

Ensemble Multiple Sequence Alignment via Advising

Dan DeBlasio^{*}
Department of Computer Science
The University of Arizona
Tucson AZ 85721, USA
deblasio@cs.arizona.edu

John Kececioğlu
Department of Computer Science
The University of Arizona
Tucson AZ 85721, USA
kece@cs.arizona.edu

ABSTRACT

The multiple sequence alignments computed by an aligner for different *settings* of its parameters, as well as the alignments computed by different *aligners* using their default settings, can differ markedly in accuracy. *Parameter advising* is the task of choosing a parameter setting for an aligner to maximize the accuracy of the resulting alignment. We extend parameter advising to *aligner advising*, which in contrast chooses among a set of aligners to maximize accuracy. In the context of aligner advising, *default* advising selects from a set of aligners that are using their default settings, while *general* advising selects both the aligner and its parameter setting.

In this paper, we apply aligner advising for the first time, to create a true *ensemble aligner*. Through cross-validation experiments on benchmark protein sequence alignments, we show that parameter advising boosts an aligner’s accuracy beyond its default setting for virtually all of the standard aligners currently used in practice. Furthermore, aligner advising with a collection of aligners further improves upon parameter advising with any single aligner, though surprisingly the performance of default advising on testing data is actually superior to general advising due to less overfitting to training data.

The new ensemble aligner that results from aligner advising is significantly more accurate than the best single default aligner, especially on hard-to-align sequences. This successfully demonstrates how to construct out of a collection of individual aligners, a more accurate ensemble aligner.

Categories and Subject Descriptors

J.3 [Life and Medical Sciences]: Biology and genetics;
F.2.2 [Nonnumerical Algorithms and Problems]: Computations on discrete structures

^{*}Corresponding author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
BCB’15, September 09–12, 2015, Atlanta, GA, USA.
Copyright 2015 ACM ISBN 978-1-4503-3853-0/15/09...\$15.00.
DOI: <http://dx.doi.org/10.1145/2808719.2808766>.

General Terms

Algorithms, Experimentation, Performance

Keywords

Multiple sequence alignment, parameter advising, aligner advising, accuracy estimation, ensemble methods.

1. INTRODUCTION

While it has long been known that the multiple sequence alignment computed by an aligner strongly depends on the settings for its tunable parameters, and that different aligners using their default settings can output markedly different alignments of the same input sequences, there has been relatively little work on how to automatically choose the best parameter settings for an aligner, or the best aligner to invoke, to obtain the most accurate alignment of a given set of input sequences.

Automatically choosing the best parameter setting for an aligner on a given input was termed by Wheeler and Kececioğlu [34], *parameter advising*. In their framework, an advisor takes a *set* of parameter settings, together with an *estimator* that estimates the accuracy of a computed alignment, and invokes the aligner on each setting, evaluates the accuracy estimator on each resulting alignment, and chooses the setting that gives the alignment of highest estimated accuracy. Analogously, we call automatically choosing the best aligner for a given input, *aligner advising*.

To make this concrete, Figure 1 shows an example of advising on a benchmark set of protein sequences for which a correct reference alignment is known, and hence for which the true accuracy of a computed alignment can be determined. In this example, the **Facet** estimator of DeBlasio and Kececioğlu [4] is used to estimate the accuracy of two alignments computed by the **Opal** [35] and **MUMMALS** [23] aligners. For these two alignments, the one of higher **Facet** value also has higher true accuracy as well, so an advisor armed with the **Facet** estimator would in fact output the more accurate alignment to a user.

For a collection of aligners, this kind of advising is akin to an *ensemble* approach to alignment, which selects a solution from those output by different methods to obtain in effect a new method that ideally is better than any individual method. Ensemble methods have been studied in machine learning [37], which combine the results of different classifiers to produce a single output classification. Typically such ensemble methods from machine learning select a

<i>d1flma</i>	8	...	fewlknegvv A IATQgedgphlvntwvnsylkv-ldgnriwvpggmhkteanva-rde	...	63
<i>d1ci0a</i>	18	...	tkv-fn-----skedpret-----lpsa T FS----- A elpsg	...	46
<i>d1nrga</i>	16	...	saw-fe-----savgqpdig----- e nam C LAT----- C t-rdg	...	43
<i>d1ejea</i>	22	...	hrltprptv M VTVdeeginaapfsftmpvsidppvafasapdhhtarnie-sth	...	78
<i>d1i0ra</i>	8	...	ykisygly I V T SEsngkcgqiant--vfgltkpvqiavclnkendthnavk- e sg	...	61

(a) Lower-accuracy alignment computed by MUMMALS

<i>d1flma</i>	1	...	-----mlpgtffevlkn-----gvv A IATQg-edgph--lvntwvnsylk---vldg	...	41
<i>d1ci0a</i>	12	...	d-dpidlftkwfneakedpretlpsa T FS A elpsg---vssrilllfc---eldh	...	59
<i>d1nrga</i>	9	...	slpvpkqfawfeavgcpdigeanam C LAT C t-rdgk---psarmllik---gfg	...	56
<i>d1ejea</i>	11	...	s-mdfedfpvesahrltpr---ptv M VTV V -eegn---inaapfsftmpvsidp	...	56
<i>d1i0ra</i>	1	...	--mdveafykisy-----gily I V T SE-sngkcgqiantvfqilt--s-kp	...	39

(b) Higher-accuracy alignment computed by `Opal`

Figure 1: *Aligner choice affects the accuracy of computed alignments.* (a) Part of an alignment of benchmark `sup_125` from the SABRE [32] suite computed by MUMMALS [23] using its default parameter choice; this alignment has accuracy value 28.9%, and Facet [14] estimator value 0.540. (b) Alignment of the same benchmark by `Opal` [34] using its default parameter choice, which has 49.9% accuracy, and higher Facet value 0.578. In both alignments, the positions that correspond to *core blocks* of the reference alignment, which should be aligned in a correct alignment, are highlighted in bold.

result by *voting*. In contrast, an advisor combines the results of aligners by selecting one via an *estimator*.

In this paper, we extend the framework of parameter advising to aligner advising, and obtain by this natural approach a true *ensemble aligner*. Moreover as our experimental results show, the resulting ensemble aligner is significantly more accurate than any individual aligner.

1.1 Related work

We very briefly summarize prior work on alignment advising. Wheeler and Kececioglu [34] first introduced the notion of *parameter advisors*; Kececioglu and DeBlasio [14] investigated the construction of alignment *accuracy estimators*, resulting in the Facet estimator [4]; DeBlasio and Kececioglu [6] investigated how to best form the set of parameter choices for an advisor, called an *advisor set*, developing an efficient approximation algorithm for finding a near-optimal advisor set for a given estimator. This prior work applied parameter advising to boosting the accuracy of the `Opal` aligner [35]. In contrast, this paper applies *parameter advising* to all commonly-used aligners, and *aligner advising* to combine them into a new, more accurate, ensemble aligner.

