A minimum reporting standard for multiple sequence alignments

Thomas KF Wong[1,2], Subha Kalyaanamoorthy[1,3], Karen Meusemann[4,5,6], David K Yeates[4], Bernhard

Misof[5], Lars S Jermiin[1,2,7,8,*]

[1]Land & Water, CSIRO, Canberra, Australian Capital Territory, Australia.

[2]Research School of Biology, Australian National University, Canberra, Australian Capital Territory,

Australia.

[3]Department of Chemistry, University of Waterloo, Waterloo, Ontario, Canada.

[4]Australian National Insect Collection, CSIRO National Research Collections Australia, Canberra,

Australian Capital Territory, Australia.

[5]Zoologisches Forschungsmuseum Alexander Koenig, Bonn, Germany.

[6] Evolutionsbiologie & Ökologie, Institut für Biologie I, Albert-Ludwigs-Universität Freiburg,

Freiburg (Brsg.), Germany.

[7]School of Biology and Environmental Science, University College Dublin, Belfield, Ireland.

[8]Earth Institute, University College Dublin, Belfield, Ireland.

[*]Correspondence should be addressed to L.S.J. (lars.jermiin@anu.edu.au)

Multiple sequence alignments (MSAs) play a pivotal role in studies of molecular sequence data,

but nobody has developed a minimum reporting standard (MRS) to quantify the completeness

of MSAs. We present an MRS that relies on four simple completeness metrics. The metrics are

implemented in AliStat, a program developed to support the MRS. A survey of published MSAs

illustrates the benefits and unprecedented transparency offered by the MRS.

MSAs are widely used during annotation and comparison of molecular sequence data, allowing us

to identify medically important substitutions[1], infer the evolution of species[2], detect lineage- and

site-specific changes in the evolutionary processes[3] and use ancestral sequence reconstruction to

engineer new enzymes[4]. There is a wide range of computational tools to obtain MSAs, and two of

these (i.e., Clustal W[5] and Clustal X[6]) are now among the 100 most cited papers in science[7].

MSAs often have gaps inserted between the nucleotides or amino acids of some of the sequences.

These gaps are inserted to maximize the homology of residues from different sequences. A correct

MSA is necessary for accurate genome annotation, phylogenetic inference and ancestral sequence

reconstruction. However, deciding where to put the alignment gaps may be more art than science.

This is because homology is defined as similarity due to historical relationships by descent[8]. Most

of these relationships belong to the unobservable distant past, so it is impossible to measure the

accuracy of most MSAs inferred from real sequence data.

Without this ability, reporting the completeness of MSAs may be the best that can be achieved. So

far, the only metric sometimes used is the *percent missing data* for a sequence[9] or an alignment[10],

but neither is sufficiently transparent and insightful. To ameliorate this, we developed a minimum

reporting standard (MRS) for MSAs.

The MRS uses four metrics to quantify the completeness of different attributes of MSAs. Given an

MSA with m sequences and n sites, we may compute four metrics: $C_a = x_a/(m \times n)$, $C_r = x_r/n$,

$C_c = x_c/m$ and $C_{ij} = x_{ij}/n$, where x_a is the number of completely specified characters[11] in the

MSA, x_r is the number of completely specified characters in the rth sequence of the MSA, x_c is

the number of completely specified characters in the cth column of the MSA and x_{ij} is the number

of homologous sites with completely specified characters in both sequences (i and j). In summary,

$\mathcal{C}_a$, $\mathcal{C}_r$, $\mathcal{C}_c$ and $\mathcal{C}_{ij}$ measure the completeness of the alignment, the rth sequence, the cth site, and

the ith and jth sequences, respectively.

