

# The Jaatha HowTo

Lisha Mathew, Paul R. Staab and Dirk Metzler

Version 2.0.2

## 1 Introduction

Jaatha is a fast composite likelihood method to estimate model parameters of the evolutionary history of (at the moment) two related species or populations. To do so, it uses SNP data from multiple individuals from both species and – optionally but highly recommended – one or more outgroup sequences. This HowTo describes the method and gives an example of using its implementation as an R package `jaatha`.

The package itself can be obtained from CRAN using

```
> install.packages('jaatha')
```

or downloaded from [http://evol.bio.lmu.de/\\_statgen/software/jaatha](http://evol.bio.lmu.de/_statgen/software/jaatha). Jaatha runs on R under Windows, OS X (Mac) and Linux with the following restrictions: On Windows the parallelization and finite sites simulation using `seq-gen` is currently not supported.

A more detailed description of the algorithm can be found in Mathew et al. [2012]. Further information about the R functions used in this document can be obtained by calling `help()` with the functions name as argument.

Please cite the above paper when using Jaatha in a publication.

## 2 A demographic model

Before we can apply Jaatha to estimate parameters, we first need to create a model of the evolutionary history of the two species. Jaatha cannot account for the effects of selection, hence we assume a neutral evolution. It can estimate the effects that either “demographic” events – like an expansion of the size of one population or migration between the two populations – or molecular events – like mutation or recombination – have on the genome of the populations. To emphasize that we can not include selection, we refer to a scenario of the evolution of the two species under such events as a “demographic model”.

For now, assume that we know that our two species are closely related; hence they must have separated at a certain time in the past. There may still be gene flow ongoing between them to which we will refer to as migration from one population into the other. Hence, we could propose the simple demographic model described in Figure 1.

To specify that model in R, we first need to load `jaatha`.

```
> library(jaatha)
```

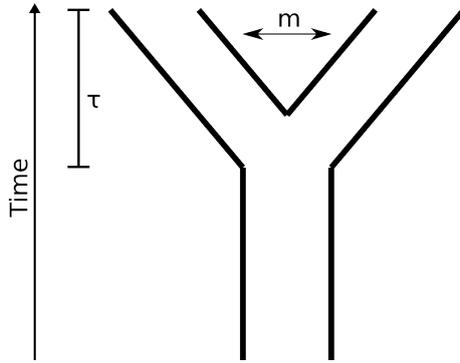


Figure 1: A simple demographic model: The ancestral population splits into two populations  $\tau$  time units ago, and afterwards individuals migrated from one population to the other with a migration rate  $M$ . Mutations are occurring with rate  $\theta$  and recombination with a known rate  $\rho$ .

We can now create an 'empty' demographic model `dm` using the `dm.createDemographicModel()` function:

```
> dm <- dm.createDemographicModel(sample.sizes=c(12,14),
+                               loci.num=50,
+                               seq.length=1000)
```

The parameter `sample.sizes` here corresponds to the number of individuals we have sampled from the first population and second population respectively. The second argument `loci.num` states that we are using data from 70 loci while `seq.length` gives the (average) length of each loci<sup>1</sup>.

We can now successively add the other assumptions of our model:

```
> dm <- dm.addSpeciationEvent(dm, .1, 5, new.time.point.name='tau')
> dm <- dm.addSymmetricMigration(dm, .01, 5, new.par.name='M')
> dm <- dm.addMutation(dm, 1, 20, new.par.name='theta')
> dm <- dm.addRecombination(dm, fixed=20)
```

The first parameter is the demographic model to which we want to add an assumption/feature. The two following numbers represent the range for the corresponding parameter. The lower border has to be strictly greater the zero, as we are using a logarithmic transformation of the parameter space. The parameters are scaled as in the popular simulation program `ms` [Hudson, 2002] that we use for simulations:

- The parameter for the *speciation* event is the split time  $\tau$ , which states how many generations ago the split of the population has occurred. As usual in population genetics, it is measured in units of  $4N_1$  generations ago, where  $N_1$  is the (diploid) effective population size of the first population.

<sup>1</sup>This is only used when a finite sites model is assumed or if intra-locus recombination is included.

- The parameter for the (symmetric) *migration* is the scaled migration rate  $M$ , which is given by  $M = 4N_1m$ , where  $m$  is the fraction of individuals of each population which are replaced by immigrants from the other population each generation.
- The *mutation* parameter  $\theta$  is  $4N_1$  times the neutral mutation rate per locus.
- Finally, the *recombination* parameter  $\rho$  is  $4N_1$  times the probability of recombination between the ends of the locus per generation.

