PharmacoGx: An R Package for Analysis of Large Pharmacogenomic Datasets

Petr Smirnov [1] Zhaleh Safikhani [1,2] Benjamin Haibe-Kains

[1]*Princess Margaret Cancer Centre, University Health Network, Toronto Canada*

[2]*Department of Medical Biophysics, University of Toronto, Toronto Canada*

[3]*Department of Computer Science, University of Toronto, Toronto Canada*

Contents

Introduction

Pharmacogenomics holds much potential to aid in discovering drug response biomarkers and developing novel targeted therapies, leading to development of precision medicine and working towards the goal of personalized therapy. Several large experiments have been conducted, both to molecularly characterize drug dose response across many cell lines, and to examine the molecular response to drug administration. However, the experiments lack a standardization of protocols and annotations, hindering meta-analysis across several experiments.

`PharmacoGx` was developed to address these challenges, by providing a unified framework for downloading and analyzing large pharmacogenomic datasets which are extensively curated to ensure maximum overlap and consistency. `PharmacoGx` is based on a level of abstraction from the raw experimental data, and allows bioinformaticians and biologists to work with data at the level of genes, drugs and cell lines. This provides a more intuitive interface and, in combination with unified curation, simplifies analyses between multiple datasets.

To organize the data released by each experiment, we developed the *PharmacoSet* class. This class efficiently stores different types of data and facilitates interrogating the data by drug or cell line. The *PharmacoSet* is also versatile in its ability to deal with two distinct types of pharmacogenomic datasets. The first type, known as *sensitivity* datasets, are datasets where cell lines were profiled on the molecular level, and then tested for drug dose response. The second type of dataset is the *perturbation* dataset. These types of datasets profile a cell line on the molecular level before and after administration of a compound, to characterize the action of the compound on the molecular level.

With the first release of *PharmacoGx* we have fully curated and created PharmacoSet objects for three publicly available large pharmacogenomic datasets. Two of these datasets are of the *sensitivity* type. These are the Genomics of Drug Sensitivity in Cancer Project (GDSC) (Garnett et al. 2012) and the Cancer Cell Line Encyclopedia (CCLE) (Barretina et al. 2012). The third dataset is of the *perturbation* type, the Connectivity Map (CMAP) project (Lamb et al. 2006).

Furthermore, `PharmacoGx` provides a suite of parallelized functions which facilitate drug response biomarker discovery, and molecular drug characterization. This vignette will provide two example analysis case studies. The first will be comparing gene expression and drug sensitivity measures across the CCLE and GDSC projects. The second case study will interrogate the CMAP database with a known signature of up and down regulated genes for HDAC inhibitors as published in [Glaser et al. (2003)}. Using the Connectivity Score as defined in (Lamb et al. 2006), it will be seen that known HDAC inhibitors have a high numerical score and high significance.

For the purpose of this vignette, an extremely minuscule subset of all three *PharmacoSet* objects are included with the package as example data. They are included for illustrative purposes only, and the results obtained with them will likely be meaningless.

Installation and Settings

`PharmacoGx` requires that several packages are installed. However, all dependencies are available from CRAN or Bioconductor.

```
if (!requireNamespace("BiocManager", quietly=TRUE))
    install.packages("BiocManager")
BiocManager::install('PharmacoGx')
```

Load `PharamacoGx` into your current workspace:

```
library(PharmacoGx)
```

Requirements

`PharmacoGx` has been tested on Windows, Mac and Cent OS platforms. The packages uses the core
R package `parallel` to preform parallel computations, and therefore if parallelization is desired, the
dependencies for the parallel package must be met.

Downloading PharmacoSet Objects

We have made the PharmacoSet objects of the curated datasets available for download using functions
provided in the package. A table of available PharmacoSet objects can be obtained by using the
`availablePSets()` function. Any of the PharmacoSets in the table can then be downloaded by calling
downloadPSet, which saves the datasets into a directory of the users choice, and returns the data into
the R session.

```
availablePSets()
GDSC <- downloadPSet("GDSC_2020(v2-8.2)")
```

Downlading Drug Signatures

The package also provides tools to compute drug perturbation and sensitivity signatures, as explained
below. However, the computation of the perturbation signatures is very lengthy, so for users' convenience
we have precomputed the signatures for our perturbation PharmacoSet objects and made them available
for download using the function `downloadPertSig()`.

