

Minia-GATB — Short manual

R. Chikhi & G. Rizk & E. Drezen & D. Lavenier
rayan.chikhi@ens-cachan.org

April 1, 2015

Abstract

Minia is a software for ultra-low memory DNA sequence assembly. It takes as input a set of short genomic sequences (typically, data produced by the Illumina DNA sequencer). Its output is a set of contigs (assembled sequences), forming an approximation of the expected genome. Minia is based on a succinct representation of the de Bruijn graph. The computational resources required to run Minia are low: a human genome can be assembled using 4 GB of memory.

Contents

1 Forewords	1
2 Installation	2
3 Parameters	2
4 Explanation of parameters	2
5 Input	3
6 Output	4
7 Memory usage	4
8 Disk usage	4
9 Data structure implementation	4

1 Forewords

Minia-GATB is the codename of a new version of Minia using the GATB library. We will drop the "-GATB" suffix and just refer to it as "Minia". In terms of features, not much has changed: Minia remains a contigs assembler using very low memory. Notable changes are:

- Different command line format

- Much faster k -mer counting step
- Bubbles are collapsed using the most covered path (coverage information is now available for each node in the graph)
- Unitigs can be generated

2 Installation

Download the binary and run Minia using the command `./minia`.

To install Minia from the sources, enter this commands in the sources folder:
`mkdir build && cd build && cmake .. && make -j 4`
 Minia has been tested on Linux and MacOS systems.

3 Parameters

The basic usage is:

```
./minia -in [input file] -kmer-size [kmer size] \
        -abundance-min [abundance value] -out [prefix]
```

An example command line is:

```
./minia -in reads.fastq -kmer-size 31 -abundance-min 3 -out minia_assembly_k31_m3
```

The main parameters are:

1. `in` – the input file(s) (see Section 5 for inputting multiple files)
2. `kmer-size` – k -mer length (integer)
3. `abundance-min` – specifies how many times a k -mer must be seen in the reads to be considered correct (integer)
4. `prefix` – any prefix string to store output contigs as well as temporary files for this assembly

4 Explanation of parameters

`kmer-size` The k -mer length is the length of the nodes in the de Bruijn graph. The default value is 31, yet it is almost certainly suboptimal. The most adequate value of k varies from one dataset to another (it depends on coverage, sequencing errors, genome structure). We recommend that you use the KmerGenie software to automatically estimate the best k for a dataset. The KmerGenie article describes in more details the role of k in assembly.

abundance-min The **abundance-min** is used to remove erroneous, low-abundance k -mers. This parameter also strongly depends on the dataset. Any k -mer that appears strictly less than *abundance-min* times in the reads will be discarded. A typical value is 3. It is also called "coverage cut-off" in other software (e.g. Velvet).

KmerGenie has an experimental feature (in the output of the command, not in the HTML report) that reports an optimal coverage cut-off (synonymous to abundance) parameter that exactly corresponds to the abundance-min parameter of Minia.

Setting *abundance-min* to 1 is not recommended, as no erroneous k -mer will be discarded, which will likely result in a very large memory usage. In particular, if KmerGenie tells you to set abundance to 1, then still, set it to 2. If the dataset has high coverage, try larger values.

prefix The **prefix** parameter is any arbitrary file name prefix, for example, `test_assembly`.

5 Input

Larger k -mer lengths

Minia supports arbitrary large k -mer lengths. To compile Minia from the source, to support k -mer lengths up to, say, 160, type this in the build folder:

```
rm -Rf CMake* && cmake -Dk4=160 .. && make
```

FASTA/FASTQ

Minia assembles any type of Illumina reads, given in the FASTA or FASTQ format. Giving paired or mate-pairs reads as input is OK, but keep in mind that Minia won't use pairing information.

Multiple Files

Minia can assemble multiple input files. Just create a text file containing the list of read files, one file name per line, and pass this list as the first parameter of Minia (instead of a FASTA/FASTQ file). Therefore the parameter **input_file** can be either (i) the read file itself (FASTA/FASTQ, gzipped or not), or (ii) a file containing a list of file names.

Line format

In FASTA files, each read can be split into multiple lines, whereas in FASTQ, each read sequence must be in a single line.

Gzip compression

Minia can directly read files compressed with gzip. Compressed files should end with `.gz`. Input files of different types can be mixed (i.e. gzipped or not, in FASTA or FASTQ)

Graph input

Minia can take as input a graph in the HDF5 format, constructed using the dbgh5 program, or a different GATB tool, or from a previous run of Minia (useful when one wants to tweak assembly parameters). Use the `-graph` parameter to specify the graph, instead of the `-in` parameter.

6 Output

The output of Minia is a set of contigs in the FASTA format, in the file `[prefix].contigs.fa`.

Creating unitigs

Minia supports the creation of unitigs. Use the following command line:

```
-starter simple -no-length-cutoff -traversal unitig
```

The `-starter simple` option is to turn off Minia's heuristics for starting node selection (which avoid starting inside a bubble). Note that there is a known bug in the `-starter simple` behavior: some of the unitigs of length `k` won't be returned (technically: the branching nodes which are neighbors of branching nodes are not returned).

7 Memory usage

We estimate that the memory usage of Minia is roughly 1 GB of RAM per gigabases in the target genome to assemble. It is independent of the coverage of the input dataset, provided that the `abundance-min` parameter is correctly set. For example, a human genome was assembled in 5.7 GB of RAM. This was using the original data structure; the current implementation relies on Cascading Bloom filters and uses ≈ 4 GB. A better estimation of the memory usage can be found in the Minia article.

8 Disk usage

Minia writes large temporary files during the k-mer counting phase. These files are written in the working directory when you launched Minia. For better performance, run Minia on a local hard drive.

9 Data structure implementation

Minia now uses the Cascading Bloom filters improvement (<http://arxiv.org/abs/1302.7278>) by default, thanks to Gustavo Sacomoto for the implementation in Minia. Launch Minia with the `-debloom original` option to revert to the original data structure.