

# hca

Wray Buntine

November 27, 2013

Version 0.4

## Abstract

*hca* is research software that does various versions of non-parametric topic models using Gibbs sampling including LDA, HDP-LDA, NP-LDA, all with/without burstiness modelling. Various diagnostics, “document completion” testing and coherence measurements with PMI are also supported.

## 1 Synopsis

*hca* [-?] [-?Arg] *DataStem* *RepStem*

## 2 Description

*hca* reads the collection of files with stem *DataStem* that form the input set of data. When checkpointing, or at termination, the output is written to files with stem *RepStem*. On restart with the **-r0** option, some of these are also read initially to restore the previous state. A log of the run is reported to **stderr** if the **-e** option is used. By default, the log goes to **RepStem.log**.

The programme is research software, so not all options or combinations of options work correctly. Note that in this release, all the more experimental features have been stripped, so this release contains only moderately well tested components.

The programme runs a Gibbs sampler for a variety of non-parametric topic models. The model selected has three parts:

**alpha:** this is the prior on topic vector (theta) for each document. LDA has a simple symmetric Dirichlet with parameter alpha and the vector has dimension  $T$  (the number of topics).

**beta:** this is the prior on word vector (phi) for each topic. LDA has a simple symmetric Dirichlet with parameter beta and the vector has dimension  $W$  (the number of words).

**burst:** this is the burstiness component which has a document specific variant of the word vector for each topic. This is not used by default.

These parts are set using the **-S**, **-A** and **-B** options.

The programme uses generalised second order Stirling numbers with the library extracted from *libstb* version 1.7 released at <http://mloss.org>. This is initialised with predefined bounds on the tables, and these can be modified with the **-N** option. This should be used for collections with more than 20,000 documents, but its best to run first and on error, increase the bounds.

## 3 Options

Options have a single letter followed by a possible single argument. Options are grouped under the following functions: *setting of hyperparameters*, *controlling sampling of hyperparameters*, *general control*, and *testing and reports*

### 3.1 Setting of hyperparameters

- Avalue** Use a symmetric Dirichlet prior for alpha with this **value** for each dimension. The value must be a positive float.
- Atype** A second version of the **-A** option with a string argument. It defines an alpha prior of **type** as follows:
- hdp** alpha is modelled with a symmetric Dirichlet and theta is modelled with a Dirichlet Process.
  - hpdd** alpha is modelled with a (truncated) GEM and theta is modelled with a Pitman-Yor Process. This is the default.
  - pdp** alpha is uniform and theta is modelled with a Pitman-Yor Process.
- Bvalue** Use a symmetric Dirichlet prior with this **value** as the total over dimensions, so each dimension gets  $\text{value}/W$  for  $W$  being the number of words. The value must be a positive float.
- Btype** Second form of the **-B** option. Use a beta prior of **type** “hdp” “hpdd” or “pdp”. Similar to the **-A** option.
- hdp** beta is modelled with a Dirichlet, by default symmetric, or its mean can be set with the **-u** option, and phi is modelled with a Dirichlet Process.
  - hpdd** beta is modelled with a (truncated) GEM and phi is modelled with a Pitman-Yor Process. This is the default.
  - pdp** beta is by default uniform, or it can be set with the **-u** option, and phi is modelled with a Pitman-Yor Process. This is not hierarchical because beta is constant.
- Svar=value** Set variable **var** to float **value**, where **var** can be one of:
- a** discount parameter for the non-parametric distribution on the theta, topic distribution per document.
  - b** concentration parameter for the non-parametric distribution on theta, the topic distribution per document.
  - a0** discount parameter for the non-parametric distribution on alpha, the prior for theta.
  - b0** concentration parameter for the non-parametric distribution on alpha, the prior for theta.
  - aw** discount parameter for the non-parametric distribution on phi, word distribution per topic.
  - bw** concentration parameter for the non-parametric distribution on phi, word distribution per topic.
  - aw0** discount parameter for the non-parametric distribution on beta, prior for phi.
  - bw0** concentration parameter for the non-parametric distribution on beta, prior for phi.
  - ad** discount parameter for burstiness.
  - bdk** concentration parameter for burstiness, a constant initially but subsequent sampling will allow a different value per topic.
- ufile** the value for the beta vector is given by the text **file**. It should contain  $W$  floats,  $W$  being the number of words.

