

ICES

Welcome to AD Model Builder and Stock Assessment

Anders Nielsen & Arni Magnusson

Purpose of this pre-course intro

- Give you a little more time to digest the first AD Model Builder example
- Give you a chance to get your laptop set up before the course (saves time)

Strategy of AD Model Builder

- Statistical environments (R, SPSS, SAS, Matlab, Gauss, Minitab, ...) have numerous standard models built in. AD Model Builder does not, but ...
- Sometimes these are not enough, and we need to build our own non-standard model.
- If the model is nonlinear and has many parameters (>5), then a tool using automatic differentiation (AD) should be considered.
- Essential parts of a statistical model
 - Read in data
 - Define parameters
 - Make predictions
 - Specify the objective function
- **Key feature** AD Model Builder finds the minimum of the objective function **fast**

Statistical modelling and likelihood

- Consider the following frequently used model:

$$y_i = \alpha + \beta x_i + \varepsilon_i, \quad \text{where } \varepsilon_i \sim \mathcal{N}(0, \sigma^2) \text{ independent.}$$

- Simple linear regression (and really AD Model Builder is overkill for this model)
- The model has three parameters α , β , and σ
- Want to find the value of the parameters that matches best with data
- Slightly rewritten the model is:

$$y_i \sim \mathcal{N}(\alpha + \beta x_i, \sigma^2) \text{ independent.}$$

- Likelihood function:** For a given set of model parameters $\theta = (\alpha, \beta, \sigma^2)$ the model assigns a measure of probability to our observations. This function is called the likelihood function:

$$L(y|\theta) = \prod_{i=1}^n \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{1}{2\sigma^2}(y_i - (\alpha + \beta x_i))^2\right)$$

(The product comes from our assumption of independent observations, n is the number of observations, and the rest is just the density of the normal distribution).

Negative log likelihood

- Often it is preferable to use $\ell(y|\theta) = -\log(L(y|\theta))$ instead:

$$\ell(y|\theta) = \frac{n}{2} \log(2\pi\sigma^2) + \frac{1}{2\sigma^2} \sum_{i=1}^n (y_i - (\alpha + \beta x_i))^2$$

- The likelihood function L expresses the likelihood of our observations for any given set of parameter values $\theta = (\alpha, \beta, \sigma^2)$, so a sensible estimate of the model parameters is to choose the values that maximize the likelihood for the actual observations, or equivalently minimizes the negative log likelihood ℓ

$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}} \ell(y|\theta)$$

- This way of estimating model parameters is called **maximum likelihood estimation**.
- In AD Model Builder this can be carried out for a wide variety of models (e.g. nonlinear, non-normal, and with and without random effects).
- What makes AD Model Builder unique is that it is possible to handle models orders of magnitude more complex (e.g. number of parameters) than standard statistical software.

Example: Linear regression in AD Model Builder

- The following data with 10 observations is assumed to follow the linear regression model:

# number of observations										
10										
# observed Y values	1.4	4.7	5.1	8.3	9.0	14.5	14.0	13.4	19.2	18
# observed x values	-1	0	1	2	3	4	5	6	7	8

- We want to estimate the three unknown model parameters.
- Entire AD Model Builder Program:

```
DATA_SECTION
  init_int N
  init_vector Y(1,N)
  init_vector x(1,N)

PARAMETER_SECTION
  init_number a
  init_number b
  init_number logSigma
  sdreport_number sigmasq
  objective_function_value nll

PROCEDURE_SECTION
  sigmasq = exp(2.0*logSigma);
  nll = N/2.0*log(2.0*M_PI*sigmasq) + 1.0/(2.0*sigmasq)*sum(square(Y-(a+b*x)));
```

- See line-by-line guide next page

- 1: First line starts the data section, which is where the data are read in. By default the data is read in from a file called the same as the model, but with extension '.dat', so if the program is called 'linreg.tpl', then the data file should be named 'linreg.dat'. In the data file, all lines starting with '#' are comments, ignored by the program.
 - 2: Reads in one integer, and stores it as 'N'
 - 3: Reads in N numbers and stores them as a vector named 'Y'
 - 4: Reads in N numbers and stores them as a vector named 'x'
-
- 5: Starts the parameter section, which is where all model parameters should be declared.
 - 6, 7, and 8: Declares three different model parameters α , β , and $\log(\sigma)$
 - 9: Declares an additional variable, which will be used to hold σ^2 , and report it in the output
 - 10: Names the function to be minimized, here 'nll' (for negative log likelihood)
-
- 11: Starts the procedure section, which is where the 'nll' function is calculated as a function of parameters and data already defined
 - 12 and 13: Implements the 'nll' function in normal C++ syntax, but helped by the built-in functions of AD Model Builder

How to compile and run

- Guides for all supported platforms are here:
<http://admb-project.org/documentation/>
 - In general the AD Model Builder program must be compiled^a, and then run.
 - Once the model has run, the results can be found in the '*.par', '*.std', and '*.cor' files produced by the program.
- ! Exercise:** Go to the website download and install AD Model Builder on your own laptop. Try to run the linear regression example shown here. Let us know how it works out.

^aThe compilation is actually a two step process 1) translate *.tpl file into a C++ program *.cpp, and 2) compile the *.cpp into a binary program with a standard C++ compiler, but these two steps are performed with one shell script, which is called 'admb'.