



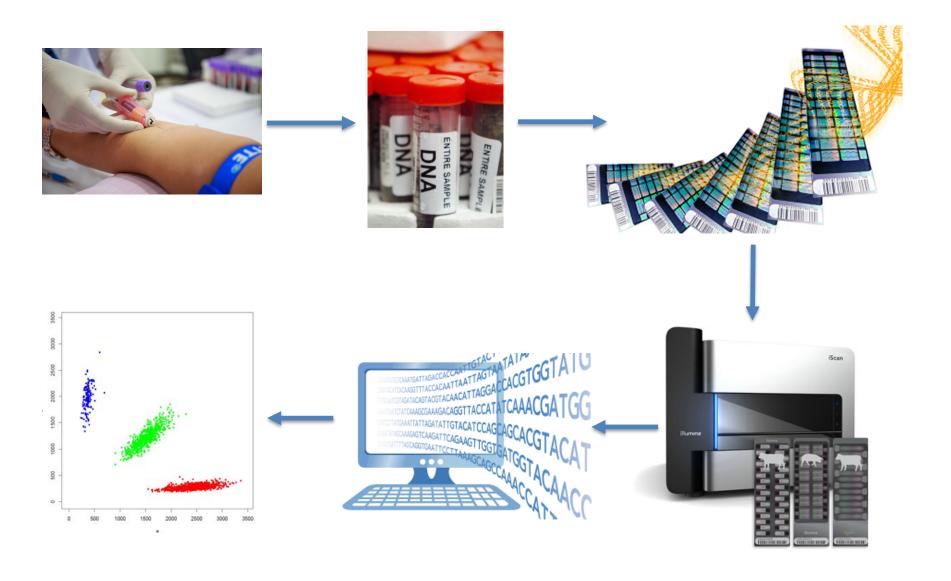
GWAS QC -theory and steps

H3ABioNet Genome Wide Association Studies Lecture Series

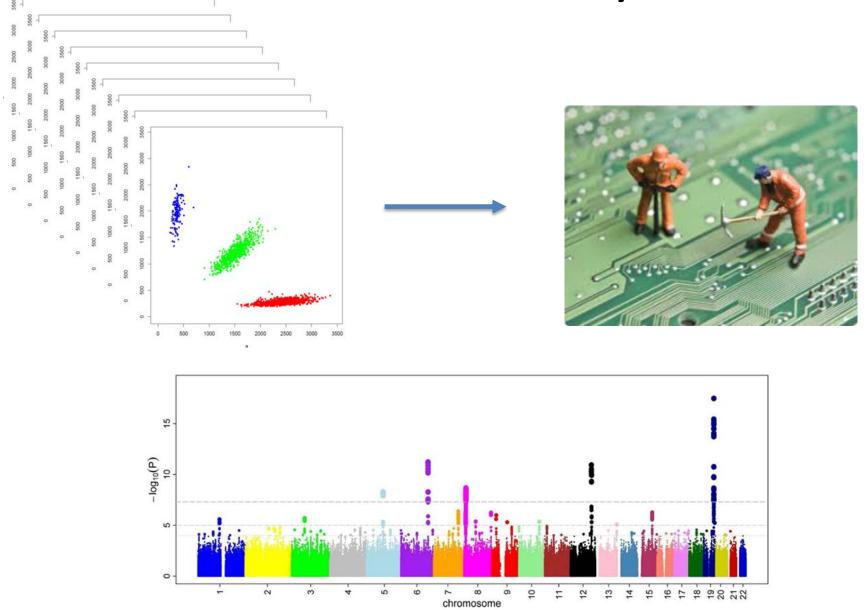
Shaun Aron

August 2018

GWAS Data Generation



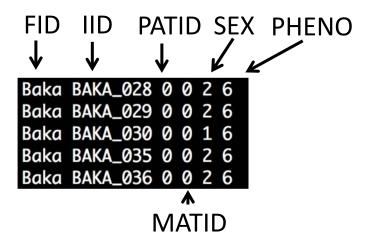
GWAS Data Analysis

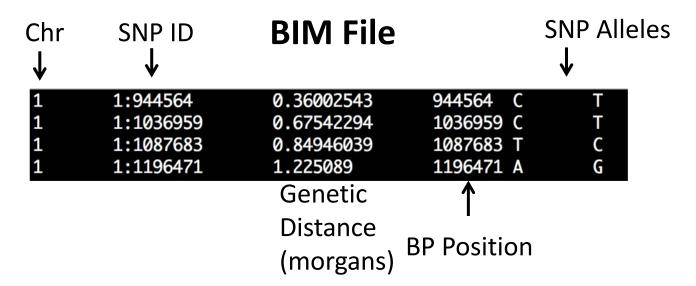


Plink format

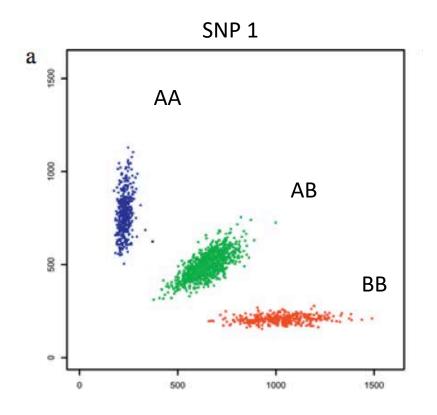
- FAM file one row per individual
- BIM file one row per SNP
- BED file one row per individual genotype calls for each individual for all SNPs – binary format
- FAM and BIM file are human readable while BED file in not

FAM File





Why Do We Need Quality Control?



In an ideal world...

our sampling practices would be perfect,

our experiments would run perfectly,

and all our SNP genotypes would look like this.

