

TITAN: Subclonal copy number and LOH prediction from whole genome sequencing of tumours

Gavin Ha

April 12, 2014

Contents

1	Introduction	1
2	Analysis Workflow Overview	2
3	Model training and parameter estimation	2
4	Extracting tumour and normal read depth	3
5	TITAN analysis	3
5.1	Loading the data	3
5.2	Loading the initial parameters	3
5.3	Correct GC content and mappability biases	5
5.4	Filter the data	6
5.5	Model training and parameter estimation	6
5.6	Find the genotype and clonal cluster state paths	8
5.7	Format and print results to file	8
5.8	Plotting the results	10
5.8.1	Copy number alterations (log ratio)	10
5.8.2	Loss of heterozygosity (allelic ratio)	10
5.8.3	Cellular prevalence and clonal clusters	11
6	Session Information	11

1 Introduction

TITAN is a probabilistic framework for predicting regions of copy number alterations (LOH) and loss of heterozygosity (LOH) events in tumour whole genome sequencing (WGS) data. The model simultaneously estimates the cellular prevalence, proportion of tumour sample containing the event, and clonal cluster memberships. The statistical framework was designed based on the principles that the observed sequencing signal is an aggregated measure of multiple heterogeneous cellular populations including normal and tumour subpopulations and sets of genetic aberrations are observed at similar cellular prevalence if these events co-occurred in the same clone due to punctuated clonal expansions [1]. The TITAN software is part of an accompanying manuscript where mathematical details are presented [3].

The observed measurements are modeled as generated from a composite of three cell populations [4], consisting normal cells at global proportion n ; tumour cells with normal genotype at proportion $(1 - n) * s_z$; and tumour cells with aberrant genotype at proportion $(1 - n) * (1 - s_z)$. s_z is the proportion of tumour cells that is diploid heterozygous (and therefore normal) at the locus. Thus, $(1 - s_z)$ is the proportion of

the tumour population containing the event, or we define as the cellular prevalence. We assume multiple somatic events share similar cellular prevalence and thus can be assigned to one of a finite number of clonal clusters, $z \in Z$.

The input data to TITAN are tumour read depth and allele counts. The read depth is corrected for GC content and mappability biases using the Bioconductor R package *HMMcopy*. Then, the log ratio between tumour and normal corrected depths is computed. The log ratio is model using a mixture of Gaussians where the mean parameter is a function of the cellular prevalence, normal proportion, and tumour ploidy; the variance is also unknown and estimated. The allele counts are transformed to allelic ratios, which is defined as the reference read count divided by depth or proportion of reference reads at a locus. Allelic ratios are modeled using a binomial distribution [2] with the success parameter being a function of cellular prevalence and normal contamination.

TITAN is implemented as a two-factor hidden Markov model (HMM) where the hidden genotypes and the hidden clonal cluster memberships are the two Markov chains. The state space is expanded as the joint genotype and clonal cluster states. The output of TITAN is a list of CNA and LOH events, with accompanying parameter estimates of normal proportion, the cellular prevalence and clonal population cluster membership for each event.

TITAN estimates the cellular prevalence (proportion of tumour cells with the tumour genotype) based on a fixed number of clonal clusters. Users are required to run TITAN on a tumour sample for a range of clonal clusters in parallel. We advise users to use a range of 1 to 5. When results for the 5 jobs are generated, one for each fixed number of clusters, each parameter output file will contain an S_Dbw Validity Index score. The run for a tumour sample that has the lowest S_Dbw score is the optimal result and the corresponding number of clonal clusters is the optimal one.

This vignette will detail an example run of TITAN for only 2 clonal clusters of chromosome 2 of a triple negative breast cancer sample [2, 3]. Once again, for real data, users should repeat the *TITAN* for additional clonal clusters and then select the optimal run using the minimum S_Dbw Validity index.

