s_1, A_s^i = a_1, R_t^j = r_2, S_t^j = s_2, A_t^j = a_2]$$

Simple+Anova2+Tensor: This alignment model is the union of all features in the first three alignment models, i.e., $\phi_{Simple}(y^k, \mathbf{s}, \mathbf{t}) + \phi_{Anova2}(y^k, \mathbf{s}, \mathbf{t}) + \phi_{Tensor}(y^k, \mathbf{s}, \mathbf{t})$.

Window: On top of the *Simple+Anova2+Tensor* feature vector, we add several terms involving the substitution score of a sliding window of features centered around positions i and j .

$$\begin{aligned}\phi_{Window}(y^k, \mathbf{s}, \mathbf{t}) &= \phi_{Simple}(y^k, \mathbf{s}, \mathbf{t}) + \phi_{Anova2}(y^k, \mathbf{s}, \mathbf{t}) + \phi_{Tensor}(y^k, \mathbf{s}, \mathbf{t}) \\ &+ \sum_{r_1, r_2, r_3 \in \mathbf{R}, r_4, r_5, r_6 \in \mathbf{R}} \mathbb{I}[R_s^{i-1} = r_1, R_s^i = r_2, R_s^{i+1} = r_3, R_t^{j-1} = r_4, R_t^j = r_5, R_t^{j+1} = r_6] \\ &+ \sum_{s_1, \dots, s_5 \in \mathbf{S}, s_6, \dots, s_{10} \in \mathbf{S}} \mathbb{I}[S_s^{i-2} = s_1, \dots, S_s^{i+2} = s_5, S_t^{j-2} = s_6, \dots, S_t^{j+2} = s_{10}] \\ &+ \sum_{a_1, \dots, a_7 \in \mathbf{A}, a_8, \dots, a_{14} \in \mathbf{A}} \mathbb{I}[A_s^{i-3} = a_1, \dots, A_s^{i+3} = a_7, A_t^{j-3} = a_8, \dots, A_t^{j+3} = a_{14}]\end{aligned}$$

The first sliding window term counts the occurrence of substituting a triplet of residues (r_1, r_2, r_3) in the target with another triplet (r_4, r_5, r_6) in the template. The other two terms counts the occurrence of substitution of two windows of secondary structures of length 5, and the occurrence of substitution of two windows of surface area type of length 7, respectively. To reduce dimensionality of these features, we bin the residues into seven groups ($\{\text{A,G,P,S,T}\}$, $\{\text{C}\}$, $\{\text{D,E,N,Q}\}$, $\{\text{F,W,Y}\}$, $\{\text{H,K,R}\}$, $\{\text{I,L,M,V}\}$, $\{\text{X}\}$, where X stands for missing value and ends of sequences), and the surface area into two values, exposed or buried.

Profile: Finally, we explore adding profile information to the alignment model, and we use P_s^i to denote the profile (a real-valued vector of dimension 20 for 20 amino acids) at the i th position of the target sequence. $P_s^i(r_1)$ denotes the weight assigned to amino acid r_1 by the profile P_s^i . The definitions for profiles in the template P_t^j are similar. The *Profile* alignment model is very similar to the *Window* alignment model, with the single interaction component ϕ_{Simple} replaced by a version that includes profile information $\phi_{SimpleWithProfile}$:

$$\begin{aligned}\phi_{SimpleWithProfile}(y^k, \mathbf{s}, \mathbf{t}) &= \phi_{Simple}(y^k, \mathbf{s}, \mathbf{t}) + \sum_{r_1 \in \mathbf{R}, r_2 \in \mathbf{R}} \mathbb{I}[R_s^i = r_1] P_t^j(r_2) + \sum_{s_1 \in \mathbf{S}, r_2 \in \mathbf{R}} \mathbb{I}[S_s^i = s_1] P_t^j(r_2) \\ &+ \sum_{a_1 \in \mathbf{A}, r_2 \in \mathbf{R}} \mathbb{I}[A_s^i = a_1] P_t^j(r_2) + \sum_{r_1 \in \mathbf{R}, r_2 \in \mathbf{R}} P_s^i(r_1) \mathbb{I}[R_t^j = r_2] + \sum_{r_1 \in \mathbf{R}, s_2 \in \mathbf{S}} P_s^i(r_1) \mathbb{I}[S_t^j = s_2] \\ &+ \sum_{r_1 \in \mathbf{R}, a_2 \in \mathbf{A}} P_s^i(r_1) \mathbb{I}[A_t^j = a_2] + \sum_{r_1 \in \mathbf{R}, r_2 \in \mathbf{R}} P_s^i(r_1) P_t^j(r_2)\end{aligned}$$

