flowCore: data structures package for flow cytometry data

N. Le Meur F. Hahne B. Ellis P. Haaland

June 9, 2020

Abstract

Background The recent application of modern automation technologies to staining and collecting flow cytometry (FCM) samples has led to many new challenges in data management and analysis. We limit our attention here to the associated problems in the analysis of the massive amounts of FCM data now being collected. From our viewpoint, see two related but substantially different problems arising. On the one hand, there is the problem of adapting existing software to apply standard methods to the increased volume of data. The second problem, which we intend to address here, is the absence of any research platform which bioinformaticians, computer scientists, and statisticians can use to develop novel methods that address both the volume and multidimensionality of the mounting tide of data. In our opinion, such a platform should be Open Source, be focused on visualization, support rapid prototyping, have a large existing base of users, and have demonstrated suitability for development of new methods. We believe that the Open Source statistical software R in conjunction with the Bioconductor Project fills all of these requirements. Consequently we have developed a Bioconductor package that we call flowCore. The flowCore package is not intended to be a complete analysis package for FCM data. Rather, we see it as providing a clear object model and a collection of standard tools that enable R as an informatics research platform for flow cytometry. One of the important issues that we have addressed in the flowCore package is that of using a standardized representation that will insure compatibility with existing technologies for data analysis and will support collaboration and interoperability of new methods as they are developed. In order to do this, we have followed the current standardized descriptions of FCM data analysis as being developed under NIH Grant xxxx [n]. We believe that researchers will find flowCore to be a solid foundation for future development of new methods to attack the many interesting open research questions in FCM data analysis.

Methods We propose a variety of data structures. We have implemented the classes and methods in the Bioconductor package flowCore. We illustrate their use with X case studies.

Results We hope that those proposed data structures will be the base for the development of many tools for the analysis of high throughput flow cytometry.

keywords Flow cytometry, high throughput, software, standard

1 Introduction

Traditionally, flow cytometry has been a tube-based technique limited to small-scale laboratory and clinical studies. High throughput methods for flow cytometry have recently been developed for drug discovery and advanced research methods (Gasparetto et al., 2004). As an example, the flow cytometry high content screening (FC-HCS) can process up to a thousand samples daily at a single workstation,

and the results have been equivalent or superior to traditional manual multi-parameter staining and analysis techniques.

The amount of information generated by high throughput technologies such as FC-HCS need to be transformed into executive summaries (which are brief enough) for creative studies by a human researcher (Brazma, 2001). Standardization is critical when developing new high throughput technologies and their associated information services (Brazma, 2001; Chicurel, 2002; Boguski and McIntosh, 2003). Standardization efforts have been made in clinical cell analysis by flow cytometry (Keeney et al., 2004), however data interpretation has not been standardized for even low throughput FCM. It is one of the most difficult and time consuming aspects of the entire analytical process as well as a primary source of variation in clinical tests, and investigators have traditionally relied on intuition rather than standardized statistical inference (Bagwell, 2004; Braylan, 2004; Parks, 1997; Suni et al., 2003). In the development of standards in high throughput FCM, little progress has been made in terms of Open Source software. In this article we propose R data structures to handle flow cytometry data through the main steps of preprocessing: compensation, transformation, filtering.

The aim is to merge both **prada** and **rflowcyt** (LeMeur and Hahne, 2006) into one core package which is compliant with the data exchange standards that are currently developed in the community (Spidlen et al., 2006).

Visualization as well as quality control will then be part of the utility packages that depend on the data structures defined in the flowCore package.

2 Representing Flow Cytometry Data

flowCore's primary task is the representation and basic manipulation of flow cytometry (or similar) data. This is accomplished through a data model very similar to that adopted by other Bioconductor packages using the `expressionSet` and `AnnotatedDataFrame` structures familiar to most Bioconductor users.

2.1 The *flowFrame* Class

The basic unit of manipulation in flowCore is the *flowFrame*, which corresponds roughly to a single "FCS" file exported from the flow cytometer's acquisition software. At the moment we support FCS file versions 2.0 through 3.1, and we expect to support FCS4/ACS1 as soon as the specification has been ratified.

2.1.1 Data elements

The primary elements of the *flowFrame* are the `exprs` and `parameters` slots, which contain the event-level information and column metadata respectively. The event information, stored as a single matrix, is accessed and manipulated via the `exprs()` and `exprs<-` methods, allowing *flowFrame*s to be stitched together if necessary (for example, if the same tube has been collected in two acquisition files for memory reasons).

The `parameters` slot is an *AnnotatedDataFrame* that contains information derived from an FCS file's "$Pn" keywords, which describe the detector and stain information. The entire list is available via the `parameter()` method, but more commonly this information is accessed through the `names`, `featureNames` and `colnames` methods. The `names` function returns a concatenated version of

`colnames` and `featureNames` using a format similar to the one employed by most flow cytometry analysis software. The `colnames` method returns the detector names, often named for the fluorochrome detected, while the `featureNames` method returns the description field of the parameters, which will typically be an identifier for the antibody.

The `keyword` method allows access to the raw FCS keywords, which are a mix of standard entries such as "SAMPLE ID," vendor specific keywords and user-defined keywords that add more information about an experiment. In the case of plate-based experiments, there are also one or more keywords that identify the specific well on the plate.

Most vendor software also includes some sort of unique identifier for the file itself. The specialized method `identifier` attempts to locate an appropriate globally unique identifier that can be used to uniquely identify a frame. Failing that, this method will return the original file name offering some assurance that this frame is at least unique to a particular session.

2.2 The *flowSet* Class

Most experiments consist of several *flowFrame* objects, which are organized using a *flowSet* object. This class provides a mechanism for efficiently hosting the *flowFrame*s with minimal copying. This reduces memory requirements and ensures that experimental metadata stay properly attached to their corresponding *flowFrame*s.

3 Reading and Manipulating **flowCore** Data Classes

3.1 Reading an FCS file into a *flowFrame*

FCS files are read into the R environment via the `read.FCS` function using the standard connection interface—allowing for the possibility of accessing FCS files hosted on a remote resource as well as those that have been compressed or even retrieved as a blob from a database interface. FCS files (version 2.0, 3.0, and 3.1) and LMD (List Mode Data) extensions are currently supported.

There are also several immediate processing options available in this function, the most important of which is the `transformation` parameter, which can "linearize" (the default), "linearize-with-PnG-scaling", or "scale" our data. To see how this works, first we will examine an FCS file without any transformation at all:

```
file.name <- system.file("extdata","0877408774.B08",
                          package="flowCore")
x <- read.FCS(file.name, transformation=FALSE)
summary(x)

##               FSC-H      SSC-H      FL1-H      FL2-H      FL3-H      FL1-A      FL4-H
## Min.        85.0000    11.0000     0.0000     0.0000     0.0000     0.0000     0.0000
## 1st Qu.    385.0000   141.0000   233.0000   277.0000    90.0000     0.0000   210.0000
## Median     441.0000   189.0000   545.5000   346.0000   193.0000    26.0000   279.0000
## Mean       491.9644   277.9105   439.1023   366.1567   179.7122    34.0766   323.5306
## 3rd Qu.    518.0000   270.0000   610.0000   437.0000   264.0000    51.0000   390.0000
## Max.      1023.0000  1023.0000   912.0000  1023.0000   900.0000  1023.0000  1022.0000
```

```
##            Time
## Min.       1.00
## 1st Qu.  122.00
## Median   288.00
## Mean     294.77
## 3rd Qu.  457.50
## Max.     626.00
```

As we can see, in this case the values from each parameter seem to run from 0 to 1023 ($2^{10} - 1$). However, inspection of the documentation of the "exponentiation" keyword ($PnE) reveals that some of the parameters (3 and 4) have been stored in a format specifying that the channel values should be scaled as $y = f_2 \times 10^{f_1 \cdot x / R}$ where x is the original channel value, y is the scaled value, f_1 and f_2 are given respectively by the first and second element of the value of the $PnE key and R is the channel range given by the value of the $PnR key. The special $PnE value of "$0, 0$" corresponds to a scale that is already linear.

```
keyword(x, c("$P1E", "$P2E", "$P3E", "$P4E"))

## $`$P1E`
## [1] "0,0"
##
## $`$P2E`
## [1] "0,0"
##
## $`$P3E`
## [1] "4,1"
##
## $`$P4E`
## [1] "4,1"
```

