



HAL
open science

A time warping approach to multiple sequence alignment

Ana Arribas-Gil, Catherine Matias

► **To cite this version:**

Ana Arribas-Gil, Catherine Matias. A time warping approach to multiple sequence alignment. 2016.
hal-01332542v1

HAL Id: hal-01332542

<https://hal.science/hal-01332542v1>

Preprint submitted on 20 Jun 2016 (v1), last revised 5 Jan 2017 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A time warping approach to multiple sequence alignment

Ana Arribas-Gil¹ and Catherine Matias²

June 16, 2016

1. Departamento de Estadística, Universidad Carlos III de Madrid, C/ Madrid, 126 - 28903 Getafe, Spain. E-mail: ana.arribas@uc3m.es
2. Sorbonne Universités, Université Pierre et Marie Curie, Université Paris Diderot, Centre National de la Recherche Scientifique, Laboratoire de Probabilités et Modèles Aléatoires, 4 place Jussieu, 75252 PARIS Cedex 05, FRANCE. E-mail: catherine.matias@math.cnrs.fr

Abstract

We propose an approach for multiple sequence alignment (MSA) derived from the dynamic time warping viewpoint and recent techniques of curve synchronization developed in the context of functional data analysis. Starting from pairwise alignments of all the sequences (viewed as paths in a certain space), we construct a median path that represents the MSA we are looking for. We establish a proof of concept that our method could be an interesting ingredient to include into refined MSA techniques. We present a simple synthetic experiment as well as the study of a benchmark dataset, together with comparisons with 2 widely used MSA softwares.

Key words and phrases: Alignment; Dynamic time warping; Multiple sequence alignment; Warping

1 Introduction

Multiple sequence alignment (MSA) is one of the most fundamental tasks in bioinformatics. While there are many attempts to handle comparative sequence analyses without relying on MSA, it still represents a starting point for most evolutionary biology methods. Pairwise sequence alignment has been conceptualized as early as the 1970's, starting with global alignments that attempt to align entire sequences [17] and then introducing a decade later local alignments that focus on the identification of subsequences sharing high similarity [23]. The standard computational formulation of both tasks is to maximize a scoring function obtained as the sum of the score for each aligned pair of residues (nucleotides or amino acids, the highest scores being attributed to pairs of residues with highest similarity), minus some gaps penalties. Since these seminal works, an abundant literature has flourished exploring the topic in many different directions, from the pairwise problem to the more complex task of aligning more than 3 sequences [one of the very first attempts appearing in 15], from exact solutions that scale exponentially with sequence lengths to faster heuristic approaches

used in the most common tools, and from the scoring formulation of the alignment problem that requires to choose the scoring parameters to probabilistic formulations in which those parameters are estimated [2, 6]. However, manually refined alignments continue to be superior to purely automated methods and there is a continuous effort to improve the accuracy of MSA tools [8]. We refer the reader to the reviews [8, 14, 18, 29] for more details on MSA.

Dynamic time warping (DTW) is a general version of the dynamic programming algorithm that solves exactly the pairwise biological sequence alignment problem. It is a well-known and general technique to find an optimal alignment between two given (time-dependent) sequences. In time series analysis, DTW is used for constructing an optimal alignment of two sequences with possible different lengths by stretching or contracting time intervals [12]. In functional data analysis, the time warping approach consists in modeling a set of curves exhibiting time and amplitude variation with respect to a common continuous process [16]. Thus, time warping techniques are used in many different areas concerned by sequence or curve comparisons, one of its most famous successes being on human-speech recognition [13].

Here, we propose a simple and fast procedure for MSA, inspired from recent techniques of curve synchronization developed in the context of functional data analysis [4, 24]. In this setup, one often observes a set of curves which are modeled as the composition of an 'amplitude process' governing their common behavior, and a 'warping process' inducing time distortion among the individuals. Specifically, $y_i(t) = x_i \circ h_i(t)$, $t \in [a, b]$, $i = 1, \dots, K$, are observed, with x_i being i.i.d. realisations of the amplitude process X , and h_i strictly monotone functions s.t. $h_i(a) = a$ and $h_i(b) = b$ being i.i.d. realisations of the warping process H . Aligning pairs of curves (that is eliminating time variation, which comes to estimating the warping functions h_i) is a first step before estimating the common amplitude process. These authors proposed to first estimate pairwise warping functions between all possible trajectories pairs which are then used to create estimators of the underlying individual warping functions in a second step. Sample means or more robust median-based estimates come into play to solve this second step. This procedure is an alternative to the widely used approach of template registration, that consists in aligning every observed curve to some given template, which should be a good estimate of the mean amplitude process. The drawback of this methodology is that it heavily relies on the choice of the template, which is not straightforward. Now, in the MSA context, the warping process is the insertion-deletion (or indel) process that stretches or contracts the initial sequence, while the amplitude process is the substitution process that modifies the value of the sequence base. The equivalent of template registration in the MSA context would be the alignment of every sequence to some estimate of the ancestral sequence, which is, of course, not available. However, exploiting the previous ideas, we show how pairwise alignments can be combined with a simple median-based approach to obtain an estimate of the multiple alignment of the sequences.