We only incorporate profile information into the single interaction component ϕ_{Simple} instead of higher-order interaction components like ϕ_{Anova2} since our goal is to test the usefulness of including profile information in our alignment models, but not to find the best alignment model that captures all the interactions between all features. This alignment model can be expressed as:

$$\begin{aligned}\phi_{Profile}(y^k, \mathbf{s}, \mathbf{t}) &= \phi_{SimpleWithProfile}(y^k, \mathbf{s}, \mathbf{t}) + \phi_{Anova2}(y^k, \mathbf{s}, \mathbf{t}) + \phi_{Tensor}(y^k, \mathbf{s}, \mathbf{t}) \\ &+ \sum_{r_1, r_2, r_3 \in \mathbf{R}, r_4, r_5, r_6 \in \mathbf{R}} \mathbb{I}[R_s^{i-1} = r_1, R_s^i = r_2, R_s^{i+1} = r_3, R_t^{j-1} = r_4, R_t^j = r_5, R_t^{j+1} = r_6] \\ &+ \sum_{s_1, \dots, s_5 \in \mathbf{S}, s_6, \dots, s_{10} \in \mathbf{S}} \mathbb{I}[S_s^{i-2} = s_1, \dots, S_s^{i+2} = s_5, S_t^{j-2} = s_6, \dots, S_t^{j+2} = s_{10}] \\ &+ \sum_{a_1, \dots, a_7 \in \mathbf{A}, a_8, \dots, a_{14} \in \mathbf{A}} \mathbb{I}[A_s^{i-3} = a_1, \dots, A_s^{i+3} = a_7, A_t^{j-3} = a_8, \dots, A_t^{j+3} = a_{14}]\end{aligned}$$

TABLE 1. Q-SCORE OF THE SVM ALGORITHM FOR DIFFERENT ALIGNMENT MODELS
(AVERAGE OF 5CV, STANDARD ERROR IN BRACKETS)

	No. of features	Training	Validation	Test
<i>Simple</i>	1020	31.93	28.53 (2.34)	42.08 (2.02)
<i>Anova2</i>	49,634	41.25	33.73 (0.87)	44.57 (0.78)
<i>Tensor</i>	203,280	55.27	33.20 (0.78)	42.41 (0.18)
<i>Simple+Anova2+Tensor</i>	253,934	47.77	34.39 (0.55)	44.59 (0.28)
<i>Window</i>	447,016	44.40	35.23 (1.05)	46.34 (0.70)
<i>Profile</i>	448,642	54.58	38.50 (0.51)	49.94 (0.11)

6.1.2. *Gap cost model.* All alignment models above share the following gap model. Consider the cost of opening a gap between position i and $i + 1$ in the target sequence $\mathbf{s}$ against position j in the template structure $\mathbf{t}$, as depicted by the following diagram

$$\begin{array}{ccccccc} \text{Target} & \mathbf{s}^i & - & - & \dots & - & \mathbf{s}^{i+1} \\ \text{Template} & \dots & \mathbf{t}^j & \mathbf{t}^{j+1} & \dots & \mathbf{t}^{j+k} & \dots \end{array}$$

We allow the cost of opening a gap to depend on the structural type at position j in the template structure. It also depends on the structural type of the target sequence immediately before the gap at position i as well as the structural type immediately after the gap at position $i + 1$. Suppose y^k is a gap operation that opens a gap between position i and $i + 1$ in the target against position j in the template sequence. The feature vector for this gap operation is:

$$\begin{aligned} \phi_{\text{Gap}}(y^k, \mathbf{s}, \mathbf{t}) = & \sum_{r_1 \in R} \mathbb{G}[R_{\mathbf{t}}^j = r_1] + \sum_{s_1 \in S, a_1 \in A} \mathbb{G}[S_{\mathbf{t}}^j = s_1, A_{\mathbf{t}}^j = a_1] \\ & + \sum_{s_1, s_2 \in S, a_1, a_2 \in A} \mathbb{G}[S_{\mathbf{s}}^i = s_1, A_{\mathbf{s}}^i = a_1, S_{\mathbf{s}}^{i+1} = s_2, A_{\mathbf{s}}^{i+1} = a_2] \end{aligned}$$

