

flagme: Fragment-level analysis of GC-MS-based metabolomics data

Mark Robinson

mrobinson@wehi.edu.au

Riccardo Romoli

riccardo.romoli@unifi.it

May 3, 2016

1 Introduction

This document gives a brief introduction to the `flagme` package, which is designed to process, visualise and statistically analyze sets of GC-MS samples. The ideas discussed here were originally designed with GC-MS-based metabolomics in mind, but indeed some of the methods and visualizations could be useful for LC-MS data sets. The *fragment-level analysis* though, takes advantage of the rich fragmentation patterns observed from electron interaction (EI) ionization.

There are many aspects of data processing for GC-MS data. Generally, algorithms are run separately on each sample to detect features, or *peaks* (e.g. AMDIS). Due to retention time shifts from run-to-run, an alignment algorithm is employed to allow the matching of the same feature across multiple samples. Alternatively, if known standards are introduced to the samples, retention *indices* can be computed for each peak and used for alignment. After peaks are matched across all samples, further processing steps are employed to create a matrix of abundances, leading into detecting differences in abundance.

Many of these data processing steps are prone to errors and they often tend to be black boxes. But, with effective exploratory data analysis, many of the pitfalls can be avoided and any problems can be fixed before proceeding to the downstream statistical analysis. The package provides various visualizations to ensure the methods applied are not black boxes.

The `flagme` package gives a complete suite of methods to go through all common stages of data processing. In addition, R is especially well suited to the downstream data analysis tasks since it is very rich in analysis tools and has excellent visualization capabilities. In addition, it is freely available (www.r-project.org), extensible and there is a growing community of users and developers. For routine analyses, graphical user interfaces could be designed.

2 Reading and visualizing GC-MS data

To run these examples, you must have the `gcspikelite` package installed. This data package contains several GC-MS samples from a spike-in experiment we designed to interrogate data processing methods. So, first, we load the packages:

To load the data and corresponding peak detection results, we simply create vectors of the file-names and create a `peakDataset` object. Note that we can speed up the import time by setting the retention time range to a subset of the elution, as below:

```
> gcmsPath <- paste(find.package("gcspikelite"), "data", sep="/")
> data(targets)
> cdfFiles <- paste(gcmsPath, targets$FileName, sep="/")
```

```

> eluFiles <- gsub("CDF", "ELU", cdfFiles)
> pd <- peaksDataset(cdfFiles, mz=seq(50,550), rtrange=c(7.5,8.5))

Reading /home/biocbuild/bbs-3.3-bioc/R/library/gcspikelite/data/0709_468.CDF
Reading /home/biocbuild/bbs-3.3-bioc/R/library/gcspikelite/data/0709_474.CDF
Reading /home/biocbuild/bbs-3.3-bioc/R/library/gcspikelite/data/0709_475.CDF
Reading /home/biocbuild/bbs-3.3-bioc/R/library/gcspikelite/data/0709_485.CDF
Reading /home/biocbuild/bbs-3.3-bioc/R/library/gcspikelite/data/0709_493.CDF
Reading /home/biocbuild/bbs-3.3-bioc/R/library/gcspikelite/data/0709_496.CDF
Reading /home/biocbuild/bbs-3.3-bioc/R/library/gcspikelite/data/0709_470.CDF
Reading /home/biocbuild/bbs-3.3-bioc/R/library/gcspikelite/data/0709_471.CDF
Reading /home/biocbuild/bbs-3.3-bioc/R/library/gcspikelite/data/0709_479.CDF

> pd <- addAMDISPeaks(pd, eluFiles)

Reading retention time range: 7.500133 8.499917
Reading /home/biocbuild/bbs-3.3-bioc/R/library/gcspikelite/data/0709_468.ELU ... Done.
Reading /home/biocbuild/bbs-3.3-bioc/R/library/gcspikelite/data/0709_474.ELU ... Done.
Reading /home/biocbuild/bbs-3.3-bioc/R/library/gcspikelite/data/0709_475.ELU ... Done.
Reading /home/biocbuild/bbs-3.3-bioc/R/library/gcspikelite/data/0709_485.ELU ... Done.
Reading /home/biocbuild/bbs-3.3-bioc/R/library/gcspikelite/data/0709_493.ELU ... Done.
Reading /home/biocbuild/bbs-3.3-bioc/R/library/gcspikelite/data/0709_496.ELU ... Done.
Reading /home/biocbuild/bbs-3.3-bioc/R/library/gcspikelite/data/0709_470.ELU ... Done.
Reading /home/biocbuild/bbs-3.3-bioc/R/library/gcspikelite/data/0709_471.ELU ... Done.
Reading /home/biocbuild/bbs-3.3-bioc/R/library/gcspikelite/data/0709_479.ELU ... Done.

> pd

An object of class "peaksDataset"
9 samples: 0709_468 0709_474 0709_475 0709_485 0709_493 0709_496 0709_470 0709_471 0709_479
501 m/z bins - range: ( 50 550 )
scans: 175 175 175 175 175 174 175 175 175
peaks: 24 23 26 20 27 24 24 25 21

