

The NestedMICA motif inference tool

Thomas A. Down
thomas.a.down@googlemail.com

Version 0.8.0 [20070425]

Abstract

NestedMICA is a sensitive, scalable, pattern-discovery system, aimed at finding transcription factor binding sites and similar motifs in biological sequence. More discussion of the principles behind NestedMICA, and an evaluation of its sensitivity can be found in a recent paper [?].

If you have any problems or questions about the system, please contact Thomas Down.

1 Theory

Motif-finding is a long standing problem in sequence bioinformatics. A typical statement of the problem would be “given a set of sequences, which motifs are significantly over-represented with respect to a given model of non-functional sequence.” The choice of non-functional sequence (background) model is an important part of this question and will be discussed in more detail later. The classical use for motif-finding software is the detection of transcription factor binding sites in promoter regions, but there are other interesting functional elements in biological sequences and elsewhere which can be found by motif-discovery methods.

Motif-finding strategies can be broadly divided into two classes: methods which rely on exhaustively enumerating a set of motifs (generally using an optimized data structure such as a suffix tree) then reporting the most frequently occurring, and methods which find significant motifs by optimizing or drawing samples from a probabilistic model. Exhaustive enumeration is a good strategy for finding perfectly conserved motifs (*i.e.* every instance is identical), but for typical transcription factor binding sites, which often have several weakly constrained positions, exhaustive enumeration becomes problematic and the results usually have to be post-processed with some kind of clustering system. NestedMICA is a probabilistic motif finder, and we will not consider exhaustive enumeration further.

Probabilistic motif-finders treat the supplied sequences as a mixture of ‘interesting’ motifs and ‘non-interesting’ background sequences. This problem has classically been reduced to a simple case: considering just one motif at a time, model each sequence with a random background model which may, or may not, contain a

single instance of the motif under consideration. This is the zero-or-one occurrence per sequence (ZOOPS) model. It can be easily represented as a hidden Markov model (figure 1)), and standard techniques such as expectation maximization [?] or Gibbs sampling [?] can be used to find good sets of model parameters (corresponding to good motif models).

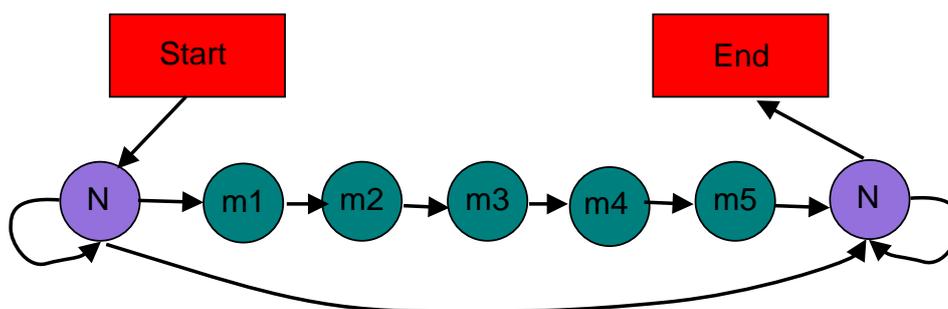


Figure 1: The ZOOPS sequence mixture model

There are two significant concerns about this strategy which NestedMICA tries to address:

- Real regulatory regions, and most other contexts where interesting motifs can be found, don't really contain just a single instance of a single motif. Programs which use the ZOOPS model work around this by finding the strongest motif in a set, then scanning for all its instances, masking them out, and re-running the process on the remaining sequence. This strategy is greedy and it is by no means clear that its behaviour will be optimal, especially when working on a system where there is a set of closely related, yet still distinct, motif types. In an environment where novel transcription factors are created by gene duplication then diverge to perform a new function, such situations seem quite probable, and have not yet been well investigated.
- Existing techniques for optimizing or exploring sequence models of this type tend to be strongly local in nature – they concentrate on regions of the probability landscape close to their starting point. This is particularly clear for expectation-maximization methods, which always move in a direction which increases the likelihood of the model. Sampling strategies don't strictly have this limitation, but in practice, crossing the low-likelihood valley between two high-likelihood peaks is unlikely, often to the point where it becomes vanishingly rare.

