

A De Novo Robust Clustering Approach for Amplicon-Based Sequence Data

Alexandre BAZIN, Didier DEBROAS and Engelbert MEPHU NGUIFO

University Clermont Auvergne, CNRS, LIMOS, F-63000 CLERMONT-FERRAND, FRANCE

University Clermont Auvergne, CNRS, LMGE, F-63000 CLERMONT-FERRAND, FRANCE

alexandre.bazin@isima.fr, mephu@isima.fr

Abstract

When analyzing microbial communities, an active and computational challenge concerns the categorization of 16S rRNA gene sequences into operational taxonomic units (OTUs). Established clustering tools use a one pass algorithm in order to tackle high numbers of gene sequences and produce OTUs in reasonable time. However, all of the current tools are based on a crisp clustering approach, where a gene sequence is assigned to one cluster. The weak quality of the output compared to more complex clustering algorithms, forces the user to post-process the obtained OTUs. Providing a membership degree when assigning a gene sequence to an OTU, will help the user during the post-processing task. Moreover it is possible to use this membership degree to automatically evaluate the quality of the obtained OTUs. So the goal of this work is to propose a new clustering approach that takes into account uncertainty when producing OTUs, and improves both the quality and the presentation of the OTUs results.

1 Introduction

Studying the structure of the communities in an ecosystem is central in environmental microbiology [Hugoni *et al.*, 2013; Roux *et al.*, 2011]. The biosphere's diversity can be determined by amplifying and sequencing specific phylogenetic markers (e.g. 16S rRNA). From there, these amplicons need to be clustered in "species" named Operational Taxonomic Units (OTUs) [Chen *et al.*, 2013; Li *et al.*, 2012; Mahé *et al.*, 2014; Westcott and Schloss, 2015]. As the volume of sequences has drastically increased in recent times, new clustering tools have emerged to treat the data in reasonable time. The currently used algorithms are, from the point of view of algorithmic complexity, the fastest available that do not produce random results. However, due to their simplicity, the reliability of the results are often discussed. These tools being essentially black boxes, their sensitivity to the sequence order, clustering threshold and structure of the data makes it that the users have no way of knowing whether

better Operational Taxonomic Units (OTUs) could have been obtained with different parameters or even whether they correctly represent the data. In these circumstances, there is no choice but to blindly trust them.

Distance-based greedy clustering algorithm such as the ones implemented in OTUclust [Albanese *et al.*, 2015], VSEARCH [Rognes *et al.*, 2016], CD-HIT [Li and Godzik, 2006] or USEARCH [Edgar, 2010] all share the same base algorithm as shown in Algorithm 1.

Algorithm 1: DBG Clustering principle

Input : A set of sequences
Output: A set of OTUs to which the sequences are assigned

```
1 Clusters =  $\emptyset$ 
2 foreach sequence S do
3   foreach known cluster C do
4     | Compute distance(S, C)
5   end
6   if a suitable cluster exists then
7     | Assign S to it
8   else
9     | Create a new cluster with S as the center
10  end
11 end
12 Return Clusters
```

While more sophisticated algorithms [Antoine *et al.*, 2014; Gath and Geva, 1989; Pérez-Suárez *et al.*, 2013; Hariz *et al.*, 2006; Antoine *et al.*, 2012] could produce better results quality-wise, their runtime would render them unusable on millions of sequences. As the quality of the OTUs is important, we have to find a way to improve it without increasing the runtime. The different available implementations use a variety of heuristics to counterbalance the simplicity of the algorithm but, to the best of our knowledge, no approach has tried to add a measure of uncertainty to the process. This is why, in order to help increase the quality and trustworthiness of the clustering, we propose to add uncertainty to this simple algorithm through the use of fuzzy clustering.

