

Score distributions of gapped multiple sequence alignments down to the low-probability tail

Pascal Fieth* and Alexander K. Hartmann

Institut für Physik, Carl von Ossietzky Universität Oldenburg, D-26111 Oldenburg, Germany

(Dated: November 24, 2015)

Assessing the significance of alignment scores of optimally aligned DNA or amino acid sequences can be achieved via the knowledge of the score distribution of random sequences. But this requires obtaining the distribution in the biologically relevant high-scoring region, where the probabilities are exponentially small. For gapless local alignments of infinitely long sequences this distribution is known analytically to follow a Gumbel distribution. Distributions for gapped local alignments and global alignments of finite lengths can only be obtained numerically. To obtain result for the small-probability region, specific statistical mechanics-based rare-event algorithms can be applied. In previous studies, this was achieved for *pairwise* alignments. They showed that, contrary to results from previous simple sampling studies, strong deviations from the Gumbel distribution occur in case of finite sequence lengths. Here we extend the studies to the for practical applications in Molecular Biology much more relevant case of *multiple* sequence alignments with gaps. We study the distributions of scores over a large range of the support, reaching probabilities as small as 10^{-160} , for global and local (sum-of-pair scores) multiple alignments. We find that even after suitable rescaling, eliminating the sequence-length dependence, the distributions for multiple alignment differ from the pairwise alignment case. Furthermore, we also show that the previously discussed Gaussian correction to the Gumbel distribution needs to be refined, also for the case of pairwise alignments.

I. INTRODUCTION

One important task in bioinformatics is the analysis of nucleotide or amino acid sequences, e.g. found in DNA, RNA, or proteins. A vast amount of sequence data exists and can be found in large data bases like PDB [1] or UniProt[2]. A particular important group of methods, widely used for queries to such data bases, is *sequence alignment* [3, 4]. Naturally appearing DNA or amino acid sequences are aligned in a way that is most likely to resemble their actual evolutionary relationship. Such alignments may contain *gaps*, where in the corresponding evolutionary processes insertion or deletion of genetic material occurred. The degree of similarity, i.e., relatedness, is quantified by a so-called (*optimum*) *alignment score*. Different alignment algorithms exist, usually employing fast heuristics, like BLAST [5].

Since the alignment score is just a natural number, one needs to assess the significance of an alignment. This is achieved typically via calculating the cumulative probability $p(S \geq S_O)$ to find a score S bigger than or equal to the observed score S_O within a suitably defined *null ensemble* of random sequences. This approach corresponds to the calculation of p -values within standard hypothesis tests. In most cases, like for standard database queries, the significance analysis is based on sequences where the letters are drawn identically and independently distributed (i.i.d.) using given probabilities. Therefore it is desirable to find the score distribution of alignment scores for the specified null ensemble. An analytical solution for the score distribution only exists for the limit-

ing case of gapless local alignments with infinite sequence lengths [6], which is not so relevant for biological analysis but of academic interest. According to this theory, the probability function $p(S)$ follows a *Gumbel* (or *extreme-value*) distribution:

$$p_G(S) = \lambda \exp\left(-\lambda(S - S_0) - e^{-\lambda(S - S_0)}\right). \quad (1)$$

For local alignments of sequences of finite length and with allowing for gaps in the alignment the distribution has to be analysed numerically. Simple sampling studies can easily randomly generate and align, e.g., 10^6 sequence pairs in computationally feasible time. This leads to sampling the distribution in the high probability region with lowest statistically reliable probabilities still at $p \approx 10^{-5}$. Early numerical studies using such a simple sampling approach indicated the Gumbel distribution to be a good estimate also for gapped alignments of finitely long sequences [7]. However, biological sequences are in most cases very similar to each other. The relevant alignments therefore generally lie in the high-scoring, extremely low-probability tail of the distribution. Unfortunately, this tail is not covered by the simple sampling approach. To obtain results for the low-probability region, one of us (AKH) applied a statistical mechanics-based rare-event simulation to the problem of pairwise local sequence alignment [8]. This work showed that the Gumbel distribution significantly deviates from the obtained score distribution in the distribution tails. Subsequent studies indicated that a Gaussian correction to this distribution suits well to gain a better description of the score distribution. Interestingly, the strength of this corrections decreases with increasing sequence length [9]. Below we will show that this Gumbel distribution with correction factor leads to a better fitting distribution, but is still

*pascal.fieth@uni-oldenburg.de

only an incomplete description of the data. These methods only have been used so far for pairwise local sequence alignment with and without gaps. It is the main purpose of this work, to extend the application of these methods to multiple local sequence alignment, where more than two sequences are compared within one alignment. Multiple sequence alignments are vastly used for analysing three or more sequences with an assumed evolutionary relation. Resulting alignments are amongst other things used to estimate the phylogenetic history of sequences or to find highly conserved protein domains. Similar protein functions can result from actual evolutionary relation or from convergent evolution, i.e. similar functions developed in independent branches of the phylogenetic tree. Significance analyses could help distinguishing the two cases. Due to the expensiveness of local multiple sequence alignment, its counterpart global multiple sequence alignment is much more common. However, local multiple sequence alignment is especially suited for finding functionally important regions within whole protein families.

Note that for global sequence alignment no analytical solution for the score distribution exists. Studies of real datasets [10] and subsequent studies following essentially a simple sampling approach [11] suggest the three-parameter Gamma distribution as a good model:

$$p_{\text{gamma}}(S) = \begin{cases} \frac{\lambda^\gamma (S-\mu)^{\gamma-1} e^{-\lambda(S-\mu)}}{\Gamma(\gamma)} & S > \mu \\ 0 & S \leq \mu, \end{cases} \quad (2)$$

with the Gamma function $\Gamma(x)$, and parameters λ , γ , and a shift μ . Due to the nature of the studies, the sampled region was again restricted to the high-probability region of the score distribution. As a second application of the rare-event simulation, we studied the score distributions of i.i.d. random sequences in the low-probability region of pairwise and multiple global sequence alignments.

The remainder of this work is organized as follows. First, we will formally define the alignments we studied and state the alignment algorithms we used. Next, we explain the statistical mechanics-based large-deviation approach, which allowed us to sample the alignments distribution of random sequences over a large range of the support. In the main section, we show our results, for gapped multiple alignments of two, three and, in case of global alignment, four sequences. The main results are that again the Gumbel distribution is not sufficient to model the data and that the distribution for multiple alignments, even more relevant for practical applications in Molecular Biology, cannot be obtained from the pairwise alignments results, justifying the present numerically demanding study. Finally, we present a summary and an outlook.