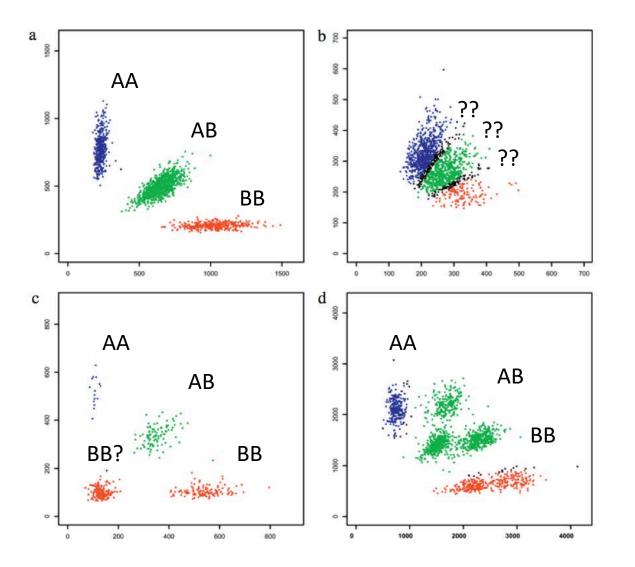
Why Do We Need Quality Control

- Large scale experiments generate true results with a certain error rate
- Errors might originate at various steps in in the process:
 - ✓ Sample selection related issues
 - ✓ Cryptic relatedness
 - ✓ Population structure
 - ✓ Sample handling related issues
 - ✓ Labeling/Plating Error
 - ✓ Genotyping array related issues
 - ✓ Genotyping error
 - ✓ Batch effect related issues
 - ✓ Difference in results due to difference in sample processing

- Not practical to visually assess the genotype plot for every SNP
- Use some biologically relevant metrics as a proxy for quality
- Steps
 - QC by SNP
 - QC by sample

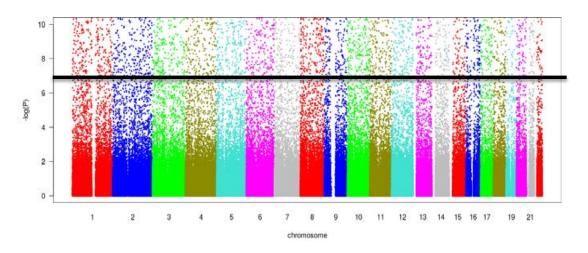


Why Do We Need Quality Control



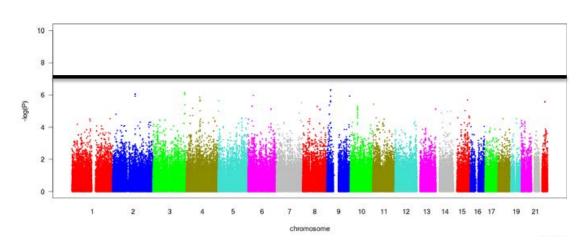
We don't live in an ideal world...

German MI family study Affymetrix 500K Array Set SNPs on chips: 493,840





SNPs passing QC: 270,701



Samani et al. 2007 N Engl J Med 357:443-53

QC Roadmap

Sample QC

Discordant sex information
High Missingness
Excess or deficiency of
heterozygosity
Duplicate or related
Divergent ancestry
Batch Effects



SNP QC

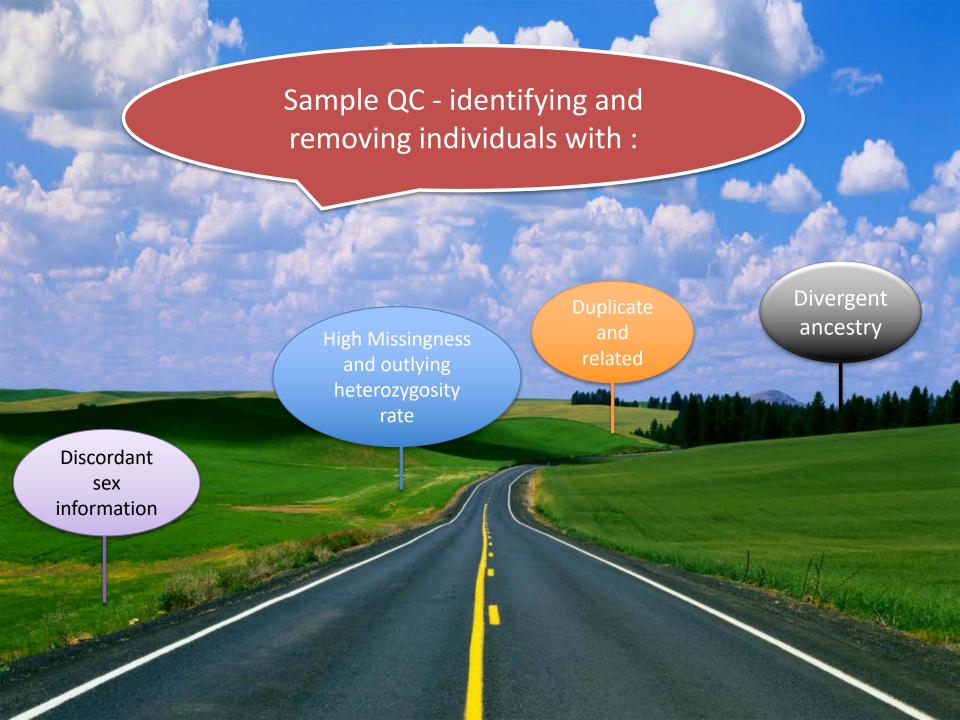
Low minor allele frequency Missingness Differential missingness Hardy-Weinberg outliers

- For the sake of simplicity I have split the QC into Sample QC and SNP QC
- The order explained here is therefore not always the best way to do they QC
- The H3ABioNet pipeline has been developed using the best generalized order

Software

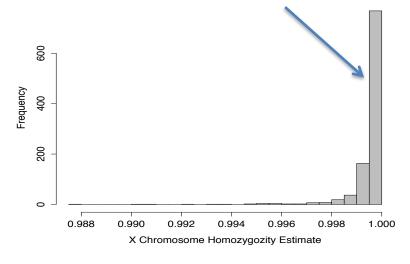
Programs required for QC

- PLINK 1.9 (Purcell, 2007)
- Scripts for processing result files
- R (Statistical Software) for plotting results
- The complete process has been built into a pipeline using NextFlow



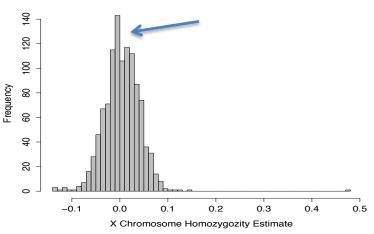
Discordant Sex Check

- Males have a single X chromosome and therefore can be estimated to be homozygous for all the X chromosome SNPs (other than those in the pseudo autosomal region(PAR)).
- Therefore, X chromosome homozygosity estimate for males (XHE) is 1
- Plink assigns sex based on XHE estimate (F or inbreeding coefficient):
- Male (1): XHE >0.80
- Female (2): XHE < 0.20
- No sex (0): 0.20 < XHE < 0.80
- Comparisons of predicted and observed sex can be used to identify miscoded sex or sample mix-ups, etc.



