

USEARCH

Software and documentation
© Copyright 2010-11 Robert C. Edgar
All rights reserved

<http://drive5.com/usearch>

robert@drive5.com

Version 4.0.41
January 20, 2011

Table of Contents

Table of Contents	2
USEARCH.....	6
License	6
Installation.....	6
Executable filename	6
USEARCH overview	7
Basic usage: clustering with UCLUST.....	8
Seed sequences	8
Sorting the input sequences.....	8
Non-redundant database: sort by decreasing length	8
OTU prediction: sort by decreasing abundance	8
The --cluster command.....	9
Output files.....	9
The --usersort option	9
Basic usage: database search with UBLAST.....	10
Database format.....	10
Search command	10
U-sorting.....	10
Number of targets tested	10
Nucleotide, protein and translated searches	10
Output files.....	11
Basic usage: Dereplication, removing identical (sub-)sequences.....	12
Search strategy	13
Database.....	13
Search order	13
Search termination (U-sorted or S-sorted only)	14
Alignment	14
Similarity measures	15
Word counting	15
Sorting sequences for clustering.....	16
Sorting by decreasing length.....	16
Sorting by decreasing abundance.....	16

Searching.....	18
Search order.....	18
Query-target comparison.....	19
Alignment parameters and heuristics	19
Similarity measure.....	19
Search termination	20
Accepts and rejects.....	20
Termination conditions	20
Maximum accepts.....	20
Maximum rejects	20
Maximum targets.....	20
Search and clustering at high identities.....	20
Search and clustering at medium identities.....	21
Search at low identities	21
Output files.....	22
The UCLUST file format	22
UCLUST to FASTA conversion.....	24
BLAST6 format.....	25
BLAST-like alignment format	25
FASTA alignment format.....	26
CD-HIT format.....	26
User-defined output.....	26
UBLASTX: Translated ORF search	30
Frame-shifts.....	30
ORF identification.....	30
Search and output	30
UIRE: hierarchical clustering, clumping and large multiple alignments.....	31
Warning	31
The .hire file format	32
Large multiple alignments.....	33
Parameter tuning	35
Improved sensitivity for distant proteins.....	35
Choose suitable quality measures	35
Quality measures for clustering.....	35

Quality measures for database search.....	35
Construct a query set that is small enough for testing.....	35
Test with increasing values of --maxtargets.....	35
Test with increasing values of --maxaccepts.....	36
U-sort word length	36
Tune alignment heuristics	36
USTAR: Fast multiple alignment of clusters	38
Gap penalties and substitution scores	39
E-value calculation.....	39
Substitution scores for nucleotides.....	39
Substitution matrix for amino acids	39
Gap penalties for local alignments	40
Gap penalties for global alignments.....	40
Considerations when using non-standard gap penalties.....	41
Sequence identity	42
Terminal gaps.....	42
Default definition of identity, --iddef 0.....	42
All-diffs definition, --iddef 1.....	43
Internal diffs definition, --iddef 2.....	43
Marine Biology Lab definition, --iddef 3.....	43
BLAST definition, --iddef 4.....	44
UCHIME: Chimeric sequence detection	45
SNPs and voting.....	45
Summary report format (--report option).....	45
Detailed report format (--reportx option).....	46
SNP classification	46
Advantages of UCHIME.....	46
Limitations of UCHIME	47
Memory requirements.....	48
Reduce redundancy	48
Trim sequence labels.....	48
Use database stepping	48
Split the database.....	48
Two-level search	49

Command line reference	51
Algorithms.....	51
Sorting.....	51
File format conversions.....	51
Output files.....	52
Database search order.....	52
Search termination.....	53
Accept / reject criteria for UCLUST and UBLAST.....	54
Weak match criteria in UBLAST.....	54
Alignment style options	55
Alignment scoring parameters	55
Alignment heuristics	56
Karlin-Altschul statistics and E-value calculation	57
Other options.....	57

USEARCH

USEARCH is a program that implements several algorithms for high-throughput clustering, database search and related tasks.

License

USEARCH is copyrighted software and generally requires a paid license. Licenses to use the 32-bit binaries are available at no charge for non-commercial use.

Installation

USEARCH is distributed as a stand-alone binary (executable file). The binary is self-contained: it does not require configuration files, environment variables, third-party libraries or other external dependencies. There is no setup script or installer because they're not needed. All you need to do is download or copy the binary file to a directory that is accessible from the computer where you want to run the code. For more information, please see <http://drive5.com/cmdline.html>.

Executable filename

The file name of the binary file includes the platform and version number, e.g. `usearch4.0.17_i86linux32` is the 32-bit binary for version 4.0.17 on Intel i86 architecture Linux. In this manual, 'usearch' (lower-case) will be used for the command name. Where you see 'usearch', you should replace this with the appropriate command name for your system. For example, this is a generic command:

```
usearch --cluster input.fasta --uc results.uc --id 0.97
```

On your system, you would use something like this:

```
usearch4.0.23_i86linux32 --cluster input.fasta --uc results.uc --id 0.97
```

Or, if you like, you can change the name of your executable file to `usearch`, which is easier to remember and type.

USEARCH overview

The USEARCH program implements several algorithms with a rich set of options. This table gives a brief summary of its most popular capabilities.

Category	Common uses
Clustering (UCLUST).	<p>Create non-redundant ("NR") and reduced-redundancy ("RR") databases.</p> <p>Dereplication: removing identical (sub-)sequences.</p> <p>Identify Operational Taxonomic Units (OTUs) from single-region environmental sequencing reads (e.g. 16S or ITS).</p>
Database search with clustering of unmatched sequences (UCLUST).	<p>Extend a previously clustered database by adding new sequences (saves compute resources compared with clustering again from scratch.)</p> <p>Classify single-region reads (e.g. 16S or ITS) by matching reads to known species in a database and generating OTUs from the unmatched reads.</p>
Database search (UBLAST).	<p>Capabilities similar to BLASTP, BLASTN and BLASTX, plus fast search options. Can achieve speed improvements over BLAST of 100 to 1000x with good sensitivity to distant relationships (down to around 75% id for nucleotides or 35% id for proteins). Used in a wide variety of sequence classification tasks.</p>
Chimera detection (UCHIME).	<p>Search for chimeric sequences <i>de novo</i> or using a trusted reference database that is believed to be chimera-free.</p>

Basic usage: clustering with UCLUST

This section describes the most popular options for the UCLUST clustering algorithm. Many advanced options are also provided; these are described in later sections.

Seed sequences

UCLUST creates clusters defined by *seed* sequences. Each cluster has exactly one seed, which is a sequence from the input set. The user specifies an identity threshold using the `--id` command-line parameter. For example with `--id 0.97`, other members of the cluster must have identity at least 97% identity with the seed. Another term for a seed is a *representative sequence*.

Sorting the input sequences

Input sequences should be sorted in an appropriate order as a pre-processing step before clustering. UCLUST processes sequences in the order they appear in the input file. If a sequence matches an existing seed, then it is assigned to that cluster; otherwise it becomes the seed of a new cluster. This means that sequences should be sorted such that the most appropriate seed for a cluster tends to appear first. Suggested sort orders for two common applications follow.

Non-redundant database: sort by decreasing length

UCLUST can create a non-redundant or reduced-redundancy database from an input set containing families of similar sequences. In this case, the generally recommended sort order is by decreasing length. Sorting by length can be done using the following `usearch` command.

```
usearch --sort seqs.fasta --output seqs.sorted.fasta --log usearch.log
```

Sorting by length is effective when some sequences are exact or approximate fragments of other sequences. A full-length sequence is usually a better choice of seed than a fragment. If the database is too large to fit in memory, the `--mergesort` command can be used.

OTU prediction: sort by decreasing abundance

UCLUST is often used to predict Operational Taxonomic Units (OTUs) from environmental sequencing of a region such as the bacterial 16S ribosomal RNA gene or fungal ITS. In this case, the generally recommended sort order is by decreasing abundance. The abundance of a read is the number of times an identical (or highly similar) sequence is found in the input. It is *not* recommended to use a decreasing length sort for OTU identification, for reasons described later. Sorting by abundance requires that cluster size information is provided from any dereplication step used in the analysis. *Dereplication* is removal of identical or highly similar sequences, which can be done using USEARCH as described later.

The --cluster command

Clustering is performed using --cluster. This is a typical command line.

```
usearch --cluster reads.sorted.fasta --id 0.97 --seedsout seeds.fasta
--uc results.uc --log usearch.log
```

The --id option specifies the identity threshold. In this example, sequences in a given cluster must have $\geq 97\%$ identity with its seed sequence. Note that the --id parameter is specified as a fractional identity in the range 0.0 to 1.0 rather than as a percentage.

Output files

The most popular output file options for clustering are --seedsout, which produces a FASTA file containing the seed sequences (i.e., a non-redundant or reduced-redundancy subset of the input), and --uc, which reports results in a custom file format designed to be easily parsed by a script or manipulated by standard Linux commands such as cut, grep and sort. The --uc file format is described in detail later. Other output file options include --userout, --blastout, --blast6out and --fastapairs.

The --usersort option

By default, the sort order is assumed to be by decreasing length, and a fatal error occurs if sequences are found to be out of order. The --usersort option specifies that another sort order was used, so is required for abundance sorting or other orders.

Basic usage: database search with UBLAST

This section describes the most popular options for the UBLAST database search algorithm. Many advanced options are also provided; these are described in later sections.

Database format

UBLAST uses FASTA files for both the query and database sequences. You don't have to use a program like formatdb or makeblastdb as for BLAST.

Search command

Here is an example of a typical database search command-line.

```
usearch --query query.fasta --db db.fasta --blastout results.blast
--blast6out results.b6 --evaluate 0.01 --log usearch.log
```

The query and database FASTA files are specified by the `--query` and `--db` parameters. The `--evaluate` option specifies the maximum E-value for a hit.

U-sorting

UBLAST is designed to quickly find one strong hit, which will often be the best hit in the database. This differs from traditional search algorithms like BLAST, which are designed to find all possible hits in the database. UBLAST tests database sequences in order of decreasing number of short words in common with a given query sequence (*U-sorted order*). Similar sequences tend to have more words in common, so the first hit found is often the strongest hit in the database, or one of the best. By default, UBLAST terminates a search when the first hit is found, which can dramatically improve search speeds. This strategy can be disabled by using the `--nousort` option. If `--nousort` is specified, USEARCH tests all database sequences using an algorithm very similar to BLASTP or BLASTN. Depending on the database size and other options specified, this may be much slower than using U-sorting.

Number of targets tested

UBLAST will abandon a search if too many database sequences (*targets*) have been tested without finding a hit. The maximum number of targets to be tested can be set using the `--maxtargets` parameter, which defaults to 9. Increasing this value improves sensitivity but reduces speed. If the database is very large, and / or if low-identity hits are desired, then sensitivity can be significantly improved by increasing this parameter. If you have, say, millions of proteins and are looking for hits that may have <50% identity, then you may get better results with `--maxtargets` values of 100 or 1000, or even larger. The `--maxtargets` option is ignored if U-sorting is disabled, i.e. if `--nousort` is specified.

Nucleotide, protein and translated searches

USEARCH can perform nucleotide searches (like BLASTN), protein searches (like BLASTP) and translated searches (like BLASTX). USEARCH automatically detects the alphabet of the query and database sequences. The type of search is determined per the following table.

Query sequences	Database sequences	Search
Nucleotide	Nucleotide	Nucleotide (like BLASTN)
Amino acid	Amino acid	Protein (like BLASTP)
Nucleotide	Amino acid	Translated ORFs (like BLASTX ¹)
Amino acid	Nucleotide	<i>Not supported in this version.</i>

¹USEARCH uses ORFs as queries, BLASTX does not.