Our aim is to establish a proof of concept that our new method could be an interesting ingredient to include into refined MSA techniques. Indeed, the method is able to align a large number K of sequences (that are assumed to share a common ancestor) in a quite simple and fast manner, although a bit rough w.r.t. accuracy. We would like to stress that we do not claim to be competitive with actual aligners. Let us recall that there already exist many competitors to solve the MSA problem whose respective performances have

been largely explored [see for e.g. 21, and the references therein]. Here, we would rather like to point out to recent developments from curve synchronization that could open the way to new improvements in MSA. While we do not pretend to propose a most accurate method, it is important to note that our approach could be used as a starting point in an iterative refinement strategy [10]. Those strategies, included in many widely used tools such as ProbCons [5], MUSCLE [7], MAFFT [11] or MUMMALS [22], mostly consists in repeatedly dividing the set of aligned sequences into two random groups and realign those groups by optimizing an objective function. Thus, our simple and fast procedure could be combined with similar refinement strategies that would improve its performance. An advantage of our method is that it only uses pairwise comparison but (unlike progressive aligners) is not sensitive to any order on which we consider the sequence pairs. Indeed, progressive aligners rely on a guiding tree and progressively build $K - 1$ pairwise alignments from the union of the initial set of K sequences with consensus sequences built at the internal nodes of the tree. As such, the consensus sequences heavily depend on the guiding tree and the sequences ordering which impacts the resulting MSA. Moreover, progressive aligners are known to tend to propagate errors appearing at the early stages of the method. For instance, ClustalW [25] is criticized for being responsible of the heuristic rule "once a gap, always a gap" that makes errors in pairwise alignments to propagate through MSA. The method proposed here is not sensitive to that issue because a gap in an alignment pair will not necessarily be selected for the MSA.

The manuscript is organized as follows. Section 2 first recalls concepts coming from the pairwise alignment and time warping problems (Section 2.1), sets our framework (Section 2.2) and describes our procedure (Section 2.3). Then Section 3 presents our experiments, starting with synthetic datasets (Section 3.1), then relying on a benchmark dataset (Section 3.2), namely the Balibase dataset [26] and concludes with possible extensions of our approach (Section 3.3).

2 Framework and procedure

In what follows, we first introduce some ideas around pairwise alignment and the time warping problem. We then explain how to use these concepts in the context of multiple sequence alignment.

2.1 Pairwise alignment and time warping

For any two sequences $X = X_1 \cdots X_n$ and $Y = Y_1 \cdots Y_m$ with values in a finite alphabet \mathcal{A} , any (global) pairwise alignment of X and Y corresponds to an increasing path in the grid $[0, n] \times [0, m]$ composed by three elementary steps $\{(1, 0), (0, 1), (1, 1)\}$, as shown in Figure 1. Note that for biological reasons, such path is often restricted to never contain two consecutive steps in $\{(1, 0), (0, 1)\}$ (a gap in one sequence may not be followed by a gap in the other sequence). For notational convenience, we extend the path to $[-\epsilon, n] \times [-\epsilon, m]$ for some small $\epsilon > 0$. Now, any such path may be viewed as an increasing function $\phi_{X,Y} : [-\epsilon, n] \rightarrow [-\epsilon, m]$ with the convention that it is càd-làg (continuous on the right, limit on the left). By letting $\phi(u-)$ denote the left limit of ϕ at u , we moreover impose that $\phi_{X,Y}(0-) = 0$ and $\phi_{X,Y}(n) = m$. Note that we can define a unique generalized inverse function $\phi_{Y,X} := \phi_{X,Y}^{-1} : [-\epsilon, m] \rightarrow [-\epsilon, n]$ constrained to be càd-làg and such that

$\phi_{X,Y} \circ \phi_{X,Y}^{-1} = \phi_{X,Y}^{-1} \circ \phi_{X,Y} = Id$ (the identity function). From a graphical point of view, the path corresponding to $\phi_{X,Y}^{-1}$ is obtained as the symmetric of the path corresponding to $\phi_{X,Y}$ with respect to the diagonal line $y = x$ (see Figure 1).




Figure 1: Graphical representation of an alignment between two sequences $X_1 \cdots X_n$ and $Y_1 \cdots Y_m$. The blue and thick line corresponds to a càd-làg function $\phi_{X,Y}$ whose generalized inverse $\phi_{Y,X} = \phi_{X,Y}^{-1}$ is shown in dashed and red.

