

# ***BAl*-Phy Tutorial (for version 3.1)**

**Benjamin Redelings**

---

## **Table of Contents**

- [1. Introduction](#)
- [2. Setting up the ~/alignment\\_files directory](#)
- [3. Command line options](#)
- [4. Analysis 1: 5S rRNA -- free alignment versus fixed alignment](#)
  - [4.1. Free alignment](#)
  - [4.2. Fixed alignment](#)
- [5. Analysis 2: ITS sequences -- multi-gene analysis](#)
  - [5.1. Questions about parameter differences between genes](#)
- [6. Analysis 3: Exons and Introns](#)
  - [6.1. Extracting parts of a gene](#)
  - [6.2. The Rates.free model](#)
  - [6.3. Fixing the alignment of some partitions](#)
  - [6.4. Linking parameters between partitions](#)
  - [6.5. Option files](#)
  - [6.6. Questions about intron alignments](#)
  - [6.7. Exons as codons + introns as nucleotides](#)
  - [6.8. Exons as amino acids + introns as nucleotides](#)
- [7. Analysis 4: ITS sequences - with a better model](#)
- [8. Complex substitution models](#)
  - [8.1. Defaults](#)
  - [8.2. Rate variation](#)
  - [8.3. Codon models](#)
  - [8.4. Fixing parameter values](#)
  - [8.5. Priors](#)
- [9. Specifying the model for each partition](#)
  - [9.1. Using different substitution models](#)
  - [9.2. Disabling alignment estimation for some partitions](#)
  - [9.3. Sharing model parameters between partitions](#)
  - [9.4. Sharing substitution rates between partitions](#)
- [10. Dataset preparation](#)
  - [10.1. Splitting and Merging Alignments](#)
  - [10.2. Shrinking the data set](#)
  - [10.3. Cleaning the data set](#)

## **1. Introduction**

Before you start this tutorial, please [download](#) and install bali-phy, following the installation instructions in the [manual](#).

## **2. Setting up the ~/alignment\_files directory**

Go to your home directory:

```
% cd ~
```

Make a directory called `alignment_files` inside it:

```
% mkdir alignment_files
```

Go into the `alignment_files` directory:

```
% cd alignment_files
```

Download the example alignment files:

```
% wget http://www.bali-phy.org/examples.tgz
```

Alternatively, you can use **curl**

```
% curl -O http://www.bali-phy.org/examples.tgz
```

Extract the compressed archive:

```
% tar -zxvf examples.tgz
```

Take a look inside the `examples` directory:

```
% ls examples
```

Take a look at an input file (you can press 'q' to exit 'less'):

```
% less examples/5S-rRNA/5d.fasta
```

Get some information about the alignment:

```
% alignment-info examples/5S-rRNA/5d.fasta
```

### 3. Command line options

What version of `bali-phy` are you running? When was it compiled? Which compiler? For what computer type?

```
% bali-phy -v  
% bali-phy --version
```

Look at the list of command line options:

```
% bali-phy -h | less  
% bali-phy --help | less
```

As you can see, many options have a long form, and a short form which is a single letter.

You can also show help for advanced options:

```
% bali-phy --help=advanced | less  
% bali-phy --help=expert | less
```

There is a shorter syntax for help topics:

```
% bali-phy help advanced | less
```

You can get help on models, distributions, functions, and command-line options:

```
% bali-phy help tn93      | less      // Tamura-Nei (1993) model
% bali-phy help normal   | less      // Normal distribution
% bali-phy help quantile | less      // quantile function
% bali-phy help iterations | less     // --iterations=<num> command line option
```

## 4. Analysis 1: 5S rRNA -- free alignment versus fixed alignment

Let's estimate a phylogeny on a very small data set under both a free alignment and a fixed alignment to see how the alignment affects our confidence in the estimated topology.