Output files

Output file options include --blastout (human-readable, BLAST-like), --blast6out (tabbed format compatible with the -m8 or -outfmt 6 option to NCBI BLAST), --fastapairs (pair-wise alignments in FASTA format), --uc (tabbed format primarily designed for clustering results) and --userout (tabbed format with fields specified by --userfields).

Basic usage: Dereplication, removing identical (sub-)sequences

Dereplication is the process of removing sequences that are identical to, or a substring of, other sequences in the input. Input should be sorted by length. The following command-line is suggested.

```
usearch --cluster input.fasta --id 1.0 --iddef 2 --band 1 --seedsout nr.fasta
```

Search strategy

A central step in most USEARCH algorithms is to search a database, which is stored and indexed in memory. The following tables summarize the most common ways in which a database can be constructed and searched. These variations are explained in more detail later.

Database

Option	Description
Search	The database is read from a FASTA file and does not change. This is done in a typical search, which is specified if the --query and --db options are both given.
Clustering	The database is initially empty. Each new seed sequence is added to the database, so the database contains one sequence per cluster. This strategy is used if the --cluster option is given but not --db.
Search + clustering	The database, which contains cluster seed sequences, is initialized from a FASTA file, and then grows as new seeds are identified. This strategy is used if the --cluster and --db options are both given.

Search order

Option	Description
U-sorted	Database is searched in order of decreasing number of words in common. This is the default.
S-sorted	Modified U-sorted order. Designed to increase sensitivity by improving the correlation of the search order with evolutionary distance. Specified by the --[no]ssort option. Amino acid databases only. S-sorting is the default if the --query and --eval options are both given, otherwise it is disabled by default.
Entire database	All database sequences are tested. Specified by the --noursort option. In this case, the --[no]ssort option is ignored.

Search termination (U-sorted or S-sorted only)

Option	Description
Max accepts	Search stops after this number of accepts (matching targets) found. Set by --maxaccepts option, default value is 1.
Max rejects	Search stops after this number of failed attempts to match a target sequence. Set by --maxrejects option, default value is 8.
Max targets	Search stops after this number of target sequences have been tested. Set by --maxtargets option, by default this option is disabled.

Alignment

Feature	Variants	Description
Heuristic/optimal	Heuristic	By default, heuristics similar to those used in the BLAST program are used to speed up calculation of an alignment.
	Optimal	Heuristics can be turned off using --nofastalign, in which case the Needleman-Wunsch or Smith-Waterman algorithms are used, i.e. full $O(L^2)$ dynamic programming. This is often much slower, and gives results similar to programs like SSEARCH and NEEDLE. Optimal alignments can be useful for benchmarking, e.g. to evaluate the effect of using heuristics, and for <i>ad hoc</i> searches of small databases where the best possible accuracy is desired; for such applications USEARCH may be convenient due to its ease of use, flexible options and rich set of output file formats.
Global/local	Global	All letters of both sequences are aligned. This is the default if --cluster is given, otherwise can be specified using the --global option.
	Local	Two segments (one substring from each sequence) are aligned. This is the default if --query is used, otherwise can be specified using the --local option.

Similarity measures

Variants	Description
Identity	Fraction of columns in the alignment that contain identical letters. Minimum identity is specified by the --id option, which ranges from 0.0 to 1.0, meaning 0% to 100% identity. Several different definitions of identity are supported, as specified by the --iddef option. Can be used for both local and global alignments.
E-value	Karlin-Altschul statistics are used to calculate bit scores and E-values. The maximum E-value is specified by the --evalue option. Applies to local alignments only.
Sequence coverage	Coverage is defined as the fraction of letters in one sequence that are aligned to letters in the other. The minimum coverage is specified by the --queryfract and --targetfract and options (range 0.0 to 1.0; default is 0.0, which effectively disables the option). Can be used for local or global alignments.

Word counting

Option	Description
Word length	Short words of fixed length (k -mers) are used for two purposes: U-sorting and seeding alignments (like in BLAST). The word length for U-sorting is set by the --w option, which defaults to 5 for amino acids and 8 for nucleotides. The seed word length is set by the --k option, which defaults to 3 for amino acids and 5 for nucleotides.
Word count rejection	By default, a target is rejected if it has too few words in common with the query. This improves speed by eliminating the expensive alignment step, but can result in some false negatives. In practice, this applies only to high identity thresholds because at lower identities the required number of words in common is one or zero. Word count rejection is disabled by --nowordcountreject.
Stepping	A subset of words in the query may be used rather than all words. This improves speed by reducing the time required to find those words in the database. The subset is obtained by extracting words at intervals > 1 letter (stepping). Specified by the --stepwords option. Default value is 8, which means that the number of words extracted from the query is chosen so that a target sequence with the required identity is expected to have at least 8 of those words in common.
Bumping	This optimization reduces the time required for word counting and U-sorting in cases where many target sequences exceed the initial word count rejection threshold. Specified by the --bump option, which defaults to 50.

Sorting sequences for clustering

Sequences should be sorted prior to clustering. The order should be chosen so that an appropriate seed sequence for a cluster tends to appear first, before other members of the same cluster. There is no need to sort query sequences for database search without clustering, because no new seeds are created.

Sorting by decreasing length

Sorting by decreasing length is effective when full-length sequences and exact or approximate fragments are both present in the input. Fragments are usually not a good choice of seed, as shown by the following example.

```
Seed:          THESEED-----  
First hit:     THESEEDINSERTED  
Second hit:   THESEEDTERMINAL
```

Here the seed is a fragment. The two hits are both 100% matches to the seed except for terminal gaps and would therefore be assigned to the same cluster. However, the hits are extended with different terminal regions (red) and therefore have only about 50% identity to each other. This issue can also be addressed by using a more appropriate definition of identity (--iddef option).

The --sort command sorts sequences by decreasing length.

```
usearch --sort input.fasta --output input_sorted.fasta
```


The current implementation of --sort loads all sequences into memory for speed. Available memory (real or virtual) must be at least as big as the input file. Larger sets can be sorted using a slower merge sort, as in the following example.

```
usearch --mergesort input.fasta --output input_sorted.fasta --split 500.0
```

The --split option (default 1000.0) specifies the number of megabytes to use for each partition of the input file. Typically, the maximum memory needed for the sort will be a little more than this, and in a worst-case scenario can be closer to 2x the --split value, so a conservative choice is to use about half the physically available memory. Smaller values of --split tend to be slower.

Sorting by decreasing abundance

Sorting by decreasing abundance is recommended for identifying OTUs from single-region environmental sequencing reads such as 16S or ITS. The most abundant sequence is likely to be a true biological sequence, while less common sequences may be artifacts due to sequencing error or PCR artifacts such as chimeras, as illustrated in the following figure. This shows the cluster for a single species; the red dot represents reads of the true sequence of the species. A dot indicates a unique sequence, the size of the dot indicates its abundance, i.e. the number of identical (or very similar) reads having that sequence. The longest sequence in the figure is likely to be one of the outliers, and will give a less accurate OTU—imagine drawing a circle of radius of size 97% around one of the outlying dots and you will see that some reads that belong to the species will be incorrectly excluded.



Sorting by abundance can be accomplished as follows. First cluster at a high identity, say --id 1.0 or 0.99. The size of the clusters can then be obtained from the C records in the .uc file (see below for more information on the .uc file format). This can be done using a series of commands like the following.

```
usearch --sort reads.fasta --output reads.sorted.fasta
usearch --cluster reads.sorted.fasta --id 0.99 --uc 99.uc
grep "^C" 99.uc | sort -nrk3 > c.uc # sort by decreasing cluster size
usearch --input reads.sorted.fasta --uc2fasta c.uc --output c.fasta
```

Grepping on the first letter (grep "^C" above) selects a single type of record from the .uc file. For details of the .uc file format, see below.

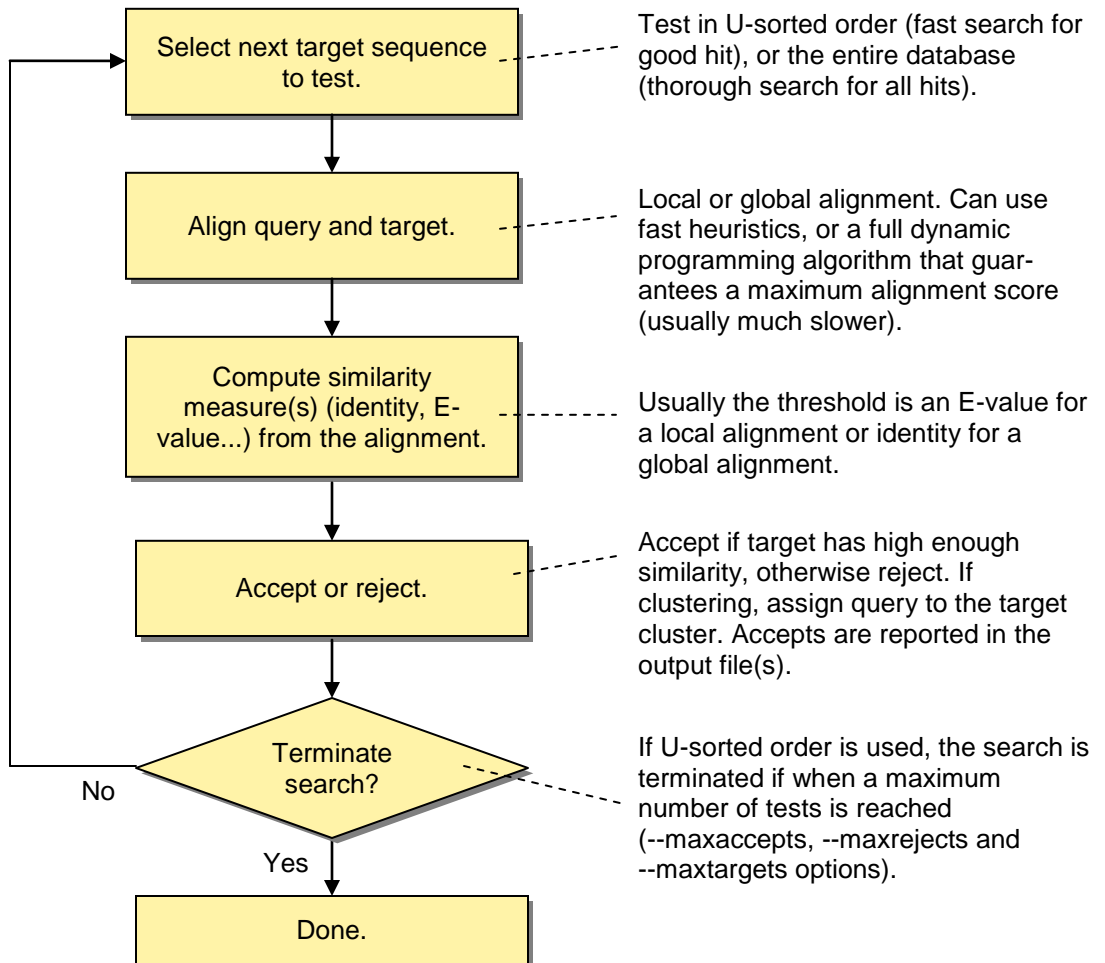
Finally, OTUs can be identified by clustering, e.g. for species 97% is typically used:

```
usearch --cluster c2.fasta --id 0.97 --uc 97.uc --usersort
```

Note the --usersort option, which is required for clustering when input sequences are not sorted by decreasing length.