Case Study

(In)Consistency across large pharmacogenomic studies

Our first case study illustrates the functions for analysis of the *sensitivity* type of dataset. The case
study will investigate the consistency between the GDSC and CCLE datasets, recreating the analysis
similar to our *Inconsistency in Large Pharmacogenomic Studies* paper (Haibe-Kains et al. 2013). In
both CCLE and GDSC, the transcriptome of cells was profiled using an Affymatrix microarray chip.
Cells were also tested for their response to increasing concentrations of various compounds, and form this
the IC50 and AUC were computed. However, the cell and drugs names used between the two datasets
were not consistent. Furthermore, two different microarray platforms were used. However, `PharmacoGx`
allows us to overcome these differences to do a comparative study between these two datasets.

GDSC was profiled using the hgu133a platform, while CCLE was profiled with the expanded hgu133plus2
platform. While in this case the hgu133a is almost a strict subset of hgu133plus2 platform, the expression
information in `PharmacoSet` objects is summarized by Ensemble Gene Ids, allowing datasets with
different platforms to be directly compared. The probe to gene mapping is done using the BrainArray
customCDF for each platform (Sabatti et al. 2002).

To begin, you would load the datasets from disk or download them using the *downloadPSet* function
above. In the following example, we use the toy datasets provided with the package to illustrate the
process, but to recreate the full analysis the full PharmacoSets have to be downloaded.

We want to investigate the consistency of the data between the two datasets. The common intersection between the datasets can then be found using *intersectPSet*. We create a summary of the gene expression and drug sensitivity measures for both datasets, so we are left with one gene expression profile and one sensitivity profile per cell line within each dataset. We can then compare the gene expression and sensitivity measures between the datasets using a standard correlation coefficient.

```r
library(Biobase)
library(SummarizedExperiment)
library(S4Vectors)
library(PharmacoGx)
data("GDSCsmall")
data("CCLEsmall")
commonGenes <- intersect(fNames(GDSCsmall, "rna"),
                         fNames(CCLEsmall,"rna"))
common <- intersectPSet(list('CCLE'=CCLEsmall,
                             'GDSC'=GDSCsmall),
                        intersectOn=c("cell.lines", "drugs"),
                        strictIntersect=TRUE)


GDSC.auc <- summarizeSensitivityProfiles(
            common$GDSC,
            sensitivity.measure='auc_published',
            summary.stat="median",
            verbose=FALSE)
CCLE.auc <- summarizeSensitivityProfiles(
            common$CCLE,
            sensitivity.measure='auc_published',
            summary.stat="median",
            verbose=FALSE)

GDSC.ic50 <- summarizeSensitivityProfiles(
            common$GDSC,
            sensitivity.measure='ic50_published',
            summary.stat="median",
            verbose=FALSE)
CCLE.ic50 <- summarizeSensitivityProfiles(
            common$CCLE,
            sensitivity.measure='ic50_published',
            summary.stat="median",
            verbose=FALSE)

GDSCexpression <- summarizeMolecularProfiles(common$GDSC,
                                    cellNames(common$GDSC),
                                    mDataType="rna",
                                    features=commonGenes,
                                    verbose=FALSE)
CCLEexpression <- summarizeMolecularProfiles(common$CCLE,
                                     cellNames(common$CCLE),
                                     mDataType="rna",
                                     features=commonGenes,
                                     verbose=FALSE)
gg <- fNames(common[[1]], 'rna')
cc <- cellNames(common[[1]])

ge.cor <- sapply(cc, function (x, d1, d2) {
  stats::cor(d1[ , x], d2[ , x], method="spearman",
             use="pairwise.complete.obs")
```

```r
## TO DO:: Ensure all assays are name so we can call by name instead of index
}, d1=assay(GDSCexpression, 1), d2=assay(CCLEexpression, 1))
ic50.cor <- sapply(cc, function (x, d1, d2) {
  stats::cor(d1[, x], d2[ , x], method="spearman",
             use="pairwise.complete.obs")
}, d1=GDSC.ic50, d2=CCLE.ic50)
auc.cor <- sapply(cc, function (x, d1, d2) {
  stats::cor(d1[ , x], d2[ , x], method="spearman",
             use="pairwise.complete.obs")
}, d1=GDSC.auc, d2=CCLE.auc)

w1 <- stats::wilcox.test(x=ge.cor, y=auc.cor,
                         conf.int=TRUE, exact=FALSE)
w2 <- stats::wilcox.test(x=ge.cor, y=ic50.cor,
                         conf.int=TRUE, exact=FALSE)
yylim <- c(-1, 1)
ss <- sprintf("GE vs. AUC = %.1E\nGE vs. IC50 = %.1E",
              w1$p.value, w2$p.value)
```

```r
boxplot(list("GE"=ge.cor,
             "AUC"=auc.cor,
             "IC50"=ic50.cor),
        main="Concordance between cell lines",
        ylab=expression(R[s]),
        sub=ss,
        ylim=yylim,
        col="lightgrey",
        pch=20,
        border="black")
```

