

I have the power and multiple testing



Boulder 2008
Benjamin Neale

Slide of Contents

- Boring Statistical principles
- Power
 - Simple example
 - Formal definition
 - Influences on Power
- Multiple testing
 - Definition
 - FDR
 - QQ-Plots and you

BOREDCAT



ISBORED

Importance of power calculation

- Help design studies that are likely to succeed
 - Determine the minimum sample size necessary to achieve the desired level of statistical power (usually > 80%), for a given effect size
 - Determine the minimum effect size that can be detected with adequate statistical power, for a fixed sample size

Importance of power calculation

- Help design studies that are likely to succeed
 - Determine the minimum sample size necessary to achieve the desired level of statistical power (usually > 80%), for a given effect size
 - Determine the minimum effect size that can be detected with adequate statistical power, for a fixed sample size

Usually obligatory for grant applications

Simple example

- Investigate the linear relationship (ρ)
- between two random variables X and Y :
 $\rho=0$ vs. $\rho\neq 0$ (correlation coefficient).
- draw a sample, measure X, Y
- calculate the measure of association ρ
(Pearson product moment corr. coeff.)
- test whether $\rho \neq 0$.

How to Test $\rho \neq 0$

- assumed the data are normally distributed
- defined a null-hypothesis ($\rho = 0$)
- chosen α level (usually .05)
- utilized the (null) distribution of the test statistic associated with $\rho=0$
- $t = \rho \sqrt{[(N-2)/(1-\rho^2)]}$

How to Test $\rho \neq 0$


- Sample $N=40$
- $r=.303$, $t=1.867$, $df=38$, $p=.06$ $\alpha=.05$
- As $p > \alpha$, we fail to reject $\rho = 0$
- have we drawn the correct conclusion?

α = type I error rate
probability of deciding $\rho \neq 0$
(while in truth $\rho=0$)


α is often chosen to equal
.05...why?

DOGMA

$N=40$, $r=0$, $nrep=1000$ – central $t(38)$,
 $\alpha=0.05$ (critical value 2.04)




Observed non-null distribution ($\rho=.2$) and null distribution



In 23% of tests of $\rho=0$, $|t|>2.024$ ($\alpha=0.05$), and thus draw the correct conclusion that of rejecting $\rho = 0$.

The probability of rejecting the null-hypothesis ($\rho=0$) correctly is $1-\beta$, or the power, when a true effect exists



Hypothesis Testing

- Correlation Coefficient hypotheses:
 - h_0 (null hypothesis) is $\rho=0$
 - h_a (alternative hypothesis) is $\rho \neq 0$
 - Two-sided test, where $\rho > 0$ or $\rho < 0$ are one-sided
- Null hypothesis usually assumes no effect
- Alternative hypothesis is the idea being tested



Summary of Possible Results








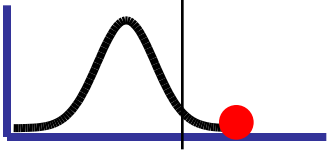
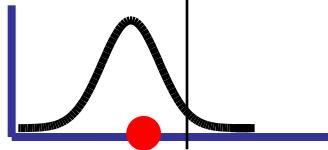
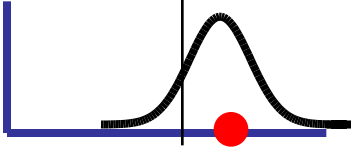
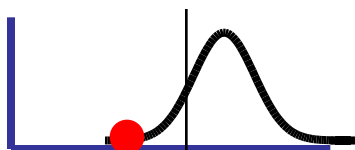
	H-0 true	H-0 false
accept H-0	$1-\alpha$	β
reject H-0	α	$1-\beta$

α =type 1 error rate

β =type 2 error rate

$1-\beta$ =statistical power

STATISTICS

		Rejection of H_0	Non-rejection of H_0
H_0 true	      	<p>Type I error at rate α</p> 	<p>Nonsignificant result $(1 - \alpha)$</p> 
H_A true		<p>Significant result $(1 - \beta)$</p> 	<p>Type II error at rate β</p> 

Power

- The probability of rejection of a false null-hypothesis depends on:
 - the significance criterion (α)
 - the sample size (N)
 - the effect size (Δ)


“The probability of detecting a given effect size in a population from a sample of size N, using significance criterion α ”



Standard Case

Sampling distribution if H_0 were true


Sampling distribution if H_A were true



Increased effect size

Sampling distribution if H_0 were true


Sampling distribution if H_A were true



Standard Case

Sampling distribution if H_0 were true

Sampling distribution if H_A were true




More conservative α

Sampling distribution if H_0 were true

Sampling distribution if H_A were true

alpha 0.01




Less conservative α

Sampling distribution if H_0 were true

Sampling distribution if H_A were true


alpha 0.01



Standard Case

Sampling distribution if H_0 were true

Sampling distribution if H_A were true

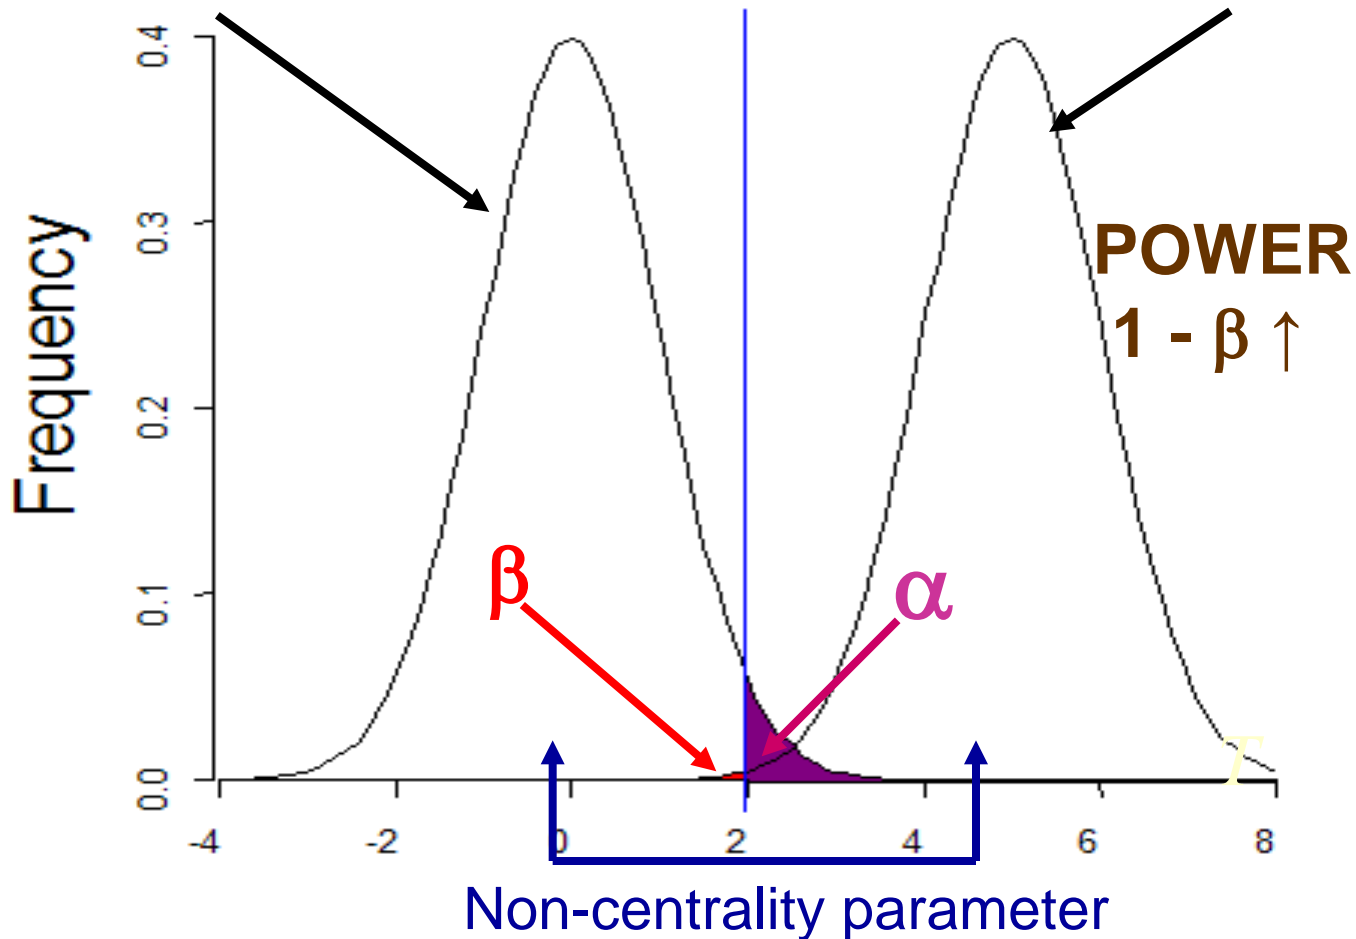


Increased sample size

Sampling distribution if H_0 were true

Sampling distribution if H_A were true

alpha 0.05




Increased sample size





Sampling distribution if H_0 were true

Sampling distribution if H_A were true

alpha 0.05



χ^2 distributions




Noncentral χ^2


- Null χ^2 has $\mu=df$ and $\sigma^2=2df$
- Noncentral χ^2 has $\mu=df + \lambda$ and $\sigma^2=2df + 4\lambda$
- Where df are degrees of freedom and λ is the noncentrality parameter