Searching

A fundamental step in most USEARCH algorithms is to search a sequence database. For example, UCLUST searches a database of seeds to find a matching cluster for an input sequence. Many different types of search can be specified via command-line options, allowing different trade-offs between speed and sensitivity, local or global alignments, and so on. The typical steps for a given query sequence are shown in the following flowchart.



Search order

Most sequence database search algorithms compare a query sequence with all database sequences (targets). By default, USEARCH algorithms test database sequences in U-sorted order and stop searching as soon as a strong enough hit ("accept") is found. This strategy is effective because U-sort order correlates well with sequence similarity, so the first hit found is often the best hit in the database, or one of the best. U-sorting can be disabled by specifying the --nousort

option, in which case all database sequences are aligned to the query and the target with the highest similarity is selected.

Query-target comparison

A query is tested against (compared with) a target sequence by first creating an alignment, then calculating a measure of similarity from the alignment. Many variants of these steps are supported, e.g. alignments can be global or local, and the measure of similarity can be identity or an E-value. The user can "mix and match" these variations as desired, e.g. the default for clustering is to use global alignment with identity as a similarity measure, but clustering can also be performed using local alignments with an E-value threshold.

Alignment parameters and heuristics

Given a pair of sequences (query and target), there are two sets of options that control the construction of an alignment: *scores* and *heuristics*. Scores include the substitution matrix and gap penalties. Changing scores will tend to change the optimal alignment of a given pair of sequences. Heuristics are approximations that reduce the time required to calculate an alignment. Ideally, changing the heuristic parameters or disabling the heuristics altogether would not change the alignment. However, by definition heuristics do not always produce an alignment with the best possible total score. They are introduced because they improve speed, at the expense of a possible reduction in accuracy. Here, accuracy should be understood in a computer science rather than a biological sense—the optimal alignment of a given pair of sequences may have biological errors despite having the best possible score.

USEARCH with `--nofastalign`, i.e. with heuristics disabled, is roughly equivalent to programs like SSEARCH and NEEDLE that are based on dynamic programming algorithms without approximations. This can be useful for benchmarking USEARCH, e.g. to evaluate the effect of using heuristics, and for searches with smaller datasets where maximum sensitivity is important. Here, USEARCH may be convenient compared with some other programs due to its ease of use, flexible options and its rich set of output file formats that are designed to be easily reviewed by a human reader or robustly parsed by scripts.

Similarity measure

One or more measures of similarity are computed from a query-target alignment. For clustering, the measure is usually identity computed from a global alignment. For database search, the measure is usually an E-value computed from a local alignment.

Search termination

By default, the database is searched in U-sorted order, and a search is terminated when either (i) a strong enough hit (accept) is found, or (ii) the maximum number of target sequences has been tested. If U-sorting is disabled (`--noursort`), the entire database is searched, and any termination options given on the command-line are ignored or cause an error.

Accepts and rejects

A target sequence that meets the threshold criteria is called an accept. Typically, this means the target sequence has a high enough identity (clustering with global alignment), or a low enough E-value (database search with local alignment). If the target sequence fails to match, it is called a reject. If the weak similarity parameters `--weak_id` or `--weak_evalue` are used, then weak matches are reported in the output files but do not count as accepts and will not cause a query to be assigned to a cluster.

Termination conditions

The following conditions terminate a search. If any condition is satisfied, the search stops. In other words, conditions are combined using "OR".

Maximum accepts

Maximum number of accepts has been found. This is set by the `--maxaccepts` parameter, which defaults to 1. In other words, by default, a search terminates immediately when the first hit is found. If `--maxaccepts` is set to zero, there is no limit on the number of accepts (so zero means infinity). Sometimes the first hit found is not the best hit in the database; increasing `--maxaccepts` increases the probability that the best hit will be found, at the expense of slower execution time. If `--maxaccepts` is increased, you should probably increase `--maxrejects` also.

Maximum rejects

Maximum number of rejections has occurred. This is set by the `--maxrejects` parameter (default 8). If `--maxrejects` is set to zero, there is no limit on the number of rejections, so the search will continue until an accept is found or the entire database has been searched (unless `--maxtargets` is set). Sometimes a hit is not found because the search is terminated too quickly; increasing `--maxaccepts` increases the probability that a hit will be found, at the expense of slower execution time.

Maximum targets

Maximum number of target sequences has been tested. This is set by the `--maxtargets` option, which defaults to zero (which disables the option, so again zero means infinity). If `--maxtargets` is set, then `--maxaccepts` and `--maxrejects` are disabled unless they are set on the command line.

Search and clustering at high identities

The default termination parameters are `--maxaccepts 1 --maxrejects 8 --maxtargets 0`. These are designed for high-identity clustering, which is one of the most common USEARCH applications, and also work well for database search when typical matches have high identity. When identity is high, word count correlates well with similarity, which means that the first accept found is usually the best, or close to it, and the probability of finding an accept falls rapidly with the

number of rejects. A U-sorted search therefore quickly reaches a point of diminishing returns if a match is not found in the first few attempts, so `--maxrejects` values larger than 8 give only small improvements in sensitivity.

Search and clustering at medium identities

When identity is lower, word count correlates less well with similarity, and sensitivity can therefore often be improved by testing more database sequences. Medium identity means, very roughly, 75% for nucleotides or 50% for proteins. Here, it may give better results to increase `--maxaccepts` over the default value of 1 because at medium identities, the first hit found is less likely to be the best hit and it may therefore be advantageous to test a few more targets. Typical parameters that might work well for medium identity applications are:

```
--maxaccepts 3 --maxrejects 32.
```

Search at low identities

When distant relationships are important, the default parameters will not work well because the number of words in common correlates poorly with similarity below around 50% identity for proteins or 80% for nucleotides. Clustering is rarely useful at such low identities, so this issue applies mainly to searching without clustering. The reduction in sensitivity can be mitigated by increasing the number of target sequences tested, which will be especially important when searching large databases, which tend to produce many spurious candidates (rejections) when tested in a U-sorted or S-sorted order. Some typical parameters are:

```
--maxtargets 1000 --maxaccepts 8 --evaluate 1e-6 --weak_evaluate 0.01.
```

Output files

The USEARCH database search and clustering algorithms support several output file formats. Most output file formats and features are supported by most of these algorithms.

Option	Format	Description
--uc	UCLUST	Tab-separated file designed primarily for clustering pipelines but can also be useful for search. There is one record for each input sequence giving its cluster assignment, identity and alignment; and one record for each cluster giving its size and average identity.
--blastout	BLAST-like	Verbose, human-readable format similar to BLAST.
--blast6out	Tab-separated	Tabbed format with one record per hit. Compatible with the -m8 or -outfmt 6 option of NCBI BLAST.
--userout	Tab-separated	Tabbed format with one record per hit, fields specified by the --userfields option.
--seedsout	FASTA	Seed sequences, i.e. the non-redundant or reduced redundancy set of sequences after clustering.
--fastapairs	FASTA	Pair-wise alignments in FASTA format.

The UCLUST file format

The native UCLUST format (.uc) is a tab-separated text file. UCLUST output is supported by clustering and database search. Each line is either a comment (starts with #) or a record. Each query sequence generates at least one record; additional record types give information about clusters. The cluster number appears in every record type except R (reject). If an input sequence matched a target sequence, then the alignment and the identity computed from that alignment are also provided. A compressed representation of the alignment is used to save disk space. Records are appended to the output file as they are generated in order to minimize memory use, and sequences therefore appear in the same order as the input file.

Some example records in .uc format are show below.

Type	Cluster	Size	%Id	Strand	Qlo	Tlo	Alignment	Query	Target
S	0	292	*	*	*	*	*	AH70_12410	*
H	0	292	99.7	+	0	0	292M	EN70_12566	AH70_12410
S	1	292	*	*	*	*	*	EX70_12567	*
H	1	292	98.2	+	0	0	292M	AH70_12410	EX70_12567

Each record has ten fields, separated by tabs, as described in the following table.

Field	Name	Description
1	Type	See table below.
2	Cluster	Cluster number
3	Size	Sequence length or cluster size
4	Id	Identity to the seed (as a percentage), or * if this is a seed.
5	Strand	+ (plus strand), - (minus strand), or . (for amino acids).
6	Qlo	0-based coordinate of alignment start in the query sequence.
7	Tlo	0-based coordinate of alignment start in target (seed) sequence. If minus strand, Tlo is relative to start of reverse-complemented target.
8	Alignment	Compressed representation of query-seed alignment, or * if a seed.
9	Query	FASTA label of query sequence.
10	Target	FASTA label of target (seed / library / database) sequence, or * if a seed.

Record types are as follows.

Type	Description
S	Seed.
H	Hit, also known as an accept; i.e. a successful match.
C	Cluster (seed is a sequence in the --cluster file).
D	Library cluster (seed is a sequence in the --db file).
N	Not matched.
R	Reject (generated only if --output_rejects is specified).
L	Library seed. There is exactly one L record for every --db sequence that has one or more hits.

Records of type C and D are used when clustering. The Size field contains the cluster size, i.e. the number of sequences in the cluster including the seed, and Id is the average identity of non-seed sequences to the seed. Otherwise, Size is the sequence length and Id is the identity of the pair-wise alignment of this sequence to the seed. For Library clusters (D), records are only output if Size > 1, i.e. library sequences with no matches are not output. A library seed record (L) is output only if a hit is found to that database sequence. This saves writing a large number of records for database sequences that are not matched, but means that cluster numbers in the .uc file may not be consecutive (because UCLUST internally assigns a cluster number to every library seed, whether or not it is matched).

Rejections (R) are sequences that were aligned to a seed but found to have an identity below the threshold. Rejections are not output unless --output_rejects is specified. Rejection records can be useful when trouble-shooting unexpected results.

The alignment is compressed using run-length encoding, as follows. Each column in the alignment is classified as M, D or I.

Class	Name	Query	Target
M	Match	Letter	Letter
D	Delete	Letter	Gap
I	Insert	Gap	Letter

Here, "match" simply means a letter-letter column; the letters may or may not be identical. Deletions and insertions are relative to the query. If there are n consecutive columns of type C, this is represented as nC . For example, 123M is 123 consecutive matches. As a special case, if $n=1$ then n is omitted. So for example, D5M2I3M represents an alignment of this form:

```

Query sequence  -XXXXXXXXXX
Seed sequence   XXXXXX--XXX
Column type     DMMMMMIIMMM

```

If a line in the output file starts with #, it is a comment and parser scripts should ignore it.

Records in the .uc file appear in the same order as the input sequences. You can sort the file by cluster number using the standard Linux sort command, as follows:

```
sort -nk2 results.uc > results_sorted.uc
```

You can also sort by cluster number using USEARCH:

```
usearch --sortuc results.uc --output results_sorted.uc
```

However, the current implementation reads the entire file into memory, so may fail for very large sequence sets.

UCLUST to FASTA conversion

You can convert UCLUST to FASTA using the --uc2fasta or --uc2fastax commands.

```

usearch --uc2fasta results.uc --input seqs.fasta --output results.fasta
usearch --uc2fastax results.uc --input seqs.fasta --output results.fasta

```

Here, seqs.fasta must be the same input file used when generating results.uc.

The --uc2fasta command outputs sequences with the same labels and sequences as found in the input file.

