Check for updates

METHOD ARTICLE

Detecting “protein words” through unsupervised word segmentation [version 1; peer review: 1 approved with reservations, 1 not approved]

Wang Liang¹, Zhao Kaiyong²¹Sogou Tech, Beijing, China²Department of Computer Science, Hong Kong Baptist University, Hong Kong, Hong Kong

V1 First published: 30 Dec 2015, 4:1517
<https://doi.org/10.12688/f1000research.7428.1>
Latest published: 30 Dec 2015, 4:1517
<https://doi.org/10.12688/f1000research.7428.1>

Abstract

Unsupervised word segmentation methods were applied to analyze protein sequences. Protein sequences, such as “MTMDKSELVQKA...,” were used as input to these methods. Segmented protein word sequences, such as “MTM DKSE LVQKA,” were then obtained. We compared the protein words derived via unsupervised segmentation and protein secondary structure segmentation. An interesting finding is that unsupervised word segmentation is more efficient than secondary structure segmentation in expressing information. Our experiment also suggests the presence of several “protein ruins” in current non-coding regions.

Keywords

Word segmentation , Protein sequence , Protein secondary structure , Unsupervised method , Soft counting , Protein word , Gene finding , Description length

Open Peer Review

Approval Status ? X

	1	2
version 1 30 Dec 2015	? view	X view
1. Rahul Banerjee , Saha Institute of Nuclear Physics, Kolkata, India		
2. Judith Klein-Seetharaman , University of Warwick, Coventry, UK		
Any reports and responses or comments on the article can be found at the end of the article.		

Corresponding author: Wang Liang (wangliang.f@gmail.com)

Competing interests: No competing interests were disclosed.

Grant information: The authors declared that no grants were involved in supporting this work.

The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Copyright: © 2015 Liang W and Kaiyong Z. This is an open access article distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

How to cite this article: Liang W and Kaiyong Z. Detecting “protein words” through unsupervised word segmentation [version 1; peer review: 1 approved with reservations, 1 not approved] F1000Research 2015, 4:1517
<https://doi.org/10.12688/f1000research.7428.1>

First published: 30 Dec 2015, 4:1517 <https://doi.org/10.12688/f1000research.7428.1>

Introduction

Word segmentation mainly refers to the process of dividing a string of written language into its component words. For some East Asian languages such as Chinese, no space or punctuation is placed between letters. The “letter” sequences must be segmented into word sequences to process the text of these languages. For example, “iloveapple” is segmented into “I love apple.”

The main concept of the segmentation method is simple. For example, the letter sequence “iloveapple” can be segmented into many forms of word sequences, including “I love apple,” “ilove apple,” “il ove apple,” “iloveapple,” “ilo vea pple,” etc. We select the “word sequences” with the highest probability as the segmentation of this letter sequence. For example:

$$P(\text{“I love apple”}) = 0.8$$

$$P(\text{“ilove apple”}) = 0.3$$

$$P(\text{“iloveapple”}) = 0.03$$

Hence, we will select “I love apple” as the segmentation of letter sequence “iloveapple.”

The probability of a word sequence is calculated by the probability of the multiplication of each word.

For example:

$$P(\text{“I love apple”}) = P(i) \times P(\text{love}) \times P(\text{apple})$$

$$P(\text{“ilove apple”}) = P(\text{ilove}) \times P(\text{apple})$$

Therefore, the key to segmentation is determining the probability of each word.

If we have many possible segmented sequences such as that for “i love apple,” then we can simply calculate the probability of each word by dividing the number of its occurrences by the number of occurrences of all the words. For example, we have two sequences, namely, “i love apple” and “i love mac;” that is, we have six word occurrences. “i” appears two times; thus, its probability $P(\text{“i”}) = 2/6$. “apple” appears one time; thus, its probability $P(\text{“apple”}) = 1/6$. This method is called supervised segmentation. For words such as “ilove” that do not appear in corpus, an extremely small probability is provided.

If we have no segmented sequence but many letter sequences such as “iloveapple,” then we can still calculate the probability of all the possible words and develop the word segmentation method. The main concept of this method is described as follows.

