

## **Analysing Genome Alignments with Gubbins**

Email: gubbins@sanger.ac.uk

### **Contents**

1. Introduction
2. Requirements
3. Installation
4. Input files
5. Running Gubbins
6. Output files
7. Viewing output files
8. Examples
9. Troubleshooting
10. Citing Gubbins
11. References

### **1. Introduction**

Gubbins (Genealogies Unbiased By recomBinations In Nucleotide Sequences) is an algorithm that iteratively identifies loci containing elevated densities of base substitutions while concurrently constructing a phylogeny based on the putative point mutations outside of these regions. Simulations demonstrate the algorithm generates highly accurate reconstructions under realistic models of short-term diversification of sequences through both point mutation and recombination, and can be run in only a few hours on alignments of hundreds of bacterial genome sequences.

Gubbins is suitable for use on alignments of hundreds of sequences. The time taken for the algorithm to converge on a stable solution increases approximately quadratically with the number of samples when RAxML is used to construct phylogenies; this increase can be ameliorated to some extent by using FastTree to construct at least the initial phylogeny. The input should be a whole genome sequence alignment; there is no need to remove accessory genome loci, as the algorithm should be able to cope with missing data.

Gubbins cannot distinguish elevated densities of polymorphisms arising through recombination from other potential causes. These may be assembly or alignment errors, mutational hotspots or regions of the genome with relaxed selection. Such false positives are more likely to arise on longer branches within a phylogeny; it is recommended that populations be subdivided into smaller groups of less diverse samples that can each be independently analysed with Gubbins. For some

discussion of potential confounding factors in the analysis of such population genomic datasets, please see (1).

## **2. Requirements**

Gubbins is a command line program designed to be run on Linux or Max OSX systems and requires Python version 2.7 or greater. Gubbins can also be run on Windows operating systems using a Bio-Linux virtual machine. Gubbins is dependent upon independent phylogenetics software (RAxML (2) and FastTree2 (3)) and sequence reconstruction software (FastML (4)) that are included in the installation package.

On large alignments (e.g. hundreds of whole genome sequences) Gubbins may take hours or days to converge and require several gigabytes of RAM. Please note that the first iteration will usually be the slowest, as the starting phylogeny is constructed using all polymorphic sites. Subsequent phylogenies do not use information from base substitutions introduced by putative recombinations and are typically faster.

In order to view the outputs of Gubbins, we would recommend:

- Seaview (<http://doua.prabi.fr/software/seaview>) for viewing sequence alignments
- FigTree (<http://tree.bio.ed.ac.uk/software>) for viewing phylogenies
- Artemis (<https://www.sanger.ac.uk/resources/software/artemis/>) for viewing annotation (.embl and .gff) files
- Any PDF viewer for viewing the figures

## **3. Installation**

Installation instructions can be found in the relevant INSTALL document included in the package.

## **4. Input files**

The main input file required for Gubbins is a FASTA format alignment. The header of each sequence should be unique. The alignment should be flush, with the '-' character used to indicate missing data or gaps. When specifying a starting tree, the Newick format should be used.

## **5. Running Gubbins**

To run Gubbins with default settings:

```
run_gubbins.py [FASTA alignment]
```

## Input options:

<b>Option</b>	<b>Description</b>
<code>--outgroup, -o</code>	The name of a sequence in the alignment on which to root the tree
<code>--starting_tree, -s</code>	A Newick-format starting tree on which to perform the first iteration analysis. The default is to compute a starting tree using RAxML
<code>--filter_percentage -f</code>	Filter out taxa with more than this percentage of missing data. Default is 25%

## Processing options:

<b>Option</b>	<b>Description</b>
<code>--tree_builder, -t</code>	The algorithm to use in the construction of phylogenies in the analysis; can be 'raxml', to use RAxML, 'fasttree', to use FastTree, or 'hybrid', to use FastTree for the first iteration and RAxML in all subsequent iterations. Default is RAxML
<code>--iterations, -i</code>	The maximum number of iterations to perform; the algorithm will stop earlier than this if it converges on the same tree in two successive iterations. Default is 5.
<code>--min_snps, -m</code>	The minimum number of base substitutions required to identify a recombination. Default is 3.

## Output options:

<b>Option</b>	<b>Description</b>
<code>--use_time_stamp, -u</code>	Include a time stamp in the name of output files to avoid overwriting previous runs on the same input file. Default is to not include a time stamp.
<code>--verbose, -v</code>	Print debugging messages. Default is off.
<code>--no_cleanup, -n</code>	Do not remove files from intermediate iterations. This option will also keep other files created by RAxML, FastML and FastTree, which would otherwise be deleted. Default is to only keep files from the final iteration.

