

ccfindR: single-cell RNA-seq analysis using non-negative matrix factorization

2018-04-30

The `ccfindR` (Cancer Clone findeR) package contains implementations and utilities for analyzing single-cell RNA-sequencing data, including quality control, unsupervised clustering for discovery of cell types, and visualization of the outcomes. It is especially suitable for analysis of transcript-count data utilizing unique molecular identifiers (UMIs), e.g., data derived from 10x Genomics platform. In these data sets, RNA counts are non-negative integers, enabling clustering using non-negative matrix factorization (NMF)¹.

Input data are UMI counts in the form of a matrix with each genetic feature (“genes”) in rows and cells (tagged by barcodes) in columns, produced by read alignment and counting pipelines. The count matrix and associated gene and cell annotation files are bundled into a main object of class `scNMFSet`, which extends the `SingleCellExperiment` class [<http://dx.doi.org/10.18129/B9.bioc.SingleCellExperiment>]. Quality control for both cells and genes can be performed via filtering steps based on UMI counts and variance of expressions, respectively. The NMF factorization is first performed for multiple values of ranks (the reduced dimension of factorization) to find the most likely value. A production run for the chosen rank then leads to factor matrices, allowing the user to identify and visualize genes representative of clusters and assign cells into clusters.

Algorithm

The NMF approach offers a means to identify cell subtypes and classify individual cells into these clusters based on clustering using expression counts. In contrast to alternatives such as principal component analyses², NMF leverages the non-negative nature of count data and factorizes the data matrix \mathbf{X} into two matrices \mathbf{W} and \mathbf{H}^1 :

$$\mathbf{X} \sim \mathbf{W}\mathbf{H}.$$

If \mathbf{X} is a $p \times n$ matrix (p genes and n cells), the basis matrix \mathbf{W} is $p \times r$ and coefficient matrix \mathbf{H} is $r \times n$ in dimension, respectively, where the rank r is a relatively small integer. A statistical inference-based interpretation of NMF is to view X_{ij} as a realization of a Poisson distribution with the mean for each matrix elements given by $(\mathbf{W}\mathbf{H})_{ij} \equiv \Lambda_{ij}$, or

$$\Pr(x_{ij}) = \frac{e^{-\Lambda_{ij}} \Lambda_{ij}^{x_{ij}}}{\Gamma(1 + x_{ij})}.$$

The maximum likelihood inference of the latter is then achieved by maximizing

$$L = \sum_{ij} \left(X_{ij} \ln \frac{\Lambda_{ij}}{X_{ij}} - \Lambda_{ij} + X_{ij} \right).$$

The Kullback-Leibler measure of the distance between \mathbf{X} and Λ , which is minimized, is equal to $-L$. Lee and Seung’s update rule¹ solves this optimization task iteratively. In addition to this classical iterative update algorithm to find basis and coefficient factors of the count matrix, the `ccfindR` package implements variational Bayesian inference developed by Cemgil³.

Key features of `ccfindR` distinguishing it from other existing implementations – NMF for generic data⁴ and NMFEM for single-cell analysis⁵ – are

- Bayesian inference allowing for a statistically well-controlled procedure to determine the most likely value of rank r .
- Procedure to derive hierarchical relationships among clusters identified under different ranks.

A traditional way (in maximum likelihood inference) to determine the rank is to evaluate the factorization quality measures (and optionally compare with those from randomized data). The Bayesian formulation of NMF algorithm³ incorporates priors for factored matrix elements W and H modeled by gamma distributions. Inference can be combined with hyperparameter update to optimize the evidence (conditional probability of data under hyperparameters and rank), which provides a statistically well-controlled means to determine the optimal rank describing data.

For large rank values, it can be challenging to interpret clusters identified. To facilitate biological interpretation, we provide a procedure where cluster assignment of cells is repeated for multiple rank values, typically ranging from 2 to the optimal rank, and a phylogenetic tree connecting different clusters at neighboring rank values are constructed. This tree gives an overview of different types of cells present in the system viewed at varying resolution.

Workflow

We illustrate a typical workflow with a single-cell count data set generated from peripheral blood mononuclear cell (PBMC) data [<https://support.10xgenomics.com/single-cell-gene-expression/datasets/1.1.0/>]. The particular data set used below was created by sampling from 5 purified immune cell subsets.

1. Installation

To install the package, download the source tar-ball and

```
$ R CMD INSTALL ccfindR_0.99-0.tar.gz
```

After installation, load the package by

```
library(ccfindR)
```

```
## Package 'ccfindR' version 1.0.0
```

2. Data input

The input data can be a simple matrix:

```
# A toy matrix for count data
set.seed(1)
mat <- matrix(rpois(n = 80, lambda = 2), nrow = 4, ncol = 20)
ABC <- LETTERS[1:4]
abc <- letters[1:20]
rownames(mat) <- ABC
colnames(mat) <- abc
```

The main S4 object containing data and subsequent analysis outcomes is of class `scNMFSet`, created by

```
# create scNMFSet object
sc <- scNMFSet(count = mat)
```

This class extends `SingleCellExperiment` class, adding extra slots for storing factorization outcomes. In particular, `assays`, `rowData`, and `colData` slots of `SingleCellExperiment` class are used to store RNA count matrix, gene, and cell annotation data frames, respectively. In the simplest initialization above, the named argument `count` is used as the count matrix and is equivalent to

```
# create scNMFSet object
sc <- scNMFSet(assays = list(counts = mat))
```

See `SingleCellExperiment` documentations for more details of these main slots. For instance, row and column names can be stored by

```
# set row and column names
suppressMessages(library(S4Vectors))
genes <- as(ABC, 'DataFrame')
rownames(genes) <- ABC
cells <- as(abc, 'DataFrame')
rownames(cells) <- abc
sc <- scNMFSet(count = mat, rowData = genes, colData = cells)
sc
```

```
## class: scNMFSet
## dim: 4 20
## metadata(0):
## assays(1): counts
## rownames(4): A B C D
## rowData names(1): X
## colnames(20): a b ... s t
## colData names(1): X
## reducedDimNames(0):
## spikeNames(0):
```