```

Here, we have added peaks from AMDIS, a well known and mature algorithm for deconvolution of GC-MS data. For GC-TOF-MS data, we have implemented a parser for the ChromaTOF output (see the analogous `addChromaTOFPeaks` function). The `addXCMSPeaks` allows to use all the XCMS peak-picking algorithms; using this approach it is also possible to elaborate the raw data file from within R instead of using an external software. In particular the function reads the raw data using XCMS, group each extracted ion according to their retention time using CAMERA and attaches them to an already created `peaksDataset` object:

```

> pd.2 <- peaksDataset(cdfFiles[1:3], mz=seq(50,550), rtrange=c(7.5,8.5))

Reading /home/biocbuild/bbs-3.3-bioc/R/library/gcspikelite/data/0709_468.CDF
Reading /home/biocbuild/bbs-3.3-bioc/R/library/gcspikelite/data/0709_474.CDF
Reading /home/biocbuild/bbs-3.3-bioc/R/library/gcspikelite/data/0709_475.CDF

> pd.2 <- addXCMSPeaks(cdfFiles[1:3], pd.2, peakPicking=c('mF'),
+                       snthresh=3, fwhm=4, step=1, steps=2, mzdiff=0.5)

Start grouping after retention time.
Created 34 pseudospectra.
Generating peak matrix!

```

```
Run isotope peak annotation
% finished: 10 20 50 60 70 80 90 100
Found isotopes: 8

Start grouping after retention time.
Created 32 pseudospectra.
Generating peak matrix!
Run isotope peak annotation
% finished: 10 20 30 50 60 70 80 90 100
Found isotopes: 9

Start grouping after retention time.
Created 38 pseudospectra.
Generating peak matrix!
Run isotope peak annotation
% finished: 10 20 30 50 60 70 80 90 100
Found isotopes: 9
```


```
> pd.2
```

```
An object of class "peaksDataset"
3 samples: 0709_468 0709_474 0709_475
501 m/z bins - range: ( 50 550 )
scans: 175 175 175
peaks: 11 10 9
```

The possibility to work using computer cluster will be added in the future.

Regardless of platform and peak detection algorithm, a useful visualization of a set of samples is the set of total ion currents (TIC), or extracted ion currents (XICs). To view TICs, you can call:

```
> plot(pd, rtrange=c(7.5,8.5), plotPeaks=TRUE, plotPeakLabels=TRUE,
+      max.near=8, how.near=0.5, col=rep(c("blue","red","black"), each=3))
```




Note here the little *hashes* represent the detected peaks and are labelled with an integer index. One of the main challenges is to match these peak detections across several samples, given that they appear at slightly different times in different runs.