Functions $\phi_{X,Y}, \phi_{X,Y}^{-1}$ may be viewed as time warping functions that describe the homology (and thus also the indels) between sequences X, Y obtained from their alignment. Let us first explain this idea on a simple example. Consider an alignment between $X = ACAGTAGT$ and $Y = CTTAAG$ given as follows

$$\begin{array}{cccccccc} A & C & A & G & T & A & - & G & T \\ - & C & T & - & T & A & A & G & - \end{array}$$

(This alignment corresponds exactly to the thick and blue path depicted in Figure 1). In this alignment, any character from one sequence may be associated to a unique character from the other sequence as follows: when the position corresponds to a *match*, the character is associated to the matching character (in the other sequence) while at positions corresponding to indels, we (arbitrarily) associate it to the character in the previous matching position in the other sequence. In order to deal with the case where no such previous matching position exists (namely when the alignment begins with an indel), we introduce extra artificial characters (X_0, Y_0) at the beginning of each sequence, these two artificial characters being aligned with each other. In other words, we consider the alignment

$$\begin{array}{cccccccc} X_0 & A & C & A & G & T & A & - & G & T \\ Y_0 & - & C & T & - & T & A & A & G & - \end{array}$$

The bare alignment (the alignment without specification of nucleotides) of these two sequences is as follows

$$\begin{array}{cccccccc} \bullet & B & B & B & B & B & - & B & B \\ \bullet & - & B & B & - & B & B & B & - \end{array}$$

Here B stands for *base* and we do not specify to which character it corresponds. The bare alignment naturally appears to separate the indel process from the substitution process [see

for e.g. 3, 28]. Given the sequences, the knowledge of the bare alignment is sufficient to recover the corresponding full alignment. In this context, the time warping process is exactly the indel process that stretches or contracts an initial sequence, while the amplitude process is the substitution process that modifies the value of the sequence. We focus on the time warping process encoded in the bare alignment.

Let \rightarrow denote this (asymmetric) relation induced by the bare alignment. In our example, we have that X_1 is associated to the artificial character Y_0 and we denote $X_1 \rightarrow Y_0$. In the same way, $X_2 \rightarrow Y_1$, $X_3 \rightarrow Y_2$, $X_4 \rightarrow Y_2$, $X_5 \rightarrow Y_3$, $X_6 \rightarrow Y_4$, $X_7 \rightarrow Y_6$ and $X_8 \rightarrow Y_6$. Note that no position of X is associated to Y_5 . More importantly, these associations are exactly encoded in the mapping $\phi_{X,Y}$ as we have $X_u \rightarrow Y_v$ if and only if $\phi_{X,Y}(u-) = v$.

More generally, for any real value $u \in [0, n]$, letting $\lceil u \rceil$ denote the smallest integer larger or equal to u , we set $X(u) := X_{\lceil u \rceil}$ (and similarly for sequence Y). From the (bare) pairwise alignment of X, Y , we obtain a time warping function $\phi_{X,Y}$ such that this alignment is entirely described by the association $X(u) \rightarrow Y \circ \phi_{X,Y}(u)$. In what follows, we denote this functional association by $X \rightarrow Y \circ \phi_{X,Y}$. Note that we equivalently have $Y \rightarrow X \circ \phi_{X,Y}^{-1} = X \circ \phi_{Y,X}$.

We also mention that when the pairwise alignment between X, Y is extracted from a multiple sequence alignment containing at least 3 sequences, say X, Y, Z , the above relation \rightarrow should be associative in the sense that whenever $X \rightarrow Y \circ \phi$ and $Y \rightarrow Z \circ \phi'$, we should also have $X \rightarrow Z \circ \phi' \circ \phi$. This property will be used in the next section.

2.2 Multiple sequence alignment

In this section, we consider a set of sequences S^1, \dots, S^K with values in a finite alphabet \mathcal{A} and respective lengths n_1, \dots, n_K and we assume that they all share some latent ancestor sequence A with values in \mathcal{A} and length N . Our goal is to estimate the alignment of each S^i to the ancestor A and thus the global alignment of the set of sequences S^1, \dots, S^K , by relying on the set of pairwise alignments of each S^i to all the other sequences S^1, \dots, S^K . To do this, we will implicitly assume that a) the multiple sequence alignment of all the sequences S^i is well approximated by the alignment we would obtain from all the sequences $\{A, S^i, 1 \leq i \leq K\}$; b) all the pairwise alignments of S^i, S^j are good approximations to the extracted pair alignments from the multiple alignment of all sequences $\{A, S^i, 1 \leq i \leq K\}$.

First, any sequence S^i is derived from the common ancestor sequence A through an evolutionary process that can be encoded in the alignment of these two sequences. This alignment induces a time warping process $\phi_i := \phi_{S^i,A} : [-\epsilon, n_i] \rightarrow [-\epsilon, N]$ such that we have $S^i \rightarrow A \circ \phi_i$. For the moment, we assume that the warping functions ϕ_i 's are i.i.d. realisations of the same random process Φ . For identifiability reasons, we constrain Φ such that $\mathbb{E}(\Phi) = Id$ (see for example [4]). Moreover, note that assuming the ϕ_i 's are i.i.d. realisations of the same process boils down to considering a star tree leading from the ancestor sequence to any observed sequence S^i . In the case of an ultra-metric tree, we would obtain non independent but identically distributed realisations. We will discuss later how to handle more realistic setups.