1.1 Motif ICA

We treat finding multiple motifs in a set of sequences as a form of independent component analysis (ICA) problem. In linear ICA, a matrix of observations, X is

approximated as a linear mixture, A of some sources, s :

$$x = As + \nu \quad (1)$$

(where ν is a noise matrix). A classical example is the “cocktail party problem” where a set of M microphones record different mixtures of the voices of N speakers. Given samples from these microphones at t timepoints, ICA attempts to factorize the $M \times t$ observation matrix into a $N \times t$ source matrix and a $M \times N$ mixing matrix.

It is possible to generalize ICA to *any* mixing situation, given a suitable mixing operator. In the classical case, the mixing operator is simply the addition operator.

In motif ICA, the sources are short sequence motifs (currently, but not necessarily, modeled as position-weight matrices [?]), while the observations are larger sequences. There are a number of interpretations of the mixing matrix. Currently, we use a boolean mixing matrix (all coefficients are either 0 or 1), and a given sequence is expected to contain a given motif if the relevant mixing coefficient is 1. The ‘noise’ part of the ICA model represents all the sequence which isn’t modelled by one of the motifs.

1.2 Nested sampling

Nested sampling is an alternative way of performing probabilistic inference in a Bayesian framework, proposed recently by John Skilling [?]. It can be considered to be a Monte Carlo method, since the process is driven forward by a series of randomly chosen events, but it is quite distinct from the classic Metropolis-Hastings method, and methods such as Gibbs sampling and slice samplings, which can all be considered as optimized implementation strategies for Metropolis-Hastings.

Nested sampling is applied to an *ensemble* of states, which represent possible solutions to the problem at hand. The ensemble is initialized by sampling uniformly from the prior distribution, then sorting the states according to their likelihoods. In nested sampling, Each state in the ensemble is considered to be a representative of the set of states with similar likelihoods. If the likelihood of each state is drawn as a contour on the likelihood distribution, we see a nested set of contour lines, converging towards the peaks of the likelihood distribution. We therefore call the ordered set of states a nested ensemble. For each cycle of nested sampling, the least likely state in the ensemble is discarded, and a new state is chosen by sampling uniformly from the prior subject to the constraint that the likelihood of the new state must be greater than or equal to the likelihood of the discarded state.

In this context, the most exciting property of nested sampling is that, given a reasonably large ensemble, the final sample drawn from a converged nested sampler can be expected to reflect the global optimum of the likelihood landscape. Moreover, in cases where more than one globally significant optimum exists, these should be represented in the sample set in direct proportion to the amount of posterior mass they represent.

1.3 Background models

The background model is an important component of the SMM framework – after all, it will usually be responsible for modeling the majority of the input sequence! The simplest strategy – and still a common one – is to treat all non-motif sites as independent and identically distributed. In HMM terms, this makes the background model a zeroth-order Markov chain. However experience shows that genomic DNA sequence, even when apparently totally non-functional, is not a good fit to the i.i.d. model. The best known deviation is perhaps the dramatic under-representation of CpG dinucleotides in most parts of vertebrate genomes, but other significant effects are known. In any case, practical experience shows that motif finders equipped with naive background models tend to report low-complexity elements rather than interesting binding sites.

The first obvious improvement is to replace the zeroth order Markov chain with a first order chain (*i.e.* the probability of observing a particular symbol at position n depends on the symbol at position $n-1$). This model is good at capturing anomalies like the CpG underrepresentation. The success of first order background models has led some researchers to investigate higher order models. One investigation of Markov chain backgrounds can be found in [?]: this concludes that pentanucleotide frequency tables (*i.e.* fourth-order background models) are optimal. However, there are two concerns about this result: firstly, it leaves an open question about what these high-order correlations in background sequence mean (and why fourth-order models appear to outperform fifth-order). Also, training a background model generally requires sequence proportional to the number of free parameters in the model. Fifth order models, with 768 parameters, therefore require large amounts of sequence. Moreover, it is desirable to train the background model on sequence which does not contain target motifs, since a fifth order model could easily capture some information about this motifs, thereby reducing the sensitivity of the motif-finding process. But it is hard to find large amounts of representative background training sequence which doesn't contain interesting motifs.