Concordance between cell lines

GE vs. AUC = 1.5E−03
GE vs. IC50 = 1.0E−03

Query the Connectivity Map

The second case study illustrates the analysis of a perturbation type datasets, where the changes in cellular molecular profiles are compared before and after administering a compound to the cell line. Of these datasets, we have currently curated and made available for download the Connectivity Map (CMAP) dataset (Lamb et al. 2006).

For this case study, we will recreate an analysis from the paper by Lamb et al., in which a known signature for HDAC inhibitors (Glaser et al. 2003) is used to recover drugs in the CMAP dataset that are also known HDAC inhibitors. For this example, the package includes this signature, already mapped to the gene level, and it can be loaded by calling `data(HDAC\_genes)`.

Once again, we load the dataset, downloading it if needed using `downloadPSet` . We then recreate drug signatures for each drug using the function `drugPerturbationSig` to preform statistical modelling of the transcriptomic response to the application of each drug. We then compare the observed up-regulated and down-regulated genes to a the known HDAC signature, using the GSEA connectivity score to determine the correlation between the two signatures.

```r
library(PharmacoGx)
library(pander)
data(CMAPsmall)
drug.perturbation <- drugPerturbationSig(CMAPsmall,
                                         mDataType="rna",
                                         verbose=FALSE)
data(HDAC_genes)

res <- apply(drug.perturbation[,,c("tstat", "fdr")],
             2, function(x, HDAC){
    return(connectivityScore(x=x,
                             y=HDAC[,2,drop=FALSE],
                             method="fgsea", nperm=100))
  }, HDAC=HDAC_genes)

rownames(res) <- c("Connectivity", "P Value")
res <- t(res)
res <- res[order(res[,1], decreasing=TRUE),]
pander::pandoc.table(res,
  caption='Connectivity Score results for HDAC inhibitor gene signature.',
  style = 'rmarkdown')
```

Table 1: Connectivity Score results for HDAC inhibitor gene signature.

	Connectivity	P Value
vorinostat	0.9449	0.003837
alvespimycin	0.8831	0.01092
acetylsalicylic acid	0	1
rosiglitazone	-0.714	0.03775
pioglitazone	-0.7335	0.1136

As we can see, the known HDAC inhibitor Varinostat has a very strong connectivity score, as well as a very high significance as determined by permutation testing, in comparison to the other drugs, which score poorly.

This example serves as a validation of the method, and demonstrates the ease with which drug perturbation analysis can be done using `PharmacoGx`. While in this case we were matching a drug signature with a drug class signature, this method can also be used in the discovery of drugs that are anti-correlated with known disease signatures, to look for potential new treatments and drug repurposing.

Estimating Drug Sensitivity Measures

PharmacoGx includes functions to calculate estimated AUC (Area Under drug response Curve) and IC50 values from drugs dose response experiments that measure cell viability after applications of varying concentrations of drug. Additionally, these measures are recomputed for every *sensitivity* `PharmacoSet` we create and included alongside any measures published with the original data. The AUC measures originally published are labelled as *auc_published*, while the recomputed measures are labelled as *auc_recomputed*, and likewise for the IC50.

While the *PharmacoSets* already contain the recomputed data, the AUC and IC50 can be calculated for arbitrary data using the *computeIC50* and *computeAUC* functions. The AUC can be calculated using either the area under the curve defined by the actual points recorded, or the area under the curve fitted to the data.