To our knowledge, the only prior work on combining aligners is by Wallace, O’Sullivan, Higgins, and Notredame [33] on *M-Coffee*, and by Muller, Creevey, Thompson, Arendt, and Bork [21] on *AQUA*. The *AQUA* tool chooses between an alignment computed by MUSCLE [9] or MAFFT [13] based on their *NorMD* [31] score; our prior work [14] shows that for choosing the more accurate alignment, the *NorMD* score used by *AQUA* is much weaker than the Facet estimator used here for aligner advising. *M-Coffee* uses a standard progressive alignment heuristic to compute an alignment under position-dependent substitution scores whose values are determined by alignments from different aligners. As Section 4.3 later shows, when run on the same set of aligners, *M-Coffee* is strongly dominated by the ensemble approach of this paper.

1.2 Plan of the paper

In the next section, we review our approach to learning an alignment advisor. An advisor selects aligners and param-

eter values from a small set of choices that is drawn from a larger universe of all possible choices. Section 3 describes how we construct this universe of aligners and their parameter choices for advisor learning. Section 4 then experimentally evaluates our approach to ensemble alignment on real biological benchmarks. Finally, Section 5 gives conclusions, and offers directions for further research.

2. LEARNING AN ALIGNMENT ADVISOR

To make the paper self-contained, we briefly review our prior work on how to learn an alignment advisor. We first review the concept of *parameter advising*, which requires an estimator of alignment accuracy and a set of parameter choices for the advisor, and then summarize our prior techniques for learning both an estimator and an advisor set. In Section 4, we apply these techniques for the first time to *aligner advising* to yield a new ensemble aligner.

2.1 Parameter advising

The goal of parameter advising is to find the parameter setting for an aligner that yields the most accurate alignment of a given set of input sequences. The accuracy of a computed alignment is measured with respect to the “correct” alignment of the sequences (which often is not known). For special benchmark sets of protein sequences, the gold-standard alignment of the proteins, called their *reference alignment*, is usually obtained through structural alignment by finding the best superposition of the known three-dimensional structures of the proteins. Columns of the reference alignment that contain a residue from every protein in the set (where a *residue* is the amino acid at a particular position in a protein), and for which the residues in the column are all mutually close in space in the superposition of the structures, are called *core columns*. Runs of consecutive core columns are called *core blocks*, which represent the regions of the structural alignment with the highest confidence of being correct. Given such a reference alignment with identified core blocks, the *accuracy* of a different, computed alignment is the fraction of the pairs of residues aligned in the core blocks of the reference alignment that are also aligned in the computed alignment. (So a computed alignment of 100% accuracy completely agrees with the reference on its core blocks, though it may disagree elsewhere.) The best computed alignment is one of highest accuracy, and the task of a parameter advisor is to find a setting of the tunable parameters of an aligner that yields an accurate output alignment. This setting can be highly input dependent, as the best choice of parameter values for an aligner can vary for different sets of input sequences.

When aligning sequences in practice, a reference alignment is almost never known, in which case the true accuracy of a computed alignment cannot be measured. Instead our parameter advisors rely on an accuracy *estimator* E that for an alignment A , gives a value $E(A)$ in the range $[0, 1]$ that estimates the true accuracy of alignment A . An estimator should be efficiently computable and positively correlated with true accuracy.

To choose a parameter setting, an advisor takes a set of choices P , where each *parameter choice* $p \in P$ is a vector that assigns values to all the tunable parameters of an aligner, and picks the choice that yields a computed alignment of highest estimated accuracy.

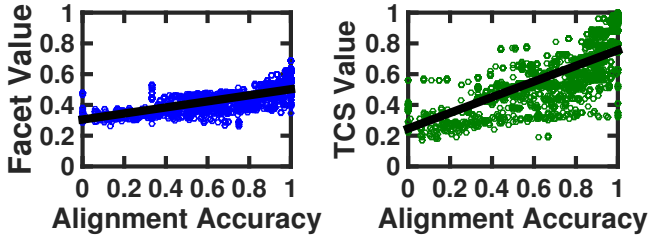


Figure 2: *Relationship of estimators to true accuracy.* Each point in a scatterplot corresponds to an alignment whose true accuracy is on the horizontal axis, and whose value under a given estimator is on the vertical axis. Both scatterplots show the same set of 3,000 alignments (randomly sampled from the more than 209,000 alignments generated by the experiments in Section 4) under the accuracy estimators *Facet* [14] and *TCS* [3].

Formally, given an accuracy estimator E and a set P of parameter choices, a *parameter advisor* tries each parameter choice $p \in P$, invokes an aligner to compute an alignment A_p using choice p , and then selects the parameter choice p^* that has maximum estimated accuracy:

$$p^* \in \operatorname{argmax}_{p \in P} \{E(A_p)\}.$$

Since the advisor runs the aligner $|P|$ times on a given set of input sequences, a crucial aspect of parameter advising is finding a small set P for which the true accuracy of the output alignment A_{p^*} is high.

To construct a good advisor, we need to find a good estimator E and a good set P . The estimator and advisor set are learned on training data consisting of benchmark sets of protein sequences for which a reference alignment is known. The learning procedure tries to find an estimator E and set P that maximize the true accuracy of the resulting advisor on this training data, which we subsequently assess on separate testing data.

Note that the process of advising is fast: for a set P of k parameter choices, advising involves computing k alignments under these choices, which can be done in parallel, evaluating the estimator on these k alignments, and taking a max. (Section 4.8 gives actual running times.) The separate process of training an advisor, by learning an estimator and an advisor set as we review next, is done once, off-line, before any advising takes place.

2.2 Learning an accuracy estimator

Kececioğlu and DeBlasio [7, 14] present an efficient approach for learning an accuracy estimator that is a linear combination of real-valued alignment feature functions, based on solving a large-scale linear programming problem. This approach resulted in the *Facet* estimator [4], which is currently the most accurate estimator for parameter advising [14, 6].

This approach assumes we have a collection of d real-valued feature functions $g_1(A), \dots, g_d(A)$ on alignments A , where these functions g_i are positively correlated with true accuracy. The alignment accuracy estimator E is a linear combination of these functions, $E(A) = \sum_{1 \leq i \leq d} c_i g_i(A)$, where the coefficients c_i specify the estimator E . When the feature functions have range $[0, 1]$ and the coefficients form a convex combination, the resulting estimator E will also have

range $[0, 1]$. *Facet* uses a collection of five feature functions, many of which make use of predicted secondary structure for the protein sequences [14]. Figure 2 shows the relationship to true accuracy of both the *Facet* and *TCS* [3] estimators.

A parameter advisor uses the estimator to effectively rank alignments, so an estimator just needs to be monotonic in true accuracy. The *difference-fitting* approach learns the coefficients of an estimator that is close to monotonic by fitting the estimator to differences in true accuracy for pairs of training alignments.

Let function $F(A)$ give the *true accuracy* of alignment A , and set \mathcal{P} be a collection of ordered pairs of alignments from training data, where every pair $(A, B) \in \mathcal{P}$ satisfies $F(A) < F(B)$. Difference fitting tries to find an estimator E that increases at least as much as accuracy F on the pairs in \mathcal{P} , by minimizing the amount that E falls short. Formally, we find the estimator E^* given by the vector of coefficients $c^* \in \mathcal{R}^d$ that minimizes

$$\sum_{(A, B) \in \mathcal{P}} w_{AB} \max \left\{ (F(B) - F(A)) - (E(B) - E(A)), 0 \right\},$$

where w_{AB} weights the above error for a pair (A, B) . Finding the optimal coefficients c^* can be reduced to solving a linear programming problem as follows.