2 Analysis Workflow Overview

1. The data points of interest are the germline heterozygous SNP loci identified in the matched normal WGS sample. This is a pre-processing step that is completed by the user, and is outside of the scope of this R package.
2. At these loci, the read depth, reference read count and non-reference read count are extracted from the tumour WGS sample. This step is completed outside of the R package
3. The read depth are normalized for GC content and mappability biases using the *HMMcopy* R package. *TITAN* uses a wrapper for the function in *HMMcopy* to accomplish this.
4. *TITAN* requires as input, the read counts and normalized read depth to estimate model parameters and infer subclonal CNA/LOH.
5. For real data, repeat *TITAN* analysis for a range of clonal clusters.

3 Model training and parameter estimation

TITAN uses the Expectation-Maximization algorithm to train the model and estimate parameters for cellular prevalence, normal proportion and tumour ploidy. In the E-step, the forwards-backwards algorithm is employed, making a call to a C implementation for faster computation. In the M-step, parameters are estimated using maximum a posteriori (MAP). The various parameters

TITAN uses the Viterbi algorithm to find the optimal state sequence path of genotype and clonal cluster membership for each data point.

4 Extracting tumour and normal read depth

TITAN requires the read depth from the tumour and normal WGS samples. This is done using a tool, outside of R, distributed as part of the *HMMcopy Suite* available at <http://compbio.bccrc.ca/software/hmmcopy/>.

In short, the suite has tools to:

- Obtain high resolution bin counts for large (≈ 250 GB) BAM files within a few hours.
- Obtain GC content for bins from standard FASTA files within minutes for a human genome.
- Obtain average mappability for bins from BigWig within minutes, or FASTA files within a day for a human genome.

For TITAN, the user will need to generate the read depth files for the tumour and normal samples in the form of WIG files. Also required, and is provided in the *HMMcopy suite* for the GRCh37-lite reference genome, are the GC content and mappability scores WIG files. Please refer to the instructions on the *HMMcopy suite* website for preparing these files.

5 TITAN analysis

Users should run TITAN once for each setting of the number of clonal clusters, ranging from 1 to 5. Although higher number of clonal clusters can be used, in practice, the sequencing coverage of 30-50X, analyzing more than 5 clonal clusters may not produce accurate results. For each run of TITAN on real data with approximately 2 million loci across all chromosomes and for 5 clonal clusters, the maximum memory requirement is approximately 24Gb.

5.1 Loading the data

Load the provided text file for the allele counts for chromosome 2 of a breast cancer sample [2]. The format of the file must be 6 columns: chromosome, position, reference base, reference read counts, non-reference base, non-reference read counts. A list object with 6 components of equal size/lengths is returned.

```
> library(TitanCNA)
> infile <- system.file("extdata", "test_alleleCounts_chr2.txt",
+                       package = "TitanCNA")
> data <- loadAlleleCountsFromFile(infile)
> names(data)

[1] "chr"           "posn"          "ref"           "refOriginal"  "nonRef"
[6] "tumDepth"     "normDepth"
```