$\mathbb{G}$ is analogous to $\mathbb{I}$, but we use a different symbol to indicate that it maps to a different set of dimensions. The first two terms create features for the residue types and joint features of secondary structure with exposed surface area at $\mathbf{t}^j$. The term $\mathbb{G}[S_{\mathbf{s}}^i = s_1, A_{\mathbf{s}}^i = a_1, S_{\mathbf{s}}^{i+1} = s_2, A_{\mathbf{s}}^{i+1} = a_2]$ considers the structure before and after the gap. For example, $\mathbb{G}[S_{\mathbf{s}}^i = \text{"}\alpha\text{"}, A_{\mathbf{s}}^i = \text{"}0\text{"}, S_{\mathbf{s}}^{i+1} = \text{"}\alpha\text{"}, A_{\mathbf{s}}^{i+1} = \text{"}1\text{"}]$ maps to the dimension for the cost of opening a gap between a position in an alpha-helix of surface type 0 with a consecutive position in the alpha-helix with surface type 1.

The case of opening a gap in the template involves exactly the same costs, with the role of target and template reversed.

6.1.3. *Results.* Table 1 shows the Q-scores of the different alignment models trained with the SVM algorithm using Q-loss. As described above, we report the fivefold cross validation results for the value of C that optimizes performance on the validation set. The table also shows the number of features in each model. Note that the training and the validation set are composed of more difficult cases than the test set, which explains the generally higher Q-scores on the test set.

Table 1 shows a general trend that, as we add features carefully, the performance of the alignment models increases.¹ The *Simple* alignment model is too simple to fit the training data, indicated by the low Q-score on the training set. The weight vector obtained in this low dimension feature space is also not stable, as the standard error from the fivefold cross validation estimates are much higher than the other alignment models. The more expressive *Anova2* model leads to substantial improvement in Q-score over *Simple*. This shows that considering pairwise interaction between structural annotations is meaningful. The *Tensor* alignment

¹In the conference version of this paper, there was an error in the feature vector encoding in the *Simple* alignment model, leading to lower performance than reported here.

model does worse than *Anova2* on the test set but has similar Q-score on the validation set. There are signs of overfitting in the relatively high Q-score on the training set. Adding the substitution costs together, as in *Simple+Anova2+Tensor*, does not give us much improvement in accuracy. The performance of the model is very close to *Anova2*. Only when we incorporate structural information in the local neighborhood, as in the alignment model *Window*, do we see another jump in the Q-score on the test set. Adding evolutionary information in the form of PSI-BLAST profile further enhances the performance as in the alignment model *Profile*, which has Q-scores close to 50 on the test set. The Q-score of 49.94 in the alignment model *Profile* is substantially better than the Q-score of 42.08 of the *Simple* alignment model that we started with. To provide a baseline, the Q-score of BLAST is 25.88 on the test set.

6.2. Is training to different loss functions beneficial?

The SVM method allows the use of different loss functions during training. The Q-loss used in the previous subsection is rather stringent and does not necessarily summarize the quality of an alignment well. For example, if all the aligned positions are shifted by just 1, the Q-loss will jump from 0 to 1, which is roughly the same Q-loss as that of a completely random alignment. Furthermore, the Q-loss does not account for the approximate nature of the training alignments, since there is typically no single exact alignment in sequence to structure alignment that is clearly correct.

Instead of Q-loss, we now consider the Q_4 -loss function. Q_4 -loss counts a residue as correctly aligned if the shift from its position in the reference alignment is no more than 4. The Q_4 -loss function captures our intuition that small shifts in alignment could be tolerated, and such alignments should be differentiated from alignments that are completely wrong. We repeat our experiments on two alignment models from the last section, *Anova2* and *Window*, but this time we train them with Q_4 as the loss function. The results on the test set are shown in Table 2. For each table entry, we select C on the validation set with respect to the performance measure that is reported. As before, the Q-scores and standard errors are estimated using models trained from fivefold cross validation.