## 6. Output files

When using RAxML in the final iteration of the analysis, output files will have the prefix:

`RAxML_result.X.Y.iteration_Z`

When using FastTree in the final iteration of the analysis, output files will have the prefix:

`X.Y.iteration_Z`

Where:

X = Prefix taken from the input FASTA file

Y = Time stamp. NOTE: This will only be included in the output file prefix if the `-u` flag has been selected

Z = Number of the iteration to which the output files relate. NOTE: this will only be added if Gubbins is run with the `"--no_cleanup"` mode.

Output file suffices:

Suffix	Description
<code>.recombination_predictions.embl</code>	Recombination predictions in EMBL file format.
<code>.recombination_predictions.gff</code>	Recombination predictions in GFF format
<code>.branch_base_reconstruction.embl</code>	Base substitution reconstruction in EMBL format.
<code>.summary_of_snp_distribution.vcf</code>	VCF file summarising the distribution of SNPs
<code>.per_branch_statistics.csv</code>	Per branch reporting of the base substitutions inside and outside recombination events.
<code>.filtered_polymorphic_sites.fasta</code>	FASTA format alignment of filtered polymorphic sites used to generate the phylogeny in the final iteration.
<code>.filtered_polymorphic_sites.phylip</code>	Phylip format alignment of filtered polymorphic sites used to generate the phylogeny in the final iteration.
<code>.final_tree.tree</code>	This file contains the final phylogeny in Newick format.

A log.txt file is also created by FastML.

## 7. Viewing output files

The “.recombination\_predictions.gff”, “.recombination\_predictions.embl” and “.branch\_base\_reconstruction.embl” files can be read onto any sequence from the alignment using the freely-available genome browsing software Artemis. Artemis can also display a summary of the SNP distribution by reading in the VCF file. The package also includes a Python script for constructing PDFs to visualise the distribution of base substitutions or recombination relative to the phylogeny:

```
gubbins_drawer.py -o [output PDF file name] -t [Newick
format tree file] [embl file]
```

The “.per\_branch\_statistics.csv” file contains summary statistics for each branch in the tree in comma delimited format. For each branch, the following statistics are reported:

<b>Statistic</b>	<b>Description</b>
Node	Name of the node subtended by the branch. This can either be one of the taxa included in the input alignment, or an internal node, which will be named with an N followed by a number.
Total SNPs	Total number of base substitutions reconstructed onto the branch
Num of SNPs inside recombinations	Number of base substitutions reconstructed onto the branch that fall within a predicted recombination (r).
Num of SNPs outside recombinations	Number of base substitutions reconstructed onto the branch that fall outside of a predicted recombination. i.e. predicted to have arisen by point mutation (m).
Num of Recombination Blocks	Total number of recombination blocks reconstructed onto the branch
Bases in recombinations	Total length of all recombination events reconstructed onto the branch
r/m	The r/m value for the branch. This value gives a measure of the relative impact of recombination and mutation on the variation accumulated on the branch.
rho/theta	The ratio of the number of recombination events to point mutations on a branch; a

	measure of the relative rates of recombination and point mutation.
Genome Length	The total number of aligned bases between the ancestral and descendent nodes for the branch excluding any missing data or gaps in either.

## 8. Examples

Two example whole genome bacterial alignments are included in the Gubbins package. These are:

File	Description
ST239.aln	An alignment of 14 methicillin resistant <i>Staphylococcus aureus</i> isolates from the MLST sequence type (ST) 239. These isolates are a subset of the isolates sequenced in (5, 6).
PMEN1.aln	An alignment of 11 <i>Streptococcus pneumoniae</i> genomes from the PMEN1 lineage. These isolates are a subset of those analysed in (7).

### Example 1: ST239


To run a Gubbins analysis of the ST239 data using the default parameters:

```
run_gubbins.py ST239.aln
```

To create a PDF showing the predicted regions of recombination against a phylogenetic reconstruction based on the final iteration of the Gubbins analysis:

```
gubbins_drawer.py -o ST239.pdf -t ST239.final_tree.tre  
RAxML_result.recombination_predictions.embl
```

The resulting PDF (below) shows the phylogeny of ST239 on the left. For each isolate, blocks representing the regions identified as recombinations by gubbins are indicated by coloured blocks. Blue blocks are unique to a single isolate while red blocks are shared by multiple isolates. The horizontal position of the blocks represents their position in the alignment.