The default "linearize" transformation option will convert these to, effectively, have a $PnE value of "$0, 0$":

```
summary(read.FCS(file.name))

##              FSC-H      SSC-H        FL1-H        FL2-H         FL3-H        FL1-A
## Min.        85.0000    11.0000     1.00000      1.00000      1.000000      0.0000
## 1st Qu.    385.0000   141.0000     8.13123     12.07901      2.246790      0.0000
## Median     441.0000   189.0000   135.16485     22.46790      5.674221     26.0000
## Mean       491.9644   277.9105   157.79417    105.98637      8.464880     34.0766
## 3rd Qu.    518.0000   270.0000   241.44182     50.93675     10.746078     51.0000
## Max.      1023.0000  1023.0000  3651.74127   9910.45856   3278.121151   1023.0000
##              FL4-H    Time
## Min.        1.00000    1.00
## 1st Qu.     6.61169  122.00
```

```
## Median    12.29826 288.00
## Mean     140.39784 294.77
## 3rd Qu.   33.37625 457.50
## Max.    9821.71889 626.00
```

The "linearize-with-PnG-scaling" option will perform the previous transformation and it will also apply a "division by gain" to parameters stored on linear scale with specified gain. The gain is specified in the $PnG keyword. This option has been introduced as part of Gating-ML 2.0 compliance.

Finally, the "scale" option will both linearize values as well as ensure that output values are contained in $[0, 1]$, which is the proposed method of data storage for the ACS1.0/FCS4.0 specification:

```
summary(read.FCS(file.name,transformation="scale"))

##              FSC-H        SSC-H        FL1-H        FL2-H        FL3-H        FL1-A
## Min.    0.08308895 0.01075269 0.0000000000 0.000000000 0.0000000000 0.00000000
## 1st Qu. 0.37634409 0.13782991 0.0007131943 0.001108012 0.0001246915 0.00000000
## Median  0.43108504 0.18475073 0.0134178268 0.002147005 0.0004674689 0.02541544
## Mean    0.48090362 0.27166227 0.0156809848 0.010499687 0.0007465626 0.03331046
## 3rd Qu. 0.50635386 0.26392962 0.0240465869 0.004994175 0.0009747053 0.04985337
## Max.    1.00000000 1.00000000 0.3651106383 0.991044961 0.3277448896 1.00000000
##              FL4-H         Time
## Min.    0.0000000000 0.0009775171
## 1st Qu. 0.0005612251 0.1192570870
## Median  0.0011299392 0.2815249267
## Mean    0.0139411784 0.2881427175
## 3rd Qu. 0.0032379485 0.4472140762
## Max.    0.9821701062 0.6119257087
```

Another parameter of interest is the `alter.names` parameter, which will convert the parameter names into more "R friendly" equivalents, usually by replacing "-" with ".":

```
read.FCS(file.name,alter.names=TRUE)

## flowFrame object '0877408774.B08'
## with 10000 cells and 8 observables:
##      name              desc range minRange maxRange
## $P1 FSC.H             FSC-H  1024 0.000000     1023
## $P2 SSC.H             SSC-H  1024 0.000000     1023
## $P3 FL1.H              <NA>  1024 1.009044    10000
## $P4 FL2.H              <NA>  1024 1.009044    10000
## $P5 FL3.H              <NA>  1024 1.009044    10000
## $P6 FL1.A              <NA>  1024 0.000000     1023
## $P7 FL4.H              <NA>  1024 1.009044    10000
## $P8  Time Time (51.20 sec.)  1024 0.000000     1023
## 164 keywords are stored in the 'description' slot
```

When only a particular subset of parameters is desired, the `column.pattern` parameter allows for the specification of a regular expression and only parameters that match the regular expression will be included in the frame. For example, to include on the Height parameters:

```
x <- read.FCS(file.name, column.pattern="-H")
x

## flowFrame object '0877408774.B08'
## with 10000 cells and 6 observables:
##      name  desc range minRange maxRange
## $P1 FSC-H FSC-H  1024 0.000000     1023
## $P2 SSC-H SSC-H  1024 0.000000     1023
## $P3 FL1-H  <NA>  1024 1.009044    10000
## $P4 FL2-H  <NA>  1024 1.009044    10000
## $P5 FL3-H  <NA>  1024 1.009044    10000
## $P7 FL4-H  <NA>  1024 1.009044    10000
## 150 keywords are stored in the 'description' slot
```

Note that `column.pattern` *is applied after* `alter.names` *if it is used.*

Finally, only a sample of lines can be read in case you need a quick overview of a large series of files.

```
lines <- sample(100:500, 50)
y <- read.FCS(file.name, which.lines = lines)
y

## flowFrame object '0877408774.B08'
## with 50 cells and 8 observables:
##      name                   desc range minRange maxRange
## $P1 FSC-H                  FSC-H  1024 0.000000     1023
## $P2 SSC-H                  SSC-H  1024 0.000000     1023
## $P3 FL1-H                   <NA>  1024 1.009044    10000
## $P4 FL2-H                   <NA>  1024 1.009044    10000
## $P5 FL3-H                   <NA>  1024 1.009044    10000
## $P6 FL1-A                   <NA>  1024 0.000000     1023
## $P7 FL4-H                   <NA>  1024 1.009044    10000
## $P8  Time Time (51.20 sec.)  1024 0.000000     1023
## 164 keywords are stored in the 'description' slot
```

3.2 Creating a *flowSet*

To facilitate the creation of *flowSet* objects from a variety of sources, we provide a means to coerce *list* and *environment* objects to a *flowSet* object using the usual coercion mechanisms. For example, if we have a directory containing FCS files, we can read in a list of those files and create a *flowSet* out of them:

6

```
fcs.dir <- system.file("extdata",
                        "compdata",
                        "data",
                        package="flowCore")
frames <- lapply(dir(fcs.dir, full.names=TRUE), read.FCS)
fs <- as(frames, "flowSet")
fs

## A flowSet with 5 experiments.
##
##    column names:
##    FSC-H SSC-H FL1-H FL2-H FL3-H FL1-A FL4-H

sampleNames(fs)

## [1] "V1" "V2" "V3" "V4" "V5"
```

Note that the original list is unnamed and that the resulting sample names are not particularly meaningful. If the list is named, the list constructed is much more meaningful. One such approach is to employ the `keyword` method for *flowFrame* objects to extract the "SAMPLE ID" keyword from each frame:

```
names(frames) <- sapply(frames, keyword, "SAMPLE ID")
fs <- as(frames, "flowSet")
fs

## A flowSet with 5 experiments.
##
##    column names:
##    FSC-H SSC-H FL1-H FL2-H FL3-H FL1-A FL4-H

sampleNames(fs)

## [1] "NA"   "fitc" "pe"   "apc"  "7AAD"
```

3.2.1 Working with experimental metadata

Like most Bioconductor organizational classes, the *flowSet* has an associated *AnnotatedDataFrame* that provides metadata not contained within the *flowFrame* objects themselves. This data frame is accessed and modified via the usual `phenoData` and `phenoData<-` methods. You can also generally treat the phenotypic data as a normal data frame to add new descriptive columns. For example, we might want to track the original filename of the frames from above in the phenotypic data for easier access:

```
phenoData(fs)$Filename <- fsApply(fs,keyword, "$FIL")
pData(phenoData(fs))

##        name   Filename
## NA       NA 060909.001
## fitc   fitc 060909.002
## pe       pe 060909.003
## apc     apc 060909.004
## 7AAD   7AAD 060909.005
```

Note that we have used the *flowSet*-specific iterator, `fsApply`, which acts much like `sapply` or `lapply` and will be discussed shortly. Additionally, we should also note that the `phenoData` data frame **must** have row names that correspond to the original names used to create the *flowSet*.

3.2.2 Reading multiple FCS files with `read.FlowSet`

The process of reading multiple FCS files in to a *flowSet* is simplified by using the `read.flowSet` function. In its simplest incarnation, this function takes a `path`, which defaults to the current working directory, and an optional `pattern` argument that allows only a subset of files contained within the working directory to be selected. For example, to read a *flowSet* of the files read in as `frames` above:

```
fs <- read.flowSet(path = fcs.dir)
```

`read.flowSet` will pass on additional arguments meant for the underlying `read.FCS` function, such as `alter.names` and `column.pattern`, but also supports several other interesting arguments for conducting initial processing:

files An alternative to the `pattern` argument, you may also supply a vector of filenames to read.

name.keyword Like the example in the previous section, you may specify a particular keyword to use in place of the filename when creating the *flowSet*.

phenoData If this is an *AnnotatedDataFrame*, then this will be used in place of the data frame that is ordinarily created. Additionally, the row names of this object will be taken to be the filenames of the FCS files in the directory specified by `path`.