First, the random or equal probability of each possible word is provided. In general, we need to set a maximal word length. For example, for “iloveapple,” if we set four as the maximal word length, then possible words include “ilov,” “love,” “ovea,” etc.

Second, given that we have the word list with probability, we can segment all the letter sequences into word sequences.

Third, we calculate the new word probability according to the word sequences in the second step.

We repeat the second and third steps until the probabilities of all the words reach a stable value. Then, we can use this word list with probability to segment letter sequences. This method is called unsupervised segmentation¹.

The basic theory of word segmentation is depicted in **Figure 1**.

Figure 1. A. Supervised word segmentation. B. Unsupervised word segmentation.

For English text, the accuracy of supervised segmentation is over 95%. Meanwhile, the unsupervised segmentation method can reach approximately 75%. Supervised methods are better than unsupervised methods; however, unsupervised approaches are highly adaptive to relatively unfamiliar “languages” for which we do not have sufficient linguistic knowledge².

Unsupervised segmentation can be described as “inputting letter sequences and outputting meaningful word sequences.” If the input is a protein sequence such as “MTMDKSELVQKA...,” then what is the output? The answer to this question is the topic of this study.

In the following sections, we will perform an unsupervised segmentation of protein sequences. Identifying the functional equivalents of words in protein sequences is a fundamental problem in bioinformatics. The elements of protein secondary structure are generally regarded as the functional building blocks of protein sequences. Therefore, we compare our “protein word” sequences and related secondary structure segmentation in Section 3. We find our “protein word segmentation” to be a better encoding method than secondary structure segmentation in terms of expressing information. We discuss this confusing result in Section 4. The last section provides a short summary of the study.

Word segmentation experiment

The soft-counting method is a representative unsupervised segmentation method; it uses the standard expectation–maximization algorithm to estimate the probability of any possible substring¹. This operation produces a vocabulary with probability. To segment a sequence, a Viterbi-style algorithm is then applied to search for the segmentations with the highest probabilities.

Unsupervised segmentation methods require only raw letter sequences. In this study, we use amino acid letter sequences. Data are selected from the Protein Data Bank (PDB, <http://www.rcsb.org/pdb/>) structure database. The PDB database contains approximately 100,000 protein sequences. These sequences are used as input in the unsupervised segmentation method. Thus, we obtain:

Definition 1: “protein word,” the segments produced by a word segmentation method.

Unsupervised segmentation only requires selecting a maximal word length. In this study, we set nine as the maximal protein word length. This value can be adjusted according to data size.

Approximately 26 million possible protein words are included in our data. Frequency and border information features are used to filter these words^{3–6}. After performing the soft-counting method for these protein sequences, we obtain a final protein vocabulary with approximately 630,598 words. We can then use this vocabulary with word probability to segment protein sequences. An example of protein word segmentation is shown as follows.

Original protein sequence:

**DADLAKKNNCIACHQVETKVVGPALKDIAAKYADKD-
DAATYLAGKIKGGSSGVWGQI PMPPNVNVSADAK LAD
WILTLK**

Corresponding segmented protein word sequence:

**DADLAKK NNCIACH QVET KV VGPAL KDIAAK YADK D
DAATYL AGKIK GGSSGV WGQI PMPPN VNVSD ADAKA
LADWI LTLK**

Comparison with secondary structure segmentation

In most human language texts, the output of segmentation is a meaningful word sequence. As expected, the same is true for a protein sequence. We compare a protein word sequence with the elements of protein secondary structure, which act as the functional building blocks of a protein.

Protein secondary structure mainly refers to the local conformation of the polypeptide backbone of proteins that is frequently and discretely classified into a few states. Word segmentation is highly similar to the protein secondary structure assignment process. For protein sequences, such as “MTMDKSELVQKA,” the corresponding consecutive amino acids of the same secondary structure can also be regarded as a structure word. This process can be called secondary structure segmentation (Figure 2). A total of 437,537 distinct structure words are included in our data set.

Definition 2: “structure word,” the protein segments produced by secondary structure segmentation.

