

MeShClust: an intelligent tool for clustering DNA sequences

Benjamin T. James^{1,2}, Brian B. Luczak^{1,2} and Hani Z. Girgis^{1,*}

¹Bioinformatics Toolsmith Laboratory, Tandy School of Computer Science, University of Tulsa, 800 South Tucker Drive, Tulsa, OK 74104, USA and ²Mathematics Department, University of Tulsa, 800 South Tucker Drive, Tulsa, OK 74104, USA

Received February 04, 2018; Revised April 06, 2018; Editorial Decision April 11, 2018; Accepted April 13, 2018

ABSTRACT

Sequence clustering is a fundamental step in analyzing DNA sequences. Widely-used software tools for sequence clustering utilize greedy approaches that are not guaranteed to produce the best results. These tools are sensitive to one parameter that determines the similarity among sequences in a cluster. Often times, a biologist may not know the exact sequence similarity. Therefore, clusters produced by these tools do not likely match the real clusters comprising the data if the provided parameter is inaccurate. To overcome this limitation, we adapted the mean shift algorithm, an unsupervised machine-learning algorithm, which has been used successfully thousands of times in fields such as image processing and computer vision. The theory behind the mean shift algorithm, unlike the greedy approaches, guarantees convergence to the modes, e.g. cluster centers. Here we describe the first application of the mean shift algorithm to clustering DNA sequences. MeShClust is one of few applications of the mean shift algorithm in bioinformatics. Further, we applied supervised machine learning to predict the identity score produced by global alignment using alignment-free methods. We demonstrate MeShClust's ability to cluster DNA sequences with high accuracy even when the sequence similarity parameter provided by the user is not very accurate.

INTRODUCTION

Clustering nucleotide sequences is an essential step in analyzing biological data. Pioneering sequence clustering tools have been proposed for reducing redundancy and correcting errors in next-generation sequencing data (1–6) and for assembling genomes *de-novo* (7). Sequence clustering tools were also proposed for barcode error correction (8) and for taxonomic profiling (9). In addition, d2_cluster (10),

wcdest (11), CD-HIT (3,12), UCLUST (13), DNACLUST (9), mBKM (14) and d2-vmc (15) are general-purpose sequence clustering tools. These tools are applied to clustering gene sequences, expressed sequence tags, RNA and reducing a set of sequences to a non-redundant group of sequences.

However, the currently available tools suffer from the following three limitations: (i) they are based on greedy algorithms, which may not produce optimal clusters, (ii) these tools are very sensitive to the sequence similarity threshold defining clusters, and (iii) the selection of the sequence representing the center of a cluster is not optimal.

Widely-used tools for sequence clustering—CD-HIT, UCLUST, and DNACLUST—depend on greedy algorithms, which are not guaranteed to find the optimal solution. Given the importance of sequence clustering in the field of computational biology, we propose a much more advanced approach. The mean shift algorithm is a general-purpose optimization technique (16), which has been widely applied in image processing and computer vision (17–19). Unlike the related greedy approaches, the mean shift algorithm is ‘guaranteed’ to converge to local optimal points, e.g. a center of a cluster. Although this algorithm has been applied successfully thousands of times in other fields, it has been applied only few times in the field of bioinformatics (20–22). Here, we propose a novel software tool, MeShClust, that utilizes the mean shift algorithm in clustering nucleotide sequences. Further, our adaptation of the algorithm utilizes a novel classifier to predict the identity score using four alignment-free sequence similarity measures.

A problem with current clustering applications is that the underlying sequence similarity that separates clusters is often unknown; therefore, a biologist may have to guess an identity score to provide to the clustering tool. If wrong, this guessed score limits the quality of the predicted clusters remarkably. For example, if the provided identity score was higher than the true identity score, a tool would produce smaller clusters; if it was much lower, a tool would produce larger clusters. In both situations, the predicted clusters do not match the real clusters. Popular tools do not account for this discrepancy.

*To whom correspondence should be addressed. Tel: +1 918 6313645; Fax: +1 918 6312927; Email: hani-girgis@utulsa.edu

Another limitation of the available tools is that the selection of the sequence representing the center of a cluster is not necessarily optimal. In these algorithms, a sequence that does not belong to any cluster is considered the center of a new cluster. Once a center is selected, it does not change. To illustrate, if the center sequence is at the periphery of the real cluster, then the predicted cluster is very likely to be a partial cluster.

Because the core of MeShClust is the mean shift algorithm, it overcomes these limitations. Specifically, the mean shift algorithm is very likely to produce optimal sequence clusters in practice. MeShClust is flexible and is capable of correcting the provided identity score to a great extent. The sequence representing a cluster *does* change, moving toward the true center of the cluster. MeShClust provides a stable clustering algorithm that is not very sensitive to the sequence similarity parameter and provides greater accuracy than its counterparts.

MATERIALS AND METHODS

Overview

Algorithms 1 and 2 give an overview of the methods underlying our software tool, MeShClust. The software consists of these two components: (i) a classifier and (ii) the mean shift algorithm.

Algorithm 1 An overview of the algorithm implemented in MeShClust

Input: A set of n nucleotide sequences $S = \{s_1, s_2, \dots, s_n\}$ sorted by decreasing length

Output: Clusters of sequences and their respective centers

If the identity is above 60%, train the classifier using a subset of S to recognize similar sequences

$center_{cur} \leftarrow s_1$
 $cluster_{cur} \leftarrow \{s_1\}$

while S is nonempty **do**

$G \leftarrow$ all sequences from S close to $center_{cur}$
 $S \leftarrow S - G$

if G is nonempty **then**

$cluster_{cur} \leftarrow cluster_{cur} \cup G$

 Run *MeanShift*($center_{cur}, cluster_{cur}$) to update $center_{cur}$ (Algorithm 2)

else

 Add $center_{cur}$ to $Centers$

 Add $cluster_{cur}$ to $Clusters$

$center_{cur} \leftarrow$ the closest sequence in S to the old $center_{cur}$ according to the Czekanowski similarity (Equation 4) or the identity score if the alignment algorithm is used instead of the classifier

$cluster_{cur} \leftarrow \{center_{cur}\}$

end if

end while

for $i = 1$ to 15 **do**

for all $center_j \in Centers, cluster_j \in Clusters$ **do**

 Run *MeanShift*($center_j, cluster_j$ along with neighboring clusters) (Algorithm 2)

end for

for $j = 1$ to $|Centers|$ **do**

for all $center_k \in Centers$ close to $center_j$ **do**

 Merge centers $center_j$ and $center_k$ if the sequences representing $center_j$ and $center_k$ are similar

end for

end for

end for

The classifier predicts whether or not two sequences are similar to each other. The similarity is measured as the identity score based on the global alignment of the two sequences (23,24). Sequences are represented as histograms of counts of short words in the sequences. The classifier predicts the identity score due to global alignment (23,24) by calculating a weighted sum consisting of few alignment-free similarity measures using a General Linear Model (GLM).