For XICs, you need to give the indices (of `pd@mz`, the grid of mass-to-charge values) that you want to plot through the `mzind` argument. This could be a single ion (e.g. `mzind=24`) or could be a range of indices if multiple ions are of interest (e.g. `mzind=c(24,25,98,99)`).

There are several other features within the `plot` command on `peaksDataset` objects that can be useful. See `?plot` (and select the `flagme` version) for full details.

Another useful visualization, at least for individual samples, is a 2D heatmap of intensity. Such plots can be enlightening, especially when peak detection results are overlaid. For example (with detected fragment peaks from AMDIS shown in white):

```
> r <- 1
> plotImage(pd, run=r, rtrange=c(7.5,8.5), main="")
> v <- which(pd@peaksdata[[r]] > 0, arr.ind=TRUE) # find detected peaks
> abline(v=pd@peaksrt[[r]])
> points(pd@peaksrt[[r]][v[,2]], pd@mz[v[,1]], pch=19, cex=.6, col="white")
```



3 Pairwise alignment with dynamic programming algorithm

One of the first challenges of GC-MS data is the matching of detected peaks (i.e. metabolites) across several samples. Although gas chromatography is quite robust, there can be some drift in the elution time of the same analyte from run to run. We have devised a strategy, based on dynamic programming, that takes into account both the similarity in spectrum (at the apex of the called peak) and the similarity in retention time, without requiring the identity of each peak; this matching uses the data alone. If each sample gives a ‘peak list’ of the detected peaks (such as that from AMDIS that we have attached to our `peaksDataset` object), the challenge is to introduce gaps into these lists such that they are best aligned. From this a matrix of retention times or a matrix of peak abundances can be extracted for further statistical analysis, visualization and interpretation. For this matching, we created a procedure analogous to a multiple *sequence* alignment.

To highlight the dynamic programming-based alignment strategy, we first illustrate a pairwise alignment of two peak lists. This example also illustrates the selection of parameters necessary for the alignment. From the data read in previously, let us consider the alignment of two samples, denoted 0709_468 and 0709_474. First, a similarity matrix for two samples is calculated. This is calculated based on a scoring function and takes into account the similarity in retention time and in the similarity of the apex spectra, according to:

$$S_{ij}(D) = \frac{\sum_{k=1}^K x_{ik} y_{jk}}{\sqrt{\sum_{k=1}^K x_{ik}^2 \cdot \sum_{k=1}^K y_{jk}^2}} \cdot \exp\left(-\frac{1}{2} \frac{(t_i - t_j)^2}{D^2}\right)$$