Segmentation performance is generally evaluated using the boundary F-score measure, $F = 2RP/(R + P)$. Recall R is the proportion of the correctly recognized word boundaries to all boundaries in the gold standard; precision P is an output of the word segmentation of a segmenter. In this study, we use structure segmentation as the gold standard (Figure 3).

For the soft-counting method, boundary precision is 39.9%, boundary recall is 28.0%, and the boundary F-score is 32.9%. A segmentation example is provided in Figure 4.

Figure 2. Secondary structure segmentation. A protein sequence is segmented by its secondary structure. The sequence in this figure has three secondary structure words: “MTM,” “DKSEL,” and “VQKA.”

Figure 3. Evaluation of word segmentation. For this example, the precision is $1/33 \approx 3\%$ and recall is $1/2 = 50\%$. The F-score is $2 \times 0.33 \times 0.5 / (0.33 + 0.5) \approx 39.8\%$.

MVLS EGEWQLVLHVWAKV EAD VAGHGQDILIRLFKS H PETLEK F DRVKH L Structure segmentation
 ---- HHHHHHHHHHHHHH GGG HHHHHHHHHHHHHH - GGGGGG - TTTT - Secondary structure
 MV LSEGE WQLV LHVW A KVEAD VAGH GQDIL IRLF K SHPET LEKFD R VKHL Soft-counting word segmentation

Figure 4. Structure segmentation and word segmentation of a protein sequence.

The boundary F-score is approximately 85% for unsegmented English text using the soft-counting method. For Chinese text, the boundary F-score is approximately 80%. For protein word segmentation, however, such value is only 32.9%. Therefore, protein word segmentation is different from secondary structure segmentation. This result does not satisfy our expectation. We will analyze the differences in the next section.

Encoding efficiency of segmentation

As a basic statistical feature, word occurrence distribution may explain the differences between the two kinds of segmentation.

In a typical English data set, if we regard the top 10% frequency words in vocabulary as high-frequency words, then the occurrences of these words generally account for approximately 90% of all the letters in the data set. For low-frequency words, for example, 1-frequency (only appear once) words account for approximately 50% of all the words in vocabulary; however, their occurrences only account for approximately 1% of the letters of the entire corpus. This law holds true for most human languages. It is also a fundamental assumption in most unsupervised segmentation methods. The related values of soft-counting segmentation and secondary structure segmentation are provided in Table 1.

Based on Table 1, secondary structure segmentation appears to be “inefficient” primarily because its vocabulary contains too many low-frequency words. We can use “description length (DL)” to describe the efficiency of segmentation⁷. The segmentation process can be regarded as an encoding process that replaces the letters of a word with related word symbols. A codebook, in which each word

is represented by a unique string, can be used to encode a corpus of words as efficiently as possible. The total number of letters required to encode the corpus (i.e., the sum of the lengths of the codebook and the encoded corpus) that uses a well-designed codebook/vocabulary will be less than that for the original corpus. Smaller units, such as morphemes or phonemes, which require fewer code words, and thus, a shorter codebook, can be encoded further. However, efficiently encoding the corpus becomes increasingly difficult as the number of code words used decreases. Meanwhile, some words may never be used when the codebook has too many words. Thus, the length of the codebook and the length of the encoded corpus must be balanced. The DL principle states that a codebook that results in the shortest total combined length must be selected. Therefore DL can be used to describe segmentation efficiency.

The DL values of structure segmentation and soft-counting segmentation are provided in Figure 5.

Table 1. Word occurrence distribution.

	Letter percentage of high-frequency words in the data set	Letter percentage of low-frequency words in the data set
English text	80%	3%
Soft-counting segmentation	34%	4%
Structure segmentation	40%	58%

Figure 5. DL values of structure segmentation and soft-counting segmentation.

From Figure 5, we can determine that the entire DL of structure segmentation is more than that of soft-counting segmentation. In structure segmentation, approximately 45% of the DL value is used for vocabulary. In soft-counting segmentation, such value is only 16%. Structure segmentation designs a large codebook, but only a few of its capabilities are used. Thus, secondary structure segmentation may not be a good encoding method. Two opposing explanations may explain this result.