Keep in mind that a ‘good’ model – which is one that approximates the real demographic history but is also as simple as possible – is crucial for obtaining meaningful estimates in the end. Jaatha will always try to find the parameters that make the model fit best to your data. If the model does not fit to the data at all, Jaatha will still return estimates, but they will not be meaningful.

### 3 Theoretical Background

It is important to understand the key concepts behind Jaatha before we can apply it. Like many estimation methods that rely on simulations, Jaatha tries to find the parameters that best fits <sup>2</sup> to your data by simulating artificial data for many different parameter combinations. It uses a learning algorithm to determine how the different parameter values influence the simulated data and uses that knowledge to find the best parameter combination for your data.

You can imagine Jaatha as a method that runs through the parameter space – the space of all possible parameter combinations, in our example a cube with borders from 0.1 to 5, 0.01 to 5 and 1 to 20 – simulating in a small part of the parameter space around the current position (we call this area a *block*). It then searches the new maximum of the current blocks and moves to it, builds a new block around it and so on. The search finally stops when the likelihood cannot be improved anymore or a maximal number of steps has been reached.

To compare the simulated data to the real one, Jaatha uses *summary statistics* of the data. As default, it calculates the Joint Site Frequency Spectrum (JSFS) of the data and further summarizes it by evaluating different sums over the JSFS. Please refer to Naduvilezhath et al. [2011] for a detailed description.

### 4 Importing Your Data

To run Jaatha, you need to calculate the JSFS of your data. To do so, you can use the function `calculateJsfs()`. This function accepts data imported with the `read.dna()` function from the package *ape*. It assumes that you provide a joined, aligned data set with multiple samples from two populations and – optional but highly recommended – one or more outgroup sequences. Additionally, you must provide the numbers of the sequences in the dataset that belong to population one, population two, and the outgroup, respectively.

Please consult the documentation of *ape* in order to get more information about `read.dna()`. For example, the import of the data could look like this:

---

<sup>2</sup>for Jaatha, the ‘best’ parameter combination is the one with the highest composite likelihood

```

> library(ape)
> # The path to the data
> sample.file <- system.file('example_fasta_files/sample.fasta',
+                             package='jaatha')
> # Reading the data
> sample.data <- read.dna(sample.file, format='fasta',
+                           as.character=TRUE)
> # Calculating the JSFS
> sample.jsfs <- calculateJsfs(sample.data,
+                               pop1.rows=3:7,
+                               pop2.rows=8:12,
+                               outgroup.row=1:2)
> sample.jsfs

```

```

      [,1] [,2] [,3] [,4] [,5] [,6]
[1,]    0    1    0    0    0    0
[2,]    1    1    1    0    0    0
[3,]    2    0    1    0    0    0
[4,]    1    1    0    2    0    1
[5,]    1    0    0    0    3    0
[6,]    1    0    0    0    0    0

```

For the purpose of this HowTo, we will use a simulated JSFS, for which we know the real parameters:

```

> # Real parameters: M = 1, tau = 1 and theta = 10
> real.pars <- c(1,1,10)
> # Simulate a JSFS with this parameters
> jsfs <- dm.simSumStats(dm, real.pars)
> # Print the upper left part of the JSFS
> jsfs[1:10, 1:10]

```

```

      [,1] [,2] [,3] [,4] [,5] [,6] [,7] [,8] [,9] [,10]
[1,]    0 469 247 171 115  63  49  52  31  18
[2,]  511  44  21  26  15  17  17  13  6  8
[3,]  262  32  10  15  8  6  6  2  7  3
[4,]  112  24  8  17  12  12  8  14  2  7
[5,]  97  23  11  9  10  6  5  5  3  1
[6,]  55  20  15  8  8  9  11  4  4  7
[7,]  56  10  11  10  15  7  5  0  6  11
[8,]  35  17  11  11  3  3  6  3  5  3
[9,]  7  11  8  5  2  2  7  2  2  0
[10,]  6  4  9  6  2  2  1  2  6  2

```

## 5 Running Jaatha

Jaatha is divided into two parts. First we find good starting positions by simulating very coarsely across the entire parameter space. We call this part *initial search*. Afterwards a more thorough *refined search* is performed starting from

the best positions of the first step. Before starting the search, we need to set some options like our demographic model, the summary statistics of the real data, and a seed to ensure reproducibility:

```
> jaatha <- Jaatha.initialize(dm, jsfs=jsfs, seed=12345)
```

For more options refer to `?Jaatha.initialize` or the Jaatha manual.