### 3.2 Controlling sampling of hyperparameters

- Dcycles,start** Start sampling **alpha** of the symmetric Dirichlet for alpha after **start** cycles and then repeat every **cycles** cycles.
- Ecycles,start** Start sampling **beta** of the symmetric Dirichlet for beta after **start** cycles and then repeat every **cycles** cycles.
- Fvar** Fix the variable **var** where it takes the value **alpha**, **beta** or one of the arguments to the **-S** option.
- Gvar,cycles,start** Sample the variable **var** where it takes the value **alpha**, **beta** or one of the arguments to the **-S** option. The **start** and **cycles** integers are used as for the **-D** option.

### 3.3 General control

- ccycles** Do a checkpoint every this many **cycles**. This saves the output statistics and the parameter file adequate to do a restart with **-r0** option.
- Ccycles** Stop after this many **cycles**. Default is 100.
- ddots** For really big batches of data, print a “.” every **dots** documents within a single cycle.
- e** Reroute logging to the **stderr**.
- fformat** Read input data from data formatted according to the type **format**. Data is expected to come from an input file with name **DataStem.Suff** where **Suff** is an appropriate suffix. These are given with Input Files below. Allowed formats are: **ldac**, **witdit**, **docword**, **bag** and **lst**.
- Ktopics** Set *T* the maximum number of topics. Default is 10.
- Mmaxtime** Quit early when total training time exceeds this many seconds.
- NmaxT,maxN** Set maximum for the Stirling number tables to count **maxN** and table count **maxT**. Default is 1000,10000. On collections with more than 20k documents, can require more.
- roffset** Restart. Currently must use **offset** equal to “0” for a normal restart.
- rphi** Second version of the **-r** option using the string “phi” as the argument. Restart but now fix the word by topic matrix to the previous value saved at **RepStem.phi**, and the beta side is held constant and not sampled. Can significantly speed up testing or querying sometimes.
- sseed** Initialise the random number seed.
- v** Up verbosity by one increment. Starts at zero and currently understands 0-3.

### 3.4 Testing and reports

- hHold** Do document completion testing on the test set. If **Hold** is in between [0,1], use it as the proportion of the tail of the document to evaluation perplexity on. The initial proportion is used to estimate latent variables (like theta) for the document. If **Hold** is greater than or equal to 2, use every **Hold**-th word in the document instead.
- lDiag,cycles,start** Do a run-time estimation of the diagnostic **Diag** starting after the **start** cycle and then taking the estimate every **cycles** cycle. Diagnostics are:
  - sp** Estimate topic sparsity in the theta matrix for the words given in **DataStem.smap**. Results placed in **RepStem.smap**. The report gives “topic/weight” for topics including the word.
  - prob** Estimate the topic probability vector for each document (train or test). Stored in the **RepStem.tprob** file.

- prog** How often to do the standard diagnostic reports (default is every 5-th cycle).
- phi** Estimate the word probability vector for each topic. Stored in the `RepStem.phi` file.
- L`Diag,cycles,start`** Do a diagnostic estimate `Diag` after all Gibbs sampling is complete. Sampling of the estimate starts after the `start` cycle and goes for a total of `cycles` cycles (including the starting ones). Diagnostics are:
- class** Estimate class probabilities with “true” classes given in `DataStem.class` and then produce confusion matrix for the test data. Output to files `DataStem.cnfs` and `DataStem.pcnfs`.
- like** Estimate likelihood/perplexity on the test set using the standard (biased) document likelihood, or document completion if the **-h** option is used. Can also be instigated during run-time with the **-P** option.
- O** Report log likelihood, not log perplexity. Both are done in base 2.
- p** Report topic coherency in the log file, and save results in the parameter file. This requires a `DataStem.pmi` or `DataStem.pmi.gz` file exist in the right format. This can be created with the `mkmat.pl` and `cooc2pmi.pl` scripts in the scripts directory of the release. The format is a simple sparse matrix form with lines of the form “N M PMI” for word indices (offset by 0) N and M and PMI value.
- P`secs`** Calculate test perplexity (using document completion) every interval in `secs` seconds. If Gibbs cycles are long, will report only after the cycle finishes.
- t`size`** Specify size of training set. It takes the first `size` entries in the data set. Default is all the set minus the test data.
- T`size`** Specify size of test set. It takes the `size` entries immediately following the training set. Default is zero.
- V** load the dictionary from the `DataStem.tokens` file for use in reporting. It has one token per line.
- X** Instigate report on naive Bayes classification using the topic model and classes given in `DataStem.class` file. The report is a confusion matrix to file `RepStem.tbyc` built on the training data.