### 4.1. Free alignment

Let's get started by going to the examples directory:

```
% cd ~/alignment_files/examples/5S-rRNA
```

Now let's start a run. BAli-Phy will create a directory called `5d-free-1` to hold output files from this analysis, as directed by the `-n 5d-free` option. Here `-n` is short for `--name`, while `-S` is a short form of `--smodel` and specifies the substitution model.

```
% bali-phy 5d-clustalw.fasta -S gtr+Rates.gamma[4]+inv -n 5d-free &
```

Yay! Now you have started your first copy of the analysis. Let's run another copy of the same analysis, which will create another directory called `5d-free-2` for its own output files. This additional run will take advantage of a second processor, and will also help detect when the runs have performed enough iterations.

```
% bali-phy 5d-clustalw.fasta -S gtr+Rates.gamma[4]+inv -n 5d-free &
```

You can take a look at jobs running in the background. There should be two `bal-phy` jobs running:

```
% jobs
```

The previous command only shows jobs started by the shell you are typing in. You can also examine all processes running on your computer. `bal-phy` should show up near the top since it is using a lot of the CPU:

```
% top      // press 'q' to exit top
```

In order to see how many iterations have completed, check the number of lines in the log files:

```
% wc -l 5d-*/C1.log
```

To summarize the results generated so far, run the summary script and then view the results in a web browser:

```
% bp-analyze 5d-free-1/ 5d-free-2/
% firefox Results/index.html &
```

#### 4.1.1. Alignment questions:

- How much difference is there between the estimated alignment and the initial alignment? (See *Alignment Distribution / Initial / Diff*)
- Which part of the alignment is most certain? (See *Alignment Distribution / Best (WPD) / AU*)

#### 4.1.2. Parameter questions:

- What is the posterior median for `inv:p_inv`? (See *Scalar Variables*)
- What is the meaning of `inv:p_inv` parameter? (Run `bal-phy help inv`)
- What is the prior distribution on `inv:p_inv`? (See *Model and priors / Substitution model*)

When you have about 2000 samples from each run, summarize the latest results:

```
% wc -l 5d-*/C1.log
% bp-analyze 5d-free-1/ 5d-free-2/
% mv Results 5d-free.Results
```

```
% firefox 5d-free.Results/index.html &
```

Now lets kill the bali-phy processes:

```
% top // use top to find the PID (process id) for each of the bali-phy processes
% kill pid1
% kill pid2
% top // check that the jobs have stopped.
```

Now let's look at the tree in figtree:

```
% figtree 5d-free.Results/c50.PP.tree &
```

1. When it says "Please select a name for these values", enter **PP**.
2. Then, enable "Branch Labels".
3. Under "Branch Labels / Display", select "PP".
4. You can increase the font size to make the figure more readable.
5. You can also increase the font size of the tip labels.

#### 4.1.3. Tree questions (free alignment):

- How many internal branches are in the 50% consensus tree?
- What is the posterior probability for each of these internal branches?

## 4.2. Fixed alignment

```
% cd ~/alignment_files/examples
% bali-phy 5d-clustalw.fasta -S gtr+Rates.gamma[4]+inv -I none -n 5d-fixed &
% bali-phy 5d-clustalw.fasta -S gtr+Rates.gamma[4]+inv -I none -n 5d-fixed &
```

The **-I none** is a short form of **--imodel=none**, where *imodel* means the insertion-deletion model. When there's no model of insertions and deletions, then the alignment must be kept fixed.

Summarize the analysis after about 2000 iterations:

```
% wc -l 5d-*/C1.log
% bp-analyze 5d-fixed-1/ 5d-fixed-2
```

Examine the majority consensus tree for the fixed alignment analysis:

```
% figtree Results/c50.PP.tree &
```

#### 4.2.1. Tree questions (fixed alignment):

- How many internal branches are in the 50% consensus tree?
- What is the posterior probability for each of these internal branches?
- For this data set, how does fixing the alignment affect our confidence in the tree? Why might this be?

You can kill the remaining runs after you answer these questions:

```
% killall bali-phy
```

However, beware: if you are running multiple analyses, this will terminate all of them.

## 5. Analysis 2: ITS sequences -- multi-gene analysis

Let's do an analysis of intergenic transcribed spacer (ITS) genes from 20 specimens of lichenized fungi (Gaya et al, 2011 analyzes 68 specimens). The ITS genes contain a lot of variation, but are hard to align. This analysis will show how to estimate the alignment, phylogeny, and evolutionary parameters using MCMC. Averaging over all alignments makes it safe to use the ITS region without throwing out ambiguously-aligned regions.