3.2.3 Manipulating a *flowSet*

You can extract a *flowFrame* from a *flowSet* object in the usual way using the `[[` or `$` extraction operators. On the other hand using the `[` extraction operator returns a new *flowSet* by **copying** the environment. However, simply assigning the *flowFrame* to a new variable will **not** copy the contained frames.

The primary iterator method for a *flowSet* is the `fsApply` method, which works more-or-less like `sapply` or `lapply` with two extra options. The first argument, `simplify`, which defaults to `TRUE`, instructs `fsApply` to attempt to simplify its results much in the same way as `sapply`. The primary difference is that if all of the return values of the iterator are *flowFrame* objects, `fsApply` will create a new *flowSet* object to hold them. The second argument, `use.exprs`, which defaults to `FALSE`

instructs `fsApply` to pass the expression matrix of each frame rather than the *flowFrame* object itself. This allows functions to operate directly on the intensity information without first having to extract it.

As an aid to this sort of operation we also introduce the `each_row` and `each_col` convenience functions that take the place of `apply` in the `fsApply` call. For example, if we wanted the median value of each parameter of each *flowFrame* we might write:

```
fsApply(fs, each_col, median)

##              FSC-H SSC-H      FL1-H      FL2-H      FL3-H FL1-A      FL4-H
## 060909.001   423   128   4.104698   4.531584   3.651741     0   7.233942
## 060909.002   436   128 930.572041 228.757320  33.376247   217   8.278826
## 060909.003   438   120  10.181517 791.475544 114.444190     0   9.305720
## 060909.004   441   129   4.371445   4.869675   4.782858     0 358.663762
## 060909.005   429   133   5.002865  14.989296  63.209339     0  20.908000
```

which is equivalent to the less readable

```
fsApply(fs,function(x) apply(x, 2, median), use.exprs=TRUE)

##              FSC-H SSC-H      FL1-H      FL2-H      FL3-H FL1-A      FL4-H
## 060909.001   423   128   4.104698   4.531584   3.651741     0   7.233942
## 060909.002   436   128 930.572041 228.757320  33.376247   217   8.278826
## 060909.003   438   120  10.181517 791.475544 114.444190     0   9.305720
## 060909.004   441   129   4.371445   4.869675   4.782858     0 358.663762
## 060909.005   429   133   5.002865  14.989296  63.209339     0  20.908000
```

In this case, the `use.exprs` argument is not required in the first case because `each_col` and `each_row` are methods and have been defined to work on *flowFrame* objects by first extracting the intensity data.

4   Visualizing Flow Cytometry Data

Much of the more sophisticated visualization of *flowFrame* and *flowSet* objects, including an interface to the ggplot2 graphics system is implemented by the ggcyto package, also included as part of Bioconductor. Here, we will only introduce the `autoplot` function for the purpose of demonstrating flowCore operations. See the vignettes of ggcyto for more examples of how to visualize flow data.

4.1   Visualizing a single *flowFrame*

It is frequently helpful to view the density of events in a *flowFrame* across one channel or two.

To create a bivariate density plot, simply provide the appropriate channels:

```
library(ggcyto)

## Loading required package:  ggplot2
## Loading required package:  ncdfFlow
```

```
## Loading required package:  RcppArmadillo
## Loading required package:  BH
## Loading required package:  flowWorkspace
## As part of improvements to flowWorkspace, some behavior of
## GatingSet objects has changed.  For details, please read the section
## titled "The cytoframe and cytoset classes" in the package vignette:
##
##  vignette("flowWorkspace-Introduction", "flowWorkspace")
```

```r
autoplot(x, "FL1-H", "FL2-H")
```


Similarly, to get a univariate densityplot:

```r
autoplot(x, "FL1-H")
```

4.2   Visualizing a *flowSet*

The syntax is basically the same for *flowSet* objects, with the output now being a grid of plots corresponding to each *flowFrame* (which may or may not be useful depending on context). The grid layout can be adjusted with additional arguments to `autoplot`

```r
fs <- read.flowSet(path = system.file("extdata",
                                      package = "flowCore"),
                   pattern = "\\.")
autoplot(fs, "FL1-H", "FL2-H")
```


11

5 Compensation

Before proceeding with further analysis of the data in a *flowFrame* or *flowSet*, it is important to properly compensate the data for spectral overlap between fluorescence channels.

5.1 Extracting and applying a pre-calculated spillover matrix

If the original FCS file contains a pre-calculated spillover matrix as the value of the $SPILLOVER or $SPILL keywords, this can be accessed from the *flowFrame*'s `description` slot using the `spillover` method:

```
fcsfiles <- list.files(pattern = "CytoTrol",
                       system.file("extdata",
                                   package = "flowWorkspaceData"),
                       full = TRUE)
fs <- read.flowSet(fcsfiles)
x <- fs[[1]]
comp_list <- spillover(x)
comp_list

## $SPILL
##              B710-A        R660-A        R780-A        V450-A        V545-A
## [1,] 1.000000e+00 3.143890e-02 0.1909655363 3.057568e-03 0.002047231
## [2,] 5.537983e-03 1.000000e+00 0.1768123886 0.000000e+00 0.000000000
## [3,] 9.958625e-05 9.847661e-03 1.0000000000 0.000000e+00 0.000000000
## [4,] 0.000000e+00 8.909845e-05 0.0000000000 1.000000e+00 0.451194675
## [5,] 2.477092e-03 5.235156e-04 0.0000000000 3.796154e-02 1.000000000
## [6,] 1.172236e-01 1.642721e-03 0.0003321532 0.000000e+00 0.000000000
## [7,] 1.420516e-02 4.568956e-04 0.1754022374 8.902497e-05 0.000000000
##              G560-A        G780-A
## [1,] 3.442413e-04 0.071933810
## [2,] 0.000000e+00 0.006618897
## [3,] 0.000000e+00 0.035399709
## [4,] 1.082746e-04 0.000000000
## [5,] 6.361807e-05 0.000000000
## [6,] 1.000000e+00 0.009219359
## [7,] 4.096870e-02 1.000000000
##
## $spillover
## NULL
##
## $`$SPILLOVER`
## NULL

comp  <- comp_list[[1]]
```

Given a *flowFrame*, `spillover` will check a few valid spillover matrix keywords. This will result in a list, usually only one of which will be valid (in this case, the first one, corresponding to the $SPILL keyword).

This spillover matrix can then be applied to the channels of the *flowFrame* using the `compensate` function.

```
x_comp <- compensate(x, comp)
```

The `compensate` function can also take a *flowSet* as its first argument and a list of named spillover matrices to compensate the corresponding *flowFrame*s. Here the `simplify = FALSE` option is necessary so that `fsApply` returns a list of named matrices.

```
comp <- fsApply(fs, function(x) spillover(x)[[1]], simplify=FALSE)
fs_comp <- compensate(fs, comp)
```

While transformation will be discussed in the next section, it is helpful to visualize the effect of compensation on transformed parameters. Quick inspection of the compensation matrix above reveals large overlap between the 450nm and 545nm channels, so we would expect to see a significant difference in each of those channels after compensation. This is in fact the case:

```
library(gridExtra)
transList <- estimateLogicle(x, c("V450-A","V545-A"))
p1 <- autoplot(transform(x, transList), "V450-A", "V545-A") +
      ggtitle("Before")
p2 <- autoplot(transform(x_comp, transList), "V450-A", "V545-A") +
      ggtitle("Before")
grid.arrange(as.ggplot(p1), as.ggplot(p2), ncol = 2)
```


Often the spillover matrix is not provided in the FCS file for each analyzed sample and instead it is necessary to work with a set of FCS files representing the controls from which the spillover matrix is to be calculated, as follows.