Alternatively, sparse matrix format (of class `dgCMatrix`) can be used. One may read a `MatrixMarket` format file directly:

```
# read sparse matrix
dir <- system.file('extdata', package = 'ccfindR')
mat <- Matrix::readMM(paste0(dir, '/matrix.mtx'))
sc <- scNMFSet(count = mat, rowData = as(1:nrow(mat), 'DataFrame'),
              colData = as(1:ncol(mat), 'DataFrame'))
sc
```

```
## class: scNMFSet
## dim: 1030 450
## metadata(0):
## assays(1): counts
## rownames: NULL
## rowData names(1): X
## colnames: NULL
## colData names(1): X
## reducedDimNames(0):
## spikeNames(0):
```

The number of rows in `assays$counts` and `rowData`, the number of columns in `assays$counts` and rows in `colData` must match.

The gene and barcode meta-data and count files resulting from 10x Genomics' Cell Ranger pipeline (<https://support.10xgenomics.com/single-cell-gene-expression/software/pipelines/latest/what-is-cell-ranger>) can also be read:

```
# read 10x files
sc <- read_10x(dir = dir, count = 'matrix.mtx', genes = 'genes.tsv',
              barcodes = 'barcodes.tsv')
sc
```

```
## class: scNMFSet
## dim: 1030 450
```

```
## metadata(0):
## assays(1): counts
## rownames(1030): ENSG00000187608 ENSG00000186891 ...
##   ENSG00000198886 ENSG00000198727
## rowData names(2): V1 V2
## colnames(450): ATGCAGTGCTTGGA-1 CATGTACTCCATGA-1 ...
##   ACTATCACTCAGTG-1 ACAATAACTAGAGA-1
## colData names(1): V1
## reducedDimNames(0):
## spikeNames(0):
```

The parameter `dir` is the directory containing the files. Filenames above are defaults and can be omitted. The function returns an `scNMFSet` object. By default, any row or column entirely consisting of zeros in `counts` and the corresponding elements in `rowData` and `colData` slots will be removed. This feature can be turned off by `remove.zeros = FALSE`.

3. Quality control

For quality control, cells and genes can be filtered manually using normal subsetting syntax of R: the slots in the object `sc` are accessed and edited using accessors and sub-setting rules; see `SingleCellExperiment`:

```
# slots and subsetting
counts(sc)[1:7,1:3]
```

```
## 7 x 3 sparse Matrix of class "dgCMatrix"
##           ATGCAGTGCTTGGA-1 CATGTACTCCATGA-1 GAGAAATGGCAAGG-1
## ENSG00000187608           .                 2                 .
## ENSG00000186891           .                 .                 .
## ENSG00000127054           .                 .                 .
## ENSG00000158109           .                 .                 .
## ENSG00000116251           .                 3                 .
## ENSG00000074800           2                 .                 .
## ENSG00000162444           .                 .                 .
```

```
head(rowData(sc))
```

```
## DataFrame with 6 rows and 2 columns
##           V1           V2
##   <character> <character>
## 1 ENSG00000187608       ISG15
## 2 ENSG00000186891       TNFRSF18
## 3 ENSG00000127054       CPSF3L
## 4 ENSG00000158109       TPRG1L
## 5 ENSG00000116251       RPL22
## 6 ENSG00000074800       ENO1
```

```
head(colData(sc))
```

```
## DataFrame with 6 rows and 1 column
##           V1
##   <character>
## ATGCAGTGCTTGGA-1 ATGCAGTGCTTGGA-1
## CATGTACTCCATGA-1 CATGTACTCCATGA-1
## GAGAAATGGCAAGG-1 GAGAAATGGCAAGG-1
## TGATATGACGTTAG-1 TGATATGACGTTAG-1
## AGTAGGCTCGGGAA-1 AGTAGGCTCGGGAA-1
```

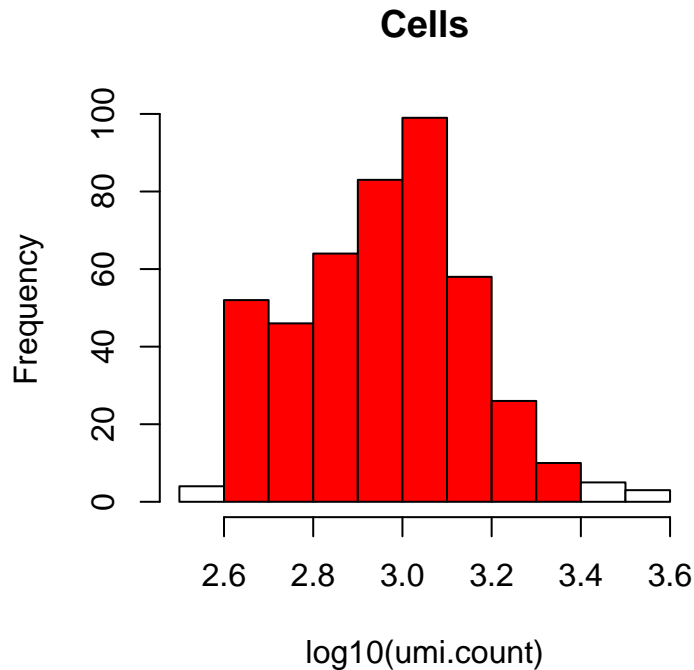


Figure 1: Quality control filtering of cells. Histogram of UMI counts is shown. Cells can be selected (red) by setting lower and upper thresholds of the UMI count.

```
## TGACCGCTGTAGCT-1 TGACCGCTGTAGCT-1
```

```
sc2 <- sc[1:20,1:70]      # subsetting of object
sc2 <- remove_zeros(sc2) # remove empty rows/columns
```

```
## 6 empty genes removed
```

```
sc2
```

```
## class: scNMFSet
## dim: 14 70
## metadata(0):
## assays(1): counts
## rownames(14): ENSG00000187608 ENSG00000186891 ... ENSG00000117318
## ENSG00000142676
## rowData names(2): V1 V2
## colnames(70): ATGCAGTGCTTGGA-1 CATGTACTCCATGA-1 ...
## ACAGTGACAGTAGA-1 TCCGAAGAAGCCAT-1
## colData names(1): V1
## reducedDimNames(0):
## spikeNames(0):
```