Algorithm 2 Sequence clustering using the mean shift

Input: The current center, $center_{cur}$, of a cluster, and a set of points, X

Output: The closest point in X to the new center

Calculate the new center, $center_{new}$, using the current center, $center_{cur}$, according to Equation 1

$$center_{new} = \frac{\sum_{i=1}^n D(x_i, center_{cur}) \cdot x_i}{\sum_{i=1}^n D(x_i, center_{cur})} \quad (1)$$

Here, D is the classifier or the alignment algorithm if the classifier is not used (Equation 2)

$$D(x, q) = \begin{cases} 1 & x \text{ is close to the query sequence, } q \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

Next, the closest point $p \in X$ to $center_{new}$ is found using the Czekanowski similarity (Equation 4) or the identity score if the classifier is not used, and p is returned as the new center

A novel adaptation of the mean shift algorithm (16) is the core of the second component. Similar to the classifier, the mean shift algorithm processes the histograms of the input sequences. The mean shift is an iterative, gradient-ascent algorithm that is capable of finding local optimal points. In this adaptation of the algorithm, a local maximum represents the center of a cluster of sequences. In each iteration, a center is recalculated as the weighted *mean* of histograms. This weighted mean is calculated only from the sequences that are similar to the center of a cluster. Similar sequences are determined by the classifier or, if the identity score is <60%, they are determined by the alignment algorithm. Once updated, a center will *shift* toward a local maximum. As these centers move, some of them converge to the same local maximum; therefore, the algorithm merges them. For this reason, the user does not need to specify the number of centers as opposed to other clustering algorithms such as k-means based applications. Once the algorithm converges, sequences that contributed to the calculation of a center are considered members of its cluster. Supplementary Data Sets 1–3 contains the source code and the executables of MeShClust.

Next, we give the details of each step of the algorithm. First, we describe how a sequence is represented as a k-mer histogram. Second, the details of the classifier are given. In the third step, the initial clusters are formed. We illustrate the construction of the final clusters in the fourth step.

Representing a sequence as a histogram of k-mers

A sequence consists of the nucleotides: A, C, G and T (or U). A k -mer is a short subsequence of length k . For example, AAA, AAC, AAT and AAG are tri-mers. To construct a histogram from a sequence, A, C, G and T are converted to

0, 1, 2 and 3, so a k -mer is built as a quaternary number of k digits. Horner's rule can be used for calculating the quaternary numbers of a long sequence efficiently (25). The count of a k -mer in the histogram is initialized to 1 instead of 0; these pseudocounts are needed to allow events that 'seem' impossible to be able to happen (26). For example, k -mers that are absent from one sequence could be present in another. Pseudocounts are important while calculating conditional probabilities. The transformation from nucleotide sequences to k -mer histograms allows for fast, alignment-free, statistical measures to be used in comparisons.

The selection of this k parameter depends on the size of the input sequences. MeShClust automatically computes the k by first taking the \log_4 of the average sequence length, then by subtracting 1. We empirically found that this formula preserves enough information to accurately determine similarity. A smaller k value decreases the amount of memory needed for each histogram and the time required to calculate the alignment-free statistics by a factor of 4 for each nucleotide (27).

Once sequences are converted to k -mer histograms, the classifier is trained in the next stage.

Identifying similar sequences

MeShClust utilizes a classifier to predict similar sequences to a query sequence. The similarity is determined according to an identity score obtained by global alignment (23,24). With regard to a query sequence, similar sequences can be viewed as one class and dissimilar sequences as the other. Therefore, this task can be represented as a classification task. To this end, we used a GLM (28) for classifying these two classes. Classifiers based on GLMs are simple, yet powerful. We have utilized GLMs in other classification tasks successfully (29–32). As a first step, MeShClust samples a roughly equal number of pairs of similar and dissimilar sequences based on a user-defined cutoff. A large number of pairs of sequences is needed to be sampled. Therefore, about 1500 sequence pairs are sampled. Similar sequences are labeled with 1's and the dissimilar sequences with -1's. After that, four features are extracted for each pair of sequences. These four features are selected according to a comprehensive evaluation of alignment-free k -mer statistics (27). The first feature is sequence length difference (Equation 3) \times Czekanowski similarity (Equation 4). Length difference² \times Manhattan distance² (Equation 5) represents the second feature. The third and the fourth features are the Pearson coefficient (Equation 6) and Kulczynski₂² (Equation 7) \times length difference².

$$LD(A, B) = |\text{length}(A) - \text{length}(B)| \quad (3)$$

$$\text{Czekanowski}(A, B) = \sum_{i=0}^N \frac{\min(A_i, B_i)}{A_i + B_i} \quad (4)$$

$$\text{Manhattan}(A, B) = \sum_{i=0}^N |A_i - B_i| \quad (5)$$

$$\text{Pearson}(A, B) = \frac{\sum_{i=0}^N (A_i - \bar{A})(B_i - \bar{B})}{\sqrt{\sum_{i=0}^N (A_i - \bar{A})^2} \sqrt{\sum_{i=0}^N (B_i - \bar{B})^2}} \quad (6)$$

$$\text{Kulczynski}_2(A, B) = \frac{4^k \times (\bar{A} + \bar{B})}{2\bar{A}\bar{B}} \sum_{i=0}^N \min(A_i, B_i) \quad (7)$$

A and B are the two histograms representing two sequences; A_i and B_i are the counts of the i th k -mer in A and B ; \bar{A} and \bar{B} are the average counts of histograms A and B . Next, the four features are scaled between 0 and 1 and converted to similarity measures if necessary by subtracting the scaled value from 1. Then MeShClust utilizes an incremental, automatic process to train the GLM. Before training, the data is divided into two sets of equal sizes. Each set has roughly equal points representing similar and dissimilar sequence pairs. One set is used for training the classifier and the other is used for testing it. First, it trains the GLM using the first two features. If the accuracy calculated on the testing set is at least 97.5%, the training process is finished. Otherwise, it continues by adding another feature followed by evaluating the testing accuracy. This accuracy is measured by the average of the true positive rate (sensitivity) and the true negative rate (specificity). Once trained, the GLM is used for predicting sequences similar to a query sequence. We have observed that if the similarity is under 60%, it is not currently possible to accurately classify sequences using alignment free statistics, so alignment is used in those cases.