This data set is divided into three gene regions, or partitions. It is assumed that all genes evolve on the same tree, but may have different rates and evolutionary parameters. Let's look at the sequences. How long are

they?

```
% cd ~/alignment_files/examples/ITS
% alignment-info ITS1.fasta
% alignment-info 5.8S.fasta
% alignment-info ITS2.fasta
```

By default, each gene gets a default substitution model based on whether it contains DNA/RNA or amino acids. By running bali-phy with the `--test` option, we can reveal what substitution models and priors will be used, without actually starting a run.

```
% bali-phy ITS1.fasta 5.8S.fasta ITS2.fasta --test
```

Let's run two copies of this analysis:

```
% bali-phy ITS1.fasta 5.8S.fasta ITS2.fasta &
% bali-phy ITS1.fasta 5.8S.fasta ITS2.fasta &
```

Since we did not specify a name for the analysis, bali-phy creates the name `ITS1-5.8S-ITS2` from the input file names. The mean or the median of these values can be used as an estimate of the parameter. The variance indicates the degree of uncertainty. Let's look at the initial parameter estimates:

```
% statreport ITS1-5.8S-ITS2-1/C1.log ITS1-5.8S-ITS2-2/C1.log | less
% statreport ITS1-5.8S-ITS2-1/C1.log ITS1-5.8S-ITS2-2/C1.log --mean | less
```

The parameter estimates are also summarized in a table in the summary report:

```
% bp-analyze ITS1-5.8S-ITS2-1/ ITS1-5.8S-ITS2-2/
% firefox Results/index.html &
```

The program Tracer graphically displays the posterior probability distribution for each parameter.

```
% tracer ITS1-5.8S-ITS2-1/C1.log ITS1-5.8S-ITS2-2/C1.log &
```

If you are using Windows or Mac, first run Tracer, and then press the `±` button to add the files.

## 5.1. Questions about parameter differences between genes

These genes have different evolutionary processes. How does the evolutionary process for these genes differ in:

1. substitution rate? (`Scale[1]`, `Scale[2]`, ...)
2. insertion-deletion rate? (`I1/rs07:log_rate`, `I2/rs07:log_rate`, ...)
3. nucleotide frequencies? (`S1/tn93:pi[A]`, `S1/tn93:pi[C]`, ...)
4. number of indels? (`#indels`)

## 6. Analysis 3: Exons and Introns

Let's look at a data set containing both exons and introns. The pattern of dividing a sequence into high-rate and low-rate regions also applies to stems and loops in RNA sequences. It is helpful to allow different insertion-deletion rates in conserved and non-conserved regions of a gene.

### 6.1. Extracting parts of a gene

First, let's split the gene into separate FASTA files for each intron and exon:

```
% cd ~/alignment_files/examples/ferns/
% alignment-cat -c 5-8 -e cleaned.fasta > exon1.fasta
% alignment-cat -c 9-453 -e cleaned.fasta > intron1.fasta
% alignment-cat -c 454-596 -e cleaned.fasta > exon2.fasta
% alignment-cat -c 597-864 -e cleaned.fasta > intron2.fasta
% alignment-cat -c 865-948 -e cleaned.fasta > exon3.fasta
% alignment-cat -c 949-1328 -e cleaned.fasta > intron3.fasta
% alignment-cat -c 1329-1393 -e cleaned.fasta > exon4.fasta
```

Now, let's put the pieces back together in case we need the complete gene in the future:

```
% alignment-cat exon1.fasta intron1.fasta exon2.fasta intron2.fasta exon3.fasta intron3.fasta exon4.fasta >
combined.fasta
```

## 6.2. The Rates.free model

Will we use the `gtr+Rates.free[n=4]` for the substitution model.

```
% bali-phy exon{1,2,3,4}.fasta intron{1,2,3}.fasta -S gtr+Rates.free[n=4] --test
```

This model is more general than `Rates.gamma`, since it estimates 4 free rates and the frequencies of the 4 rates. These 8 parameters have 6 degrees of freedom, which is a lot more than the 1 degree of freedom for the `Rates.gamma` model. Furthermore, each partition by default gets a separate copy of the model, which leads to 42 degrees of freedom just for the `Rates.free` part of the model!