We provide two streamlined functions each for cell and gene filtering as shown below:

```
sc <- filter_cells(sc, umi.min = 10^2.6, umi.max = 10^3.4)
```

```
## 438 cells out of 450 selected
## 21 empty genes removed
```

```
markers <- c('CD4', 'CD8A', 'CD8B', 'CD19', 'CD3G', 'CD3D',
             'CD3Z', 'CD14')
```

```
sc0 <- filter_genes(sc, markers = markers, vmr.min = 1.5,
```




Figure 2: Selection of genes for clustering. The scatter plot shows distributions of expression variance to mean ratio (VMR) and the number of cells expressed. Minimum VMR and a range of cell number can be set to select genes (red). Symbols in orange are marker genes provided as input, selected irrespective of expression variance.

```
min.cells.expressed = 50, rescue.genes = FALSE)
```

```
## 5 marker genes found
## 297 variable genes out of 1009
## 299 genes selected
```

The function `filter_cells()` plots histogram of UMI counts for all cells when called without threshold parameters (Fig. 1). This plot can be used to set desirable thresholds, `umi.min` and `umi.max`. Cells with UMI counts outside will be filtered out. The function `filter_genes()` displays scatter plot of the total number of cells with nonzero count and VMR (variance-to-mean ratio) for each gene (Fig. 2). In both plots, selected cells and genes are shown in red. Note that the above example has thresholds that are too stringent, which is intended to speed up the subsequent illustrative runs. A list of pre-selected marker genes can be provided to help identify clusters via the `markers` parameter in `filter_genes()`. Here, we use a set of classical PBMC marker genes (shown in orange).

Gene-filtering can also be augmented by scanning for those genes whose count distributions among cells are non-trivial: most have zero count as its maximum; some have one or more distinct peaks at nonzero count values. These may signify the existence of groups of cells in which the genes are expressed in distinguishable fashion. The selection of genes by `filter_genes()` will be set as the union of threshold-based group and those with such nonzero-count modes by setting `rescue.genes = TRUE` (default):

```
sc_rescue <- filter_genes(sc, markers = markers, vmr.min = 1.5, min.cells.expressed = 50,
  rescue.genes = TRUE, progress.bar = FALSE)
```

```
## Looking for genes with modes ...
## 5 marker genes found
## 297 variable genes out of 1009
## 36 additional genes rescued
```




Figure 3: Additional selection of genes with modes at nonzero counts. Symbols in blue represent genes rescued.

```
## 333 genes selected
```

This “gene rescue” scan will take some time and a progress bar is displayed if `progress.bar = TRUE`.

For subsequent analysis, we will use the latter selection and also name rows with gene symbols:

```
rownames(sc_rescue) <- rowData(sc_rescue)[,2]
sc <- sc_rescue
```

4. Rank determination

The main function for maximum likelihood NMF on a count matrix is `factorize()`. It performs a series of iterative updates to matrices W and H . Since the global optimum of likelihood function is not directly accessible, computational inference relies on local maxima, which depends on initializations. We adopt the randomized initialization scheme, where the factor matrix elements are drawn from uniform distributions. To make the inference reproducible, one can set the random number seed by `set.seed(seed)`, where `seed` is a positive integer, prior to calling `factorize()`. Updates continue until convergence is reached, defined by either the fractional change in likelihood being smaller than `Tol` (`criterion = likelihood`) or a set number (`ncnn.step`) of steps observed during which the connectivity matrix remains unchanged (`criterion = connectivity`). The connectivity matrix C is a symmetric $n \times n$ matrix with elements $C_{jl} = 1$ if j and l cells belong to the same cluster and 0 otherwise. The cluster membership is dynamically checked by finding the row index k for which the coefficient matrix element H_{kj} is maximum for each cell indexed by j .