Efficient data storage

Because the algorithm repeatedly selects and removes similar sequences to a query sequence, the largest time bottleneck was finding and removing multiple sequences from the input list. As a remedy, a data structure similar to a separate chaining hash table was implemented that allowed for faster retrieval and deletion of elements. This data structure consists of a list of smaller bins which specify a range of sequence lengths. Therefore, searching for similar sequences within a certain length range only affects a few bins when sequences are removed from the data structure. Using this data structure, initial centers can be found with relative ease, as discussed in the next section.

Finding initial clusters

MeShClust aims at clustering the input sequences into distinct groups. Figure 1 diagrams this step. MeShClust gathers similar sequences into initial clusters. To start, input sequences are sorted based on length. The shortest sequence is the center of the first initial cluster. Then the classifier or the alignment algorithm, if applicable, is used for finding similar sequences to that center. After that, the mean shift is applied on all sequences in the current cluster to calculate the updated mean. Next, the sequence closest to the new mean becomes the center of the cluster. Therefore, at each iteration of the algorithm, a better representative center of the cluster is found. This new center is used for the addition of more similar sequences to the cluster. This step is repeated until no similar sequences are left. At this point, the current

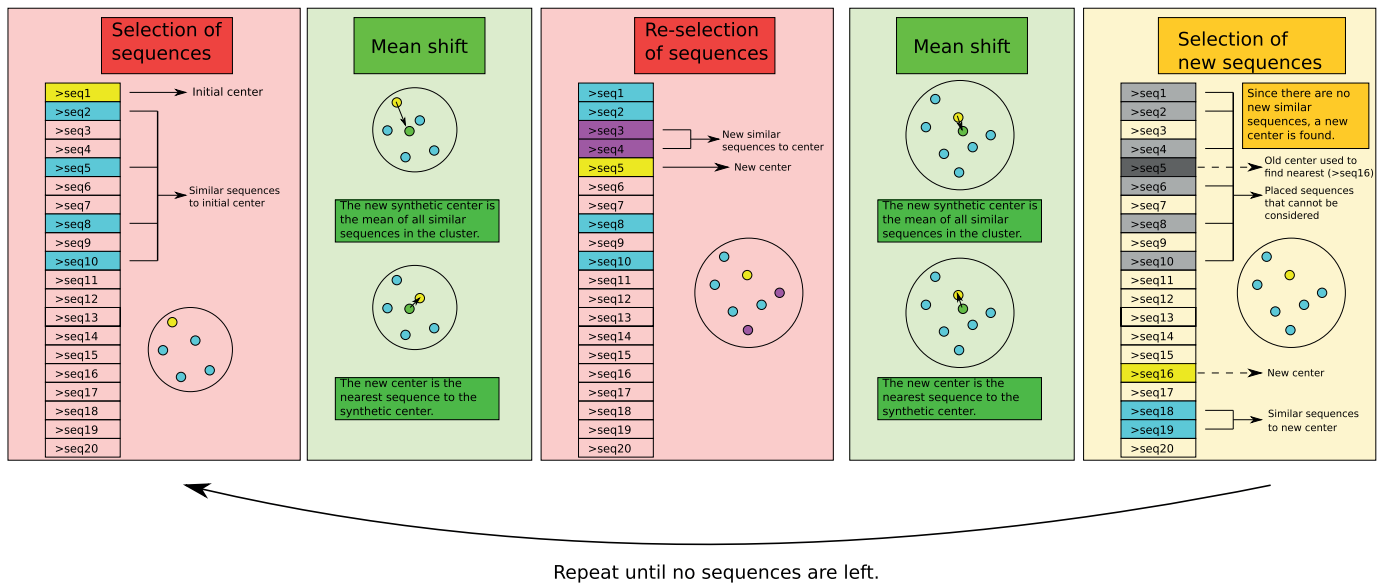


Figure 1. Overview of the initialization step of MeShClust, the first part of Algorithm 1. Before running this step the input sequences are sorted according to their lengths. The center of the first initial cluster is the shortest sequence in the input. The diagram shows the iterative process of finding initial clusters of sequences, getting better clusters from the selected sequences, and then finding the next closest cluster.

cluster is set aside, and a new cluster is formed using the next closest sequence to the last center. The selection of the next center improves clustering by producing a semi-sorted list of sequences; neighboring clusters that may be merged later are grouped near each other. In effect, the combination of using a binary classifier and running the mean shift represents a ‘flat kernel’ (16), except it only considers sequences not already placed in the initial clusters.

After grouping the sequences into initial clusters, the mean shift is used once again for forming the final clusters.

Clustering

In the previous step, the classifier only considers the unplaced sequences; therefore, some of the initial clusters may have similar sequences in the already placed clusters. Further, the center of a cluster is updated in the initialization step by considering its sequences only. In this step, unlike the initial clustering, the mean shift considers sequences in neighboring clusters. These neighboring clusters include five clusters above and below the current cluster. Recall that these clusters are placed in a semi-sorted list. Centers that are close to each other, as determined by the classifier or by the alignment algorithm, are merged. If two centers are merged, the sequences that belong to each center are also merged.

RESULTS

We start with defining multiple evaluation measures in order to evaluate MeShClust. These measures are intended to evaluate the quality of the predicted clusters as well as the time and the memory requirements of each software tool. After that, we discuss the results of comparing five sequence clustering software tools, including MeShClust.