The first of these metrics ($\mathcal{C}_a$) is related to the *percent missing data* used previously, but it is also,

as shown in **Figure 1a**, the least useful completeness metric considered here: Alignments A and B

differ greatly, but they have the same $\mathcal{C}_a$ value (i.e., 0.7). The $\mathcal{C}_r$, $\mathcal{C}_c$ and $\mathcal{C}_{ij}$ metrics, on the other

hand, are able to detect these differences. For example, the $\mathcal{C}_r$ values range from 0.3 to 1.0 for

Alignment A and from 0.4 to 1.0 for Alignment B, raising greater concern, from a sequence-centric

perspective, about Alignment A than about Alignment B. If we were to omit any sequence from

Alignment A, then it would be sensible to omit the one with the smallest $\mathcal{C}_r$ value. The $\mathcal{C}_c$ values

range from 0.2 to 1.0 for Alignment A and from 0.5 to 0.8 for Alignment B. Again, there is greater

concern about Alignment A than about Alignment B (due to the lower $\mathcal{C}_c$ scores and the greater

range of values). The $\mathcal{C}_{ij}$ values range from 0.3 to 1.0 for Alignment A and from 0.0 to 0.9 for

Alignment B. There is cause for great concern if $\mathcal{C}_{ij} = 0.0$ is detected as it means that sequences i

and j have no shared homologous sites with completely specified characters in both sequences.

Evolutionary distances between such sequences cannot be estimated unless the MSA contains at

least one other sequence that overlaps both i and j. When such a case occurs, the evolutionary

distance between sequences i and j is called *inferred by proxy*. Currently, the prevalence of this

problem is unknown.

Figures 1b and **1c** reveal the distributions of $\mathcal{C}_r$ and $\mathcal{C}_c$ for Alignments A and B, offering additional

insight into the alignments' completeness. Conveniently, the $\mathcal{C}_c$ scores may be used to selectively

omit the least complete sites. This *masking of sites* in MSAs is popular in phylogenetics and many

methods[12-20] are now available. Additional information can be obtained by analyzing heat maps

generated from the $\mathcal{C}_{ij}$ values. **Figure 1d** shows the heat maps obtained from Alignments A and B.

The most obvious things to note are that in Alignment A *Tagliatelle* stands out as being the least

complete sequence whereas *Capellini* and *Spaghetti* share no homologous sites with completely

specified nucleotides in both sequences in Alignment B. Although this was easy to detect in **Figure**

1, it will be more difficult to do if n and/or m were larger, as is typically the case in phylogenomic

data.

The benefits offered by the new completeness metrics are clear, but embedding figures like those

in **Figure 1** in publications may be impractical. Alternatively, the essential details may be reported

in a table (**Table 1**) or in one line (e.g., Alignment B: $m = 10$, $n = 100$, $C_a = 0.7$, $C_r = [0.4, 1.0]$,

$C_c = [0.4, 0.8]$, and $C_{ij} = [0.0, 0.9]$). The closer to 1.0 the four C scores are, the more complete

an alignment is. If, on the other hand, the values are closer to 0.0 than to 1.0, users may consider

masking some of the sequences and/or sites before starting a phylogenetic analysis of the data.

Given their potential to inform researchers across a wide range of scientific disciplines, we argue

that m, n, C_a, C_r, C_c and C_{ij} should be combined into what we henceforth call an MRS for MSAs,

and that publications that report all of these values be labelled *compliant with the MRS for MSAs*.

To our knowledge, this has never been done beforehand, leading to widespread ignorance about

the MSAs that are relied upon in ground-breaking bio-medical research.

The MRS may be used to identify dubious MSAs in phylogenomic projects. These alignments occur

in large phylogenomic projects because of a lack of awareness concerning a multitude of problems

in the assembly, orthology assignment, and alignment procedures. Methods to identify odd MSAs

are available[2], but they are not yet fully reliable. The MRS enables a better outlier check for MSA.

Typically, MSAs comprise more sequences and sites than those in **Figure 1a**, so to facilitate using

the MRS, we implemented AliStat, a fast, flexible, and user-friendly program for surveying MSAs.