The linear program has a variable c_i for each estimator coefficient, and an error variable e_{AB} for each pair $(A, B) \in \mathcal{P}$. The constraints are $c_i \geq 0$ and $\sum_i c_i = 1$, which ensure the coefficients form a convex combination, together with $e_{AB} \geq 0$, and

$$e_{AB} \geq (F(B) - F(A)) - (E(B) - E(A)).$$

(Note that the expressions $E(A)$ and $E(B)$ are linear in the variables c_i , while the quantities $F(A)$ and $F(B)$ are constants.) The linear program then minimizes the objective function $\sum_{(A, B) \in \mathcal{P}} w_{AB} e_{AB}$.

For the linear program to be of manageable size for a large number of training alignments, the set \mathcal{P} of pairs must be quite sparse. Kececioğlu and DeBlasio [14] describe how to find a good sparse set \mathcal{P} together with a good set of weights w_{AB} by an efficient graph algorithm.

Learning an accuracy estimator with d feature functions using a set \mathcal{P} of p pairs, involves solving the above linear program with $p+d$ variables and $\Theta(p+d)$ inequalities. Evaluating the *Facet* estimator on an alignment with m sequences and n columns, after secondary structure has been predicted for the protein sequences, takes $\Theta(m^2 n)$ time.

2.3 Learning an advisor set

DeBlasio and Kececioğlu [5, 6] present an efficient approximation algorithm for learning a near-optimal set of parameter choices for an advisor that uses a given estimator. The approximation algorithm follows a greedy strategy, so we call the sets found by the approximation algorithm *greedy sets*, in contrast to *exact sets* that are optimal for the training data, and which can be found by exhaustive search for small instances. These greedy sets tend to generalize better than exact sets, with the remarkable behavior that the greedy sets often outperform exact sets on testing data [6].

The problem of learning an optimal set P of parameter choices for an advisor is formulated as follows. Let U be the *universe* of possible parameter choices that might be included in advisor set P . (Section 3 describes how we con-

struct the universe U for aligner advising.) The training data is a collection of *reference alignments* R_i , one for each benchmark B_i , and a collection of *alternate alignments* A_{ij} , where each alignment A_{ij} is computed by running the aligner on the sequences in benchmark B_i using parameter choice $j \in U$. By comparing each alternate alignment to the reference alignment for its benchmark, we can measure the *true accuracy* a_{ij} of each alignment A_{ij} .

For a candidate set $P \subseteq U$ of parameter choices for an advisor that uses estimator E , the set of parameter choices from P that could potentially be *output* by the advisor on benchmark B_i is

$$\mathcal{O}_i(P) = \operatorname{argmax}_{j \in P} \{E(A_{ij})\},$$

where the argmax gives the set of parameter choices $j \in P$ that are tied for maximizing the estimator E on the benchmark. The advisor could output any parameter choice from $\mathcal{O}_i(P)$, as all of them appear equally good under the estimator, so we assume the advisor selects a choice uniformly at random from this set. Then the *expected accuracy* achieved by the advisor on benchmark B_i using parameter set P is

$$\mathcal{A}_i(P) = \frac{1}{|\mathcal{O}_i(P)|} \sum_{j \in \mathcal{O}_i(P)} a_{ij},$$

where again a_{ij} is the true accuracy of alignment A_{ij} .

In learning an advisor set P , we seek a set P that maximizes the advisor’s expected accuracy $\mathcal{A}_i(P)$ on the training benchmarks B_i . Formally, we want a set P that maximizes the objective function

$$f(P) = \sum_i w_i \mathcal{A}_i(P),$$

where i indexes the benchmarks, and w_i is the weight placed on benchmark B_i . (The benchmark weights correct for sampling bias in the training data, as discussed in Section 4.) In words, we want to find an advisor set $P \subseteq U$ that maximizes the expected accuracy of the parameter choices selected by the advisor, averaged across weighted training benchmarks.

DeBlasio and Kececioglu [6] prove that for a given bound k on the size of advisor set P , finding an optimal set $P \subseteq U$ with $|P| \leq k$ that maximizes objective $f(P)$ is NP-complete.

2.3.1 Greedy sets

While this NP-completeness result implies it is unlikely we can efficiently find an *optimal* advisor set, there is a natural greedy algorithm that is guaranteed to efficiently find a *near-optimal* set. For any constant ℓ , the optimal advisor set of cardinality at most ℓ can be found in polynomial time by exhaustive search. The following procedure **Greedy** builds on this idea to find a near-optimal advisor set for cardinality bound k , by starting with an optimal set of size at most ℓ , where $\ell \leq k$, and greedily augmenting it.

procedure Greedy(k, ℓ) begin

Find an optimal subset $P \subseteq U$ of size $|P| \leq \ell$ that maximizes $f(P)$.

$\tilde{P} := P$

$\tilde{\ell} := |P|$

for cardinalities $\tilde{\ell}+1, \dots, k$ **do begin**

Find parameter choice $j^* \in U - \tilde{P}$ that maximizes $f(\tilde{P} \cup \{j^*\})$.

$\tilde{P} := \tilde{P} \cup \{j^*\}$

if $f(\tilde{P}) \geq f(P)$ **then**

$P := \tilde{P}$

end

output P

end

DeBlasio and Kececioglu [6] prove that procedure **Greedy** is an (ℓ/k) -approximation algorithm for finding an optimal advisor set, for any constant ℓ with $\ell \leq k$.

Learning a greedy advisor set for cardinality bound k when $\ell = 1$, which is the value we use in practice, for the **Facet** estimator on a universe of u parameter choices and a training set of t benchmarks, takes $\Theta(k^2 ut)$ time.

2.3.2 Oracle sets

A useful notion in parameter advising introduced by Wheeler and Kececioglu [34] is the concept of an oracle, which is a perfect advisor that has access to the true accuracy of an alignment. For a given advisor set P , an *oracle* selects parameter choice $\operatorname{argmax}_{p \in P} \{F(A_p)\}$, where again function F gives the true accuracy of an alignment. (Equivalently, an oracle is an advisor that uses the perfect estimator F .) An oracle always picks the parameter choice that yields the highest accuracy alignment.

While an oracle is impossible to construct in practice, it gives a theoretical limit on the accuracy achievable by advising with a given set. Furthermore, if we can find the optimal advisor set for an oracle for a given cardinality bound k , which we call an *oracle set*, then the performance of an oracle on an oracle set gives a theoretical limit on how well advising can perform for a given bound k on the number of parameter choices.