Evaluation criteria

We applied the following seven evaluation measures in evaluating MeShClust and four related tools: (i) intra-cluster similarity, (ii) inter-cluster similarity, (iii) silhouette score, (iv) purity, (v) normalized mutual information (NMI), (vi) time requirement and (vii) memory usage. Only clusters of at least five sequences are considered in our analysis of the first three measures, except when it is indicated otherwise. The intra-cluster similarity is the average similarity between the sequence representing the center of a cluster and the other sequences in the same cluster. Sequence similarity is determined by calculating the identity score. To measure the dissimilarity between different clusters, we applied the inter-cluster similarity measure, which is the average similarity between different centers. Our third criterion is a variant of the silhouette score (33). This measure compares the suitability of placing a sequence in its current cluster to the suitability of placing this sequence in the closest cluster. We define $d_c(s)$ as the distance between the sequence s and the sequence representing its own cluster and $d_n(s)$ as the distance between s and the sequence representing the closest ‘neighboring’ cluster. The distance is measured by subtracting the identity score from 100. Equation (8) defines the silhouette score.

$$\text{Silhouette} = \frac{1}{M} \sum_{s \in \text{sequences}} \frac{d_n(s) - d_c(s)}{\max(d_n(s), d_c(s))} \quad (8)$$

Here, M is the number of sequences in the set. The silhouette score ranges between 1 and -1 ; when the score is high, clusters are tighter and more separated from each other. A perceived problem with the silhouette score is that clusters consisting of a single element always have the value 1. Therefore, when the silhouette is very high, single element clusters may be inflating that value. A possible remedy for this issue is the redefinition of a cluster as a collection of

at least five elements. However, a good silhouette value may not match actual—according to the ground truth—clusters. The silhouette is only a measure for combined separation and tightness, not correctness.

Our fourth and fifth evaluation measures do measure correctness. They are purity (Equation 9) (34) and NMI (Equation 12) (34), both of which are applied when the true clusters are known. These measures compare all clusters found, F , to the actual clusters, A , in a set of N sequences.

Purity (Equation 9) measures how mixed each cluster is; if a predicted cluster only includes items from one real cluster, the purity is high. A disadvantage of this measure is that if every cluster is a single element, the purity will be 1.

$$\text{purity}(F; A) = \frac{1}{N} \sum_{f \in F} \max_{a \in A} |f \cap a| \quad (9)$$

NMI corrects this by considering the probability that F and A contain the same data. To make these values comparable, this mutual information is normalized by the average entropy of F and of A . Equations (10–12) describe mutual information, entropy, and NMI.

$$I(F; A) = \sum_{f \in F} \sum_{a \in A} \frac{|f \cap a|}{N} \log \left(\frac{N|f \cap a|}{|f||a|} \right) \quad (10)$$

$$H(X) = - \sum_{x \in X} \frac{|x|}{N} \log \left(\frac{|x|}{N} \right) \quad (11)$$

$$\text{NMI}(F; A) = \frac{I(F; A)}{[H(F) + H(A)]/2} \quad (12)$$

Next, we apply these evaluation measures to assessing the performance of MeShClust and four widely used tools.

Comparison to related tools

We utilized both synthetic and real data sets in comparing the performance of MeShClust to the performances of UCLUST (13), CD-HIT (3,12), DNACLUSt (9) and wcdst (11), which are widely used clustering tools. All related tools were ran using the default values for all parameters except the sequence identity parameter, for which we tried multiple values. However, wcdst does not have an identity parameter; the authors of wcdst provide an example parameter values equivalent to 96% identity, given as '-N 4 -l 100 -T 40 -H 72 -show_clusters'. Therefore, we compared to wcdst only using the 96% identity when applicable. All other tools were also evaluated at 96% for comparison. Evaluations were done on a Dell OptiPlex 990 with 4 Intel i5-2500 processors running GNU/Linux (CentOS 7).

Evaluation on synthetic data sets

As a start, we generated three data sets, which we call the 10%, the 25%, and the 100-centers sets (Supplementary Data Set 4–6). The three synthetic data sets were generated at 10%, 25%, and 25% mutation rates, respectively. Each of the 10% and the 25% data sets has 10 clusters, each of which consists of about 25 sequences. The 100-centers set has 100

clusters each of size 10. The length of a sequence is 1000 base pairs (bp) approximately.

To test the tolerance of the five tools to the identity parameter, we ran the tools using five identity scores (up to 10% above and below the actual identity score). Two of these identity scores are above the real identity score, and two are below it, and one approximately matches the identity score used for generating the clusters. These identities were selected to demonstrate to what degree a tool is tolerant to an inaccurate identity score. Table 1 shows the results of evaluating the four tools on the 10% data set using these identity scores: 75%, 80%, 85%, 90% and 95%. The true identity score is between 85% and 90%.

When the true clusters are known, the performance is best measured by purity and NMI. On most synthetic data sets, the purity values achieved by all tools were 1, and the lowest value found was above 0.9. Thus, NMI becomes the de-facto measure. On four out of the five tests, MeShClust obtained perfect NMI and purity scores. On those tests, MeShClust found perfect clusters, even though the identity parameter was inaccurate. Among the three related tools, UCLUST obtained almost perfect NMI score in one test only. CD-HIT and DNACLUSt did not obtain perfect or close to perfect results on any of the five tests. Further, MeShClust achieved much better results in terms of the silhouette and the intra-cluster similarity scores. Compared to MeShClust's intra-cluster score of 90 and silhouette of 0.81, the next highest of any tool on any identity was UCLUST with 84 similarity on the 80% identity and UCLUST with a silhouette of 0.66 on the 75% identity. Similar results were obtained on the 25% and the 100-centers data sets (see Supplementary Table S1). These results demonstrate that MeShClust is tolerant to inaccurate identity score. This tolerance is evident by the consistency of the high quality clusters obtained on different synthetic data sets at different identity levels.

Evaluation on a comprehensive microbiome

Next, we aimed at evaluating the tools on real data; therefore, we obtained sequences from a microbiome study (35). We call this set the Costello set. About 1.1 million sequences comprise this set. Sequences in this set range between 200 and 400 bp. Before evaluating the tools on the Costello set, we generated similar, smaller synthetic sets because the real clusters are unknown. We generated the 15K and the 150K sets consisting of 15 000 and 150 000 sequences, respectively (Supplementary Data Sets 7 and 8). A synthetic cluster contains around 75 sequences, forming 200 and 2000 actual clusters, respectively. These clusters were generated using 3% mutation rate; however, the actual mutation rate is usually higher than 3% due to randomization.