Noncentral χ^2 3 degrees of freedom


$\lambda=1$




$\lambda=4$



$\lambda=9$



$\lambda=16$



Short practical on GPC

- Genetic Power Calculator is an online resource for carrying out basic power calculations
- For our 1st example we will use the probability function calculator to play with power
- <http://pngu.mgh.harvard.edu/~purcell/gpc/>



GPC Power Practical

Genetic Power Calculator - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://pngu.mgh.harvard.edu/~purcell/gpc/

Most Visited Smart Bookmarks Free Hotmail Amazon.com: Statistic... The Arreat Summit - It... Share on Facebook

Gmail - GENEVA I... IBG workshop:2009:... IBG Index of /worksh... Wiley InterScien... Hulu - Videos Cypriots and U.N... Genetic Powe...

This site provides automated power analysis for variance components (VC) quantitative trait locus (QTL) linkage and association tests in sibships, and other common tests. It is currently under construction - suggestions, comments to [Shaun Purcell](#). If you use this site, please reference the following [Bioinformatics article](#):

Purcell S, Cherny SS, Sham PC. (2003) Genetic Power Calculator: design of linkage and association genetic mapping studies of complex traits. *Bioinformatics*, 19(1):149-150.

Modules

VC QTL linkage for sibships	Notes
VC QTL association for sibships and singletons	Notes
TDT for discrete traits	Notes
TDT and parenTDT with ascertainment	Notes
Case-control for discrete traits (new output features)	Notes
TDT for threshold-selected quantitative traits	Notes
Case-control for threshold-selected quantitative traits	Notes
Epistasis power calculator	Notes
Probability Function Calculator	Notes
Various miscellaneous utilities	

Click this link [no this isn't a banner ad]

Instructions for VC power calculations

All calculations are based upon formula derived in Sham et al (2000) [[AJHG, 66, 1616-1630](#)]. Users of this site who are unsure of the nature of the VC tests and power calculations are **strongly** advised to consult this article.

Find: exclude Next Previous Highlight all Match case

Bosoton: Wed 11:41 UK: Wed 16:41 Netherlands: Wed 17:41 Hong Kong: Thu 00:41 Los Angeles: Wed 08:41 http://pngu.mgh.harvard.edu/~purcell/gpc/pdf.html


start 2 F.. 2 W.. 2 M.. 2 M.. 5 M.. RGU... mac... 54% 11:41 AM

Parameters in probability function calculator

- Click on the link to probability function calculator
- 4 main terms:
- X : critical value of the chi-square
- $P(X > x)$: Power
- df: degrees of freedom
- NCP: non-centrality parameter



GPC



Probability Function Calculator

S. Purcell, 2000, 2005.

This site is designed to provide a calculator for the chi-squared and normal distributions. See below for notes on how to use the forms.


X	P(X>x)	df	NCP	
<input type="text"/>	?	<input type="text"/>	<input type="text"/>	<input type="button" value="Non-central chi-squared"/>
?	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="button" value="Inverse non-central chi-squared"/>
<input type="text"/>	<input type="text"/>	<input type="text"/>	?	<input type="button" value="NCP non-central chi-squared"/>
x	P(X>x)	mean	SD	
<input type="text"/>	?	<input type="text"/>	<input type="text"/>	<input type="button" value="Normal cumulative"/>
?	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="button" value="Inverse normal cumulative"/>

X : critical value for statistic
P(X>x) : probability of being above the critical value
df : degrees of freedom
NCP : noncentrality parameter

For central chi-squared distribution, NCP = 0.
i.e. df=1, P(X>x) = 0.05, NCP = 0 gives X = 3.84146

For standard normal distribution, mean = 0, sd = 1
i.e. mean = 0, sd = 1, P(X>x) = 0.025 give X = 1.95996

- 1) Fill in three
- 2) Click the button
- 3) Reveals the fourth





GPC

Statistical Genetics Group - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://pngu.mgh.harvard.edu/~purcell/gpc/pdf.html

Most Visited Smart Bookmarks Free Hotmail Amazon.com: Statistic... The Arreat Summit - It... Share on Facebook

Gmail - GENEVA I... IBG workshop:2009:... IBG Index of /worksh... Wiley InterScien... Hulu - Videos Cypriots and U.N... Statistical Ge...

Probability Function Calculator

S. Purcell, 2000, 2005.

This site is designed to provide a calculator for the chi-squared and normal distributions. See below for notes on how to use the forms.

X	P(X>x)	df	NCP	
<input type="text" value="3.85"/>	?	<input type="text" value="1"/>	<input type="text" value="10"/>	<input type="button" value="Non-central chi-squared"/>
<input type="text" value="?"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="button" value="Inverse non-central chi-squared"/>
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text" value="?"/>	<input type="button" value="NCP non-central chi-squared"/>

x	P(X>x)	mean	SD	
<input type="text"/>	?	<input type="text"/>	<input type="text"/>	<input type="button" value="Normal cumulative"/>
<input type="text" value="?"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="button" value="Inverse normal cumulative"/>

X : critical value for statistic
P(X>x) : probability of being above the critical value
df : degrees of freedom
NCP : noncentrality parameter

For central chi-squared distribution, NCP = 0.
i.e. df=1, P(X>x) = 0.05, NCP = 0 gives X = 3.84146

For standard normal distribution, mean = 0, sd = 1
i.e. mean = 0, sd = 1, P(X>x) = 0.025 give X = 1.95996

- 1) Fill in three
- 2) Click the button
- 3) Reveals the fourth

Find: Next Previous Highlight all Match case

Boston: Wed 11:49 UK: Wed 16:49 Netherlands: Wed 17:49 Hong Kong: Thu 00:49 Los Angeles: Wed 08:49 Done

start 2 F... 2 W... 2 M... 2 M... 5 M... RG... ma... 61% 11:49 AM



GPC

Statistical Genetics Group - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://pngu.mgh.harvard.edu/~purcell/cgi-bin/pdf.cgi

Most Visited Smart Bookmarks Free Hotmail Amazon.com: Statistic... The Arreat Summit - It... Share on Facebook

Gmail - GENEVA I... IBG workshop:2009:... IBG Index of /worksh... Wiley InterScien... Hulu - Videos Cypriots and U.N... Statistical Ge...

Probability Function Calculator

This site is designed to provide a calculator for the chi-squared and normal distributions. See below for notes on how to use the forms.