Curve Fitting

While AUC can be numerically calculated without curve fitting, to estimate the IC50 a drug dose response curve must be fit to the data.The dose-response curves are fitted to the equation

$$y = E_\infty + \frac{1 - E_\infty}{1 + (\frac{x}{IC50})^{HS}}$$

where the maximum viability is normalized to be $y = y(0) = 1$, E_∞ denotes the minimum possible viability achieved by administering any amount of the drug, $IC50$ is the concentration at which viability is reduced to half of the viability observed in the presence of an arbitrarily large concentration of drug, and HS is a parameter describing the cooperativity of binding. $HS < 1$ denotes negative binding cooperativity, $HS = 1$ denotes noncooperative binding, and $HS > 1$ denotes positive binding cooperativity. The parameters of the curves are fitted using the least squares optimization framework. The fitting of these curves to arbitrary points is implemented by the *logLogisticRegression* function.

Plotting Drug-Dose Response Data

Drug-Dose response data included in the PharmacoSet objects can be conviniently plotted using the `drugDoseResponseCurve` function. Given a list of PharmacoSets, a drug name and a cell name, it will plot the drug dose response curves for the given cell-drug combination in each dataset, allowing direct comparisons of data between datasets.

Gene-Drug Association Modelling

`PharmacoGx` provides methods to model the association between drugs and molecular data such as transcriptomics, genomics and proteomics.

Sensitivity studies allow the discovery of molecular features that improve or inhibit the sensitivity of cell lines to various compounds, by looking at the association between the expression of the feature and the response towards each compound. This allows the selection of drugs to be tailored to each specific patient based on the expressed known sensitivity biomarkers. The function *drugSensitivitySig* models these associations.

Perturbation studies on the other hand look at the molecular profiles of a cell before and after application of a drug, and therefore can characterize the action of a drug on the molecular level. It is hypothesized that drugs which act to down-regulate expression of known disease biomarkers would be effective in reversing the cell from a diseased to healthy state. The function *drugPerturbationSig* models the molecular profiles of drugs tested in a *perturbation* dataset.

Sensitivity Modelling

The association between molecular features and response to a given drug is modelled using a linear regression model adjusted for tissue source:

$$Y = \beta_0 + \beta_i G_i + \beta_t T + \beta_b B$$

where Y denotes the drug sensitivity variable, G_i, T and B denote the expression of gene i, the tissue source and the experimental batch respectively, and βs are the regression coefficients. The strength of gene-drug association is quantified by β_i, above and beyond the relationship between drug sensitivity and tissue source. The variables Y and G are scaled (standard deviation equals to 1) to estimate standardized coefficients from the linear model. Significance of the gene-drug association is estimated by the statistical significance of β_i (two-sided t test). P-values are then corrected for multiple testing using the false discovery rate (FDR) approach.

As an example of the efficacy of the modelling approach, we can model the significance of the association between two drugs and their known biomarkers in CCLE. We examine the association between drug *17-AAG* and gene *NQO1*, as well as drug *PD-0325901* and gene *BRAF*:

```
library(pander)
data(CCLEsmall)
features <- fNames(CCLEsmall, "rna")[
                    which(featureInfo(CCLEsmall,
                                "rna")$Symbol == "NQO1")]
```

```
sig.rna <- drugSensitivitySig(object=CCLEsmall,
                              mDataType="rna",
                              drugs=c("17-AAG"),
                              features=features,
                              sensitivity.measure="auc_published",
                              molecular.summary.stat="median",
                              sensitivity.summary.stat="median",
                              verbose=FALSE)
sig.mut <- drugSensitivitySig(object=CCLEsmall,
                              mDataType="mutation",
                              drugs=c("PD-0325901"),
                              features="BRAF",
                              sensitivity.measure="auc_published",
                              molecular.summary.stat="and",
                              sensitivity.summary.stat="median",
                              verbose=FALSE)
sig <- rbind(sig.rna, sig.mut)
rownames(sig) <- c("17-AAG + NQO1","PD-0325901 + BRAF")
colnames(sig) <- dimnames(sig.mut)[[3]]
pander::pandoc.table(t(sig), style = "rmarkdown", caption='P Value of Gene-Drug Association' )
```