A different way to generalize the naive background model is to allow several different classes of sequence, each with its own particular base distribution (which could be zeroth-order or higher-order). We call these *mosaic models*, since their underlying assumption is that genome evolution includes some set of constraints which act non-uniformly, even on background sequence.

To investigate the benefit of mosaic models, we took a set of human upstream flanking regions, and split the set in half, using one subset to optimize the background model parameters and the second subset for testing. Test likelihoods for a variety of class numbers and Markov chain orders are shown in figure 2. Considering one class 'mosaics' (which are equivalent to uniform background models), we repeat the previously reported observation that higher order Markov chains are better models of genomic DNA. However, we also see large increases in likelihood when moving to larger numbers of mosaic classes. Interestingly, the lines for zeroth-order and first-order models run almost parallel: this suggests that the benefits of

mosaic models are almost orthogonal to the benefits of first-order models. However, this is not true when moving beyond first-order models.

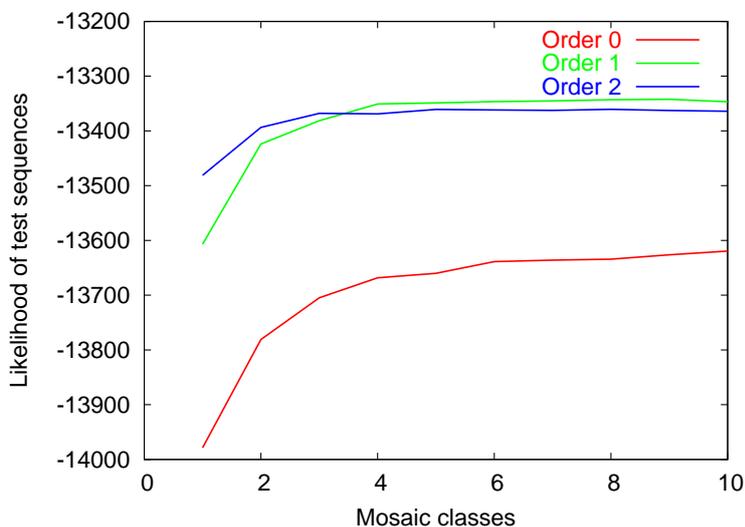


Figure 2: Comparison of mosaic backgrounds

Based on these results, we recommend the use of a four class, first order, mosaic background model for most motif-finding applications on mammalian genomic sequence. In practice, the four classes appear to be:

- A relatively neutral class
- A C+G rich class (CpG islands?)
- Purine rich
- Pyrimidine rich

There are still open questions about the biological significance of the mosaic classes we have observed.

2 Running NestedMICA

NestedMICA is open source software, released under the GNU Lesser GPL. The main program is written in Java, but there are a few small pieces of native C code which are required. The C code is simple and (hopefully) portable, so it should be possible to run NestedMICA on most modern computer platforms.

The NestedMICA source distribution can be downloaded from

<http://www.sanger.ac.uk/Software/analysis/nmica/>

This distribution includes all the required library code. To compile the Java code, you'll need version 1.7.0 or later of the ANT java build tool, which can be downloaded from <http://ant.apache.org/>.

As of NestedMICA 0.7, a Java 5 runtime environment is required. In general, we recommend that you use the latest available Java from Sun Microsystems. Initial benchmarks suggest that Java 6 runtimes run NestedMICA about 10% faster, so use the latest version if it's available. Mac OS X is a supported platform, but you'll need to run version 10.4.0 ("Tiger") or later, and separately download the Java 5 packages from Apple. You'll also need to reconfigure your system to make Java 5 the default version for command-line applications (note that Apple's Java Preferences tool doesn't do this, it only affects programs started via your web browser). The best way to do this is to add the following lines to the `.bashrc` file in your home directory:

```
export JAVA_HOME=/System/Library/Frameworks/JavaVM.framework/  
    Versions/1.5.0/Home  
export PATH=${JAVA_HOME}/bin:$PATH
```

2.1 Building and installing NestedMICA

NestedMICA is written in the Java programming language, with a small amount of C used for performance reasons. On supported platforms, it should be possible to compile the whole system by typing `ant` at the command line. If this does not complete cleanly, contact Thomas Down with the complete error message and details of the platform you're running on.