X	P(X>x)	df	NCP	
<input type="text" value="3.85"/>	<input type="text" value="0.884957"/>	<input type="text" value="1"/>	<input type="text" value="10"/>	<input type="button" value="Non-central chi-squared"/>
<input type="text" value="?"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="button" value="Inverse non-central chi-squared"/>
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text" value="?"/>	<input type="button" value="NCP non-central chi-squared"/>

X	P(X>x)	mean	SD	
<input type="text"/>	<input type="text" value="?"/>	<input type="text"/>	<input type="text"/>	<input type="button" value="Normal cumulative"/>
<input type="text" value="?"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="button" value="Inverse normal cumulative"/>

X : critical value for statistic
P(X>x) : probability of being above the critical value
df : degrees of freedom
NCP : noncentrality parameter

For central chi-squared distribution, NCP = 0.
i.e. df=1, P(X>x) = 0.05, NCP = 0 gives X = 3.84146

For standard normal distribution, mean = 0, sd = 1
i.e. mean = 0, sd = 1, P(X>x) = 0.025 give X = 1.95996

- 1) Fill in three
- 2) Click the button
- 3) Reveals the fourth

Find: exclude Next Previous Highlight all Match case

Boston: Wed 11:50 UK: Wed 16:50 Netherlands: Wed 17:50 Hong Kong: Thu 00:50 Los Angeles: Wed 08:50 Done

start 2 F... 2 W... 2 M... 2 M... 5 M... RG... ma... 62% 11:50 AM

Exercises

- 1) Find the power when $NCP=5$, degrees of freedom=1, and the critical X is 3.84
- 2) Find the NCP for power of .8, degrees of freedom=1 and critical X is 13.8

Answers

- 1) Power=0.608922, when NCP=5, degrees of freedom=1, and the critical X is 3.84
- 2) NCP=20.7613 when power of .8, degrees of freedom=1 and critical X is 13.8

Additional Factors

- Type of Data:
 - Continuous > Ordinal > Binary
 - Do not turn “true” binary into continuous
- Multivariate analysis
- Remove confounding/bias

Effects on Power Recap

- Larger Effect Size
- Larger Sample Size
- Alpha Level shifts
 - Beware the False Positive!!!
- Empirical significance/permutation

When To Do Power Calculations?

- Generally study planning stages of study
- Occasionally with negative result
- No need if significance is achieved
- Computed to determine chances of success

Steps in power calculation

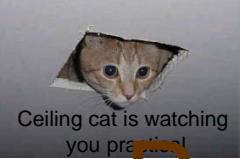
- Specify
 - Study design (e.g. case-control)
 - Statistical test

Steps in power calculation

- Specify
 - Study design (e.g. case-control)
 - Statistical test
- Assume hypothetical values for 2 of the 3 parameters:
 - Sample size
 - Effect size (including “exposure” frequency)
 - Statistical power

Steps in power calculation

- Specify
 - Study design (e.g. case-control)
 - Statistical test
- Assume hypothetical values for 2 of the 3 parameters:
 - Sample size
 - Effect size (including effect frequency)
 - Statistical power
- Calculate the remaining parameter



Practical using GPC for association

Genetic Power Calculator - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://pngu.mgh.harvard.edu/~purcell/gpc/

Most Visited Smart Bookmarks Free Hotmail Amazon.com: Statistic... The Arreat Summit - It... Share on Facebook

Gmail - Inbox (85... IBG workshop:2009:... IBG Index of /worksh... Wiley InterScien... Hulu - Videos Cypriots and U.N... Genetic Powe...

This site provides automated power analysis for variance components (VC) quantitative trait locus (QTL) linkage and association tests in sibships, and other common tests. It is currently under construction - suggestions, comments to [Shaun Purcell](#). If you use this site, please reference the following [Bioinformatics article](#):

Purcell S, Cherny SS, Sham PC. (2003) Genetic Power Calculator: design of linkage and association genetic mapping studies of complex traits. *Bioinformatics*, 19(1):149-150.

Modules

VC QTL linkage for sibships	Notes
VC QTL association for sibships and singletons	Notes
TDT for discrete traits	Notes
TDT and parenTDT with ascertainment	Notes
Case-control for discrete traits (new output features)	Notes
TDT for threshold-selected quantitative traits	Notes
Case-control for threshold-selected quantitative traits	Notes
Epistasis power calculator	Notes
Probability Function Calculator	Notes
Various miscellaneous utilities	

These are all association tests
Each refers to a different study design

Instructions for VC power calculations

All calculations are based upon formula derived in Sham et al (2000) [[AJHG, 66, 1616-1630](#)]. Users of this site who are unsure of the nature of the VC tests and power calculations are **strongly** advised to consult this article.

Find: exclude Next Previous Highlight all Match case

Boston: Wed 12:11 UK: Wed 17:11 Netherlands: Wed 18:11 Hong Kong: Thu 01:11 Los Angeles: Wed 09:11 Done

start 2 F.. 2 W.. 2 M.. 3 M.. 5 M.. RGU... mac... 78% 12:11 PM

Question 1

- What case control sample size do we need to achieve genome-wide significance for an odds ratio of 1.2 in a multiplicative model and an allele frequency of 20% when we directly type the locus for a disease with 5% prevalence?



Question 1

Genetic Power Calculator - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://pngu.mgh.harvard.edu/~purcell/gpc/

Most Visited Smart Bookmarks Free Hotmail Amazon.com: Statistic... The Arreat Summit - It... Share on Facebook

Gmail - Inbox (85... IBG workshop:2009:... IBG Index of /worksh... Wiley InterScien... Hulu - Videos Cypriots and U.N... Genetic Powe...

This site provides automated power analysis for variance components (VC) quantitative trait locus (QTL) linkage and association tests in sibships, and other common tests. It is currently under construction - suggestions, comments to [Shaun Purcell](#). If you use this site, please reference the following [Bioinformatics article](#):

Purcell S, Cherny SS, Sham PC. (2003) Genetic Power Calculator: design of linkage and association genetic mapping studies of complex traits. *Bioinformatics*, 19(1):149-150.

Modules

VC QTL linkage for sibships	Notes
VC QTL association for sibships and singletons	Notes
TDT for discrete traits	Notes
TDT and parenTDT with ascertainment	Notes
Case-control for discrete traits (new output features)	Notes
TDT for threshold-selected quantitative traits	Notes
Case-control for threshold-selected quantitative traits	Notes
Epistasis power calculator	Notes
Probability Function Calculator	Notes
Various miscellaneous utilities	

Click this link

Instructions for VC power calculations

All calculations are based upon formula derived in Sham et al (2000) [[AJHG, 66, 1616-1630](#)]. Users of this site who are unsure of the nature of the VC tests and power calculations are **strongly** advised to consult this article.

Find: exclude Next Previous Highlight all Match case

Boston: Wed 12:11 UK: Wed 17:11 Netherlands: Wed 18:11 Hong Kong: Thu 01:11 Los Angeles: Wed 09:11 Done

start 2 F.. 2 W. 2 M.. 3 M.. 5 M.. RGU... mac... 78% 12:11 PM



Question 1

Statistical Genetics Group - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://pngu.mgh.harvard.edu/~purcell/gpc/cc2.html

Most Visited Smart Bookmarks Free Hotmail Amazon.com: Statistic... The Arreat Summit - It... Share on Facebook

Gmail - Update: ... IBG workshop:2009:... IBG Index of /worksh... Wiley InterScien... Hulu - Videos Cypriots and U.N... Statistical Ge...

Genetic Power Calculator

Case - control for discrete traits

Allele frequency at the risk locus

High risk allele frequency (A)	: 0.2	(0 - 1)
Prevalence	: .05	(0.0001 - 0.9999)
Genotype relative risk Aa	: 1.2	(>1)
Genotype relative risk AA	: 1.44	(>1)
D-prime	: 1	(0 - 1)
Marker allele frequency (B)	: 0.2	(0 - 1)
Number of cases	: 1000	(0 - 10000000)
Control : case ratio	: 1	(>0) (1 = equal number of cases and controls)
<input type="checkbox"/> Unselected controls? (* see below)		
User-defined type I error rate	: 5e-8	(0.00000001 - 0.5)
User-defined power: determine N (1 - type II error rate)	: 0.80	(0 - 1)

Created by *Shaun Purcell* 24.Oct.2008

Find: exclude Next Previous Highlight all Match case

Boston: Wed 12:49 UK: Wed 17:49 Netherlands: Wed 18:49 Hong Kong: Thu 01:49 Los Angeles: Wed 09:49 Stopped

start 2 F. 2 W. 2 M. 3 M. 5 M. RG... ma... silv... 92% 12:49 PM



Question 1

Statistical Genetics Group - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://pngu.mgh.harvard.edu/~purcell/gpc/cc2.html

Most Visited Smart Bookmarks Free Hotmail Amazon.com: Statistic... The Arreat Summit - It... Share on Facebook

Gmail - Update: ... IBG workshop:2009:... IBG Index of /worksh... Wiley InterScien... Hulu - Videos Cypriots and U.N... Statistical Ge...