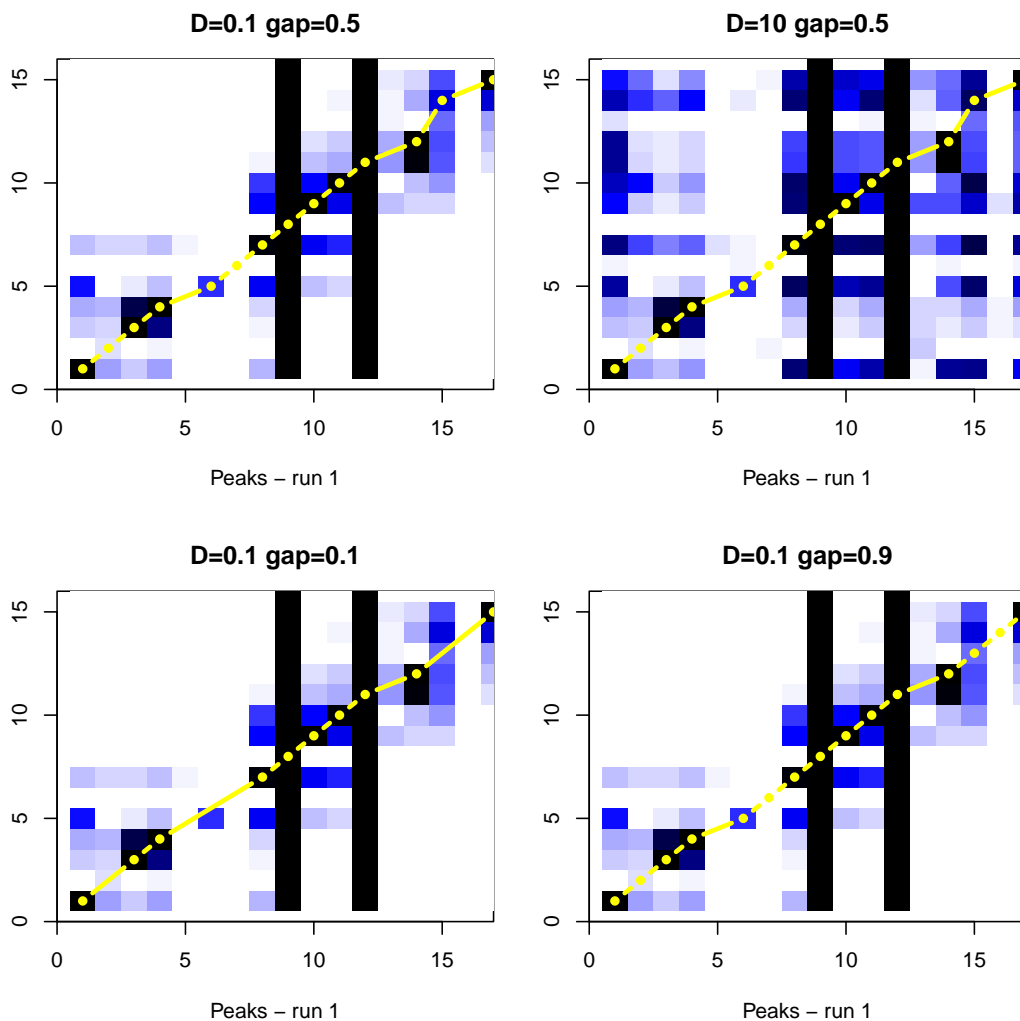
where i is the index of the peak in the first sample and j is the index of the peak in the second sample, \mathbf{x}_i and \mathbf{y}_j are the spectra vectors and t_i and t_j are their respective retention times. As you can see, there are two components to the similarity: spectra similarity (left term) and similarity in retention time (right term). Of course, other metrics

for spectra similarity are feasible. Ask the author if you want to see other metrics implemented. We have some non-optimized code for a few alternative metrics.

The peak alignment algorithm, much like sequence alignments, requires a `gap` parameter to be set, here a number between 0 and 1. A high gap penalty discourages gaps when matching the two lists of peaks and a low gap penalty allows gaps at a very low *cost*. We find that a gap penalty in the middle range (0.4-0.6) works well for GC-MS peak matching. Another parameter, `D`, modulates the impact of the difference in retention time penalty. A large value for `D` essentially eliminates the effect. Generally, we set this parameter to be a bit larger than the average width of a peak, allowing a little flexibility for retention time shifts between samples. Keep in mind the `D` parameter should be set on the scale (i.e. seconds or minutes) of the `peaksrt` slot of the `peaksDataset` object. The next example shows the effect of the `gap` and `D` penalty on the matching of a small ranges of peaks.

```
> Ds <- c(0.1, 10, 0.1, 0.1)
> gaps <- c(0.5, 0.5, 0.1, 0.9)
> par(mfrow=c(2,2), mai=c(0.8466,0.4806,0.4806,0.1486))
> for(i in 1:4){
+   pa <- peaksAlignment(pd@peaksdata[[1]], pd@peaksdata[[2]],
+                       pd@peaksrt[[1]], pd@peaksrt[[2]], D=Ds[i],
+                       gap=gaps[i])
+   plot(pa, xlim=c(0, 17), ylim=c(0, 16), matchCol="yellow",
+        main=paste("D=", Ds[i], " gap=", gaps[i], sep=""))
+ }
```

```
[peaksAlignment] Comparing 24 peaks to 23 peaks -- gap= 0.5 D= 0.1
[peaksAlignment] 21 matched. Similarity= 0.2308905
[peaksAlignment] Comparing 24 peaks to 23 peaks -- gap= 0.5 D= 10
[peaksAlignment] 21 matched. Similarity= 0.2180835
[peaksAlignment] Comparing 24 peaks to 23 peaks -- gap= 0.1 D= 0.1
[peaksAlignment] 15 matched. Similarity= 0.01170268
[peaksAlignment] Comparing 24 peaks to 23 peaks -- gap= 0.9 D= 0.1
[peaksAlignment] 22 matched. Similarity= 0.2785359
```