During iteration, with `verbose = 3`, step number, log likelihood per elements, and the number of terms in the upper-diagonal part of C that changed from the previous step are printed:

```
set.seed(1)
sc <- factorize(sc, ranks = 3, nrun = 1, ncnn.step = 1,
               criterion='connectivity', verbose = 3)
```

```

## Rank 3
## Run # 1 :
## 1 : likelihood = -0.569847 , connectivity change = 95703
## 2 : likelihood = -0.5571901 , connectivity change = 5038
## 3 : likelihood = -0.5490381 , connectivity change = 4126
## 4 : likelihood = -0.5427167 , connectivity change = 6999
## 5 : likelihood = -0.5367327 , connectivity change = 9493
## 6 : likelihood = -0.5297582 , connectivity change = 6316
## 7 : likelihood = -0.5204045 , connectivity change = 7176
## 8 : likelihood = -0.5075277 , connectivity change = 8037
## 9 : likelihood = -0.4911819 , connectivity change = 4352
## 10 : likelihood = -0.4731239 , connectivity change = 3916
## 11 : likelihood = -0.4555406 , connectivity change = 4466
## 12 : likelihood = -0.4397135 , connectivity change = 3695
## 13 : likelihood = -0.4260338 , connectivity change = 1746
## 14 : likelihood = -0.414567 , connectivity change = 3739
## 15 : likelihood = -0.4053164 , connectivity change = 3393
## 16 : likelihood = -0.3980694 , connectivity change = 1665
## 17 : likelihood = -0.3924104 , connectivity change = 2375
## 18 : likelihood = -0.3879222 , connectivity change = 1814
## 19 : likelihood = -0.3843011 , connectivity change = 1613
## 20 : likelihood = -0.3813472 , connectivity change = 1920
## 21 : likelihood = -0.3789228 , connectivity change = 3184
## 22 : likelihood = -0.3769243 , connectivity change = 1557
## 23 : likelihood = -0.3752681 , connectivity change = 341
## 24 : likelihood = -0.3738842 , connectivity change = 959
## 25 : likelihood = -0.3727151 , connectivity change = 1228
## 26 : likelihood = -0.3717172 , connectivity change = 1272
## 27 : likelihood = -0.3708584 , connectivity change = 1314
## 28 : likelihood = -0.3701144 , connectivity change = 1631
## 29 : likelihood = -0.3694659 , connectivity change = 1289
## 30 : likelihood = -0.3688975 , connectivity change = 516
## 31 : likelihood = -0.3683969 , connectivity change = 1029
## 32 : likelihood = -0.3679529 , connectivity change = 0
## Nsteps = 32 , likelihood = -0.3679529 , dispersion = 1
##
## Sample# 1 : Max(likelihood) = -0.3679529 , dispersion = 1 , copenetic = 1

```

The function `factorize()` returns the same object `sc` with extra slots `ranks` (the rank value for which factorization was performed), `basis` (a list containing the basis matrix W), `coeff` (a list containing the coefficient matrix H), and `measure` (a data frame containing the factorization quality measure; see below). The criterion used to stop iteration is either `connectivity` (no changes to connectivity matrix for `ncnn.steps`) or `likelihood` (changes to likelihood smaller than `Tol`).

To reduce the dependence of final estimates for W and H on initial guess, inferences need to be repeated for many different initializations:

```
sc <- factorize(sc, ranks = 3, nrun = 5, verbose = 2)
```

```

## Rank 3
## Run # 1 :
## Nsteps = 64 , likelihood = -0.3622432 , dispersion = 1
##
## Run # 2 :
## Nsteps = 55 , likelihood = -0.376752 , dispersion = 0.7215237
##

```



```

## Run # 3 :
## Nsteps = 76 , likelihood = -0.3627259 , dispersion = 0.7489812
##
## Run # 4 :
## Nsteps = 75 , likelihood = -0.3642667 , dispersion = 0.7813901
##
## Run # 5 :
## Nsteps = 104 , likelihood = -0.3630267 , dispersion = 0.7391948
##
## Sample# 1 : Max(likelihood) = -0.3622432 , dispersion = 0.7391948 , cophenetic = 0.9842846

```

After each run, the residual and dispersion are printed, and the global minimum of residual as well as the corresponding matrices W and H are stored. The dispersion ρ is a scalar measure of how close the consistency matrix $\bar{C} \equiv \text{Mean}(C)$ elements, where C is the connectivity matrix, are to binary values 0, 1. The mean is over multiple runs:

$$\rho = \frac{4}{n^2} \sum_{jl} (\bar{C}_{jl} - 1/2)^2.$$

Note in the output above that ρ decays from 1 as the number of runs increases and then stabilizes. This degree of convergence of ρ is a good indication for the adequacy of `nrun`. The cophenetic is the correlation between the distance $1 - \bar{C}$ and the height matrix of hierarchical clustering⁶.

To discover clusters of cells, the reduced dimensionality of factorization, or the rank r , must be estimated. The examples above used a single rank value. If the parameter `ranks` is a vector, the set of inferences will be repeated for each rank value.

```
sc <- factorize(sc, ranks = seq(3,7), nrun = 5, verbose = 1, progress.bar = FALSE)
```

```

## Rank 3
## Sample# 1 : Max(likelihood) = -0.3625648 , dispersion = 0.7917775 , cophenetic = 0.9853564
## Rank 4
## Sample# 1 : Max(likelihood) = -0.3268542 , dispersion = 0.9676104 , cophenetic = 0.9957969
## Rank 5
## Sample# 1 : Max(likelihood) = -0.3081752 , dispersion = 0.9147574 , cophenetic = 0.9845725
## Rank 6
## Sample# 1 : Max(likelihood) = -0.3018445 , dispersion = 0.9051196 , cophenetic = 0.9851415
## Rank 7
## Sample# 1 : Max(likelihood) = -0.2960238 , dispersion = 0.9228406 , cophenetic = 0.9872926

```

Note that `nrun` parameter above is set to a small value for illustration. In a real application, typical values of `nrun` would be larger. The progress bar shown by default under `verbose = 1` for overall `nrun` runs is turned off above. It can be set to `TRUE` here (and below) to monitor the progress. After factorization, the `measure` slot has been filled:

```
measure(sc)
```

```

##   rank likelihood dispersion cophenetic
## 1    3 -0.3625648  0.7917775  0.9853564
## 2    4 -0.3268542  0.9676104  0.9957969
## 3    5 -0.3081752  0.9147574  0.9845725
## 4    6 -0.3018445  0.9051196  0.9851415
## 5    7 -0.2960238  0.9228406  0.9872926

```

These measures can be plotted (Fig. 4):

```
plot(sc)
```




Figure 4: Factorization quality measures as functions of the rank. Dispersion measures the degree of bimodality in consistency matrix. Cophenetic correlation measures the degree of agreement between consistency matrix and hierarchical clustering.