All Male Samples





Identify individuals with discordant sex information

plink --bfile example --check-sex --out sexstat

Creates a file named sexstat.sexcheck

FID	IID	PEDSEX		SNPSEX		STATUS	F	
P554	P554		2		2	OK	-0.02654	
P555	P555		1		0	PROBLEM	0.5685	
P557	P557		2		2	OK	0.1264	
P558	P558		2		2	OK	-0.0007684	

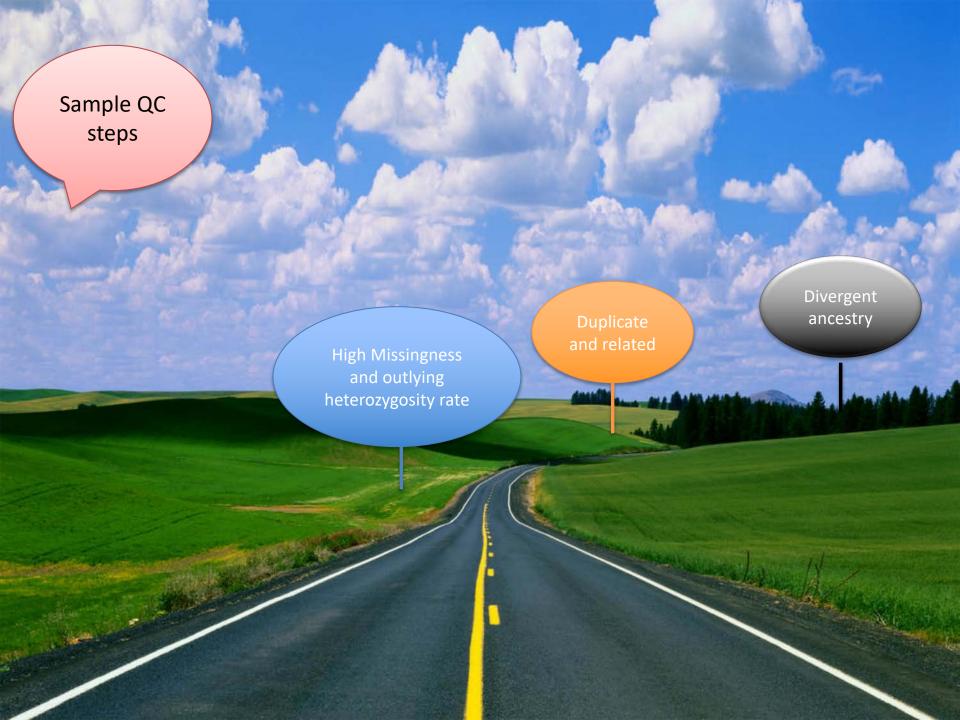
Select individuals with Status="PROBLEM" in the file sexstat.sexcheck

Try to identify the problem. If the problem cannot be resolved write the IDs of the individuals with discordant sex information to a file "fail_sex_check.txt"

grep "PROBLEM" sexstat.sexcheck > fail_sex_check.txt

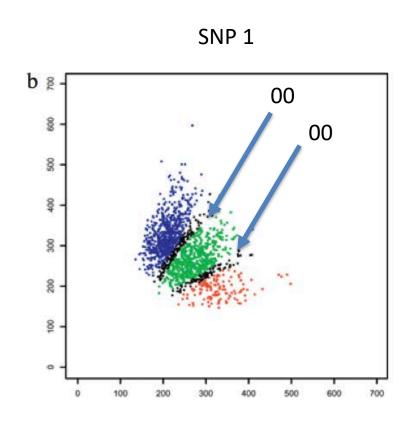
nextflow Parameters

- sex_info_available = true
- f low male = 0.8
- f high female = 0.2



Missingness

- Unable to call a genotype for a particular SNP – will be called as missing
- Per sample missingness
 - Percentage of SNPs with missing data per sample
- Per SNP missingness
 - Percentage of missing calls for a SNP



Missingness and Heterozygosity Rate

Genotyping call rate

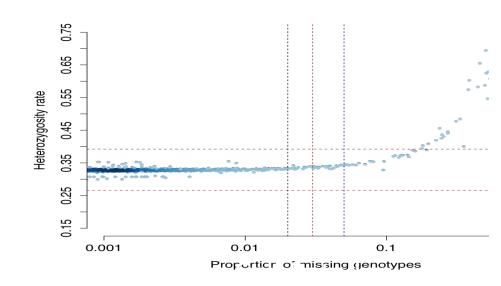
- Per sample (individual) rate
- Number of non-missing genotypes divided by the total number of genotyped markers.
- Low genotyping call rate indicate problem with sample DNA like low concentration.
- Thresholds used generally vary between 3% and 7%

Genotyping call rate and heterozygosity rate are generally plotted together.

Cutoffs are selected so as to identify outlier individuals based on both the statistics

Heterozygosity Rate

- Per sample (individual) rate
- Number of (total non-missing genotypes(N) homozygous(0)) genotypes divided by total nonmissing genotypes(N)
- Excess heterozygosity Possible sample contamination
- Less than expected heterozygosity Possibly inbreeding
- Threshold is usually to remove any individual with +- 3 standard deviations from the mean heterozygosity rate of all samples.

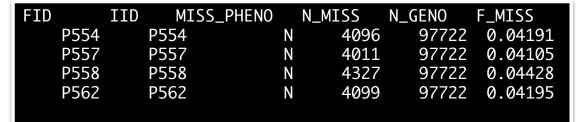


Missingness per Individual and per SNP

Missing phenotype (Y/NN)

Missingness per individual

plink --bfile example -missing -out example_miss





```
Before frequency and genotyping pruning, there are 98604 SNPs 646 founders and 0 non-founders found 34704 heterozygous haploid genotypes; set to missing Writing list of heterozygous haploid genotypes to [example_miss.hh] 3452 SNPs with no founder genotypes observed Warning, MAF set to 0 for these SNPs (see --nonfounders) Writing list of these SNPs to [example_miss.nof] Writing individual missingness information to [example_miss.imiss] Writing locus missingness information to [example_miss.lmiss]
```

Missingness per SNP (Come back to later)

HR	SNP	N_MISS	N_GENO	F_MISS
1	vh_1_1108138	9	646	0.01393
1	vh_1_1110294	4	646	0.006192
1	rs7515488	1	646	0.001548
1	rs6603785	9	646	0.01393