Table 2: P Value of Gene-Drug Association

	17-AAG + NQO1	PD-0325901 + BRAF
estimate	0.601	0.8252
se	0.05368	0.1339
n	492	472
tstat	11.2	6.161
fstat	125.3	37.95
pvalue	6.034e-26	1.613e-09
df	469	449
fdr	6.034e-26	1.613e-09

Perturbation Modelling

The molecular response profile of a given drug is modelled as a linear regression model adjusted experimental batch effects, cell specific differences, and duration of experiment to isolate the effect of the concentration of the drug applied:

$$G = \beta_0 + \beta_i C_i + \beta_t T + \beta_d D + \beta_b B$$

where G denotes the molecular feature expression (gene), C_i denotes the concentration of the compound applied, T the cell line identity, D denotes the duration of the experiment, B denotes the experimental batch, and βs are the regression coefficients. The strength of feature response is quantified by β_i. Unlike the sensitivity signatures, the G and C variables are unscaled. Significance of the gene-drug association is estimated by the statistical significance of β_i, calculated using an F-test on the improvement in fit after the inclusion of the term. P-values are then corrected for multiple testing using the false discovery rate (FDR) approach.

Connectivity Scoring

The package also provides two methods for quantifying the similarity between two molecular signatures of the form returned by *drugPerturbationSig* and *drugSensitivitySig*, or a set of up and down regulated genes as part of a disease signature. The two methods are the *GSEA* method as introduced by Subramanian et at (Subramanian et al. 2005), and *GWC*, a method based on a weighted Spearman correlation coefficient. Both methods are implemented by the *connectivityScore* function.

GSEA

The *GSEA* method is implemented to compare a signature returned by *drugPerturbationSig* with a known set of up and down regulated genes in a disease state. For the disease signature, the function expects a vector of features with a value, either binary (1, -1) or continuous, where the sign signifies if the gene is up or down regulated in the disease. The names of the vector are expected to be the feature names, matching the feature names used in the drug signature. The function then returns a GSEA score measuring the concordance of the disease signature to the drug signature, as well as an estimated P-Value for the significance of the connectivity determined by permutation testing.

GWC

The GWC method (Genome Wide Correlation) is implemented to compare two signatures of the same length, such as two drug signatures returned by *drugPerturbationSig*. The score is a Spearman correlation weighted by the normalized negative logarithm significance of each value. The normalization is done so that datasets of different size can be compared without concern for lower p-values in one dataset due to larger sample sizes.

More precisely, take X_i and Y_i to be the ranks of the first and second set of data respectively, and Px_i, Py_i to be the p-values of those observations. The weight for each pair of observations is:

$$Wx_i = \frac{-\log_{10}(Px_i)}{\sum_i -\log_{10}(Px_i)}$$

$$Wy_i = \frac{-\log_{10}(Py_i)}{\sum_i -\log_{10}(Py_i)}$$

$$W_i = Wx_i + Wy_i$$

If we further define the weighted mean as follows:

$$m(X;W) = \frac{\sum_i W_i X_i}{\sum_i W_i}$$

Then the weighted correlation is given by:

$$cor(X,Y,W) = \frac{\frac{\sum_i W_i(X_i - m(X;W))(Y_i - m(Y,W))}{\sum_i W_i}}{\sqrt{(\frac{\sum_i W_i(X_i - m(X;W))^2}{\sum_i W_i})(\frac{\sum_i W_i(Y_i - m(Y;W))^2}{\sum_i W_i})}}$$

This correlation therefore takes into account not only the ranking of each feature in a signature, but also of the significance of each rank.

Acknowledgements

The authors of the package would like to thank the investigators of the Genomics of Drug Sensitivity in Cancer Project, the Cancer Cell Line Encyclopedia and the Connectivity Map Project who have made their invaluable data available to the scientific community. We would also like to thank Mark Freeman for contributing the code for MLE fitting drug-dose response curves towards this package. We are indebted to Nehme El-Hachem, Donald Wang and Adrian She for their contributions towards the accurate curation of the datasets. Finally, it is important to acknowledge all the members of the Benjamin Haibe-Kains lab for their contribution towards testing and feedback during the design of the package.