5.2 Computing a spillover matrix from a set of compensation controls

The spillover matrix can also be calculated using a set of FCS files that contain data for an unstained sample and singly-stained samples for each of the measurement channels. If these files are not already incorporated as *flowFrame*s in a *flowSet* object, this must be done first. Here we also are assigning the appropriate channel name for each control file before coercion of the frames to a *flowSet*. This spillover calculation method is now part of the flowStats package.

```
library(flowStats)
fcs.dir <- system.file("extdata", "compdata", "data",
                       package="flowCore")
frames <- lapply(dir(fcs.dir, full.names=TRUE), read.FCS)
names(frames) <- c("UNSTAINED", "FL1-H", "FL2-H", "FL4-H", "FL3-H")
frames <- as(frames, "flowSet")
comp <- spillover(frames, unstained="UNSTAINED", patt = "-H",
                  fsc = "FSC-H", ssc = "SSC-H",
                  stain_match = "ordered")
comp
```

```
##               FL1-H        FL2-H       FL3-H        FL4-H
## FL1-H 1.0000000000 0.2420222776 0.032083706 0.001127816
## FL2-H 0.0077220477 1.0000000000 0.140788232 0.002632689
## FL3-H 0.0007590319 0.0009620459 0.003218614 1.000000000
## FL4-H 0.0150806322 0.1755899032 1.000000000 0.229593860
```

Note that when using the stain_match = "ordered" option of spillover, we are specifying that the method should assume the singly-stained *flowFrame*s in the *flowSet* are in the order of the channel columns. In this case this behavior is incorrect as can be seen by quick inspection of the resulting matrix. To override this behavior and use the names we have provided, use the stain_match = "regexpr" option instead. This will use regular expression matching of the channel names to the *flowFrame* names to achieve the proper ordering.

```
sampleNames(frames)
```

```
## [1] "UNSTAINED" "FL1-H"     "FL2-H"     "FL4-H"     "FL3-H"
```

```
comp <- spillover(frames, unstained="UNSTAINED", patt = "-H",
                  fsc = "FSC-H", ssc = "SSC-H",
                  stain_match = "regexpr")
comp
```

```
##               FL1-H        FL2-H       FL3-H        FL4-H
## FL1-H 1.0000000000 0.2420222776 0.032083706 0.001127816
## FL2-H 0.0077220477 1.0000000000 0.140788232 0.002632689
## FL3-H 0.0150806322 0.1755899032 1.000000000 0.229593860
## FL4-H 0.0007590319 0.0009620459 0.003218614 1.000000000
```

As an alternative way to match channels to their corresponding compensation controls, we provide the `spillover_match` method. This allows for the matching of control files to specific channels via a simple csv file, resulting in a *flowSet* ready for `spillover`. Lastly, it is acceptable to omit the *flowSet* argument x to both `spillover` and `spillover_match` if you provide the `path` argument to point to a directory containing the files to be used in the compensation *flowSet*:

```r
comp_match <- system.file("extdata", "compdata", "comp_match",
                          package="flowCore")
# The matchfile has a simple format
writeLines(readLines(comp_match))

## filename,channel
## 060909.001,unstained
## 060909.002,FL1-H
## 060909.004,FL4-H
## 060909.005,FL3-H
## 060909.003,FL2-H

control_path <- system.file("extdata", "compdata", "data",
                            package="flowCore")
# Using path rather than pre-constructed flowSet
matched_fs <- spillover_match(path=control_path,
                              fsc = "FSC-H", ssc = "SSC-H",
                              matchfile = comp_match)
comp <- spillover(matched_fs, fsc = "FSC-H", ssc = "SSC-H",
                  prematched = TRUE)
```

This spillover matrix can then be used to compensate a sample *flowFrame* or an entire *flowSet*.

6   Transformation

flowCore features two methods of transforming parameters within a *flowFrame* or *flowSet*: inline and out-of-line. The inline method, discussed in the next section, has been developed primarily to support filtering features and is strictly more limited than the out-of-line transformation method, which uses R's `transform` function to accomplish the filtering. Like the normal `transform` function, the *flowFrame* is considered to be a data frame with columns named for parameters of the FCS file. For example, if we wished to plot our first *flowFrame*'s first two fluorescence parameters on the log scale we might write:

```r
fs <- read.flowSet(path=system.file("extdata", "compdata", "data",
                   package="flowCore"), name.keyword="SAMPLE ID")
autoplot(transform(fs[[1]],
                   `FL1-H`=log(`FL1-H`),
                   `FL2-H`=log(`FL2-H`)
                   ),
```

```
"FL1-H","FL2-H")
```


060909.001

Like the usual `transform` function, we can also create new parameters based on the old parameters, without destroying the old

```
autoplot(transform(fs[[1]],
                   log.FL1.H=log(`FL1-H`),
                   log.FL2.H=log(`FL2-H`)
                   ),
        "log.FL1.H", "log.FL2.H")
```


060909.001

6.1 Standard Transforms

Though any function can be used as a transform in both the out-of-line and inline transformation techniques, flowCore provides a number of parameterized transform generators that correspond to the transforms commonly found in flow cytometry and defined in the Transformation Markup Language (Transformation-ML, see `http://www.ficcs.org/` and Spidlen et al. (2006) for more details). Briefly, the predefined transforms are:

truncateTransform $y = \begin{cases} a & x < a \\ x & x \geq a \end{cases}$

scaleTransform $f(x) = \frac{x-a}{b-a}$

linearTransform $f(x) = a + bx$

quadraticTransform $f(x) = ax^2 + bx + c$

lnTransform $f(x) = \log(x)\frac{r}{d}$

logTransform $f(x) = \log_b(x)\frac{r}{d}$

biexponentialTransform $f^{-1}(x) = ae^{bx} - ce^{dx} + f$

logicleTransform A special form of the biexponential transform with parameters selected by the data.

arcsinhTransform $f(x) = asinh(a + bx) + c$

To use a standard transform, first we create a transform function via the constructors supplied by flowCore:

```
aTrans <- truncateTransform("truncate at 1", a=1)
aTrans

## transform object 'truncate at 1'
```

which we can then use on the parameter of interest in the usual way, in this case for the full *flowSet* using the same syntax

```
transform(fs,`FL1-H`=aTrans(`FL1-H`))

## A flowSet with 5 experiments.
##
##   column names:
##   FSC-H SSC-H FL1-H FL2-H FL3-H FL1-A FL4-H
```

However this form of transform call is not intended to be used in the programmatic context because a locally defined transform function (e.g. 'aTrans') may not always be visible to the non-standard evaluation environment. .e.g

```
f1 <- function(fs,...){
  transform(fs, ...)[,'FL1-H']
}

f2 <- function(fs){
  aTrans <- truncateTransform("truncate at 1", a=1)
  f1(fs, `FL1-H` = aTrans(`FL1-H`))
}
res <- try(f2(fs), silent = TRUE)
res

## A flowSet with 5 experiments.
##
##    column names:
##    FL1-H
```

So this form of usage of 'transform' method is only useful for the interactive exploratory. we highly recommend the usage of `transformList` instead for more robust and reproducible code.

```
myTrans <- transformList('FL1-H', aTrans)
transform(fs, myTrans)

## A flowSet with 5 experiments.
##
##    column names:
##    FSC-H SSC-H FL1-H FL2-H FL3-H FL1-A FL4-H
```

7   Gating

The most common task in the analysis of flow cytometry data is some form of filtering operation, also known as gating, either to obtain summary statistics about the number of events that meet certain criteria or to perform further analysis on a subset of the data. Most filtering operations are a composition of one or more common filtering operations. The definition of gates in flowCore follows the Gating Markup Language Candidate Recommendation Spidlen et al. (2008), thus any flowCore gating strategy can be reproduced by any other software that also adheres to the standard and *vice versa*.

7.1   Standard gates and filters

Like transformations, flowCore includes a number of built-in common flow cytometry gates. The simplest of these gates are the geometric gates, which correspond to those typically found in interactive flow cytometry software:

rectangleGate Describes a cubic shape in one or more dimensions–a rectangle in one dimension is simply an interval gate.

polygonGate Describes an arbitrary two dimensional polygonal gate.

polytopeGate Describes a region that is the convex hull of the given points. This gate can exist in dimensions higher than 2, unlike the `polygonGate`.

ellipsoidGate Describes an ellipsoidal region in two or more dimensions

These gates are all described in more or less the same manner (see man pages for more details):

```
rectGate <- rectangleGate(filterId="Fluorescence Region",
                          "FL1-H"=c(0, 12), "FL2-H"=c(0, 12))
```

In addition, we introduce the notion of data-driven gates, or filters, not usually found in flow cytometry software. In these approaches, the necessary parameters are computed based on the properties of the underlying data, for instance by modeling data distribution or by density estimation:

norm2Filter A robust method for finding a region that most resembles a bivariate Normal distribution. Note that this method now resides in the flowStats package

kmeansFilter Identifies populations based on a one dimensional k-means clustering operation. Allows the specification of **multiple** populations.