Genetic Power Calculator

Case - control for discrete traits

How common disease is

High risk allele frequency (A) : 0.2 (0 - 1)

Prevalence : 0.05 (0.0001 - 0.9999)

Genotype relative risk Aa : 1.2 (>1)

Genotype relative risk AA : 1.44 (>1)

D-prime : 1 (0 - 1)

Marker allele frequency (B) : 0.2 (0 - 1)

Number of cases : 1000 (0 - 10000000)

Control : case ratio : 1 (>0)
(1 = equal number of cases and controls)

Unselected controls? (* see below)

User-defined type I error rate : 5e-8 (0.00000001 - 0.5)

User-defined power: determine N : 0.80 (0 - 1)
(1 - type II error rate)

Created by *Shaun Purcell* 24.Oct.2008

Find: exclude Next Previous Highlight all Match case

Boston: Wed 12:49 UK: Wed 17:49 Netherlands: Wed 18:49 Hong Kong: Thu 01:49 Los Angeles: Wed 09:49 Stopped

start 2 F. 2 W. 2 M. 3 M. 5 M. RG... ma... silv... 92% 12:49 PM



Question 1

Statistical Genetics Group - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://pngu.mgh.harvard.edu/~purcell/gpc/cc2.html

Most Visited Smart Bookmarks Free Hotmail Amazon.com: Statistic... The Arreat Summit - It... Share on Facebook

Gmail - Update: ... IBG workshop:2009:... IBG Index of /worksh... Wiley InterScien... Hulu - Videos Cypriots and U.N... Statistical Ge...

Genetic Power Calculator

Case - control for discrete traits

High risk allele frequency (A) : 0.2 (0 - 1)
Prevalence : .05 (0.0001 - 0.9999)
Genotype relative risk Aa : 1.2 (>1)
Genotype relative risk AA : 1.44 (>1)

D-prime : 1 (0 - 1)
Marker allele frequency (B) : 0.2 (0 - 1)

Number of cases : 1000 (0 - 10000000)
Control : case ratio : 1 (>0)
(1 = equal number of cases and controls)

Unselected controls? (* see below)

User-defined type I error rate : 5e-8 (0.00000001 - 0.5)
User-defined power: determine N : 0.80 (0 - 1)
(1 - type II error rate)

Created by *Shaun Purcell* 24.Oct.2008

Find: exclude Next Previous Highlight all Match case

Boston: Wed 12:49 UK: Wed 17:49 Netherlands: Wed 18:49 Hong Kong: Thu 01:49 Los Angeles: Wed 09:49 Stopped

start 2 F. 2 W. 2 M. 3 M. 5 M. RG... ma... silv... 92% 12:49 PM

This is the relative risk—not the odds ratio. The OR is approximately equivalent to the RR for small values of RR.



Question 1

Statistical Genetics Group - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://pngu.mgh.harvard.edu/~purcell/gpc/cc2.html

Most Visited Smart Bookmarks Free Hotmail Amazon.com: Statistic... The Arreat Summit - It... Share on Facebook

Gmail - Update: ... IBG workshop:2009:... IBG Index of /worksh... Wiley InterScien... Hulu - Videos Cypriots and U.N... Statistical Ge...

Genetic Power Calculator

Case - control for discrete traits

High risk allele frequency (A)	: 0.2	(0 - 1)
Prevalence	: .05	(0.0001 - 0.9999)
Genotype relative risk Aa	: 1.2	(>1)
Genotype relative risk AA	: 1.44	(>1)
D-prime	: 1	(0 - 1)
Marker allele frequency (B)	: 0.2	(0 - 1)
Number of cases	: 1000	(0 - 10000000)
Control : case ratio	: 1	(>0) (1 = equal number of cases and controls)

Unselected controls? (* see below)

User-defined type I error rate	: 5e-8	(0.00000001 - 0.5)
User-defined power: determine N (1 - type II error rate)	: 0.80	(0 - 1)

Created by *Shaun Purcell* 24.Oct.2008

Find: exclude Next Previous Highlight all Match case

Boston: Wed 12:49 UK: Wed 17:49 Netherlands: Wed 18:49 Hong Kong: Thu 01:49 Los Angeles: Wed 09:49 Stopped

start 2 F. 2 W. 2 M. 3 M. 5 M. RG... ma... silv... 92% 12:49 PM

Risk of the AA genotype. Note that the model of risk is defined by the relationship between Aa and AA. We have a multiplicative model because $1.44 = 1.2 * 1.2$.



Question 1

Statistical Genetics Group - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://pngu.mgh.harvard.edu/~purcell/gpc/cc2.html

Most Visited Smart Bookmarks Free Hotmail Amazon.com: Statistic... The Arreat Summit - It... Share on Facebook

Gmail - Update: ... IBG workshop:2009:... IBG Index of /worksh... Wiley InterScien... Hulu - Videos Cypriots and U.N... Statistical Ge...

Genetic Power Calculator

Case - control for discrete traits

High risk allele frequency (A)	: 0.2	(0 - 1)
Prevalence	: .05	(0.0001 - 0.9999)
Genotype relative risk Aa	: 1.2	(> 1)
Genotype relative risk AA	: 1.44	(> 1)
D-prime	: 1	(0 - 1)
Marker allele frequency (B)	: 0.2	(0 - 1)
Number of cases	: 1000	(0 - 10000000)
Control : case ratio	: 1	(> 0) (1 = equal number of cases and controls)
<input type="checkbox"/> Unselected controls? (* see below)		
User-defined type I error rate	: 5e-8	(0.00000001 - 0.5)
User-defined power: determine N (1 - type II error rate)	: 0.80	(0 - 1)

Created by *Shaun Purcell* 24.Oct.2008

Find: exclude Next Previous Highlight all Match case

Boston: Wed 12:49 UK: Wed 17:49 Netherlands: Wed 18:49 Hong Kong: Thu 01:49 Los Angeles: Wed 09:49 Stopped

start 2 F. 2 W. 2 M. 3 M. 5 M. RG... ma... silv... 92% 12:49 PM

The LD statistic D' which represents recombination patterns historically. $D' +$ allele frequency at the typed locus information yields r^2



Question 1

Statistical Genetics Group - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://pngu.mgh.harvard.edu/~purcell/gpc/cc2.html

Most Visited Smart Bookmarks Free Hotmail Amazon.com: Statistic... The Arreat Summit - It... Share on Facebook

Gmail - Update: ... IBG workshop:2009:... IBG Index of /worksh... Wiley InterScien... Hulu - Videos Cypriots and U.N... Statistical Ge...

Genetic Power Calculator

Case - control for discrete traits

Sample size for cases

High risk allele frequency (A)	: 0.2	(0 - 1)
Prevalence	: .05	(0.0001 - 0.9999)
Genotype relative risk Aa	: 1.2	(>1)
Genotype relative risk AA	: 1.44	(>1)
D-prime	: 1	(0 - 1)
Marker allele frequency (B)	: 0.2	(0 - 1)
Number of cases	: 1000	(0 - 10000000)
Control : case ratio	: 1	(>0) (1 = equal number of cases and controls)

Unselected controls? (* see below)

User-defined type I error rate	: 5e-8	(0.00000001 - 0.5)
User-defined power: determine N (1 - type II error rate)	: 0.80	(0 - 1)

Created by *Shaun Purcell* 24.Oct.2008

Find: exclude Next Previous Highlight all Match case

Boston: Wed 12:49 UK: Wed 17:49 Netherlands: Wed 18:49 Hong Kong: Thu 01:49 Los Angeles: Wed 09:49 Stopped

start 2 F. 2 W. 2 M. 3 M. 5 M. RG... ma... silv... 92% 12:49 PM



Question 1

Statistical Genetics Group - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://pngu.mgh.harvard.edu/~purcell/gpc/cc2.html

Most Visited Smart Bookmarks Free Hotmail Amazon.com: Statistic... The Arreat Summit - It... Share on Facebook

Gmail - Update: ... IBG workshop:2009:... IBG Index of /worksh... Wiley InterScien... Hulu - Videos Cypriots and U.N... Statistical Ge...

