


AIC vs Goodness of Fit

Comparing models

The model which is the best fit to the *data* is not necessarily the best representation of *reality*.



To see the difference, let's look at a hypothetical population of elephants. The graph shows the numbers for eight years. The crosses indicate the true population, which has an underlying upward trend. The black dots are the estimates from annual surveys of elephants. The lines represent two models calculated from the survey data. The straight line is a simple model with just 2 parameters; the wiggly line has 8 parameters and fits the data perfectly. In fact, the wiggleness of the wiggly line is due mainly to errors in estimating the population, and the simpler model is closer to the true situation. The wiggly model is described as “over-fitted” or “over-parameterized”.

In reality of course we don't know the true population, all we have are estimates which are never quite exact – we see the dots but not the crosses in the diagram. How do we decide which is the best model?

Akaike's Information Criterion (AIC) balances number of parameters and fit to the data (likelihood):

$$\text{AIC} = 2 * \text{No. of Parameters} - 2 * \log(\text{Likelihood})$$

and a small value for AIC indicates a better combination of simplicity and fit to data.