Heterozygosity rate per Individual

plink --bfile example --het --out example_het



Before frequency and genotyping pruning, there are 98604 SNPs 646 founders and 0 non-founders found

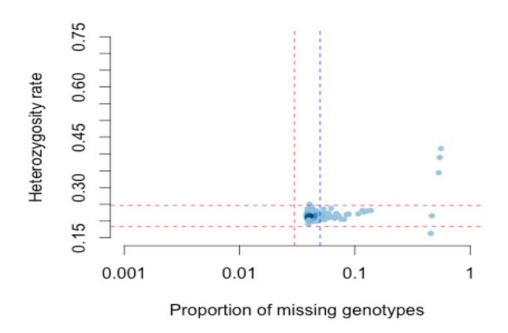
Detected that binary PED file is v1.00 SNP-major mode
Before frequency and genotyping pruning, there are 98604 SNPs
646 founders and 0 non-founders found
34704 heterozygous haploid genotypes; set to missing
Writing list of heterozygous haploid genotypes to [example_het.hh]
3452 SNPs with no founder genotypes observed
Warning, MAF set to 0 for these SNPs (see --nonfounders)
Writing individual heterozygosity information to [example_het.het]

Observed number of homozygous genotypes

Expected number of homozygous genotypes

Inbreeding coefficient estimate

FID	IID	O(HOM)	E(HOM)	N(NM)	F	
P554	P554	67663	6.725e+04	86305	0.02173	
P557	P557	66873	6.731e+04	86388	-0.02301	
P558	P558	67155	6.707e+04	86091	0.004538	
P562	P562	68367	6.724e+04	86306	0.05891	



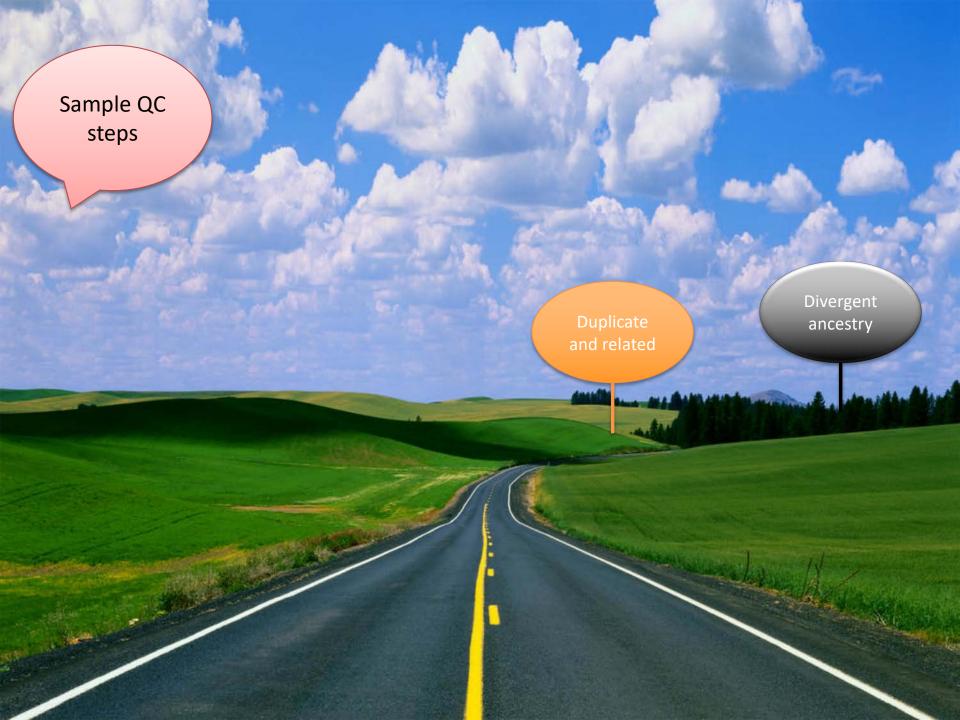
Based on the plot you can decide on reasonable thresholds at which to exclude individuals based on elevated missing or extreme heterozygosity.



We decided to exclude all individuals with a genotype failure rate ≥ 0.06 and heterozygosity rate ± 3 standard deviations from the mean heterozygosity rate of all samples

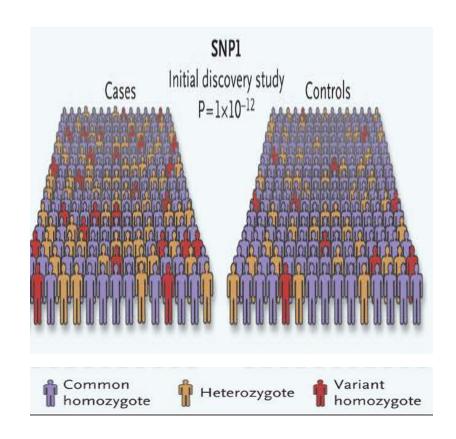
nextflow Parameters

- cut mind = 0.02
- cut het high = 0.343
- cut het low = 0.15

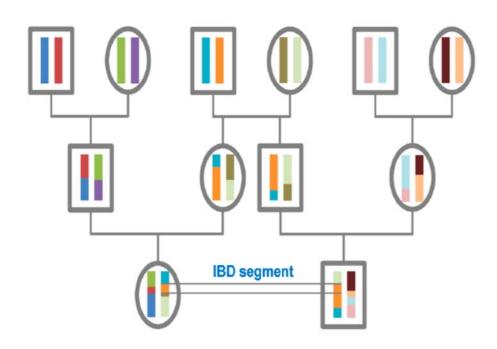


Identify related and duplicate individuals

- A basic assumption of standard population-based association studies is that all the samples are unrelated (i.e. the maximum relatedness between any pair of individuals is less than a second degree relative)
- Presence of duplicate and related individuals in the dataset may introduce bias and cause genotypes in families to be over-represented.
- To identify duplicate and related individuals, a metric (identity by state, IBS) is calculated for each pair of individuals based on the average proportion of alleles shared in common at genotyped SNPs (excluding the sex chromosomes)



- Identity by Descent (IBD) is a measure of the recent shared ancestry between two individuals based on genome wide IBS
- IBD calculations works best when only independent SNPs are included in the analysis.
- An independent SNP set for IBD calculations is generally prepared by removing regions of extended LD and pruning the remaining regions so that no pair of SNPs within a given window (say, 50kb) is correlated.
- IBD is calculated and denoted in Plink as Pi-hat
- Convention is to remove one individual from a pair with a Pi-hat > ~0.2