II. SEQUENCE ALIGNMENTS

Sequence alignment algorithms aim to find the optimal scoring alignment of two or more sequences. DNA and amino acid sequences are given by a representing sequence of letters from an alphabet Σ . It is $|\Sigma| = 4$ for DNA sequences consisting only of the four bases. In contrast, there are 20 different amino acids, leading to $|\Sigma| = 20$ for the respective sequences. Let $x^{(i)} = x_1^{(i)}, x_2^{(i)}, \dots, x_{L_i}^{(i)} \in \Sigma^{L_i}$ be the i^{th} sequence of length L_i in a set of N_{seq} sequences. A multiple alignment \mathcal{A} is defined as a set of tuples of indices

$$\mathcal{A} = \{(l_k^{(1)}, l_k^{(2)}, \dots, l_k^{(N_{\text{seq}})})\}; k = 1, 2, \dots, K \quad (3)$$

$$1 \leq l_k^{(i)} < l_{k+1}^{(i)} \leq L_i.$$

A pair of letters $(x_{l_k^{(i)}}^{(i)}, x_{l_k^{(j)}}^{(j)})$ is called a *match* if $x_{l_k^{(i)}}^{(i)} = x_{l_k^{(j)}}^{(j)}$. Otherwise it is called a *mismatch*. Note that in each tuple N_{seq} letters are joined, so some pairs in this tuple may match while other pairs may form a mismatch. If $l_{k+1}^{(i)} = l_k^{(i)} + 1$ and $l_{k+1}^{(j)} = l_k^{(j)} + 1 + M$ with $M > 0$, i.e., the indices $l_k^{(j)} + 1, \dots, l_k^{(j)} + M$ do not appear in any tuple, sequence $x^{(i)}$ is said to have a *gap* of length M in respect to sequence $x^{(j)}$. Also if $l_1^{(k)} = 1 + M > 1$ or $l_K^{(k)} = L_i - M < L_i$ one speaks of a gap of the k^{th} sequence.

To find the alignment most likely resembling the actual evolutionary relationship, an objective *score* function $S(\mathcal{A}, \{x^{(i)}\})$ is used. Usually, the score is based on sums of pairwise scores $s(x_{l_k^{(i)}}^{(i)}, x_{l_k^{(j)}}^{(j)})$, which are taken from so-called *substitution matrices* given by biologists. Typically, the scores are proportional to log of mutation (mismatch) or conservation (match) probabilities in evolutionary models. Therefore, scores for matches are positive, while scores for mismatches are smaller, usually negative. Two widely used sets of substitution matrices are the PAM [12] and BLOSUM [13] substitution matrix sets. Gaps are penalised depending on their lengths with some length-dependent function $g(M)$. The gap penalties are intended to model log of probabilities of the occurrence of insertion or deletion events in evolutionary processes.

To wrap this up, the score of a *global* pairwise alignment is the sum over all matches, mismatches, and the penalties of all gaps g . A multiple alignment is scored by the sum-of-pairs score, i.e., its score is the sum over all pairwise scores:

$$S(\mathcal{A}, \{x^{(i)}\}) = \sum_{i < j} \left(\sum_k s(x_{l_k^{(i)}}^{(i)}, x_{l_k^{(j)}}^{(j)}) - \sum_g g(M_g) \right), \quad (4)$$

where the sum \sum_g runs over all gaps.

The aim of the alignment comparison is that the more similar two sequences are, the higher the resulting score.

Nevertheless, the score does not depend only on the sequences but also on the alignment. For example, even for two equal sequences, one can find alignments with a very small score. Thus, one is seeking for the *optimal alignment* \mathcal{A}_O which is the alignment with the maximum score

$$S_{\max}(\{\vec{x}^{(i)}\}) = \max_{\mathcal{A}} S(\mathcal{A}, \{\vec{x}^{(i)}\}),$$

maximized over all possible alignments. In subsequent sections S is used synonymously with $S_{\max}(\{\vec{x}^{(i)}\})$. In this work, affine gap costs

$$g(M) = \alpha + \beta(M - 1) \quad (5)$$

were used. Thus, a high penalty α can be given for opening a gap and a smaller penalty β for extending one. Although even for two sequences there are exponentially many alignments, an optimal pairwise global alignment, i.e., $N_{\text{seq}} = 2$, with affine gap costs can easily be found in polynomial time by the algorithm by Needleman and Wunsch[14] with Gotoh's extension[15]. For multiple global alignments the *progressive* heuristics by Feng and Doolittle [16] was used in this work: All possible sequence pairs are aligned with pairwise alignment first. A guide tree is constructed using the different obtained pairwise scores. Sub-alignments are then aligned to each other by aligning the highest-scoring sequence pair, taking into account the existing alignments.

Another important alignment method next to global sequence alignment is *local* sequence alignment. Here, from each sequence a subsequence is chosen and only the subsequences are aligned. This means for each sequence start- and endpoints $l_s^{(i)}, l_e^{(i)}$ of the subsequences are subject to optimization as well. Thus, the optimal local alignment has to maximise the score over all possible alignments of all possible subsequences. This is also equivalent to not penalising gaps at the beginning and the end of sequences. With the algorithm by Smith and Waterman[17] this is possible also in polynomial time $\mathcal{O}(L^{N_{\text{seq}}})$ for a set of N_{seq} sequences, if each has the length $L_i = L$.