First, current structure segmentation is really inefficient. Therefore, we can find a different structure assignment method whose structure segments can express the functions of genes better. Our data set uses the DSSP secondary structure definition. We also try other kinds of secondary structure definitions such as STRIDE⁸, but we obtain similar low efficiency segmentation.

Second, the structure segmentation scheme is truly efficient. Given that the segmentation method is appropriate, some problems may exist in the data set. A highly common problem in the data set construction process is the unbalanced data problem. For example, if we only select several paragraphs of a long book to build a data set, then this problem will occur.

We develop the experimental data by filtering similar protein sequences in PDB. The PDB protein data can also be regarded as randomly selected from all available protein data. No obvious problem is observed in our data construction process. Hence, if our

data set is unbalanced, then all current protein data are unbalanced. If yes, then more protein words are present in current non-coding regions. For example, several short DNA regions can be translated into protein words; however, such protein words that cover a region may not comply with the strict definition of a protein sequence. Therefore, we have:

Definition 3: “protein word coverage region,” the DNA region that can be translated into protein words.

We can also detect a protein word coverage region using segmentation methods. We back-translate the protein words into DNA forms and obtain a DNA vocabulary. Then, we can use this vocabulary to segment DNA sequences using supervised segmentation methods. For example, a DNA sequence “ATGGTTTTGTGTCTTGAGGGT” is segmented as Figure 6.

Figure 6. Segment DNA sequence based on a protein word vocabulary. The green parts can be translated into protein words. The red parts cannot be translated. (Note: This sequence is not a protein sequence. Only some regions can be covered by “protein words”).

The following is a typical example of DNA sequence segmentation:

TTGTGG G TTGTCTATTGATGTT TTTGGT C TATTCTAAG A
 A TTGGAG A GAG A GAGGTT A A AAT C TCT G ACT ATG A
 TTGTGG A TTGTCT GCTGAT GTT TTTGGT C TATTGTTCT
 AAGAAT TGG A GAG A GAG A GAGGTT A A A A TCT C C
 G ACT ATG A TTGTGG A TTGTCT ATTGCTGTTTTT GGT C
 TATTGTTCT AAGAAT.

This sequence is part of “gi|157704448|ref|AC_000133.1| *Homo sapiens* chromosome 1, alternate assembly HuRef, whole genome shotgun sequence:307501-308000.” The green segments can be translated into “protein words.” Although 87% of the region of this sequence is a protein word coverage region, no distinct intron, exon, and stop codon is found. Hence, such sequence cannot be regarded as a coding sequence.

We use the human genome to judge whether the data imbalance can be improved by adding all the protein word coverage regions. The genome is divided into 500 bps equal-length sequences. We use two

vocabularies, namely, soft-counting word vocabulary and secondary structure word vocabulary, to segment the aforementioned DNA sequences. Then, we calculate their DL individually. The results are presented in Figure 7.

We determine that the two kinds of segmentation methods have similar DL values. Then, we convert all the DNA protein word coverage regions into protein word sequences. These sequences are regarded as gold standard segmentation. Then, we delete the space between words and perform unsupervised segmentation test again. Boundary precision reaches 67%, boundary recall reaches 60%, and the boundary F-score reaches 63%. These results concur with our conjecture. Adding protein word coverage regions can address the data imbalance and improve structure segmentation efficiency.

The distribution of protein word coverage regions is shown in Figure 8.

We determine that approximately 16% of the DNA sequences have a “protein word coverage region” percentage of over 70%. Among

Figure 7. DL values of structure segmentation and soft-counting segmentation.

Figure 8. Distribution of the protein word coverage sequence in *Homo sapiens* chromosome 1. For example, DNA sequences whose percentage of protein word encoding region is within [0.6, 0.7] are approximately 75%.

these sequences, only approximately 10% are actual gene areas, the remaining sequences may be regarded as protein ruins, which may be abandoned in evolution.

Conclusion

We use protein sequences as input to perform a word segmentation test. As expected, the output is several meaningful protein word sequences. Thus, we compare them with secondary structure segments, which are generally regarded as the basic functional blocks of a protein. However, we determine that the two kinds of segmentation methods are different. Subsequently, we use DL to analyze their difference and determine that word segmentation is more efficient than structure segmentation. This result is irrational. We suppose that this problem is mainly caused by the data imbalance. That is, the gene is similar to a large book and current protein sequences are only several paragraphs in this book. Our DNA segmentation experiments also support this conjecture.