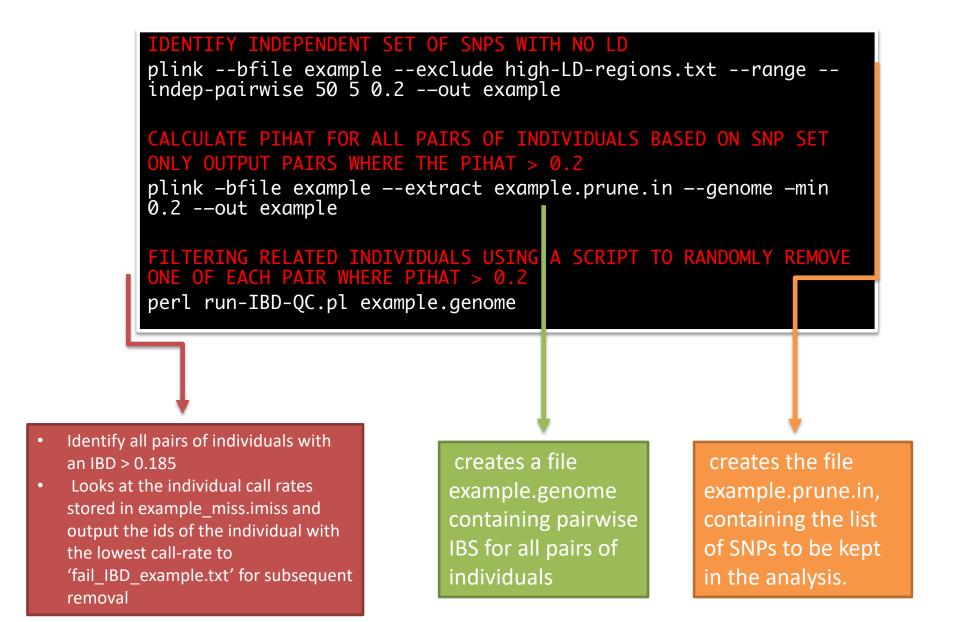
Pi – hat ~ 1 (Duplicate sample or Monozygotic twins)

Pi – hat ~ 0.5 (First degree relative)

Pi – hat ~ 0.25 (Second degree relative)

Pi – hat ~ 0.125 (Third degree relative)

Identification of duplicated or related individuals



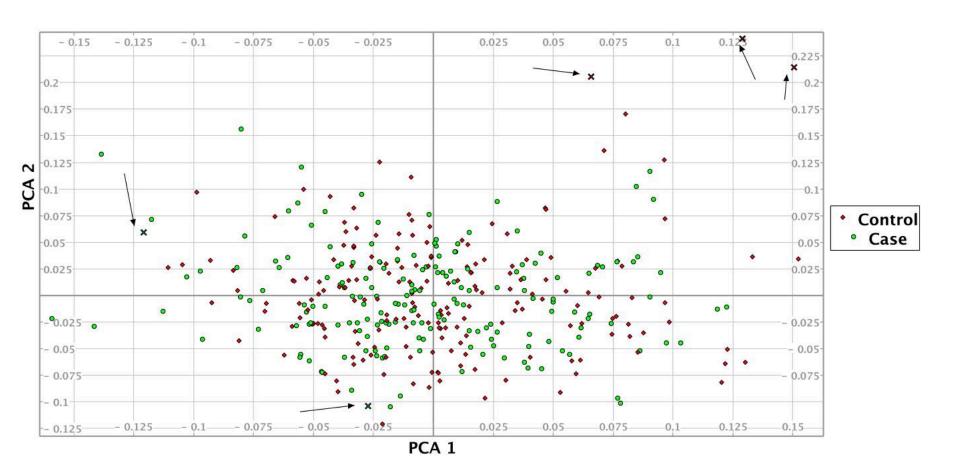
nextflow Parameters

- pi-hat = 0.11
- super_pi_hat = 0.7

Population structure

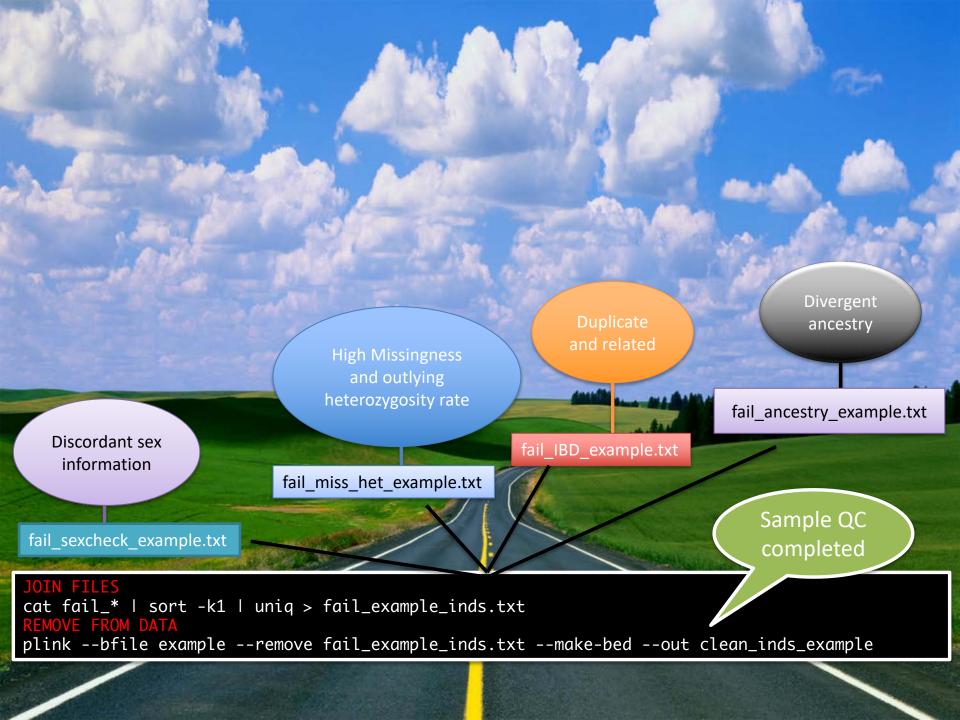
- Population substructure or stratification occurs when samples have different genetic ancestries
- Can lead to spurious associations due to differences in ancestry rather than true associations
- Imperative to check for population structure within samples
- Can control for structure if identified, in downstream analysis