AliStat computes the C_a, C_r, C_c, and C_{ij} values from MSAs of nucleotides, di-nucleotides, codons

and amino acids., AliStat lists the results on the command-line or in files that can be accessed by

other programs.

The benefit of new MRS for MSAs is underlined in two surveys of large MSAs (**Table 2**). In the first

case, surveying an MSA of the enzyme carboxyl/cholinesterase[21] revealed that some of the $\mathcal{C}_r$ and

$\mathcal{C}_c$ scores are closer to 0.0 than 1.0, and that at least two sequences have no homologous sites in

common with completely specified characters in both sequences. Further inspection of the output

files revealed large proportions of low $\mathcal{C}_r$ and $\mathcal{C}_c$ scores (**Supplementary material**), so it might be

wise to mask some of the sequences and/or sites before phylogenetic analysis of these data.

In the second case, surveying a massive concatenation of MSAs of nuclear genes[22] revealed a more

complete alignment but also low $\mathcal{C}_r$, $\mathcal{C}_c$ and $\mathcal{C}_{ij}$ values. The presence of these values indicates that

additional masking of this MSA might have been wise (see **Supplementary material**). For example,

omitting the two least-complete sequences (i.e., the genera *Leucoptera* and *Pseudopostega*) could

have been considered.

The MRS for MSAs is a sound solution to a large and so-far-neglected problem: how do we report,

as transparently and informatively as possible, the completeness of the MSAs used in bio-medical

research? Better transparency about the completeness of MSAs is clearly needed, because MSAs

represent a foundational cornerstone in many bio-medical research projects and, as revealed by

the example in **Figure 1**, MSAs may look different but have the same percentage of missing data.

So far, information on the completeness of MSAs used in bio-medical research has been largely

absent, leaving readers unable to critically evaluate the merits of scientific discoveries made on

the basis of MSAs. It is critical to recognize, and acknowledge, that many MSAs are the result of

scientific procedures. Therefore, it is necessary to present the results of these procedures more

transparently and comprehensively. Many scientific papers now include links to the MSAs used,

but the MSAs are often so large that it is impossible to form a comprehensive picture about the

completeness of these MSAs.

Our MRS enables a radical change in scientific behavior, allowing authors to report their results

more transparently, and readers the ability to critically assess discoveries made from analyses of

sequence data stored in MSAs.

METHODS

137   Further details about AliStat are available in the online version of this paper.

DATA

140   Data, and code used to analyze the data, are available from http://github.com/thomaskf/AliStat.

ACKNOWLEDGEMENTS

143   We thank staff at the Australian National University and University College Dublin for their helpful

144   feedback on the color scheme used in the heat map. Many of the respondents were color-blind.

AUTHOR CONTRIBUTIONS

147   L.S.J. conceived the project and wrote the first version of AliStat (in C) to conduct a pilot study of

148   the merits of using the metrics. B.M. implemented the first Perl script to produce the heat maps,

149   and T.K.F.W. implemented the final version of AliStat (in C++). S.K., K.M., D.K. Y. and L.S.J. tested

150   the C++ version of AliStat, and provided constructive feedback on the software. L.S.J. drafted the

151   paper with input from the other authors.

COMPETING FINANCIAL INTERESTS

154   The authors declare not competing financial interests.

156 **FIGURE LEGENDS**

157 **Figure 1**. Example, based on two multiple sequences alignments (**a**), illustrating the corresponding

158 distributions of completeness scores for rows (**b**), columns (**c**), and pairs of sequences (**d**).