Currently supported platforms are Linux and Mac OS X, but other Unix-like platforms may also work, and if not it should be reasonably easy to add support. Note that we've had reports of problems when building on Mac OS with old developer tools (before gcc version 3.1). If in doubt, please install the latest developer tools.

Once you have compiled NestedMICA successfully, you can either add it directly to your PATH:

```
export PATH=/home/thomas/nmica-X.Y.Z/bin:$PATH
```

or copy it to some other location. If you are packaging NestedMICA or building it for other users, only the `bin`, `lib`, and `native` subdirectories are required once NestedMICA has been compiled.

2.2 Basic operation

This section describes a (very) small set of test sequences named `micatest100.fa` which are included with the NestedMICA source distribution. The file contains 100 synthetic DNA sequences, each of 50 bases long. The set contains two spiked-in motifs.

The main part of NestedMICA is the motif inference tool, `nminfer`. You should be able to recover the two spiked motifs using a command like:

```
nminfer -seqs micatest100.fa -numMotifs 2
        -minLength 4 -maxLength 12 -ensembleSize 50
```

This runs the motif finding system until the nested sampling process is close to convergence, then writes the optimal set of motifs to a file named `motifs.xml`. In this case, we're telling `nminfer` that we expect to find 2 motifs, but setting a fairly broad range of possible motif lengths.

You can view your discovered motifs using the MotifExplorer tool, available from:

<http://www.sanger.ac.uk/Software/analysis/nmica/mxt.shtml>

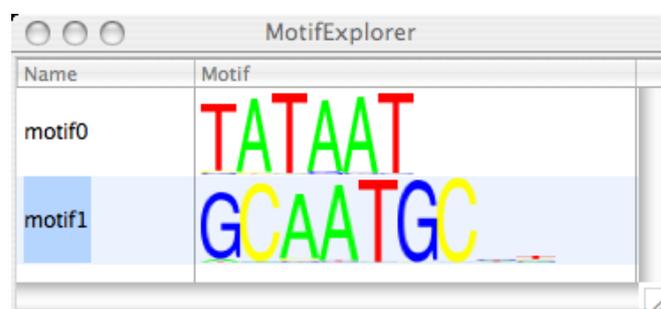


Figure 3: Motifs discovered from the demonstration set.

One word of warning here: the `-ensembleSize 50` option is recommended here in order to get quick results on this simple test data set, but will lead to lower sensitivity than leaving `-ensembleSize` at its default value. For normal use, we do not recommend altering `-ensembleSize` from its default value.

2.3 Increasing the memory limits

Some Java implementations may, by default, limit NestedMICA's memory usage to an unreasonably small value. If NestedMICA is crashing with `OutOfMemoryError`, you can explicitly set a new memory limit by setting an environment variable. To set the memory limit to 500Mb, use:

```
export NMICA_JVMOPTS=-Xmx500M
```

As a rough rule-of-thumb, a good memory limit for large datasets is 1Mb of memory for every 1Kb of input sequence. Note that just because you have allocated NestedMICA 500Mb of memory using a command like that shown above, it won't immediately use all that memory – it just specifies how much memory the program is allowed to use before exiting.

2.4 Background-model options

By default, `nminfer` will build a background model from the supplied sequences. The default model is quite simple: a single zeroth order (*i.e.* mononucleotide frequencies) model for all the input sequences. You can specify more complex models using the `-backgroundOrder` and `-backgroundClasses` options. So to use a four-class dinucleotide model (which seems to be a good choice for vertebrate sequences) you could use:

```
nminfer ... -backgroundOrder 1 -backgroundClasses 4 ...
```

However, complex background models can take some time to train, and it is useful to be able to run `nminfer` multiple times using exactly the same background. Therefore, we recommend building background models as a separate step, *e.g.*:

```
nmmakebg -seqs seqfile.fa -order 1 -classes 4  
-out background42.xml
```

The resulting background model is stored as an XML file which can be inspected if you want to see the exact parameters that have been chosen. Having built a background model in this way, you can use it for any motif finding run:

```
nminfer ... -backgroundModel background42.xml ...
```

Remember that if you run `nminfer` in this way, the `-backgroundOrder` and `-backgroundClasses` options will be silently ignored – the background model architecture is hardwired when the model is trained.