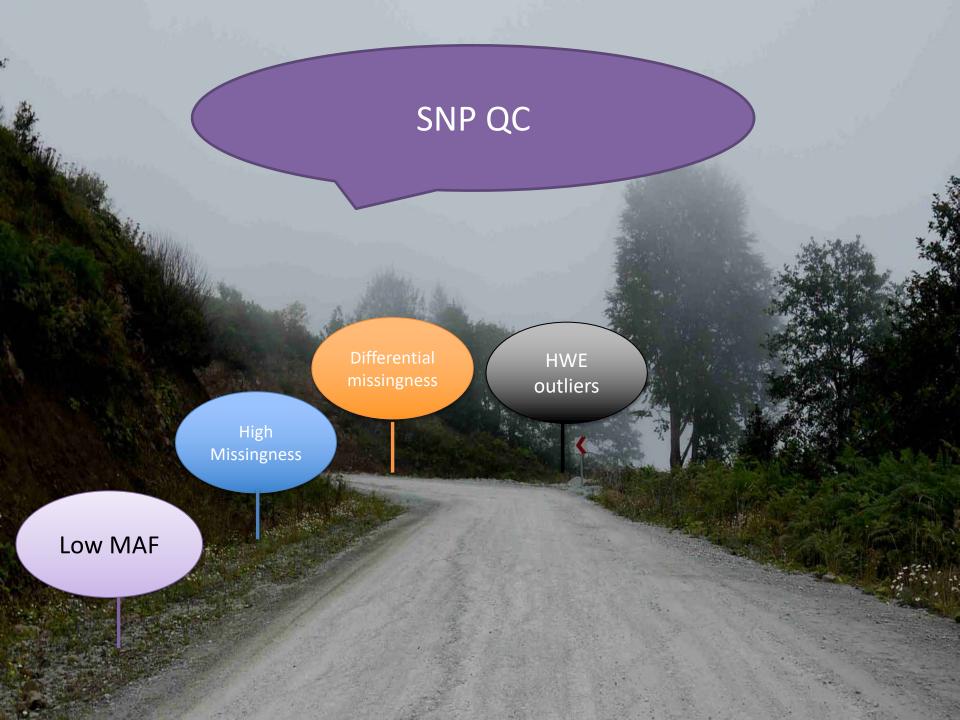
Population structure - PCA



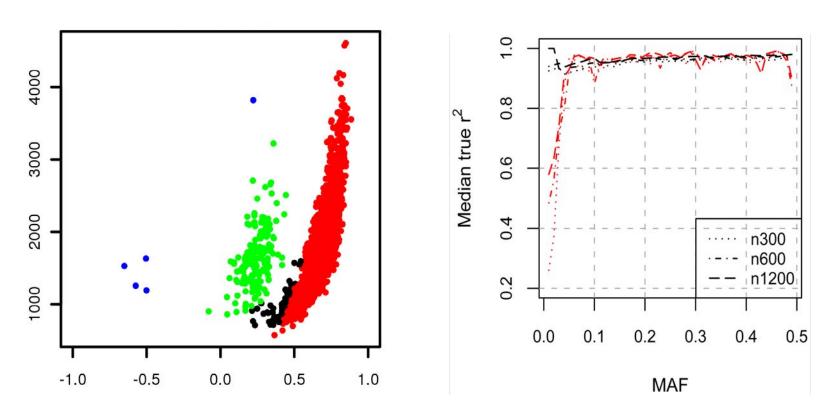
nextflow Parameters

- case control = sample.phe
- case control col = PHE
- batch = sample.phe
- batch_col = batch_no





Low minor allele frequency SNPs



- Genotype calling algorithms perform poorly for SNPs with low MAF and low samples sizes.
- Power for detecting associations with SNPs with low MAF is low unless the sample size is very large
- Commonly used exclusion threshold are SNPs with a MAF 0.01 to 0.05 (dependent on sample size)

Identify low minor allele frequency SNPs



plink --bfile clean-inds-example --freq --out
clean_inds_example_freq

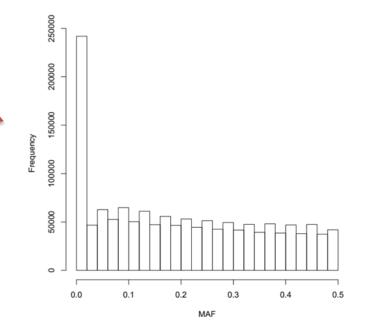
Generates the file "clean_inds_example_freq.frq" containing minor allele frequency of each SNP

GENERATE PLOT OF MAF DISTRIBUTION

PLOT MAF.R

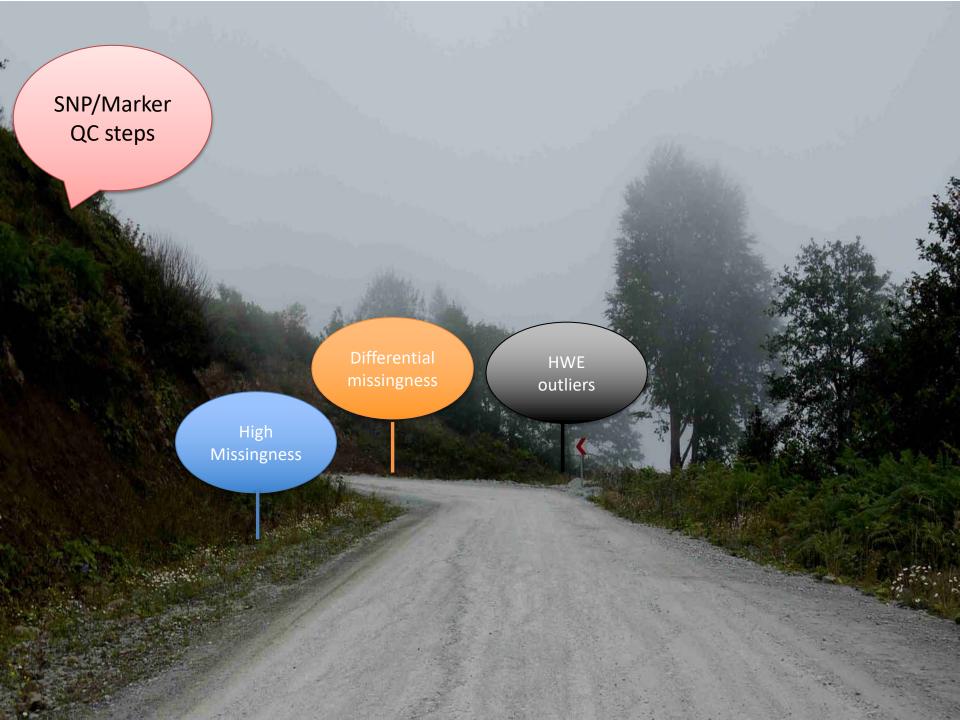
Choose standard MAF threshold (MAF > 0.01 OR MAF > 0.5) or base on distribution