Overall, “protein word” and “protein word coverage region” may have no biological meaning. Word segmentation is a fundamental technology in processing Chinese and other languages without space or punctuation between words, similar to that of protein sequences. Word segmentation is the preliminary step in search engine systems, translation systems, and proofreading systems, etc. The main objective of this study is to inspire new ideas that apply these technologies in analyzing protein sequences. Word segmentation simply acts as a bridge.

Data and software availability

Protein data

Unsupervised segmenting methods only need raw letter sequence. We mainly use the data of PDB (<http://www.rcsb.org/pdb/>) as our experiment data. This dataset contains about 100,000 pieces of protein sequences. We use CD-HIT algorithms to delete the similar protein sequence. The code is available from: <http://weizhong-lab.ucsd.edu/cd-hit/download.php>.

Instructions and source code

Latest source code and experimental instructions: https://github.com/maris205/secondary_structure_detection

Archived source code as at the time of publication: <http://dx.doi.org/10.5281/zenodo.346639>

License: MIT (<http://opensource.org/licenses/MIT>)

Author contributions

Wang Liang designed the experiments, wrote the codes. Zhao Kaiyong prepared the experimental data and provided the biology knowledge supports.

Competing interests

No competing interests were disclosed.

Grant information

The author(s) declared that no grants were involved in supporting this work.

References

1. Ge X, Prat W, Smyth P: **Discovering Chinese words from unsegmented text.** *Proceedings on the 22 Annual International ACM SIGIR Conference On Research and Development in Information Retrieval*. USA, 1999; 271–272.
[Publisher Full Text](#)
2. Wang H, Zhu J, Tang S, *et al.*: **A new unsupervised approach to word segmentation.** *Comput Linguist*. 2011; 37(3): 421–454.
[Publisher Full Text](#)
3. Cohen P, Adams N, Heeringa B: **Voting Experts: An unsupervised algorithm for segmenting sequences.** *Intell Data Anal*. 2007; 11(6): 607–625.
[Reference Source](#)
4. Hewlett D, Cohen P: **Word segmentation as general Chunking.** *Proceedings of the Fifteenth Conference on Computational Natural Language Learning*. USA. 2011; 39–47.
[Reference Source](#)
5. Zhao H, Kit C: **Exploiting unlabeled text with different unsupervised segmentation criteria for Chinese word segmentation.** *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, Israel. 2008; 17–23.
[Reference Source](#)
6. Zhao H, Kit C: **An empirical comparison of goodness measures for unsupervised Chinese word segmentation with a unified framework.** *The Third International Joint Conference on Natural Language Processing*, India. 1: 2008; 9–16.
[Reference Source](#)
7. Kit C, Wilks Y: **Unsupervised learning of word boundary with description length gain.** *CoNLL*, Norway. 1999; 1–6.
[Reference Source](#)
8. Zhang W, Dunker AK, Zhou Y: **Assessing secondary structure assignment of protein structures by using pairwise sequence-alignment benchmarks.** *Proteins*. 2008; 71(1): 61–7.
[PubMed Abstract](#) | [Publisher Full Text](#)
9. Liang W: **Detecting “protein words” through unsupervised word segmentation: First release.** *Zenodo*. 2015.
[Data Source](#)

Open Peer Review

Current Peer Review Status: ? ✖

Version 1

Reviewer Report 21 March 2016

<https://doi.org/10.5256/f1000research.8005.r12977>

© 2016 Klein-Seetharaman J. This is an open access peer review report distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Judith Klein-Seetharaman

Division of Metabolic and Vascular Health, Warwick Medical School, University of Warwick, Coventry, UK

This is an interesting article and approach, but the results presented are too preliminary to assess how useful this approach may be for biologists. The title is fine but the abstract is not sufficient and reflects the deficiencies of the paper. While the idea is interesting and could potentially lead to discovery of new “words” in protein sequences, in its current state, the paper is restricted entirely on the language aspects. There is no expert knowledge on protein sequence/structure/bioinformatics that has gone into the design of the study or the analysis of the results. There are no references at all on prior work on this topic and the methods are described in too brief format. For example, when were the 100,000 sequences downloaded from pdb and does this represent all sequences. What secondary structure elements were included from DSSP/Stride? What cut-off was used to arrive at the reduced dataset, and how many protein sequences were included in the final data analysed?