Kececioglu and DeBlasio [14] show that finding an oracle set is NP-complete, and give the following integer linear program for finding an optimal oracle set. In our prior notation, an optimal oracle set for cardinality bound k is a set $P \subseteq U$ with $|P| \leq k$ that maximizes $\sum_i w_i \max_{j \in P} a_{ij}$. To formulate this as an integer linear programming problem, let y_j for all $j \in U$, and x_{ij} for all benchmarks B_i and all $j \in U$, be integer variables that either have the value 0 or 1, where $y_j = 1$ iff parameter choice j is selected for oracle set P , and $x_{ij} = 1$ iff choice j is used by the oracle on benchmark B_i . The constraints are $0 \leq x_{ij} \leq y_j \leq 1$, $\sum_j y_j \leq k$, and for each benchmark B_i the constraint $\sum_j x_{ij} = 1$. The objective function is to maximize $\sum_i w_i \sum_j x_{ij} a_{ij}$. An optimal solution to this integer linear program gives an optimal oracle set $P^* = \{j \in U : y_j = 1\}$.

Learning an optimal oracle set of cardinality k , for a universe of u parameter choices and a training set of t benchmarks, involves solving the above integer linear program with $\Theta(ut)$ variables and $\Theta(ut)$ inequalities. Kececioglu and DeBlasio [14] show that using the CPLEX integer linear programming solver, this formulation permits finding optimal oracle sets in practice even for cardinalities up to $k = 25$.

3. CONSTRUCTING THE UNIVERSE FOR ALIGNER ADVISING

We extend *parameter advising* with a single aligner to *aligner advising* with a collection of aligners, by having the choices in the advisor set now specify both a particular aligner and a parameter setting for that aligner. To specify the *universe* that such an advisor set is drawn from during learning, we must determine what aligners to consider, and what parameter settings to consider for those aligners.

3.1 Determining the universe of aligners

For *default aligner advising*, where the advisor set consists of distinct aligners, each using their default parameter setting, we learned advisor sets over a universe containing as many of the commonly-used aligners from the literature as possible. Specifically, the universe for default advising consisted of the following 17 aligners: Clustal [30], Clustal2 [15], Clustal Omega [27], DIALIGN [28], FSA [2], Kalign [16], MAFFT [13], MUMMALS [23], MUSCLE [10], MSAProbs [19], Opal [34], POA [17], PRANK [20], Probalign [25], ProbCons [8], SATé [18], and T-Coffee [22].

3.2 Determining the universe of parameter settings

For *general aligner advising*, we selected a subset of the above aligners on which we enumerated values for their tunable parameters, to form a universe of parameter settings. We selected this subset of aligners by the following process. First, we computed an optimal oracle set of cardinality $k=5$ (as described in Section 2.3.2) over the universe of 17 aligners for default advising listed above. This set consisted of Kalign, MUMMALS, Opal, Probalign, and T-Coffee. We then expanded this set further by adding four aligners that are used extensively in the literature: Clustal Omega, MAFFT, MUSCLE, and ProbCons. In the experiments described later in Section 4.2, we constructed greedy advisor sets over the universe of 17 aligners for default aligner advising, and noticed a large increase in advising accuracy at cardinality [6, 8] (which can be seen in Figure 8). The greedy advisor sets at these cardinalities contained all of the aligners already chosen so far, with the addition of the PRANK aligner. Finally, we added PRANK to our set for this reason. The above 10 aligners comprise the set we considered for general aligner advising.

Table 1 lists the universe of parameter settings for these aligners for general advising. For each aligner, we enumerated parameter settings by forming a cross product of values for each of its tunable parameters. We determined the values for each tunable parameter by one of two ways. For aligners with web-server versions (namely Clustal Omega and ProbCons), we used all values recommended for each parameter. For all other aligners, we chose either one or two values above and below the default value for each parameter, to attain a cross product with less than 200 parameter settings. If a range was specified for a numeric parameter, values were chosen to cover this range as evenly as possible. For non-numeric parameters, we used all available options. Table 1 summarizes the resulting universe for general advising of over 800 parameter settings.

4. EVALUATING ENSEMBLE ALIGNMENT

We evaluate the performance of advising through experiments on a collection of protein multiple sequence alignment

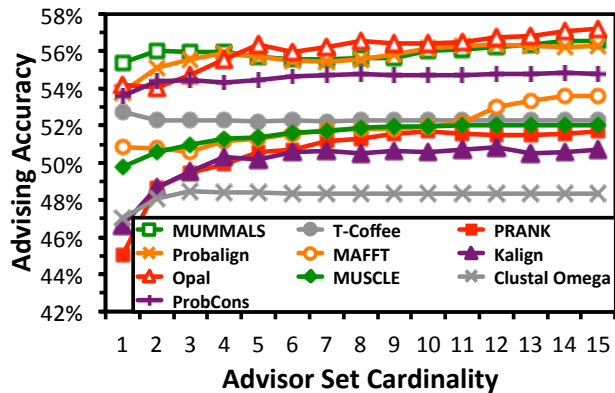


Figure 3: Accuracy of parameter advising using Facet. The plot shows advising accuracy for each aligner from Table 1, using parameter advising on greedy sets with the Facet estimator learned by difference fitting. The horizontal axis is the cardinality of the advisor set, and the vertical axis is the advising accuracy on testing data averaged over all benchmarks and folds, under 12-fold cross-validation.

benchmarks. A full description of the benchmark collection is given in [14], and is briefly summarized below. The experiments compare the accuracy of parameter and aligner advising to the accuracy of individual aligners using their default parameter settings.

The benchmark suites used in our experiments consist of reference alignments that are largely induced by performing structural alignment of the known three-dimensional structures of the proteins. Specifically, we use the BENCH suite of Edgar [11], supplemented by a selection of benchmarks from the PALI suite [1]. The entire benchmark collection consists of 861 reference alignments.

As is common in benchmark suites, easy-to-align benchmarks are highly over-represented in this collection, compared to hard-to-align benchmarks. To correct for this bias when evaluating average advising accuracy, we binned the 861 benchmarks in our collection by *difficulty*, where the difficulty of a benchmark is its average accuracy under three commonly-used aligners, namely Clustal Omega, MAFFT, and ProbCons, using their default parameter settings. We then divided the full range $[0, 1]$ of accuracies into 10 bins with difficulties $[(j-1)/10, j/10]$ for $j = 1, \dots, 10$. The weight w_i of benchmark B_i falling in bin j that we used for training is $w_i = (1/10)(1/n_j)$, where n_j is the number of benchmarks in bin j . These weights w_i are such that each difficulty bin contributes equally to the advising objective function $f(P)$. Note that with this weighting, an aligner that on every benchmark gets an accuracy equal to its difficulty, will achieve an average advising accuracy of roughly 50%.

4.1 Parameter advising

We first examine the results of parameter advising for a single aligner using the Facet estimator. We learned the coefficients for Facet by difference fitting on computed alignments obtained using the oracle set of cardinality $k=17$ found for the parameter universe for each aligner. (We trained the estimator on an oracle set of this cardinality to match the size of the universe for default aligner advising.) Given this estimator, we constructed greedy advisor sets for each aligner.