Note: earlier versions of NestedMICA used a different file format for storing background models. You can convert these old files into the new XML format using `nmconvertbg`:

```
nmconvertbg -in background42.sbg -out background42.xml
```

2.5 Background model architecture optimization

Our experience suggests that a background model trained with the parameters `-order 1` and `-classes 4` is generally a good choice when working with vertebrate genomes. However, genome composition varies between species and it is likely that other architectures may be better choices in some cases: for example, a six-class model was used in a recent study of the insect *Drosophila melanogaster* [?]. One principled approach to background selection might be to repeat the experiment depicted in figure 2 using sequences from your species of interest. To do this, take a set of sequences (ideally around 200), then:

```
nmevaluatebg -seqs seqs.fa -order 1 >eval.dat
```

The `nmevaluatebg` splits the sequences into two sets – by default it places 50 sequences in the “test” set and the rest in the “training” set. It fits background models with varying numbers of classes to the training sequences, then calculates the likelihood of observing the test sequences under the newly-trained model. If you plot the output using a tool like `gnuplot`, you should see a plot similar to one of the lines shown in figure 2. Look for a plateau in this plot: the class number on the shoulder of the plateau should represent a good trade-off between background model complexity and fitting the data well.

It is also possible to use `nmevaluatebg` to compare different values of the `-order` parameter, by running the program several times then comparing the output. This has a few complications: firstly, you must ensure that the “test” and “training” sets are the same for all runs. To do this, split the sequences up manually then use the `-trainSeqs` and `-testSeqs` options. Also, it is necessary to compensate for edge effects: when working with an order 2 background model, NestedMICA can’t “see” the first two bases of each sequence, therefore the likelihood of the sequence set will appear to be higher. You can compensate for this using the `-trim` option when testing low-order background models, e.g.:

```
nmevaluatebg ... -order 0 -trim 2 >order0.dat
nmevaluatebg ... -order 1 -trim 1 >order1.dat
nmevaluatebg ... -order 2 >order2.dat
```

2.6 Motif-finding options

The `nfinfer` program has many options. You can see a list of common options by typing `nminfer -help`.

2.7 Everyday options

These are options that you’ll use all the time, and are generally well-behaved.

- `-numMotifs` *n* specifies the number of motifs you want to find (default=1).
- `-maxLength` *n* the maximum length of motifs to find (default=10).
- `-minLength` *n* the minimum length of motifs to find (if not specified, a warning will be printed and this will be set equal to `-maxLength`).
- `-seqs` *filename* a FASTA-formatted database of sequences to analyse.
- `-backgroundModel` *filename* the background model to use, normally generated by running the `nmmakebg` program.
- `-out` *filename* name of a file to write the final motif set (default `motifs.xml`)
- `-sampleFile` *filename* name of a file to write periodic samples during the motif-finding process. If you run the finder with the option `-sampleFile foo`, it will write sample files `foo.1000.xml`, `foo.2000.xml`, etc.
- `-sampleInterval` *n* number of cycles between sample files (default=1000).

`-maxCycles` *n* the maximum number of cycles to run the sampler. The program will then exit regardless of convergence status. Currently the criteria for deciding when the process has converged are quite conservative, so this option may be useful.

`-ensembleSize` *n* the size of the nested ensemble used to explore the probability landscape. Larger values improve reproducibility at the expense of speed. Default behaviour is to set the ensembleSize to 200 for large problems (≥ 5 motifs), but to use larger ensemble sizes when learning less motifs, up to a maximum of 1000 when learning a single motif. This should give maximum sensitivity under most circumstances. When learning one or two motifs from a very simple dataset, you may want to explicitly force a small ensembleSize for performance reasons.

`-alphabet` (`dna|protein`) the alphabet of input sequences. Default is DNA.

`-revComp` allow motifs to occur in either orientation (off by default, slows the program down by a factor of about 2).