III. METHOD

In the approach used [8], a sequence pair, correspondingly here a set of K sequences, is viewed as a “state” \mathcal{C} of a physical system with “energy” $E = -S$. A “temperature” T is introduced and the states are sampled according to the rules of the canonical ensemble in statistical mechanics. Specifically, a Markov chain $\mathcal{C}_0 \rightarrow \mathcal{C}_1 \rightarrow \dots$ is generated. In each step t a trial state \mathcal{C}' is generated from the current state \mathcal{C}_t by randomly choosing and replacing one residue in the sequence set [25]. The alignment score $S(\mathcal{C}')$ is calculated and the trial state accepted, i.e., $\mathcal{C}_{t+1} = \mathcal{C}'$, with the Metropolis probability [18] $P(\mathcal{C}_t \rightarrow \mathcal{C}') = \min[1, \exp(\Delta S/T)]$

with $\Delta S = S(\mathcal{C}') - S(\mathcal{C}_t)$. If not accepted, the current configuration is kept, i.e., $\mathcal{C}_{t+1} = \mathcal{C}_t$. The equilibrium distribution of the sampled states is known to be $Q(\mathcal{C}) = P(\mathcal{C}) \exp(S(\mathcal{C})/T) / Z_T$ with the partition function $Z_T = \sum_{\mathcal{C}} P(\mathcal{C}) \exp(S(\mathcal{C})/T)$. The score distribution, biased by the scale parameter (or “temperature”) T is then:

$$\begin{aligned} p_T(S) &= \sum_{\{\mathcal{C}|S(\mathcal{C})=S\}} Q(\mathcal{C}) \\ &= \sum_{\{\mathcal{C}|S(\mathcal{C})=S\}} \frac{\exp(S/T)}{Z_T} P(\mathcal{C}) \\ &= \frac{\exp(S/T)}{Z_T} p(S) \end{aligned}$$

The unbiased distribution then is $p(S) = p_T(S) Z_T \exp(-S/T)$. After rescaling via $p_T^R = p_T \exp(-S/T)$ only the correct values for the partition functions Z_T remain unknown. Determining them is possible after covering different probability regions, starting with the one for $T = \infty$. To cover the entire score range, simulations are done at different temperatures. The efficiency of simulations is then improved further by using the parallel tempering technique [19–21], which works as follows. Simulations of the system are done for N_T different temperatures $T_1 < T_2 < \dots < T_{N_T}$, i.e., an independent configuration \mathcal{C}_i is simulated for each temperature T_i , $i = 1, \dots, N_T$. The configurations at neighbouring temperatures (T_i, T_{i+1}) are swapped in suitably chosen simulation time intervals with a swap probability $P_{\text{sw}}(\mathcal{C}_i \leftrightarrow \mathcal{C}_{i+1}) = \min(1, \exp[-\Delta\beta\Delta S])$ with $\Delta\beta = 1/T_{i+1} - 1/T_i$ and $\Delta S = S(\mathcal{C}_{i+1}) - S(\mathcal{C}_i)$. One time step t consists here of one Monte Carlo sweep, i.e., a number of Metropolis steps equal to the total number of residues in a configuration, and a subsequent sweep of $N_T - 1$ exchange attempts of randomly selected temperature pairs. To sample only weakly correlated data, the autocorrelation $c_S(t) = \{\langle S(t_0)S(t_0+t) \rangle - \langle S \rangle^2\} / \{\langle S^2 \rangle - \langle S \rangle^2\}$ is calculated and a relaxation time τ determined for which $c_S(\tau) = 1/e$. Only every τ^{th} value is considered when sampling the data. Equilibrium is ensured by starting to sample after an equilibration time t_0 . To determine this time, two sets of initial conditions are simulated. One set of simulations is initialised with randomly generated sequence sets, i.e., with low scoring sequence sets. The other set of simulations is initialised with sets of identical sequences, only consisting of the letter a with the highest pairwise score $s(a, a)$ [26], i.e., with maximum scoring sequence sets. The equilibration time t_0 is reached when the average scores of the different regions have reached the same value within error bars. Compare Fig. 1 which shows the first 10,000 time steps of the simulations for the simulations done for $N_{\text{seq}} = 3$ sequences of lengths $L = 40$ for different temperatures T .

Having covered the whole distribution range, it is now possible to determine the partition functions Z_T . As-




FIG. 1: Equilibration times for different temperatures: First 10,000 Monte Carlo sweeps in the parallel tempering time series averaged over 30 different realisations for random (bottom branch) and high scoring (top branch) initial conditions, respectively. System: $N_{\text{seq}} = 3$ sequences of length $L = 40$, local alignment. A selection of the used temperatures is shown. At lower temperatures it takes more time for equilibration.

suming for $T = \infty$ the biased distribution approximates the true distribution $P_T(S) \approx P(S)$ the partition function, or normalisation constant, becomes $Z_\infty \approx 1$. After rescaling, the other distributions differ from the unbiased distributions only by Z_T which can be determined by a shift between neighboring functions $T_{i-1,i}$ on the logarithmic scale, $\log Z_i$. The shift between $T_1 = \infty$ and the unbiased distribution is defined as $\log Z_1 = 0$. The consecutive shifts are determined by minimising the function

$$f(\log Z_i) = \sum_{S \in [S_s, S_e]} \{ (\log p_{T_i}^R(S) + \log Z_i) - (\log p_{T_{i-1}}^R(S) + \log Z_{i-1}) \}^2. \quad (6)$$

in the overlapping score region $[S_s, S_e]$ of the two neighboring distributions. The shifted distributions are then $p_T^S = p_T^R \cdot Z_T$. For most score ranges, several values from the different biased distributions are available. For each value S of the score distribution a weighted average over all available data is calculated:

$$p(S) = \frac{1}{\sum_T w_T(S)} \sum_T w_T(S) p_T^S. \quad (7)$$

w_T is the inverse relative error of the corresponding distribution $p_T(S)$ before rescaling and shifting:

$$w_T(S) = \sqrt{\frac{n_T}{p_T(S)(1 - p_T(S))}} \quad (8)$$

where n_T is the number of used samples for temperature T . This finally gives values for a wide range of the distribution $p(S)$ down to regions of very low probabilities.

IV. RESULTS

The rare-event simulations were performed for local and global alignments. If not mentioned otherwise, the alignments were calculated using the BLOSUM62 substitution matrices and gap penalties according to (5) with $(\alpha = 12, \beta = 1)$. These values were used in the previous large-deviation studies [8, 9] and therefore allow a direct comparison between results for pairwise and multiple sequence alignment.

A. Local multiple sequence alignments

Rare-event simulations for local pairwise sequence alignments have shown that the Gumbel distribution alone is not a good description of the score distribution of gapped local sequence alignments. The introduction of a Gaussian correction to the Gumbel distribution (1) improved the agreement between analytical distribution and data. The corrected distribution is

$$p_C = p_G \cdot \exp[-\lambda_2(S - S_0)^2] = \lambda \exp\left(-\lambda(S - S_0) - \lambda_2(S - S_0)^2 - e^{-\lambda(S - S_0)}\right). \quad (9)$$

The strength of the Gaussian correction is indicated by the fit parameter λ_2 and has been shown to decrease with increasing sequence length for pairwise alignments. λ_2 was observed to decrease with a power-law for small gap costs and faster than a power-law for high gap costs. In the case analysed here, $(\alpha = 12, \beta = 1)$, the decrease is expected to be just in the power-law region. Subsequently the distributions obtained with the rare-event simulations for multiple local sequence alignment of $N_{\text{seq}} = 3$ sequences will be presented and later on compared to the results for pairwise alignment.