As before, we evaluated the tools using the following identity scores: 83%, 87%, 90%, 93%, 96% and 97% (Supplementary Table S2). On the 15K set, MeShClust obtained perfect or almost perfect NMI in six tests, demonstrating its ability to find the real clusters even when the sequence similarity parameter is inaccurate. CD-HIT obtained perfect or almost perfect NMI in four tests, whereas UCLUST obtained its best, 0.98 and 0.95 NMI, in two tests. DNACLUSt performed well in one test only, its best

Table 1. Comparison of the performances of MeShClust, UCLUST, CD-HIT and DNACLUST on the 10%-mutation-rate synthetic data set

Tool	Identity (%)	NMI	Intra (%)	Inter (%)	Silhouette (%)	Real time (min:sec)	Max memory (MB)
CD-HIT	75	NA	NA	NA	NA	NA	NA
DNACLUST	75	0.60	-	-	-	0:04.6	392
MeShClust	75	1.00	90	48	0.81	0:07.7	79
UCLUST	75	0.98	81	48	0.66	0:00.2	4
CD-HIT	80	0.89	82	48	0.34	0:00.2	35
DNACLUST	80	0.60	-	-	-	0:03.2	392
MeShClust	80	1.00	90	48	0.81	0:07.8	79
UCLUST	80	0.79	84	48	0.40	0:00.3	6
CD-HIT	85	0.83	82	48	0.29	0:00.2	35
DNACLUST	85	0.59	-	-	-	0:02.0	392
MeShClust	85	1.00	90	48	0.81	0:05.8	79
UCLUST	85	0.60	-	-	-	0:00.6	9
CD-HIT	90	0.60	-	-	-	0:01.0	39
DNACLUST	90	0.59	-	-	-	0:01.1	392
MeShClust	90	1.00	90	48	0.81	0:05.6	79
UCLUST	90	0.60	-	-	-	0:00.3	9
CD-HIT	95	0.60	-	-	-	0:00.1	39
DNACLUST	95	0.59	-	-	-	0:00.5	392
MeShClust	95	0.59	100	46	1.00	0:05.5	80
UCLUST	95	0.59	-	-	-	0:00.2	9

The 10% data set contains 10 clusters, each of which is generated by mutating approximately 10% of the bases comprising a template sequence. Each cluster consists of around 25 sequences. The silhouette score ranges between -1 and 1. MeShClust was the only tool that was capable of finding perfect clusters according to the NMI in four tests, whereas UCLUST succeeded in finding almost correct clusters in one test only. CD-HIT is not designed to cluster sequences that have identity scores less than 80%; therefore we list results of CD-HIT using 75% identity as Not Applicable (NA). Since a cluster is considered when there are at least five sequences, a '-' indicates that the software tool did not produce any clusters with at least five sequences to measure.

NMI was 0.89. When the identity threshold is 96%, wcdet has 0.88 NMI, outperforming all other tools except MeShClust, which achieved nearly 1.00 NMI. With respect to time required, MeShClust takes more time than UCLUST or CD-HIT, but MeShClust is faster than DNACLUST and wcdet. In terms of memory use, MeShClust takes more memory than UCLUST, CD-HIT, and wcdet, but it uses only ~90MB, which is readily available on modern hardware. MeShClust has the highest silhouette scores except for a few cases where DNACLUST or wcdet has higher values. However, DNACLUST at 97% only had 0.2% of its clusters containing at least five sequences. Thus, DNACLUST's silhouette of 0.99 is due to a small number of clusters because the majority of its clusters are too small. A similar phenomenon happens with wcdet.

MeShClust achieved similar, even better, results on the 150k data set (Supplementary Table S2). Specifically, it obtained an NMI of 1 in all but one test, where it has NMI of 0.98, outperforming all related tools on this data set. The second best performance was achieved by CD-HIT, which scored perfect or almost perfect NMI in four tests. The silhouette scores obtained by wcdet and MeShClust were comparable (0.94 versus 0.93). Even though wcdet had a slightly higher silhouette score, again wcdet only had 0.6% of its clusters with at least five sequences; therefore, its silhouette of 0.94 is due to 0.6% only of wcdet's output. In contrast, MeShClust had >95% of its clusters containing at least 5 sequences except on the 97% identity test. Interestingly, the clusters produced by MeShClust were stable across all identity cutoffs demonstrated by almost the same values of NMI, silhouette, intra-clustering similarity, and the inter-clustering similarity.

After that, we evaluated the five tools on the Costello set (Table 2 and Supplementary Table S2), which consists of

1.1 million real sequences. MeShClust obtained higher silhouette and intra-cluster similarity scores than the four related tools in the six tests. With respect to the time required, MeShClust took 8–17 min to cluster the Costello set. On average, MeShClust takes 20.4% longer time than the other tools, and ~6 min more than UCLUST. However, on computers with many CPU cores, this time gap will be lowered because of MeShClust's parallel algorithm. Regarding the memory usage, MeShClust takes a manageable amount of memory of ~1GB, which is available on almost all personal computers. Concerning the quality of the clusters, MeShClust has the highest silhouette of any tool besides wcdet, averaging 0.19 higher silhouette than the next highest silhouette on the Costello data set (excluding wcdet). However, the percent of clusters produced by wcdet at or above five sequences was 2.4%, implying that wcdet missed the majority of the real clusters. In contrast, all the other tools had at least 72 times as many clusters at or above five sequences as wcdet. These results show that the clusters produced by MeShClust are more separable and more compact than those produced by the related tools.