`-cluster` prefer motifs which occur in clusters. Causes a small slowdown in training. Currently not sure how helpful this is.

2.8 Advanced options

More complex options, that we'd prefer weren't there. If you're interested in these, it's probably best to discuss them with Thomas Down.

`-mixtureType` (`binary|flat|logit|weighted`) Wierd switch which effects the interpretation of the mixing matrix. Default is `binary`, leave it that way unless you know what you're doing.

`-mixtureUpdate` (`resample|weakResample|max|queue|random`) Strategy for updating the mixing matrix during training. `resample` is the best in theory and works well for most problems. `weakResample` gives better results on large sets of human promoters. Don't use any of the others.

`-expectedUsageFraction` *d* number between 0.0 and 1.0 giving the prior belief of motif frequency. Default is 0.5. Try 0.1 if you think motifs are rare, 0.9 is you think almost all the sequences are likely to contain an instance.

(`-counted| -uncounted`) switch between a model where a motif can occur any number of times in a sequence (`uncounted`), and a more conventional motif-inference strategy where motif-containing sequences are expected to contain exactly one occurrence. Note that the `-counted` model is slower, and is likely to become unusably slow with values of `-numMotifs` greater than about 5. However, it may be preferred in some situations where sensitivity is paramount. Default is `-uncounted`.

2.9 Checkpointing

To minimize the risk of data loss during long runs, NestedMICA can generate checkpoint files at regular intervals during the training process. If you want to keep checkpoint files, add the following arguments to your `nminfer` command line

```
-checkpoint mycp -checkpointInterval 1000
```

Like sample files, checkpoints are automatically numbered with the cycle number at which they are taken. Unlike samples they files are rather large (typically several megabytes, increasing rapidly with dataset and nested ensemble size). By default, `nminfer` deletes old checkpoints: in the standard configuration only two checkpoints are kept on disk, so after the third is written, the first will be deleted. This behaviour can be changed using the `-keepCheckpoints` option.

To restart from a checkpoint:

```
nminfer -restartFromCheckpoint mycp.12000.jos
        -maxCycles 100000 -sampleFile sample -sampleInterval 10000
        -checkpoint mycp2 -checkpointInterval 1000
```

Note that, while the entire trainer state is restored from the checkpoint file, currently none of the ‘housekeeping’ options (termination, sampling, and checkpointing) are, so it is necessary to re-specify them when restarting the checkpoint. This is potentially useful if you want to continue the training of a run which hasn’t quite converged for a few extra cycles.

2.10 Multithreaded operation (SMP machines)

If you are running on a multi-processor machine, add the option `-threads N` to the `nminfer` command line, where `N` is the number of processors in the machine, or the number of CPUs in a large multiprocessor server that you want to devote to `nminfer`. This mode of operation is quite efficient, usually giving better than 95% scaling from one to four processors. This option should also work efficiently on large SMP machines, at least when working with large datasets, but has not been tested extensively with more than 4 CPUs.

2.11 Distributed operation (clusters, farms, Beowulfs, etc.)

An alternative approach to paralelizing a large `nminfer` run is distribute the workload over multiple nodes. Start `nminfer` as usual, but add the options:

```
-distributed -port XXXXX
```

You may use any port number between 1024 and 65535, but if you are running multiple `nminfer` jobs on one cluster, you should give each job a unique port number. This process will be your *master node*. Once it has started up, you can run any number of *worker nodes*, which will evaluate sub-problems for the master node. To run a worker node:

```
nmworker -server YYYYY -port XXXXX
```

Where `YYYYY` is the name (or IP address) of the master node, and the port number matches that used to start the master. To remove a worker node from the cluster you can simply kill the process. It's fine to add and remove workers at any point during the run, although there is currently a short (~10 seconds) glitch after a worker dies during which no work will be done, so very rapid worker turnover can be inefficient. If the `nminfer` run ends normally, all associated worker nodes will be sent a shutdown message, but under other circumstances it may be necessary to kill worker nodes manually.