Since the ancestral sequence A is not observed, the idea of [4] is to estimate the warping function ϕ_i from all the pairwise alignments of S^i with the other sequences. More precisely, when considering the pairwise alignment of S^i with S^j , we obtain a warping function

$g_{ji} = \phi_j^{-1} \circ \phi_i$ satisfying

$$S^i \rightarrow A \circ \phi_i \rightarrow S^j \circ \phi_j^{-1} \circ \phi_i = S^j \circ g_{ji}.$$

As a consequence, relying on the desired associative property of \rightarrow , we should have $S^i \rightarrow S^j \circ g_{ji}$ so that we can estimate the warping function g_{ji} from a pairwise alignment of $\{S^i, S^j\}$. Note that since pairwise alignment of $\{S^i, S^j\}$ does not depend on the order of the sequences, the proposed estimators satisfy $\hat{g}_{ji} = \hat{g}_{ij}^{-1}$. Then, we want to combine these estimates $\{\hat{g}_{ji}, 1 \leq j \leq K\}$ to obtain an estimate of ϕ_i . In [24], the authors use the empirical mean and notice that this makes sense since $\mathbb{E}(g_{ji}|\phi_i) = \phi_i$ and under the assumption $\mathbb{E}(\Phi) = Id$ [see also 4]. Let us recall that these works come from functional data analysis where the functions ϕ_i 's are real-valued functions. Here, we are dealing with discrete sequences and the empirical mean of a set of paths g_{ji} would not give a proper path and thus would not correspond to an alignment of S^i to the ancestor A . That is why instead of using empirical mean, we will rather rely on median values. Note that – in general – medians have the advantage of being more robust than mean values.

From an algorithmic perspective, a global alignment of $X = X_1 \cdots X_n$ and $Y = Y_1 \cdots Y_m$ may be encoded through the sequence of coordinates Z_1, \dots, Z_T , where T is the length of the alignment, $Z_1 = (0, 0)$, $Z_T = (n, m)$ and for any $1 \leq t \leq T - 1$, we have $Z_{t+1} - Z_t \in \{(1, 0), (0, 1), (1, 1)\}$. We let

$$\begin{aligned} \mathcal{Z}_{n,m} = \{ & (Z_1, \dots, Z_T), T \geq 1, Z_1 = (0, 0), Z_T = (n, m), \\ & \text{such that } \forall 1 \leq t \leq T - 1, Z_{t+1} - Z_t \in \{(1, 0), (0, 1), (1, 1)\} \}, \end{aligned} \quad (1)$$

be the set of possible alignments between X and Y . The link between the path sequence Z and the warping function $\phi_{X,Y}$ is as follows

$$\exists t \geq 1 \text{ such that } Z_t = (k, l) \iff \phi_{X,Y}(k-) = l.$$

Now, considering a median path boils down to computing median coordinates of points Z_t as explained in the next section.

2.3 The median procedure

In this section, we explain how for each sequence S^i , we combine the set of pairwise alignments of S^i with $S^j, 1 \leq j \leq K$ to obtain an estimate of the warping function ϕ_i . We thus fix a sequence S^i of length n_i and consider its pairwise alignments with all the other sequences $S^j, 1 \leq j \leq K$, including itself (the alignment path is then the identical function from $[0, n_i]$ to $[0, n_i]$). We first define an estimate of the true number N of homologous positions in our MSA (that is positions that were present in the ancestral sequence A , which are the only positions of A that may be recovered), as a median value (constrained to be an integer number) of the set of lengths $\{n_j, 1 \leq j \leq K\}$, namely

$$\hat{N} = \text{median}\{n_j, 1 \leq j \leq K\}.$$

More precisely, as a convention and for the rest of this work, we set the median of a sequence of integers (n_1, \dots, n_K) with ordered values $(n_{(1)}, \dots, n_{(K)})$ as follows. Whenever K is even

and the mean value m between the ordered statistics $n_{(\lfloor K/2 \rfloor)}$ and $n_{(\lfloor K/2 \rfloor + 1)}$ is not an integer, we (arbitrarily) set the median as $\lfloor m \rfloor$.

Now, we aim at defining a path from $[0, n_i]$ to $[0, \hat{N}]$ by using all the estimated paths \hat{g}_{ji} obtained by aligning $\{S^i, S^j\}$. For any position $1 \leq u \leq n_i$, we consider the character S_u^i throughout all the pairwise alignments. The character S_u^i may either be aligned with another character S_v^j or to a gap after position v in sequence S^j , as may be viewed in Figure 2 (left part). As already explained, each pairwise alignment of S^i with S^j corresponds to a sequence of coordinates $Z^j = (Z_1^j, \dots, Z_{T_j}^j) \in \mathcal{Z}_{n_i, n_j}$ as defined above, see Equation (1). Note that we dropped the index i on which the sequence Z^j also depends. Now, for each value $1 \leq u \leq n_i$, there is a unique time point t_u^j and two integers $v_{u,1}^j \leq v_{u,2}^j$ such that

$$Z_{t_u^j}^j = (u - 1, v_{u,1}^j) \text{ and } Z_{t_{u+1}^j}^j = (u, v_{u,2}^j).$$