## 5.1 The Initial Search

For the initial search, we divide the parameter space into equally-sized blocks by dividing each of the  $n$  parameters ranges into `blocks.per.par` intervals such that we obtain  $(\text{blocks.per.par})^n$  blocks. Within each block we simulate `sim` data sets with – on a logarithmic scale – uniformly drawn parameter values within each block. To ensure a better sampling of the edges, we simulate in addition data sets for all corner points of each parameter block.

For these data sets we then fit the GLMs and estimate the parameter combination with the maximal score<sup>3</sup>. Each of the blocks provides a single best parameter combination.

In R, the initial search is performed with the command

```
> jaatha <- Jaatha.initialSearch(jaatha,
+                               sim=100,
+                               blocks.per.par=3)
```

```
externalTheta set to TRUE for initial search.
```

```
*** Starting position is being determined ***
```

```
Creating 9 initial blocks ...
```

```
*** Block 1 (lowerB: 0.1 0.01 upperB: 0.368 0.079 )
```

```
Best parameters: 0.175 0.079 1 | Score: 15909.24
```

```
*** Block 2 (lowerB: 0.1 0.079 upperB: 0.368 0.63 )
```

```
Best parameters: 0.229 0.475 1 | Score: 15952.82
```

```
*** Block 3 (lowerB: 0.1 0.63 upperB: 0.368 5 )
```

```
Best parameters: 0.281 0.63 1 | Score: 15997.87
```

```
*** Block 4 (lowerB: 0.368 0.01 upperB: 1.357 0.079 )
```

```
Best parameters: 0.368 0.079 1 | Score: 15722.36
```

```
*** Block 5 (lowerB: 0.368 0.079 upperB: 1.357 0.63 )
```

```
Best parameters: 0.41 0.63 1 | Score: 16011.53
```

```
*** Block 6 (lowerB: 0.368 0.63 upperB: 1.357 5 )
```

```
Best parameters: 0.804 0.918 1 | Score: 16040.45
```

```
*** Block 7 (lowerB: 1.357 0.01 upperB: 5 0.079 )
```

```
Best parameters: 1.357 0.079 1 | Score: 14231.25
```

---

<sup>3</sup>In this phase, Jaatha uses a score instead of the likelihood for computational reasons. The likelihood is proportional to  $\exp(\text{score})$ . The higher the score, the higher the likelihood.

```
*** Block 8 (lowerB: 1.357 0.079 upperB: 5 0.63 )
Best parameters: 1.398 0.63 1 | Score: 15934.2
```

```
*** Block 9 (lowerB: 1.357 0.63 upperB: 5 5 )
Best parameters: 1.357 0.876 1 | Score: 16039.62
```

	score	tau	M	theta
[1,]	16040.45	0.804	0.918	20
[2,]	16039.62	1.357	0.876	20
[3,]	16011.53	0.410	0.630	20
[4,]	15997.87	0.281	0.630	20
[5,]	15952.82	0.229	0.475	20
[6,]	15934.20	1.398	0.630	20
[7,]	15909.24	0.175	0.079	20
[8,]	15722.36	0.368	0.079	20
[9,]	14231.25	1.357	0.079	20

To visualise the estimates for good starting positions sorted by score, type:

```
> Jaatha.getStartingPoints(jaatha)
```

	score	tau	M	theta
[1,]	16040.45	0.804	0.918	20
[2,]	16039.62	1.357	0.876	20
[3,]	16011.53	0.410	0.630	20
[4,]	15997.87	0.281	0.630	20
[5,]	15952.82	0.229	0.475	20
[6,]	15934.20	1.398	0.630	20
[7,]	15909.24	0.175	0.079	20
[8,]	15722.36	0.368	0.079	20
[9,]	14231.25	1.357	0.079	20

Here, there is a big reduction in the scores after the first seven blocks, and a smaller one after the first two. This is suggesting that we either use the first two or the first seven blocks as starting positions for the refined search, depending on how much time we want to spend. For now, we will just use the first two points.

## 5.2 The Refined Search

Now we can conduct the more thorough refined search described above to improve the likelihood approximations.

```
> jaatha <- Jaatha.refinedSearch(jaatha, best.start.pos=2, sim=100)
```

```
*** Search with starting Point in Block 1 of 2 ***
```

```
-----
```

```
Step No 1
```

```
Number of blocks kept: 0 / Number of blocks deleted: 0
```

```
Best parameters: 0.661 0.898 14.823 | Score: 15811.67
```

-----  
Step No 2  
Number of blocks kept: 1 / Number of blocks deleted: 0  
Best parameters: 0.544 0.906 12.761 | Score: 15990.72