Session Info

```
# set eval = FALSE if you don't want this info (useful for reproducibility) to appear
sessionInfo()
```

```
R version 4.2.0 RC (2022-04-19 r82224)
Platform: x86_64-pc-linux-gnu (64-bit)
Running under: Ubuntu 20.04.4 LTS

Matrix products: default
BLAS:   /home/biocbuild/bbs-3.15-bioc/R/lib/libRblas.so
LAPACK: /home/biocbuild/bbs-3.15-bioc/R/lib/libRlapack.so

locale:
 [1] LC_CTYPE=en_US.UTF-8       LC_NUMERIC=C
 [3] LC_TIME=en_GB              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] stats4    stats     graphics  grDevices utils     datasets  methods
[8] base

other attached packages:
 [1] pander_0.6.5               PharmacoGx_3.0.0
 [3] CoreGx_2.0.0               SummarizedExperiment_1.26.0
 [5] GenomicRanges_1.48.0       GenomeInfoDb_1.32.0
 [7] IRanges_2.30.0             S4Vectors_0.34.0
 [9] MatrixGenerics_1.8.0       matrixStats_0.62.0
[11] Biobase_2.56.0             BiocGenerics_0.42.0
[13] knitcitations_1.0.12       knitr_1.38

loaded via a namespace (and not attached):
 [1] fgsea_1.22.0               colorspace_2.0-3
 [3] ellipsis_0.3.2             lsa_0.73.2
 [5] XVector_0.36.0             SnowballC_0.7.0
 [7] MultiAssayExperiment_1.22.0 DT_0.22
 [9] fansi_1.0.3                lubridate_1.8.0
[11] RefManageR_1.3.0           xml2_1.3.3
[13] jsonlite_1.8.0             magicaxis_2.2.14
[15] cluster_2.1.3              shinydashboard_0.7.2
[17] shiny_1.7.1                mapproj_1.2.8
[19] compiler_4.2.0             httr_1.4.2
[21] backports_1.4.1            assertthat_0.2.1
[23] Matrix_1.4-1               fastmap_1.1.0
[25] limma_3.52.0               cli_3.3.0
[27] later_1.3.0                visNetwork_2.1.0
[29] htmltools_0.5.2            tools_4.2.0
[31] igraph_1.3.1               gtable_0.3.0
[33] glue_1.6.2                 GenomeInfoDbData_1.2.8
[35] reshape2_1.4.4             RANN_2.6.1
[37] dplyr_1.0.8                maps_3.4.0
[39] tinytex_0.38               fastmatch_1.1-3
[41] Rcpp_1.0.8.3               slam_0.1-50
[43] vctrs_0.4.1                BumpyMatrix_1.4.0
[45] xfun_0.30                  stringr_1.4.0
[47] mime_0.12                  lifecycle_1.0.1
[49] gtools_3.9.2               MASS_7.3-57
[51] zlibbioc_1.42.0            scales_1.2.0
[53] promises_1.2.0.1           relations_0.6-12
[55] parallel_4.2.0             RColorBrewer_1.1-3
[57] sets_1.0-21                yaml_2.3.5
```

```
 [59] gridExtra_2.3          ggplot2_3.3.5
 [61] downloader_0.4         stringi_1.7.6
 [63] highr_0.9              NISTunits_1.0.1
 [65] plotrix_3.8-2          checkmate_2.1.0
 [67] caTools_1.18.2         boot_1.3-28
 [69] BiocParallel_1.30.0    rlang_1.0.2
 [71] pkgconfig_2.0.3        bitops_1.0-7
 [73] pracma_2.3.8           evaluate_0.15
 [75] lattice_0.20-45        purrr_0.3.4
 [77] htmlwidgets_1.5.4      tidyselect_1.1.2
 [79] plyr_1.8.7             magrittr_2.0.3
 [81] R6_2.5.1               gplots_3.1.3
 [83] generics_0.1.2         sm_2.2-5.7
 [85] DelayedArray_0.22.0    DBI_1.1.2
 [87] pillar_1.7.0           RCurl_1.98-1.6
 [89] tibble_3.1.6           crayon_1.5.1
 [91] KernSmooth_2.23-20     utf8_1.2.2
 [93] rmarkdown_2.14         grid_4.2.0
 [95] data.table_1.14.2      marray_1.74.0
 [97] piano_2.12.0           digest_0.6.29
 [99] xtable_1.8-4           httpuv_1.6.5
[101] coop_0.6-3             munsell_0.5.0
[103] celestial_1.4.6        shinyjs_2.1.0
```