2 Adding uncertainty to clustering

2.1 Motivation

Distance-based greedy clustering algorithms, such as the one in VSEARCH, produce a number of OTUs and assign each sequence to one of them. The OTU to which a sequence is said to belong to is usually the first one to be encountered that is sufficiently close, i.e. within the specified threshold. This creates two problems :

- A sequence can only belong to a single OTU
- An OTU either includes or does not include a sequence

Having a sequence associated to a single OTU is expected as the ultimate output of the algorithm. For this reason, algorithms can stop after finding the first OTU that is close enough to a sequence, which speeds the computation up. However, not considering all the OTUs a sequence could be assigned to increases the sensitivity to the order - a weakness of these algorithms - and reduces the quality of the clustering. Indeed, what if two different OTUs are close enough ? Giving priority to the first generated OTU only creates a bias that no heuristic - such as sorting the sequences - could hope to overcome.

Moreover, by using strict thresholds, it is possible to have two nearly identical sequences such that one belongs to a particular OTU while the other does not. This strictness makes it so an OTU partitions the set of sequences into two sets inside of which sequences are considered the same regardless of their distance to the center of the OTU. This lack of distinction between sequences that are isolated and sequences on the border of OTUs hides information that could help understand the data.

While these would not be problems were the clustering optimal, the need for fast algorithms gives rise to results that are not always trustworthy. The OTUs being presented as absolute, the end user has no choice, should consider them correct and cannot know whether the algorithm has encountered ambiguity. We believe that being less strict in the way the OTUs partition sequences would help produce better results from the end user's point of view.

2.2 Fuzzy Clustering

To help increase the quality of the clustering and maximize the information that can be gathered from the data, we propose to add uncertainty to the clustering by means of fuzzy sets.

We define a membership function $f_C(S)$ that, for an OTU C , associates a membership value to a sequence S . Usually, this value is either 0 or 1. Here, we propose to have $f_C(S)$ take its value in $\{\frac{n}{10} \mid n = 0..10\}$. This value represents the degree of membership and, as such, 1 means that the sequence **certainly** belongs to the OTU while 0 means that the sequence **certainly** does not belong to it. Other values represent uncertainty and are used to express that the sequence


nearly belongs to the OTU. This membership value can easily be computed from the distance between the sequence and the center of the OTU using two thresholds t_1 and t_2 such that $t_1 \geq t_2$. If the distance is less than the threshold t_1 , the membership value is 1. If the distance is greater than t_2 the value is 0. If the distance is between t_1 and t_2 , it increases gradually.

Algorithm 2: Fuzzy DBG Clustering

Input : A set of sequences
Output: A set of OTUs to which the sequences are assigned

```
1  $Clusters = \emptyset$ 
2 foreach sequence  $S$  do
3   foreach known cluster  $C$  do
4     Compute  $distance(S, C)$ 
5     Assign  $S$  to  $C$  with value  $f_C(S)$ 
6   end
7   if  $S$  has not been sufficiently assigned then
8     Create a new cluster with  $S$  as the center
9   end
10 end
11 Return  $Clusters$ 
```

Figure 1: Representations of a Crisp (Left) and a Fuzzy (Right) Cluster.



Using fuzzy OTUs allows us to discern the difference between sequences close to the OTU and sequences extremely far. Using the parameters t_1 and t_2 , we can tune the “detection radius” around OTUs to gather information that would normally be discarded by the clustering algorithm.

3 Evaluating fuzzy OTUs

Having a non-binary membership function produces OTUs that partition the sequences into multiple sets. If we consider only the sequences that belong (more or less) to an OTU, the repartition of their membership values provides information on the topology of the OTU. An ideal OTU would contain only sequences with a membership value of 1, meaning a group of sequences has been perfectly regrouped with a good threshold and no sequence lies ambiguously on the border. More realistically, a good OTU would contain many sequences with high membership values and little sequences with low values. A bad OTU with the majority of its sequences having low membership values could mean that the

	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
OTU1	6	4	1	1	0	3	8	13	29	88
OTU2	70	41	30	41	34	19	11	6	5	16

Table 1: Two example OTUs with the number of sequences that belong to them with each possible membership value.