Table 2 shows that the models trained on Q_4 show better Q_4 performance on the test set. More surprisingly, the models trained on Q_4 also show better Q-score on the test set. Figure 2 shows the Q-score of the *Anova2* alignment model for different values of the regularization parameter C . We can see that models trained on the Q_4 -loss function dominate those trained on the Q-loss across the whole range of values of C . Similar trends are observed in other alignment models. This gives evidence that Q_4 can indeed effectively account for the inaccuracy of the training alignments, instead of trying to model the noise. However, in situations where the alignments have higher sequence similarity or we are more confident of the alignments, the use of Q-loss or reducing the allowable shift of 4 in Q_4 to lower values could be beneficial. The flexibility of the SVM regarding the selection of loss function would cater either of these situations.

6.3. How does the accuracy of SVM models compare to conventional methods?

As selected by validation performance, the best alignment model is *Profile* trained on Q_4 , while the best alignment model without the use of evolutionary information is *Window* trained on Q_4 . Table 3 shows the test set performance of various other methods in comparison. SSALN (Qiu and Elber, 2006) is one

TABLE 2. COMPARING TRAINING FOR Q-SCORE WITH TRAINING FOR Q_4 -SCORE BY TEST SET PERFORMANCE

	<i>Test Q</i>	<i>Test Q_4</i>
<i>Anova2</i>		
Train Q	44.57 (0.77)	66.30 (0.33)
Train Q_4	46.32 (0.06)	68.87 (0.44)
<i>Window</i>		
Train Q	46.34 (0.70)	67.58 (1.30)
Train Q_4	47.31 (0.21)	70.52 (0.33)

FIG. 2. Q-score of the *Anova2* alignment model against C .

TABLE 3. COMPARING TRAINING FOR Q-SCORE WITH TRAINING FOR Q_4 -SCORE BY TEST SET PERFORMANCE

	Q on test	Q_4 on test
SVM (<i>Window</i> , Q_4)	47.31 (0.21)	70.53 (0.33)
SVM (<i>Profile</i> , Q_4)	50.15 (0.19)	72.19 (0.26)
SSALN	47.49	67.51
BLAST	25.88	30.66
TM-align	70.21	85.45

of the best current alignment algorithm trained using generative methods, and it outperforms alignment and threading algorithms like CLUSTALW, GenTHREADER, and FUGUE on a variety of benchmarks. It incorporates structural information in its substitution matrices, and contains a hand-tuned gap model. The training set we use is derived from the same training set used by SSALN (we are using about 4500 alignments out of more than 5200 alignments in the SSALN training set), and the same set of structural annotations (with the exception of the last alignment model *Profile*), so a direct comparison is particularly meaningful. Using exactly the same set of features the SVM model *Window* substantially outperforms SSALN with respect to Q_4 -score, and essentially ties with SSALN on Q-score. Incorporating profile information makes the SVM model *Profile* performs even better. The performance of BLAST is included to provide a baseline. The performance of the structural alignment program TM-align (Zhang and Skolnick, 2005) is reported here to show its agreement with the CE alignments, and demonstrates the rather high inherent noise in the data.

7. DISCUSSION

This paper explores an SVM method for learning complex alignment models for sequence to structure alignment. We show that the algorithm can learn high-dimensional models that include many features

beyond residue identity while effectively controlling overfitting. Unlike generative methods, it does not require independence assumptions between features. The SVM method provides great modeling flexibility to biologists, allowing the estimation of models that include all available information without having to worry about statistical dependencies between features. Furthermore, we show that one can incorporate different loss functions during training, which provides the flexibility to specify the costs of different alignment errors. The empirical results show that the SVM algorithm outperforms one of the best current generative models, and is practical to train on large datasets.

ACKNOWLEDGMENTS

We thank Jian Qiu for his detailed answers on our many questions on SSALN, Tamara Galor-Neah for her help with the LOOPP program, and Phil Zigoris for his work on an early prototype of the SVM software. This work is supported by the NIH (grants IS10RR020889 and GM67823) and NSF (award IIS-0412894).

DISCLOSURE STATEMENT

No competing financial interests exist.