1. Score distributions for local multiple sequence alignments

Simulations for local multiple sequence alignments of $N_{\text{seq}} = 3$ sequences of length $L = 40$ were performed for 60 different realisations of the driving randomness, each over $5 \cdot 10^4$ sweeps. The score distribution was obtained as described in sec. III. A fit of the Gumbel distribution without (1) and with (9) Gaussian correction was performed. The data and the fits are shown in Fig. 2. The deviation of the Gumbel distribution from the data in the tail of the distribution is clearly visible. The better performance of the fit of the Gumbel distribution with Gaussian correction is also indicated by the χ^2 value per degree of freedom, $\chi^2/\text{ndf} = 1.5$ in contrast to the value for the distribution without Gaussian correction, $\chi^2/\text{ndf} = 343.1$. Thus, the behavior for the low-probability tail for pairwise alignment could be confirmed for the alignment of $N_{\text{seq}} = 3$ sequences. We performed another fit where the Gumbel distribution




FIG. 2: The distribution for local alignments of $N_{\text{seq}} = 3$ sequences of length $L = 40$ as obtained by the rare-event simulation. Fits of the Gumbel distribution to the whole distribution as well as constrained to the score range $[S_l = 21, S_u = 115]$, i.e., $P(S) \leq 10^{-10}$, are shown. Also shown is the better suited fit of the Gumbel distribution with a Gaussian correction. The inset shows the high probability region.

without Gaussian correction was restricted to the distribution range with $P(S) \geq 10^{-10}$, which corresponds to the interval $[S_l = 21, S_u = 115]$. Simple sampling would only cover probabilities $P(S) > N_{\text{samples}}^{-1}$ by creating and aligning N_{samples} i.i.d. random sequence sets. Therefore, the restricted distribution range is still generously large, requiring more than $N_{\text{samples}} > 10^{10}$ samples if it were to be obtained by simple sampling. The restricted fit agrees with the data in the high-probability region, visible in the inset. This is also indicated by a somehow better χ^2 -value of $\chi^2/\text{ndf} = 63$. Of course the fit now obviously deviates strongly in the distribution tail. Indeed, the Gumbel distribution with correction constrained to the same score range still performs better than without correction ($\chi^2/\text{ndf} = 9.45$). Over the complete obtained distribution range this results in a significant overestimation of λ_2 and a strong underestimation of $P(S)$ in the tail of the distribution. The results of these fits show that the behaviour of the distribution can not be fully studied with just a simple sampling approach.

Figure 3 shows the fit of the Gumbel distribution with Gaussian correction to the data for local alignment of $N_{\text{seq}} = 3$ sequences for different sequence lengths. The fit performed as well as for the example case of $L = 40$. However, varying the range of the distribution on which the fit is performed results in a change of parameter values not accounted for by the standard error calculated during the fit. Therefore parameter fits for every system were performed with varying window sizes. All windows start at the minimum score S_l for which the simulations

L	λ	$10^4 \lambda_2$	S_0	$\frac{\chi^2}{\text{ndf}}$	S_l [S_l/S_{max}]
40	0.25570(7)	0.968(4)	32.163(4)	8.2	200[0.152]
50	0.25378(5)	0.7441(15)	35.568(4)	7.3	225[0.14]
60	0.25164(6)	0.6045(17)	37.9981(13)	6.7	225[0.114]
80	0.247685(25)	0.4889(4)	41.744(7)	3.7	750[0.284]

TABLE I: Parameters for the fit of the Gumbel distribution with Gaussian correction to the obtained data for local multiple sequence alignments ($N_{\text{seq}} = 3$). Parameter values are convergence values: A fit was performed for each parameter value as a function of the window size of the distribution. χ^2 values were calculated for these convergence values of the parameters in the Gumbel distribution in respect to the obtained distribution data.

yielded a data point (i.e., this point is fixed for each system) and ends at a variablescore $S_u \leq S_{\text{max}}$, where S_{max} is the maximum possible score, i.e., for two equal sequences with the highest scoring letter. For the analysed system sizes for multiple local sequence alignment with $N_{\text{seq}} = 3$ the parameters seemed to converge. Figure 4 exemplarily shows the parameter curve for $\lambda(S_u)$. For sequence lengths up to $L = 60$ it seems reasonable to estimate the parameters by fitting an exponential function with a constant. For a parameter g we use $g(S_u) = g_b + C \exp(-k \cdot S_u)$ with fit parameters g_b , C , and k for an appropriate part of the value curve obtained. The new fit parameter g_b approximates the value against which parameter $g(S_u)$ converges. For $L = 80$ this approach already appears less promising. Restricting the window for the fit is necessary to gain a feasible value for the parameters. Parameter values were obtained in the way described and are shown in Tab. I, including the range of windows chosen, but should be taken with the according caution. We will see that this lack of confidence is not the result of expanding the analyses to multiple sequence alignment. The same analysis of the distributions found for pairwise alignments rather shows that the fit performs even worse there. The analytical solution [6] does assume infinitely long sequences. Generally a decrease of the parameter λ_2 , indicating the strength of the Gaussian correction, can be observed, see Tab. I. The decrease of λ_2 with increasing sequence length indicates that the correction is partially due to a finite size effect, disappearing for long sequences. This decrease has already been found in the study of pairwise alignment distributions [9] and can be confirmed here for multiple alignments. As even the corrected function and the acquisition of its parameters is still quite makeshift, as visible from the χ^2/nfd values shown in tab. I, beginning considerably larger than one, we did not conduct a more quantitative analysis of the decrease of λ_2 .




FIG. 3: The score distributions for local multiple alignments for several sequence lengths and the fitted Gumbel distributions with Gaussian correction. The inset shows the high-probability region.




FIG. 4: The fit parameters found for different window sizes (S_{\min}, S_{end}). Several sequence lengths L for $N_{\text{seq}} = 3$ are shown.