Genetic Power Calculator

Case - control for discrete traits

Ratio of Controls to Cases

High risk allele frequency (A)	: 0.2	(0 - 1)
Prevalence	: .05	(0.0001 - 0.9999)
Genotype relative risk Aa	: 1.2	(>1)
Genotype relative risk AA	: 1.44	(>1)
D-prime	: 1	(0 - 1)
Marker allele frequency (B)	: 0.2	(0 - 1)
Number of cases	: 1000	(0 - 10000000)
Control : case ratio	: 1	(>0) (1 = equal number of cases and controls)

Unselected controls? (* see below)

User-defined type I error rate	: 5e-8	(0.00000001 - 0.5)
User-defined power: determine N (1 - type II error rate)	: 0.80	(0 - 1)

Created by *Shaun Purcell* 24.Oct.2008

Find: exclude Next Previous Highlight all Match case

Boston: Wed 12:49 UK: Wed 17:49 Netherlands: Wed 18:49 Hong Kong: Thu 01:49 Los Angeles: Wed 09:49 Stopped

start 2 F. 2 W. 2 M. 3 M. 5 M. RG... ma... silv... 92% 12:49 PM



Question 1

Statistical Genetics Group - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://pngu.mgh.harvard.edu/~purcell/gpc/cc2.html

Most Visited Smart Bookmarks Free Hotmail Amazon.com: Statistic... The Arreat Summit - It... Share on Facebook

Gmail - Update: ... IBG workshop:2009:... IBG Index of /worksh... Wiley InterScien... Hulu - Videos Cypriots and U.N... Statistical Ge...

Genetic Power Calculator

Case - control for discrete traits

High risk allele frequency (A) : 0.2 (0 - 1)
Prevalence : .05 (0.0001 - 0.9999)
Genotype relative risk Aa : 1.2 (>1)
Genotype relative risk AA : 1.44 (>1)

D-prime : 1 (0 - 1)
Marker allele frequency (B) : 0.2 (0 - 1)

Number of cases : 1000 (0 - 10000000)
Control : case ratio : 1 (>0)
(1 = equal number of cases and controls)

Unselected controls? (* see below)

User-defined type I error rate : 5e-8 (0.00000001 - 0.5)
User-defined power: determine N : 0.80 (0 - 1)
(1 - type II error rate)

Created by *Shaun Purcell* 24.Oct.2008

Find: exclude Next Previous Highlight all Match case

Boston: Wed 12:49 UK: Wed 17:49 Netherlands: Wed 18:49 Hong Kong: Thu 01:49 Los Angeles: Wed 09:49 Stopped

start 2 F. 2 W. 2 M. 3 M. 5 M. RG... ma... silv... 92% 12:49 PM

Genome-wide significance threshold
We'll learn about this later in the session



Question 1

Statistical Genetics Group - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://pngu.mgh.harvard.edu/~purcell/gpc/cc2.html

Most Visited Smart Bookmarks Free Hotmail Amazon.com: Statistic... The Arreat Summit - It... Share on Facebook

Gmail - Update: ... IBG workshop:2009:... IBG Index of /worksh... Wiley InterScien... Hulu - Videos Cypriots and U.N... Statistical Ge...

Genetic Power Calculator

Case - control for discrete traits

High risk allele frequency (A) : 0.2 (0 - 1)
Prevalence : .05 (0.0001 - 0.9999)
Genotype relative risk Aa : 1.2 (>1)
Genotype relative risk AA : 1.44 (>1)

D-prime : 1 (0 - 1)
Marker allele frequency (B) : 0.2 (0 - 1)

Number of cases : 1000 (0 - 10000000)
Control : case ratio : 1 (>0)
(1 = equal number of cases and controls)

Unselected controls? (* see below)

User-defined type I error rate : 5e-8 (0.00000001 - 0.5)
User-defined power: determine N : 0.80 (0 - 1)
(1 - type II error rate)

Created by *Shaun Purcell* 24.Oct.2008

Find: exclude Next Previous Highlight all Match case

Boston: Wed 12:49 UK: Wed 17:49 Netherlands: Wed 18:49 Hong Kong: Thu 01:49 Los Angeles: Wed 09:49 Stopped

start 2 F. 2 W. 2 M. 3 M. 5 M. RG... ma... silv... 92% 12:49 PM

Power level—what we're interested in observing



Question 1

Statistical Genetics Group - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://pngu.mgh.harvard.edu/~purcell/gpc/cc2.html

Most Visited Smart Bookmarks Free Hotmail Amazon.com: Statistic... The Arreat Summit - It... Share on Facebook

Gmail - Update: ... IBG workshop:2009:... IBG Index of /worksh... Wiley InterScien... Hulu - Videos Cypriots and U.N... Statistical Ge...

Genetic Power Calculator

Case - control for discrete traits

High risk allele frequency (A) : 0.2 (0 - 1)
Prevalence : .05 (0.0001 - 0.9999)
Genotype relative risk Aa : 1.2 (>1)
Genotype relative risk AA : 1.44 (>1)

D-prime : 1 (0 - 1)
Marker allele frequency (B) : 0.2 (0 - 1)

Number of cases : 1000 (0 - 10000000)
Control : case ratio : 1 (>0)
(1 = equal number of cases and controls)

Unselected controls? (* see below)

User-defined type I error rate : 5e-8 (0.00000001 - 0.5)
User-defined power: determine N : 0.80 (0 - 1)
(1 - type II error rate)

Click here to process

Created by *Shaun Purcell* 24.Oct.2008

Find: exclude Next Previous Highlight all Match case

Boston: Wed 12:49 UK: Wed 17:49 Netherlands: Wed 18:49 Hong Kong: Thu 01:49 Los Angeles: Wed 09:49 Stopped

start 2 F. 2 W. 2 M. 3 M. 5 M. RG... ma... silv... 92% 12:49 PM



Answer 1

Genetic Power Calculator - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://pngu.mgh.harvard.edu/~purcell/cgi-bin/cc2k.cgi

Most Visited Smart Bookmarks Free Hotmail Amazon.com: Statistic... The Arreat Summit - It... Share on Facebook

Gmail - Update: ... IBG workshop:2009:... IBG Index of /worksh... Wiley InterScien... Hulu - Videos Cypriots and U.N... Genetic Powe...

Alpha	Power	N cases for 80% power
0.1	0.4374	2803
0.05	0.3177	3559
0.01	0.1377	5296
0.001	0.0355	7742
$5e-08$	3.674e-05	17958

Case-control statistics: general 2 df test (BB versus Bb versus bb)
Sample NCP = 6.216

Alpha	Power	N cases for 80% power
0.1	0.716	1240
0.05	0.6002	1550
0.01	0.3609	2233
0.001	0.1451	3163
$5e-08$	0.0007464	6920

Case-control statistics: allelic 1 df test (B versus b)
Sample NCP = 6.224

Alpha	Power	N cases for 80% power
0.1	0.8024	993
0.05	0.7037	1260
0.01	0.4677	1876
0.001	0.2131	2743
$5e-08$	0.001557	6362

Scroll to the bottom for answer

Controls are selected (i.e. screened for not being a case)

Find: exclude Next Previous Highlight all Match case

Boston: Wed 13:02 UK: Wed 18:02 Netherlands: Wed 19:02 Hong Kong: Thu 02:02 Los Angeles: Wed 10:02 Done

start 2 F... 2 W... 2 M... 2 M... 2 M... RG... mac... silv... 94% 1:02 PM



Answer 1

Genetic Power Calculator - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://pngu.mgh.harvard.edu/~purcell/cgi-bin/cc2k.cgi

Most Visited Smart Bookmarks Free Hotmail Amazon.com: Statistic... The Arreat Summit - It... Share on Facebook

Gmail - Update: ... IBG workshop:2009:... IBG Index of /worksh... Wiley InterScien... Hulu - Videos Cypriots and U.N... Genetic Powe...