160 **Table 1**. Example of the MRS for the alignments in **Figure 1a**

Feature	Alignment A	Alignment B
Sequences	10	10
Sites	100	100
Alphabet	Nucleotides	Nucleotides
$\mathcal{C}_a$	0.7	0.7
$\mathcal{C}_r$ [min – max]	0.3 – 1.0	0.4 – 1.0
$\mathcal{C}_c$ [min – max]	0.1 – 1.0	0.4 – 0.8
$\mathcal{C}_{ij}$ [min – max]	0.3 – 1.0	0.0 – 0.9

162 **Table 2**. Example of the MRS for two published MSAs

Feature	Carboxyl/Colineesterase[21]	Lepidoptera[22]
Sequences	364	203
Sites	2,645	749,791
Alphabet	Amino acids	Amino acids
$\mathcal{C}_a$	0.2262	0.6422
$\mathcal{C}_r$ [min – max]	0.0106 – 0.5550	0.0609 – 0.9738
$\mathcal{C}_c$ [min – max]	0.0027 – 0.9972	0.0000 – 0.9655
$\mathcal{C}_{ij}$ [min – max]	0.0000 – 0.5550	0.0084 – 0.9672

ONLINE METHODS

The MRS for MSAs is first-of-a-kind, and AliStat, which enables compliance with the MRS for MSAs,

is, to our knowledge, the first program to compute the four completeness scores presented above.

AliStat is written in C++ and is available, under an CSIRO Open Source Software License Agreement

(variation of the BSD / MIT License), from http://github.com/thomaskf/AliStat/).

AliStat reads a text file with nucleotide, di-nucleotide, codon or amino-acid sequences, which are

aligned and saved in the FASTA format. In other words, AliStat considers alphabets with four, 16,

20, and 64 states. If the sequences comprise single nucleotides, the characters may be lumped to

form six 3-state alphabets (i.e., CRT, AGY, ACK, GMT, AST and CGW) and seven 2-state alphabets

(i.e., RY, KM, SW, AB, CD, GH and TV). If the 3- and 2-state alphabets are used, the letters R, Y, K,

M, S, W, B, D, H and V are considered completely specified characters (unlike normal practice[11]).

AliStat can be run in two modes: Brief mode or Full mode. Execution in brief mode is done using

the following command:

180
```
alistat <infile> <data type> -b
```

and results in the following output format are printed to the terminal:

File name, #seqs, #sites, $\mathcal{C}_a$, max $\mathcal{C}_r$, min $\mathcal{C}_r$, max $\mathcal{C}_c$, min $\mathcal{C}_c$, max $\mathcal{C}_{ij}$, min $\mathcal{C}_{ij}$

The brief-mode execution was included to allow users to quickly obtain the essential values from a

great number of alignments (e.g., when comparing genomes phylogenetically).

The full-mode execution (default option) allows other options to be used and is intended when a

more detailed examination of an MSA is required. For example, the $-t$ option is used to indicate

what types of $\mathcal{C}$ scores should be printed in output files, the $-m$ option is used to set a threshold

for masking sites, and the $-i$ option is used to indicate that a heat map is needed. Other options

and how all of the options may be used are described in the AliStat manual. The same information

can be obtained by typing

`alistat -h`

in the command-line.

The output files appear in the .txt, .csv, .R, .dis, .svg, and .fst formats, which can be processed by

other software packages. The .txt file summarizes the results. The .csv files present the $\mathcal{C}$ scores

and may be examined using R. For example, if a user wishes to generate a histogram of the $\mathcal{C}_c$

scores, the Table_2.csv file may be analyzed using the Histogram_Cr.R file. In some cases, users

may want to infer a tree or network based on the $\mathcal{C}_{ij}$ score (or the $\mathcal{I}_{ij}$ score, where $\mathcal{I}_{ij} = 1.0 -$

$\mathcal{C}_{ij}$). In such cases, .dis files may be analyzed by, for example, SplitsTree[23]. The heat map, which

may be triangular or square, is stored in the .svg file and may be opened using Adobe Illustrator™.

If the $-\text{m}$ option is used, the original MSA is split into two, with all sites having a $\mathcal{C}_c$ score larger

than a user-specified threshold saved in a file called Mask.fst and the other sites saved in a file

called Disc.fst. The two .fst files may be analyzed separately by other means (e.g., phylogenetic

programs).