Table 1: *Universe of Parameter Settings for General Aligner Advising*

Aligner	Parameter settings	Tunable parameters	Version	Parameter name	Default value, v	Alternate values
Clustal Omega [27]	120 ¹	5	1.2.0	Number of guide tree iterations	0	1, 3, 5
				Number of HMM iterations	0	1, 3, 5
				Number of combined iterations	0	1, 3, 5
				Distance matrix calculations, initial	mBed	Full alignments
				Distance matrix calculations, iterations	mBed	Full alignments
Kalign [16]	162	4	2.04	Gap open penalty	55	40, 70
				Gap extension penalty	8.5	7, 10
				Terminal gap penalty	4.25	3.5, 5
				Bonus	No	Yes
MAFFT [13]	100	3	6.923b	Substitution matrix	BLSM62	BLSM80, VTML120, VTML200
				Gap open penalty	1.53	$\frac{1}{4}v, \frac{1}{2}v, \frac{3}{2}v, 2v$
				Gap extension penalty	0.123	$\frac{1}{2}v, 2v, 4v$
MUSCLE [10]	80	3	3.8.31	Profile score	Log-expectation: VTML240	Sum-of-pairs: PAM200, VTML240
				Objective function ²	spm	dp, ps, sp, spf, xp
				Gap open penalty, profile dependent	v^3	$\frac{1}{2}v, \frac{3}{4}v, \frac{5}{4}v, \frac{3}{2}v$
MUMMALS [23]	29	3 ⁴	1.01	Differentiate match states in unaligned regions	Yes	No
				Solvent accessibility categories	1	2, 3
				Secondary structure types	3	1
Opal [34]	162	5	3.0b	Substitution matrix	VTML200 ⁵	BLSM62 ⁵ , VTML40 ⁵
				Internal gap open penalty	$\gamma = 45$	70, 95
				Terminal gap open penalty	0.4γ	$0.05\gamma, 0.75\gamma$
				Internal gap extension penalty	$\lambda = 42$	40, 45
PRANK [20]	50	3	.140603	Terminal gap extension penalty	$\lambda - 3$	λ
				Gap rate	0.005	$\frac{1}{5}v, \frac{1}{2}v, \frac{3}{2}v, 2v$
				Gap extension	0.5	$\frac{1}{5}v, \frac{1}{2}v, \frac{3}{2}v, 2v$
				Terminal gaps	Alternate scoring	Normal scoring
Probalign [25]	64	3	1.4	Consistency repetitions	2	0, 1, 3
				Iterative refinement repetitions	100	0, 500
				Pre-training repetitions	0	1, 2, 3, 4, 5, 20
ProbCons [8]	48 ⁶	3	1.12	Thermodynamic temperature	5	3, 5
				Gap open	22	11, 33
				Gap extension	1	0.5, 1.5
T-Coffee [22]	36	3	10.00.r1613	Substitution matrix	BLSM62	BLSM40, BLSM80
				Gap open	0	-50, -500, -1000
				Gap extension	0	-5, -10
Total	856					

¹Parameter settings retrieved from the Clustal Omega web-server at EBI (www.ebi.ac.uk/Tools/msa/clustalo).

²sp: sum-of-pairs score; spf: dimer approximation of sum-of-pairs score; spm: input dependent (sp if input is less than 100 sequences, spf otherwise); dp: dynamic programming score; ps: average profile sequence score; xp: cross profile score.

³Default values for the gap open penalty are -2.9 when the log-expectation profile is chosen, -1439 for sum-of-pairs using PAM200, and -300 for sum-of-pairs using VTML240. Alternate values are multiples of this default value.

⁴MUMMALS is distributed with 29 precomputed hidden Markov models, each of which is associated with a setting of three tunable parameters.

⁵The substitution matrices used by Opal are shifted, scaled, and rounded to integer values in the range [0, 100].

⁶Parameter settings retrieved from the ProbCons web-server at Stanford (probcons.stanford.edu).

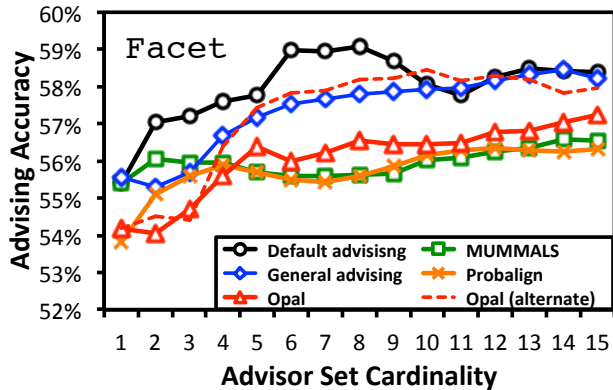


Figure 4: *Aligner advising and parameter advising using Facet.* The plot shows default and general aligner advising accuracy, and parameter advising accuracy for *Opal*, *MUMMALS*, and *Probalign*, using the *Facet* estimator. The horizontal axis is the cardinality of the advisor set, and the vertical axis is advising accuracy on testing data averaged over all benchmarks and folds under 12-fold cross-validation.

Figure 3 shows the accuracy of parameter advising using greedy advisor sets of cardinality $k \leq 15$, for each of the 10 aligners in Table 1, under 12-fold cross-validation. The plot shows advising accuracy on the testing data, averaged over all benchmarks and folds.

Almost all aligners benefit from parameter advising, though their advising accuracy eventually reaches a plateau. While our prior work [6] showed that parameter advising boosts the accuracy of the *Opal* aligner, Figure 3 shows this result is not aligner dependent.

4.2 Aligner advising

To evaluate aligner advising, we followed a similar approach, constructing an oracle set of cardinality $k=17$ over the union of the universe for default advising from Section 3.1 and the universe for general advising from Section 3.2, learning coefficients for *Facet* using difference fitting, and constructing greedy sets using *Facet* for default and general advising.

Figure 4 shows the accuracy of default and general advising using greedy sets of cardinality $k \leq 15$, along with the three best parameter advising curves from Figure 3, for *Opal*, *Probalign*, and *MUMMALS*. The plot shows advising accuracy on testing data, averaged over benchmarks and folds.

The dashed red curve in Figure 4 also shows the accuracy of *Opal* for parameter advising with greedy sets computed over an alternate universe of much more fine-grained parameter choices. To construct this alternate universe, we first started with a set of over 16,000 parameter settings, and for each training fold chose the 15 parameters with highest average accuracy on each difficulty bin. Unioning these choices across the bins, and removing duplicates, gave a universe of 147 to 150 settings for each fold. Note that the dashed curve for parameter advising with *Opal*, using greedy sets from these finer universes for each fold, essentially matches the accuracy of general advising at cardinality $k \geq 4$.