2. Comparison to pairwise alignment

As the simulations for multiple sequence alignment were computationally expensive especially in the case of local alignments, only short sequences up to $L = 80$ could be used for score distribution analysis in the present work. For a better comparison with the data for pairwise alignment, where mostly longer sequences were studied [9] we have performed additional simulations for shorter sequences also for pairwise alignments. As we implemented a parallel version of the algorithm for the MSA simulations, we could also obtain the dis-




FIG. 5: The fit parameters found for different window sizes (S_{\min}, S_{end}). Several sequence lengths L for $N_{\text{seq}} = 2$ are shown.

tribution of scores over a larger range of the support, compared to the previous work. We fitted the data to the Gaussian-corrected Gumbel distribution. Obtained values are shown in Tab. I. Restricting the fit to the range of the support, which was addressed in the previous work where known, reproduces the parameter values found before. But when we extended the distribution range for the fit further, it yielded different parameter values. Figure 5 shows the fit value of λ for $N_{\text{seq}} = 2$ for different window sizes. This resulted in a small but visible linear decrease of the parameter value, showing that in fact even the corrected Gumbel distribution is not sufficient to describe the obtained data over a large range of the support. The lower boundary of the window, S_l , was varied as well, eliminating small S values. The linear decrease of parameters could be observed for varying S_l while fixing $S_u = S_{\text{max}}$ as well as for varying S_u again, but with different, higher, boundaries S_l . Furthermore, χ^2 values were obtained for actually performing the fit to the whole support obtained and are shown together with parameter values in Tab. II. They also indicate that the fit of the corrected Gumbel distribution does not perform as well for pairwise local sequence alignment as for multiple local sequence alignment.

Thus, instead of comparing parameters of fitted distributions, we rather compare the data itself. This requires some scaling to get rid of sequence-length dependencies. The results in Ref. [22] indicate that for lengths (L, L') of a sequence pair, by rescaling the S -axis with S_{max} and by rescaling the log probabilities by $\log p(S_{\text{max}}|L, L')$ the curves for different pairs of sequence lengths fall approximately on top of each other. This means

$$\frac{\log p(S|L, L')}{\log p(S_{\text{max}}|L, L')} \approx f(S/S_{\text{max}}) \quad (10)$$

L	λ	$10^4 \lambda_2$	S_0	χ^2_{ndf}
30	0.3243(4)	12.2506(13)	12.755(29)	21.2
40	0.31857(18)	9.175(7)	15.096(19)	26.5
50	0.30084(2)	7.597(5)	15.50(3)	34.5
60	0.29452(12)	6.3636(27)	16.675(23)	40.9
80	0.2824(12)	4.8321(17)	17.42(4)	43.4

TABLE II: Parameters for the fit of the Gumbel distribution with Gaussian correction to the obtained data for local pairwise sequence alignments ($N_{\text{seq}} = 2$). As no convergence of parameter values were observed, all values, including χ^2 , are given as obtained by fitting over the whole range of the support for which data has been obtained.

where $f(\cdot)$ is a universal function of the relative score $Q = S/S_{\text{max}}$. Thus, for sequence lengths \hat{L}, \hat{L}' and a different score \hat{S} but with the same relative score Q , i.e., $\hat{S} = S/S_{\text{max}} \hat{S}_{\text{max}}$ one obtains

$$\log p(S|L, L') \approx \log p\left(\frac{S\hat{S}_{\text{max}}}{S_{\text{max}}}\bigg|\hat{L}, \hat{L}'\right) \frac{\log p(S_{\text{max}}|L, L')}{\log p(\hat{S}_{\text{max}}|\hat{L}, \hat{L}')} \quad (11)$$

Figure 6 shows selected distributions for $N_{\text{seq}} = 2$ and $N_{\text{seq}} = 3$ with rescaled score- and probability axes. The distributions for pairwise alignments coincide in the low probability region as well as the distributions for multiple alignments. There is only a deviation in the small region of high-probabilities which disappears with increasing sequences length. There is however an overall difference between pairwise and multiple alignments. The distributions do not coincide and the curvature is significantly stronger for $N_{\text{seq}} = 3$. Thus, the deviation from the Gumbel distributions is stronger for multiple alignments and distributions of scores for multiple alignments can not be estimated easily from the data obtained for pairwise alignment distributions. Nevertheless, the knowledge of one distribution of arbitrary length L makes it possible to estimate distributions for other sequence lengths L' also for multiple alignment of $N_{\text{seq}} = 3$ sequences. The same can be expected for $N_{\text{seq}} > 3$. But a more precise analysis would require more simulation work for multiple sequence alignments with more sequences. These analyses could potentially yield a method to estimate the distributions for $N_{\text{seq}} > 3$ by readily obtained distributions with less sequences. Nevertheless, for an actual study the numerical demand for the alignment of even more sequences appears to be too high. Here, the implementation of a good working heuristics for multiple local sequence alignments with gaps would be helpful.

B. Global multiple sequence alignments

We could adopt the large-deviation approach easily to global pairwise and multiple sequence alignment with and without gaps. We below first present the analysis of




FIG. 6: The distributions for various L and N_{seq} with rescaled scores and probabilities. The distributions coincide for low probabilities and identical number of sequences N_{seq} but differ among pairwise and multiple alignments. Inset: High probability region in which the distributions deviate from each other (even for same N_{seq}).

the high-probability tail for pairwise sequence alignment and compare it to previous simple sampling results. Furthermore, we expanded the method to multiple sequence alignment and compare to the results for pairwise sequence alignment.

1. Pairwise global sequence alignments

We are not aware of any previous large-deviation study for multiple or for pairwise global alignment. Here we first show our results for pairwise global alignment. For a direct comparison to the results of Pang et al. [11], simulations were performed for alignments with identical gap costs of ($\alpha = 7, \beta = 1$), see Fig. 7. First we performed a fit of the Gamma distribution (2), which was found in Ref. [11] to fit the data well, to the high-probability region $P(S) > 10^{-3}$ of the obtained distributions. For pairwise alignment of sequences of lengths $L = 50, 100, 200$ we obtained parameters similar to the sampling results [11], as shown in tab. III. As visible in Fig. 7, this fit does not match the data in the tail of the distribution. Extending the fit to regions with lower probabilities yielded significantly different parameter values. Nevertheless, this leads to a strong deviation in the high-probability region. For a more elaborate study one could also obtain data for the low-scoring side of the distributions (by means of negative “temperatures” in the large deviation simulations). Nevertheless, the results indicate that the assumption of scores distributed according to the Gamma distribution is somewhat flawed. Thus, we also applied a Gaussian