1   Higgs, D. R. & Wood, W. G. *Proc. Natl. Acad. Sci. USA* **105**, 11595-11596 (2008).

2   Misof, B. *et al. Science* **346**, 763-767 (2014).

3   Jayaswal, V., Wong, T. K. F., Robinson, J., Poladian, L. & Jermiin, L. S. *Syst. Biol.* **63**, 726-742
(2014).

4   Wilding, M. *et al. Green Chem.* **19**, 5375-5380 (2017).

5   Thompson, J. D., Higgins, D. G. & Gibson, T. J. *Nucl. Acid. Res.* **22**, 4673-4680 (1994).

6   Thompson, J. D., Gibson, T. J., Plewniak, F., Jeanmougin, F. & Higgins, D. G. *Nucl. Acid. Res.*
25, 4876-4882 (1997).

7   Van Noorden, R., Maher, B. & Nuzzo, R. *Nature* **514**, 550-553 (2014).

8   Morrison, D. A. *Syst. Bot.* **40**, 14-26 (2015).

9   Wiens, J. J. *Syst. Biol.* **52**, 528-538 (2003).

10   Driskell, A. C. *et al. Science* **306**, 1172-1174 (2004).

11   Cornish-Bowden, A. *Nucl. Acid. Res.* **13**, 3021-3030 (1985).

12   Castresana, J. *Mol. Biol. Evol.* **17**, 540-552 (2000).

13   Talavera, G. & Castresana, J. *Syst. Biol.* **56**, 564-577 (2007).

14   Dress, A. W. M. *et al. Algor. Mol. Biol.* **3**, 7 (2008).

15   Hartmann, S. & Vision, T. J. *BMC Evol. Biol.* **8**, 95 (2008).

16   Capella-Gutierrez, S., Silla-Martinez, J. M. & Gabaldon, T. *Bioinformatics* **25**, 1972-1973
(2009).

17   Misof, B. & Misof, K. *Syst. Biol.* **58**, 21-34 (2009).

18   Kück, P. *et al. Front. Zool.* **7**, 10 (2010).

19   Criscuolo, A. & Gribaldo, S. *BMC Evol. Biol.* **10**, 210 (2010).

20   Wu, M. T., Chatterji, S. & Eisen, J. A. *PLoS One* **7**, e30288 (2012).

21   Pearce, S. L. *et al. BMC Biol.* **15**, 63 (2017).

22   Kawahara, A. Y. *et al. Proc. Natl. Acad. Sci. USA* **116**, 22657-22663 (2019).

23   Huson, D. H. & Bryant, D. *Mol. Biol. Evol.* **23**, 254-267 (2006).

Supplementary Material:

A minimum reporting standard for multiple sequence alignments

Thomas KF Wong[1,2,#], Subha Kalyaanamoorthy[1,3,#], Karen Meusemann[4,5,6], David K Yeates[4],

Bernhard Misof[5], Lars S Jermiin[1,2,7,8,*]

[1]Land & Water, CSIRO, Canberra, Australian Capital Territory, Australia.

[2]Research School of Biology, Australian National University, Canberra, Australian Capital

Territory, Australia.

[3] Department of Chemistry, University of Waterloo, Waterloo, Ontario, Canada.

[4]Australian National Insect Collection, CSIRO National Research Collections

Australia, Canberra, Australian Capital Territory, Australia.

[5]Zoologisches Forschungsmuseum Alexander Koenig, Bonn, Germany.

[6]Evolutionsbiologie & Ökologie, Institut für Biologie I, Albert-Ludwigs-Universität Freiburg,

Freiburg (Brsg.), Germany.

[7]School of Biology and Environmental Science, University College Dublin, Belfield, Ireland.

[8]Earth Institute, University College Dublin, Belfield, Ireland.