4.2.1 Testing the significance of improvement

To test the statistical significance of the improvement in default advising accuracy over using a single default aligner,

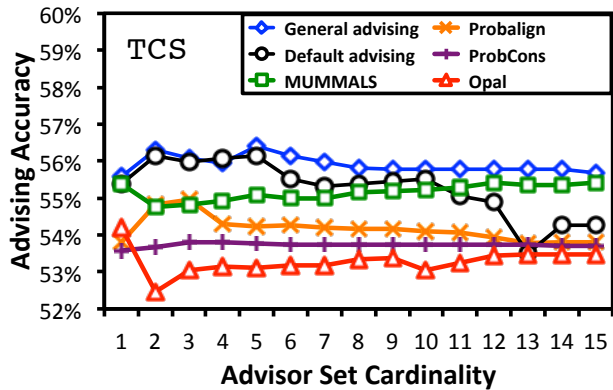


Figure 5: *Aligner advising and parameter advising using TCS.* The plot shows default and general aligner advising accuracy, and parameter advising accuracy for *Opal*, *MUMMALS*, *Probalign*, and *ProbCons*, using the *TCS* estimator. The horizontal axis is the cardinality of the advisor set, and the vertical axis is advising accuracy on testing data averaged over all benchmarks and folds under 12-fold cross-validation.

we used the one-tailed Wilcoxon sign test [36]. Performing this test in each difficulty bin, we found a significant improvement in accuracy ($p < 0.05$) on benchmarks with difficulty $(0.3, 0.4]$ at all cardinalities $2 \leq k \leq 15$, and on benchmarks with difficulty at most 0.4 at cardinality $6 \leq k \leq 9$.

We also tested the significance of the improvement of default advising over the best parameter advisor at each cardinality k (namely *MUMMALS* for $k \leq 4$ and *Opal* for $k \geq 5$), and found that at cardinality $k \geq 5$ there is again significant improvement ($p < 0.05$) on benchmarks with difficulty $(0.3, 0.4]$.

4.2.2 Advising with an alternate estimator

We also evaluated in the same way parameter advising and aligner advising on greedy sets using the *TCS* estimator [3] (the best other estimator for advising from the literature). Figure 5 shows results using *TCS* for parameter advising (on the four most accurate aligners), and for general and default aligner advising. Note that while *TCS* is sometimes able to increase accuracy above using a single default parameter, this increase is smaller than for *Facet*; moreover, *TCS* often has a decreasing trend in accuracy for increasing cardinality.

4.3 Comparing ensemble alignment to meta-alignment

Another approach to combining aligners is the so-called *meta-alignment* approach of *M-Coffee* [33] (described in Section 1.1). *M-Coffee* computes a multiple alignment using position-dependent substitution scores obtained from alternate alignments generated by a collection of aligners. By default, *M-Coffee* uses the following six aligners: *Clustal2*, *T-Coffee*, *POA*, *MUSCLE*, *MAFFT*, *Dialign-T* [29], *PCMA* [24], and *ProbCons*. The tool also allows use of *Clustal*, *Clustal Omega*, *Kalign*, *AMAP* [26], and *Dialign-TX*. Figure 6 shows the average accuracy of both *M-Coffee* and our ensemble approach with *Facet*, using the default aligner set of *M-Coffee* (the dotted vertical line with large circles), as well as oracle sets constructed over this *M-Coffee* universe of 13 aligners.

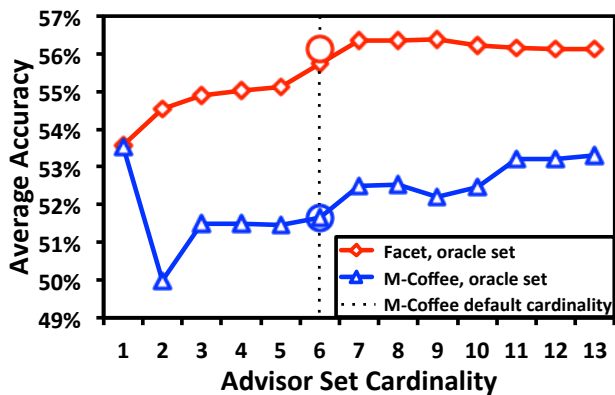


Figure 6: Accuracy of aligner advising compared to M-Coffee. The plot shows average accuracy for aligner advising using Facet, and meta-alignment using M-Coffee, on oracle sets of aligners. Performance on the default M-Coffee set of six aligners is indicated by large circles on the dotted vertical line. The horizontal axis is cardinality of the oracle sets, and the vertical axis is average accuracy on testing data over all benchmarks and folds under 12-fold cross-validation.

4.4 Advising accuracy within difficulty bins

Figure 7 shows advising accuracy within difficulty bins for default aligner advising compared to using the default parameter settings for the three aligners with highest average accuracy, namely MUMMALS, Opal, and Probalign. The figure displays the default advising result from Section 4.2 at cardinality $k=5$. The bars in the chart show average accuracy over the benchmarks in each difficulty bin, as well as the average accuracy across all bins. (The number of benchmarks in each bin is in parentheses above the bars.) Note that aligner advising gives the greatest boost for the hardest-to-align benchmarks: for the bottom two bins, advising yields an 8% increase in accuracy over the best aligner using its default parameter setting.

4.5 Generalization of aligner advising

The results thus far have shown advising accuracy averaged over the testing data associated with each fold. We now compare the training and testing advising accuracy to assess how our method might generalize to data not in our benchmark set.

Figure 8 shows the average accuracy of default and general aligner advising on both training and testing data. Note that the drop between training and testing accuracy is much larger for general advising than for default advising, resulting in general advising performing worse than default advising though its training accuracy is much higher. This indicates that general advising is strongly overfitting to the training data, but could potentially achieve much higher testing accuracy. Additionally, there is a drop in training accuracy for default advising with increasing cardinality, though after its peak an advisor using greedy sets should remain flat in training accuracy as cardinality increases when using a strong estimator.

4.6 Theoretical limit on advising accuracy

An *oracle* is an advisor that uses a perfect estimator, always choosing the alignment from a set that has highest true ac-

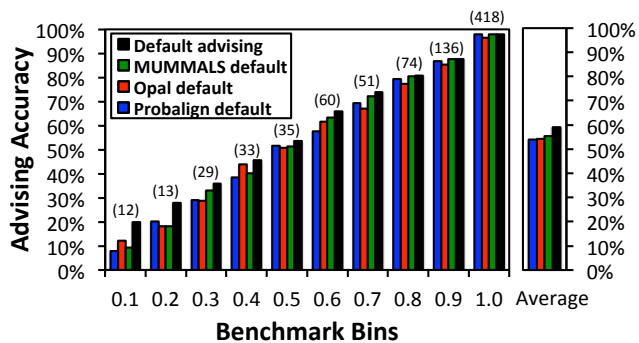


Figure 7: Accuracy of default aligner advising, and aligners with their default settings, within difficulty bins. In the bar chart on the left, the horizontal axis shows all ten benchmark bins, and the vertical bars show accuracy averaged over just the benchmarks in each bin. The accuracy of default advising using the Facet estimator is shown for the greedy sets of cardinality $k=5$, along with the accuracy of the default settings for Probalign, Opal, and MUMMALS. The bar chart on the right shows accuracy uniformly averaged over the bins. In parentheses above the bars are the number of benchmarks in each bin.

curacy. To examine the theoretical limit on how well aligner advising can perform, we compare the accuracy of aligner advising using Facet with the performance of an oracle. Figure 9 shows the accuracy of both default and general aligner advising using greedy sets, as well as the performance of an oracle using oracle sets computed on the default and general advising universes. (Recall an *oracle set* is an optimal advisor set for an oracle.) The plot shows advising accuracy on testing data, averaged over all benchmarks and folds. The large gap in performance between the oracle and an advisor using Facet shows the increase in accuracy that could potentially be achieved by developing an improved estimator.