L	Pang[11]	high $P(S)$	
50	γ	41.00	33(6)
	λ	0.63	0.60(5)
	μ	-84.60	-66(5)
	$\frac{\chi^2}{\text{ndf}}$		1.43
100	γ	49.16	48(6)
	λ	0.55	0.57(4)
	μ	-115.44	-95(5)
	$\frac{\chi^2}{\text{ndf}}$		1.05
200	γ	52.24	59(12)
	λ	0.44	0.50(5)
	μ	-153.25	-123(12)
	$\frac{\chi^2}{\text{ndf}}$		0.74

TABLE III: Parameters for the fit of the Gamma distribution to the obtained data for global multiple sequence alignments as found in [11] and as found with the large-deviation simulations with the fit restricted to $P(S) > 10^{-3}$.

correction

$$p_{\text{gc}}(S) = \begin{cases} \frac{\lambda^\gamma (S-\mu)^{\gamma-1} e^{-\lambda(S-\mu)}}{\Gamma(\gamma)} e^{\lambda_2 S^2} & S > \mu \\ 0 & S \leq \mu, \end{cases} \quad (12)$$

with the Gamma distribution (2) and the Gaussian correction with the parameter λ_2 indicating its strength. Fig. 7 shows the fits of the different distributions. Clearly, the Gamma distribution with the Gaussian correction covers the distribution best over the whole probability range.

2. Multiple global sequence alignments

The results for the multiple global alignment of $N_{\text{seq}} = 3$ sequences of length $L = 40$ are shown in Fig. 8. Included are the distributions obtained with the progressive heuristics as well as with the computationally much more expensive exact algorithm. The significant difference in the low-scoring region of the distributions is clearly visible. In the high-scoring region the heuristics seems to work better, i.e., approaches the exact optimum alignments, which is visible in the better consensus between the two probability distributions. Shown are also fits to the (uncorrected) Gamma distribution. The distribution fits only well for the distribution obtained with the heuristics, but not for the distribution obtained with the exact algorithm, i.e. the Gaussian correction does not seem to be necessary for multiple alignments of $N_{\text{seq}} > 2$ sequences *when using the heuristics*. For $N_{\text{seq}} = 2$ the correction was found to be necessary in any case. But for pairwise alignment, there is only one possible pair of sequences to be aligned, which is done by the dynamic programming algorithm. The true optimal alignment is found. In contrast, when using the heuristics, for $N_{\text{seq}} > 2$ sequences the first sequence pair is aligned




FIG. 7: The alignment score distribution as obtained with the rare-event simulations for the global alignment of $N_{\text{seq}} = 2$ sequences of length $L = 100$, here with gap costs ($\alpha = 7, \beta = -1$) for direct comparison with [11]. Shown are the fits of the Gamma distribution with and without Gaussian correction as well as the fit of the Gamma distribution to the high-probability region ($P(S) > 10^{-3}$, $S = [R_{\text{min}} = -40 : R_{\text{max}} = 23]$) only. Inset: High probability-region.

with the dynamic programming algorithm. The following sequences are aligned to this initial alignment. If the sequences are very dissimilar, the initial alignment does not reflect the structure of the multiple alignment very well. This means, especially in the low-scoring region scores lower than the true optimal scores are found with progressive alignment. The different behaviour in the high-probability region between pairwise and multiple alignments is one possible explanation for the difference in parameter behaviour. Generally, the exact algorithm is computationally too expensive for multiple sequence alignment and the heuristics is used. It should be noted that for the score statistics this use of the heuristics renders the Gaussian correction unnecessary. Further observations for global multiple sequence alignment were made for the case of using the heuristics if not stated otherwise.

3. Comparison of pairwise and multiple sequence alignment

Figure 9 shows the values of fit parameter λ_2 for different sequence lengths and different number of sequences. The most obvious observation to be made is that for $N_{\text{seq}} = 2$ the λ_2 values are positive, i.e., correcting the curvature of the distribution to lower probabilities, converging strongly towards zero. For multiple alignments of $N_{\text{seq}} > 2$, however, λ_2 values are relatively close to 0.




FIG. 8: Distributions obtained for the global alignment of $N_{\text{seq}} = 3$ sequences of length $L = 40$ with the exact algorithm and the heuristics. The lines show the fit of the *uncorrected* Gamma function against the data. The inset shows the high-probability region, in which both distributions differ significantly from each other and the fit for exact alignments deviates significantly from the data.

This is compatible with the fact that a Gamma distribution without Gaussian correction seems to fit the data well when using the heuristics. Thus, when obtaining score distributions for significance analysis of sequence alignments, one actually needs to simulate multiple alignments explicitly, including large-deviations techniques. It is not possible to deduce the distributions from the (large-deviation) results of pairwise alignment.

The behavior of the parameter γ as function of sequence length is shown in Fig. 10 for the three cases $N_{\text{seq}} = 2, 3$ and 4. While the correction parameter λ_2 decreases with sequence length or is basically zero anyway, the parameter γ increases. For $\gamma \rightarrow \infty$ the Gamma distributions converges to the normal distribution. Thus, for long enough sequences, the distribution of global alignment scores converges to a normal distribution. This can be understood in the following way: Neglecting gaps, which is in particular true for very large scores, the pairwise scores can approximately assumed to be i.i.d. random numbers. The overall sequence score is a sum over these values. For $L \rightarrow \infty$ the central limit theorem holds, resulting in a Gaussian distribution.

Same as for local sequence alignments, the probability distributions can be rescaled according to (11). The rescaled distributions are found in Fig. 11. Again, for identical number of sequences, the distributions coincide again for different sequence lengths L in the low-probability region. Again, there is a significant difference between the distributions for different number of sequences N_{seq} , indicating that indeed significance analysis for multiple sequence alignments cannot be based on




FIG. 9: The parameter λ_2 as found for the distributions of global alignment for different sequence lengths L and different number of sequences N_{seq} . Drawn lines are guides for the eye.