In practice, either $v_{u,2}^j = v_{u,1}^j + 1$ when S_u^i is aligned to a character or $v_{u,2}^j = v_{u,1}^j$ when S_u^i is aligned to a gap. We then compute two median points

$$\tilde{Z}_1(u) = (u - 1, \text{median}_{1 \leq j \leq K} \{v_{u,1}^j\}) \quad \text{and} \quad \tilde{Z}_2(u) = (u, \text{median}_{1 \leq j \leq K} \{v_{u,2}^j\}), \quad (2)$$

which gives the part of the median path that concerns character S_u^i (Figure 2, left part). Note that at this step, we did not take into account the possible vertical steps in the alignments \hat{g}_{ji} . We consider this now. Indeed, the concatenation of the positions

$$(\tilde{Z}_1(1), \tilde{Z}_2(1), \tilde{Z}_1(2), \tilde{Z}_2(2), \dots, \tilde{Z}_1(n_i), \tilde{Z}_2(n_i))$$

is almost what we wanted, namely the median path that gives our estimate of the alignment between S^i and ancestral sequence A . However one should notice that for each value $1 \leq u \leq n_i - 1$, we either have

$$\tilde{Z}_2(u) = \tilde{Z}_1(u + 1) \text{ or } \tilde{Z}_2(u) < \tilde{Z}_1(u + 1).$$

In the first case ($\tilde{Z}_2(u) = \tilde{Z}_1(u + 1)$), one of the two points $\tilde{Z}_2(u)$ or $\tilde{Z}_1(u + 1)$ is redundant and should be removed to obtain a proper path in $\mathcal{Z}_{n_i, \hat{N}}$. Otherwise, vertical movements are included in the median path and we keep both values $\tilde{Z}_2(u)$ and $\tilde{Z}_1(u + 1)$.

Let us note that at the last position $u = n_i$, we have $\tilde{Z}_2(n_i) = \text{median}_{1 \leq j \leq K} n_j = \hat{N}$ so that the median path correctly ends at (n_i, \hat{N}) .

2.4 MSA from median warping

We now combine the different steps to obtain our global alignment procedure as explained in pseudocode in Algorithm 1. This algorithm outputs a set of homologous positions (Hom) as well as a set of inserted positions (Ins). Let us recall that in each sequence, homologous positions correspond to positions that are aligned with an ancestor position A_u , $1 \leq u \leq \hat{N}$. Then two homologous positions may be separated by insert runs which are not aligned. Indeed, the choice of how to put the letters in the insert regions is arbitrary and most MSA implementations simply left-justify insert regions.

Thus, the table Ins used in the algorithm has K rows and $\hat{N} + 1$ columns. Each entry (i, u) with $2 \leq u \leq \hat{N}$, of this table contains the number of insertions in S^i between two




Figure 2: On the left: a set of $K = 5$ affine vectors (thicked and blue lines) extracted from the K alignments of S^i to the set of sequences S^j , $1 \leq j \leq K$, considered only at position S_u^i . The median value is shown in dashed and red. On the right: a set of $K = 3$ paths (in blue, including diagonal path) standing for the alignment of S^i with all other S^j , $1 \leq j \leq K$, and the median path (in dashed and red).

homologous positions with respective ancestor positions A_{u-1} and A_u , while entry $(i, 1)$ (resp. entry $(i, \hat{N} + 1)$) is the number of insertions before first (resp. after last) homologous position. The table Hom has K rows and \hat{N} columns. Each entry (i, u) of this table contains the position in sequence S^i that is homologous to ancestor position A_u . When there is no such position (a deletion in sequence S^i), the entry is set to 0.

With respect to the algorithmic complexity of the proposed procedure, we can assume that the K sequences have average length similar to N . Then, the time complexity of the MSA from median warping (once the pairwise alignments are given) is $O(NK^2)$ since we need $K \cdot N$ iterations in which we compute (twice) a median over K values (the complexity of the median calculation being $O(K)$ relying for e.g. on the `Quickselect` algorithm [9]).

3 Results

3.1 Synthetic experiments

In this section, we propose a simple synthetic experiment in order to assess the performances of our approach. We start with an ancestral sequence A on the set of nucleotides $\mathcal{A} = \{A, C, G, T\}$ with length $N = 100$. From this ancestral sequence, we simulate i.i.d. sequences S^1, \dots, S^K as follows. We rely on the simple Thorne-Kishino-Felsenstein [28, hereafter TKF] model that includes an insertion-deletion process and that can be combined with any substitution process. Here we set parameters for TKF $\lambda = \mu = 0.03$ and use the Jukes-Cantor substitution model [see for e.g. 30] with substitution rate $\alpha = 0.1$ (all nucleotide frequencies being set to $1/4$). That is, we are simulating K nucleotide sequences related by a star tree and whose branch lengths are set to be equal, as done for instance in [1]. We repeat this experiment $M = 100$ times for each K , and we let the number of sequences vary in $\{10, 20, 30, 40, 50\}$.