In NestedMICA 0.7.x, the master node did not perform much computation, and was restricted to housekeeping tasks. As of NestedMICA 0.8.0, you can *optionally* use the master node as an extra compute node. To do this, you should explicitly specify a `-threads` option when starting the `nminfer` process. Depending on the hardware you are using and the number of worker nodes you have attached, the optimum value will be either n or $n - 1$ where n is the number of CPUs in your computer. Using a value of n will maximize CPU utilization, while a value of $n - 1$ leaves one CPU core free at all times for housekeeping tasks. If you run a lot of distributed `nminfer` processes, it is worth doing some testing to find the optimum value for your cluster.

2.12 Comparative genomics

If groups of orthologous regulatory regions are available, NestedMICA can take advantage of the known relationships between species. Unlike phylogenetic footprinting and other 'infinite monkeys' techniques, comparative NestedMICA does not consider the alignments of the sequences, but simply assumes that orthologous regulators are likely to contain similar complements of motifs. In practice, this is implemented by using a single row of the ICA mixing matrix to model all the orthologs.

To run NestedMICA in comparative mode, just add more than one `-seqs` option to the `nminfer` command line. If two sequences in different files share the same name, they will be treated as orthologues.

There are two options for providing background models in comparative mode. If only one `-backgroundModel` switch is present on the command line, that single background model will be used for all species. Alternatively, you can train a separate background model for each species, and specify them with multiple `-backgroundModel` switches, e.g.

```
-seqs human.fa -backgroundModel human.sbg
-seqs mouse.fa -backgroundModel mouse.sbg
```

If you use multiple background models they *must* be listed in the same order as the sequence databases, and there must be exactly equal numbers of sequence databases and background models. If no background models are supplied, a separate background model will be constructed for each supplied sequence file.

If ortholog information is only available for some sequences, that's fine.

There is no limit to the number of species that can be used (except for practical memory limits). We would expect to see diminishing returns beyond about 4 species, but much more testing is needed to determine the optimal configurations.

2.13 Scanning with motifs

A simple program is supplied to scan a sequence using weight matrices inferred by NestedMICA. Hits are output in GFF format. This program uses bits-sub-optimal scoring: this means that the best possible match to a given weight matrix will always score 0.0, while other sequences will receive negative scores.

Typical usage:

```
nmscan -seqs seqfile.fa -motifs motifs.xmls
      -scoreThreshold -5.0 -strand both
```

3 Using NestedMICA for protein motif discovery

--by Mutlu Doğruel, md5@sanger.ac.uk

As of version 0.7.3, NestedMICA is capable of finding motifs in protein sequences [?]. For protein analysis, each NestedMICA tool (`nminfer`, `nmevaluatebg`, `nmmakebg`, etc.) must be run with the `-alphabet protein` option. While the underlying principle remains the same as DNA motif finding, some extra care must be taken during background model training and parameter choice.

3.1 Background optimization and training

Experience suggests that protein sequences are complicated enough to prevent us from easily saying that a pre-determined number of mosaic classes or background order would be optimal. Most of the time, training a dedicated background model for each protein sequence set is needed to maximise performance and sensitivity.

As described in Section 2.5, in order to find a reasonable pair of order and class number for background creation, it is possible to use a build-in application called `nmevaluatebg`. If you divide your data into two parts, you can use `nmevaluatebg` in the following way:

```
nmevaluatebg -trainSeqs proteinsTrain.fa
             -testSeqs proteinsTest.fa
             -alphabet protein -order 0
             -minClasses 1 -maxClasses 7
```

Unless `-trainSeqs` and `-testSeqs` parameters are used, NestedMICA will randomly use a certain number of sequences from a dataset provided with the