You might ask: is the `flame` package useful without peak detection results? Possibly. There have been some developments in alignment (generally on LC-MS proteomics experiments) without peak/feature detection, such as Prince et al. 2006, where a very similar dynamic programming is used for a pairwise alignment. We have experimented with alignments without using the peaks, but do not have any convincing results. It does introduce a new set of challenges in terms of highlighting differentially abundant metabolites. However, in the `peaksAlignment` routine above (and those mentioned below), you can set `usePeaks=FALSE` in order to do *scan*-based alignments instead of *peak*-based alignments. In addition, the `flame` package may be useful simply for its bare-bones dynamic programming algorithm.

3.1 Normalizing retention time score to drift estimates

In what is mentioned above for pairwise alignments, we are penalizing for differences in retention times that are non-zero. But, as you can see from the TICs, some differences in retention time are consistent. For example, all of the peaks from sample 0709_485 elute at later times than peaks from sample 0709_496. We should be able to estimate this drift and normalize the time penalty to that estimate, instead of penalizing to zero. That is, we should replace $t_i - t_j$ with $t_i - t_j - \hat{d}_{ij}$ where \hat{d}_{ij} is the expected drift between peak i of the first sample and peak j of the second sample.

More details coming soon.

3.2 Imputing location of undetected peaks

One goal of the alignment leading into downstream data analyses is the generation of a table of abundances for each metabolite across all samples. As you can see from the TICs above, there are some low intensity peaks that fall below detection in some but not all samples. Our view is that instead of inserting arbitrary low constants (such as 0 or half the detection limit) or imputing the intensities post-hoc or having missing data in the data matrices, it is best to return to the area of the where the peak should be and give some kind of abundance. The alignments themselves are rich in information with respect to the location of undetected peaks. We feel this is a more conservative and statistically valid approach than introducing arbitrary values.

More details coming soon.

4 Multiple alignment of several experimental groups

Next, we discuss the multiple alignment of peaks across many samples. With replicates, we typically do an alignment within replicates, then combine these together into a summarized form. This cuts down on the computational cost. For example, consider 2 sets of samples, each with 5 replicates. Aligning first within replicates requires 10+10+1 total alignments whereas an all-pairwise alignment requires 45 pairwise alignments. In addition, this allows some flexibility in setting different gap and distance penalty parameters for the *within* alignment and *between* alignment. An example follows.

```
> print(targets)
```


```
      FileName Group
1 0709_468.CDF  mmA
2 0709_474.CDF  mmA
3 0709_475.CDF  mmA
4 0709_485.CDF  mmC
5 0709_493.CDF  mmC
6 0709_496.CDF  mmC
7 0709_470.CDF  mmD
8 0709_471.CDF  mmD
9 0709_479.CDF  mmD
```

```
> ma <- multipleAlignment(pd, group=targets$Group, wn.gap=0.5, wn.D=.05,
+                          bw.gap=.6, bw.D=0.05, usePeaks=TRUE,
+                          filterMin=2, df=50, verbose=FALSE)
> ma
```

```
An object of class "multipleAlignment"
3 groups: 3 3 3 samples, respectively.
21 merged peaks
```