[#]Joint first authors (these authors contributed equally to this work)

[*]Correspondence should be addressed to L.S.J. (lars.jermiin@anu.edu.au)

Analysis of an alignment of carboxyl/cholinesterases (CCEs) from a paper by Pearce et al.[1]

The alignment of amino acids used to annotate the CCE genes from *Helicoverpa armigera*, *H. zea*, *Manduca sexta* and *Bombyx mori* was surveyed using AliStat v1.11. The alignment comprised 364 sequences and 2645 sites. **Figure S1** presents the distribution of C_r values, with several sequences having values close to 0.0. Because the objective of the study by Pearce et al.[1] was to annotate the genes, it was not possible to remove any sequences from the data.

Figure S1. Histogram showing the distribution of C_r scores from the CCE genes. The arrows point to the lowest C_r scores.

Figure S2 reveals the distribution of C_c scores, with a high proportion of sites with low C_c values. Based on this distribution, sites with $C_c \leq 0.5$ were masked in the study by Pearce et al.[1].

43 **Figure S2**. Histogram showing the distribution
44 of C_c scores from the CCE genes.

46 **Figure S3** shows the heat map of C_{ij} scores, revealing a high proportion of sequence pairs

47 with low C_{ij} values. Because the aim of the study by Pearce et al.[1] was to annotate the genes,

48 it was not possible to remove any sequences from the data.

50
Figure S3. Heat map showing the distribution of C_{ij} scores from the CCE genes. The benefit of
this heat map is realized by enlarging the image, at which point it become obvious that two
sequences, labelled HzeaCCE016h and MsexCCE001s, have no homologous sites with
unambiguous characters in both sequences.

Pearce et al.[1] used a threshold of $C_c = 0.5$ to mask the alignment of amino acids. **Table 1**

reveals the effect of doing so. In this case, the table complies with the minimum reporting

standard (MRS) for multiple sequence alignments (MSAs).

Table 1. Example highlighting the effect of using $C_c = $
0.5 to mask the sites in the alignment of CCEs.

Feature	Before masking	After masking
Sequences	364	364
Sites	2,645	546
Alphabet	Amino acids	Amino acids
C_a	0.2262	0.9562
C_r [min − max]	0.0106 − 0.5550	0.0458 − 0.9890
C_c [min − max]	0.0027 − 0.9972	0.5055 − 0.9972
C_{ij} [min − max]	0.0000 − 0.5550	0.0000 − 0.9890

63 **Analysis of lepidopteran nuclear data from a study by Kawahara et al.[2]**
64 The alignment of amino acids labelled SA4_aminoacid_supermatrix_resorted_renamed.fas
65 was surveyed using AliStat v1.11. The alignment comprised 203 sequences and 749,791 sites.
66 **Figure S4** reveals the distribution of C_r values, with most sequences having values close to
67 1.0. In this case, only a few sequences had very low C_r scores.

69
Figure S4. Histogram showing the distribution of
C_r scores from the super-alignment (i.e., 2098
MSAs of concatenated single-copy genes). The
arrows point to the two lowest C_r scores.

Figure S5 reveals the distribution of C_c scores, with a high proportion of sites with high C_c
values. Based on this distribution, omitting sites with $C_c \leq 0.2$ might have been sufficient.

83   **Figure S6** shows the heat map of C_{ij} scores, revealing a low proportion of sequence pairs with

84   low C_{ij} values. In this case, the sequences found to be least complete are from the genera

85   *Leucoptera* and *Pseudopostega*.

87
88 **Figure S6**. Heat map showing the distribution of $\mathcal{C}_{ij}$ scores from the concatenate gene
89 alignments. Again, the benefit of this heat map is realized by enlarging the image. In this case,
90 the two most incomplete sequences are identified as being from the genera *Leucoptera* and
91 *Pseudopostega*.
92

93 **References**

94 1   Pearce, S. L. *et al. BMC Biol.* **15**, 63 (2017).
95 2   Kawahara, A. Y. *et al. Proc. Natl. Acad. Sci. USA* **116**, 22657-22663 (2019).

96
97