Non-Parametric Dose Response Models in Bayesian Adaptive Designs

Andrew P. Grieve (School of Medicine, King's College)

Questions & Answers from February 18, 2010 Webinar

Why are you using the term non- parametric yet in NDLM , we have parameters in observration and state inovations variances (V, W)

You are correct that there are parameters. I should probably have referred to these approaches as semiparametric in the spirit of the book by Rupper, Wand and Carroll (Semiparametric Regression, Cambridge University Press, 2003) although to quote the preface of that book "Flexible nonlinear regression is traditionally known as nonparametric regression; it differs from parametric regression in that the shape of the functional relationships are not pre-determined but can adjust to capture unusual or unexpected features of the data"

Will these NDLM curve-fitting approaches work well for small sample (e.g. nonclinical) studies? Would informative priors be required for good performance?

Clearly performance will depend on sample size, but also on the number of doses, remember the ASTIN trial used 15 doses + placebo. Informative priors will help to some extent but ones has to be careful if the sample size is small that the prior doesn't dominate the data. Also dogmatic views expressed in the prior may reduce the ability of the approach to "capture un unusual or unexpected features of the data"

In terms of the computational time, does WinBugs take a bit loger? or a lot longer comparing to C++ code?

WinBUGS will definitely take a lot longer than C++ code or Fortran or other compiled languages. For example there are fast, efficient algorithms for fitting the NDLM, for example the Frühwirth-Schnatter algorithm, which can be used in a purpose built system. On the other hand WinBUGS uses a general purpose approach that can be used for many different applications, and the applications can be developed rapidly. The reason for looking for ways to make WinBUGS more efficient is an acknowledgement of this fundamental difference.

When you apply MCMC to genearte the predicted observations, did you find that the algorithm converged relatively quickly? This will be a consideration when one is considering implementing Bayesian adaptive designs.

In general this is not a major problem. Since the previous MCMC run will provide a good starting value for the next generation. Of course as I pointed out both M and T are important parameters and these may depend on the amount of data that are currently available. If the generation is being done early in the trials you may need a large number of iterates, and convergence may take longer than later in the trial. In terms of implementing a single trial this is not a problem, but when simulating trials in the planning phase this may be an issue hence the search for efficient methods.