The --uc2fastax format reformats both labels and sequences when generating FASTA format output. Labels look like this:

```
>43|99.7%|AH70_12410
```

Here, 43 is the cluster number and 99.7% is the identity to the seed. The identity will be shown as * for the seed:

```
>43|*|AH70_12200
```


If a .uc record has an alignment, then the query sequence is re-formatted to indicate its pair-wise alignment to the seed. Gaps indicate deletions relative to the seed, lower-case indicates insertions relative to the seed. Here is an example:

```
>43|99.7%|TheSeed
SEQUENCE
>43|96.0%|NonSeed
S-QLENNCE
```

This represents the following pair-wise alignment:

```
TheSeed   SEQVEN-CE
NonSeed   S-QLENNCE
```

BLAST6 format

Local alignments (HSPs) are reported if the --blast6out filename option is given. The format is compatible with the NCBI BLAST -m8 or -outfmt 6 options. It is a tab-separated text file with one line per HSPs. By convention I use the .b6 extension for files in this format. There are twelve fields, as shown in the following table.

Field	Description
1	Query label
2	Target label
3	Percent identity
4	Alignment length
5	Number of mismatches
6	Number of gap-opens
7	1-based position of start in query
8	1-based position of end in query
9	1-based position of start in target
10	1-based position of end in target
11	E-value
12	Bit score

BLAST-like alignment format

Alignments generated during clustering or database search can be saved in a human-readable BLAST-like format by using the --blastout option, e.g.:

```
usearch --query seqs.fasta --db genes.fasta --blastout hits.blast
```

Since this format is rather verbose, the file size will be much larger than the corresponding .b6 or .uc file. The details of the formatting are subject to change between versions. It is therefore recommended that parsers use --userout or --blast6out. If full alignments are required, --fastpairs can also be used, though see also the grow and trow fields for --userout. An example --blastout alignment is shown below.

```

Query    280aa >Q4HFD3_CAMCO
Target   337aa >A6CVA5_9VIBR

  1 LCLGVFGLISMELGVMGIIPLISEKFGVSVSDAGWSVSIFALIVMCCAPIAPMLCANFNPKKLM 64
    | | | : . | : ||| :. :. ||: :| :|| :. | | . | | | . :
  1 LTLAAFAIGTAEFIIAGILPQVATSLSITEGQAGYLISAYALAIVIGGPILTIYLARFNKKMVL 64

 65 LFCLAIFSLSSLASMFVNDFWLHLILRAIPAFFHPIYLALAFSTAANLADDKSKVPHIVAKIFM 128
    : :||:| . .| | | : : | | | : : | | | :| | | . .|
 65 IGLMALFIVGNLMSAFSPSYDILFISRIISGLVQGPFGYIGAVVATNLVSEKM-AGRAVGQMFA 127

129 AISAGLVLVGPLSSYFVGGNFSFEMAMAFYVVINSLAFFITLFFMPEFKKTSRIKVGKQLLSLRY 192
    :. ||||| .: | | . .| | | . | | | . :
128 GLTLANVLGVPGGTWIGVEFGWHTTFIVVAAFVVALFAILAAIHSTGHGEAKNVKAQLAAAFKN 191

193 ALLWISMLAVFCISTGYLGFYSYSEFLFSVSKMSFTNISLALFIYGFASIIGNNIAGKTLVNH 256
    | ||: : ||:: | | . | . . :. | | | |||| : | .
192 PKLLISLAITAVVWTGFMTLYGYIAPIAMHVAGYGESAVTWILVIVGLGLIIGNTLGGHSSDKD 255

257 SNQTLIFASIAMILYALIFV 277
    |.. :| .||| | : |
256 LNKSSLFWAIAMIASLVLVGV 276

277 cols, 69 ids (24.9%), 207 diffs (74.7%), 1 gaps (0.4%)
Score 256.0 (103.2 bits), Evalue 9.1e-023

```

FASTA alignment format

Alignments can be saved in FASTA format by using the `--fastapairs` option, e.g.:

```
usearch --query seqs.fasta --db greengenes.fasta --fastapairs hits.fasta
```

The query sequence is first, the target (seed, database) sequence is second. If the input sequences are nucleotides, then a + or - is appended to the label of the target sequence to indicate the strand. If the strand is - (reverse strand match), then the target sequence is reverse-complemented.

CD-HIT format

The CD-HIT `.clstr` format is supported for the benefit of code already written for that format. You can convert UCLUST format to and from `.clstr` as follows:

```

usearch --uc2clstr results.uc --output results.clstr
usearch --clstr2uc results.clstr --output results.uc

```

User-defined output

Tabbed output in a user-defined format is produced by using the `--userout` and `--userfields` options. For example,

```

usearch --query query.fasta --db db.fasta --userout results.user
--userfields query+target+evaluate

```

The `--userout` option specifies the filename, and the `--userfields` option specifies one or more field names separated by +.

The output file is tab-separated. The first line contains the field names as specified by the --userfields option; each subsequent line contains one hit. Fields are output in the order given by --userfields. An example output file produced by --userfields query+target+evaluate is as follows:

```

query      target      evaluate
FQ76998    PF01023    1.2e-12
AZT77876   PF10922    6.7e-23
...etc...
```

Supported user fields are described in the following table.

User field	Description
query	Query sequence label.
target	Target (database, seed) sequence label.
evaluate	E-value computed using Karlin-Altschul statistics.
id	%id as reported in other output files, i.e. calculated according to the --iddef option.
id0	%id as if --iddef 0 was specified.
id1	%id as if --iddef 1 was specified.
id2	%id as if --iddef 2 was specified.
id3	%id as if --iddef 3 was specified.
id4	%id as if --iddef 4 was specified.
pctpv	% alignment columns that contain a pair of letters with score > 0 per the substitution matrix.
pctpvz	% alignment columns that contain a pair of letters with score >= 0 per the substitution matrix.
pctgaps	% alignment columns that contain a gap.
pairs	Number of alignment columns containing a pair of letters.
gaps	Number of alignment columns that contain a gap.
ins	Number of alignment columns that contain an insertion (gap in query).
del	Number of alignment columns that contain a deletion (gap in target).
intgaps	Number of internal gaps.
tgaps	Number of terminal gaps.
ltgaps	Number of terminal gaps at the left of the alignment.
rtgaps	Number of terminal gaps at the right of the alignment.
qlo	Start coordinate in query (one-based relative to start of sequence).
qhi	End coordinate in query (one-based relative to start of sequence).

User field	Description
tlo	Start coordinate in target (one-based relative to start of sequence or reverse-complemented sequence).
thi	End coordinate in target (one-based relative to start of sequence or reverse-complemented sequence).
qloz	Start coordinate in query (zero-based relative to start of sequence).
qhiz	End coordinate in query (zero-based relative to start of sequence).
tloz	Start coordinate in target (zero-based relative to start of sequence or reverse-complemented sequence).
thiz	End coordinate in target (zero-based relative to start of sequence or reverse-complemented sequence).
pv	Number of alignment columns that contain a pair of letters with score > 0 per the substitution matrix.
pvz	Number of alignment columns that contain a pair of letters with score >= 0 per the substitution matrix.
ql	Full length of query sequence.
tl	Full length of target sequence.
qs	Length of query segment appearing in the alignment.
ts	Length of target segment appearing in the alignment.
cols	Number of alignment columns.
intcols	Internal columns, i.e. number of columns that are not terminal gaps.
opens	Number of gap opens.
exts	Number of gap extensions.
qi	Query sequence index, i.e. the zero-based number 0, 1, 2... of the sequence in the query file.
ti	Target sequence index, i.e. the zero-based number 0, 1, 2... of the sequence in the database.
raw	Raw score = sum of substitution scores minus gap penalties.
bits	Bit score computed from the raw score using Karlin-Altschul statistics.
strand	Strand: one letter '+' (forward strand), '-' (backward strand), or '.' for amino acid sequences.
frame	Signed integer -3, -2, -1, +1, +2 or +3 indicating the frame. For translated searches only, otherwise appears as ".".

User field	Description
aln	Alignment, coded as a string with one letter for each column: M is a pair of letters, D is a delete (gap in target), I is insert (gap in query).
caln	Alignment compressed using run-length encoding, exactly as in the .uc file format.
qrow	Query alignment row, i.e. the aligned segment of the query sequence with gap characters '-' inserted as appropriate.
trow	Target alignment row, i.e. the aligned segment of the target sequence with gap characters '-' inserted as appropriate.

UBLASTX: Translated ORF search

UBLAST supports translated searches of nucleotide sequences against a protein database containing amino acid sequences. This is somewhat similar to BLASTX, except that ORFs are used as queries. This makes more effective use of the U-sort and S-sort heuristics.

Frame-shifts

The current UBLASTX implementation does not allow frame-shifts within an alignment. However, frame-shifts can be inferred from hits to ORFs in different frames on the same strand.

ORF identification

An ORF begins at the start of the sequence or with a START codon (ATG), and ends at a STOP codon (TAA, TAG or TGA) or the end of the sequence. Please [let me know](#) if you would like support for non-standard genetic codes. The minimum number of amino acids in the ORF is set by the `--mincodons` option (default 20).

The `--orfstyle` option controls how ORFs are defined. The value is created by adding up the following integers.

Value	Description
1	Allow an ORF to start at the beginning of a sequence, even if this is not a START codon (default cannot start before the first START codon).
2	Allow an ORF to start immediately following a STOP codon (default cannot start before the first START following a STOP).
4	Allow an ORF to end at the end of the nucleotide sequence (default must be terminated by a STOP codon).
8	Include the terminating STOP codon, if any, in the translated sequence (default do not include the STOP).

Default is `--orfstyle 7`. Since $7=1+2+4$, ORFs are identified as subsequences that do not contain a STOP, which is appropriate for shotgun metagenomic reads that may only partially cover a gene and may contain errors such as frameshifts.

Search and output

The translated amino acid sequence for each ORF is used as a query sequence to search the target database. The `--maxtargets` option is especially important here if the database is large and / or if low-identity matches are designed, typically you will need to specify a larger value of `--maxtargets` (say, 100 or 1000) to achieve good sensitivity with low-identity proteins.

UIRE: hierarchical clustering, clumping and large multiple alignments

The `--uhire` command performs hierarchical clustering with the goal of generating clusters of approximately a predetermined size (*clumps*). Sequences within a clump should be more similar to each other than to sequences in other clumps. This is intended to reduce the dataset size to be tractable for more expensive algorithms, such as multiple alignments. Basic usage is as follows.

```
usearch --uhire reads.sorted.fasta --hireout results.hire
--clumpout results.clump --clumpfasta filenameprefix --maxclump 1000
--ids id1,id2...,idN
```

At least one output option must be given, i.e. at least one of `--hireout`, `--clumpout` or `--clumpfasta`. The `--maxclump` option gives the maximum number of sequences in a clump (default 1000). The `--ids` option gives the percent identities of each level in the hierarchy. The default is 99,98,95,90,85,80,70,50,35. Note that `--ids` uses percentages (0 to 100), unlike `--id` which uses fractional identities (0.0 to 1.0). Values are separated by commas. Since commas are significant to most command shells, the value of the `--ids` argument should usually be quoted. If the `--clumpfasta` option is given, each clump is written to a file named `clump.0`, `clump.1`, `clump.2` etc., prefixed by the `--filenameprefix` option. This will typically be a directory name. E.g., you might do this:

```
usearch --uhire reads.sorted.fasta --clumpfasta myclumpdir/ --maxclump 256
```

Note the `'/'` at the end of the prefix. This is not required, but if present specifies that clump files are to be stored in the given directory, which must exist. Sequences for each clump will be stored in these files:

```
myclumpdir/clump.0
myclumpdir/clump.1
..etc..
```

In addition to the clumps, a file named 'master' will also be written. This contains the longest sequence in each clump. It can be used for creating large multiple alignments, as explained shortly below.