4.7 Composition of advisor sets

Table 2 lists the greedy advisor sets for both default and general advising for all cardinalities $k \leq 10$. A consequence of the construction of greedy advisor sets is that the greedy set of cardinality k consists of the entries in a column in the first k rows of the table. The table shows these sets for just one fold from the 12-fold cross-validation. For general advising sets, an entry specifies the aligner that is used, and for aligners from the general advising universe, a tuple of parameter values in the order listed in Table 1. The two exceptions are MUMMALS, whose 6-tuple comes from its predefined settings file, and MSAProbs, whose empty tuple stands for its default setting. It is interesting that other than MSAProbs, the general advising set does not contain any aligner’s default parameter settings, though its values are close to the default setting.

4.8 Running time for advising

We compared the time to evaluate the Facet estimator on an alignment to the time needed to compute that alignment by the three aligners used for determining alignment difficulty: Clustal Omega, MAFFT, and ProbCons. To compute the average running time for these aligners on a benchmark,

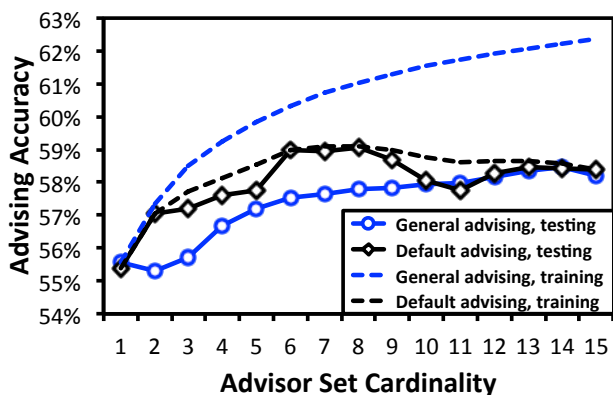


Figure 8: *General and default aligner advising on training and testing data.* The plot shows general and default aligner advising accuracy using *Facet*. Accuracy on the training data is shown with dashed lines, and on the testing data with solid lines. The horizontal axis is cardinality of the advisor set, and the vertical axis is advising accuracy averaged over all benchmarks and folds under 12-fold cross-validation.

Table 2: *Greedy Default and General Advising Sets*

	Default advising	General advising
1	MUMMALS	MUMMALS (0.2, 0.4, 0.6, 1, 2, 3)
2	Opal	Opal (VTML200, 45, 2, 45, 45)
3	Probalign	Opal (BLSM62, 70, 3, 45, 42)
4	Kalign	MUMMALS (0.15, 0.2, 0.6, 1, 1, 3)
5	MUSCLE	Opal (BLSM62, 45, 33, 42, 42)
6	T-Coffee	MSAProbs ()
7	PRANK	Kalign (55, 8.5, 4.25, 0)
8	Clustal Omega	MAFFT (VTML200, 0.7515, 0.492)
9	DIALIGN	Opal (BLSM62, 95, 4, 45, 42)
10	ProbCons	Opal (BLSM62, 45, 2, 45, 42)

we measured the total time for each of these aligners to align all 861 benchmarks on a desktop computer with a 2.4 GHz Intel i7 8-core processor and 8 Gb of RAM. The average running time for *Clustal Omega*, *MAFFT*, and *ProbCons* was less than 1 second per benchmark, as was the average running time for *Facet*. As stated in Section 2.2, the time complexity for *Facet* is linear in the number of columns in an alignment, and should take relatively less time than computing an alignment for benchmarks with long sequences; the standard benchmark suites tend to include short sequences, however, which are fast to align. This time to evaluate *Facet* does not include the time to predict protein secondary structure, which is done once for the sequences in a benchmark, and was performed using *PSIPRED* [12] version 3.2 with its standard settings. Secondary structure prediction with a tool like *PSIPRED* has a considerably longer running time than alignment, due to an internal *PSI-BLAST* search during prediction; on average, *PSIPRED* took just under 6 minutes per benchmark to predict secondary structure.

5. CONCLUSION

In this work, we have extended parameter advising to *aligner advising*, to yield a true *ensemble aligner*. Parameter advising gives a substantial boost in accuracy for nearly all aligners currently used in practice. Furthermore, default and general aligner advising both yield further boosts in accu-

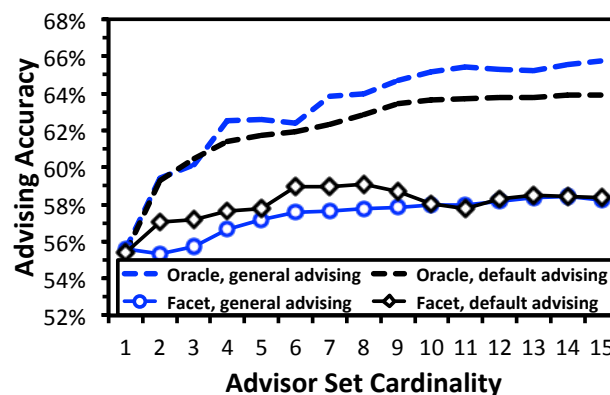


Figure 9: *Accuracy of aligner advising using a perfect estimator.* The plot shows advising accuracy for default and general aligner advising, both on oracle sets for a perfect estimator, and on greedy sets for the *Facet* estimator. The horizontal axis is the cardinality of the advisor set, and the vertical axis is advising accuracy on testing data averaged over all benchmarks and folds under 12-fold cross-validation.

racy, with default advising having better generalization. As these results indicate, ensemble alignment by aligner advising is a promising approach for exploiting future advances in aligner technology.

5.1 Further research

An important question left to explore is how to learn advisor sets that have improved *generalization*. While greedy advisor sets for general aligner advising achieve very high accuracy on training data, this does not translate to similar accuracy on testing data due to overfitting. Standard techniques from machine learning for improving generalization like regularization do not apply here, as the number of parameters for each aligner and the number of choices in the advisor set are both fixed. Applying this advising framework to *DNA* and *RNA* sequence alignment also seems fruitful.

6. ACKNOWLEDGEMENTS

We thank the reviewers for their helpful comments. This work was supported by NSF Grant IIS-1217886 to J.K., and a PhD fellowship to D.D. from NSF Grant DGE-0654435.

7. REFERENCES

- [1] S. Balaji, S. Sujatha, S. S. C. Kumar, and N. Srinivasan. *PALI*: a database of Phylogeny and ALignment of homologous protein structures. *Nucleic Acids Research*, 29(1):61, Jan. 2001.
- [2] R. K. Bradley, A. Roberts, M. Smoot, S. Juvekar, J. Do, C. Dewey, I. Holmes, and L. Pachter. Fast statistical alignment. *PLoS Computational Biology*, 5(5):e1000392, May 2009.
- [3] J. M. Chang, P. D. Tommaso, and C. Notredame. *TCS*: A new multiple sequence alignment reliability measure to estimate alignment accuracy and improve phylogenetic tree reconstruction. *Molecular Biology and Evolution*, Apr. 2014.
- [4] D. F. DeBlasio and J. D. Kececioglu. *Facet*: software for accuracy estimation of protein multiple sequence alignments. facet.cs.arizona.edu, Jan. 2014.