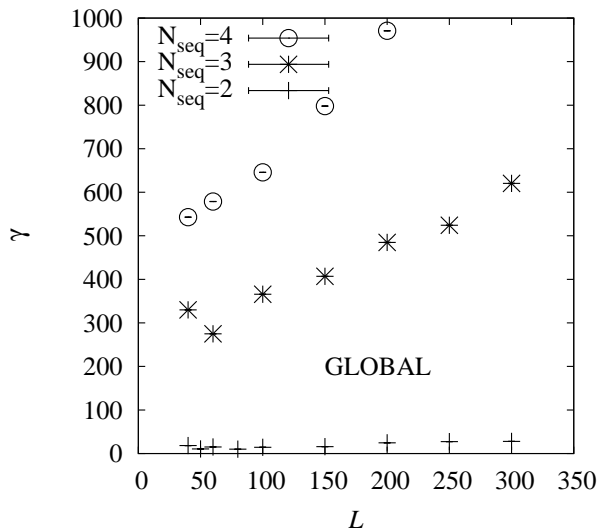


FIG. 10: The parameter γ as found for the distributions for global alignment of different sequence lengths L and different number of sequences N_{seq} .

the results for pairwise alignments, i.e., dedicated multiple sequence alignment studies have to be performed explicitly, as done in this work. Note that the initial deviation for higher probabilities seems more distinct than for local alignments, but it should be kept in mind that the sequence lengths shown for global alignments differ more than for local alignments, because more data was available.

The strong difference of the resulting distribution of pairwise and multiple alignments could be influenced by the use of a heuristics for multiple alignments. For small




FIG. 11: Distributions rescaled with (11) for different number of sequences N_{seq} and different sequence lengths L . The rescaled distributions coincide for low-probabilities and identical N_{seq} but differing L .

sequence length, we were able to compare the results of both algorithms. Figure 12 shows the rescaled distribution obtained for $L = 40$ for pairwise alignment and in case of multiple sequence alignment with $N_{\text{seq}} = 3$ the rescaled distributions obtained with the heuristics as well as with the exact algorithm. The distributions obtained from heuristics and exact algorithm, respectively, differ only in a small region, while the results from pairwise and multiple alignment differ everywhere. Thus, the different development of distributions does not seem to be a result of the application of a heuristics.

C. Comparison of local and global alignments

Finally, the score distributions of local and global sequence alignment can be compared. For mostly dissimilar sequences the selection of best-scoring subsequences can increase the alignment score. Since the probability of the sampling of a specific sequence set is independent of the alignment algorithm, dissimilar sequence sets would score higher in local sequence alignment, where no negative scores S are possible. This means that the low-scoring region of the distribution for local alignments should be skewed towards higher alignment scores compared to the distribution of global alignments. This can be observed in a comparison of score distributions with the different alignment types for the same sequence set properties, as shown in Fig. 13. Here, as an example, the obtained distributions for $N_{\text{seq}} = 3$ sequences of length $L = 60$ are shown. As can be seen, the maximum of the distribution for global alignments is found for negative scores, the maximum for local alignments is positive (as




FIG. 12: Distributions obtained for global alignment of $N_{\text{seq}} = 2$ and $N_{\text{seq}} = 3$ sequences of length $L = 40$ (with the exact algorithm as well as with the heuristics for $N_{\text{seq}} = 3$). The $\log(P)$ values are rescaled with (11) and the scores with S/S_{max} . The inherent difference between $N_{\text{seq}} = 2$ and $N_{\text{seq}} = 3$ is obviously not based the use of an heuristics. The inset show the high-probability region, in which all three obtained distributions differ. In case of $N_{\text{seq}} = 3$ this difference is small compared to the exact algorithm and is due to the bad performance of the heuristics in the low-scoring region.

local alignments can only have scores $S > 0$).

On the other hand, most high-scoring global alignments, which have only few negative scoring residue-pairings or gaps, would not yield better scores by selecting (smaller) subsequences. This means that in the high-scoring region, many sequence sets yield the same score in local and global alignment. Therefore the distributions of local and global sequence alignment can be expected to agree for higher scores. This can exactly be seen in Fig. 13. This agreement might allow to estimate statistics of computationally expensive local alignments in the high-scoring region by means of the cheaper global alignment. However, for practical applications in Molecular Biology, local alignment results typically in scores which are located in the intermediate region, where it differs from global alignments. Thus, knowing the distribution of random-sequence scores for global alignment alone would not be sufficient to estimate the local-alignment score distribution here.

V. SUMMARY AND OUTLOOK

We have studied the distribution of scores of multiple sequence alignments in the ensemble of i.i.d. random protein sequences. The statistical mechanics-based large deviation simulations used here allowed us to numerically measure the distribution of alignment scores in the biologically most relevant region of small probabilities, e.g.,




FIG. 13: Obtained distributions for local as well as global alignment of $N_{\text{seq}} = 3$ sequences of length $L = 60$. The inset shows the high-scoring region, where both distributions approach each other.

$\sim 10^{-100}$. Data for these low-probability regions was not available before.

Analysis of the distribution for multiple *local* sequence alignment showed that the Gumbel distribution alone does not describe the data as found previously for pairwise alignments [9]. Further analysis showed that the suggested Gaussian correction improves the behaviour, but still is not sufficient to describe the whole distribution. Nevertheless the results indicate that for a known distribution for a certain number of sequences N_{seq} and fixed sequence lengths L , distributions for other sequence lengths but the same N_{seq} can be found by rescaling.

Also for pairwise *global* sequence alignment, were previously only simple-sampling results were available, we could obtain the distribution of scores over a large range of the support. We could reproduce the previous simple-sampling results [11] only by restricting the fits to the high-probability region. Extending the analysis to the low-probability region showed significant deviations of the Gamma distribution fit from the data. Again, a Gaussian correction improves the fit of the distribution. The comparison between multiple and pairwise global alignments revealed different characteristics of the fit parameters. While the parameter λ_2 indicating the

strength of the Gaussian correction was found to be positive but decreasing with sequence length for pairwise alignments, it was almost negligible (actually slightly negative) for multiple alignments, also converging towards zero. The application of fast heuristics seems to be the reason for this difference in behaviour, because the behavior for multiple alignment was comparable to the behavior of pairwise alignment when using the progressive heuristics. However, in the biologically relevant low-probability region the heuristics performs well and for small sequence length the distributions found with the exact algorithm and the heuristics do not differ.

It is interesting to note the behavior in the limit of infinitely long sequences, which is for biological applications not so relevant but of fundamental interest, in particular with respect to the analytical results. Similar to previous studies for pairwise alignments, in this limit the distribution of scores are compatible with standard distributions: For long sequences the score distribution of local alignments appears to converge to the Gumbel distribution, the score distribution of global alignments to a normal distribution. Since in many actual applications, multiple alignments with more than ten sequences are not uncommon, simulation of sequence sets with a higher number of sequences ($N_{\text{seq}} > 4$) appears sensible in future studies. Nevertheless, due to the computational complexity of the exact algorithm which depends strongly on N_{seq} , this requires the design and implementation of faster algorithms, e.g., in particular efficient heuristics for multiple local sequence alignment.