Warning

This method was primarily designed to support clumping (see below). Clusters at levels below the first (highest identity) level will tend to be more diverse than clusters obtained in a single step. Say the first two levels are 99% and 98%. The 98% step uses seeds from the 99% step as input. Suppose a cluster at 99% includes two sequences S and A where S is the seed and A is another sequence such that $\text{pctid}(A,S) \geq 99\%$. A is discarded when the 98% clustering is done. Now suppose T is the seed at 98%, so $\text{pctid}(S,T) \geq 98\%$. There is no guarantee that A has $\geq 98\%$ id with T, it may be less, and in fact we should expect such cases because A can be 'further' from T than S is. So clustering all sequences including A at 98% will tend to give different numbers of clusters than the hierarchical method.

The .hire file format

The .hire format is designed to be easily parsed by a scripting language and to avoid very long lines as found in [mothur](#) files. If you would like a script to convert .hire to mothur format, please let me know.

A .hire file is a text file.

The first line is the number of levels (K), i.e. the number of ids specified in the --ids option.

The second line is the number of sequences (N).

The following N lines specify sequences. Each line contains three tab-separated fields, for example:

```
37261 167 GF2FOAC01BL6E9
```

The first field is the sequence ID, an integer 0, 1 ... (N-1). This is redundant, but should be used by parsers to check that they are in synch with the file.

The second field is the sequence length in letters.

The third field is the sequence label from the FASTA file.

Following the last sequence (ID=N-1) will be K levels. Each level is specified as follows.

The first line in a LEVEL is a record with four fields, for example:

```
LEVEL 6 9 70.0
```

The first field is always the text "LEVEL". The remaining fields are:

6 Level number, a zero-based level number 0, 1 ... K-1.
9 Number of input sequences at this level.
70.0 Percent identity for this level.

This is followed by one line per sequence. Each line has three fields. Here is a complete example of a level.

```
LEVEL 6 9 70.0
6 0 *
6 61 *
6 565 0
6 726 *
6 1542 *
6 4408 61
6 4858 0
6 4879 *
6 9366 *
```


In the lines following the LEVEL record, the first field is the level number. This field is also redundant, but should be used by parsers to verify consistency with the file. The second field is the sequence ID, referring back to the sequence records at the beginning of the file. Sequence IDs are the same for all levels. A given level will have only the subset of IDs that correspond to seeds discovered in the previous level. The third field is either a second sequence ID, indicating a match, or an asterisk '*', indicating no match. A match means that the sequence was assigned to a cluster, an asterisk means that this sequence becomes a new seed at this level. So the above example has six seeds that would be passed down to the next level and three matches, two to seed ID=0 and one to seed ID=61. If there is a 7th level, it will have six input sequences which are the seeds identified at level 6.

Large multiple alignments

[MUSCLE](#) can create alignments of up to perhaps 10,000 to 20,000 sequences, depending on the available memory and sequence lengths. Larger sets can be aligned using a divide and conquer strategy based on clumping. This may be advantageous even in cases where MUSCLE can align the complete set as the resulting alignments tend to be more compact, having fewer columns and thus fewer gaps, which may be preferred for some types of analysis.

In outline, the strategy is as follows.

1. Create clumps, i.e. clusters that are small enough for MUSCLE to align.
2. Create a 'master' set containing the longest sequence from each clump.
3. Align each clump.
4. Align the master set.
5. Merge the clumps into a final alignment, using the master alignment as a guide.

The first step is to create clumps. Anecdotally, I have found that a clump size of around 5000 gives good results, but this may vary depending on your data. I recommend experimenting with different clump sizes and examining the results. Typical commands would be:

```
mkdir myclumpdir
usearch --uhire seqs.sorted.fasta --clumpfasta myclumpdir/ --maxclump 5000
```

The clumps and the master set are then aligned using MUSCLE. For example (bash syntax):

```
mkdir clumpalns
cd myclumpdir
for filename in clump.* master
do
    muscle -in $filename -out ../clumpalns/$filename -maxiters 2
done
cd ..
```

I recommend the `-maxiters 2` option to MUSCLE as a good compromise between speed and accuracy for larger sets. Any multiple alignment method can be used in place of MUSCLE if desired.

The alignments are combined using the `--mergeclumps` command, as follows.

```
usearch --mergeclumps clumpalns/ --output aligned.fasta
```

Sequences in the master file are required to have their labels formatted to indicate the clump number. This is done automatically if the `--clumpfasta` option is used; if you use some other method to select the master set then you must take care to follow the label formatting convention. The clump ID (0, 1... N-1) is indicated by a prefix like `>M123|` where 123 is the clump ID. For example, this is a valid FASTA label for the master sequence of clump 28:

```
>M28|GF2FOAC01AU7TA
```

Clump 28 must contain an identical sequence with label `>GF2FOAC01AU7TA`, this correspondence is used to merge the alignments of each clump into a single multiple alignment.

Parameter tuning

Where possible, I recommended that you tune parameters to obtain a good trade-off between speed and sensitivity. Following are some suggestions for how this can be achieved.

Improved sensitivity for distant proteins

Two options to try when clustering or searching with distantly-related proteins are `--nb` and `--ssort`. If `--noursort` is specified, then `--nb` is the default, otherwise it may give improved sensitivity with only a small cost in speed.

Choose suitable quality measures

Typical goals of tuning are find parameters that give high-quality results with the shortest possible execution times. This requires a measure of quality. The log file (`--log` option) reports execution time, memory use and some statistics on search and clustering which could be used as quality measures. Alternatively, you could write a script to parse one of the output files: the `--uc`, `--blast6out` and `--userout` files are well suited for this purpose.

Quality measures for clustering

For clustering, sensitivity can be measured by (i) the number of clusters or, equivalently, by the average cluster size, and (ii) the average identity of a cluster member to the seed. Fewer clusters (larger clusters) indicate higher sensitivity, and higher average identity with the seed indicates that a better cluster assignment is made in cases where more than one seed matches.

Quality measures for database search

For database searching, sensitivity can be measured by (i) the fraction of query sequences that are matched to the database at the given E-value or identity threshold, and (ii) the average similarity of a hit. It is generally better to measure similarity by identity even if an E-value threshold is used, because E-values range over many orders of magnitude so the mean or median is not very informative.

Construct a query set that is small enough for testing

If a typical query set is so large that repeated testing is unreasonably slow, then the size of the set can be reduced. For a database search application, this can be done by taking a random sample. For clustering, a random sample is not suitable because this tends to reduce the average size of a cluster but not the number of clusters, which increases the average number of rejections per query. A smaller set can be obtained by clustering the input sequences and taking a subset of the clusters. This should give a subset with similar redundancy to the original.

Test with increasing values of `--maxtargets`

First test with `--noursort`, if possible. This causes the entire database to be searched and thus guarantees the best possible hit for a given query, but may be unreasonably slow. Either way, try a range of values of `--maxtargets`, for example 8, 16, 32, 64, 128, etc. Increase the value until your quality measure(s) do not increase significantly. It might be helpful to make graphs that plot execution time and your chosen quality measures as a function of `--maxtargets`. Regardless of whether you were able to use `--noursort`, you now have an estimate of the best possible results and the minimum value of `--maxtargets` that gives you good enough results.

Test with increasing values of --maxaccepts

Once you know the highest values of your quality measures that can be achieved on your test data, you can experiment with changing parameters and obtain an acceptable compromise: e.g., you might be satisfied with achieving 90% of the best possible sensitivity if the speed is improved by a factor of 100. The minimum value of --maxtargets that gives good enough quality is a first guess at a suitable value for --maxrejects, while much smaller values of --maxaccepts often perform well. For example, suppose --maxtargets 64 gives good enough results. Then set --maxtargets to zero (the default), use --maxrejects 64, and test increasing values of --maxaccepts. Try 1, 2, 4, 8 etc. until the quality is high enough.

U-sort word length

The --w option sets the word length used for U-sorting. The default is 5 for amino acids and 8 for nucleotides. Try different values to check the effect on memory use and speed. Start by adding and subtracting one. If adding or subtracting one gives a better result, try changing by two, and so on.

Tune alignment heuristics

You can measure the impact of alignment heuristics by testing with and without --nofastalign. If --nofastalign is specified, heuristics are disabled, and the execution time may be tens or hundreds of times slower. Measures such as the number of clusters may not change much despite a significant reduction in biological accuracy. (This happens when over-aggressive heuristics produce many bad alignments without this causing a bias in the quality measure). Therefore, it is best to use biologically informed reference data if possible in order to test the effects of the heuristics. (Of course, biological reference data are preferred for tuning all parameters). The recommended heuristics to try are summarized in the following table. For all numerical parameters except seed word length, larger values tend to increase execution times and smaller values are faster but may degrade accuracy, though often the effect on accuracy is negligible. The effect of the seed word length is less predictable. Reducing the band radius is often an effective way to improve speed without a significant loss in quality.

Option	Heuristic
--wordcountreject, --nowordcountreject	Enables / disables word count rejection. For higher identities, tends to improve speed when enabled, but may induce false negatives.
--k	Seed word length.
--nb, --nonb	Use / don't use word neighborhoods (amino acids only). Using neighborhoods improves sensitivity; effect on speed varies.

Option	Heuristic
--seedt	Seed score threshold (applies only for amino acids and if --nb is specified).
--xdrop_u, --xdrop_g, --xdtop_ug, --xdrop_nw	X-drop.
--band	Radius for banded dynamic programming.

USTAR: Fast multiple alignment of clusters

USEARCH can create a multiple alignment of each cluster found by UCLUST. This requires three steps: 1. clustering (--cluster), 2. extraction of S (seed) and H (hit) records, 3. conversion to FASTA (--uc2fastax) and 4. inserting additional gaps (--staralign).

```
usearch --cluster seqs_sorted.fasta --uc results.uc --id 0.97
grep "^[SH]" results.uc > sh.uc
usearch --uc2fastax sh.uc --input seqs_sorted.fasta --output sh.fasta
usearch --staralign sh.fasta --output aligned.fasta
```

The algorithm creates a 'star' alignment using pair-wise alignments to the seed, so the seed is the center of the star. This method emphasizes very high speed over alignment quality. It is not intended to replace slower but more accurate methods like [MUSCLE](#). When sequence identity is reasonably high, the alignment will be good enough to be informative, e.g. for identifying highly conserved segments. Note that in addition to creating a multiple alignment, a consensus sequence is generated for each cluster. This can be useful for high-throughput evaluation of cluster quality. See the UHIRE algorithm for a method that can create high-quality alignments of very large sets.

Gap penalties and substitution scores

USEARCH supports a comprehensive set of gap penalty and substitution score options. Different options apply to local vs. global alignments. All alignment scores and penalties in USEARCH can be specified as integer, floating point or real values.

E-value calculation

E-values are calculated by Karlin-Altschul statistics assuming default values for substitution scores and gap penalties. If you change the alignment scoring parameters, then E-value parameters must be adjusted accordingly. This is not a trivial exercise; the easiest way is usually to borrow parameters from some other program, such as BLAST. [Contact me](#) if you need more information.

Substitution scores for nucleotides

Two substitution scores are used for nucleotide sequences: match and mismatch. The match score must be positive and the mismatch score must be negative. For local alignments, the absolute value of the mismatch score should be greater than the match score. If you use non-default substitution scores, you should probably also specify appropriate gap penalties for those scores.

Score	Option	Default
Match	--match	1.0
Mismatch	--mismatch	-2.0

Substitution matrix for amino acids

By default, the BLOSUM62 matrix is used for amino acid sequences. The user can specify a different matrix by using the --matrix *filename* option. The matrix should be formatted as for NCBI BLAST. Integer or floating-point values may be used. If a different matrix is specified, you should probably also specify appropriate gap penalties for that matrix.