5. Bayesian NMF

The maximum likelihood-based inference must rely on quality measures to choose optimal rank. Bayesian NMF allows for the statistical comparison of different models, namely those with different ranks. The quantity compared is the log probability (“evidence”) of data conditional to models (defined by rank and hyperparameters). The main function for Bayesian factorization is `vb_factorize()`:

```

sb <- sc_rescue
set.seed(2)
sb <- vb_factorize(sb, ranks =3, verbose = 3, Tol = 2e-4, hyper.update.n0 = 5)

## Rank 3
## Run # 1 :
## 1: log(evidence) = -1.505411, aw = 1, bw = 1, ah = 1, bh = 1
## 2: log(evidence) = -1.575041, aw = 1, bw = 1, ah = 1, bh = 1
## 3: log(evidence) = -1.608358, aw = 1, bw = 1, ah = 1, bh = 1
## 4: log(evidence) = -1.627154, aw = 1, bw = 1, ah = 1, bh = 1
## 5: log(evidence) = -1.639133, aw = 1, bw = 1, ah = 1, bh = 1
## 6: log(evidence) = -1.647411, aw = 0.5245557, bw = 0.9404055, ah = 2.800559, bh = 0.9890178
## 7: log(evidence) = -1.647935, aw = 0.5165306, bw = 0.9402905, ah = 2.876746, bh = 0.9890304
## 8: log(evidence) = -1.651081, aw = 0.5087818, bw = 0.9401596, ah = 2.87484, bh = 0.9890439
## 9: log(evidence) = -1.651345, aw = 0.4976702, bw = 0.9399658, ah = 2.80461, bh = 0.9890581
## 10: log(evidence) = -1.648552, aw = 0.4828468, bw = 0.9397002, ah = 2.679995, bh = 0.9890725
## 11: log(evidence) = -1.642816, aw = 0.4649239, bw = 0.939366, ah = 2.520109, bh = 0.9890867
## 12: log(evidence) = -1.634198, aw = 0.4452534, bw = 0.938978, ah = 2.345064, bh = 0.9891001
## 13: log(evidence) = -1.622659, aw = 0.4256377, bw = 0.9385576, ah = 2.171214, bh = 0.989112
## 14: log(evidence) = -1.608519, aw = 0.4072656, bw = 0.9381276, ah = 2.009665, bh = 0.989122
## 15: log(evidence) = -1.59282, aw = 0.3902048, bw = 0.9377068, ah = 1.866502, bh = 0.9891296
## 16: log(evidence) = -1.577083, aw = 0.3743271, bw = 0.9373079, ah = 1.743545, bh = 0.9891348
## 17: log(evidence) = -1.56265, aw = 0.3599941, bw = 0.9369359, ah = 1.63954, bh = 0.9891376
## 18: log(evidence) = -1.550257, aw = 0.3476776, bw = 0.9365903, ah = 1.551651, bh = 0.9891381
## 19: log(evidence) = -1.54004, aw = 0.337401, bw = 0.9362669, ah = 1.476737, bh = 0.9891362
## 20: log(evidence) = -1.531779, aw = 0.3289474, bw = 0.9359607, ah = 1.412033, bh = 0.989132
## 21: log(evidence) = -1.525123, aw = 0.3218827, bw = 0.9356671, ah = 1.355323, bh = 0.9891255
## 22: log(evidence) = -1.519728, aw = 0.3157648, bw = 0.9353826, ah = 1.304886, bh = 0.9891167