### Example 2: PMEN1

To run an analysis of the PMEN1 data with a maximum of 10 iterations:

```
run_gubbins.py -i 10 PMEN1.aln
```

To create a PDF showing the predicted regions of recombination against a phylogenetic reconstruction based on the final iteration of the Gubbins analysis:

```
gubbins_drawer.py -o PMEN1.pdf -t PMEN1.final_tree.tre  
PMEN1.recombination_predictions.embl
```



Again, you may find that your analysis does not converge on iteration 5, so substitute the tab file with that from the last iteration of your analysis.

To create a PDF showing the reconstructed SNPs against the same tree:

```
gubbins_drawer.py -o PMEN1_SNPs.pdf -t PMEN1.final_tree.tre  
PMEN1.branch_base_reconstruction.embl
```

The two PDF files, PMEN1.pdf and PMEN1\_SNPs.pdf are shown below, illustrating the correlation between reconstructed recombination regions and SNP density.

*Gubbins manual*





## 9. Troubleshooting

*RAxML returns a non-zero exit status* – check the sequence identifiers in the FASTA alignment don't contain any unusual characters, such as '#', that are not accepted by RAxML. There may also be problems with pre-existing files having output file names; Gubbins will not automatically overwrite these, so they must be removed manually before running the analysis.

*Gubbins fails to converge on a stable solution* – this may be due to the limit of genetic resolution between very closely related samples. It may be necessary to obtain higher resolution data, if possible, or otherwise prune the dataset to remove very closely related isolates in order to obtain a stable solution.

*Gubbins takes a long time to analyse the data* – if Gubbins is taking a long time to analyse a large dataset, it is worth using FastTree rather than RAxML to construct the initial, or all, phylogenies; removing identical sequences from the dataset; removing sequences with a high proportion of missing data that may be difficult to reconstruction, or splitting the dataset into multiple sets of more closely related samples using a fast, phylogeny-independent approach (e.g. BAPS (8)).

## 10. Citing Gubbins

Croucher N. J., Page A. J., Connor T. R., Delaney A. J., McQuillan J. A., Bentley S. D., Parkhill J., Harris S.R. Rapid phylogenetic analysis of large samples of recombinant bacterial whole genome sequences using Gubbins. (submitted)

## 11. References

1. **Croucher NJ, Harris SR, Grad YH, Hanage WP.** 2013. Bacterial genomes in epidemiology—present and future. *Philos. Trans. R. Soc. B Biol. Sci.* **368**.
2. **Stamatakis A, Ludwig T, Meier H.** 2005. RAxML-III: a fast program for maximum likelihood-based inference of large phylogenetic trees. *Bioinformatics*, 2004/12/21 ed. **21**:456–463.
3. **Price MN, Dehal PS, Arkin AP.** 2010. FastTree 2—approximately maximum-likelihood trees for large alignments. *PLoS One* **5**:e9490.
4. **Ashkenazy H, Penn O, Doron-Faigenboim A, Cohen O, Cannarozzi G, Zomer O, Pupko T.** 2012. FastML: a web server for probabilistic reconstruction of ancestral sequences. *Nucleic Acids Res.* **40**:W580–W584.
5. **Harris SR, Feil EJ, Holden MT, Quail MA, Nickerson EK, Chantratita N, Gardete S, Tavares A, Day N, Lindsay JA, Edgeworth JD, de Lencastre H, Parkhill J, Peacock SJ, Bentley SD.** 2010. Evolution of MRSA during hospital

transmission and intercontinental spread. *Science*, 2010/01/23 ed. **327**:469–474.

6. **Castillo-Ramírez S, Corander J, Marttinen P, Aldeljawi M, Hanage WP, Westh H, Boye K, Gulay Z, Bentley SD, Parkhill J.** 2012. Phylogeographic variation in recombination rates within a global clone of Methicillin-Resistant *Staphylococcus aureus* (MRSA). *Genome Biol.* **13**:R126.
7. **Croucher NJ, Harris SR, Fraser C, Quail MA, Burton J, van der Linden M, McGee L, von Gottberg A, Song JH, Ko KS, Pichon B, Baker S, Parry CM, Lambertsen LM, Shahinas D, Pillai DR, Mitchell TJ, Dougan G, Tomasz A, Klugman KP, Parkhill J, Hanage WP, Bentley SD.** 2011. Rapid pneumococcal evolution in response to clinical interventions. *Science*, 2011/01/29 ed. **331**:430–434.
8. **Tang J, Hanage WP, Fraser C, Corander J.** 2009. Identifying currents in the gene pool for bacterial populations using an integrative approach. *PLoS Comput Biol*, **5**:e1000455.