CHR	SNP	A1	A2	MAF	NCHROBS
1	vh_1_1108138	A	G	0.4443	1292
1	vh_1_1110294	. A	G	0.362	1304
1	rs7515488	A	G	0.2893	1310
1	rs6603785	T	A	0.4985	1292
1	rs6603788	G	A	0.2221	1306
1	1_1209245	G	C	0.04609	1150
1	rs2274264	A	G	0.1336	1302
1	rs12103	G	A	0.09862	1308

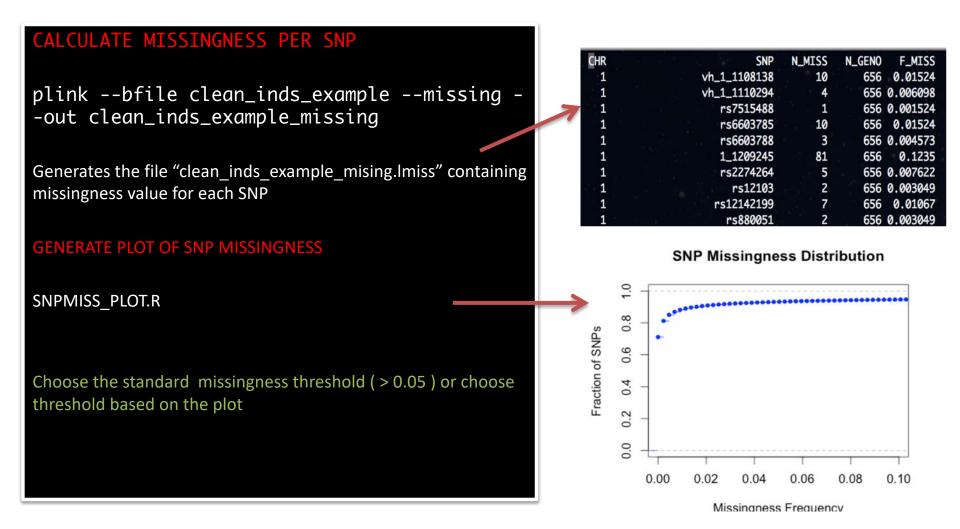


nextflow Parameters

• $cut_maf = 0.01$



Identify SNPS with high missingness

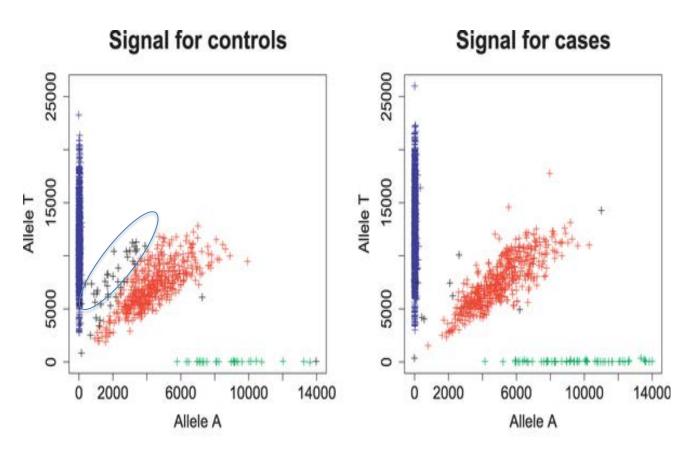


nextflow Parameters

• cut geno = 0.01



Differential missingness



- Missing frequency is also assessed separately in cases and in controls because differential missingness is a common source of false positive associations.
- SNPs showing highly differential missingness (P<0.00001) are excluded

Identify SNPS with high differential missingness between cases and controls

CALCULATE DIFFERENTIAL MISSINGNESS

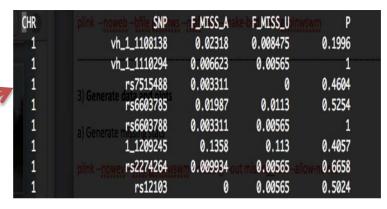
plink --bfile clean_inds_example --test-missing -out clean_inds_example_test_missing

Generates the file "example_test_missing.missing" containing differential missingness statistics for each SNP

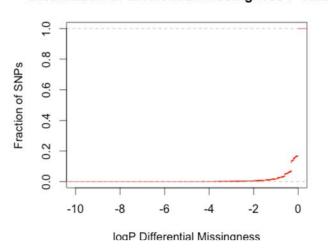
GENERATE PLOT OF P-VALUE DISTRIBUTION

DIFFMISS PLOT.R

Choose standard differential missingness p-value threshold (0.00001) or choose on the basis of the plot



Distribution of differential missingness P-value



nextflow Parameters

• cut diff miss = 0.05

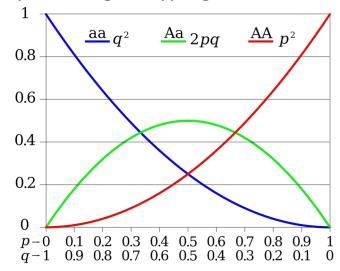


Hardy Weinberg Equilibrium

- Expected relationship between allele and genotype frequencies under certain assumptions
- Allele frequencies and genotypes remains constant over generations
- Deviations from HWE are used as a proxy for possible genotyping errors

Assumptions

- Diploid organisms
- Infinite population size
- Non-overlapping generations
- Random mating
- No selection, mutation or migration



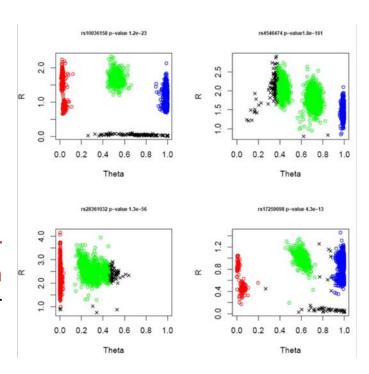
Testing for HWE

- Calculate the allele frequency (p)
 - Using observed genotype counts
- Calculate the expected genotype counts
 - Using the allele frequency (p)
- Compare the observed to the expected counts
 - χ^2 test