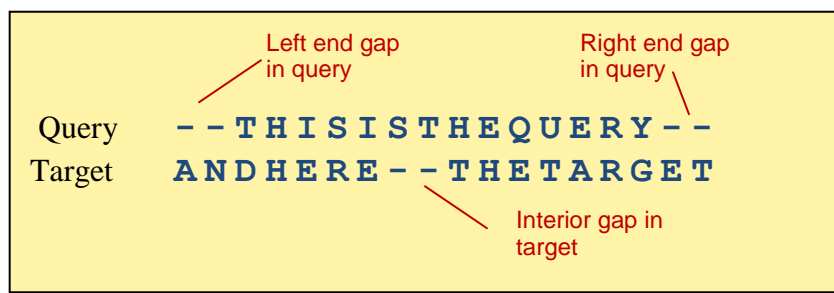
Gap penalties for local alignments

The `--loopen` and `--lertext` options specify open and extend penalties for local alignments.

Penalty	Option	Default
Local gap open	<code>--loopen</code>	10.0
Local gap extend	<code>--lertext</code>	1.0

Gap penalties for global alignments

Up to 12 separate penalties can be specified: all combinations of query / target, left / interior / terminal, and open / extend can be assigned different penalties.



Default penalties are as follows.

Penalty	Default
Interior gap open	10.0 nt, 17.0 aa
End gap open	1.0
Interior gap extend	1.0
End gap extend	0.5

End gaps (also called terminal gaps) are penalized much less than interior gaps, which is typically appropriate when fragments are aligned to full-length sequences. These defaults can be changed using the `--gapopen` and `--gapext` options. The nucleotide defaults would be set using these options:

```
--gapopen 10.0I/1.0E --gapext 0.5
```

A numerical value for a penalty is optionally followed by one or more letters that specify particular types of gap. Here, "10.0I" means "Interior gap=10.0", and "1.0E" means "End gap=1.0". If no letters are given after the numerical value, then the penalty applies to all gaps. More than one letter can be specified, so for example "0.5IE" means "Interior and End gap=0.5", which is the same as all gaps. Following are valid letters: I=Interior, E=End, L=Left, R=Right, Q=Query and T=Target. If more than one numerical value is specified, then they must be separated by a slash character '/'. White space is not allowed. If a star (*) is used as the numerical value, then the gap is forbidden. Using * in an open penalty means that the gap will never be allowed, using * in an extension penalty means that gaps longer than one will be forbidden. So, for example, *LQ in `--gapopen` means "left end-gaps in the query are not

allowed". A sign (plus or minus) is not allowed in the numerical value, which can be integer or floating-point (in which case a period '.' must be used for the decimal point). The --gapopen and --gapext options are interpreted first by setting the defaults, then by scanning the string left-to-right. Later values override previous values.

The final settings are written to the --log file, and I strongly recommend that you use this information to check that your options are correctly formatted. Here is another set of example options.

```
--gapopen 10.0QL/*QL/2.0TE/1.0QR --gapext 0.5I/0.1E
```

The resulting penalties appear as follows in the log file.

```
10.00 Open penalty (query, internal)
    * Open penalty (query, left end)
    1.00 Open penalty (query, right end)
10.00 Open penalty (target, internal)
    2.00 Open penalty (target, left end)
    2.00 Open penalty (target, right end)
    0.50 Ext. penalty (query, internal)
    0.10 Ext. penalty (query, left end)
    0.10 Ext. penalty (query, right end)
    0.50 Ext. penalty (target, internal)
    0.10 Ext. penalty (target, left end)
    0.10 Ext. penalty (target, right end)
```

Considerations when using non-standard gap penalties

The --gapopen and --gapext options do not always work well with the fast alignment heuristics that are enabled by default. In some cases, especially if some gap types are forbidden, then this can cause USEARCH to crash because no alignment is possible, and this condition is currently not handled gracefully (this is a really a bug; better would be to reject the target, but this is hard to implement).

If possible, the best thing to do is to disable the heuristics by using --nofastalign. Then the gap penalties should work well. If you have very large datasets and heuristics are needed, then I recommend testing on a small subset and reviewing the --blastout file to make sure that the alignments look reasonable for your application.

Sequence identity

Sequence identity can be defined in many different ways; see for example this web page and its literature references: http://openwetware.org/wiki/Wikiomics:Percentage_identity. Identity is usually defined to be a ratio where the numerator is the number of identities (columns containing the same letter) in an alignment. Many choices are possible for the denominator, each of which has pros and cons in different applications. Common choices include:

- The number of columns in the alignment (terminal gaps may be included or excluded).
- The length of the shorter sequence.
- The length of the longer sequence.
- The average sequence length.
- The number of columns containing letter pairs (i.e., gaps are ignored).

Terminal gaps

Some definitions of identity treat terminal gaps as special cases. This can be important, e.g. if fragments are being aligned to full-length sequences, in which case terminal gaps are experimental artifacts rather than evidence of insertions or deletions. It should be noted that definitions of identity that count terminal gaps differently from internal gaps are more sensitive to details of the algorithm used to generate the alignment, and in particular to gap penalties. Problems may be caused if a short motif is misaligned close to a terminal, like this.

```
Query:  -XX-----XXXXXXXXXXXXXXXX-----  
Target:  XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
```

Presumably, the correct alignment would look more like this:

```
Query:  -----XXXXXXXXXXXXXXXX-----  
Target:  XXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
```

If gapped columns count as differences and terminal gaps are discarded, then the first alignment may have much lower identity.

The `--iddef n` option specifies how identity should be calculated, where *n* is 0, 1 ... etc. The default is `--iddef 0`.

Default definition of identity, `--iddef 0`

The default definition, `--iddef 0`, uses the length of the shorter sequence as the denominator. This definition is the one used by the CD-HIT program and was originally used by UCLUST to facilitate comparison of the two programs.

Since all gap columns are discarded, this definition can report 100% identity despite gaps in the shorter sequence. Consider the following example.

```
Query:  SEQ-ENCE
Target: SEQUENCE
```

Here, there are 7 identities and the length of the shorter sequence is also 7, giving $\text{Id} = 7/7 = 100\%$.

All-diffs definition, --iddef 1

The all-diffs definition (--iddef 1) considers every gap column and every mismatch to be a difference, which is achieved by using the number of columns in the alignment as the denominator. This is the same as $1 - (\text{edit distance})$, where edit distance is the smallest number of insertions, deletions and substitutions that transform one sequence into the other. In the above example, there are 8 columns in the alignment, so $\text{Id} = 7/8 = 87.5\%$.

Internal diffs definition, --iddef 2

The internal diffs definition (--iddef 2) is similar to all-diffs, except that terminal gaps are not included in the alignment length. See above (*Terminal gaps*) for a discussion of a potential problem with this definition. This may be more appropriate if fragment sequences (e.g., partial 16S genes from a short-read sequencing experiment) are aligned to full-length sequences (complete genes in a reference database). Consider this example.

```
Query:  ---V-NC-
Target: SEQUENCE
```

Here, there are 4 columns after terminal gaps are discarded, so the internal diffs $\text{Id} = 3/4 = 75\%$, while the default $\text{Id} = 3/3 = 100\%$ and the all-diffs $\text{Id} = 3/8 = 37.5\%$.

Marine Biology Lab definition, --iddef 3

The MBL definition (--iddef 3) is similar to all-diffs, except that a gap of any length (i.e., consecutive series of gap columns) counts as a single difference. Both internal and terminal gaps are counted. Identity is defined as:

$$1.0 - [(\text{mismatches} + \text{gaps})/(\text{longer_sequence_length})]$$

Notice that unlike other definitions, this does not use the number of identities as the numerator. Consider the following example.

```
Query:  --QVDNC-
Target: SEQUENCE
```

Here, mismatches = 1 and gaps = 2 so $\text{Id} = 1 - (1 + 2)/8 = 72.5\%$. In theory, this expression can be negative, in which case it is considered to be zero.

BLAST definition, --iddef 4

This definition is the one used by NCBI BLAST: the number of columns in the alignment that contain identical matches divided by the total number of columns. If the alignment is global, terminal gaps are included in the total number of columns, so this definition is probably not suitable if fragment sequences are aligned to full-length sequences using global alignments.

UCHIME: Chimeric sequence detection

UCHIME searches for chimeric sequences. While some people have found UCHIME to be useful in practice, please note that the present version of UCHIME is a work in progress; I do not consider it mature enough for publication. I am hoping to improve the algorithm in the near future.

The basic command-line is:

```
usearch --uchime myseqs.fasta --report myseqs.rep --reportx myseqs.rep  
[--db ref.fasta] [--minh 0.6] [--minsnp 8]
```

The `--db` option specifies a trusted reference database. If not specified, *de novo* detection is performed. In *de novo* mode, chimeric triples are reported and additional evidence is required to determine which of the three is the true chimera. A good heuristic is to designate the least abundant sequence as the chimera. Input and reference sequences should be in unaligned FASTA format (no gaps). The `--minh` option specifies the minimum score to be considered a hit. Default is `--minh 0.6`, which gives an error rate $< 1\%$ on full-length simulated sequences.

The `--minsnp` option specifies the minimum number of SNPs that must be present on each side of the cross-over. The default value is 8. For short reads, this value may be too large, for example for a typical 454 read of length ~ 250 nt, this means that the divergence on each side must be at least $8/125 = 6.4\%$. So for short reads, a smaller value such as `--minsnp 3` or `--minsnp 2` is recommended, though detection is harder in shorter reads and this may produce more false positives.

SNPs and voting

UCHIME constructs a three-way alignment of two putative step-parents (A and B) to the query sequence (Q). A is always the first step-parent that models the left side of the query, B is the second (right side). UCHIME counts columns in the three-way alignment containing differences (SNPs). SNPs in which two sequences agree and the third differs either support or contradict the model. SNPs in which all three sequences differ do not provide direct evidence but could reduce confidence by indicating noise (sequence errors, inaccurate alignment or divergence between the true parent and the step-parent).

Summary report format (`--report` option)

In the summary report produced by the `--report <filename>` option, the top hit (T) is also reported. The top hit is the most similar sequence to the query, and acts as a control: the closer the top hit is to the model, the lower the confidence in the prediction. Often, but not always, the top hit is one of the step-parents in the model. In the Labels column, there are either three or four labels: Q(A,B) or Q(A,B,T). The top hit is indicated by a + prefix on the label. Here is some sample output.