7.2 Count Statistics

When we have constructed a filter, we can apply it in two basic ways. The first is to collect simple summary statistics on the number and proportion of events considered to be contained within the gate or filter. This is done using the `filter` method. The first step is to apply our filter to some data

```
result = filter(fs[[1]],rectGate)
result

## A filterResult produced by the filter named 'Fluorescence Region'
```

As we can see, we have returned a *filterResult* object, which is in turn a filter allowing for reuse in, for example, subsetting operations. To obtain count and proportion statistics, we take the summary of this *filterResult*, which returns a list of summary values:

```
summary(result)

## Fluorescence Region+: 9811 of 10000 events (98.11%)

summary(result)$n

## [1] 10000

summary(result)$true

## [1] 9811

summary(result)$p

## [1] 0.9811
```

A filter which contains multiple populations, such as the *kmeansFilter*, can return a list of summary lists:

```
summary(filter(fs[[1]],
               kmeansFilter("FSC-H"=c("Low", "Medium", "High"),
                             filterId="myKMeans")))

## Low: 2518 of 10000 events (25.18%)
## Medium: 5109 of 10000 events (51.09%)
## High: 2373 of 10000 events (23.73%)
```

A filter may also be applied to an entire *flowSet*, in which case it returns a list of *filterResult* objects:

```
filter(fs,rectGate)

## A list of filterResults for a flowSet containing 5 frames
## produced by the filter named 'Fluorescence Region'
```

7.3   Subsetting

To subset or split a *flowFrame* or *flowSet*, we use the `Subset` and `split` methods respectively. The first, `Subset`, behaves similarly to the standard R `subset` function, which unfortunately could not be used. For example, the morphology parameters, Forward Scatter and Side Scatter, contain a large more-or-less ellipse shaped population:

```
autoplot(fs[[1]], "FSC-H", "SSC-H")
```


If we wished to deal only with that population, we might use `Subset` along with a *norm2Filter* object as follows:

```
library(flowStats)
morphGate <- norm2Filter("FSC-H", "SSC-H", filterId="MorphologyGate",
                         scale=2)
smaller <- Subset(fs, morphGate)
fs[[1]]

## flowFrame object '060909.001'
## with 10000 cells and 7 observables:
##      name        desc range minRange maxRange
## $P1 FSC-H FSC-Height  1024        0     1023
## $P2 SSC-H SSC-Height  1024        0     1023
## $P3 FL1-H       <NA>  1024        1    10000
## $P4 FL2-H       <NA>  1024        1    10000
## $P5 FL3-H       <NA>  1024        1    10000
## $P6 FL1-A       <NA>  1024        0     1023
## $P7 FL4-H       <NA>  1024        1    10000
## 141 keywords are stored in the 'description' slot

smaller[[1]]

## flowFrame object '060909.001'
## with 8312 cells and 7 observables:
##      name        desc range minRange maxRange
## $P1 FSC-H FSC-Height  1024        0     1023
## $P2 SSC-H SSC-Height  1024        0     1023
## $P3 FL1-H       <NA>  1024        1    10000
## $P4 FL2-H       <NA>  1024        1    10000
## $P5 FL3-H       <NA>  1024        1    10000
## $P6 FL1-A       <NA>  1024        0     1023
## $P7 FL4-H       <NA>  1024        1    10000
## 141 keywords are stored in the 'description' slot
```

Notice how the smaller *flowFrame* objects contain fewer events. Now imagine we wanted to use a *kmeansFilter* as before to split our first fluorescence parameter into three populations. To do this we employ the split function:

```
split(smaller[[1]], kmeansFilter("FSC-H"=c("Low","Medium","High"),
                                  filterId="myKMeans"))

## $Low
## flowFrame object '060909.001 (Low)'
## with 2422 cells and 7 observables:
##      name        desc range minRange maxRange
## $P1 FSC-H FSC-Height  1024        0     1023
## $P2 SSC-H SSC-Height  1024        0     1023
```

```
## $P3 FL1-H         <NA>  1024       1    10000
## $P4 FL2-H         <NA>  1024       1    10000
## $P5 FL3-H         <NA>  1024       1    10000
## $P6 FL1-A         <NA>  1024       0     1023
## $P7 FL4-H         <NA>  1024       1    10000
## 141 keywords are stored in the 'description' slot
##
## $Medium
## flowFrame object '060909.001 (Medium)'
## with 3563 cells and 7 observables:
##      name        desc range minRange maxRange
## $P1 FSC-H FSC-Height  1024       0     1023
## $P2 SSC-H SSC-Height  1024       0     1023
## $P3 FL1-H         <NA>  1024       1    10000
## $P4 FL2-H         <NA>  1024       1    10000
## $P5 FL3-H         <NA>  1024       1    10000
## $P6 FL1-A         <NA>  1024       0     1023
## $P7 FL4-H         <NA>  1024       1    10000
## 141 keywords are stored in the 'description' slot
##
## $High
## flowFrame object '060909.001 (High)'
## with 2327 cells and 7 observables:
##      name        desc range minRange maxRange
## $P1 FSC-H FSC-Height  1024       0     1023
## $P2 SSC-H SSC-Height  1024       0     1023
## $P3 FL1-H         <NA>  1024       1    10000
## $P4 FL2-H         <NA>  1024       1    10000
## $P5 FL3-H         <NA>  1024       1    10000
## $P6 FL1-A         <NA>  1024       0     1023
## $P7 FL4-H         <NA>  1024       1    10000
## 141 keywords are stored in the 'description' slot
```

or for an entire *flowSet*

```
split(smaller, kmeansFilter("FSC-H"=c("Low", "Medium", "High"),
                            filterId="myKMeans"))

## $Low
## A flowSet with 5 experiments.
##
## An object of class 'AnnotatedDataFrame'
##   rowNames: 060909.001 060909.002 ... 060909.005 (5 total)
##   varLabels: name population
##   varMetadata: labelDescription
```

```
##
##   column names:
##    FSC-H SSC-H FL1-H FL2-H FL3-H FL1-A FL4-H
##
## $Medium
## A flowSet with 5 experiments.
##
## An object of class 'AnnotatedDataFrame'
##    rowNames: 060909.001 060909.002 ... 060909.005 (5 total)
##    varLabels: name population
##    varMetadata: labelDescription
##
##   column names:
##    FSC-H SSC-H FL1-H FL2-H FL3-H FL1-A FL4-H
##
## $High
## A flowSet with 5 experiments.
##
## An object of class 'AnnotatedDataFrame'
##    rowNames: 060909.001 060909.002 ... 060909.005 (5 total)
##    varLabels: name population
##    varMetadata: labelDescription
##
##   column names:
##    FSC-H SSC-H FL1-H FL2-H FL3-H FL1-A FL4-H
```

7.4 Combining Filters

Of course, most filtering operations consist of more than one gate. To combine gates and filters we use the standard R Boolean operators: &, — and ! to construct an intersection, union and complement respectively:

```
rectGate & morphGate

## filter 'Fluorescence Region and MorphologyGate'
## the intersection between the 2 filters
##
## Rectangular gate 'Fluorescence Region' with dimensions:
##   FL1-H: (0,12)
##   FL2-H: (0,12)
##
## norm2Filter 'MorphologyGate' in dimensions FSC-H and SSC-H with parameters:
##   method: covMcd
##   scale.factor: 2
##   n: 50000
```

```
rectGate | morphGate

## filter 'Fluorescence Region or MorphologyGate'
## the union of the 2 filters
##
## Rectangular gate 'Fluorescence Region' with dimensions:
##    FL1-H: (0,12)
##    FL2-H: (0,12)
##
## norm2Filter 'MorphologyGate' in dimensions FSC-H and SSC-H with parameters:
##    method: covMcd
##    scale.factor: 2
##    n: 50000

!morphGate

## filter 'not MorphologyGate', the complement of
## norm2Filter 'MorphologyGate' in dimensions FSC-H and SSC-H with parameters:
##    method: covMcd
##    scale.factor: 2
##    n: 50000
```

We also introduce the notion of the subset operation, denoted by either %subset% or %&%. This combination of two gates first performs a subsetting operation on the input *flowFrame* using the right-hand filter and then applies the left-hand filter. For example,

```
summary(filter(smaller[[1]],rectGate %&% morphGate))

## Fluorescence Region in MorphologyGate+: 7194 of 8312 events (86.55%)
```

first calculates a subset based on the `morphGate` filter and then applies the `rectGate`.