To obtain multiple alignments with our procedure, we first conduct all the pairwise alignments between pairs of sequences through the Needleman-Wunsch algorithm for global alignment [17] as implemented in the `pairwiseAlignment` function of the R `Biostrings`

Algorithm 1: Pseudocode for MSA from median warping

```

// Pairwise alignments  $\text{Align}(i, j)$ ;
For any pair of sequences  $\{S^i, S^j\}$ , compute their pairwise alignment  $\text{Align}(i, j)$ .

// Reconstructing alignment  $\text{Path}(i)$  of each  $S^i$  with ancestor  $A$ ;
 $\tilde{Z}_2(0) \leftarrow ((0, 0))$ ;
Ins  $\leftarrow$  table of 0's with  $K$  rows and  $\hat{N} + 1$  columns;
Hom  $\leftarrow$  table of 0's with  $K$  rows and  $\hat{N}$  columns;
for  $i = 1$  to  $K$  do
  // Alignment initialization;
  Path( $i$ )  $\leftarrow \tilde{Z}_2(0)$ 
  for  $u = 1$  to  $n_i$  do
    Compute the sequence of median values  $\tilde{Z}_1(u), \tilde{Z}_2(u)$  at position  $u$  from
    Equation (2);
    // Update alignment;
    if  $\tilde{Z}_1(u) = \tilde{Z}_2(u - 1)$  then
      Path( $i$ )  $\leftarrow$  concatenate (Path( $i$ ),  $\tilde{Z}_2(u)$ )
    else
      Path( $i$ )  $\leftarrow$  concatenate (Path( $i$ ),  $\tilde{Z}_1(u), \tilde{Z}_2(u)$ )
    // Update insertion and homologous positions tables;
    if  $\tilde{Z}_1(u)[2] = \tilde{Z}_2(u)[2]$  then
      Ins[ $i, \tilde{Z}_1(u)[2] + 1$ ]  $\leftarrow$  Ins[ $i, \tilde{Z}_1(u)[2] + 1$ ] + 1
    else
      Hom[ $i, \tilde{Z}_2(u)[2]$ ]  $\leftarrow u$ 

// Compute the maximal number of insertions before first, between 2
and after last homologous positions, respectively;
NbIns[.]  $\leftarrow \max_{1 \leq i \leq K} \text{Ins}[i, \cdot]$ 

// Compute size of multiple alignment;
 $L \leftarrow \hat{N} + \sum_{j=1}^{\hat{N}+1} \text{NbIns}[j]$ ;

// Return the multiple alignment;
 $T \leftarrow$  table with  $K$  rows and  $L$  columns filled with gaps symbols;
Insert the homologous positions from table Hom at correct positions in table  $T$ ;
Insert the inserted positions from table Ins at correct positions in table  $T$ ;
Return( $T$ ) ;

```

library [20]. We set the parameters for the pairwise alignment as: gap opening penalty equal to -10 , gap extension penalty equal to -0.5 , nucleotide substitution matrix with diagonal values equal to 5 and non-diagonal values equal to -4 . These are commonly used as default parameter values for the global alignment of DNA sequences. That is, we do not look at optimising the alignment parameters, but just at showing that our procedure can produce reasonable results under general conditions.

In order to assess the performance of the method, we compute for each alignment two scores that measure its overall concordance with the simulated one, SP (sum of pairs) and TC (total column) scores (see next section for details), and we compare them to the scores obtained by two well established multiple alignment softwares, namely `ClustalW` [25] in its 2.1 version and `T-coffee` [19] in its 10.00.r1613 version. Finally, we also run our procedure on the reference pairwise alignments, that is, those extracted from the simulated multiple alignment for each pair of sequences. This is done to assess the performance of the method in the best case scenario, as a way to validate the proposed algorithm for combining pairwise alignments.

Results are presented in Figure 3. The first thing to notice is that our method, when used with the reference pairwise alignments, reaches the maximum score possible in almost every case. Although these are not realistic results since reference pairwise alignments are unknown in practice, they confirm the soundness of our method that combines pairwise alignments from median warping. With respect to the other three methods, in all the scenarios `T-coffee` always provides the best performance. In general, TC scores tend to get worse as the number of sequences increases whereas SP scores are more stable. However, when comparing our procedure (with estimated pairwise alignments) and `ClustalW` (which is a widely used software), we see that `ClustalW` is better for a small number of sequences, but from $K = 20$, the performance of our method is superior.

3.2 Benchmark results

MSA methods are typically benchmarked on sets of reference alignments, the most widely used being the `BALiBASE` [26]. It is a large scale benchmark specifically designed for multiple sequence alignment. It contains test cases based on 3D structural superpositions of protein sequences that are manually refined to ensure the correct alignment of conserved residues. It is organized into several reference sets, designed to represent real multiple alignment problems. Reference 1 contains alignments of equidistant sequences with 2 different levels of conservation (RV11 and RV12); Reference 2 (RV20) contains families aligned with one or more highly divergent 'orphan' sequences; Reference 3 (RV 30) contains divergent subfamilies; Reference 4 (RV 40) contains sequences with large N/C-terminal extensions; and finally Reference 5 (RV 50) contains sequences with large internal insertions. In addition, three separate Reference Sets (6–8), are devoted to the particular problems posed by sequences with transmembrane regions, repeats, and inverted domains [26]. These last 3 sets are rarely included in benchmark analyses and we exclude them from ours. For each reference set among the first five ones, except for RV 40, two versions of the same multiple alignments are provided: one with the original sequences (which is noted BB), and one with a shorter version of the sequences contained in the alignment (which is noted BBS). In general, BBS alignments should be easier to recover than the original ones, since the less alignable parts of the original sequences have been removed.