## 6.3. Fixing the alignment of some partitions

We would like to disable alignment estimation, but only for the exons:

```
% bali-phy exon{1,2,3,4}.fasta intron{1,2,3}.fasta -I 1,2,3,4:none -S gtr+Rates.free[n=4] --test
```

## 6.4. Linking parameters between partitions

We can reduce the number of parameters by using one set of parameters for the exons, and one set for the introns. This is called *linking* the models.

```
% bali-phy exon{1,2,3,4}.fasta intron{1,2,3}.fasta -I 1,2,3,4:none -S 1,2,3,4:gtr+Rates.free[n=4] -S
5,6,7:gtr+Rates.free[n=4] --test
```

We would also like to share the indel model between introns:

```
% bali-phy exon{1,2,3,4}.fasta intron{1,2,3}.fasta -I 1,2,3,4:none -S 1,2,3,4:gtr+Rates.free[n=4] -S
5,6,7:gtr+Rates.free[n=4] -I 5,6,7: --test
```

However, the command line is getting so long that it is difficult to manage.

## 6.5. Option files

We can put the options into a text file that is easier to manage and edit. Let's make a text file called `options1.txt`:

```
align = exon1.fasta
align = exon2.fasta
align = exon3.fasta
align = exon4.fasta
align = intron1.fasta
align = intron2.fasta
align = intron3.fasta

smodel = 1,2,3,4:gtr+Rates.free[n=4]
imodel = 1,2,3,4:none

smodel = 5,6,7:gtr+Rates.free[n=4]
imodel = 5,6,7:
```

We can then run

```
% bali-phy -c options1.txt --test
```

Let's run two copies of this analysis:

```
% bali-phy -c options1.txt &
% bali-phy -c options1.txt &
```

## 6.6. Questions about intron alignments

- How do bali-phy alignments of the introns differ from muscle and mafft alignments of the introns?
- How much certainty or uncertainty do the intron alignments have?
- How do the evolutionary rates of the exons compare to the evolutionary rates of the introns?
- Would it be reasonable to link the intron rates? How about the exon rates?

## 6.7. Exons as codons + introns as nucleotides

Another thing that we can do is to use a codon model on the combined exons, while using a nucleotide model on the introns. In order to load nucleotide sequences as codons, three conditions must be met:

### 1. Sequence lengths must be multiples of three.

AAA --- TTT is OK

AAA --- T-- is not OK.

### 2. Gaps must be codon-aligned.

AAA --- TTT is OK

AA- --A TT- is not OK.

### 3. The sequences must be in-frame and not have stop codons.

ATG AAT AAA is OK, because it translates to MNK.

AAT GAA TAA is not OK, because it translates to NE\*, where\* is a stop codon.

Note that spaces in FASTA files are allowed, but not required.

Let's first extract the coding data:

```
% alignment-cat exon{1,2,3,4}.fasta -c2-295 > coding.fasta // Concatenate exons and chop off partial codons at
beginning and end
% sed -i 's/-/N/g' coding.fasta // Change codons like A-- to ANN at end of sequence
```

We can then construct an option file options2.txt:

```
align = intron1.fasta
align = intron2.fasta
align = intron3.fasta
align = coding.fasta

smodel = 1,2,3:gtr+Rates.free[n=4]
imodel = 1,2,3:rs07

smodel = 4:yn94
#smodel = 4:fMutSel0
#smodel = 4:m3[gtr_sym, f3x4]
imodel = 4:none
```

We can choose any one of the three codon models for the coding sequences in the 4th partition.

```
% bali-phy -c options2.txt --test
```

Try changing the model to **fMutSel0**. Does the number of parameters increase or decrease?

## 6.8. Exons as amino acids + introns as nucleotides

We can also translate the coding sequence into a protein sequence, while using a nucleotide model on the introns. Let's first extract the coding data:

```
% alignment-translate < coding.fasta > protein.fasta
```

We can then construct an option file `options3.txt`:

```
align = intron1.fasta
align = intron2.fasta
align = intron3.fasta
align = protein.fasta

smodel = 1,2,3:gtr+Rates.free[n=4]
imodel = 1,2,3:rs07

smodel = 4:lg08+Rates.log_normal[n=4]
imodel = 4:none
```

```
% bali-phy -c options3.txt --test
```

## 7. Analysis 4: ITS sequences - with a better model

Let's revisit the ITS sequences used in analysis 2. Now that you've been exposed to different models in different partitions, let's use a more complex substitution model for the ITS partitions (`--smodel 1,3:tn93+Rates.free[n=3]`). Note that this also links the substitution models for the ITS partitions.