7.5 Transformation Filters

Finally, it is sometimes desirable to construct a filter with respect to transformed parameters. To allow for this in our filtering constructs we introduce a special form of the `transform` method along with another filter combination operator %on%, which can be applied to both filters and *flowFrame* or *flowSet* objects. To specify our transform filter we must first construct a transform list using a simplified version of the `transform` function:

```
tFilter <- transform("FL1-H"=log,"FL2-H"=log)
tFilter

## An object of class "transformList"
## Slot "transforms":
## [[1]]
```

24

```
## transformMap for parameter 'FL1-H' mapping to 'FL1-H'
##
## [[2]]
## transformMap for parameter 'FL2-H' mapping to 'FL2-H'
##
##
## Slot "transformationId":
## [1] "defaultTransformation"
```

Note that this version of the transform filter does not take parameters on the right-hand side–the functions can only take a single vector that is specified by the parameter on the left-hand side. In this case those parameters are "FL1-H" and "FL2-H." The function also does not take a specific *flowFrame* or *flowSet* allowing us to use this with any appropriate data. We can then construct a filter with respect to the transform as follows:

```
rect2 <- rectangleGate(filterId="Another Rect", "FL1-H"=c(1,2),
"FL2-H"=c(2,3)) %on% tFilter
rect2

## transformed filter 'Another Rect on transformed values of FL1-H,FL2-H'
```

Additionally, we can use this construct directly on a *flowFrame* or *flowSet* by moving the transform to the left-hand side and placing the data on the right-hand side:

```
autoplot(tFilter %on% smaller[[1]], "FL1-H","FL2-H")
```


which has the same effect as the log transform used earlier.

8 Extensions from flowWorkspace: *GatingHierarchy* and *GatingSet*

The now-defunct *filterSet* class was very limited in its use for complex analysis work flows. It was result-centric and it was hard to access intermediate results. flowWorkspace and the openCyto framework (http://opencyto.org) offer much more versatile tools for such tasks through the *GatingSet* class. The general idea is to let the software handle the organization of intermediate results and operations and to provide a unified API to access and summarize these operations.

8.1 Abstraction of GatingSet

There are two classes in flowWorkspace that are used to abstract work flows: *GatingSet* objects are the basic container holding all the necessary bits and pieces and they are the main structure for user interaction. It is the container storing multiple *GatingHierarchy* objects which are associated with individual samples. One can think of *GatingSet* as corresponding to *(*flowSet) and *GatingHierarchy* as corresponding to *(*flowFrame).

It is important to know that *GatingSet* objects use an external pointer to store the gating tree and thus most of its accessors have a reference semantic instead of the pass-by-value semantic that is usually found in the R language. The main consequence on the user-level is the fact that direct assignments to a *GatingSet* object are usually not necessary; i.e., functions that operate on the *GatingSet* have the potential side-effect of modifying the object.

8.2 Creating *GatingSet* objects

Before creating a *GatingSet*, we need to have flow data loaded into R as a *flowSet* or *ncdfFlowSet* (the disk-based version of *flowSet*, to handle large data sets that are too big for memory; see ncdfFlow package).

```r
library(flowWorkspace)
fcsfiles <- list.files(pattern = "CytoTrol",
                        system.file("extdata",
                                    package = "flowWorkspaceData"),
                        full = TRUE)
fs <- read.flowSet(fcsfiles)
```

Then a *GatingSet* can be created using the constructor `GatingSet`.

```r
gs <- GatingSet(fs)
gs

## A GatingSet with 2 samples
```

Normally, we want to compensate the data first by using a user-supplied compensation matrix:

```r
comp

## Compensation object 'defaultCompensation':
```

26

```
##            B710-A     G560-A        G780-A       R660-A      R780-A      V450-A    V545-A
## B710-A 1.000000 0.0009476  0.071170002  0.0362400 0.1800000 0.007104 0.007608
## G560-A 0.115400 1.0000000  0.009097001  0.0018360 0.0000000 0.000000 0.000000
## G780-A 0.014280 0.0380000  1.000000000  0.0006481 0.1500000 0.000000 0.000000
## R660-A 0.005621 0.0000000  0.006604000  1.0000000 0.1786000 0.000000 0.000000
## R780-A 0.000000 0.0000000  0.035340000  0.0102100 1.0000000 0.000000 0.000000
## V450-A 0.000000 0.0000000 -0.059999999 -0.0400000 0.0000000 1.000000 0.410000
## V545-A 0.002749 0.0000000  0.000000000  0.0000000 0.0006963 0.035000 1.000000

gs <- compensate(gs, comp)
fs_comp <- getData(gs)
```

We can query the available nodes in the *GatingSet* using the `gs_get_pop_paths` method:

```
gs_get_pop_paths(gs)
```

```
## [1] "root"
```

It shows the only node "root" which corresponds to the raw flow data just added.

8.3   transform the data

Transformation can be either done on a *flowSet* before using it to construct a *GatingSet* or a *transformerList* can be directly added to a *GatingSet*:

```
biexpTrans <- flowJo_biexp_trans(channelRange=4096, maxValue=262144
                      , pos=4.5,neg=0, widthBasis=-10)
chnls <- parameters(comp)
tf <- transformerList(chnls, biexpTrans)

# or use estimateLogicle directly on GatingHierarchy
# object to generate transformerList automatically:
# tf <- estimateLogicle(gs[[1]], chnls)

gs <- transform(gs, tf)
```

```
p1 <- autoplot(fs_comp[[1]], "B710-A") + ggtitle("raw")
p2 <- autoplot(gs_cyto_data(gs)[[1]], "B710-A") +
          ggtitle("trans") +
          ggcyto_par_set(limits = "instrument")
```

```
## Coordinate system already present.  Adding new coordinate system, which
will replace the existing one.
```

```
grid.arrange(as.ggplot(p1), as.ggplot(p2), ncol = 2)
```

Note that we did assign the return value of `transform` back to `gs`. This is because the flow data is stored as an R object and thus transforming the data still follows the pass-by-value semantics.

8.4   Add the gates

Some basic flowCore *filter*s can be added to a *GatingSet*:

```
rg1 <- rectangleGate("FSC-A"=c(50000, Inf), filterId="NonDebris")
gs_pop_add(gs, rg1, parent = "root")

## [1] 2

gs_get_pop_paths(gs)

## [1] "root"        "/NonDebris"

# gate the data
recompute(gs)

## done!
```

As we see, here we don't need to assign *GatingSet* back because all the modifications are made in place to the *external pointer* rather than the R object itself. And now there is one new population node under the "root" node called "NonDebris". The node is named after the `filterId` of the gate if not explictly supplied. After the gates are added, the actual gating process is done by explictly calling the `recompute` method. Note that the numeric value it returns is the internal ID for the new population just added, which can normally be ignored since the gating path is the recommended way to refer to population nodes later.

To view the gate we just added,

```
autoplot(gs, "NonDebris")

## Coordinate system already present.  Adding new coordinate system, which
will replace the existing one.
```


Since They are 1-D gates, we can also display them in density plots

```
ggcyto(gs, aes(x = `FSC-A`)) + geom_density() + geom_gate("NonDebris")
```


To get population statistics for the given populatuion

```
gh_pop_get_stats(gs[[1]], "NonDebris")#counts

##          pop count
## 1: NonDebris 94764

gh_pop_get_stats(gs[[1]], "NonDebris", type = "percent")#proportion

##          pop    percent
## 1: NonDebris 0.7927985
```

Now we add two more gates:

```
# add the second gate
mat <- matrix(c(54272,59392,259071.99382782
                ,255999.994277954,62464,43008,70656
                ,234495.997428894,169983.997344971,34816)
              , nrow = 5)
colnames(mat) <-c("FSC-A", "FSC-H")
mat

##        FSC-A  FSC-H
## [1,]   54272  43008
## [2,]   59392  70656
## [3,] 259072 234496
## [4,] 256000 169984
## [5,]  62464  34816

pg <- polygonGate(mat)
gs_pop_add(gs, pg, parent = "NonDebris", name = "singlets")

## [1] 3

# add the third gate
rg2 <- rectangleGate("V450-A"=c(2000, Inf))
gs_pop_add(gs, rg2, parent = "singlets", name = "CD3")

## [1] 4

gs_get_pop_paths(gs)

## [1] "root"                      "/NonDebris"
## [3] "/NonDebris/singlets"       "/NonDebris/singlets/CD3"
```

We see two more nodes are added to 'GatingSet' and the population names are explicitly specified during the addition this time.