If you set `verbose=TRUE`, many nitty-gritty details of the alignment procedure are given. Next, we can take the alignment results and overlay it onto the TICs, allowing a visual inspection.

```
> plot(pd, rtrange=c(7.5,8.5), runs=ma@betweenAlignment@runs,
+       mind=ma@betweenAlignment@ind, plotPeaks=TRUE,
+       plotPeakLabels=TRUE, max.near=8, how.near=.5,
+       col=rep(c("blue","red","black"), each=3))
```

5 Correlation Alignment algorithm

Another approach, represented by the `correlationAlignment` function, is to use a modified form of the Pearson correlation algorithm. After the correlation between two samples is calculated, a penalization coefficient, based on the retention time differences, is applied to the result. It is also possible to set a retention time range in which the penalization is 0, this because in gas chromatography we can have a little deviation in the retention time of the metabolite so, based on the experimental data, we can choose the retention time window for the penalization coefficient being applied.

```
> mp <- correlationAlignment(object=pd.2, thr=0.85, D=20, penalty=0.2,
+                             normalize=TRUE, minFilter=1)
> mp
```

```
An object of class correlationAlignment
Aligned Samples:3
Aligned Peaks:9
```

where `thr` represent correlation threshold from 0 (min) to 1 (max); `D` represent the retention time window in seconds; `penalty` represent the penalty inflicted to a match between two peaks when the retention time difference exceed the parameter `D`; `normalize` is about the peak normalization-to-100 before the correlation is calculated; `minFilter` give the opportunity to exclude from the resulting correlation matrix each feature that in represented in our samples less time than this value. The value of `minFilter` must be smaller than the number of samples.

The correlation-based peak alignment for multiple GC-MS peak lists uses a center-star technique to the alignment of the peaks. The combination of the `D` and `penalty` parameters allow the users to force the algorithm to match the peaks close to the reference. The `thr` parameter control the matching factor.

5.1 Gathering results

The alignment results can be extracted from the `multipleAlignment` object as:

```
> ma@betweenAlignment@runs
[1] 5 4 6 2 1 3 7 8 9

> ma@betweenAlignment@ind
      [,1] [,2] [,3] [,4] [,5] [,6] [,7] [,8] [,9]
[1,]    1    1    1    1    1    1    1    1    1
[2,]   NA   NA   NA    2    2    2   NA   NA   NA
[3,]    3    2    2    3    3    3    3    3    2
[4,]    4    3    3    4    4    4    4    4    3
[5,]    6    4    5    5    6    5    6    6    4
[6,]    7    5    6    7    8    6    7    7    5
[7,]    9    7    7    8    9    7   NA   NA   NA
[8,]   10    8    8    9   10    8    9    9    8
[9,]   NA   NA   NA   11   12    9   NA   NA   NA
[10,]  13    9    9   12   14   10   10   12   11
[11,]  15   11   10   13   15   11   11   13   12
[12,]  18   12   13   NA   NA   NA   13   15   13
[13,]   NA   NA   NA   NA   NA   NA   14   16   14
[14,]  20   13   15   15   17   15   15   17   15
[15,]  21   14   16   17   18   16   17   18   16
[16,]   NA   NA   NA   18   19   17   18   19   17
[17,]  23   16   19   19   20   19   19   20   18
[18,]  24   17   21   NA   NA   NA   NA   NA   NA
[19,]  25   18   22   21   22   22   NA   NA   NA
[20,]  26   19   23   22   23   24   23   24   20
[21,]  27   20   24   23   24   26   24   25   21
```

This table would suggest that matched peak 8 (see numbers below the TICs in the figure above) corresponds to detected peaks 9, 12, 11 in runs 4, 5, 6 and so on, same as shown in the above plot.