We can also fix the alignment for 5.8S partition, since it has almost no indels.

```
align = ITS1.fasta
align = 5.8S.fasta
align = ITS2.fasta
smodel = 1,3:tn93+Rates.free[n=3]
smodel = 2:tn93
imodel = 2:none
scale = 1,3:
```

This file additionally links the *scale* for the ITS partitions (`--scale=1,3:`). This forces them to share the same evolutionary rate, and allows more precise estimates of the shared scale. You can run the analysis with unlinked scales by removing this option in order to determine if (i) the scales are similar enough to share and (ii) the scales have a high enough variance that they need to be linked.

## 8. Complex substitution models

While those analyses are running, let's look at how to specify more complex substitution models in `bali-phy`.

### 8.1. Defaults

When you don't specify values for parameters like *imodel*, `bali-phy` uses sensible defaults. For example, these two commands are equivalent:

```
% cd ~/alignment_files/examples/
% bali-phy 5S-rRNA/25-muscle.fasta --test
% bali-phy 5S-rRNA/25-muscle.fasta --test --alphabet=RNA --smodel=tn93 --imodel=rs07
```

You can change the substitution model from the Tamura-Nei model to the General Time-Reversible model:

```
% bali-phy 5S-rRNA/25-muscle.fasta --test -S gtr
```

Here the `-S gtr` is a short form of `--smodel=gtr`, where *smodel* means the substitution model.

### 8.2. Rate variation

You can also allow different sites to evolve at 5 different rates using the `gamma[4]+INV` model of rate heterogeneity:

```
% bali-phy 5S-rRNA/25-muscle.fasta --test -S gtr+Rates.gamma[4]+inv
```

You can also use a `log_normal` instead of a `gamma`. `log_normal+INV` is sometimes better behaved than

gamma+INV, because the smallest rate bin under log\_normal is not quite as close to 0.

```
% bali-phy 5S-rRNA/25-muscle.fasta --test -S gtr+Rates.log_normal[4]+inv
```

You can allow 5 different rates that are all independently estimated:

```
% bali-phy 5S-rRNA/25-muscle.fasta --test -S gtr+Rates.free[n=5]
```

### 8.3. Codon models

We can also conduct codon-based analyses using the Nielsen and Yang (1998) model of diversifying positive selection (dN/dS):

```
% bali-phy Globins/bglobin.fasta --test -S yn94+f1x4
```

The yn94 model takes a nucleotide exchange model as a parameter. This parameter is optional, and the default is hky85, which you could specify as yn94[,hky85\_sym]. You can change this to be more flexible:

```
% bali-phy Globins/bglobin.fasta --test -S yn94[,gtr_sym]+f1x4
```

You can make the codon frequencies to be generated from a single set of nucleotide frequencies:

```
% bali-phy Globins/bglobin.fasta --test -S yn94[,gtr_sym]+mg94
```

The M7 model allows different sites to have different dN/dS values, where the probability of dN/dS values follows a beta distribution:

```
% bali-phy Globins/bglobin.fasta --test -S m7
```

The M7 model has parameters as well. Here are the defaults:

```
% bali-phy Globins/bglobin.fasta --test -S m7[4,hky85_sym,f61]
```

The M3 model allows different sites to have different dN/dS values, but directly estimates what these values are:

```
% bali-phy Globins/bglobin.fasta --test -S m3[n=3]
```

The M8a\_Test model allows testing for positive selection in some fraction of the sites:

```
% bali-phy Globins/bglobin.fasta --test -S m8a_test[4,hky85_sym,f3x4]
```

### 8.4. Fixing parameter values

We can use the TN93+Gamma[4]+INV model without specifying parameters:

```
% bali-phy Globins/bglobin.fasta --test -S tn93+Rates.gamma+inv
```

However, we can also fix parameter values:

```
% bali-phy Globins/bglobin.fasta --test -S tn93+Rates.gamma[n=4,alpha=1]+inv[p_inv=0.2]
```

Here we have set the shape parameter for the Gamma distribution to 1, and the fraction of invariant sites to 20%. Since these parameters are fixed, they will not be estimated and their values will not be shown in the log file.

You can see the parameters for a model by using the `help` command, as in:

```
% bali-phy help Rates.gamma
```

This will show the default value or default prior for each parameter, if there is one.