-----  
Step No 3  
Number of blocks kept: 1 / Number of blocks deleted: 1  
Best parameters: 0.661 0.936 10.986 | Score: 16043.4

-----  
Step No 4  
Number of blocks kept: 1 / Number of blocks deleted: 1  
Best parameters: 0.804 0.962 10.025 | Score: 16054.62

-----  
Step No 5  
Number of blocks kept: 1 / Number of blocks deleted: 1  
Best parameters: 0.978 0.923 9.594 | Score: 16058.53

-----  
Step No 6  
Number of blocks kept: 1 / Number of blocks deleted: 1  
Best parameters: 1.189 0.93 9.225 | Score: 16060.94

-----  
Step No 7  
Number of blocks kept: 1 / Number of blocks deleted: 1  
Best parameters: 1.104 0.92 9.371 | Score: 16060.74

-----  
Step No 8  
Number of blocks kept: 2 / Number of blocks deleted: 0  
Best parameters: 1.124 0.924 9.342 | Score: 16061.16  
No significant score changes in the last 1 Step(s)

-----  
Step No 9  
Number of blocks kept: 3 / Number of blocks deleted: 0  
Best parameters: 1.117 0.92 9.345 | Score: 16061.44  
No significant score changes in the last 2 Step(s)

-----  
Step No 10  
Number of blocks kept: 4 / Number of blocks deleted: 0  
Best parameters: 1.116 0.924 9.346 | Score: 16061.61  
No significant score changes in the last 3 Step(s)

-----

Step No 11  
Number of blocks kept: 5 / Number of blocks deleted: 0  
Best parameters: 1.098 0.925 9.394 | Score: 16061.77  
No significant score changes in the last 4 Step(s)

-----  
Step No 12  
Number of blocks kept: 6 / Number of blocks deleted: 0  
Best parameters: 1.087 0.921 9.416 | Score: 16061.7  
No significant score changes in the last 5 Step(s)

\*\*\* Finished search \*\*\*  
Score has not change much in the last 5 steps.  
Seems we have converged.

Calculating log-composite-likelihoods for best estimates:

- \* Parameter combination 1 of 10
- \* Parameter combination 2 of 10
- \* Parameter combination 3 of 10
- \* Parameter combination 4 of 10
- \* Parameter combination 5 of 10
- \* Parameter combination 6 of 10
- \* Parameter combination 7 of 10
- \* Parameter combination 8 of 10
- \* Parameter combination 9 of 10
- \* Parameter combination 10 of 10

\*\*\* Search with starting Point in Block 2 of 2 \*\*\*\*

-----  
Step No 1  
Number of blocks kept: 0 / Number of blocks deleted: 0  
Best parameters: 1.116 0.931 14.823 | Score: 15650.69

-----  
Step No 2  
Number of blocks kept: 1 / Number of blocks deleted: 0  
Best parameters: 0.918 0.933 12.761 | Score: 15925.43

-----  
Step No 3  
Number of blocks kept: 1 / Number of blocks deleted: 1  
Best parameters: 0.755 0.94 10.986 | Score: 16040.68

-----  
Step No 4  
Number of blocks kept: 1 / Number of blocks deleted: 1  
Best parameters: 0.918 0.903 9.697 | Score: 16059.07

-----  
Step No 5

Number of blocks kept: 1 / Number of blocks deleted: 1  
Best parameters: 1.099 0.923 9.33 | Score: 16061.05

-----  
Step No 6  
Number of blocks kept: 1 / Number of blocks deleted: 1  
Best parameters: 1.096 0.927 9.401 | Score: 16061.67  
No significant score changes in the last 1 Step(s)

-----  
Step No 7  
Number of blocks kept: 2 / Number of blocks deleted: 0  
Best parameters: 1.12 0.927 9.354 | Score: 16061.54  
No significant score changes in the last 2 Step(s)

-----  
Step No 8  
Number of blocks kept: 2 / Number of blocks deleted: 1  
Best parameters: 1.175 0.92 9.257 | Score: 16061.45  
No significant score changes in the last 3 Step(s)

-----  
Step No 9  
Number of blocks kept: 3 / Number of blocks deleted: 0  
Best parameters: 1.128 0.92 9.346 | Score: 16061.2  
No significant score changes in the last 4 Step(s)

-----  
Step No 10  
Number of blocks kept: 4 / Number of blocks deleted: 0  
Best parameters: 1.135 0.916 9.326 | Score: 16061.32  
No significant score changes in the last 5 Step(s)

\*\*\* Finished search \*\*\*  
Score has not change much in the last 5 steps.  
Seems we have converged.