Alpha	Power	N cases for 80% power
0.1	0.4374	2803
0.05	0.3177	3559
0.01	0.1377	5296
0.001	0.0355	7742
5e-08	3.674e-05	17958

Case-control statistics: general 2 df test (BB versus Bb versus bb)
Sample NCP = 6.216

Alpha	Power	N cases for 80% power
0.1	0.716	1240
0.05	0.6002	1550
0.01	0.3609	2233
0.001	0.1451	3163
5e-08	0.0007464	6920

Case-control statistics: allelic 1 df test (B versus b)
Sample NCP = 6.224

Alpha	Power	N cases for 80% power
0.1	0.8024	993
0.05	0.7037	1260
0.01	0.4677	1876
0.001	0.2131	2743
5e-08	0.001557	6362

Controls are selected (i.e. screened for not being a case)

Find: exclude Next Previous Highlight all Match case

Boston: Wed 13:02 UK: Wed 18:02 Netherlands: Wed 19:02 Hong Kong: Thu 02:02 Los Angeles: Wed 10:02 Done

start 2 F... 2 W... 2 M... 2 M... RG... mac... silv... 94% 1:02 PM

Scroll to the bottom for answer

6,362 case samples required: total sample size 12,724

Questions on your own

- For the same model as above, find the total sample size required for a TDT
 - Hint: use TDT for discrete traits
 - Try for different effect sizes and models (e.g. dominance)
- What is the effect of degrading LD in case-control data?
 - Change the D' and keep allele freq the same
 - Change allele freq and keep D' the same
- How well does the additive model capture a dominance only effect?
- Should you use 2x population controls vs 1x screened controls
 - For a prevalence of 5% and for a prevalence of 25%?

Answers

- Additive
 - Total case number for CC: 6,362
 - Total case number for TDT: 7,079
- Dominance only
 - RR: 1; 1; 1.44
 - 30,595 cases for CC
 - 33,950 cases for TDT

Impact of indirect association

If a direct association study of a causal SNP would provide an NCP of λ

Then an indirect association study of a SNP in LD with the causal SNP has NCP of $R^2\lambda$

Impact of indirect association

If a direct association study of a causal SNP would provide an NCP of λ

Then an indirect association study of a SNP in LD with the causal SNP has NCP of $R^2\lambda$

i.e. NCP is linearly related to the magnitude of R^2 between causal and genotyped SNP

Hence the appropriateness of using R^2 as an LD metric for the selection of tag SNPs.

- Case-control for discrete traits

Disease $K = 0.1$

Locus $R_{AA} = R_{Aa} = 2$



$MAF = 0.05$

Marker1 $MAF = 0.05$ $D' = \{ 1, 0.8, 0.6, 0.4, 0.2, 0 \}$

Marker2 $MAF = 0.25$ $D' = \{ 1, 0.8, 0.6, 0.4, 0.2, 0 \}$

Sample 500 cases, 500 controls

- Genotypic risk at marker1 (left) and marker2 (right) as a function of D'



Additivity vs Dominance

- Recall from the biometrical model that
 - $V_a =$

Additivity vs Dominance

- Recall from the biometrical model that
 - $V_a = 2pq[a + (q-p)d]^2$

Additivity vs Dominance

- Recall from the biometrical model that
 - $V_a = 2pq[a + (q-p)d]^2$
 - $V_d =$

Additivity vs Dominance

- Recall from the biometrical model that
 - $V_a = 2pq[a + (q-p)d]^2$
 - $V_d = \text{Chlamydia}$

Additivity vs Dominance

- Recall from the biometrical model that
 - $V_a = 2pq[a + (q-p)d]^2$
 - $V_d = [2pqd]^2$

Additivity vs Dominance

- Recall from the biometrical model that
 - $V_a = 2pq[a + (q-p)d]^2$
 - $V_d = [2pqd]^2$
- Therefore, there can still be association evidence when the two homozygous classes have the same trait value mean and the heterozygous class does not equal the homozygotes

Additivity vs Dominance


- Recall from the biometrical model that
 - $V_a = 2pq[a + (q-p)d]^2$
 - $V_d = [2pqd]^2$
- $V_a = 0$ can only be achieved if $a = 0$ and $p = q$ or $a = (p-q)d$

Answers

- Should you use 2x population controls vs 1x screened controls
 - For a prevalence of 5% and for a prevalence of 25%?
- For a prevalence of 5% the 2x population controls are more powerful; for a prevalence of 25% the 1x screened controls are more powerful


Answers

The effect of adding misclassified cases




Answers

Increase in NCP as a function of including unselected controls



Power and NCP (df=1)

$\alpha = 0.01, 0.001, 0.0001, 0.00001, 0.000001, 0.0000001$



Ways to enhance power

- Increase sample size
- Combination of studies: meta-analysis
- Increase marker density
- Increase accuracy of phenotype measurements
- Increase accuracy of genotyping
- Rigorous quality control and error checking
- Collect and adjust for environmental covariates

Ways to enhance power

- Appropriate treatment of heterogeneity (including gene-environment interaction)
- Appropriate treatment of population substructure
- Select individuals with highest genetic loading as cases, and individuals with lowest genetic loading as controls (e.g. use quantitative endophenotypes and select individuals in the two extreme tails of the trait distribution)
- Well thought-through and sophisticated analysis plan, making full use of phenotype and genotype information

Simulation using PLINK

- PLINK simulation file-format

```
#SNPs  label  lower-freq  upper-freq  geno-rel-risk
```

- Exercise, to replicate result of analytic power calculation
 - See PLINK web-page on simulation
 - 600 cases, 600 controls
 - disease variant 20% MAF and GRR of 1.5
 - simulate and test 100,000 markers under the model
 - calculate power for type I error rate of 1×10^{-4}
 - *Hint.* To determine how many passing SNPs, you have several options:
 - Load results into R
 - Use `--pfilter` and `wc` Unix command
 - Use `awk` / `gawk`
 - Use *Haploview* to examine PLINK results file

- File *sim1.txt*

```
100000 alt 0.2 0.2 1.5
```

- **Generate and test SNPs on-the-fly**

```
./plink --simulate sim1.txt  
--simulate-ncases 600  
--simulate-ncontrols 600  
--simulate-prevalence 0.01  
--assoc
```

- **Calculate power**

```
awk ' $9 < 1e-4 ' plink.assoc | wc -l
```

Simulation using PLINK

- To specify 2-SNP haplotypes, given SNP frequencies and D' (not documented on current [www](http://www.plink.org) yet) add the flag `--simulate-tags` also

#SNPs	label	<u>l-freq</u>	<u>u-freq</u>	<u>l-freq</u>	<u>u-freq</u>	d-prime	geno-rel-risk
		Disease variant		Marker locus			↑
		(not observed)		(genotyped)			At disease locus

- Now simulate a 2% disease allele, with 5-fold (multiplicative) effect, that is in complete LD with a marker allele of 20% MAF
 - what is power now at the 20% genotype?
 - verify this using the GPC calculator
 - what is the apparent odds-ratio at the genotyped SNPs
 - what is the LD in terms of r^2 between the two loci (from GPC)?

- File *sim2.txt*

```
100000 alt 0.02 0.02 0.2 0.2 1.0 5.0
```

- Generate and test SNPs on-the-fly

```
./plink --simulate sim2.txt  
--simulate-tags  
--simulate-ncases 600  
--simulate-ncontrols 600  
--simulate-prevalence 0.01  
--assoc
```

- Calculate power

```
awk ' $9 < 1e-4 ' plink.assoc | wc -l
```

Working with NCPs

- Expected chi-square = NCP + df
- The NCP scales linearly with sample N
 - for same case/control ratio
- Two useful properties
 - combine independent tests by summing NCPs
 - NCP at marker $\sim r^2 \times$ NCP at disease locus
- To calculate power given NCP
 - using R

```
> 1 - pchisq( qchisq( 1 - 1e-4 , df=1 ) , df = 1 , ncp = 17.96 )
[1] 0.6358291
```
 - or PDF utility in GPC

Hodgepodge anyone?

- Multiple testing
 - Where it comes from
 - Why is it a problem
- False discovery
 - Theory & practice

Hodgepodge anyone?

