Adaptive Dose Finding in Clinical Trials

Christopher Jennison

Department of Mathematical Sciences,

University of Bath, UK

http://people.bath.ac.uk/mascj

Föreningen för Medicinsk Statistik

Göteberg,

September 24, 2007

Motivation

This talk draws heavily on the survey paper “Adaptive dose response studies” by Gaydos et al. (*Drug Information Journal*, 2006).

These authors state:

> “Insufficient exploration of the dose response . . . is often cited as a key contributor to the high late-stage attrition rate currently faced by the industry.”
>
> “Efficient learning about the dose response will ultimately reduce overall costs and provide better information on dose in the filing package.”
>
> “Initial proof-of-concept (PoC) studies often rely on testing just one dose level (e.g. the maximum tolerated dose) . . . hoping the right dose was selected.”
>
> “It is both feasible and advantageous to design a PoC study as an adaptive dose response trial,”

Topics for consideration

1. *Toxicity response:*

   Finding the Maximum Tolerated Dose (MTD)

   Typical for a Phase I study of a cancer therapy

   The need to avoid exposing subjects to doses with high toxicity leads automatically to adaptive designs.

2. *Efficacy response:*

   Comparing response against control (Phase IIa/Proof-of-concept)

   Selecting the best dose for further study (Phase IIb)

   Possibly, modelling the dose response curve.

1. Toxicity response: finding the MTD

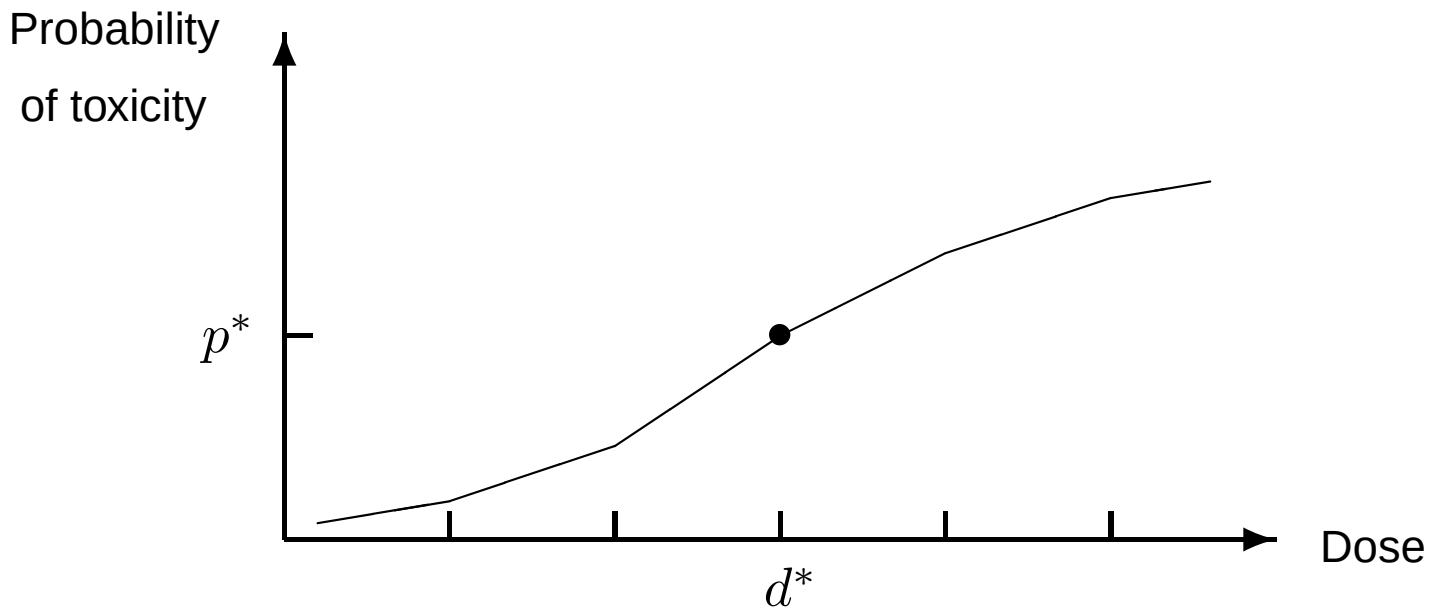

Aims are:

To find the dose d^* for which $Pr\{\text{Toxicity}\} = p^*$,

for example, $p^* = 0.3$.

To avoid exposing subjects to doses higher than d^* in the process.

Toxicity response: finding the MTD

a) Up-and-Down designs

"Toxicity" is defined as a binary outcome.

In 3+3 designs, patients are treated in groups of 3.

The algorithm moves the dose up and down, depending on the number of toxicities observed.

The MTD is taken to be the highest dose studied with less than, say, one third of subjects experiencing toxicity.

Remarks

Toxicity is assumed to increase with dose, but no specific model is used.

Reiner et al. (*Computational Statistics and Data Analysis*, 1999) show there is a significant probability of estimating the MTD incorrectly.