ω_1	ω_2	ω_3	ω_4	ω_5	ω_6	ω_7	ω_8	ω_9
1	0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.2

Table 2: Example of weight values.

algorithm has chosen as a center a sequence on the border of a group or, even worse, between two distinct groups.

We can quickly evaluate the quality of an OTU with this repartition. If we suppose that each sequence lowers the quality of the OTU depending on its membership value, we can use the following formula :

$$Quality(OTU) = 1 - \sum_{i=1}^9 \omega_i \times \frac{\# \text{ sequences with membership value } i \times 0.1}{\# \text{ sequences in the OTU}}$$

with ω_i being the ‘‘cost’’ of having a sequence with membership value $i \times 0.1$. In our previous examples, and with the following values of ω_i

we obtain a quality of respectively 0.71 and 0.26 for OTU1 and OTU2, showing OTU1 is better.

A problem arises with singletons that always have perfect quality but these can safely be treated separately.

4 Choosing an OTU

A sequence can belong to multiple OTUs due to fuzzy membership. However, in the end, we want each sequence to be assigned to a single OTU. Hence, we have to choose one of the possible OTUs. We have two types of values left from the clustering process : membership and quality. The first one is based on the distance between the OTU and the sequence and the second one is used to recognize bad OTUs. Choosing the OTU with the best membership value is akin to running VSEARCH. Choosing the OTU with the best quality tends to create bigger OTUs that absorb distant sequences. To better compromise, we can use a linear combination of both values :


$$\alpha \times quality + \beta \times membership$$

Increasing the importance of the quality reduces the number of OTUs containing sequences. When α is low, the ‘‘best’’ OTUs quality-wise absorb very close sequences that would have been attributed to other OTUs. When α gets too high, the best OTUs start absorbing all the sequences around them, effectively acting like an increase of the distance threshold.

5 Identifying ambiguous sequences

Distance-based greedy algorithms are good at clustering objects that are easy to cluster. Groups of very similar sequences that are different from the rest of the dataset are supposed to birth a new OTU while isolated singletons should be identified to be either removed or treated separately. A problem arises when groups of sequences are close to each other but not enough to be the same OTU. In this case and supposing the algorithm ideally chooses the centers of the OTUs, sequences can lie just between these OTUs. In the current implementations, these ambiguous sequences that must be assigned are usually put in OTUs of their own, increasing the number of OTUs and reducing the overall quality of the clustering.

Figure 2: A Case of Ambiguous Sequences



Using fuzzy clustering allows us to identify these ambiguous sequences. Using the previously mentioned choice strategy, they can be assigned to a good OTU even though they lie slightly outside of the distance threshold. However, their ambiguousness may be significant for the user. It is thus important to highlight their existence and the various fuzzy OTUs they could have alternatively been assigned to.

6 Experimental Results

6.1 Data

We used our algorithm on a dataset containing 5977 sequences of length between 900 and 3081 for an average of 1442 and taxonomies extracted from the SILVA database. We used a threshold of 0.97 (97% similarity) for determining new OTUs and a threshold of 0.95 for fuzzy membership. For the choice of the OTU for each sequence, we present the results of three strategies : best quality ($\alpha = 1$ and $\beta = 0$), compromise ($\alpha = 0.5$ and $\beta = 0.5$) and distance ($\alpha = 0$ and $\beta = 1$). The comparison with VSEARCH is done using identical parameters when applicable.

The program, dataset and corresponding taxonomy are available on <http://projets.isima.fr/sclust/Expe.html>.

6.2 Relevant Metrics

To measure the effects of introducing uncertainty to the clustering, we consider the following metrics :

Method	Time (min)	Memory	#OTUs	#Singletons	#Doubletons	Distance
Fuzzy (best quality)	1:06	652744	3461	2581	442	0.75
Fuzzy (compromise)	1:06	651980	3596	2776	413	0.54
Fuzzy (distance)	1:06	683772	3631	2837	395	0.59
VSEARCH	0:21	632832	3716	2935	388	0.57

Table 3: Results of the clustering using default *maxaccepts* and *maxrejects*.