A *quadrantGate*, which results in four sub-populations, is also supported.

30

```
qg <- quadGate("B710-A" = 2000, "R780-A" = 3000)
gs_pop_add(gs, qg, parent="CD3", names = c("CD8", "DPT", "CD4", "DNT"))

## [1] 5 6 7 8

gs_pop_get_children(gs[[1]], "CD3")

## [1] "/NonDebris/singlets/CD3/CD8" "/NonDebris/singlets/CD3/DPT"
## [3] "/NonDebris/singlets/CD3/CD4" "/NonDebris/singlets/CD3/DNT"

# gate the data from "singlets"
recompute(gs, "singlets")

## done!
```

Here we see four child nodes are added to the "CD3" parent node. Four quadrants are named explicitly through the names argument by clock-wise order (start from top-left quadrant). Recomputing only needs to be done once from the first ungated node, which will automatically compute all of its descendants.

To plot the underlying tree:

```
plot(gs)
```

```
## Loading required package:  Rgraphviz
## Loading required package:  graph
## Loading required package:  BiocGenerics
## Loading required package:  parallel
##
## Attaching package:  'BiocGenerics'
## The following objects are masked from 'package:parallel':
##
##     clusterApply, clusterApplyLB, clusterCall, clusterEvalQ,
##     clusterExport, clusterMap, parApply, parCapply, parLapply,
##     parLapplyLB, parRapply, parSapply, parSapplyLB
## The following object is masked from 'package:flowStats':
##
##     normalize
## The following object is masked from 'package:gridExtra':
##
##     combine
## The following object is masked from 'package:flowCore':
##
##     normalize
```

31

```
## The following objects are masked from 'package:stats':
##
##     IQR, mad, sd, var, xtabs
## The following objects are masked from 'package:base':
##
##     Filter, Find, Map, Position, Reduce, anyDuplicated, append,
##     as.data.frame, basename, cbind, colnames, dirname, do.call,
##     duplicated, eval, evalq, get, grep, grepl, intersect, is.unsorted,
##     lapply, mapply, match, mget, order, paste, pmax, pmax.int, pmin,
##     pmin.int, rank, rbind, rownames, sapply, setdiff, sort, table,
##     tapply, union, unique, unsplit, which, which.max, which.min
## Loading required package:  grid
```


To plot all gates for one sample:

```
autoplot(gs[[1]])

## Coordinate system already present.  Adding new coordinate system, which
will replace the existing one.
## Coordinate system already present.  Adding new coordinate system, which
will replace the existing one.
## Coordinate system already present.  Adding new coordinate system, which
will replace the existing one.
## Coordinate system already present.  Adding new coordinate system, which
will replace the existing one.
```
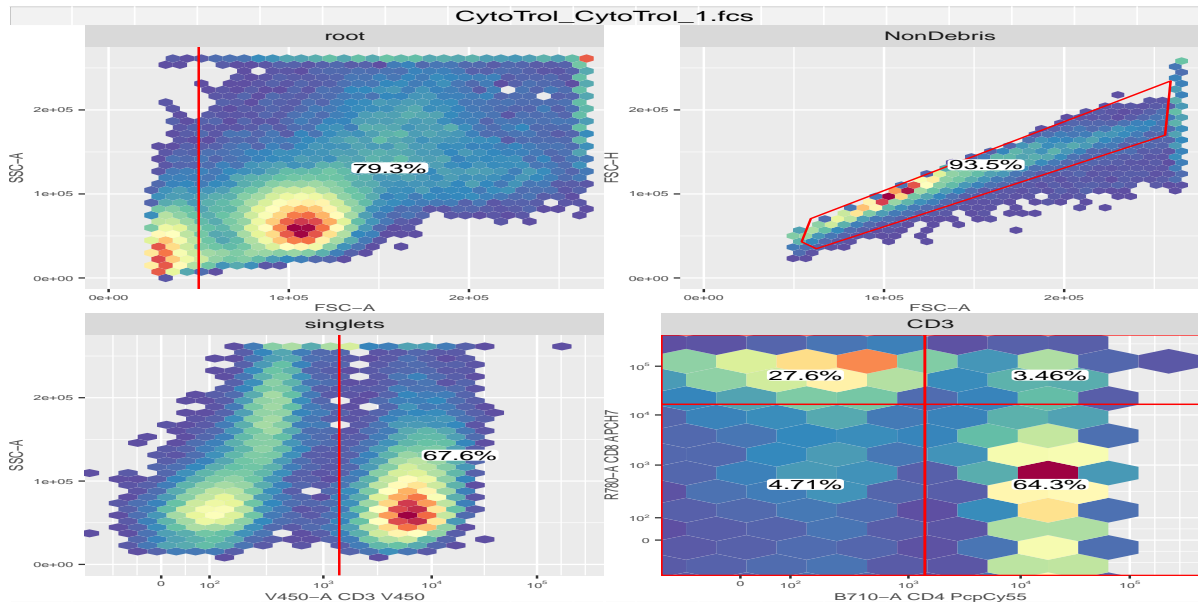
To retreive the underlying flow data for a gated *population*:

```
fs_nonDebris <- getData(gs, "NonDebris")
fs_nonDebris

## A cytoset with 2 samples.
##
##    column names:
##      FSC-A, FSC-H, FSC-W, SSC-A, B710-A, R660-A, R780-A, V450-A, V545-A, G560-A,

nrow(fs_nonDebris[[1]])

## [1] 94764

nrow(fs[[1]])

## [1] 119531
```

To get all the population statistics:

```
gs_pop_get_count_fast(gs)

##                          name                Population                 Parent
##  1: CytoTrol_CytoTrol_1.fcs                /NonDebris                   root
##  2: CytoTrol_CytoTrol_1.fcs        /NonDebris/singlets              /NonDebris
##  3: CytoTrol_CytoTrol_1.fcs     /NonDebris/singlets/CD3     /NonDebris/singlets
##  4: CytoTrol_CytoTrol_1.fcs /NonDebris/singlets/CD3/CD8 /NonDebris/singlets/CD3
##  5: CytoTrol_CytoTrol_1.fcs /NonDebris/singlets/CD3/DPT /NonDebris/singlets/CD3
##  6: CytoTrol_CytoTrol_1.fcs /NonDebris/singlets/CD3/CD4 /NonDebris/singlets/CD3
```

```
##  7: CytoTrol_CytoTrol_1.fcs /NonDebris/singlets/CD3/DNT /NonDebris/singlets/CD3
##  8: CytoTrol_CytoTrol_2.fcs                        /NonDebris                        root
##  9: CytoTrol_CytoTrol_2.fcs           /NonDebris/singlets            /NonDebris
## 10: CytoTrol_CytoTrol_2.fcs     /NonDebris/singlets/CD3    /NonDebris/singlets
## 11: CytoTrol_CytoTrol_2.fcs /NonDebris/singlets/CD3/CD8 /NonDebris/singlets/CD3
## 12: CytoTrol_CytoTrol_2.fcs /NonDebris/singlets/CD3/DPT /NonDebris/singlets/CD3
## 13: CytoTrol_CytoTrol_2.fcs /NonDebris/singlets/CD3/CD4 /NonDebris/singlets/CD3
## 14: CytoTrol_CytoTrol_2.fcs /NonDebris/singlets/CD3/DNT /NonDebris/singlets/CD3
##      Count ParentCount
##  1: 94764      119531
##  2: 88586       94764
##  3: 59911       88586
##  4: 16515       59911
##  5:  2070       59911
##  6: 38506       59911
##  7:  2820       59911
##  8: 94290      115728
##  9: 88334       94290
## 10: 59845       88334
## 11: 16774       59845
## 12:  2200       59845
## 13: 38127       59845
## 14:  2744       59845
```

8.5 Removing nodes from a *GatingSet* object

There are dependencies between *nodes* in the hierarchical structure of a *GatingSet* object. Thus, removing a particular node means also removing all of its associated child *nodes*.

```
Rm('CD3', gs)
gs_get_pop_paths(gs)

## [1] "root"                  "/NonDebris"          "/NonDebris/singlets"

Rm('NonDebris', gs)
gs_get_pop_paths(gs)

## [1] "root"
```

Now for the larger data set, it would be either inaccurate to apply the same hard-coded gate to all samples or impractical to manually set the gate coordinates for each individual sample. openCyto((G. et al., 2014)) provides some data-driven gating functions to automatically generate these gates.