Finding the MTD

b) Random-Walk-Rule (RWR) designs

The MTD is treated as the dose for which the probability of serious toxicity is equal to a specified value. The aim is to estimate this dose.

The algorithm tosses a “biased coin” to select the next dose, with probabilities depending on observed toxicities.

RWR designs generate a cluster of doses with probability of toxicity around the value of interest.

Remarks

Again, no specific parametric model is employed.

RWR designs have a workable distribution theory: this can be used to check they achieve their aims effectively.

RWR designs can be efficient in finding the MTD.

Finding the MTD

c) The Continual-Reassessment Method (CRM)

A parametric model is assumed for the binary toxicity response, e.g., logistic regression.

Following a Bayesian approach, CRM updates the posterior distribution of model parameters after each response.

The next dose is chosen from a pre-defined set of doses to be that with posterior probability of toxicity closest to the target level for the MTD.

Remarks

The assignment of doses converges to the MTD.

There is a danger of escalating doses too quickly, but modifications can be made to mitigate the risk of unacceptable toxicity exposure.

Finding the MTD

d) Bayesian D-optimal designs

A particular parametric model is assumed for the dose response curve.

A “D-optimal” design minimises the determinant of the variance-covariance matrix of the model parameter estimates.

A prior distribution is specified for the model parameters and a small initial stage conducted to home in on the true parameter values.

Then, each new patient is allocated the dose which maximises the expected gain in information under the current posterior distribution.

Remarks

These designs aim to estimate the overall dose response curve rather than just the MTD.

Additional constraints can be incorporated to avoid assigning patients to doses with high toxicity.

General comments on “Finding the MTD”

Adaptive designs are essential in order to control the risk of exposing patients to highly toxic doses — and such designs have been around for some time.

Up-and-Down designs and Random-Walk-Rules make minimal model assumptions: their frequentist properties can be assessed by theory and simulation.

Continual-Reassessment and related methods follow a Bayesian approach: evaluation of their frequentist properties requires extensive simulation.

Results will not play a pivotal role in a New Drug Application, so a more relaxed approach to P-values and confidence intervals is acceptable at this stage.

But, decisions at this point have a crucial effect on the likelihood of developing a successful product. Objectives should be chosen carefully and adequate sample size allowed to attain these reliably.

2. Efficacy response: selecting the dose

a) Fixed-dose, parallel group designs

In this non-adaptive design, equal numbers of patients are allocated to each pre-specified dose level, plus a placebo or control treatment.

Efficacy can be a continuous or binary outcome.

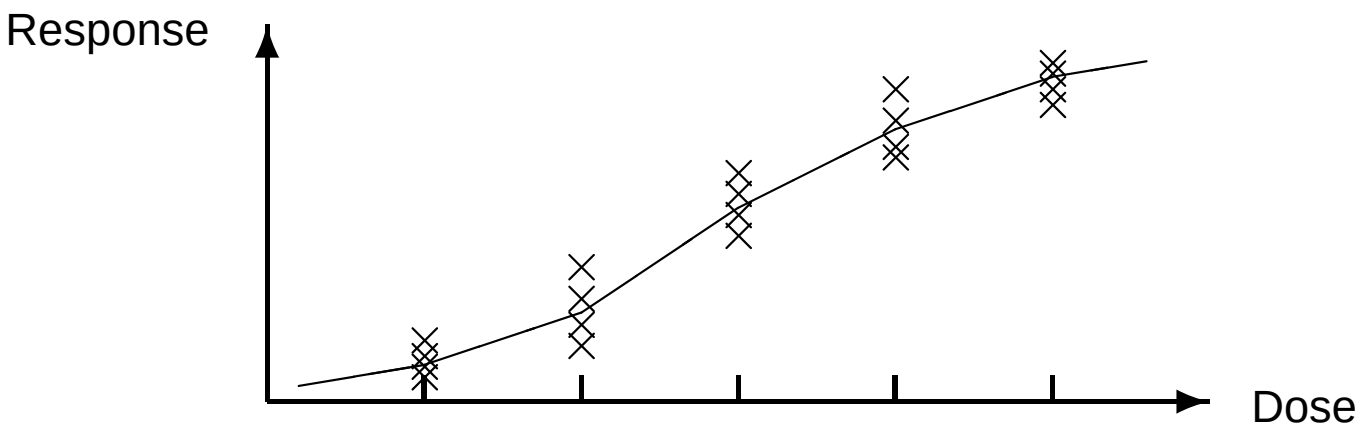

Goals are:

To test for a treatment effect vs placebo or control,

To select a dose to use in subsequent phases of testing.

Fixed-dose, parallel group designs

The “selected dose” could be:

- The smallest dose with discernible benefit,
- The maximum dose beyond which no further benefit is seen.

Remarks

- Since all doses are treated equally,
  - A considerable number of patients may be allocated to doses which turn out to be of little interest,
  - Sample size restrictions can limit the number of doses tested.
- Sequential methods might be used to stop allocation to some dose levels,
  - Observations “saved” can then be re-allocated to other dose levels,
  - Computing properties of the overall procedure can be a challenge.