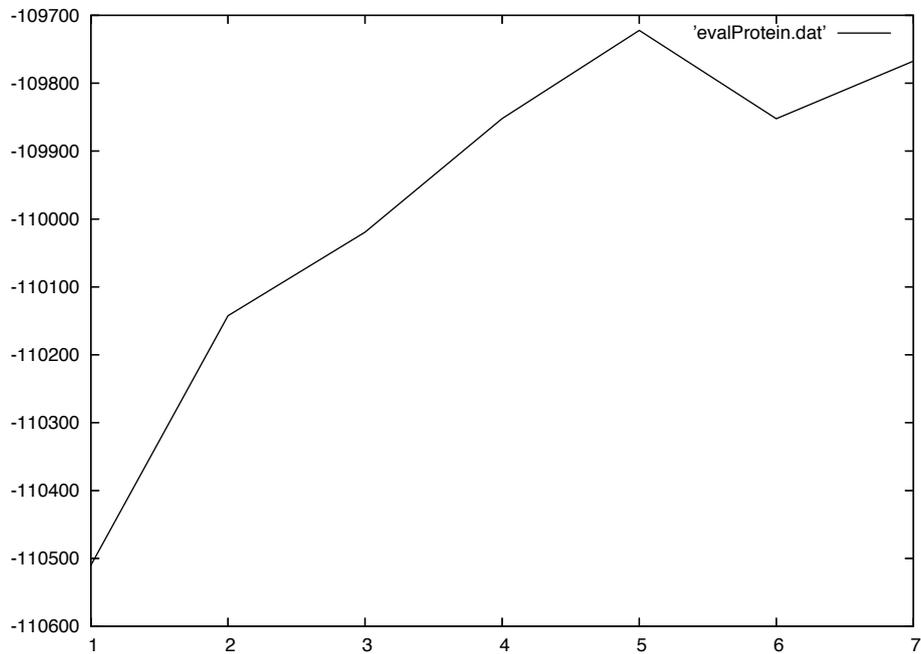


Figure 4: Likelihood curve for protein background models with different number of mosaic classes.

`-seqs` option. This would not be ideal if, for example, you would like to compare the likelihood of an order-1 background with a previous run.

While, in principle, using higher orders and larger number of mosaic classes would fit the data better and better, this will not hold true after a while when there is no sufficient data to train a more complicated model properly. Generally, it's a good idea to pick a class number before the likelihood curve starts to drop, no matter of it is increasing at a later stage. Figure 4, for example, shows the values of an `nmevaluatebg` optimization for a small protein set. In this particular case, choosing 5 or 4 mosaic classes would be a safer option and more optimal.

Having found a good pair of parameters, a background model is trained in a similar way as described in Section 2.4:

```
nmakebg -seqs proteins.fa -order 0 -classes 4
        -alphabet protein -out proteinbg.xml
```

3.2 Motif finding options

While all user parameters remain the same, a small number of NestedMICA parameters have different default values for proteins. `-ensembleSize` is the most notable one:

When searching for up to 3 motifs, the default behaviour is to set the `ensembleSize` to 1500 divided by the number of motifs. When more number of motifs are searched, it is set automatically to 500, unless there is a user specified value. This parameter can affect the speed of the program dramatically.

References

- [1] Down TA, Hubbard TJP: **NestedMICA: sensitive inference of overrepresented motifs in biological sequence**. *Nucleic Acids Res.* 2005, **33**:1445–1453.
- [2] Bailey T, Elkan C: **Fitting a mixture model by expectation maximization to discover motifs in biopolymers**. In *Proceedings of the Second International Conference on Intelligent Systems for Molecular Biology* 1994:28–36.
- [3] Thompson W, Rouchka E, Lawrence C: **Gibbs Recursive Sampler: finding transcription factor binding sites**. *Nucleic Acids Research* 2003, **31**:358–3585.
- [4] Bucher P: **Weight matrix descriptions of four eukaryotic RNA polymerase II promoter elements derived from 502 unrelated promoter sequences**. *Journal of Molecular Biology* 1990, **212**:563–578.
- [5] Skilling J: **Nested Sampling**[[<http://www.inference.phy.cam.ac.uk/bayesys/>]].
- [6] Thijs G, Lescot M, Marchal K, Rombauts S, De Moor B, Rouze P, Moreau Y: **A higher-order background model improves the detection of promoter regulatory elements by Gibbs sampling**. *Bioinformatics* 2001, **17**:1113–1122.
- [7] Down TA, Bergman CM, Su J, Hubbard TJP: **Large scale discovery of promoter motifs in *Drosophila melanogaster***. *PLoS Comp Biol.* 2007, **3**:e7.
- [8] Dogruel M, Down TA, Hubbard TJP: **NestedMICA as an ab initio protein motif finder**. *Submitted* 2007.