## 4 Input Files

The following files provide details about the dataset. The filenames are constructed by adding a suffix to the data stem. The data (document+word) format itself can be one of four different formats and is specified with the **-f** option.

**DataStem.class** Class index for each document, one per line. Optional file used with some reports instigated by **-X** or **-L`class`** options.

**DataStem.dit+DataStem.wit** Simple document index and word index files, both indices offset by 1, one index per line. Words in the collection are listed by document. The `DataStem.dit` file gives the document index, and the corresponding line in `DataStem.wit` gives the dictionary index.

**DataStem.docword** This format appears in some UCI data sets at <http://archive.ics.uci.edu/ml/datasets/Bag+of+Words>. Word indices offset by 1.

**DataStem.ldac** Standard LdaC format. Word indices to the dictionary are offset by 1.

**DataStem.smap** A list of word indices (offset by 0) about which one wants a sparsity report generated. The report is instigated by the **-lsp** option.

**DataStem.tokens** tokens/words in the dictionary, one per line. Optional file used with **-V** option.

**DataStem.txtbag** default bag or list format for *linkBags(1)* command of DCABags. Word indices offset by 0.

The various output files such as **RepStem.par** (Parameter and dimension output file) are also read on restart with the **-r0** option.

## 5 Output Files

The following files are output when the system checkpoints or at the end of the run. These are built by adding a suffix to the report stem, **RepStem**. The first set of files are:

**RepStem.beta** If a constant beta vector is specified using the **-u** option, this saves the value, for possible use in a restart.

**RepStem.cnfs+RepStem.pcnfs** Best prediction and probability vector confusion matrices built on the test data with the **-Lclass** command.

**RepStem.log** Log file created if **-e** option not used.

**RepStem.par** Parameter and dimensions file in simple “var = value” format. These are detailed in the next section.

**RepStem.phi** The Phi matrix written as a binary file: first  $W$  (dictionary size),  $T$  (topics),  $C$  (sample size) are written as 32 bit integers and then the full Phi matrix as native floats with  $W$  as the minor index. Only generated with appropriate use of the **-lphi** option.

**RepStem.smap** Optional sparsity report on the word indices listed in **DataStem.smap**. The report is instigated by the **-lsp** option.

**RepStem.tbyc** Optional confusion matrix printed when the **-X** option is used.

**RepStem.top** A simple text report giving the top word indices for each topic. If a hierarchical model in use, then the “-1” topic is for the base distribution of words. Word indices are offset from 0.

**RepStem.tprob** Topic probabilities for each document written in a simple sparse form. The class index (“-1” or “+1” for binary classes, otherwise just the index) is also added if it exists. Topic indices are offset by 1. Only generated with appropriate use of the **-lprob** option.

The second set of files gives the actual runtime statistics. Output matrices are in a simple readable sparse vector format the same as the **DataStem.docword** format.

**RepStem.ndt** Document by topic counts.

**RepStem.nwt** Word by topic counts.

**RepStem.tdt** Document by topic table counts if the Alpha side of the model is non-parametric.

**RepStem.twt** Word by topic table counts if the Beta side of the model is non-parametric.

**RepStem.zt** With no burstiness, gives topic index (offset by 0), one per line. With burstiness, gives one “z,r” per line where “z” is the topic index (offset by 0) and “r” is the burst table indicator, which is 1 if the word contributes to standard LDA statistics, and 0 if burstiness modelling says the word is a burst so does not contribute to LDA statistics.

These files along with **RepStem.par** are input on a restart using **-r0**.

## 6 The Parameter File

The parameter file has the following *dimensions*:

N – number of words in the full collection, summed over all documents.

NT – number of words in the training set, summed over all training documents.

W – number of words in the dictionary.

D – number of documents in total.

TRAIN – number of documents to train on, is always the the first ones in the file.

TEST – number of documents to test on, is always the the last ones in the file.

T – maximum number of topics.

ITER – number of major cycles made last.

In addition, the float parameters allowed to be specified with the **-F** and **-G** options are also given. Finally, the type of model for alpha as specified by the **-A** option is coded in the `PYalpha` variable. It is 0 if the mode is a symmetric Dirichlet, the LDA default. It is 1 for hdp, 2 for hpdd and 3 for pdp. Likewise for the `PYbeta` variable and the **-B** option.