Calculating log-composite-likelihoods for best estimates:  
\* Parameter combination 1 of 10  
\* Parameter combination 2 of 10  
\* Parameter combination 3 of 10  
\* Parameter combination 4 of 10  
\* Parameter combination 5 of 10  
\* Parameter combination 6 of 10  
\* Parameter combination 7 of 10  
\* Parameter combination 8 of 10  
\* Parameter combination 9 of 10  
\* Parameter combination 10 of 10

Best log-composite-likelihood values are:

	log.cl	block	tau	M	theta
6	-75.36915	2	1.095642	0.9272804	9.401442
7	-75.59973	2	1.119801	0.9269984	9.354007
7	-75.64526	1	1.103654	0.9204776	9.371305
8	-76.02556	1	1.123687	0.9241057	9.341661
1	-76.03412	2	1.135029	0.9157692	9.325577

Hence we perform two independent searches, starting in the `best.start.pos` best starting positions we choose before and according to the general options we choose during initialization, which are stored in the `jaatha` object. In each step, we build a block with `half.block.size` in each direction (on a logarithmic scale) and perform simulations for `sim` random parameter combinations within this block (plus one for very corner). We use this information to estimate the composite maximum likelihood parameters within this block and take this value as new starting position for the next step.

The algorithm stops when the score has not changed more than `epsilon` for five consecutive steps or step `max.steps` is reached. To avoid getting stuck in local maxima, the `weight` option decreases the weight of simulations of previous blocks.

Finally the log composite likelihoods for the best ten parameter combinations are approximated using `sim.final` simulations. These are values printed at the end of the search. This matrix can also be accessed via

```
> likelihoods <- Jaatha.getLikelihoods(jaatha)
> print(likelihoods[1:3, ])
```

	log.cl	block	tau	M	theta
6	-75.36915	2	1.095642	0.9272804	9.401442
7	-75.59973	2	1.119801	0.9269984	9.354007
7	-75.64526	1	1.103654	0.9204776	9.371305

## 6 Parallelization

On Linux and OS X, Jaatha can distribute the simulations on multiple CPU cores. To use this feature, set the `cores` option during initialization. The value of the option specifies the number of cores you want to use. As default, Jaatha will execute 'packages' of 10 simulation on a core in a row to reduce the overhead created by inter-process communication. This is fine for demographic models that can be quickly simulated. However, if you have simulations that takes multiple seconds and/or many cores available, you may want smaller packages. This can be archived by setting `sim.package.size = 1`.

The use of the `cores` option does affect the seeding system. A run without `cores` will give different results as the identical run with the option set at a value greater than one, even when using the same seed. However, if you use `cores`, the actual number of cores you use does not change the results.

## 7 Finite sites models

As described in Mathew et al. [2012], you can use finite sites mutation models in Jaatha. However, this currently only works on Linux and OS X and requires that seq-gen [Rambaut and Grassly, 1997] is installed on your system. On Debian GNU/Linux and its derivatives (Ubuntu, Mint) you can install it running `apt-get install seq-gen` as root. Otherwise download the current version from <http://tree.bio.ed.ac.uk/software/seqgen> and compile it according to the instruction within the package. Jaatha will search for the seq-gen executable using your PATH variable. If you get an error that it is not able to find it, you can specify the path to the executable using the `Jaatha.setSeqgenExecutable` function.

To create a finite sites model, you must specify a mutation model using the `dm.setMutationModel` function and add an outgroup to your model using `dm.addOutgroup`. Optionally, you can then add a mutation rate heterogeneity to your model using `dm.addMutationRateHeterogeneity`.

## References

- Richard R. Hudson. Generating samples under a wright-fisher neutral model of genetic variation. *Bioinformatics*, 18(2):337–338, February 2002. doi: 10.1093/bioinformatics/18.2.337.
- Lisha Mathew, Paul R. Staab, Laura E Rose, and Dirk Metzler. A composite likelihood approach to distinguish gene flow from violations of the infinite-sites model. *Submitted.*, December 2012.
- Lisha Naduvilezhath, Laura E Rose, and Dirk Metzler. Jaatha: a fast composite likelihood approach to estimate demographic parameters. *Molecular Ecology*, 20(13):2709–2723, July 2011. doi: 10.1111/j.1365-294X.2011.05131.x.
- Andrew Rambaut and Nicholas C Grassly. Seq-gen: An application for the monte carlo simulation of DNA sequence evolution along phylogenetic trees. *Computer applications in the biosciences : CABIOS*, 13(3):235–238, June 1997. ISSN 1367-4803, 1460-2059. doi: 10.1093/bioinformatics/13.3.235.