Efficacy response

b) Model based procedures

It is natural to use a parametric model for the dose response curve, although the form of this relationship may not be known for certain at the outset.

Bretz, Pinheiro and Branson (*Biometrics*, 2005) combine model selection and inference about the most suitable dose using the selected model.

They allow for the multiple comparisons in identifying those models under which the data show a treatment effect vs control.

They consider several rules for the choice of dose and assess these by simulation.

Sonesson and Burman (2007) derive D-optimal designs for a variety of Emax models with up to four parameters.

Efficacy response

c) Bayesian adaptive designs

A model is specified for the dose response curve. This can be flexible and allow monotonicity if this is indicated by the data.

Normal Dynamic Linear Models can incorporate properties expected in a dose response curve.

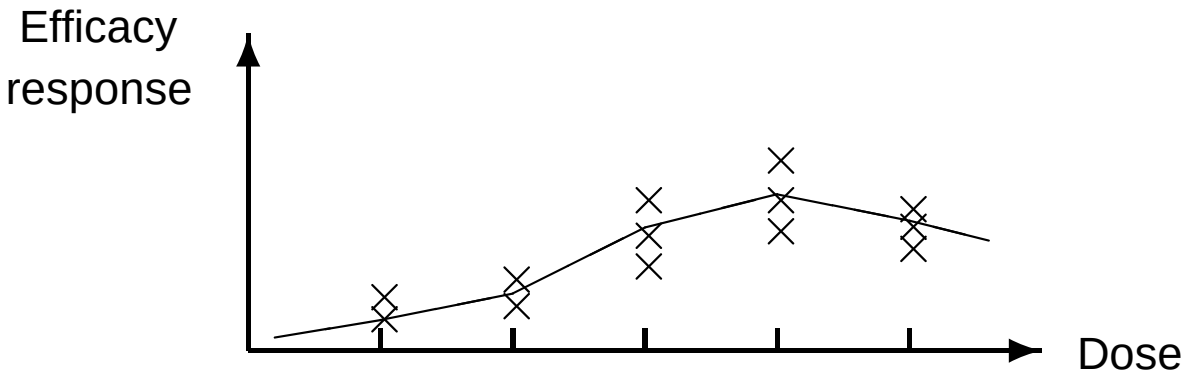

A prior on the model parameters is required. The posterior distribution is updated as data accrue, by MCMC simulation if necessary.

Bayesian adaptive designs

Assignment of patients to doses can be optimised for efficiency in meeting the study's objectives.

These may be:

To test for evidence of a treatment effect,

To find, say, an ED95 or E50 dose level.

Frequentist properties of the whole procedure can be investigated by simulation studies under a variety of true model parameters.

The ASTIN trial

An application to a stroke treatment is described by Krams et al. (*Stroke*, 2003).

Short term and long term responses were modelled jointly to gain information on the key long term response as rapidly as possible.

Rolling Dose Studies

Gaydos et al. (*Drug Information Journal*, 2006) use this term for a broad class of methods that allow flexible, dynamic allocation of patients to dose levels.

Dose level arms may be started or discontinued during the trial.

Allocation fractions and stopping rules are geared to the specific study objectives.

These designs are intended to maximise:

Learning about the dose response relationship,

Precision of estimated doses to achieve certain effects.

Some, but not all, of these methods are Bayesian.

PhRMA has constituted a working group on Rolling Dose Studies to evaluate and develop these methods.

3. Further Topics

Seamless designs

It may be possible to combine Phase II or Phase IIb data with later Phase III data on the selected dose level. We shall consider this issue in a later session.

Surrogate endpoints

When the primary endpoint is a long term response, a Phase II study may use a more immediate endpoint.

For cancer therapies, it is common to use tumour response or progression free survival as the short term endpoint when the primary endpoint is overall survival.

Studies of both safety and efficacy

Some Phase II studies consider safety and efficacy endpoints in seeking a dose that is both safe and efficacious.

4. Recommendations

Gaydos et al. recommend:

Consider adaptive dose response designs in exploratory development and to establish proof of concept.

Use an approach that incorporates a model for dose response.

Define the dose assignment mechanism prospectively and evaluate its operational characteristics through simulation, prior to the study.

Stop the trial as soon as there is enough information to make a decision.

Establish a monitoring committee to verify the design performs as expected.

Engage this committee in scenario simulations prior to protocol approval.

Use disease state and exposure-response model data to design studies.

Consider seamless approaches to improve the efficiency of learning.

Recommendations

I would pick out one further point:

> Remember the importance of decisions at Phase II for the future success of the therapy under investigation.
>
> This should affect power, and thus sample size, chosen for a Phase II study.

Paradoxically, Phase II studies tend to be small in comparison to Phase III trials.

However, differences in effect size between doses compared in Phase IIb are substantially smaller than the treatment effect investigators hope will be present for the selected dose level in Phase III.

This point is in keeping with Gaydos et al's motivating comment that:

> "Insufficient exploration of the dose response . . . is often cited as a key contributor to the high late-stage attrition rate currently faced by the industry."