Evaluation on viral genomes

Entire viral genomes (36) were used for testing MeShClust on much longer sequences, with sequence lengths averaging 6625 base pairs and with some sequences over 13 000 bp. The data set contained seven different families or genera, totaling 96 sequences (Table 3). Genomes from genera Badnavirus, Caulimovirus, Reptarenavirus, Soymovirus and Spumavirus were used, and genomes from families Amalgaviridae and Birnaviridae were combined to create this data set. Since the taxonomic data was known, the real clusters were created by grouping every genera/family into

Table 2. Comparison of the performances of the five tools on the Costello data

Tool	Identity (%)	Silhouette	Intra (%)	Inter (%)	Real time (min:sec)	Max memory (MB)
CD-HIT	83	0.17	82	64	10:25.1	491
DNACLUST	83	0.17	82	63	2:26.5	1127
MeShClust	83	0.45	89	65	8:23.3	910
UCLUST	83	0.27	86	65	2:24.6	412
CD-HIT	87	0.21	86	65	10:28.0	492
DNACLUST	87	0.22	86	64	3:30.1	1123
MeShClust	87	0.42	90	64	8:49.7	916
UCLUST	87	0.27	88	66	3:35.5	413
CD-HIT	90	0.20	88	66	12:03.7	494
DNACLUST	90	0.20	88	65	4:23.9	1120
MeShClust	90	0.44	92	65	9:08.7	922
UCLUST	90	0.26	90	66	1:09.4	412
CD-HIT	93	0.18	90	66	9:18.6	498
DNACLUST	93	0.19	91	66	1:33.5	1126
MeShClust	93	0.45	93	66	11:41.9	968
UCLUST	93	0.20	92	67	8:05.9	412
CD-HIT	96	0.21	93	67	14:34.4	510
DNACLUST	96	0.18	94	67	4:28.9	1123
MeShClust	96	0.48	95	66	11:56.1	997
UCLUST	96	0.16	95	67	2:09.6	412
WCDEST	96	0.68	91	64	39:11.7	205
CD-HIT	97	0.16	94	67	17:30.7	521
DNACLUST	97	0.16	95	67	2:48.1	1123
MeShClust	97	0.50	96	67	16:58.2	1029
UCLUST	97	0.06	96	68	7:44.1	416

The Costello data set was obtained from a popular microbiome study (35). About 1.1 million sequences comprise this data set. Sequence lengths range between 200 and 400 base pairs. The performances were compared using these six identity scores: 83%, 87%, 90%, 93%, 96% and 97%. MeShClust achieved the highest Silhouette scores and the highest intra-cluster sequence similarity scores in the six tests. Clusters with fewer than 5 elements were not considered for evaluation.

Table 3. The viral data set

Genome	Average length	Sequences in cluster
Complete genome Badnavirus	7574	42
Segment B Birnaviridae	2836	8
Complete genome Spumavirus	12307	6
Complete genome Amalgaviridae	3369	8
Complete genome Caulimovirus	7940	11
Complete genome Soyrovirus	8088	5
Segment S Reptarenavirus	3400	4
Segment A Birnaviridae	3173	8
Segment L Reptarenavirus	6900	4

The viral data set consists of 9 clusters and 96 viral genomes. Viral genomes were obtained from viruSITE (36).

a cluster. Therefore, purity and NMI values can be calculated for this test.

An analysis of the real clusters showed an intra-cluster similarity of 59% and an inter-cluster similarity of 40%. Therefore, identities 43%, 47%, 50%, 53%, and 57% were used for comparing MeShClust and UCLUST. These identities were too low for CD-HIT or DNACLUST; therefore they were not included in this test.

Table 4 shows the results on the viral set. MeShClust and UCLUST had very similar purity values over the viral data sets. However, MeShClust outperformed UCLUST on those tests via NMI by an average difference of 0.25 on a scale from 0 to 1, while using similar memory requirements and a manageable amount of time.

On the 53% test, MeShClust found an NMI of 0.87 and a silhouette of 0.59, while UCLUST found an NMI of 0.60

and a silhouette of 0.84. The silhouette of the real clusters is 0.27. Thus, a higher silhouette score does not directly imply having more accurate clusters in this case. Using a cut-off of 53%, UCLUST found 55 clusters, each averaging 2 sequences, whereas MeShClust found 19 clusters, each averaging 5; this data set consists of 9 real clusters (Genomes of Reptarenavirus and Birdnavirus are broken up into two segments). These statistics can be found in Supplementary Table S3. Because UCLUST's median cluster size is 1, many of its sequences will get a silhouette value of 1 and very high intra cluster values. Therefore, a higher silhouette score does not indicate better clusters in this case. However, MeShClust's clusters overlap with the real clusters more than the ones produced by UCLUST, evident by higher NMI.

DISCUSSION

Added benefits of the classifier

Recall that when the identity cutoff is greater than 60%, MeShClust uses a classifier to predict the identity score from a combination of alignment-free k-mer statistics. To assess the added benefits of the classifier over the alignment algorithm, we evaluated the performances of the two versions of MeShClust on the 15k data set. As expected, the alignment-only version is 16 minutes (198 times) slower, on average, than the classifier-based version (Supplementary Table S4). Using only alignment does not give higher quality clusters; in no case does the alignment-only version have higher NMI values than the classifier-based version. These results demonstrate the remarkable time reduction due to the classifier.

Table 4. MeShClust and UCLUST were used for clustering viral genomes

Tool	Identity (%)	Purity	NMI	Intra (%)	Inter (%)	Silhouette	Real time (min:sec)	Max memory (MB)
MeShClust	43	0.57	0.58	52	50	0.30	0:25.0	81
UCLUST	43	0.47	0.00	38	-	0.38	0:11.1	72
MeShClust	47	0.67	0.68	55	45	0.27	0:31.3	81
UCLUST	47	0.70	0.49	63	43	0.30	0:19.0	74
MeShClust	50	0.92	0.87	72	39	0.50	1:28.1	82
UCLUST	50	1.00	0.73	91	38	0.83	0:32.1	76
MeShClust	53	1.00	0.87	78	37	0.59	2:48.8	82
UCLUST	53	1.00	0.60	93	38	0.84	0:34.1	77
MeShClust	57	1.00	0.69	91	38	0.82	6:15.5	82
UCLUST	57	1.00	0.58	96	39	0.91	0:35.6	78
Ground truth	NA	1.00	1.00	59	40	0.27	NA	NA

The viral data set containing 96 genomes was obtained from viruSITE (36). Since the clusters were so small, the intra, inter, and silhouette considered any cluster, not just clusters of at least 5 elements. The purity was not 1 all around, getting as low as 0.47. Overall, MeShClust had similar purity values; however, MeShClust had 61% higher NMI values than UCLUST, on average. On the 43% data set, UCLUST only found 1 cluster; therefore, its inter-cluster distance could not be measured (labeled as '-') and its NMI was 0.