Figure 3: SP and TC scores on the synthetic data sets. Distribution over $M=100$ simulated multiple alignments for each value of K .

The BALiBASE dataset comes with a function called `baliscore` used to assess the quality of a MSA. Two different criteria are used to score an alignment: SP (sum of pairs) and TC (total column) scores [27]. These criteria only use core blocks of the alignment. SP score is the percentage of the correctly aligned residue pairs in the alignment. It is defined as the number of correctly aligned residue pairs found in the test alignment divided by the total number of aligned residue pairs in core blocks of the reference alignment. TC score is the percentage of the correct columns in the alignment. It is defined as the number of correctly aligned columns found in the test alignment divided by the total number of aligned columns in core blocks of the reference alignment.

In order to assess the performance of the method on the Balibase data sets, we proceed as for the simulated alignments of the preceding section. The only differences being that we now consider BLOSUM62 as the default substitution matrix for pairwise alignments (since we now have protein sequences) and that for the last two reference sets (RV40 and RV50) we consider “overlap” pairwise alignments instead of global ones. Indeed, these two sets are characterized by large differences in sequences lengths so it is convenient to allow for gaps at the beginning and the end of pairwise alignments. In this section we also report the results of `ClustalW` and `T-coffee` for reference. We refer to [21] for a full comparison of available methods on Balibase.

Figure 4 shows that while our method is less performant than `ClustalW` and `T-coffee`, it still provides reasonable results. These results should be put in perspective with the level of complexity of the method that is exceptionally low (only pairwise alignments are needed as input and a simple median path is computed). Moreover, as it was the case for the simulated alignments, the scores based on the reference pairwise alignments are almost always equal to 1 in all sets for both SP and TC scores. Again, these are not realistic results since the pairwise alignments extracted from the reference multiple alignment might be far from the optimal pairwise alignment between the two sequences for which no extra information on the remaining sequences is available. Nevertheless, these results are encouraging since once again they serve to validate our method from a theoretical point of view in the sense that the median warping approach provides sound results.

3.3 Conclusions

In this work, we propose a proof of concept that a simple method derived from recent techniques of curve synchronization in the context of functional data analysis could be of potential interest to MSA. Our method is able to align a large number of sequences in a quite simple and fast manner, although a bit rough w.r.t. accuracy. While we do not claim to be competitive with actual aligners, we believe that our procedure could be successfully included (for e.g. as a starting point) in more refined MSA techniques. While it is out of the scope of the current work to provide such a refined tool, our simulations as well as the use of the Balibase dataset establish that the method has good potential, particularly when looking at the almost perfect results obtained by relying on the (unknown) reference pairwise alignments (that is, the pairwise alignments extracted from the MSA). In particular, while our experiments rely here on exact scoring alignment with default parameter values, the method could be combined with more refined pairwise alignments, such as probabilistic methods that automatically select optimal scoring parameters for the sequences at stake.

While the method implicitly assumes that a star ultra-metric tree describes the evolution




Figure 4: SP and TC scores on the Balibase data sets. Distribution over all the multiple alignments of each reference set. Reference set RV 40 does not provide a short version of the sequences (BBS).

of the sequences from their common ancestor, it could be improved to gain in robustness w.r.t. this assumption. In particular, let us assume that additional to the sequences one has access to a non ultra-metric guide tree describing the common evolution of these sequences. Then, we propose to weight the sequences in a way inversely proportional to their distance to the root. Let d_i denote the evolutionary distance (i.e. branch length) from sequence S^i to root and $d_{\max} = \max_{1 \leq i \leq K} d_i$. We fix $\epsilon > 0$ and set the weight values to

$$w_i = \frac{1 - d_i / (d_{\max} + \epsilon)}{\sum_{1 \leq k \leq K} 1 - d_k / (d_{\max} + \epsilon)} \in (0, 1).$$

Then our method could be generalized to the computation of a weighted median path, namely relying on weighted median values instead of simple ones. In such a way, sequences farther from the ancestor will have a lower weight in the MSA. We leave this for later investigation.