```

```

## 23: log(evidence) = -1.5153, aw = 0.3104483, bw = 0.9351046, ah = 1.259403, bh = 0.9891055
## 24: log(evidence) = -1.511599, aw = 0.3058824, bw = 0.9348314, ah = 1.217874, bh = 0.9890919
## 25: log(evidence) = -1.508429, aw = 0.3019294, bw = 0.9345619, ah = 1.179552, bh = 0.9890758
## 26: log(evidence) = -1.505648, aw = 0.2984252, bw = 0.9342954, ah = 1.143879, bh = 0.9890572
## 27: log(evidence) = -1.503154, aw = 0.2952752, bw = 0.9340311, ah = 1.11044, bh = 0.989036
## 28: log(evidence) = -1.500872, aw = 0.2924123, bw = 0.9337688, ah = 1.077925, bh = 0.9890122
## 29: log(evidence) = -1.498747, aw = 0.2897952, bw = 0.933508, ah = 1.048233, bh = 0.9889857
## 30: log(evidence) = -1.496719, aw = 0.2873186, bw = 0.9332484, ah = 1.02002, bh = 0.9889565
## 31: log(evidence) = -1.494763, aw = 0.2849743, bw = 0.9329899, ah = 0.9932674, bh = 0.9889245
## 32: log(evidence) = -1.492861, aw = 0.2828196, bw = 0.9327324, ah = 0.9679112, bh = 0.9888899
## 33: log(evidence) = -1.491012, aw = 0.280967, bw = 0.932476, ah = 0.9438943, bh = 0.9888524
## 34: log(evidence) = -1.489226, aw = 0.2794903, bw = 0.9322207, ah = 0.9211381, bh = 0.988812
## 35: log(evidence) = -1.487523, aw = 0.2783321, bw = 0.9319667, ah = 0.8995274, bh = 0.9887687
## 36: log(evidence) = -1.485911, aw = 0.2773973, bw = 0.9317142, ah = 0.8789431, bh = 0.9887223
## 37: log(evidence) = -1.484367, aw = 0.2766264, bw = 0.9314632, ah = 0.8593442, bh = 0.9886727
## 38: log(evidence) = -1.482874, aw = 0.2759599, bw = 0.9312139, ah = 0.8407493, bh = 0.9886199
## 39: log(evidence) = -1.481428, aw = 0.2753403, bw = 0.9309665, ah = 0.8231597, bh = 0.9885638
## 40: log(evidence) = -1.480027, aw = 0.2747491, bw = 0.9307212, ah = 0.8065463, bh = 0.9885045
## 41: log(evidence) = -1.478667, aw = 0.2742057, bw = 0.9304785, ah = 0.7908618, bh = 0.9884418
## 42: log(evidence) = -1.477339, aw = 0.2737161, bw = 0.9302386, ah = 0.7760516, bh = 0.9883758
## 43: log(evidence) = -1.476032, aw = 0.2732587, bw = 0.9300023, ah = 0.7620592, bh = 0.9883064
## 44: log(evidence) = -1.474736, aw = 0.2728171, bw = 0.9297701, ah = 0.748827, bh = 0.9882337
## 45: log(evidence) = -1.473444, aw = 0.272395, bw = 0.9295428, ah = 0.7362946, bh = 0.9881578
## 46: log(evidence) = -1.472157, aw = 0.2720142, bw = 0.929321, ah = 0.7243985, bh = 0.9880785
## 47: log(evidence) = -1.470875, aw = 0.2716963, bw = 0.9291052, ah = 0.7130831, bh = 0.9879961
## 48: log(evidence) = -1.469597, aw = 0.2714583, bw = 0.9288959, ah = 0.7023144, bh = 0.9879104
## 49: log(evidence) = -1.468322, aw = 0.2713008, bw = 0.9286936, ah = 0.6920732, bh = 0.9878216
## 50: log(evidence) = -1.46706, aw = 0.2712097, bw = 0.9284983, ah = 0.6823365, bh = 0.9877297
## 51: log(evidence) = -1.465826, aw = 0.2711722, bw = 0.9283101, ah = 0.6730705, bh = 0.9876347
## 52: log(evidence) = -1.464635, aw = 0.2711773, bw = 0.9281289, ah = 0.6642343, bh = 0.9875367
## 53: log(evidence) = -1.463494, aw = 0.2712114, bw = 0.9279545, ah = 0.6557894, bh = 0.9874359
## 54: log(evidence) = -1.462395, aw = 0.2712614, bw = 0.9277868, ah = 0.6477044, bh = 0.9873321
## 55: log(evidence) = -1.461326, aw = 0.271326, bw = 0.9276257, ah = 0.6399515, bh = 0.9872256
## 56: log(evidence) = -1.460276, aw = 0.2714255, bw = 0.9274711, ah = 0.632499, bh = 0.9871164
## 57: log(evidence) = -1.459241, aw = 0.2715837, bw = 0.9273225, ah = 0.6253032, bh = 0.9870045
## 58: log(evidence) = -1.458231, aw = 0.2718038, bw = 0.9271793, ah = 0.6182993, bh = 0.9868901
## 59: log(evidence) = -1.457287, aw = 0.2720781, bw = 0.9270405, ah = 0.6113954, bh = 0.9867731
## 60: log(evidence) = -1.456486, aw = 0.2723851, bw = 0.9269046, ah = 0.6044818, bh = 0.9866535
## 61: log(evidence) = -1.455889, aw = 0.2726905, bw = 0.9267699, ah = 0.5974664, bh = 0.9865313
## 62: log(evidence) = -1.455475, aw = 0.2729715, bw = 0.9266353, ah = 0.5903158, bh = 0.9864064
## 63: log(evidence) = -1.455123, aw = 0.273229, bw = 0.9265011, ah = 0.5830669, bh = 0.9862789
## 64: log(evidence) = -1.454697, aw = 0.2734769, bw = 0.9263686, ah = 0.5757986, bh = 0.9861488
## 65: log(evidence) = -1.454134, aw = 0.2737247, bw = 0.9262393, ah = 0.5685917, bh = 0.9860161
## 66: log(evidence) = -1.453448, aw = 0.2739733, bw = 0.9261144, ah = 0.5615047, bh = 0.9858808
## 67: log(evidence) = -1.452687, aw = 0.2742239, bw = 0.9259944, ah = 0.5545732, bh = 0.985743
## 68: log(evidence) = -1.451889, aw = 0.2744845, bw = 0.9258796, ah = 0.5478191, bh = 0.9856029
## 69: log(evidence) = -1.451065, aw = 0.2747695, bw = 0.92577, ah = 0.5412558, bh = 0.9854604
## 70: log(evidence) = -1.450232, aw = 0.2750925, bw = 0.9256658, ah = 0.534884, bh = 0.9853157
## 71: log(evidence) = -1.449424, aw = 0.2754554, bw = 0.9255665, ah = 0.5286835, bh = 0.9851688
## 72: log(evidence) = -1.448692, aw = 0.2758423, bw = 0.9254716, ah = 0.5226112, bh = 0.9850198
## 73: log(evidence) = -1.448076, aw = 0.2762252, bw = 0.9253801, ah = 0.5166092, bh = 0.9848687
## 74: log(evidence) = -1.447556, aw = 0.2765769, bw = 0.9252913, ah = 0.5106332, bh = 0.9847156
## 75: log(evidence) = -1.447057, aw = 0.2768889, bw = 0.9252052, ah = 0.5046794, bh = 0.9845605
## 76: log(evidence) = -1.446523, aw = 0.2771799, bw = 0.9251224, ah = 0.4987779, bh = 0.9844035