Also it would be desirable, to extend these alignment simulations to more refined ensembles of random sequences, e.g., for biologically more specific ensembles like transmembrane proteins [23], or sequences with correlations.

Acknowledgments

PF acknowledges financial support from the German Science Foundation (DPG) within the Graduiertenkolleg GRK 1885. The simulations were performed at the HPC Cluster HERO, located at the University of Oldenburg (Germany) and funded by the DFG through its Major Research Instrumentation Programme (INST 184/108-1 FUGG) and the Ministry of Science and Culture (MWK) of the Lower Saxony State.

[1] H. M. Berman, J. Westbrook, Z. Feng, G. Gilliland, T. N. Bhat, H. Weissig, I. N. Shindyalov, and P. E. Bourne, Nucl. Acids Res. **28**, 235 (2000), ISSN 0305-1048, 1362-4962, URL <http://nar.oxfordjournals.org/content/28/1/235>.

[2] The UniProt Consortium, Nucl. Acids Res. **43**, D204 (2015), ISSN 0305-1048, 1362-4962, URL <http://nar.oxfordjournals.org/content/43/D1/D204>.

[3] R. Durbin, S. Eddy, A. Krogh, and G. Mitchison, *Biological Sequence Analysis: Probabilistic Models of Proteins and Nucleic Acids* (Cambridge University Press, 1998),

- ISBN 978-0-521-62971-3.
- [4] P. Clote and R. Backofen, *Computational molecular biology: an introduction*, Wiley series in mathematical and computational biology (Wiley, New York, 2000), ISBN 0-471-87251-2.
- [5] S. F. Altschul, W. Gish, W. Miller, E. W. Myers, and D. J. Lipman, *Journal of Molecular Biology* **215**, 403 (1990), ISSN 0022-2836, URL <http://www.sciencedirect.com/science/article/pii/S0022283605803602>.
- [6] S. Karlin and S. F. Altschul, *PNAS* **87**, 2264 (1990), ISSN 0027-8424, 1091-6490, URL <http://www.pnas.org/content/87/6/2264>.
- [7] S. F. Altschul and W. Gish, in *Methods in Enzymology*, edited by Russell F. Doolittle (Academic Press, 1996), vol. Volume 266 of *Computer Methods for Macromolecular Sequence Analysis*, pp. 460–480, ISBN 0076-6879, URL <http://www.sciencedirect.com/science/article/pii/S0076687996660297>.
- [8] A. K. Hartmann, *Phys. Rev. E* **65**, 056102 (2002), URL <http://link.aps.org/doi/10.1103/PhysRevE.65.056102>.
- [9] S. Wolfsheimer, B. Burghardt, and A. K. Hartmann, *Algorithms for Molecular Biology* **2**, 9 (2007), ISSN 1748-7188, URL <http://www.almob.org/content/2/1/9/abstract>.
- [10] C. Webber and G. J. Barton, *Bioinformatics* **17**, 1158 (2001), ISSN 1367-4803, 1460-2059, URL <http://bioinformatics.oxfordjournals.org/content/17/12/1158>.
- [11] H. Pang, J. Tang, S.-S. Chen, and S. Tao, *BMC Bioinformatics* **6**, 257 (2005), ISSN 1471-2105, URL <http://www.biomedcentral.com/1471-2105/6/257/abstract>.
- [12] M. O. Dayhoff and R. M. Schwartz, in *Atlas of Protein Sequence and Structure* (1978).
- [13] S. Henikoff and J. G. Henikoff, *PNAS* **89**, 10915 (1992), ISSN 0027-8424, 1091-6490, URL <http://www.pnas.org/content/89/22/10915>.
- [14] S. B. Needleman and C. D. Wunsch, *Journal of Molecular Biology* **48**, 443 (1970), ISSN 0022-2836, URL <http://www.sciencedirect.com/science/article/pii/0022283670900574>.
- [15] O. Gotoh, *Journal of Molecular Biology* **162**, 705 (1982), ISSN 0022-2836, URL <http://www.sciencedirect.com/science/article/pii/0022283682903989>.
- [16] D.-F. Feng and R. F. Doolittle, *J Mol Evol* **25**, 351 (1987), ISSN 0022-2844, 1432-1432, URL <http://link.springer.com/article/10.1007/BF02603120>.
- [17] T. F. Smith and M. S. Waterman, *Journal of Molecular Biology* **147**, 195 (1981), ISSN 0022-2836, URL <http://www.sciencedirect.com/science/article/pii/0022283681900875>.
- [18] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, *The Journal of Chemical Physics* **21**, 1087 (1953), ISSN 0021-9606, 1089-7690, URL <http://scitation.aip.org/content/aip/journal/jcp/21/6/10.1063/1.1699114>.
- [19] C. Geyer, in *23rd Symposium on the Interface between Computing Science and Statistics* (Interface Foundation North America, Fairfax, 1991), p. 156.
- [20] E. Marinari and G. Parisi, *EPL* **19**, 451 (1992), ISSN 0295-5075, URL <http://iopscience.iop.org/0295-5075/19/6/002>.
- [21] K. Hukushima and K. Nemoto, *Journal of the Physical Society of Japan* **65**, 1604 (1996), ISSN 0031-9015, 1347-4073, URL <http://journals.jps.jp/doi/abs/10.1143/JPSJ.65.1604>.
- [22] L. A. Newberg, *Journal of Computational Biology* **15**, 1187 (2008), ISSN 1066-5277, URL <http://online.liebertpub.com/doi/abs/10.1089/cmb.2008.0125>.
- [23] S. Wolfsheimer, I. Herms, A. Hartmann, and S. Rahmann, *BMC Bioinformatics* **12**, 47 (2011).
- [24] A. B. Robinson and L. R. Robinson, *PNAS* **88**, 8880 (1991), ISSN 0027-8424, 1091-6490, URL <http://www.pnas.org/content/88/20/8880>.
- [25] For amino acid sequences the new residue is drawn according to the naturally occurring background frequency as found by Robinson & Robinson [24]
- [26] For the BLOSUM62 matrix this is the amino acid triptophan (W) with $s(W, W) = 11$