28 February 2003

Power of tests to compare two means & power of QTL mapping

Assume that we have *n* observations, in two groups of size *n/2*. We wish to test whether the group mean are different from each, and calculate the probability of detecting a difference when it exists, the statistical power. Such a power calculation might arise in the context of a new 'treatment', when the interest is in the mean effect of the treatment relative to some control. In the application of QTL mapping the question of power could arise before a crossing experiment is conducted.

The usual way of doing the power calculations is to assume two parameters, the size of the difference (d) and the within population (group) variance (σ_w^2). The mean difference is assumed fixed and the within-population observations are random. The model is:

$$y_{ij} \quad = \quad \mu_i \quad + \quad e_{ij}$$

Under the null hypothesis (H_0),

$\mu_1 = \mu_2$, and var(e) = σ_w^2

Under the alternative hypothesis (H_1),

$\mu_1 \neq \mu_2$, and var(e) = σ_w^2 (same as before)

We could calculate the 'overall' variance across both populations (groups), but this is usually not done because d is considered fixed. The overall variance, ignoring the between group difference, is

$$\text{var}(y) = \tfrac{1}{4}d^2 + \sigma_w^2$$

This variance is not constant but a function of d and σ_w^2. The proportion of variance in the combined population due to the mean difference is,

$$r^2 = \tfrac{1}{4}d^2 / (\tfrac{1}{4}d^2 + \sigma_w^2)$$

Statistical tests

Within-group variance is not known

Usually, the within-group variance is not known and has to be estimated from the data. In that case, the test statistic is:

$$T = \frac{(\bar{y}_2 - \bar{y}_1)}{\sqrt{\hat{\sigma}_w^2(\frac{2}{n} + \frac{2}{n})}}\,.$$

with SED = $\hat{\sigma}_w \sqrt{(4/n)}$ the estimated standard error of the difference. Under the null distribution, the mean and variance of the test statistic are 0 and n-2/(n-4), respectively, because T follows a standard t-distribution with df = n-2 degrees of freedom. Under the alternative hypothesis, the distribution of the test statistic is,

$T \sim t'(df,\lambda)$

a non-central t-distribution with $df = n-2$ degrees of freedom and non-centrality parameter

$\lambda = [(\mu_2 - \mu_1)/\sigma_w]\sqrt{(n/4)} = [d/\sigma_w]\sqrt{(n/4)}$

The mean and variance of the non-central t-distribution are

$E(t') = \sqrt{(½df)}\ [\Gamma(½(df-1))/\Gamma(½df)]\ \lambda$

$var(t') = (df/(df-2)[1 + \lambda^2] – E(t')^2$

An exact power calculation would proceed by calculation the probability that a test statistic from the non-central t-distribution would exceed the type-I error threshold from the central t-distribution.

An approximation to the mean and variance of the non-central t-distribution is, for large df,

$E(t') \approx \lambda$

$var(t') \approx 1 + ½\lambda^2/df = 1 + [d^2/\sigma_w^2]/8$

This approximation was verified using the exact expressions, and was found to be accurate for large n (n > 500). Note that the variance asympotically only goes to 1 for a fixed λ. This is consistent with considering the asymptotic variance of $(t')^2$. For large n, this is the variance of a non-central χ^2 with 1 degree of freedom and non-centrality parameter γ, which is $2(1 + 2\gamma)$, and also depends on the non-centrality parameter.

The variance of t' can also be written as a function of r^2,

$var(t') \approx (2-r^2) / (2(1-r^2) = (1 - ½r^2)/(1-r^2)$

If we make the further approximation that r^2 is small ($½r^2 \sim 0$), then,

$var(t') \approx 1 / (1 – r^2)$

This is the ratio that appears in the power calculations for QTL detection in a backcross of F_2 design in L&W (Chapter 14 & Appendix 5).

28 February 2003

Within-group variance known

Normal test

If the within-group variance (σ_w^2) is known, then the test statistic is,

$$T = \frac{(\bar{y}_2 - \bar{y}_1)}{\sqrt{\sigma_w^2(\frac{2}{n} + \frac{2}{n})}}$$

Under the null distribution this test statistic is distributed ~N(0,1) and under the alternative distribution as ~N(λ,1). Note that the variance is the same under H_0 and H_1, so that the required sample size for a given power is a function of $(z_{(1-\alpha/2)} + z_{(1-\beta)})^2$, with Pr(standard normal variate $<z_{(\alpha)}=\alpha$). Published work by Soller and co-workers on power calculations for QTL mapping also contain the term $(z_{(1-\alpha/2)} + z_{(1-\beta)})^2$, and not a scaled version depending on the ratio of variances under the null and alternative hypothesis.

Likelihood-ratio-test

A likelihood-ratio-test statistic (LRT) can also be used to conduct power studies. Under H_0, the asymptotic distribution of LRT is ~ $\chi^2_{(1)}$. Under the alternative hypothesis the distribution of LRT is ~ $\chi^2_{(1,\gamma)}$, a non-central χ^2 with non-centrality parameter γ. Since the distributions are based upon large sample theory, the likelihood ratio test should be very similar to the test assuming that the within-group variance is known, because both the central and non-central χ^2 are transformations of a normal distribution with mean 0 and *d*, respectively. The non-centrality parameter is,

$$\gamma = -n \log(1 - r^2)$$

If r^2 is small, then $\gamma \approx nr^2 = \lambda^2 (1-r^2) \approx \lambda^2$, as expected.

Lynch & Walsh

L&W perform their power calculations assuming that the total variance across the two groups (populations) is fixed. This is an unusual approach, because it implicity assumes that we know the maximum size of the effect before we do an experiment ($d^2 < 4\text{var(y)}$). However, perhaps this is reasonable approach when we've already observed the total variance (in, for example, an F_2 or halfsib population) and have to make a decision whether it's worthwhile to do a QTL experiment. L&W appear to use a hybrid between the normal and t-test in their power calculations (e.g., Example 1, page 873). They assume that the test statistic is normally distributed, consistent with knowing the within-group variance, but assume that the variance of the test statistic is different under the null and alternative hypothesis, consistent with the estimation of the within-group variance. Although for practical purposes these approaches will give the same answer (when r^2 is small), the explanation is somewhat confusing, because it is not clear what assumptions and approximations have been used.

Peter Visscher
Ian White
Andrew Carothers

Edinburgh, February 2003