```

```

## 77: log(evidence) = -1.445952, aw = 0.2774808, bw = 0.9250434, ah = 0.4929637, bh = 0.9842447
## 78: log(evidence) = -1.445368, aw = 0.2778135, bw = 0.9249682, ah = 0.4872582, bh = 0.9840841
## 79: log(evidence) = -1.444795, aw = 0.2781872, bw = 0.9248967, ah = 0.481665, bh = 0.9839218
## 80: log(evidence) = -1.444259, aw = 0.2785996, bw = 0.9248287, ah = 0.4761712, bh = 0.9837577
## 81: log(evidence) = -1.443774, aw = 0.2790389, bw = 0.9247638, ah = 0.4707554, bh = 0.9835921
## 82: log(evidence) = -1.443327, aw = 0.279494, bw = 0.9247017, ah = 0.4654015, bh = 0.9834248
## 83: log(evidence) = -1.442895, aw = 0.2799635, bw = 0.9246425, ah = 0.4601068, bh = 0.9832559
## 84: log(evidence) = -1.442461, aw = 0.280454, bw = 0.924586, ah = 0.4548759, bh = 0.9830855
## 85: log(evidence) = -1.442025, aw = 0.2809715, bw = 0.9245323, ah = 0.449711, bh = 0.9829136
## 86: log(evidence) = -1.441599, aw = 0.2815148, bw = 0.9244811, ah = 0.4446053, bh = 0.9827401
## 87: log(evidence) = -1.441197, aw = 0.2820768, bw = 0.9244321, ah = 0.4395422, bh = 0.9825652
## 88: log(evidence) = -1.440831, aw = 0.2826519, bw = 0.924385, ah = 0.4344983, bh = 0.9823889
## 89: log(evidence) = -1.440505, aw = 0.2832388, bw = 0.9243393, ah = 0.429445, bh = 0.9822111
## Nsteps =90, log(evidence) =-1.440505, hyper = (0.2838348,0.9242947,0.4243521,0.9820318), dispersion =
##
## Max(evidence) = -1.440505

```

The iteration maximizes (log) evidence (per matrix elements) and terminates when its fractional change becomes smaller than Tol. The option `criterion = connectivity` can also be used. By default, hyperparameters of priors are also updated after `hyper.update.n0` steps. As in maximum likelihood, multiple ranks can be specified:

```
sb <- vb_factorize(sb, ranks = seq(2,7), nrun = 5, verbose = 1, Tol = 1e-4, progress.bar = FALSE)
```

```

## Rank 2
## Max(evidence) = -1.478329
##
## Rank 3
## Max(evidence) = -1.408516
##
## Rank 4
## Max(evidence) = -1.358924
##
## Rank 5
## Max(evidence) = -1.354162
##
## Rank 6
## Max(evidence) = -1.360185
##
## Rank 7
## Max(evidence) = -1.365712

```

With `nrun` larger than 1, multiple inferences will be performed for each rank with different initial conditions and the solution with the highest evidence will be chosen. The object after a `vb_factorize` run will have its `measure` slot filled:

```
measure(sb)
```

```

##   rank  evidence      aw      bw      ah      bh
## 1    2 -1.478329 0.3426491 1.4310219 0.5580671 0.9579385
## 2    3 -1.408516 0.3160585 0.9476028 0.1893169 0.9707558
## 3    4 -1.358924 0.3010692 0.7061936 0.1131652 0.9610053
## 4    5 -1.354162 0.2336142 0.5707557 0.1694986 0.9719610
## 5    6 -1.360185 0.2037749 0.4720986 0.2108323 0.9816994
## 6    7 -1.365712 0.1744965 0.4067136 0.2002811 0.9721059

```

Plotting the object displays the log evidence as a function of rank (Fig. 5):




Figure 5: Dependence of log evidence with rank.

```
plot(sb)
```

6. Visualization

The rank scan above using Bayesian inference correctly identifies $r = 5$ as the optimal rank. The fit results for each rank—from either maximum likelihood or Bayesian inference—are stored in `sb@basis` and `sb@coeff`. Both are lists of matrices of length equal to the number of rank values scanned. One can access them by, e.g.,

```
ranks(sb)
```

```
## [1] 2 3 4 5 6 7
```

```
head(basis(sb)[ranks(sb)==5][[1]]) # basis matrix W for rank 5
```

```
##           1           2           3           4           5
## ISG15    0.0005566516 0.04400669 0.09123646 0.017173204 0.123080868
## EN01     0.0120951859 0.08010730 0.10411271 0.072666639 0.293073976
## EFHD2    0.0005440680 0.12801528 0.08177918 0.005303569 0.000613638
## RPL11    5.1911055900 0.94968087 1.13717449 2.122015576 1.887603772
## SH3BGRL3 0.1199674917 0.42496643 0.75639247 0.190543609 1.237799982
## CD52     0.4720400786 0.11896619 0.04190325 0.625517988 0.883824419
```

Heatmaps of W and H matrices are displayed by `gene_map()` and `cell_map()`, respectively (Figs. 6-7):

```
gene_map(sb, markers = markers, rank = 5, max.per.cluster = 4, gene.name = rowData(sb)[,2],
         cexRow = 0.7)
```

In addition to the marker gene list provided as a parameter, the representative groups of genes for clusters are selected by the “max” scheme⁷: genes are sorted for each cluster with decreasing magnitudes of coefficient matrix elements, and among the top members of the list, those for which the magnitude is the actual maximum over all clusters are chosen. Based on the marker-metagene map in Fig. 6, we rename the clusters 1-5 as follows:




Figure 6: Heatmap of basis matrix elements. Marker genes selected in rows, other than those provided as input, are based on the degree to which each features strongly in a particular cluster only and not in the rest. Columns represent the clusters.

```
cell_type <- c('CD8+_T', 'NK', 'Monocytes', 'B_cell', 'CD4+_T')
colnames(basis(sb)[ranks(sb) == 5][[1]]) <- cell_type
rownames(coeff(sb)[ranks(sb) == 5][[1]]) <- cell_type
```

```
cell_map(sb, rank = 5)
```

In `visualize_clusters()`, each column of H matrix is used to assign cells into clusters, and inter/intra-cluster separations are visualized using tSNE algorithm⁸. It uses the `Rtsne()` function of the `Rtsne` package. A barplot of cluster cell counts are also displayed (Fig. 8):

```
visualize_clusters(sb, rank = 5, cex = 0.7)
```

It is useful to extract hierarchical relationships among the clusters identified. This feature requires a series of inference outcomes for an uninterrupted range of rank values, e.g., from 2 to 7:

```
tree <- build_tree(sb, rmax = 5)
tree <- rename_tips(tree, rank = 5, tip.labels = cell_type)
plot_tree(tree, cex = 0.8, show.node.label = TRUE)
```

The `build_tree` function returns a list containing the tree. The second command above renames the label of terminal nodes by our cell type label. In Fig. 9, the relative distance between clusters can be seen to be consistent with the tSNE plot in Fig. 8.