## 7 Examples

These examples work as is on late model Linux, Macs and Windows. However, you need to replace the executable, `hca`, by the system dependent one, from the install directory where the `data/` directory is. For instance, on Windows that might be `hca/hca.exe`.

Run basic LDA with parameter fitting on the full dataset and no testing, sending logging to `stderr`.

```
hca -v -e -V -K20 -C100 data/ch c1
```

The parameters as they are sampled will be printed on a line beginning with “Par:”. At the end, the top 20 words will be printed and the final training perplexity printed. This is based on the posterior probability, not word probability estimates. Restart and build a topic probability vector for each document, as well as sparsity mappings for the words in `data/ch.smap` file.

```
hca -v -r0 -e -lsp,2,1 -lprob,2,1 -C20 data/ch c1
```

Now view the sparsity report at `c1.smap` and the topic probabilities at `c1.tprob`, and the values saved in the parameter file `c1.par`.

Rerun basic LDA with parameter fitting on the full dataset and testing on the final 100 documents. Logging now to `c1.log`. Checkpoint every 20 cycles (note, we usually only do this for cycles taking over 10 minutes each).

```
hca -v -V -K20 -C100 -c20 -T100 data/ch c1
```

View the end of the log file to get the test perplexity, which is printed after “log\_2(test perpML)”. Now rerun but use document completion (every 4th word), not the default likelihood calculation.

```
hca -v -V -r0 -h4 data/ch c1
```

View the end of the log file to get the test perplexity, which is printed after “log\_2(test perpHold)”. Note it is also recorded in the parameter file. Restart, no sampling this time using **-C0**, just record the PMI and the classification details on test data.

```
hca -v -V -r0 -C0 -Llike,0,0 -Lclass,10,1 -p data/ch c1
```

Note the **-Llike,0,0** option is used to prevent it doing test likelihood calculations, which are potentially slow on larger data sets. The PMI data has a value printed for each topic as well as a final average. It bases its calculations on the matrix `data/ch/pmi.gz` created explicitly for this test set. For other datasets, you will need to download prepared PMI matrices from the project homepage.

The burstiness version is new. Our best bet, currently is to run with optimisation of the hyperparameters:

```
hca -v -e -V -K20 -C100 -Sbdk=200 -Sad=0.5 data/ch c1
```

So the initial discount for the bursty topics is 0.5 and its a shared parameter for all. The concentration we set quite high initially with “bdk=200”, and these will be sampled separately with each topic, so bdk is a vector in the parameter file. The hyperparameter sampling slows it down quite a bit but seems to make a significant difference. Unused topics sometimes get a very low concentration.

## 8 Errors

There is some error reporting on failure.

If the software quits during a run on larger data with an error message like:

```
S_V(N,M,A) tagged 'XXX' hit bounds (BN,BM)
```

for integers N,M and label XXX then you need to increase the bounds. If the tag XXX is “SX, docXtopic PYP” then increase the bound BM using the option **-NBM** with an increased BM. If the tag XXX is “SX, docXtopic PYP” then increase both bounds BN,BM using the option **-NBM,BN** (note the order of the bounds).

For other errors, please report to the maintainer.

## 9 See Also

The command *linkBags*(1) is available from <http://mloss.org> in the software *DCABags*. The extended library *libstb*, parts of which are included, is available individually from <http://mloss.org>.

Other supporting software and data sets are available at the project home page at <http://www.nicta.com.au/people/buntinew/softwareanddata>

## 10 Version

This programme is version 0.4 of November 27, 2013. This incorporates parts of the library *libstb* version 1.7 also of November 27, 2013.

## 11 License and Copyright

**Copyright** © 2011-2013, Dr. Wray Buntine, NICTA, Canberra, Australia,  
[wray.buntine@nicta.com.au](mailto:wray.buntine@nicta.com.au)

**License** This code for the programme are subject to the Mozilla Public License Version 1.1 (the “License”); you may not use this file except in compliance with the License. You may obtain a copy of the License at <http://www.mozilla.org/MPL/>.

## 12 Author

Dr. Wray Buntine  
Email: [Wray.Buntine@NICTA.com.au](mailto:Wray.Buntine@NICTA.com.au)  
WWW: <http://www.nicta.com.au/people/buntinew>.