- Multiple testing
 - Where it comes from
 - Why is it a problem
- False discovery
 - Theory & practice

What do we test

- Raise your hand if:

What do we test

- Raise your hand if:
 - You have analyzed more than 1 phenotype on a dataset

What do we test

- Raise your hand if:
 - You have analyzed more than 1 phenotype on a dataset
 - Used more than one analytic technique on a dataset (e.g. single marker association and haplotype association)

What do we test

- **Raise your hand if:**
 - You have analyzed more than 1 phenotype on a dataset
 - Used more than one analytic technique on a dataset (e.g. single marker association and haplotype association)
 - **Picked your best result from the bunch**

Genome-wide association



High throughput genotyping

Other multiple testing considerations

- Genome-wide association is really bad
 - At 1 test per SNP for 500,000 SNPs
 - 25,000 expected to be significant at $p < 0.05$, by chance alone

Other multiple testing considerations

- Genome-wide association is really bad
 - At 1 test per SNP for 500,000 SNPs
 - 25,000 expected to be significant at $p < 0.05$, by chance alone
- To make things worse
 - Dominance (additive/dominant/recessive)
 - Epistasis (multiple combinations of SNPs)
 - Multiple phenotype definitions
 - Subgroup analyses
 - Multiple analytic methods

Bonferroni correction

- For testing 500,000 SNPs
 - 5,000 expected to be significant at $p < 0.01$
 - 500 expected to be significant at $p < 0.001$
 -
 - 0.05 expected to be significant at $p < 0.0000001$
- Suggests setting significance level to $\alpha = 10^{-7}$ *
- Bonferroni correction for m tests
 - set significance level for p-values to $\alpha = 0.05 / m$
 - (or adjust the p-values to $m \times p$, before applying the usual $\alpha = 0.05$ significance level)
- *See Risch and Merikangas 1999

Genome-wide significance

- Multiple testing theory requires an estimate of the number of 'independent tests'
- Risch and Merikangas 1996 estimated a threshold of $10^{-6} = (0.05/(5*10,000))$
- HapMap 2005 estimate 10^{-8} based on encode deep sequencing in ENCODE regions
- Dudbridge and Gusnato, and Pe'er et al. 2008 Genetic Epidemiology estimate based on 'infinite density' like Lander and Kruglyak 1995 generate 5×10^{-8}

Implication for sample size

Genetic Power Calculator

m	α	χ^2	NCP (80% power)	Ratio
1	0.05	3.84	7.85	1
500	10^{-4}	15.14	22.39	2.85
500×10^3	10^{-7}	28.37	38.05	4.85
500×10^6	10^{-10}	41.82	53.42	6.81

Large but not “impossible” increase in sample size

Technical objection

Conservative when tests are non-independent

Nyholt (2004)

Spectral decomposition of correlation matrix

Effective number of independent tests

May be conservative: Salyakina et al (2005)

False Discovery

Permutation procedure

Philosophical objection

“Bonferroni adjustments are, at best, unnecessary and, at worst, deleterious to sound statistical inference”
Perneger (1998)

- Counter-intuitive: interpretation of finding depends on the number of other tests performed
- The general null hypothesis (that all the null hypotheses are true) is rarely of interest
- High probability of type 2 errors, i.e. of not rejecting the general null hypothesis when important effects exist

A Bayesian perspective

For each significant test, we can consider the probability that H_0 is in fact true (i.e. false positive probability)

Prob (H_0 True | H_0 Rejected)

Using Bayes' rule

$$\begin{aligned} P(H_0 | p \leq \alpha) &= \frac{P(p \leq \alpha | H_0)P(H_0)}{P(p \leq \alpha | H_0)P(H_0) + P(p \leq \alpha | H_1)P(H_1)} \\ &= \frac{\alpha\pi_0}{\alpha\pi_0 + (1 - \beta)(1 - \pi_0)} \end{aligned}$$

Taking the formula apart

$$\begin{aligned} P(H_0 | p \leq \alpha) &= \frac{P(p \leq \alpha | H_0)P(H_0)}{P(p \leq \alpha | H_0)P(H_0) + P(p \leq \alpha | H_1)P(H_1)} \\ &= \frac{\alpha\pi_0}{\alpha\pi_0 + (1 - \beta)(1 - \pi_0)} \end{aligned}$$

False Discovery Rate


Taking the formula apart

$$\begin{aligned} P(H_0 | p \leq \alpha) &= \frac{P(p \leq \alpha | H_0)P(H_0)}{P(p \leq \alpha | H_0)P(H_0) + P(p \leq \alpha | H_1)P(H_1)} \\ &= \frac{\alpha\pi_0}{\alpha\pi_0 + (1 - \beta)(1 - \pi_0)} \end{aligned}$$

Alpha level:
Rate of false positives

Taking the formula apart


$$\begin{aligned} P(H_0 | p \leq \alpha) &= \frac{P(p \leq \alpha | H_0)P(H_0)}{P(p \leq \alpha | H_0)P(H_0) + P(p \leq \alpha | H_1)P(H_1)} \\ &= \frac{\alpha\pi_0}{\alpha\pi_0 + (1 - \beta)(1 - \pi_0)} \end{aligned}$$



Proportion of tests
that follow the null distribution

Taking the formula apart

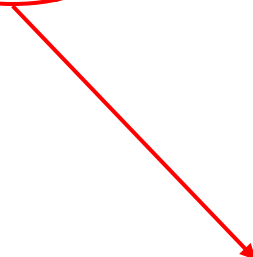
$$\begin{aligned} P(H_0 | p \leq \alpha) &= \frac{P(p \leq \alpha | H_0)P(H_0)}{P(p \leq \alpha | H_0)P(H_0) + P(p \leq \alpha | H_1)P(H_1)} \\ &= \frac{\alpha\pi_0}{\alpha\pi_0 + (1 - \beta)(1 - \pi_0)} \end{aligned}$$



Power to detect
association

Taking the formula apart

$$\begin{aligned} P(H_0 | p \leq \alpha) &= \frac{P(p \leq \alpha | H_0)P(H_0)}{P(p \leq \alpha | H_0)P(H_0) + P(p \leq \alpha | H_1)P(H_1)} \\ &= \frac{\alpha\pi_0}{\alpha\pi_0 + (1 - \beta)(1 - \pi_0)} \end{aligned}$$



Proportion of tests
that follow the
alternative distribution

Taking the formula apart

$$P(H_0 | p \leq \alpha) = \frac{P(p \leq \alpha | H_0)P(H_0)}{P(p \leq \alpha | H_0)P(H_0) + P(p \leq \alpha | H_1)P(H_1)}$$

False Discovery Rate = $\frac{\alpha\pi_0}{\alpha\pi_0 + (1-\beta)(1-\pi_0)}$

Alpha level:
Rate of false positives

Proportion of tests
that follow the null distribution

Power to detect
association

Proportion of tests
that follow the
alternative distribution

The diagram illustrates the components of the False Discovery Rate (FDR) formula. The formula is shown as $P(H_0 | p \leq \alpha) = \frac{P(p \leq \alpha | H_0)P(H_0)}{P(p \leq \alpha | H_0)P(H_0) + P(p \leq \alpha | H_1)P(H_1)}$. Below this, the FDR is simplified to $\frac{\alpha\pi_0}{\alpha\pi_0 + (1-\beta)(1-\pi_0)}$. Red circles highlight the terms α , π_0 , $(1-\beta)$, and $(1-\pi_0)$ in the simplified formula. Red arrows point from these terms to their definitions: α is the Alpha level (Rate of false positives), π_0 is the Proportion of tests that follow the null distribution, $(1-\beta)$ is the Power to detect association, and $(1-\pi_0)$ is the Proportion of tests that follow the alternative distribution.