5.2 Loading the initial parameters

Load the default parameters using a specified maximum copy number *TITAN* will consider. Here, 5 is used but up to 8 copies can be specified, however, running time and memory performance hits will be incurred. Here, you will also specify the number of clonal clusters to use for the TITAN run. This is where the number of clusters will be specified when running for 1,2,3,4, or 5 clonal cluster for real data. In general, the default parameters will work well with real, full datasets of whole genome sequencing of tumour samples.

```
> numClusters <- 2
> params <- loadDefaultParameters(copyNumber = 5,
+                                numberClonalClusters = numClusters)
> params
```

```

$genotypeParams
$genotypeParams$rt
[1] 0.5000000 1.0000000 1.0000000 0.5000000 1.0000000 0.6666667 1.0000000
[8] 0.7500000 0.5000000 1.0000000 0.8000000 0.6000000

$genotypeParams$rn
[1] 0.5

$genotypeParams$ZS
[1] 0 1 2 3 4 5 6 7 8 9 10 11

$genotypeParams$ct
[1] 0 1 2 2 3 3 4 4 4 5 5 5

$genotypeParams$var_0
[1] 0.05 0.05 0.05 0.05 0.05 0.05 0.05 0.05 0.05 0.05 0.05 0.05 0.05

$genotypeParams$alphaKHyper
[1] 5000 5000 5000 5000 5000 5000 5000 5000 5000 5000 5000 5000 15000 15000

$genotypeParams$betaKHyper
[1] 25 25 25 25 25 25 25 25 25 25 25 25 25 25

$genotypeParams$kappaGHyper
[1] 2 2 2 2 5 2 2 2 2 2 2 2 2

$genotypeParams$outlierVar
[1] 10000

$ploidyParams
$ploidyParams$phi_0
[1] 2

$ploidyParams$alphaPHyper
[1] 20

$ploidyParams$betaPHyper
[1] 42

$normalParams
$normalParams$n_0
[1] 0.5

$normalParams$alphaNHyper
[1] 2

$normalParams$betaNHyper
[1] 2

```

```
$cellPrevParams
$cellPrevParams$s_0
[1] 0.001 0.200
```

```
$cellPrevParams$alphaSHyper
[1] 2 2
```

```
$cellPrevParams$betaSHyper
[1] 2 2
```

```
$cellPrevParams$kappaZHyper
[1] 2 2
```

`params` is a list object containing 4 sets of parameters, each as a component:

- `genotypeParams`: Parameters for copy number and allelic ratios genotype states
- `normalParams`: Parameters for normal contamination
- `ploidyParams`: Parameters for average tumour ploidy
- `sParams`: Parameters for modeling subclonality: clonal clusters and cellular prevalence

Analysis of copy number by *TITAN* is subject to tumour ploidy. *TITAN* attempts to estimate this value globally and is accurate in most cases. However, the issue is that, from the signals observed in the data, *TITAN* cannot distinguish between diploid (2 global copies) and tetraploidy (4 global copies) because the data can explain both situations. Therefore, in general, if the user is aware that a sample is ployploid, then *TITAN* should be run with the ploidy initialization of 2 and also for initialization of 4. If the user is unaware of the ploidy status of the sample, then inspection of a *TITAN* run at ploidy initialization of 2 can help. If this run has many large, prominent regions of inferred homozygous deletions, then it is likely that this sample is triploid or tetraploid.

```
> params$ploidyParams$phi_0 <- 2 # for diploid or
> params$ploidyParams$phi_0 <- 4 # for tetraploid/ployploid
```

Because this vignette is an example involving only one chromosome, the fewer loci in the analysis will give different results than for a sample with all chromosomes. For this example, we will modify the Gaussian variance hyperparameter (prior used in modeling the log ratios) to a less influential setting, making the analysis more suitable for one chromosome. Also, because we know the average tumour ploidy of the genome-wide sample is slightly less than 2, we will initialize the `phi_0` to 1.5 so that it does not overfit to this one chromosome.

```
> K <- length(params$genotypeParams$alphaKHyper)
> params$genotypeParams$alphaKHyper <- rep(500, K)
> params$ploidyParams$phi_0 <- 1.5
```

5.3 Correct GC content and mappability biases

Correct GC content and mappability biases using a wrapper function that calls *HMMcopy* R-specific routines. Here, 4 wig files are required: tumour, normal, GC content, and mappability score. The wrapper will apply the correction sub-routines and compute the log ratio as $\log_2(\text{tumour}/\text{normal})$.

```

> tumWig <- system.file("extdata", "test_tum_chr2.wig", package = "TitanCNA")
> normWig <- system.file("extdata", "test_norm_chr2.wig", package = "TitanCNA")
> gc <- system.file("extdata", "gc_chr2.wig", package = "TitanCNA")
> map <- system.file("extdata", "map_chr2.wig", package = "TitanCNA")
> cnData <- correctReadDepth(tumWig, normWig, gc, map)
> head(cnData)