References

- [1] A. Arribas-Gil. Parameter estimation in multiple hidden i.i.d. models from biological multiple alignment. *Statistical Applications in Genetics and Molecular Biology*, 9(1):10, 2010.
- [2] A. Arribas-Gil, E. Gassiat, and C. Matias. Parameter estimation in pair-hidden Markov models. *Scand J Stat*, 33(4):651–671, 2006.
- [3] A. Arribas-Gil, D. Metzler, and J.-L. Plouhinec. Statistical alignment with a sequence evolution model allowing rate heterogeneity along the sequence. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 6(2):281–295, 2009.
- [4] A. Arribas-Gil and H.-G. Müller. Pairwise dynamic time warping for event data. *Computational Statistics & Data Analysis*, 69(0):255 – 268, 2014.
- [5] C. B. Do, M. S. Mahabhashyam, M. Brudno, and S. Batzoglou. ProbCons: Probabilistic consistency-based multiple sequence alignment. *Genome Research*, 15(2):330–340, 2005.
- [6] R. Durbin, S. Eddy, A. Krogh, and G. Mitchison. *Biological sequence analysis: probabilistic models of proteins and nucleic acids*. Cambridge University Press, Cambridge, UK, 1998.
- [7] R. C. Edgar. MUSCLE: multiple sequence alignment with high accuracy and high throughput. *Nucleic Acids Research*, 32(5):1792–1797, 2004.
- [8] R. C. Edgar and S. Batzoglou. Multiple sequence alignment. *Current Opinion in Structural Biology*, 16(3):368 – 373, 2006.
- [9] R. W. Floyd and R. L. Rivest. Algorithm 489: The algorithm select—for finding the i th smallest of n elements [m1]. *Commun. ACM*, 18(3):173–, 1975.

- [10] O. Gotoh. Significant improvement in accuracy of multiple protein sequence alignments by iterative refinement as assessed by reference to structural alignments. *Journal of Molecular Biology*, 264(4):823 – 838, 1996.
- [11] K. Katoh and D. M. Standley. MAFFT multiple sequence alignment software version 7: Improvements in performance and usability. *Molecular Biology and Evolution*, 30(4):772–780, 2013.
- [12] E. Keogh and A. C. Ratanamahatana. Exact indexing of dynamic time warping. *Knowledge and Information Systems*, 7(3):358–386, 2005.
- [13] J. B. Kruskal. An overview of sequence comparison: Time warps, string edits, and macromolecules. *SIAM Review*, 25(2):201–237, 1983.
- [14] S. Kumar and A. Filipinski. Multiple sequence alignment: In pursuit of homologous DNA positions. *Genome Research*, 17(2):127–135, 2007.
- [15] D. J. Lipman, S. F. Altschul, and J. D. Kececioglu. A tool for multiple sequence alignment. *Proc Natl Acad Sci U S A*, 86(12):4412–5, 1989.
- [16] X. Liu and H.-G. Müller. Functional convex averaging and synchronization for time-warped random curves. *Journal of the American Statistical Association*, 99(467):687–699, 2004.
- [17] S. B. Needleman and C. D. Wunsch. A general method applicable to the search for similarities in the amino acid sequence of two proteins. *J Mol Biol*, 48, 1970.
- [18] C. Notredame. Recent evolutions of multiple sequence alignment algorithms. *PLoS Comput Biol*, 3(8):e123, 2007.
- [19] C. Notredame, D. Higgins, and J. Heringa. T-coffee: a novel method for fast and accurate multiple sequence alignment. *Journal of Molecular Biology*, 302(1):205–217, 2000.
- [20] H. Pages, P. Aboyoun, R. Gentleman, and S. DebRoy. *Biostrings: String objects representing biological sequences, and matching algorithms*. R package version 2.28.0, 2016.
- [21] F. S.-M. Pais, P. d. C. Ruy, G. Oliveira, and R. S. Coimbra. Assessing the efficiency of multiple sequence alignment programs. *Algorithms for Molecular Biology*, 9(1):1–8, 2014.
- [22] 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, 2006.
- [23] T. F. Smith, M. S. Waterman, and W. M. Fitch. Comparative biosequence metrics. *J Mol Evol*, 18, 1981.
- [24] R. Tang and H.-G. Müller. Pairwise curve synchronization for functional data. *Biometrika*, 95(4):875–889, 2008.

- [25] J. D. Thompson, D. G. Higgins, and T. J. Gibson. CLUSTAL W: 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.
- [26] J. D. Thompson, P. Koehl, R. Ripp, and O. Poch. BALiBASE 3.0: Latest developments of the multiple sequence alignment benchmark. *Proteins: Structure, Function, and Bioinformatics*, 61(1):127–136, 2005.
- [27] J. D. Thompson, F. Plewniak, and O. Poch. A comprehensive comparison of multiple sequence alignment programs. *Nucleic Acids Research*, 27(13):2682–2690, 1999.
- [28] J. Thorne, H. Kishino, and J. Felsenstein. An evolutionary model for maximum likelihood alignment of DNA sequences. *Journal of Molecular Evolution*, 33:114–124, 1991.
- [29] I. M. Wallace, G. Blackshields, and D. G. Higgins. Multiple sequence alignments. *Current Opinion in Structural Biology*, 15(3):261 – 266, 2005.
- [30] Z. Yang. *Computational Molecular Evolution*. Oxford Series in Ecology and Evolution. Oxford University Press, 2006.