A Bayesian perspective

Re-expressing the equation in term of α :

$$\alpha = \frac{P(H_0 | p \leq \alpha)}{1 - P(H_0 | p \leq \alpha)} \frac{1 - \pi_0}{\pi_0} \frac{1 - \beta}{1}$$

A Bayesian perspective

Re-expressing the equation in term of α :

$$\alpha = \frac{P(H_0 | p \leq \alpha)}{1 - P(H_0 | p \leq \alpha)} \frac{1 - \pi_0}{\pi_0} \frac{1 - \beta}{1}$$

False Discovery Rate

Proportion of tests that follows the null distribution

Power

Implications

- Justification of traditional choice $\alpha=0.05$
 - False positive rate $\sim \alpha$, when $\pi_0 \sim \frac{1}{2}$ and $1-\beta \rightarrow 1$

Implications

- Justification of traditional choice $\alpha=0.05$
 - False positive rate $\sim \alpha$, when $\pi_0 \sim \frac{1}{2}$ and $1-\beta \rightarrow 1$
- Maintenance of low false positive rate requires α to be set proportional to
 - $1-\beta$ (power)
 - $(1-\pi_0)/\pi_0$ (proportion of tests that follow the null)

Implications

- Justification of traditional choice $\alpha=0.05$
 - False positive rate $\sim \alpha$, when $\pi_0 \sim 1/2$ and $1-\beta \rightarrow 1$
- Maintenance of low false positive rate requires α to be set proportional to
 - $1-\beta$ (power)
 - $(1-\pi_0)/\pi_0$ (proportion of tests that follow the null)
- Multiple testing usually reflects lack of strong hypotheses and therefore associated with high π_0
 - Bonferroni adjustment effectively sets $\alpha \propto 1/m$, which is equivalent to assuming $\pi_0 = m/(m+1)$. **But is this reasonable?**

Fixed significance level

- Use fixed value of π_0 based on a guesstimate of the proportion of SNPs in the genome that have an effect, e.g. $1 - \pi_0 = 25/10^7 = 2.5 \times 10^{-6}$
- Power = 0.8
- False positive rate = 0.05
- Then $\alpha \sim 10^{-7}$ (regardless of m)

Adaptive significance level

- Use the availability of multiple tests to our advantage, because the empirical distribution of p-values can inform us about the suitable significance level
- Suppose that out of 500,000 SNPs, 100 are observed to be significant at $\alpha=0.00001$. Since the expected number of significant SNPs occurring by chance is 5, the false positive rate given by setting $\alpha=0.00001$ is 5/100
- Therefore a desired false positive rate can be obtained by setting α appropriately, according to the observed distribution of p-values (False Discovery Rate method)

Hodgepodge anyone?

- Multiple testing
 - Where it comes from
 - Why is it a problem
- **False discovery**
 - Theory & practice

Benjamini-Hochberg FDR method

Benjamini & Hochberg (1995) Procedure:

1. Set FDR (e.g. to 0.05)
2. Rank the tests in ascending order of p-value, giving $p_1 \leq p_2 \leq \dots \leq p_r \leq \dots \leq p_m$
3. Then find the test with the highest rank, r , for which the p-value, p_r , is less than or equal to $(r/m) \times \text{FDR}$
4. Declare the tests of rank 1, 2, ..., r as significant

A minor modification is to replace m by $m\pi_0$

B & H FDR method


FDR=0.05

Rank	P-value	$(\text{Rank}/m) \times \text{FDR}$	Reject H_0 ?
1	.008	.005	1
2	.009	.010	1
3	.165	.015	0
4	.205	.020	0
5	.396	.025	0
6	.450	.030	0
7	.641	.035	0
8	.781	.040	0
9	.901	.045	0
10	.953	.050	0

Modified FDR methods

Storey 2002 procedure:

Under the null P-values look like:




Modified FDR methods

Storey 2002 procedure:

Under the alternative P-values look like:

Distribution of P-values under alternative




Modified FDR methods

Storey 2002 procedure:

Under the alternative P-values look like:

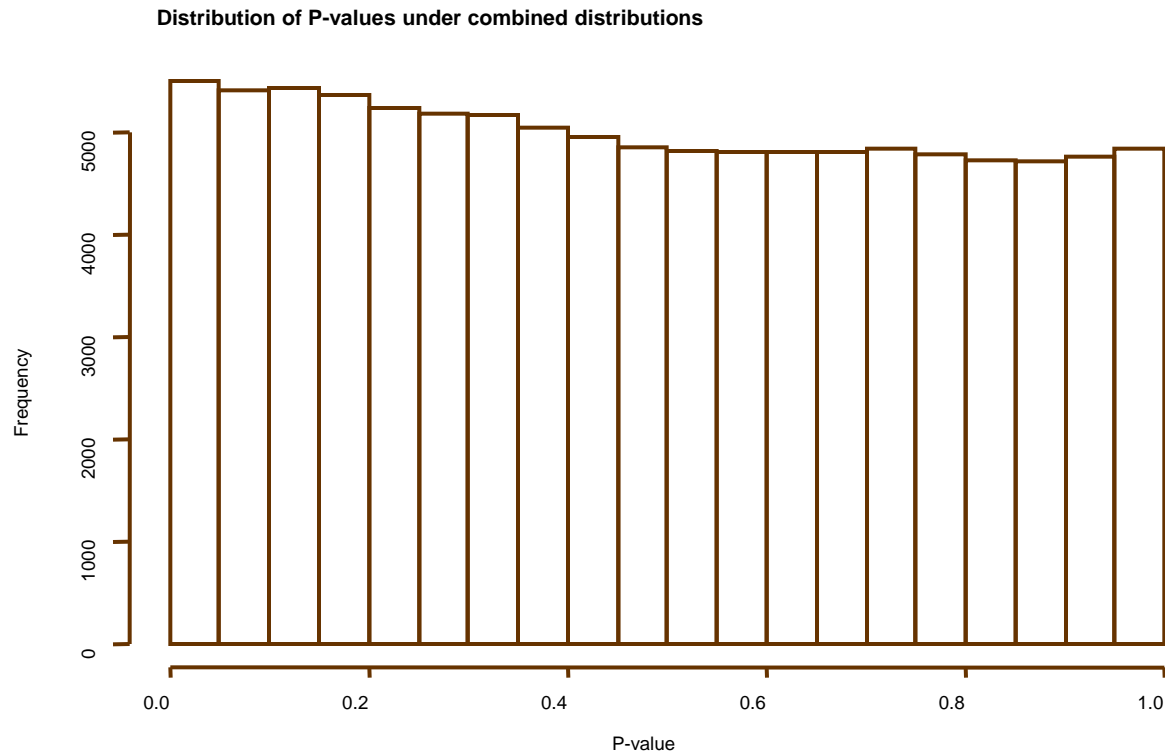
Distribution of P-values under alternative



Modified FDR methods

Storey 2002 procedure:


Combined distribution of P-values look like:



Modified FDR methods

Storey 2002 procedure:


Combined distribution of P-values look like:



Modified FDR methods

Storey 2002 procedure:

Combined distribution of P-values look like:




This is the information that FDR detects

Modified FDR methods

Storey 2002 procedure:


Combined distribution of P-values look like:



Modified FDR methods

Storey 2002 procedure:

Combined distribution of P-values look like:




The number of tests above $p = .5$ is 47651 out of 100000

Modified FDR methods

Storey 2002 procedure:

Combined distribution of P-values look like:




The number of tests above $p = .5$ is 47651 out of 100000

So the proportion of tests that follows the null: $47651/50000$ or $.95302 = \pi_0$

Modified FDR methods

Storey 2002 procedure:

Combined distribution of P-values look like:



The number of tests above $p = .5$ is 47651 out of 100000

So the proportion of tests that follows the null: $47651/50000$ or $.95302 = \pi_0$

So we replace the number of tests with the number of tests times π_0 or 95302.

“Parametric FDR” methods

Mixture model: some test statistics follow the null distribution, while others follow a **specified** alternative distribution

Special cases:

- Central and non-central chi-square distributions (Everitt & Bullmore, 1999)

- Central and non-central normal distributions (Cox & Wong, 2004)

- Uniform and beta distributions (Allison et al, 2002)

From fitted model, calculates the posterior probability of each test belonging to the null distribution (i.e. of being a false discovery if declared significant)


Pitfalls of the FDR method

- Assumption: p-values are distributed as $U[0,1]$ under H_0
 - If untrue (e.g. biased genotyping, population substructure) then this could lead to an excess of small p-values and hence misleading FDR results
- Requires a large number of tests to work
- The accuracy of the FDR is not easy to determine
- Requires a distribution (detectable number) of tests under the alternative


QQ Plots and You

- Q-Q plots stand for quantile-quantile plots
- A quantile is the value of a distribution at a given percentage
- The 100-quantiles are called percentile
- The 20-quantiles are called vigiciles
- The 12-quantiles are called duo-deciles
- The 10-quantiles are called decile
- The 9-quantiles are called noniles
- The 5-quantiles are called quintiles
- The 4-quantiles are called quartiles
- The 3-quantiles are called tertiles or terciles

QQ Plots and You




QQ Plots and You



QQ Plots and You

Five True GWS Effects Fifteen Extra Effects QQ plot



Effects spiked in:

1e-8, 2e-8, 3e-8,
4e-8, 5e-8, 1e-7,
1e-6, 2e-6, 3e-6,
4e-6, 5e-6, 6e-6,
7e-6, 8e-6, 9e-6