```

```

  chr start  end      logR
1   2     1 1000  1.23017573
2   2    1001 2000 -0.59165317
3   2    2001 3000 -0.11715274
4   2    3001 4000 -0.11250534
5   2    4001 5000  0.10922812
6   2    5001 6000  0.03968276

```

For real data, users can use the genome-wide GC content and mappability score for 1kb windows. The two files can be found compressed in `data/GRCh37-lite.tar.gz` of the git repository. These WIG files were generated from the reference `GRCh37-lite.fasta`; the tumour and normal BAM files used for generating the tumour and normal WIG files in the preprocessing steps MUST also use this same reference. Otherwise, new GC and map wig files will need to be created as per instructions on <http://compbio.bccrc.ca/software/hmmcopy/>

Then, find the log ratio at each position of interest (germline heterozygous SNPs) given in the object "data". Then transform the log ratios to natural logs instead of log base 2. Remove "logR" and "cnData" objects to save memory.

```

> logR <- getPositionOverlap(data$chr, data$posn, cnData)
> data$logR <- log(2^logR) #transform the log ratio to natural logs

```

5.4 Filter the data

Filter the data for low and high depth positions and positions with NA's. For ease of analysis, TITAN converts chromosome X->23 and Y->24. Do not worry, in the final output, "X" and "Y" will be in the output. Optionally, a data.frame containing a list of positions can be passed as an argument, specifying which positions to use.

```

> data <- filterData(data, 1:24, minDepth = 10, maxDepth = 200,
+                   positionList = NULL)

```

5.5 Model training and parameter estimation

Parameter estimation in *TITAN* uses the Expectation Maximization algorithm (EM) and forwards-backwards algorithm. This step generally requires parallelization to increase time performance. The *TITAN* package uses the *foreach* package which conveniently enables the user to use libraries, such as *doMC* and *doMPI*, to parallelize the training by chromosome. It is up to the user to decide which is the best library to use depending on the current system. For example, if one wishes to use multiple cores via forking on a single machine,

```

> library(doMC)
> registerDoMC(cores = 10) #use 10 cores on a single machine

```

Here is a little more detail regarding some important arguments:

- **maxiter**: maximum number of EM iterations allowed. In practice, users should set it at 20 since it does not exceed 20. For this vignette to finish in a shorter time, we use 3.

- `maxiterUpdate`: maximum number of coordinate descent iterations during the M-step when parameters are estimated. In practice, 1500 iterations should be used.
- `txnExpLen`: influences prior probability of genotype transitions in the HMM. The smaller, the lower tendency to change state; however, too small and it produces underflow problems.
- `txnZstrength`: influences prior probability of clonal cluster transitions in the HMM. Smaller values means lower tendency to change clonal cluster state.
- `useOutlierState`: logical indicating whether an additional outlier state should be used. In practice, this usually is not necessary.
- `normalEstimateMethod`: specifies how to handle normal proportion estimation. To estimate, use "map" which is maximum a posteriori. If you wish to not estimate this parameter, then use "fixed". This will default the normal proportion to whatever is specified in `params$normalParams$n_0`. For example, if you know that this sample has absolutely no normal contamination, then use `params$normalParams$n_0 <- 0` and set `normalEstimateMethod="fixed"`.
- `estimateS`: logical indicating whether to account for clonality and estimate subclonal events
- `estimatePloidy`: logical indicating whether to estimate and account for tumour ploidy

```
> convergeParams <- runEMclonalCN(data, gParams = params$genotypeParams,
+                               nParams = params$normalParams,
+                               pParams = params$ploidyParams,
+                               sParams = params$cellPrevParams,
+                               maxiter = 3, maxiterUpdate = 50,
+                               useOutlierState = FALSE, txnExpLen = 1e9,
+                               txnZstrength = 1e9,
+                               normalEstimateMethod = "map",
+                               estimateS = TRUE, estimatePloidy = TRUE)
> names(convergeParams)

[1] "n"           "s"           "var"         "phi"
[5] "piG"        "piZ"        "muR"        "muC"
[9] "loglik"     "rhoG"       "rhoZ"       "txnExpLen"
[13] "txnZstrength" "useOutlierState" "genotypeParams" "ploidyParams"
[17] "normalParams" "clonalParams"
```

`convergeParams` is a list object with components containing the converged parameters from the EM training, including posterior marginal responsibilities, log likelihood, and original parameter settings.