Method	Time (min)	Memory	#OTUs	#Singletons	#Doubletons	Distance
Fuzzy (best quality)	27:01	720968	3431	2575	413	0.60
Fuzzy (compromise)	29:14	734604	3566	2767	398	0.47
Fuzzy (distance)	28:27	723693	3631	2835	391	0.48
VSEARCH	27:51	648052	3631	2859	394	0.52

Table 4: Results of the clustering using *maxaccepts* 10000 and *maxrejects* 10000.

- Computation time in minutes
- Memory usage
- Number of OTUs containing at least a sequence
- Number of OTUs containing a single sequence
- Number of OTUs containing only two sequences
- Average distance in the taxonomy between sequences in the same cluster

The distance between two sequences in the taxonomy is defined as the sum of the lengths of the path from their nearest commonality. For example, if a sequence is classified as "bacteria;proteobacteria;betaproteobacteria" and the other is classified as "bacteria;proteobacteria;alphaproteobacteria", their distance is 2 as each of them is at a distance 1 from their commonality "bacteria;proteobacteria".

6.3 Results

First, let us begin with the results obtained using the default values for *-maxaccepts* and *-maxrejects* in Table 3.

Then, the results obtained using *-maxaccepts 10000* and *-maxrejects 10000* in Table 4.

6.4 Analysis

Results show that the choice strategy affects every metric relevant to the quality of the clustering : number of OTUs, singletons and pairs, average misclassification. The fuzzy approach uses slightly more memory than VSEARCH but all choice strategies are similar on this metric. When using the default *-maxaccepts* and *-maxrejects* values, computation time is lower for VSEARCH. However, when using higher values for these parameters – and thus more precise clustering - the computation time is the same for both approaches.

We observe that increasing the importance of the quality in the OTU choice strategy lowers the final number of OTUs. This is due to the fact that some OTUs are initially created

centered on isolated sequences near good OTUs. That isolation lowers their quality and the good OTUs absorb their sequences.

Using the quality also lowers the number of singletons and increases the number of pairs. This most likely means that singletons were created close to either good clusters or one another. The fuzzy approach allows the algorithm to merge those sequences that were slightly too far from the center with their corresponding OTU. The increase in the number of pairs appears to be due to the merging of singletons lying too close to one another.

The average taxonomy distance in OTUs is shown to vary wildly. Using only the quality to choose OTUs increases this number as the "best" OTUs attract all the sequences in their fuzzy surroundings. This causes some sequences belonging to different species to be classified together. However, using a compromise between quality and distance lowers this metric as the best clusters only absorb sequences that are sufficiently close to them and should probably be together while rejecting the sequences that are too different.

7 Discussion

We observe that the experimental results confirm that adding uncertainty to the clustering helps improve the quality of the output by reducing the number of singletons. Using fuzzy clusters, we are able to extend the clustering threshold to gather additional information on the OTUs's surroundings and use it to quickly assess their quality. This quality can be used together with the distance to choose an OTU for each sequence. The resulting output contains less singletons and misclassifications. Being able to choose the weight of both distance and quality allows for additional tuning.

As previously mentioned, the fuzziness also makes it possible to detect ambiguous sequences and clusters. In our opinion, this is where further work is required. An ambiguous sequence could be arbitrarily assigned to a nearby OTU, become the center of its own OTU or even be considered as an error and deleted but these operations imply such a knowledge of the domain that interactions with the human user become necessary. However, on datasets containing millions of sequences, the number of alerts would render manual treatment impractical or even impossible. Automatizing this treatment would require being able to adapt to the type of data, domain and preferences of the user. We suggest that machine learning techniques be introduced in the process to automatically learn how to handle these ambiguities.

Acknowledgements

This work was supported by the European Union's "Fonds Européen de Développement Régional (FEDER)" program and the Auvergne-Rhone-Alpes region.