Related algorithms

DBSCAN (37) is a 'density-based' clustering algorithm. Like the mean shift, it is based upon the principle that clusters consist of densely-packed points. Both the mean shift and DBSCAN can discover the number of clusters on their own. DBSCAN depends on two parameters, the minimum number of close points and the minimum distance between close points. In contrast, the mean shift algorithm depends on only one parameter, the cluster bandwidth, which is the distance among points in the same cluster. The average time of DBSCAN is $\mathcal{O}(n \log(n))$ when using an efficient data structure such as R* tree. However, this degrades into a quadratic time if an efficient data structure is not used. Similarly, the mean shift is a quadratic algorithm in theory because each point represents the center of an initial cluster. However, using a reasonable similarity threshold, our implementation should take $\mathcal{O}(mn)$, where m is the number of clusters in the data set. This run time may degenerate to quadratic time if the similarity threshold is too high. DBSCAN considers adding a new point to a cluster if it is close to a at least a minimum number of points in the cluster. This new point and the points close to it are added to the cluster. For this reason, DBSCAN can find clusters of many different shapes. Overall, the structure of the algorithm is similar in nature to MeShClust's implementation of the mean shift, except MeShClust uses the average histogram as the new center; points that are close to the average are added to the cluster.

MCL (38,39) is a graph based clustering algorithm. It uses random Markov walks to collect 'flow' in a graph, that is, if a random walk is performed in a graph, most likely that walk will end in the same cluster as it started. Using an all-vs-all adjacency matrix, MCL finds clusters by converting the distances into transition probabilities, i.e. constructing a stochastic matrix, in which each column sums to 1. The core of this algorithm is a process involving squaring the stochastic matrix. An entry in the newly-squared matrix represents the probability of reaching a node from another directly or via one more node. In other words, squaring the matrix simulates a step along another edge. This process is performed until each element is either 0 or 1. Although it consists of

matrix operations, efficient structures such as sparse matrices allow for fast computation. Using a BLAST-style input, the alignments can be pre-computed for reuse, and the parameters are determined by BLAST, not by the algorithm itself. In sum, MCL uses a different approach that is based on graph theory, whereas the mean shift is based on density-based optimization.

Convergence

We have developed a test to assess the convergence of the mean shift algorithm. If all centers do not change from the previous iteration, convergence is achieved. On the small, synthetic data sets, convergence was immediate; nullifying the effect of the second part of the mean shift algorithm. On larger data sets, such as the 15k data set (Supplementary Data Set 7), the second part of the algorithm, which involves the merging step, is needed. These merges would sometimes be done in the first few iterations—15 iterations are the default—of the algorithm. In 37 out of 100 tests performed, the algorithm converged in a few iterations. However, oscillations occurred in 63 out of the 100 tests, making convergence impossible automatically. These oscillations were sequences that were considered close to a center, but when the new center was calculated from the close points, a new point displaces the old one, which will change the center in the next iteration. Out of the 63 times that oscillations occurred on the 15k data set (15 000 sequences), 33 of them had 1 center oscillating, 26 trials had 2 centers and 4 trials had 3 centers. When the number of threads was limited to 1, no oscillations occurred, and the algorithm converged in a few iterations. Therefore, the limit of 15 iterations is enough for the algorithm to converge.

CONCLUSION

DNA sequence clustering algorithms have many applications. Nonetheless, the widely applicable tools depend on greedy algorithms, which do not necessarily produce the best results. Further, the related tools are sensitive to the sequence similarity parameter provided by the user. Often

times, the exact value of this parameter is not known, resulting in inaccurate clusters. Our clustering software, MeShClust, is a novel tool that utilizes the mean shift algorithm. MeShClust is the first application of the mean shift to clustering DNA sequences and one of few applications of the mean shift algorithm in bioinformatics. Unlike the related greedy approaches, the theory behind the mean shift guarantees convergence to local optimal points, resulting in higher quality clusters. Further, most sequence clustering tools use a slow quadratic algorithm for sequence alignment. In contrast, MeShClust reduces the dependency on alignment algorithms by using a novel, machine-learning-based, alignment-assisted method for computing sequence similarity. Furthermore, this is the first attempt to formulate the task of identifying similar sequences as a classification task. When tested on multiple synthetic and real data sets, MeShClust outperformed the related tools with a clear margin, advancing the methodology in the field of sequence analysis.

SUPPLEMENTARY DATA

[Supplementary Data](#) are available at NAR Online.

ACKNOWLEDGEMENTS

We are thankful to the anonymous reviewers for taking the time to review this manuscript. Their comments and suggestions have improved the software and the manuscript. We would like to thank Joseph Valencia and Robert Geraghty for their help with coding the GLM and the alignment algorithm.

FUNDING

College of Engineering and Natural Sciences and the Faculty Research Grant Program at the University of Tulsa; Oklahoma Center for the Advancement of Science and Technology [PS17-015]. Funding for open access charge: The University of Tulsa and OCAST.

Conflict of interest statement. None declared.