- `n`: Converged estimate for normal contamination parameter. `numeric array` containing estimates at each EM iteration.
- `s`: Converged estimate(s) for cellular prevalence parameter(s). This value is defined as the proportion of tumour sample that does *not* contain the aberrant genotype. This will contrast what is output in `outputTitanResults`. `numeric array` containing estimates at each EM iteration. If more than one cluster is specified, then `s` is a `numeric matrix`.
- `var`: Converged estimates for variance parameter of the Gaussian mixtures used to model the log ratio data. `numeric matrix` containing estimates at each EM iteration.
- `phi`: Converged estimate for tumour ploidy parameter. `numeric array` containing estimates at each EM iteration.

- **piG**: Converged estimate for initial genotype state distribution. **numeric matrix** containing estimates at each EM iteration.
- **piZ**: Converged estimate for initial clonal cluster state distribution. **numeric matrix** containing estimates at each EM iteration.
- **muR**: Mean of binomial mixtures computed as a function of **s** and **n**. **numeric matrix** containing estimates at each EM iteration. See References for mathematical details.
- **muC**: Mean of Gaussian mixtures computed as a function of **s**, **n**, and **phi**. **numeric matrix** containing estimates at each EM iteration. See References for mathematical details.
- **loglik**: Posterior Log-likelihood that includes data likelihood and the priors. **numeric array** containing estimates at each EM iteration.

5.6 Find the genotype and clonal cluster state paths

Viterbi algorithm is used to return the optimal genotype/clonal cluster state path. After running EM, use the converge parameters and the input data to compute the optimal segmentation results.

```
> optimalPath <- viterbiClonalCN(data, convergeParams)
> head(optimalPath)

[1] 14 14 14 14 14 14
```

It is difficult to interpret the output of this function directly. The user should use the function `outputTitanResults` to format the results.

5.7 Format and print results to file

Two files are generated:

Position-specific results that includes genotype, clonal cluster, and cellular prevalence estimates. The posterior probabilities at each position can optionally be returned. The results can be output to a file if a file name is specified for the `filename` argument.

```
> results <- outputTitanResults(data, convergeParams, optimalPath,
+                               filename = NULL, posteriorProbs = FALSE)
> head(results)
```

Chr	Position	RefCount	NRefCount	Depth	AllelicRatio	LogRatio	CopyNumber	
1	2	55802	13	12	25	0.52	0.13	1
2	2	56232	3	7	10	0.30	0.02	1
3	2	56619	7	6	13	0.54	0.02	1
4	2	57093	10	2	12	0.83	-0.39	1
5	2	60324	6	10	16	0.38	-0.04	1
6	2	60906	6	5	11	0.55	-0.04	1

TITANstate	TITANcall	ClonalCluster	CellularPrevalence	
1	1	DLOH	2	0.53
2	1	DLOH	2	0.53
3	1	DLOH	2	0.53
4	1	DLOH	2	0.53
5	1	DLOH	2	0.53
6	1	DLOH	2	0.53

The format of this output file and data.frame is the following:

1. **Chr**
2. **Position**
3. **RefCount**: number of reads matching the reference base
4. **NRefCount**: number of reads matching the non-reference base
5. **Depth**: total read depth at the position
6. **AllelicRatio**: RefCount/Depth
7. **LogRatio**: log₂ ratio between normalized tumour and normal read depths
8. **CopyNumber**: predicted TITAN copy number
9. **TITANstate**: internal state number used by TITAN; see supplementary table 2 in manuscript
10. **TITANcall**: interpretable TITAN state; string {HOMD,DLOH,HET,NLOH,ALOH,ASCNA,BCNA,UBCNA}, see supplementary table 2 in manuscript
11. **ClonalCluster**: predicted TITAN clonal cluster; lower cluster numbers represent clusters with higher cellular prevalence
12. **CellularPrevalence**: proportion of tumour cells containing event; not to be mistaken as proportion of sample (including normal)

Model parameters and summary values such as the number of clonal clusters and their corresponding cellular prevalence, normal contamination, ploidy, and the S_Dbw validity index for model selection later.

```
> outparam <- paste("test_cluster02_params.txt", sep = "")
> outputModelParameters(convergeParams, results, outparam)
```

The format of the parameter output file is the following:

1. Normal contamination estimate: proportion of normal content in the sample; tumour content is 1 minus this number
2. Average tumour ploidy estimate: average number of estimated copies in the genome; 2 represents diploid
3. Clonal cluster cellular prevalence: Z denotes the number of clonal clusters; each value (space-delimited) following are the cellular prevalence estimates for each cluster
4. Genotype binomial means for clonal cluster Z : set of 21 binomial estimated parameters for each specified cluster
5. Genotype Gaussian means for clonal cluster Z : set of 21 Gaussian estimated means for each specified cluster
6. Genotype Gaussian variance: set of 21 Gaussian estimated variances; variances are shared for across all clusters
7. Number of iterations: number of EM iterations needed for convergence
8. Log likelihood: complete data log-likelihood for current cluster run
9. S_Dbw dens.bw: density component of S_Dbw index
10. S_Dbw scat: scatter component of S_Dbw index
11. S_Dbw validity index: used for model selection; choose run with optimal number of clusters based on lowest S_Dbw index

5.8 Plotting the results

Plot the results using 3 built in functions. For each chromosome, there is a separate plot function for the log ratios (CNA), allelic ratios (LOH), and cellular prevalence. Each function adds to an existing plot.


5.8.1 Copy number alterations (log ratio)

Generate the figure of copy number alteration results by plotting the log ratios. Optionally, the data can be corrected by using the estimated ploidy. Otherwise, use `ploidy=NULL`.

```
> ploidy <- tail(convergeParams$phi, 1)
> ploidy

[1] 1.504029

> plotCNlogRByChr(results, chr = 2, ploidy = ploidy, ylim = c(-2, 2), cex = 0.25,
+                 xlab = "", main = "Chr 2")
```




The Y-axis is based on log ratios. Log ratios are computed ratios between normalized tumour and normal read depths. Data points close to 0 represent diploid, above 0 are copy gains, below 0 are deletions. Bright Green - HOMD Green - DLOH Blue - HET, NLOH Dark Red - GAIN Red - ASCNA, UBCNA, BCNA

5.8.2 Loss of heterozygosity (allelic ratio)

The loss of heterozygosity (LOH) results is shown by plotting the allelic ratio.

```
> plotAllelicRatio(results, chr = 2, ylim = c(0, 1), cex = 0.25,
+                  xlab = "", main = "Chr 2")
```



The Y-axis is based on allelic ratios. Allelic ratios are computed as RefCount/Depth. Data points close to 1 represent homozygous reference base, close to 0 represent homozygous non-reference base, and close to 0.5 represent heterozygous. Normal contamination influences the divergence away from 0.5 for LOH events. Grey - HET, BCNA Bright Green - HOMD Green - DLOH, ALOH Blue - NLOH Dark Red - GAIN Red - ASCNA, UBCNA