In addition, you can gather a list of all the merged peaks with the `gatherInfo` function, giving elements for the retention times, the detected fragment ions and their intensities. The example below also shows the how to construct a table of retention times of the matched peaks (No attempt is made here to adjust retention times onto a common scale. Instead, the peaks are matched to each other on their original scale). For example:

```
> outList <- gatherInfo(pd,ma)
> outList[[8]]

$rt
  mmC.5  mmC.4  mmC.6  mmA.2  mmA.1  mmA.3  mmD.7  mmD.8
7.809283 7.825483 7.802867 7.799283 7.800000 7.803050 7.811567 7.816033
  mmD.9
7.822517

$mz
 [1] 56 58 69 73 74 75 77 110 134 147 184 228 229 281

$data
```

```

      mmC.5 mmC.4 mmC.6 mmA.2 mmA.1 mmA.3 mmD.7 mmD.8 mmD.9
[1,] 3840 2731 2544 1697 1263 1437 2669 3028 2563
[2,] 14731 14105 14167 13902 8982 11768 13300 13784 14439
[3,] 6196 6228 5271 5373 3953 4426 5802 6230 6154
[4,] 47480 24064 34816 15176 13949 18072 26624 30088 22712
[5,] 4982 2895 3965 2315 1619 3094 3272 3220 2546
[6,] 26312 22352 23472 20744 16033 20272 23064 21888 22024
[7,] 6372 4619 3575 2375 1615 1910 3883 4707 3119
[8,] 5068 3062 2551 552 411 570 1937 2809 2072
[9,] 5299 2714 2766 176 263 415 2311 3431 2247
[10,] 22064 11447 20160 13162 11084 16504 14499 12684 11305
[11,] 6664 3159 3503 0 255 0 2828 4457 2419
[12,] 16075 6703 7660 0 0 0 6645 10229 5183
[13,] 8357 3080 4097 0 153 0 3123 5018 2492
[14,] 15915 4195 11465 6110 5653 9659 10519 7392 4722

```

```

> rtmat <- matrix(unlist(lapply(outList,.subset,"rt"), use.names=FALSE),
+               nr=length(outList), byrow=TRUE)
> colnames(rtmat) <- names(outList[[1]]$rt); rownames(rtmat) <- 1:nrow(rtmat)
> round(rtmat, 3)

```

```

      mmC.5 mmC.4 mmC.6 mmA.2 mmA.1 mmA.3 mmD.7 mmD.8 mmD.9
1  7.512 7.534 7.506 7.531 7.526 7.540 7.520 7.519 7.531
2    NA    NA    NA 7.559 7.549 7.557    NA    NA    NA
3  7.558 7.580 7.551 7.576 7.566 7.574 7.560 7.565 7.577
4  7.575 7.597 7.569 7.588 7.583 7.592 7.577 7.582 7.594
5  7.615 7.614 7.614 7.616 7.617 7.614 7.617 7.610 7.617
6  7.695 7.717 7.694 7.714 7.709 7.649 7.703 7.702 7.714
7  7.804 7.803 7.711 7.736 7.783 7.712    NA    NA    NA
8  7.809 7.825 7.803 7.799 7.800 7.803 7.812 7.816 7.823
9    NA    NA    NA 7.942 7.880 7.826    NA    NA    NA
10 7.958 7.946 7.951 7.976 7.966 7.975 7.966 7.965 7.977
11 7.986 8.008 7.980 7.999 7.994 7.997 7.989 7.993 8.000
12 8.061 8.077 8.060    NA    NA    NA 8.069 8.068 8.080
13    NA    NA    NA    NA    NA    NA 8.086 8.085 8.091
14 8.204 8.111 8.111 8.114 8.109 8.112 8.109 8.108 8.114
15 8.244 8.254 8.237 8.251 8.246 8.249 8.246 8.245 8.251
16    NA    NA    NA 8.285 8.263 8.283 8.280 8.262 8.263
17 8.324 8.334 8.323 8.337 8.332 8.335 8.326 8.330 8.337
18 8.352 8.363 8.352    NA    NA    NA    NA    NA    NA
19 8.392 8.403 8.386 8.399 8.394 8.403    NA    NA    NA
20 8.432 8.437 8.432 8.434 8.434 8.437 8.435 8.433 8.440
21 8.461 8.477 8.460 8.474 8.469 8.472 8.469 8.468 8.474

```

6 Future improvements and extension

There are many procedures that we have implemented in our investigation of GC-MS data, but have not made part of the package just yet. Some of the most useful procedures will be released, such as:

1. Parsers for other peak detection algorithms (e.g. XCMS, MzMine) and parsers for other alignment procedures (e.g. SpectConnect) and perhaps retention indices procedures.