REFERENCES

- Adamczak, R., Porollo, A., and Meller, J. 2004. Accurate prediction of solvent accessibility using neural networks-based regression. *Proteins* 56, 753–767.
- Bucher, P., and Hofmann, K. 1996. A sequence similarity search algorithm based on a probabilistic interpretation of an alignment scoring system. *Int. Conf. ISMB*, pgs. 44–51.
- Dayhoff, M.O., Schwartz, R.M., and Orcutt, B.C. 1978. A model of evolutionary change in proteins. *Atlas Protein Seq. Struct.* 5, 345–352.
- Do, C.B., Gross, S.S., and Batzoglou, S. 2006. CONTRAlign: discriminative training for protein sequence alignment. *RECOMB*.
- Durbin, R., Eddy, S., Krogh, A., et al. 1998. *Biological Sequence Analysis*. Cambridge University Press, Cambridge, UK.
- Gusfield, D., and Stelling, P. 1996. Parametric and inverse-parametric sequence alignment with XPARAL. *Methods Enzymol.* 266, 481–494.
- Henikoff, S., and Henikoff, J.G. 1992. Amino acid substitution matrices from protein blocks. *Proc. Natl. Acad. Sci. USA* 89, 10915–10919.
- Joachims, T. 1998. Text categorization with support vector machines: learning with many relevant features. *Proc. Eur. Conf. Machine Learn.* 137–142.
- Joachims, T. 2003. Learning to align sequences: a maximum-margin approach. Available at: www.joachims.org. Accessed July 6, 2008.
- Joachims, T., Galor, T., and Elber, R. 2005. Learning to align sequences: a maximum-margin approach. *Lect. Notes Comput. Sci.* 49, 57–68.
- Kabsch, W., and Sander, C. 1983. Dictionary of protein secondary structure: pattern recognition of hydrogen bond and geometrical features. *Biopolymers* 22, 2577–2637.
- Kececioğlu, J.D., and Kim, E. 2006. Simple and fast inverse alignment. *Lect. Notes Comput. Sci.* 3909, 441–455.
- Lafferty, J., McCallum, A., and Pereira, F. 2001. Conditional random fields: probabilistic models for segmenting and labeling sequence data. *ICML*, pgs. 282–289.
- McCallum, A., Bellare, K., and Pereira, F. 2005. A conditional random field for discriminatively-trained finite-state string edit distance. *Conf. Uncertainty Artif. Intellig.*
- Pachter, L., and Sturmfelds, B. 2004. Parametric inference for biological sequence analysis. *Proc. Natl. Acad. Sci. USA* 101, 16138–16143.
- Qiu, J., and Elber, R. 2006. SSALN: an alignment algorithm using structure-dependent substitution matrices and gap penalties learned from structurally aligned protein pairs. *Proteins* 62, 881–891.
- Ristad, S.E., and Yianilos, P.N. 1998. Learning string edit distance. *IEEE Trans. Pattern Recogn. Machine Intellig.* 20, 522–532.

- Shindyalov, I.N., and Bourne, P.E. 1998. Protein structure alignment by incremental combinatorial extension (CE) of the optimal path. *Protein Eng.* 11, 739–747.
- Smith, T., and Waterman, M. 1981. Identification of common molecular subsequences. *J. Mol. Biol.* 147, 195–197.
- Sun, F., Fernandez-Baca, D., and Yu, W. 2002. Inverse parametric sequence alignment. *Int. Comput. Combinatorics Conf.*, pgs. 97–106.
- Taskar, B., Guestrin, C., and Koller, D. 2003. Maximum-margin Markov networks. *Neural Inform. Process. Syst.*
- Tsochantaridis, I., Hofmann, T., Joachims, T., et al. 2004. Support vector machine learning for interdependent and structured output spaces. *ICML*.
- Tsochantaridis, I., Joachims, T., Hofmann, T., et al. 2005. Large margin methods for structured and interdependent output variables. *J. Machine Learning Res.* 6, 1453–1484.
- Vapnik, V. 1998. *Statistical Learning Theory*. Wiley, Chichester, UK.
- Zhang, Y., and Skolnick, J. 2005. TM-align: a protein structure alignment algorithm based on TM-score. *Nucleic Acids Res.* 33, 2302–2309.

Address reprint requests to:
Mr. Chun-Nam John Yu
Department of Computer Science
4130 Upson Hall
Ithaca, NY 14850

E-mail: cnyu@cs.cornell.edu