>	H	Div	L.SNPs y/n/?	R.SNPs y/n/?	IdP	IdT	IdAB	IdM	Labels
Y	1.172	1.010	A14/b0/N0/?0	B187/a0/N0/?	99.0%	99.1%	86.2%	100.0%	7000004131498723 (S000129426,+7000004131498722)
Y	1.100	1.023	A35/b0/N2/?0	B32/a0/N1/?0	97.6%	99.0%	95.4%	99.8%	7000004131499276 (7000004131499535,S000004313,+7000004131499279)
Y	1.058	1.009	A17/b0/N0/?0	B13/a0/N0/?0	99.2%	99.2%	98.0%	100.0%	7000004131500055 (+7000004131500053,7000004131500110)
Y	0.966	1.008	A136/b0/N0/?	B12/a0/N0/?0	99.2%	99.4%	90.5%	100.0%	7000004131313504 (7000004131313502,S000546730,+7000004131319294)
Y	0.933	1.008	A55/b0/N0/?0	B11/a0/N0/?0	99.2%	99.2%	95.2%	100.0%	S000439503 (+S000439506,S000364348)
Y	0.791	1.007	A76/b0/N1/?0	B10/a0/N0/?0	99.3%	99.3%	94.3%	99.9%	7000004130820864 (+7000004130820865,S000436035)

```

Y 0.790 1.007 A54/b1/N0/?0 B10/a0/N0/?0 99.3% 99.3% 95.7% 99.9% 7000004130820865 (+7000004130820864,S000390059)
Y 0.748 1.007 A10/b0/N0/?0 B191/a2/N0/? 99.0% 98.4% 85.9% 99.7% S000474146 (S000541573,+S000539440)
Y 0.736 1.006 A52/b0/N0/?0 B9/a0/N0/?0 99.4% 99.4% 95.8% 100.0% 7000004128190036 (S000391339,S000388296,+S000391451)
Y 0.691 1.019 A30/b0/N1/?0 B28/a2/N0/?0 97.9% 97.8% 95.7% 99.8% 7000004131499453 (S000470254,7000004131499640,+7000004131499276)
Y 0.651 1.006 A21/b0/N0/?0 B8/a0/N0/?0 99.5% 98.4% 98.0% 100.0% S000008957 (S000544247,+S000005612)
Y 0.641 1.005 A23/b0/N0/?0 B8/a0/N0/?0 99.5% 99.2% 98.0% 100.0% 7000004131499449 (+7000004131499640,7000004129457926)
Y 0.628 1.005 A8/b0/N0/?0 B75/a0/N1/?0 99.4% 99.4% 94.5% 99.9% 7000004131503051 (7000004129457926,7000004131503121,+7000004131503117)
Y 0.604 1.006 A8/b0/N0/?0 B10/a0/N0/?0 99.2% 98.8% 98.5% 99.7% S000544231 (S000142900,+S000003569)

```

Detailed report format (--reportx option)

The detailed report shows three-way alignments and summary statistics for each model. Here is an example alignment.

```

Query ( 209 nt) FW1022_2nd_NO3Lac_F3SGFHA01AFHGD
ParentA ( 207 nt) FW1023_2nd_NO3Lac_F3SGFHA01DNBN3
ParentB ( 208 nt) FW1022_2nd_NO3Lac_F3SGFHA01CZS9S

A 1 tgcgcaggcgggttatataagacagatgtgaaat-ccccgggctcaacctgggaactgcattagtgtgactgtatagctagag 79
Q 1 tgcgcaggcgggttatataagacagatgtgaaatccccgggctcaacctgggaactgcattagtgtgactgtatagctagag 80
B 1 tgcgcaggcgggtTtGtaagacagGCgtgaaat-ccccgggctcaacctgggaactgcGCtTgtgactgCaCagctagag 79
SNPs A A AA AA A A
Model .....AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA

A 80 tgcggcagagggggatggaattccGcgtgtagcagtgaaatgcgtGgaTatgCggaggaacaccgatggcgaaggca-AT 158
Q 81 tgcggcagagggggatggaattccacgtgtagcagtgaaatgcgtagagatgtggaggaacaccgatggcgaaggcagcc 160
B 80 tAcggcagagggggGtggaattccacgtgtagcagtgaaatgcgtagagatgtggaggaacaccgatggcgaaggcagcc 159
SNPs A A B B B BB
Model AAAAAAAAAAAAAAxxxxxxxxxBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB

A 159 cccctgggccTGCactgacgctcatgcacgaaagcgtgggGagcaaaaca 207
Q 161 cccctgggcgatactgacgctcatgcacgaaagcgtgggtagcaaaaca 209
B 160 cccctgggcgatactgacgctcatgcacgaaagcgtgggtagcaaaaca 208
SNPs BBB B
Model BBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBBB.....

Score 2.838
Ids. QA 95.2%, QB 94.7%, AB 89.9%, QModel 100.0%, Div. +5.1%
SNPs Left A=11 (100%)/b=0/N=0/?=0, Right B=10 (100%)/a=0/N=0/?=0

```

SNP classification

Letters that agree are shown in lower-case, disagreements are shown in upper-case (in order to stand out -- I think upper case stands out better against lower-case than vice versa). SNPs are classified as follows. Here, a column is written as AQB, and upper case indicates disagreement. These are the possible types of column.

aqb	Q=A=B	Not a SNP.
aqB	Q=A, Q!=B	Votes for model if on the left (symbol 'A') or against model if on right (symbol 'a').
Aqb	Q=B, Q!=A	Votes for model if on the right (symbol 'B') or against model if on left (symbol 'b').
aQb	A=B, Q!=A, Q!=B	Indicates SNP in query vs. (step-)parents. Could be a read error, or a change in one of the true parents vs. the closest step-parent. Symbol 'N'.
AQB	All different.	Not directly informative, interpreted as noise. Symbol '?'.

Advantages of UCHIME

Compared with previous methods, UCHIME is very fast: it can process hundreds or thousands of sequences per minute, depending on length. Also, unlike some other methods it does not require a trusted (i.e., chimera-free) reference database or a pre-built multiple alignment of trusted

sequences. Preliminary tests suggest that on full-length 16S rRNA sequences, UCHIME has better sensitivity and lower error rates than any other method. UCHIME provides a score for each predicted chimera which allows the user to set a score threshold. Adjusting the threshold increases sensitivity or reduces the false-positive rate, similar to a BLAST e-value. At reasonable thresholds, the false-positive rate is consistently low.

Limitations of UCHIME

Like most other methods, UCHIME currently only searches for "bimeras" (two parent sequences), but "multimeras" can also be important in practice [[doi 10.2144/000113219](https://doi.org/10.2144/000113219)]. Preliminary tests suggest that ChimeraSlayer has significantly better sensitivity than UCHIME in two scenarios: (i) with short reads, and (ii) when parent sequences are not available in the reference database so that related sequences ("step-parents") are needed to construct a model of the query sequence. Here, "short" means less than about 500 nucleotides.

While several avenues for improving detection accuracy remain to be explored, I believe UCHIME can, and therefore should, at least match the performance of ChimeraSlayer on bimeras before moving on. This is because UCHIME currently uses a similar strategy to ChimeraSlayer that constructs a three-way alignment of putative step-parents to a query sequence. If UCHIME is less sensitive than ChimeraSlayer, then one or more phases of UCHIME must be sub-optimal (search for candidate step-parents, alignment construction or model scoring). Optimizing the "three-way / bimeras" algorithm will inform attempts to develop a better "multi-way / multimeras" algorithm, which is my long-term goal.

If you are working with short reads, then a practical compromise could be to use UCHIME to find candidate chimeras using a threshold designed to maximize sensitivity at the expense of a higher error rate. The candidates could then be passed to the much slower ChimeraSlayer for final discrimination. (Note that the complete CMCS pipeline should be used, not the ChimeraSlayer perl script alone).

Memory requirements

The amount of memory needed is approximately 10x the size of a FASTA file containing the database. If you are clustering, the database is the final set of seed sequences, which can be written to a FASTA file by using `--seedsout`. A more accurate estimate is:

$$(9 \times \text{the number of letters in all sequences}) + (1 \times \text{the number of letters in all labels})$$

The amount of memory required can be reduced in a number of ways, as follows.

Reduce redundancy

If you have very similar sequences in your database, then it could pay to reduce redundancy by clustering at a high identity, say 98% or 99%. This, of course, can be done using UCLUST to pre-process your database. For sure, if you have 100% identical sequences these should be deleted since they can adversely affect sensitivity in a U-sorted search.

Trim sequence labels

Sequence labels, i.e. the characters following `>` in a FASTA file, are stored as-is in memory. If your labels are long and your sequences are short, then the amount of memory required for labels may be a significant fraction of the total memory requirement. This is true for example of the NCBI NR protein database, which has many very long labels. In such cases it pays to reduce the label size. For example, you could label your sequences with an integer or some other short string that can be used as a key for retrieving longer annotations in a post-processing step. You can use the `--trunclabels` option to trim labels by discarding any text after the first white space (blank or tab).

Use database stepping


The `--dbstep n` option reduces the memory required by a factor of roughly *n* for large databases (less for smaller database). However, sensitivity tends to be reduced when clustering or searching at lower identities (say, below 80%). Using `--dbstep` reduces the number of processor operations required to search the in-memory database index, which might be expected to improve speed, but in practice execution times are often slower due to a reduction in cache coherence.

Split the database

You can split the database into smaller pieces. This allows you to parallelize a search (e.g. by running the query against *N* pieces in parallel on *N* machines in a cluster, or to serialize (by running one piece after the other on a single machine). Splitting the database may also have the advantageous side-effect of improving sensitivity. The very high speed of the USEARCH algorithm is achieved by limiting explicit sequence comparisons to a small subset of the database having the most unique words in common with the query sequence. As the database size grows, more spurious sequences will tend to appear in this subset and sensitivity may be reduced as a result.

Two-level search

If finding the closest possible match to a very large database is important in your application, then you can combine the "reduce redundancy" and "split" strategies to achieve improved speed, reduced memory use and (usually, but not always) higher sensitivity. The idea is to search first in a low-redundancy database (LRD). Sequences in the LRD are annotated with the name of a second-level database (SLD) which has more closely-related sequences. There are several SLDs that, when combined, contain the full set of sequences. In the second pass, the query is searched against the SLD identified in the first search.



This picture is over-simplified: we don't want a separate SLD for every sequence in the low-redundancy DB. There are two reasons for this: if the SLD is too small, we lose the advantage of the high search speed of USEARCH because there will be too much overhead setting up each query. Also, we want to group related families into a single SLD because otherwise the hit to the LRD may not correctly identify the SLD with the closest possible match.

To create the databases, I suggest the following approach.

1. Cluster at a fairly low identity; say 50% for proteins or 80% for nucleotides.
2. Pick a desired size for an SLD, say $1/N$ of the full database. If a cluster from step 1 is larger than this, you can split it by clustering at a higher identity, or go back and re-cluster the entire database at a higher identity.
3. Merge clusters from step 1 to create the SLDs (SLD1, SLD2 ... SLDN). This can be done by a simple greedy algorithm which can be implemented in a script, let me know if you'd like help with this. Label each sequence with the name of its SLD (this is so that the SLD name is available in step 5 below where the LRD is created).
4. Cluster each SLD at a high identity; say 98% for nucleotides or 90% for proteins.
5. Combine all the seeds from step 4 above, this produces the LRD.

To run a two-pass query, first search the query sequences against the LRD. Then divide them according to the SLD identified by the LRD hit and run each subset against its SLD; this of course can be done serially or in parallel.