Reasons for HW Deviations

- Genotyping Error
- Subdivided Population
- Excess homozygotes= "Allele dropout in old samples"
- Any violations of the HW assumptions

- SNPs are excluded if substantially more or fewer samples heterozygous at a SNP than expected (excess heterozygosity or heterozygote deficiency)
- Threshold for significance 10⁻³ to 10⁻⁶
- Can only remove SNPs in controls which deviate from HWE or use less stringent HWE threshold in SNPs in cases over controls



GENEVA alcohol-dependence project: Quality control report

Identify SNPS which show extreme HWE deviations

GET DISTRIBUTION OF P-VALUES FOR ALL SNPS

plink --bfile clean_inds_example --hardy --out clean_inds_example_hwe

Generates the file "clean_inds_example_hwe.hwe" containing Hardy Weinberg statistics for each SNP separately in cases, controls and all. samples

SELECT UNAFFECTED

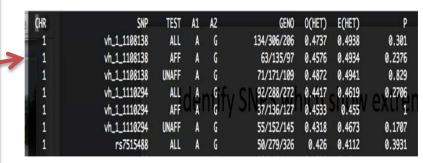
head -1 clean_inds_example_hwe.hwe >
example_clean_inds_example_hweu.hwe | grep
"UNAFF" clean_inds_example_hwe.hwe >>
example_clean_inds_example_hweu.hwe

GENERATE PLOT USING R SCRIPT

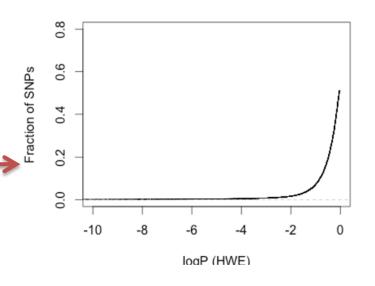
HWE_PLOT.R

(based only on controls)

Choose the standard HWE P-value threshold (0.00001) or select one on the basis of the plot

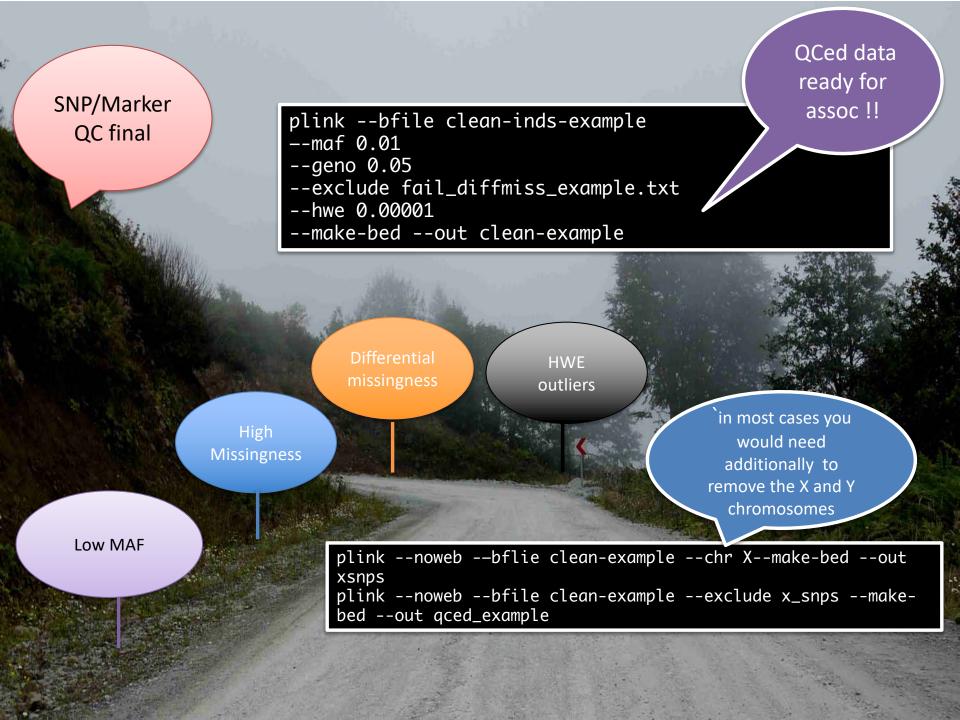


HWE P-value

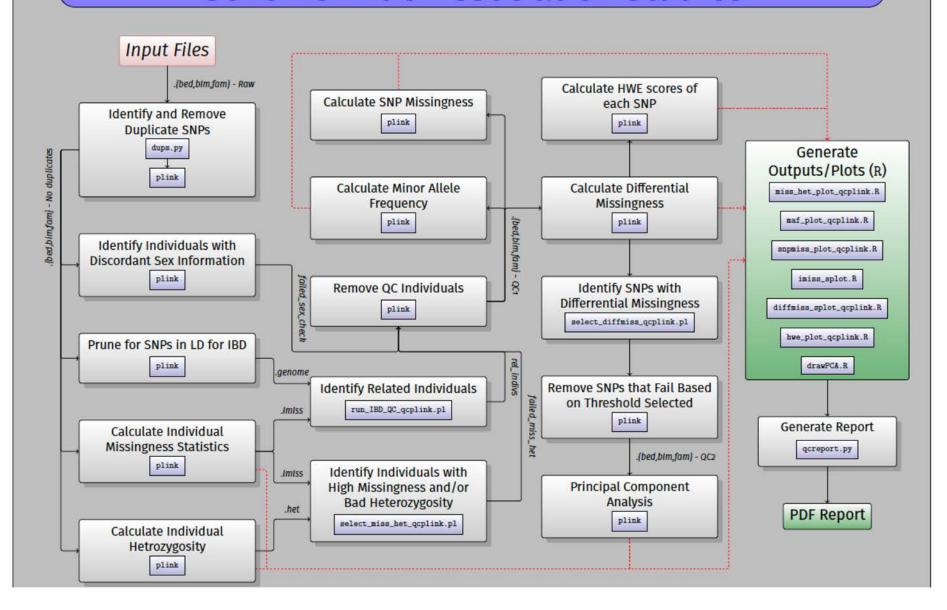


nextflow Parameters

• cut hwe = 0.00001



Genome Wide-Association Studies



Acknowledgments

- Ananyo Choudhury
- H3ABioNet funded by NHGRI grant number U41HG006941