- [5] D. F. DeBlasio and J. D. Kececioglu. Learning parameter sets for alignment advising. *Proceedings of the 5th ACM Conference on Bioinformatics, Computational Biology, and Health Informatics (BCB)*, pages 230–239, 2014.
- [6] D. F. DeBlasio and J. D. Kececioglu. Learning parameter-advising sets for multiple sequence alignment. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 2015.
- [7] D. F. DeBlasio, T. J. Wheeler, and J. D. Kececioglu. Estimating the accuracy of multiple alignments and its use in parameter advising. *Proceedings of the 16th Conference on Research in Computational Molecular Biology (RECOMB)*, pages 45–59, 2012.
- [8] C. B. Do, M. S. P. Mahabhashyam, M. Brudno, and S. Batzoglu. **ProbCons**: probabilistic consistency-based multiple sequence alignment. *Genome Research*, 15(2):330–340, Jan. 2005.
- [9] R. C. Edgar. **MUSCLE**: a multiple sequence alignment method with reduced time and space complexity. *BMC Bioinformatics*, 5(1):113, 2004.
- [10] R. C. Edgar. **MUSCLE**: multiple sequence alignment with high accuracy and high throughput. *Nucleic Acids Research*, 32(5):1792–1797, May 2004.
- [11] R. C. Edgar. **BENCH**. drive5.com/bench, Jan. 2009.
- [12] D. T. Jones. Protein secondary structure prediction based on position-specific scoring matrices. *Journal of Molecular Biology*, 292(2):195 – 202, 1999.
- [13] K. Katoh, K.-i. Kuma, H. Toh, and T. Miyata. **MAFFT** version 5: improvement in accuracy of multiple sequence alignment. *Nucleic Acids Research*, 33(2):511–518, Jan. 2005.
- [14] J. D. Kececioglu and D. F. DeBlasio. Accuracy estimation and parameter advising for protein multiple sequence alignment. *Journal of Computational Biology*, 20(4):259–279, Apr. 2013.
- [15] M. A. Larkin et al. **ClustalW** and **ClustalX** version 2.0. *Bioinformatics*, 23(21):2947–2948, Nov. 2007.
- [16] T. Lassmann and E. Sonnhammer. **Kalign**: an accurate and fast multiple sequence alignment algorithm. *BMC Bioinformatics*, 6(1):298, 2005.
- [17] C. Lee, C. Grasso, and M. F. Sharlow. Multiple sequence alignment using partial order graphs. *Bioinformatics*, 18(3):452–464, Mar. 2002.
- [18] K. Liu, T. J. Warnow, M. T. Holder, S. M. Nelesen, J. Yu, A. P. Stamatakis, and C. R. Linder. **SATé-II**: Very fast and accurate simultaneous estimation of multiple sequence alignments and phylogenetic trees. *Systematic Biology*, 61(1):90–106, Dec. 2011.
- [19] Y. Liu, B. Schmidt, and D. L. Maskell. **MSAProbs**: multiple sequence alignment based on pair hidden markov models and partition function posterior probabilities. *Bioinformatics*, 26(16):1958–1964, 2010.
- [20] A. Loytynoja and N. Goldman. An algorithm for progressive multiple alignment of sequences with insertions. *Proceedings of the National Academy of Sciences*, 102(30):10557–10562, July 2005.
- [21] J. Muller, C. J. Creevey, J. D. Thompson, D. Arendt, and P. Bork. **AQUA**: automated quality improvement for multiple sequence alignments. *Bioinformatics*, 26(2):263–265, 2010.
- [22] C. Notredame, D. G. Higgins, and J. Heringa. **T-Coffee**: A novel method for fast and accurate multiple sequence alignment. *Journal of Molecular Biology*, 302(1):205–217, Sept. 2000.
- [23] J. Pei and N. V. Grishin. **MUMMALS**: multiple sequence alignment improved by using hidden Markov models with local structural information. *Nucleic Acids Research*, 34(16):4364–4374, Aug. 2006.
- [24] J. Pei, R. Sadreyev, and N. V. Grishin. **PCMA**: fast and accurate multiple sequence alignment based on profile consistency. *Bioinformatics*, 19(3):427–428, 2003.
- [25] U. Roshan and D. R. Livesay. **Probalign**: multiple sequence alignment using partition function posterior probabilities. *Bioinformatics*, 22(22):2715–2721, Nov. 2006.
- [26] A. S. Schwartz and L. Pachter. Multiple alignment by sequence annealing. *Bioinformatics*, 23(2):e24–e29, 2007.
- [27] F. Sievers et al. Fast, scalable generation of high-quality protein multiple sequence alignments using **Clustal Omega**. *Molecular Systems Biology*, 7(1):539–539, Jan. 2011.
- [28] A. R. Subramanian, M. Kaufmann, and B. Morgenstern. **DIALIGN-TX**: greedy and progressive approaches for segment-based multiple sequence alignment. *Algorithms for Mol. Biology*, 3(1):6, 2008.
- [29] A. R. Subramanian, J. Weyer-Menkhoff, M. Kaufmann, and B. Morgenstern. **DIALIGN-T**: An improved algorithm for segment-based multiple sequence alignment. *BMC Bioinformatics*, 6(1), 2005.
- [30] J. D. Thompson, D. G. Higgins, and T. J. Gibson. **ClustalW**: improving the sensitivity of progressive multiple sequence alignment through sequence weighting, position-specific gap penalties and weight matrix choice. *Nucleic Acids Research*, 22(22):4673–4680, 1994.
- [31] J. D. Thompson, F. Plewniak, R. Ripp, J.-C. Thierry, and O. Poch. Towards a reliable objective function for multiple sequence alignments1. *Journal of Molecular Biology*, 314(4):937–951, Dec. 2001.
- [32] I. Van Walle, I. Lasters, and L. Wyns. **SABmark**: a benchmark for sequence alignment that covers the entire known fold space. *Bioinformatics*, 21(7):1267–1268, 2005.
- [33] I. M. Wallace, O. O’Sullivan, D. G. Higgins, and C. Notredame. **M-Coffee**: combining multiple sequence alignment methods with **T-Coffee**. *Nucleic Acids Research*, 34(6):1692–1699, 2006.
- [34] T. J. Wheeler and J. D. Kececioglu. Multiple alignment by aligning alignments. *Bioinformatics*, 23(13):i559–68, July 2007.
- [35] T. J. Wheeler and J. D. Kececioglu. **Opal**: software for sum-of-pairs multiple sequence alignment. opal.cs.arizona.edu, Jan. 2012.
- [36] F. Wilcoxon. Individual comparisons by ranking methods. *Biometrics Bulletin*, 1(6):pp. 80–83, 1945.
- [37] Z. Zhihua. *Ensemble Methods: Foundations and Algorithms*. Chapman and Hall, 2012.