References

Barretina, J., G. Caponigro, N. Stransky, K. Venkatesan, A. A. Margolin, S. Kim, C. J. Wilson, J. Lehár, G. V. Kryukov, D. Sonkin, A. Reddy, M. Liu, L. Murray, M. F. Berger, J. E. Monahan, P. Morais, J. Meltzer, A. Korejwa, J. Jané-Valbuena, F. A. Mapa, J. Thibault, E. Bric-Furlong, P. Raman, A. Shipway, I. H. Engels, J. Cheng, G. K. Yu, J. Yu, P. Aspesi, M. de Silva, K. Jagtap, M. D. Jones, L. Wang, C. Hatton, E. Palescandolo, S. Gupta, S. Mahan, C. Sougnez, R. C. Onofrio, T. Liefeld, L. MacConaill, W. Winckler, M. Reich, N. Li, J. P. Mesirov, S. B. Gabriel, G. Getz, K. Ardlie, V. Chan, V. E. Myer, B. L. Weber, J. Porter, M. Warmuth, P. Finan, J. L. Harris, M. Meyerson, T. R. Golub, M. P. Morrissey, W. R. Sellers, R. Schlegel, and L. A. Garraway. 2012. The Cancer Cell Line Encyclopedia enables predictive modelling of anticancer drug sensitivity. Nature 483:603–607.

Garnett, M. J., E. J. Edelman, S. J. Heidorn, C. D. Greenman, A. Dastur, K. W. Lau, P. Greninger, I. R. Thompson, X. Luo, J. Soares, Q. Liu, F. Iorio, D. Surdez, L. Chen, R. J. Milano, G. R. Bignell, A. T. Tam, H. Davies, J. A. Stevenson, S. Barthorpe, S. R. Lutz, F. Kogera, K. Lawrence, A. McLaren-Douglas, X. Mitropoulos, T. Mironenko, H. Thi, L. Richardson, W. Zhou, F. Jewitt, T. Zhang, P. O'Brien, J. L. Boisvert, S. Price, W. Hur, W. Yang, X. Deng, A. Butler, H. G. Choi, J. W. Chang, J. Baselga, I. Stamenkovic, J. A. Engelman, S. V. Sharma, O. Delattre, J. Saez-Rodriguez, N. S. Gray, J. Settleman, P. A. Futreal, D. A. Haber, M. R. Stratton, S. Ramaswamy, U. McDermott, and C. H. Benes. 2012. Systematic identification of genomic markers of drug sensitivity in cancer cells. Nature 483:570–575.

Glaser, K. B., M. J. Staver, J. F. Waring, J. Stender, R. G. Ulrich, and S. K. Davidsen. 2003. Gene expression profiling of multiple histone deacetylase (HDAC) inhibitors: Defining a common gene set produced by HDAC inhibition in T24 and MDA carcinoma cell lines. Molecular cancer therapeutics 2:151–163.

Haibe-Kains, B., N. El-Hachem, N. J. Birkbak, A. C. Jin, A. H. Beck, H. J. W. L. Aerts, and J. Quackenbush. 2013. Inconsistency in large pharmacogenomic studies. Nature 504:389–393.

Lamb, J., E. D. Crawford, D. Peck, J. W. Modell, I. C. Blat, M. J. Wrobel, J. Lerner, J.-P. Brunet, A. Subramanian, K. N. Ross, M. Reich, H. Hieronymus, G. Wei, S. A. Armstrong, S. J. Haggarty, P. A. Clemons, R. Wei, S. A. Carr, E. S. Lander, and T. R. Golub. 2006. The Connectivity Map: Using Gene-Expression Signatures to Connect Small Molecules, Genes, and Disease. Science 313:1929–1935.

Sabatti, C., S. L. Karsten, and D. H. Geschwind. 2002. Thresholding rules for recovering a sparse signal from microarray experiments. Mathematical biosciences 176:17–34.

Subramanian, A., P. Tamayo, V. K. Mootha, S. Mukherjee, B. L. Ebert, M. A. Gillette, A. Paulovich, S. L. Pomeroy, T. R. Golub, E. S. Lander, and J. P. Mesirov. 2005. Gene set enrichment analysis: A knowledge-based approach for interpreting genome-wide expression profiles. Proceedings of the National Academy of Sciences 102:15545–15550.