REFERENCES

- Bao, E., Jiang, T., Kaloshian, I. and Girke, T. (2011) SEED: efficient clustering of next-generation sequences. *Bioinformatics*, **27**, 2502.
- Chong, Z., Ruan, J. and Wu, C.-I. (2012) Rainbow: an integrated tool for efficient clustering and assembling RAD-seq reads. *Bioinformatics*, **28**, 2732.
- Fu, L., Niu, B., Zhu, Z., Wu, S. and Li, W. (2012) CD-HIT: accelerated for clustering the next-generation sequencing data. *Bioinformatics*, **28**, 3150.
- Comin, M., Leoni, A. and Schmid, M. (2015) Clustering of reads with alignment-free measures and quality values. *Algorithms Mol. Biol.*, **10**, 4.
- Shimizu, K. and Tsuda, K. (2011) SlideSort: all pairs similarity search for short reads. *Bioinformatics*, **27**, 464.
- Solovyov, A. and Lipkin, W.I. (2013) Centroid based clustering of high throughput sequencing reads based on n-mer counts. *BMC Bioinformatics*, **14**, 268.
- Warren, R.L., Sutton, G.G., Jones, S. J.M. and Holt, R.A. (2007) Assembling millions of short DNA sequences using SSAKE. *Bioinformatics*, **23**, 500.
- Zorita, E., Cusc, P. and Fillion, G.J. (2015) Starcode: sequence clustering based on all-pairs search. *Bioinformatics*, **31**, 1913.
- Ghods, M., Liu, B. and Pop, M. (2011) DNACLUSt: accurate and efficient clustering of phylogenetic marker genes. *BMC Bioinformatics*, **12**, 271.
- Burke, J., Davison, D. and Hide, W. (1999) d2_cluster: a validated method for clustering EST and full-length cDNA sequences. *Genome Res.*, **9**, 1135–1142.
- Hazelhurst, S., Hide, W., Lipták, Z., Nogueira, R. and Starfield, R. (2008) An overview of the wcd EST clustering tool. *Bioinformatics*, **24**, 1542–1546.
- Li, W. and Godzik, A. (2006) Cd-hit: a fast program for clustering and comparing large sets of protein or nucleotide sequences. *Bioinformatics*, **22**, 1658.
- Edgar, R.C. (2010) Search and clustering orders of magnitude faster than BLAST. *Bioinformatics*, **26**, 2460–2461.
- Wei, D., Jiang, Q., Wei, Y. and Wang, S. (2012) A novel hierarchical clustering algorithm for gene sequences. *BMC Bioinformatics*, **13**, 174.
- Liao, W., Ren, J., Wang, K., Wang, S., Zeng, F., Wang, Y. and Sun, F. (2016) Alignment-free transcriptomic and metatranscriptomic comparison using sequencing signatures with variable length Markov chains. *Sci. Rep.*, **6**, 37243.
- Cheng, Y. (1995) Mean shift, mode seeking, and clustering. *IEEE Trans. Pattern Anal. Mach. Intell.*, **17**, 790–799.
- Comaniciu, D. and Meer, P. (1999) Mean shift analysis and applications. In: *Proc IEEE Int Conf Comput Vis*. Kerkyra, pp. 1197–1203.
- Comaniciu, D. and Meer, P. (2002) Mean shift: a robust approach toward feature space analysis. *IEEE Trans. Pattern Anal. Mach. Intell.*, **24**, 603–619.
- Girgis, H.Z., Mitchell, B.R., Dassopoulos, T., Mullin, G. and Hager, G. (2010) An intelligent system to detect Crohn's disease inflammation in Wireless Capsule Endoscopy videos. In: *Proc IEEE Int Symp Biomed Imaging*. pp. 1373–1376.
- Barash, D. and Comaniciu, D. (2004) Meanshift clustering for DNA microarray analysis. In: *IEEE Computational Systems Bioinformatics Conference*. pp. 578–579.
- Wang, L.-y., Abyzov, A., Korbel, J.O., Snyder, M. and Gerstein, M. (2009) MSB: a mean-shift-based approach for the analysis of structural variation in the genome. *Genome Res.*, **19**, 106–117.
- de, Brito D., Maracaja-Coutinho, V., de, Farias S., Batista, L. and do, Rêgo T. (2016) A novel method to predict genomic islands based on mean shift clustering algorithm. *PLoS ONE*, **11**, e0146352.
- Needleman, S.B. and Wunsch, C.D. (1970) A general method applicable to the search for similarities in the amino acid sequence of two proteins. *J. Mol. Biol.*, **48**, 443–453.
- Gotoh, O. (1982) An improved algorithm for matching biological sequences. *J. Mol. Biol.*, **162**, 705–708.
- Girgis, H.Z. (2015) Red: an intelligent, rapid, accurate tool for detecting repeats de-novo on the genomic scale. *BMC Bioinformatics*, **16**, 227.
- Compeau, P. and Pevzner, P. (2015) *Bioinformatics Algorithms: An Active Learning Approach*. Active Learning Publishers, La Jolla.
- Luczak, B.B., James, B.T. and Girgis, H.Z. (2017) A survey and evaluations of histogram-based statistics in alignment-free sequence comparison. *Brief. Bioinformatics*, doi:10.1093/bib/bbx161.
- McCullagh, P. (1984) Generalized linear models. *Eur. J. Oper. Res.*, **16**, 285–292.
- Girgis, H.Z. (2008) *Machine-learning-based meta approaches to protein structure prediction*, PhD thesis, The State University of New York, Buffalo.
- Girgis, H.Z. and Corso, J.J. (2008) *Stp: the sample-train-predict algorithm and its application to protein structure meta-selection*. Technical Report 16, The State University of New York, Buffalo.
- Girgis, H.Z., Corso, J.J. and Fischer, D. (2009) On-line hierarchy of general linear models for selecting and ranking the best predicted protein structures. *Conf. Proc. IEEE Eng. Med. Biol. Soc.*, **2009**, 4949–4953.
- Girgis, H.Z. and Sheetlin, S.L. (2013) MsDetector: toward a standard computational tool for DNA microsatellites detection. *Nucleic Acids Res.*, **41**, e22.
- Rousseeuw, P.J. (1987) Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *J. Comput. Appl. Math.*, **20**, 53–65.

34. Manning,C.D., Raghavan,P. and Schütze,H. (2008) *Introduction to Information Retrieval*. Cambridge University Press, Cambridge.
35. Costello,E.K., Lauber,C.L., Hamady,M., Fierer,N., Gordon,J.I. and Knight,R. (2009) Bacterial community variation in human body habitats across space and time. *Science*, **326**, 1694–1697.
36. Stano,M., Beke,G. and Klucar,L. (2016) viruSITE - integrated database for viral genomics. *Database*, **2016**, baw162.
37. Ester,M., Kriegel,H.-P., Sander,J. and Xu,X. (1996) *A Density-based Algorithm for Discovering Clusters in Large Spatial Databases with Noise*. AAAI Press, Portland , pp. 226–231.
38. van Dongen,S. (2000) *Graph Clustering by Flow Simulation*, PhD thesis, University of Utrecht.
39. Enright,A.J., Van,Dongen S. and Ouzounis,C.A. (2002) An efficient algorithm for large-scale detection of protein families. *Nucleic Acids Res.*, **30**, 1575–1584.