References

1. Lee, D. D. & Seung, H. S. Learning the parts of objects by non-negative matrix factorization. *Nature* **401**, 788–791 (1999).
2. Hastie, T., Tibshirani, R. & Friedman, J. *The elements of statistical learning: data mining, inference, and prediction*. (2009).




Figure 7: Heatmap of cluster coefficient matrix elements. Rows indicate clusters and columns the cells.

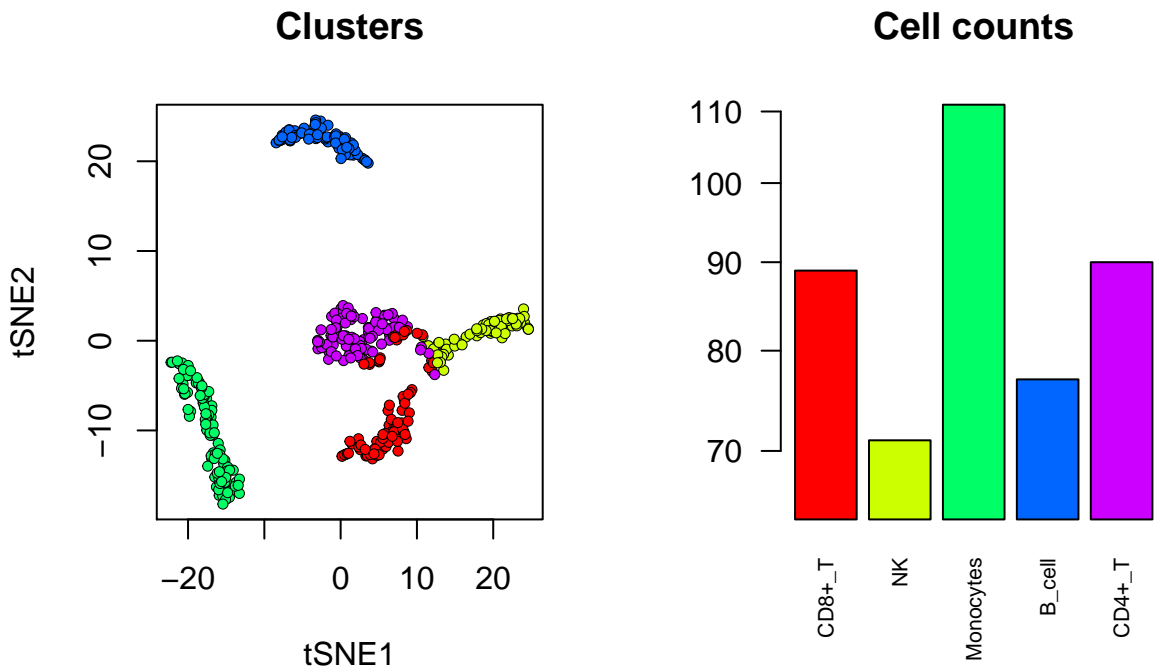


Figure 8: tSNE-based visualization of coefficient matrix elements of cells with colors indicating predicted cluster assignment. The bar plot shows the cell counts of each cluster.

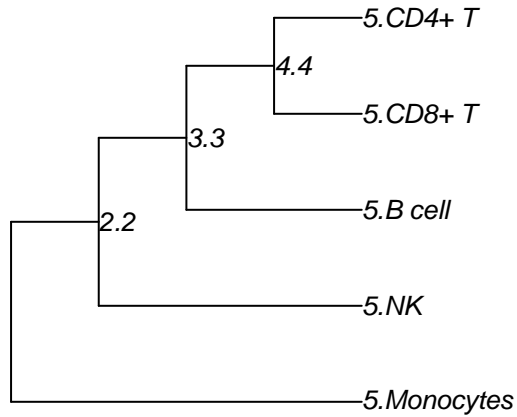


Figure 9: Hierarchical tree of clusters derived from varying ranks. The rank increases from 2 to 5 horizontally and nodes are labeled by cluster IDs which bifurcated in each rank.

3. Cemgil, A. T. Bayesian inference for nonnegative matrix factorisation models. *Comput. Intell. Neurosci.* 785152 (2009).
4. Gaujoux, R. & Seoighe, C. A flexible R package for nonnegative matrix factorization. *BMC Bioinformatics* **11**, 367 (2010).
5. Zhu, X., Ching, T., Pan, X., Weissman, S. M. & Garmire, L. Detecting heterogeneity in single-cell RNA-Seq data by non-negative matrix factorization. *PeerJ* **5**, e2888 (2017).
6. Brunet, J. P., Tamayo, P., Golub, T. R. & Mesirov, J. P. Metagenes and molecular pattern discovery using matrix factorization. *Proc. Natl. Acad. Sci. U.S.A.* **101**, 4164–4169 (2004).
7. Carmona-Saez, P., Pascual-Marqui, R. D., Tirado, F., Carazo, J. M. & Pascual-Montano, A. Biclustering of gene expression data by non-smooth non-negative matrix factorization. *BMC Bioinformatics* **7**, 78 (2006).
8. Maaten, L. van der & Hinton, G. Visualizing high-dimensional data using t-SNE. *J. Mach. Learn. Res.* **9**, 2579–2605 (2008).