While it is reasonable to compare unsupervised protein words with secondary structure elements, it does not make sense to assume that the boundaries have to be fixed. In protein structures, the exact boundaries are not precise and that contributes to difficulties in further improvement on accuracy of secondary structure prediction. Furthermore, there are differences between the types of boundaries predicted, so discovery of a coil-helix versus a coil-sheet would have different error rates. Secondary structures are only one type of segmentation, there is others such as protein interaction recognition elements, enzyme active sites, posttranslational modification indicators etc that also carry “meaning” as a “word”. Last but not least, the paper neglects the contributions of tertiary structure in the definition of meaning in biology. All of these issues are much more likely to contribute to the low performance than the speculation on data imbalance. Of course, the pdb structure dataset is extremely imbalanced, and this is noted wrongly in the text that “all current protein data are unbalanced”. The proteins in the pdb are a very biased subset of proteins that crystallize, thus highly over representing small soluble proteins.

Competing Interests: No competing interests were disclosed.

I confirm that I have read this submission and believe that I have an appropriate level of expertise to state that I do not consider it to be of an acceptable scientific standard, for reasons outlined above.

Reviewer Report 01 February 2016

<https://doi.org/10.5256/f1000research.8005.r12078>

© 2016 Banerjee R. This is an open access peer review report distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Rahul Banerjee

Crystallography and Molecular Biology Division, Saha Institute of Nuclear Physics, Kolkata, India

The authors have applied unsupervised word segmentation methods to protein sequences which in itself is a very novel undertaking and overall the manuscript is well written. However there are some outstanding points which the authors could consider to improve the quality of the manuscript.

1. Some justification should be provided as to why 9 was selected as the maximal protein word length. How are the results affected as a function of maximal word length.
2. On page 3, the authors write 'In human language texts, the output of segmentation is a meaningful word sequence. As expected, the same is true for a protein sequence.'

It is not clear what the authors mean by 'meaningful' in the case of protein sequences. What kind of 'meaning' do the authors expect on applying an ostensibly linguistic method to protein sequences.

3. Again on the same page the authors note that 'word segmentation is highly similar to the protein secondary structure assignment process.' Secondary structural assignment in proteins is based on the successive values of the main chain phi,psi angles and the hydrogen bonding patterns. The authors however could mean the propensity based assignment of secondary structure in case of protein structure prediction which is known to be inaccurate, the propensities being estimated from appropriate databases. Although the authors cite back references (3-6) how the structural segmentation was performed should be discussed in much greater detail in the present manuscript.
4. If the structural segmentation was performed on the basis of actual secondary structure present in a protein (phi,psi and hydrogen bonding pattern) then there does not appear to be any connection between the two processes of segmentation, in terms of word and structure. So they should naturally be expected to be different. On fails to understand the contrary expectations of the authors when they note on page 4, 'Therefore, protein word segmentation is different from secondary structure segmentation. This result does not satisfy our expectation'.
5. On should also note that identical sequences have been known to adopt variable secondary

structures in proteins due to other factors such as tertiary contacts and its position with respect to the rest of the polypeptide chain. How does this fact reflect on the overall results of the manuscripts? Therefore although the work concerns a novel implementation of the word segmentation method in the context of protein sequences, some of the 'expectations' and interpretations of the authors possibly require greater clarity or stronger justification.

Competing Interests: No competing interests were disclosed.

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard, however I have significant reservations, as outlined above.

The benefits of publishing with F1000Research:

- Your article is published within days, with no editorial bias
- You can publish traditional articles, null/negative results, case reports, data notes and more
- The peer review process is transparent and collaborative
- Your article is indexed in PubMed after passing peer review
- Dedicated customer support at every stage

For pre-submission enquiries, contact research@f1000.com

F1000Research