5.8.3 Cellular prevalence and clonal clusters

One of the key features of *TITAN* is the estimation of cellular prevalence and the inference of clonal cluster membership. The estimated normal proportion is required for this plot. This can be obtained from `convergeParams`.


```
> norm <- tail(convergeParams$n, 1)
> norm # estimated normal contamination

[1] 0.1824862

> 1 - convergeParams$s[, ncol(convergeParams$s)] # estimated cellular prevalence

[1] 0.9611795 0.5257538

> plotClonalFrequency(results, chr = 2, normal = norm, ylim = c(0, 1),
+                       cex = 0.25, xlab = "", main = "Chr 2")
```



The Y-axis is the cellular prevalence that includes the normal proportion. Therefore, the cellular prevalence here refers to the proportion in the sample (including normal). Lines are drawn for each data point indicating the cellular prevalence. Heterozygous diploid are not shown because it is a normal genotype and is not categorized as being subclonal (this means 100% of cells are normal). The black horizontal line represents the tumour content labeled as "T". Each horizontal grey line represents the cellular prevalence of the clonal clusters labeled as Z1, Z2, etc. Colours are the same for allelic ratio plots.

6 Session Information

The version number of R and packages loaded for generating the vignette

- R version 3.1.0 (2014-04-10), x86_64-unknown-linux-gnu
- Locale: LC_CTYPE=en_US.UTF-8, LC_NUMERIC=C, LC_TIME=en_US.UTF-8, LC_COLLATE=C, LC_MONETARY=en_US.UTF-8, LC_MESSAGES=en_US.UTF-8, LC_PAPER=en_US.UTF-8, LC_NAME=C, LC_ADDRESS=C, LC_TELEPHONE=C, LC_MEASUREMENT=en_US.UTF-8, LC_IDENTIFICATION=C

- Base packages: base, datasets, grDevices, graphics, methods, parallel, stats, utils
- Other packages: AnnotationDbi 1.26.0, Biobase 2.24.0, BiocGenerics 0.10.0, GenomeInfoDb 1.0.0, HMMcopy 1.6.0, IRanges 1.21.45, TitanCNA 1.0.0, annotate 1.42.0, foreach 1.4.2, geneplotter 1.42.0, lattice 0.20-29
- Loaded via a namespace (and not attached): DBI 0.2-7, RColorBrewer 1.0-5, RSQLite 0.11.4, XML 3.98-1.1, codetools 0.2-8, compiler 3.1.0, grid 3.1.0, iterators 1.0.7, stats4 3.1.0, tools 3.1.0, xtable 1.7-3

References

- [1] Mel Greaves and Carlo C Maley. Clonal evolution in cancer. *Nature*, 481(7381):306–313, Jan 2012.
- [2] Gavin Ha, Andrew Roth, Daniel Lai, Ali Bashashati, Jiarui Ding, Rodrigo Goya, Ryan Giuliany, Jamie Rosner, Arusha Oloumi, Karey Shumansky, Suet-Feung Chin, Gulisa Turashvili, Martin Hirst, Carlos Caldas, Marco A. Marra, Samuel Aparicio, and Sohrab P. Shah. Integrative analysis of genome-wide loss of heterozygosity and monoallelic expression at nucleotide resolution reveals disrupted pathways in triple-negative breast cancer. *Genome research*, 22(10):1995–2007, Oct 2012.
- [3] Gavin Ha, Andrew Roth, Nataliya Melnyk, Andrew McPherson, Leah M Prentice, Ali Bashashati, Jiarui Ding, Alan Le, Jamie Rosner, Karey Shumansky, Marco A Marra, David G Huntsman, Jessica N McAlpine, Samuel A J R Aparicio, and Sohrab P Shah. Modeling the evolution of copy number architecture in tumour whole genome sequencing data. *Under Review*, 2013.
- [4] C Yau, D Mouradov, R N Jorissen, S Colella, G Mirza, G Steers, A Harris, J Ragoussis, O Sieber, and C C Holmes. A statistical approach for detecting genomic aberrations in heterogeneous tumor samples from single nucleotide polymorphism genotyping data. *Genome Biol*, 11(9), 2010.