Command line reference

Algorithms

Algorithm	Description	Required options
UCLUST	<i>De novo</i> clustering	--cluster <i>fastafile</i>
UCLUST	Search + clustering	--cluster <i>fastafile</i> --db <i>fastafle</i>
UBLAST	Database search	--query <i>fastafile</i> --db <i>fastafile</i>
UCHIME	<i>De novo</i> chimera detection	--uchime <i>fastafile</i>
UCHIME	Chimera detection with reference database	--uchime <i>fastafile</i> --db <i>fastafile</i>

Sorting

Command	Command line
Sort sequences by length	--sort <i>fastafile</i> --output <i>fastafile</i> --mergesort <i>fastafile</i> --output <i>fastafile</i> [--split <i>size</i>]
Sort UCLUST file by cluster nr.	--sortuc <i>ufile</i> --output <i>ucfile</i>

File format conversions

From	To	Command line
UCLUST (.uc)	FASTA	--uc2fasta <i>ucfile</i> --output <i>fastafile</i>
UCLUST (.uc)	Annotated FASTA	--uc2fastax <i>ucfile</i> --output <i>fastafile</i>
UCLUST (.uc)	CD-HIT (.clstr)	--uc2clstr <i>ucfile</i> --output <i>clstrfile</i>
CD-HIT (.clstr)	UCLUST (.uc)	--clstr2uc <i>clstrfile</i> --output <i>ucfile</i>

Output files

Option	Format	Description
--uc <i>filename</i>	UCLUST	Tab-separated file designed primarily for clustering pipelines but can also be useful for search. One record for each input sequence giving its cluster assignment, identity and alignment; and one record for each cluster giving its size and average identity. Supported by UCLUST and UBLAST.
--blastout <i>filename</i>	BLAST-like	Human-readable format similar to BLAST. Supported by UCLUST and UBLAST.
--blast6out <i>filename</i>	Tab-separated	Tabbed format with one record per hit. Compatible with the -m8 or -outfmt 6 option of NCBI BLAST. Supported by UCLUST and UBLAST.
--userout <i>filename</i>	Tab-separated	Tabbed format with one record per hit, fields specified by the --userfields option (see manual). Supported by UCLUST and UBLAST.
--seedsout <i>filename</i>	FASTA	Seed sequences, i.e. the non-redundant or reduced redundancy set of sequences after clustering. Supported by UCLUST only.
--fastapairs <i>filename</i>	FASTA	Pair-wise alignments in FASTA format. Supported by UCLUST and UBLAST.

Database search order

Option	Description
--[no]usort	[Do not] test database sequences in U-sorted order, i.e. in order of decreasing number of words in common. If --nousort is specified, the entire database is tested and search termination options are ignored or give an error. Default is --usort.
--[no]ssort	Change U-sort order to better correlate with evolutionary distance. Applies to amino acid databases only. If --query and --evaluate are used, then --ssort is the default, otherwise --nossort is the default.
--stable_sort	Specifies that a stable algorithm should be used for U-sorting. This may be a little slower, but gives reproducible results when a given query has the same word count with more than one target sequence, which can cause the accepted target to change if a non-stable sort is used. Default is to use a non-stable sort.
--w n	Word length for U-sorting. Default 5 for amino acids, 8 for nucleotides. Changing the word length changes speed, sensitivity and memory requirements in different ways depending on the input data.

Search termination

These options determine when a U-sorted or S-sorted search terminates. These options are ignored if `--noursort` is specified, in which case the entire database is searched.

Option	Description
<code>--maxaccepts</code> <i>n</i>	Maximum number of accepted targets. Zero means that this option is ignored (i.e., zero means infinity). Default 1, unless <code>--maxtargets</code> is specified, in which case the default is zero. Increasing <code>--maxaccepts</code> improves sensitivity and also the probability that the best possible hit is found, at the expense of slower times. If <code>--maxaccepts</code> is increased, you should generally increase <code>--maxrejects</code> also.
<code>--maxrejects</code> <i>n</i>	Maximum number of rejected targets. Zero means that this option is ignored (i.e., zero means infinity). Default is 8, unless <code>--maxtargets</code> is specified, in which case the default is zero. Increasing <code>--maxrejects</code> improves sensitivity by reducing the probability of a false negative, i.e. failing to find a possible hit, at the expense of slower times.
<code>--maxtargets</code> <i>n</i>	Maximum number of targets to test. Zero means that this option is ignored (i.e., zero means infinity). If this option is specified, the default <code>--maxaccepts</code> and <code>--maxrejects</code> parameters are changed to zero (i.e., are ignored by default). This option is useful when searching for low-identity relationships, especially in large databases, in which case U-sort order correlates poorly with similarity and it is important to test a large number of target sequences. In such cases, values such as <code>--maxtargets</code> 100 or 1000 might be used.

Accept / reject criteria for UCLUST and UBLAST

Accept / reject criteria specify one or more sequence similarity thresholds. At least one of --id or --evalue must be specified. Thresholds are combined with AND, so all must be satisfied for a target to be accepted.

Option	Description
--id <i>f</i>	Minimum identity, as a value 0.0 to 1.0, meaning 0% to 100% identity. The --iddef option determines how identity is defined. There is no default value.
--evalue <i>E</i>	Maximum E-value. There is no default value.
--queryfract <i>f</i>	Minimum fraction of the query sequence that is covered by its alignment to the target, as a value 0.0 to 1.0, meaning 0% to 100% coverage. Coverage is defined as the number of letters in the query that are aligned to letters in the target, divided by the length of the query sequence. Default is 0.0.
--targetfract <i>f</i>	Minimum fraction of the target sequence that is covered by its alignment to the query, as a value 0.0 to 1.0, meaning 0% to 100% coverage. Coverage is defined as the number of letters in the target that are aligned to letters in the query, divided by the length of the target sequence. Default is 0.0.
--[no]wordcountreject	[Do not] reject a target sequence if it has too few words in common. The number of words in common is used to estimate identity, which is faster than calculating identity from an alignment, but can give some false negatives. Applies only to global alignments and on if --id is used as an accept threshold. Default is --wordcountreject.
--iddef <i>n</i>	Definition of sequence identity. See manual for details. Default 0.

Weak match criteria in UBLAST

Weak matches are reported in output files but are not considered accepts, will not terminate a U-sorted search, and do not match a query to a cluster. Weak criteria are also combined with AND. Weak matches will also be reported by UCLUST, though this is rarely useful in practice.

Option	Description
--weak_id <i>f</i>	Minimum identity, as a value 0.0 to 1.0, meaning 0% to 100% identity. The --iddef option determines how identity is defined.
--weak_evalue <i>E</i>	Maximum E-value. There is no default value.

Alignment style options

Option	Description
--global	Global alignments. This is the default for UCLUST, i.e. if --cluster is specified.
--local	Local alignments. This is the default for UBLAST, i.e. if --query is specified.
--[no]gapped	[Do not] make gapped alignments. If --nogapped is specified, ungapped alignments will be created. The --nogapped option cannot be used if --global is specified. Default is --gapped.

Alignment scoring parameters

Note that if you change these parameters, then E-values will not be calculated correctly unless the K and Lambda parameters for E-value calculation are adjusted accordingly. If you don't need E-values, e.g. because you use global identity as your similarity measure, then you don't need to adjust K and Lambda.

Option	Description
--match <i>s</i>	Match score for nucleotides. Default 1.0. Must be > 0.
--mismatch <i>s</i>	Mismatch score for nucleotides. Default -2.0. Must be < 0. For local alignments, absolute value should be greater than --match.
--lopen <i>s</i>	Gap open penalty for local alignments. Default 10.0. Must be > 0.
--lEXT <i>s</i>	Gap extension penalty for local alignments. Default 1.0. Must be > 0.
--gapopen <i>spec</i>	Specifies gap open penalties for global alignments.
--gapEXT <i>spec</i>	Specifies gap extension penalties for global alignments.
--matrix <i>filename</i>	File name of amino acid substitution matrix in NCBI BLAST format. Default is BLOSUM62.

Alignment heuristics

Option	Description
--k <i>n</i>	Word length for alignment seeds. Default 3 for amino acids, 4 for nucleotides.
--minhsp <i>n</i>	Minimum length of HSP. Default 32. In versions up to 4.0.40, this could prevent short sequences from matching. In v4.0.41 and later, the minimum length of an HSP is automatically set to half the shorter sequence length if this is < <i>n</i> .
xdrop_u <i>s</i> xdrop_g <i>s</i> xdrop_ug <i>s</i> xdrop_nw <i>s</i>	X-drop parameters for extending alignments. If the value of (maximum alignment score found so far) – (current score) > X-drop, alignment extension is terminated. Smaller values are faster, but will miss more opportunities to find a higher scoring alignment by continuing to extend. With global alignments, smaller values are faster in the HSP-finding stage, but may result in slower overall times due to longer regions between HSPs that must be aligned by banding. xdrop_u : for ungapped local alignments, default 16.0. xdrop_ug : for ungapped local alignments used to trigger gapped extensions, default 16.0. xdrop_g : for gapped extensions of local alignments, default 32.0. xdrop_nw : for finding local HSPs in a global alignment, default 16.0.
--band <i>n</i>	Radius of band for banded dynamic programming, which is used to align regions between HSPs in a global alignment. Smaller values are faster, but may tend to produce less accurate alignments. Default 16.
--[no]twohit	[Do not] use two-hit word seeding. Two-hit seeding requires that two matching words are found on a single diagonal with maximum distance set by the --max2 option. With --notwohit , a single word match triggers an extension. Default is --twohit .
--[no]nb	[Do not] use word neighborhoods for seeding alignments. Applies to amino acids only. If U-sorting is enabled, default is --nonb , otherwise the default is --nb . Using --nb is typically a little slower, but not by much and is more sensitive for low-identity matches.
--max2 <i>n</i>	Maximum distance between two word seeds on a diagonal. Ignored if --notwohit is set. Default 40.
--seedt1 <i>t</i>	Minimum score of a word seed. Used if single-hit seeding is specified (--notwohit) and word neighborhoods are enabled. Default 13.0.
--seedt2 <i>t</i>	Minimum score of a word seed. Used if word neighborhoods are enabled and two-hit seeding is used. Default 11.0.

Karlin-Altschul statistics and E-value calculation

Option	Description
--ka_dbsize	Effective database size, in letters. If the database has high redundancy, the effective size should be set to a value smaller than the actual size. Default is the actual size of the database.
--ka_gapped_k	K parameter for gapped local alignments. Default 0.041 for amino acids, 0.460 for nucleotides.
--ka_ungapped_k	K parameter for ungapped local alignments. Default 0.128 for amino acids, 0.621 for nucleotides.
--ka_gapped_lambda	Lambda parameter for gapped local alignments. Default 0.267 for amino acids, 0.128 for nucleotides.
--ka_ungapped_lambda	Lambda parameter for ungapped local alignments. Default 0.331 for amino acids, 1.330 for nucleotides.

Other options

Option	Description
--rev	Search reverse strand. Default search plus strand only. Applies to nucleotide databases only.
--[no]output_rejects	[Do not] output rejects to the --uc file. Useful for trouble-shooting cases where an expected match is not found. Default is --nooutput_rejects.
--mincodons <i>n</i>	Minimum number of amino acids in a predicted ORF. Default 20.
--usersort	Specifies that a user-defined sort order is used for the input sequences. Applies to clustering only. Default is to require that input sequences are sorted by decreasing length.
--stepwords <i>n</i>	Select a subset of words from query so that the expected number of words in common with an accepted target is <i>n</i> . Default 8. Zero means that query stepping is disabled, so all query words will be used. Larger values (or zero) tend to improve sensitivity at the expense of slower speed.
--dbstep <i>n</i>	Index every <i>n</i> th word in the database. Default is <i>n</i> =1, i.e. all words are indexed. Using this option reduces the memory required by a factor of roughly <i>n</i> . However, sensitivity tends to be reduced when clustering or searching at lower identities (say, below 80%). Using --dbstep reduces the number of processor operations required to search the in-memory database index, which might be expected to improve speed, but in practice execution times are often slower due to a reduction in cache coherence.
--bump <i>n</i>	Optimization for U-sorting, specified as integer percentage. Default 50. Zero means disabled. Larger values tend to improve speed at the expense of lower sensitivity.

Option	Description
--split <i>s</i>	Size of partition for --mergesort, in Mb. Default 1000.0, i.e. 1Gb.
--quiet	Do not show progress messages to standard error output while executing.
--log <i>logfile</i>	Log file name. Contains information about parameters and performance.
--version	Write program version number and exit.
--help	Write summary of command line options and exit.