For example, `mindensity` can be used for estimating a "nonDebris" gate for each sample.

```
if(require(openCyto)){
thisData <- getData(gs)
nonDebris_gate <- fsApply(thisData,
                          function(fr)
                            openCyto:::.mindensity(fr, channels = "FSC-A"))
gs_pop_add(gs, nonDebris_gate, parent = "root", name = "nonDebris")
recompute(gs)
}

## Loading required package:   openCyto
## done!
```

singletGate can be used for estimating "singlets"

```
if(require(openCyto)){
thisData <- getData(gs, "nonDebris") #get parent data
singlet_gate <- fsApply(thisData,
                        function(fr)
                          openCyto:::.singletGate(fr, channels =c("FSC-A", "FSC-H'
gs_pop_add(gs, singlet_gate, parent = "nonDebris", name = "singlets")
recompute(gs)
}

## done!
```

and then mindensity can be used again for the "CD3" gate

```
if(require(openCyto)){
thisData <- getData(gs, "singlets") #get parent data
CD3_gate <- fsApply(thisData,
                    function(fr)
                      openCyto:::.mindensity(fr, channels ="V450-A"))
gs_pop_add(gs, CD3_gate, parent = "singlets", name = "CD3")
recompute(gs)
}

## done!
```

and then we can use the more adanced version of quadrantGate: quadGate.seq for gating "CD4" and "CD8" sequentially:

```
if(require(openCyto)){
thisData <- getData(gs, "CD3") #get parent data
Tsub_gate <- fsApply(thisData,
                     function(fr)
```

```
                    openCyto::quadGate.seq(fr,
                        channels = c("B710-A", "R780-A"),
                        gFunc = 'mindensity')
                )
gs_pop_add(gs, Tsub_gate, parent = "CD3", names = c("CD8", "DPT", "CD4", "DNT"))
recompute(gs)
}

## done!
```

and then plot the gates

```
autoplot(gs[[1]])

## Coordinate system already present.  Adding new coordinate system, which
will replace the existing one.
## Coordinate system already present.  Adding new coordinate system, which
will replace the existing one.
## Coordinate system already present.  Adding new coordinate system, which
will replace the existing one.
## Coordinate system already present.  Adding new coordinate system, which
will replace the existing one.
```
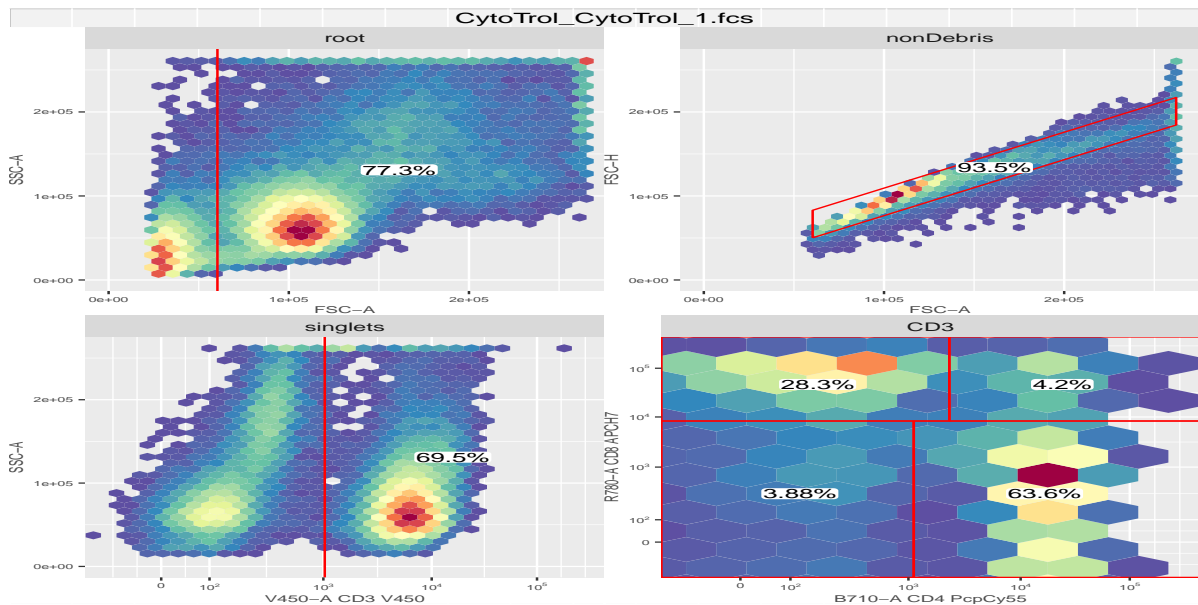

Note that in order to get parent gated data by `getData`, we have to call `recompute` after adding each gate. The result is very similar to the manual gates but the gating process is more data-driven and more consistent across samples.

To further automate the process, a gating pipeline can be established through OpenCyto((G. et al., 2014)) that defines the hierarchical gating template in a text-based csv file. (See more details from http://opencyto.org)

References

C. Bruce Bagwell. DNA histogram analysis for node-negative breast cancer. *Cytometry A*, 58:76–78, 2004.

Mark S Boguski and Martin W McIntosh. Biomedical informatics for proteomics. *Nature*, 422:233–237, 2003.

Raul C Braylan. Impact of flow cytometry on the diagnosis and characterization of lymphomas, chronic lymphoproliferative disorders and plasma cell neoplasias. *Cytometry A*, 58:57–61, 2004.

A. Brazma. On the importance of standardisation in life sciences. *Bioinformatics*, 17:113–114, 2001.

M. Chicurel. Bioinformatics: bringing it all together. *Nature*, 419:751–755, 2002.

Finak G., Frelinger J., Newell E.W., Ramey J., Davis M.M., Kalams S.A., De Rosa S.C., and Gottardo R. *OpenCyto: An Open Source Infrastructure for Scalable, Robust, Reproducible, and Automated, End-to-End Flow Cytometry Data Analysis*, volume 10. Public Library of Science, Aug 2014.

Maura Gasparetto, Tracy Gentry, Said Sebti, Erica O'Bryan, Ramadevi Nimmanapalli, Michelle A Blaskovich, Kapil Bhalla, David Rizzieri, Perry Haaland, Jack Dunne, and Clay Smith. Identification of compounds that enhance the anti-lymphoma activity of rituximab using flow cytometric high-content screening. *J Immunol Methods*, 292:59–71, 2004.

M. Keeney, D. Barnett, and J. W. Gratama. Impact of standardization on clinical cell analysis by flow cytometry. *J Biol Regul Homeost Agents*, 18:305–312, 2004.

N. LeMeur and F. Hahne. Analyzing flow cytometry data with bioconductor. *Rnews*, 6:27–32, 2006.

DR Parks. *Data Processing and Analysis: Data Management.*, volume 1 of *Current Protocols in Cytometry*. John Wiley & Sons, Inc, New York, 1997.

J. Spidlen, R.C. Gentleman, P.D. Haaland, M. Langille, N. Le Meur N, M.F. Ochs, C. Schmitt, C.A. Smith, A.S. Treister, and R.R. Brinkman. Data standards for flow cytometry. *OMICS*, 10(2):209–214, 2006.

J. Spidlen, R.C. Leif, W. Moore, M. Roederer, International Society for the Advancement of Cytometry Data Standards Task Force, and R.R. Brinkman. Gating-ml: Xml-based gating descriptions in flow cytometry. *Cytometry A*, 73A(12):1151–1157, 2008.

Maria A Suni, Holli S Dunn, Patricia L Orr, Rian de Laat, Elizabeth Sinclair, Smita A Ghanekar, Barry M Bredt, John F Dunne, Vernon C Maino, and Holden T Maecker. Performance of plate-based cytokine flow cytometry with automated data analysis. *BMC Immunol*, 4:9, 2003.