Mean Values and Standard Deviations

We often want to determine the average value of a set of measurements, and to see how much they differ from each other. Let's introduce the concept of the mean value (μ, pronounced mu or mew), and the spread of a distribution (σ, pronounced sigma), also called the standard deviation. We will work with data produced by measuring the weights of seven pinto beans (chosen randomly from a large sack).

The mean value of a sample is simply the average of all measurements. If we measured seven pinto bean weights, 0.124, 0.351, 0.300, 0.323, 0.377, 0.402, and 0.356 grams, then the mean value is simply

$$\mu = \frac{1}{7}\left(m_1 + m_2 + m_3 + \ldots + m_7\right) = \frac{1}{N}\sum_{i=1}^{N} m_i$$

where the summation symbol $\sum$ means that we are adding together the seven m_i measurements, as i runs from 1 to N (and $N = 7$ because we have seven measurements). For this sample,

$$\mu = \frac{1}{7}\left(0.124 + 0.351 + 0.300 + 0.323 + 0.377 + 0.402 + 0.356\right) = 0.319 \text{ gm.}$$

Because we have only seven measurements, we want to make sure that an error in measurement does not skew our results dramatically. If you accidentally placed your thumb on the scale while making one measurement, for example, you might end up with an artificially high weight for one pinto bean. We will thus perform the averaging process again, but first discard the lowest and highest measured values.

We begin by sorting the seven values in order from lowest to highest, and then remove the top and bottom values from the list. Take care to remove the *lowest* and the *highest* values, not the *first* and *last* values. (The lowest and highest values are only the first and last values once you sort your list to run from from lowest to highest.)

$$\mu = \frac{1}{5}\left(\cancel{0.124} + 0.300 + 0.323 + 0.351 + 0.356 + 0.377 + \cancel{0.402}\right) = 0.341 \text{ gm.}$$

Our revised estimate of the average weight of a pinto bean is thus 0.341 grams.

We next want to know how wide the scatter is between values – are most pinto beans within a tenth of a gram, or ten grams, of the mean value? We can calculate the spread, or the standard deviation, as follows.

$$\sigma = \sqrt{\frac{1}{N-1}\sum_{i=1}^{N}\left(m_i - \mu\right)^2}$$

Let's step through each part of this formula, so that it makes sense.

The $(m_i - \mu)$ term means that we take each individual weight measurement m_i, and subtract off the mean value of the set. This tells us how far off each measurement lies from the average value. We then sum up the square of these differences. By squaring each difference before summing them, we make sure that points which lie above and below the average do not "cancel each other out." Put another way, the spread in the data set [2, 2, 2] is clearly not the same as in the data set [1, 2, 3], though the average value is equal to two, and the *sum* of the deviations, $0 + 0 + 0$ versus $-1 + 0 + 1$, from the mean is zero in each case.

The next step is to divide the sum of these squared differences by the term $(N - 1)$, in essence to get the average deviation per data point. If we did not do this, then the spread would increase as we took additional measurements, and there is no reason that measuring more points should cause them to spread out. Think of this as summing up the deviations from N measurements, and then dividing by N to get an average deviation per point. (The difference between N and $N - 1$ in this context is due to a subtle argument concerning degrees of freedom, one that we won't discuss here.)

Our final step is to take the square-root of this value, to give it the same unit as μ and our individual measurements. We have found the average value of $(m_i - \mu)^2$, so we take a square-root to get something which scales as m_i and μ do, in this case in units of grams.

Let's try this for our five remaining pinto bean data points. Get out your own calculator, and work through the calculation with us, so that you are confident that you understand each step. We first subtract off the mean value from each point, transforming the m_i values of 0.300, 0.323, 0.351, 0.356, and 0.377 grams to $(m_i - \mu)$ values of -0.041 (because $0.300 - 0.341$ = -0.041), -0.018, 0.010, 0.050, and 0.036 grams. Placing these $(m_i - \mu)$ values into the expression for σ, we observe that the standard deviation for our five points is

$$\sigma = \sqrt{\frac{1}{5 - 1} \left(-0.041^2 + -0.018^2 + 0.010^2 + 0.015^2 + 0.036^2\right)}$$

$$= \sqrt{\frac{1}{4} \left(0.001681 + 0.000324 + 0.000100 + 0.000225 + 0.001296\right)}$$

$$= \sqrt{0.25 \times 0.003626} = \sqrt{0.000907} = 0.030 \text{ gm.}$$

We can now combine our values for the mean value μ and the spread σ to say that our observed sample of pinto beans has a weight per bean of $\mu \pm \sigma$, or 0.341 ± 0.030 grams. In general you should find that two-thirds of all data points lie within $\pm\sigma$ of the mean value for a distribution (and 95% should lie with $\pm 2\sigma$), so we expect that two-thirds of a large sample of pinto bean weights should lie between $\mu - \sigma = 0.311$ and $\mu + \sigma = 0.371$ grams.

Do roughly 67% of our weights lie between 0.311 and 0.371 grams? Three of our five points do, or 60%, so we are doing fine (and note that all five values lie within 2σ of the mean).

It is a bit of work to calculate σ values, so we will usually use a computer program for this purpose. We will also plot our data, so that we can examine the distribution. You can work on both of these activities at

http://astronomy.nmsu.edu/geas/labs/html/plotter.shtml