2. More convenient access to the alignment information and abundance table.
3. Statistical analysis of differential metabolite abundance.
4. Fragment-level analysis, an alternative method to summarize abundance across all detected fragments of a metabolite peak.

7 References

See the following for further details:

1. Robinson MD. *Methods for the analysis of gas chromatography - mass spectrometry data*. **Ph.D. Thesis**. October 2008. Department of Medical Biology (Walter and Eliza Hall Institute of Medical Research), University of Melbourne.
2. Robinson MD, De Souza DP, Keen WW, Saunders EC, McConville MJ, Speed TP, Likić VA. (2007) *A dynamic programming approach for the alignment of signal peaks in multiple gas chromatography-mass spectrometry experiments*. **BMC Bioinformatics**. 8:419.
3. Prince JT, Marcotte EM (2006) *Chromatographic alignment of ESI-LC-MS proteomics data sets by ordered bijective interpolated warping*. **Anal Chem**. 78(17):6140-52.

8 This vignette was built with/at ...

```
> sessionInfo()
```

```
R version 3.3.0 (2016-05-03)
```

```
Platform: x86_64-pc-linux-gnu (64-bit)
```

```
Running under: Ubuntu 14.04.4 LTS
```

```
locale:
```

```
[1] LC_CTYPE=en_US.UTF-8      LC_NUMERIC=C
[3] LC_TIME=en_US.UTF-8      LC_COLLATE=C
[5] LC_MONETARY=en_US.UTF-8  LC_MESSAGES=en_US.UTF-8
[7] LC_PAPER=en_US.UTF-8     LC_NAME=C
[9] LC_ADDRESS=C             LC_TELEPHONE=C
[11] LC_MEASUREMENT=en_US.UTF-8 LC_IDENTIFICATION=C
```

```
attached base packages:
```

```
[1] parallel stats graphics grDevices utils datasets methods
[8] base
```

```
other attached packages:
```

```
[1] flagme_1.28.0      CAMERA_1.28.0      xcms_1.48.0
[4] Biobase_2.32.0     ProtGenerics_1.4.0 BiocGenerics_0.18.0
[7] mzR_2.6.0          Rcpp_0.12.4.5      gcspikelite_1.9.0
```

```
loaded via a namespace (and not attached):
```

```
[1] RColorBrewer_1.1-2  plyr_1.8.3         bitops_1.0-6
[4] tools_3.3.0         rpart_4.1-10      gtable_0.2.0
[7] lattice_0.20-33     Matrix_1.2-6      graph_1.50.0
[10] igraph_1.0.1        SparseM_1.7        gridExtra_2.2.1
```

```
[13] cluster_2.0.4      gtools_3.5.0      caTools_1.17.1
[16] stats4_3.3.0      grid_3.3.0        nnet_7.3-12
[19] data.table_1.9.6  RBGL_1.48.0       survival_2.39-2
[22] foreign_0.8-66    gdata_2.17.0     latticeExtra_0.6-28
[25] Formula_1.2-1    ggplot2_2.1.0     magrittr_1.5
[28] Hmisc_3.17-4     scales_0.4.0      gplots_3.0.1
[31] codetools_0.2-14 splines_3.3.0     MASS_7.3-45
[34] colorspace_1.2-6 KernSmooth_2.23-15 acepack_1.3-3.3
[37] munsell_0.4.3    chron_2.3-47
```

```
> date()
```

```
[1] "Tue May 3 20:11:11 2016"
```