References

- [Albanese *et al.*, 2015] Davide Albanese, Paolo Fontana, Carlotta De Filippo, Duccio Cavalieri, and Claudio Donati. Micca: a complete and accurate software for taxonomic profiling of metagenomic data. *Scientific reports*, 5:9743, 2015.
- [Antoine *et al.*, 2012] Violaine Antoine, Benjamin Quost, Marie-Hélène Masson, and Thierry Denoeux. CECM: constrained evidential c-means algorithm. *Computational Statistics & Data Analysis*, 56(4):894–914, 2012.
- [Antoine *et al.*, 2014] Violaine Antoine, Benjamin Quost, Marie-Hélène Masson, and Thierry Denoeux. CEVCLUS: evidential clustering with instance-level constraints for relational data. *Soft Comput.*, 18(7):1321–1335, 2014.
- [Chen *et al.*, 2013] Wei Chen, Clarence K Zhang, Yongmei Cheng, Shaowu Zhang, and Hongyu Zhao. A comparison of methods for clustering 16s rRNA sequences into OTUs. *PLoS one*, 8(8):e70837, 2013.
- [Edgar, 2010] Robert C Edgar. Search and clustering orders of magnitude faster than blast. *Bioinformatics*, 26(19):2460–2461, 2010.
- [Gath and Geva, 1989] Isak Gath and Amir B. Geva. Unsupervised optimal fuzzy clustering. *IEEE Transactions on pattern analysis and machine intelligence*, 11(7):773–780, 1989.
- [Hariz *et al.*, 2006] Sarra Ben Hariz, Zied Elouedi, and Khaled Mellouli. Clustering approach using belief function theory. In *International Conference on Artificial Intelligence: Methodology, Systems, and Applications*, pages 162–171. Springer, 2006.
- [Hugoni *et al.*, 2013] Mylène Hugoni, Najwa Taib, Didier Debroas, Isabelle Domaizon, Isabelle Jouan Dufournel, Gisèle Bronner, Ian Salter, Hélène Agogué, Isabelle Mary, and Pierre E Galand. Structure of the rare archaeal biosphere and seasonal dynamics of active ecotypes in surface coastal waters. *Proceedings of the National Academy of Sciences*, 110(15):6004–6009, 2013.
- [Li and Godzik, 2006] Weizhong Li and Adam Godzik. Cd-hit: a fast program for clustering and comparing large sets of protein or nucleotide sequences. *Bioinformatics*, 22(13):1658–1659, 2006.
- [Li *et al.*, 2012] Weizhong Li, Limin Fu, Beifang Niu, Sitao Wu, and John Wooley. Ultrafast clustering algorithms for metagenomic sequence analysis. *Briefings in bioinformatics*, page bbs035, 2012.
- [Mahé *et al.*, 2014] Frédéric Mahé, Torbjørn Rognes, Christopher Quince, Colomán de Vargas, and Micah Dunthorn. Swarm: robust and fast clustering method for amplicon-based studies. *PeerJ*, 2:e593, 2014.
- [Pérez-Suárez *et al.*, 2013] Airel Pérez-Suárez, José F Martínez-Trinidad, Jesús A Carrasco-Ochoa, and José E Medina-Pagola. Oclustr: A new graph-based algorithm for overlapping clustering. *Neurocomputing*, 121:234–247, 2013.
- [Rognes *et al.*, 2016] Torbjørn Rognes, Tomáš Flouri, Ben Nichols, Christopher Quince, and Frédéric Mahé. Vsearch: a versatile open source tool for metagenomics. *PeerJ*, 4:e2584, 2016.
- [Roux *et al.*, 2011] Simon Roux, Michaël Faubladier, Antoine Mahul, Nils Paulhe, Aurélien Bernard, Didier Debroas, and François Enault. Metavir: a web server dedicated to virome analysis. *Bioinformatics*, 27(21):3074–3075, 2011.
- [Westcott and Schloss, 2015] Sarah L Westcott and Patrick D Schloss. De novo clustering methods outperform reference-based methods for assigning 16s rRNA gene sequences to operational taxonomic units. *PeerJ*, 3:e1487, 2015.