## 8.5. Priors

By default the fraction of invariant sites follows a uniform[0,1] distribution:

```
% bali-phy help inv
```

However, we can specify an alternative prior:

```
% bali-phy Globins/bglobin.fasta --test -S tn93+Rates.gamma[n=4]+inv[p_inv~uniform[0,0.2]]
```

We can also specify parameters as positional arguments instead of using variable names:

```
% bali-phy Globins/bglobin.fasta --test -S tn93+Rates.gamma[4]+inv[~uniform[0,0.2]]
```

Here "~" indicates a sample from the uniform distribution instead of the distribution itself.

The insertion-deletion model also has parameters.

```
% bali-phy help rs07
```

Here the default value for rs07:mean\_length is exponential[10,1]. This indicates a random value that is obtained by sampling an Exponential random variable with mean 10 and then adding 1 to it.

## 9. Specifying the model for each partition

For analyses with multiple partitions, we might want to use different models for different partitions. When two partitions have the same model, we might also want them to have the same parameters. This is described in more detail in section 4.3 of the [manual](#).

### 9.1. Using different substitution models

Now lets specify different substitution models for different partitions.

```
% cd ~/alignment_files/examples/ITS
% bali-phy {ITS1,5.8S,ITS2}.fasta -S 1:gtr -S 2:hky85 -S 3:tn93 --test
```

### 9.2. Disabling alignment estimation for some partitions

We can also disable alignment estimation for some, but not all, partitions:

```
% bali-phy {ITS1,5.8S,ITS2}.fasta -I 1:rs07 -I 2:none -I 3:rs07 --test
```

Specifying **-I none** removes the insertion-deletion model and parameters for partition 2 and also disables alignment estimation for that partition.

Note that there is no longer an I3 indel model. Partition #3 now has the I2 indel model.

### 9.3. Sharing model parameters between partitions

We can also specify that some partitions with the same model also share the same parameters for that model:

```
% bali-phy {ITS1,5.8S,ITS2}.fasta -S 1,3:gtr -I 1,3:rs07 -S 2:tn93 -I 2:none --test
```

This means that the information is *pooled* between the partitions to better estimate the shared parameters.

Take a look at the model parameters, and the parenthesis after the model descriptions. You should see that there is no longer an S3 substitution model or an I3 indel model. Instead, partitions #1 and #3 share the S1 substitution model and the I1 indel model.

### 9.4. Sharing substitution rates between partitions

We can also specify that some partitions share the same scaling factor for branch lengths:

```
% bali-phy {ITS1,5.8S,ITS2}.fasta -S 1,3:gtr -I 1,3:rs07 -S 2:tn93 -I 2:none --scale=1,3: --test
```

This means that the branch lengths for partitions 1 and 3 are the same, instead of being independently estimated.

Take a look at the model parameters. There is no longer a Scale[3] parameter. Instead, partitions 1 and 3 share Scale[1].

## 10. Dataset preparation

### 10.1. Splitting and Merging Alignments

BAlI-Phy generally wants you to split concatenated gene regions in order to analyze them.

```
% cd ~/alignment_files/examples/ITS-many/  
% alignment-cat -c1-223 ITS-region.fasta > 1.fasta  
% alignment-cat -c224-379 ITS-region.fasta > 2.fasta  
% alignment-cat -c378-551 ITS-region.fasta > 3.fasta
```

Later you might want to put them back together again:

```
% alignment-cat 1.fasta 2.fasta 3.fasta > 123.fasta
```

### 10.2. Shrinking the data set

You might want to reduce the number of taxa while attempting to preserve phylogenetic diversity:

```
% alignment-thin --down-to=30 ITS-region.fasta > ITS-region-thinned.fasta
```

You can specify that certain sequences should not be removed:

```
% alignment-thin --down-to=30 --keep=Csaxicola420 ITS-region.fasta > ITS-region-thinned.fasta
```

### 10.3. Cleaning the data set

Keep only sequences that are not too short:

```
% alignment-thin --longer-than=250 ITS-region.fasta > ITS-region-long.fasta
```

Remove 10 sequences with the smallest number of conserved residues:

```
% alignment-thin --remove-crazy=10 ITS-region.fasta > ITS-region-sane.fasta
```

Keep only columns with a minimum number of residues:

```
% alignment-thin --min-letters=5 ITS-region.fasta > ITS-region-censored.fasta
```