Random Variables: Cheat Sheet

393	400	379	394	393	411	418	405	415	419

(experiment by A.C. Norman, 10–x–2006)

We could imagine an infinite set of such results, from which each measurement result x_i above has been randomly selected. Each single measurement returns a value x_i of the random variable x, which has a probability distribution $p(x)$ (in this case number of counts from a γ source in 0.1 s, but it could be anything we want to measure in physics).

Single measurement x_i of quantity x

		Best estimate from n measurements
Expected value	μ	$\overline{x} = \frac{1}{n}\sum_{i=1}^{n} x_i$
Variance	σ^2	$\frac{\sum_{i=1}^{n}(x_i-\overline{x})^2}{n-1} = \frac{\sum x_i^2 - \frac{1}{n}\left(\sum x_i\right)^2}{n-1}$
Standard deviation	σ	$\sqrt{\frac{\sum_{i=1}^{n}(x_i-\overline{x})^2}{n-1}} = \sqrt{\frac{\sum x_i^2 - \frac{1}{n}\left(\sum x_i\right)^2}{n-1}}$

Mean of n measurements of x

If n measurements of x have been made, the best estimate of the 'true' value of x is the mean of these, $\overline{x}$ (see μ above). This is also the expected value of a single measurement of x, and the expected value of the mean of a set of measurements.

		Best estimate from n measurements
Expected value	μ	$\overline{x}$
Variance	σ_m^2	$\sigma_m^2 = \frac{\sigma^2}{n} \approx \frac{\sum(x_i-\overline{x})^2}{n(n-1)} = \frac{\sum x_i^2 - \frac{1}{n}\left(\sum x_i\right)^2}{n(n-1)}$
Standard deviation	σ_m	$\sigma_m = \frac{\sigma}{\sqrt{n}} \approx \sqrt{\frac{\sum(x_i-\overline{x})^2}{n(n-1)}} = \sqrt{\frac{\sum x_i^2 - \frac{1}{n}\left(\sum x_i\right)^2}{n(n-1)}}$

Quoting experimental results

The result of a set of measurements is quoted as $\overline{x} \pm \sigma_m$. In graphical work, error bars normally run from $\overline{x} - \sigma_m$ to $\overline{x} + \sigma_m$
e.g. $R = 4.625 \pm 0.007\ \Omega$.

Range method

The standard method of estimating σ and σ_m is via the standard deviation of the sample. However, if no calculating aids are used, this can be tedious. The range method is particularly simple, almost as reliable as the standard method, and for values of n up to about 12, is quite adequate for most purposes.

If r is the difference between the highest and lowest of n readings, an estimate of σ is given by

$$\sigma \approx \frac{r}{\sqrt{n}},$$

and since $\sigma_m = \frac{\sigma}{\sqrt{n}}$,

$$\sigma_m \approx \frac{r}{n}.$$

NB This doesn't imply that σ_m is proportional to $1/n$, since $r \propto \sqrt{n}$.

For values of n greater than 12, the range method becomes increasingly unreliable, tending to underestimate σ.

Chebyshev's inequality and the Central Limit Theorem

Chebyshev's inequality tells us that $\overline{x}$ is unlikely to differ from μ, the true value, by more than a few multiples of $\sigma_m = \sigma/\sqrt{n}$. This means that $\overline{x}$ is an increasingly good estimate of μ as n increases (but since σ_m only decreases as $1/\sqrt{n}$, it becomes increasingly unprofitable to keep taking readings of the same quantity – better to reduce σ_m by reducing σ, i.e. by taking a more precise set of readings).

The Central Limit Theorem states that whatever the form of $p(x)$ (the probability distribution of x) with expected value μ and variance σ^2, the probability distribution of $\overline{x}$ will tend towards a normal or gaussian distribution with mean μ and standard deviation $\frac{\sigma}{\sqrt{n}}$ as n becomes large. We can therefore model the result of any measurement statistically as being normally distributed:

$$p(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(\frac{-x^2}{2\sigma^2}\right)$$

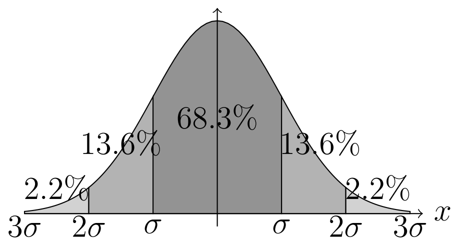

n	% of x_i within $\pm n\sigma$	Approximate fraction of readings outside $n\sigma$
0	0%	1 out of 1
1	68.3%	3
2	95.4%	20
3	99.73%	400
4	99.994%	16 000

For a set of readings with mean $\overline{x}$, about two-thirds of the individual readings should lie within $\overline{x} \pm \sigma$.

For a result quoted as $\overline{x} \pm \sigma_m$, the probability that the true value lies in the quoted range is roughly two-thirds.

ACN, 11–v–2011