Response-adaptive dose-finding under model uncertainty

Frank Bretz (Novartis)

Based on: Bornkamp, Bretz, Dette, Pinheiro (2011) Response-adaptive dose-finding under model uncertainty. *Annals of Applied Statistics* Vol. 5, No. 2B, 1611–1631

Design of Experiments in Healthcare
Isaac Newton Institute for Mathematical Sciences
Cambridge – August 15-19, 2011

Notation and framework

- Assume that the primary endpoint Y is observed for k parallel groups corresponding to doses $d_1, \ldots, d_k$ (d_1 typically placebo) and

$$Y_{ij} = f(d_i, \boldsymbol{\theta}) + \epsilon_{ij}, \ \epsilon_{ij} \overset{\text{ind}}{\sim} \mathcal{N}(0, \sigma^2) \ j = 1, \ldots, n_i, i = 1, \ldots, k$$

- Here $f(d_i, \boldsymbol{\theta})$ denotes a (typically) non-linear dose-response model function (*e.g.* logistic or the sigmoid Emax model)

- Model uncertainty regarding the correct model f (often even monotonicity questionable).
  $\hookrightarrow$ Use candidate set of models $\mathcal{M} = \{M_1, \ldots, M_M\}$ with model functions $f_1(d, \boldsymbol{\theta}_1), \ldots, f_M(d, \boldsymbol{\theta}_M)$.

Notation and framework

- Let Δ denote clinically relevant effect; we consider estimating the MED, i.e., the smallest dose achieving Δ on top of placebo:

$$\text{MED} = \min_{d \in (d_1, d_k]} \{ f(d, \boldsymbol{\theta}) > f(d_1, \boldsymbol{\theta}) + \Delta \}$$

  (other objectives are possible here of course)

- How to choose the allocation weights $\boldsymbol{w} = (w_1, \ldots, w_k)'$ of the total number of patients N to the dose-groups of the study to obtain most information of the MED?

Optimal design for MED estimation

- Dette et al. (2008) derive asymptotic approximations $V_m(\boldsymbol{w}, \boldsymbol{\theta}_m)$ for the variance of the MED estimate under a dose response model M_m.

- To acknowledge the inherent model uncertainty, Dette et al. (2008) consider as design criterion:

$$\Psi(\boldsymbol{w}) = \sum_{m=1}^{M} \alpha_m \log V(\boldsymbol{w}, \boldsymbol{\theta}_m | M_m),$$

  where α_m is the probability for M_m.

- Minimize average variance over all candidate models

Optimal design for MED estimation

- The weights $\boldsymbol{w}^*$ minimizing $\Psi(\boldsymbol{w})$ (and hence expected variance of the MED) depend on model parameters $\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_M$ and model probabilities $\alpha_1, \ldots, \alpha_M$
- How to obtain parameter values and model probabilities?

Adaptive Design for MED estimation

- Idea: Perform interim analyses to adapt the design
  $\rightarrow$ Use available data to update $\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_M$ and $\alpha_1, \ldots, \alpha_M$
  $\rightarrow$ Calculate optimal sample size allocation for next cohort
  based on current estimates
- Challenge in dose-finding studies: Large variability in response
  $\Rightarrow$ Highly variable ML-estimates (especially for initial cohorts)
- How to prevent that estimates (and resulting allocation) vary too wildly?
  Shrinkage towards a-priori plausible parameter values in a Bayesian framework.
- Another advantage of this approach:
  One can use posterior model probabilities for α_m.

Priors for parameters

- Typically dose-response models can be written as

$$f_m(x, \boldsymbol{\theta}_m) = \theta_{m0} + \theta_{m1} f^0(x, \boldsymbol{\theta}_m^0)$$

- Set up priors for baseline and maximum effect
  $\rightarrow$ can be translated to priors for linear parameters θ_{m0} and θ_{m1} for all presented models.
- Use beta priors for $\boldsymbol{\theta}_m^0$ (e.g. mode of beta distribution equal to parameter guess).

Procedure: 1) Before Trial Start

1. Selection of candidate models $f_1(d, \boldsymbol{\theta}_1), \ldots, f_M(d, \boldsymbol{\theta}_M)$
2. Choose prior model probabilities $p(M_m)$
3. Choose prior for $\boldsymbol{\theta}_m = (\theta_{m0}, \theta_{m1}, \boldsymbol{\theta}_m^0)'$ and σ^2
   - For $\theta_{m0}, \theta_{m1}, \sigma^2$ use normal-inverse gamma priors (conditionally conjugate)
   - For $\boldsymbol{\theta}_m^0$ use beta distribution with mode equal to best guess. The sum of both parameters of the beta distribution determines variability

Procedure: 2a) At Interim

Given current allocations to doses $n_1^{cur}, \ldots, n_k^{cur}$, with total number of patients $N^{cur} = \sum_{i=1}^{k} n_i^{cur}$.
How to allocate the N^{next} patients of the next cohort?
First calculate

1. Posterior Model Probabilities

$$p(M_m|\boldsymbol{y}) \propto p(M_m) \int p(\boldsymbol{y}|\boldsymbol{\theta}_m, M_m) p(\boldsymbol{\theta}_m|M_m) d\boldsymbol{\theta}_m$$

Note: At most 1d or 2d integration problems; straightforward to calculate with quadrature rules.

2. Posterior estimate of $\boldsymbol{\theta}_m$, e.g.

$$\widehat{\boldsymbol{\theta}}_m = \text{argmax}\{p(\boldsymbol{y}|\boldsymbol{\theta}_m, M_m) p(\boldsymbol{\theta}_m|M_m)\}$$

Procedure: 2b) At Interim

3. Plug in $\widehat{\boldsymbol{\theta}}_1, \ldots, \widehat{\boldsymbol{\theta}}_M$ and $p(M_1|\boldsymbol{y}), \ldots, p(M_M|\boldsymbol{y})$ into

$$\Psi(\boldsymbol{w}) = \sum_{m=1}^{M} \alpha_m \log V(\boldsymbol{w}^{next}, \boldsymbol{\theta}_m | M_m),$$

where $w_i^{next} = (w_i N^{next} + n_i^{cur})/(N^{next} + N^{cur})$ and minimize with respect to $\boldsymbol{w}$ (subject to $w_i \geq 0$ and $\sum_{i=1}^{k} w_i = 1$).

Note: This constrained optimization can be reparameterized into an unconstrained problem; an optimality check is possible.

4. Round the found optimal design $\boldsymbol{w}_{opt}$ so that the sum equals N^{next}; allocate next cohort.

Procedure: 3) At Trial End

- Use, for example, MCP-Mod approach (Bretz et al., 2005), which combines multiple comparison procedures with non-linear regression modeling techniques and has frequentist properties
  - Use Bayesian approach (and priors) only for design
- Alternatively, fit Bayesian model and estimate MED
- Extensive software implemented with the DoseFinding package in R, freely available on CRAN

Simulation Study

Perform simulations in setting of dose-finding trial for compound in asthma (each scenario simulated 5000 times):

- 6 dose-response models (placebo effect 100, max. effect 300)
- Overall sample size: 300, $\sigma = 350$, $\Delta = 200$
- 2 different dose allocations (for dose range $[0, 50]$):
  4 doses: $0, 2.5, 10, 20, 50$ and
  7 doses: $0, 0.5, 1, 2.5, 5, 10, 20, 50$
- Number of interim analyses: $0, 1, 2, 4, 9$
- Performance measures compared here:
  mean abs error in MED estimate: $E|MED_{true} - \widehat{MED}|$

Dose Response Shapes

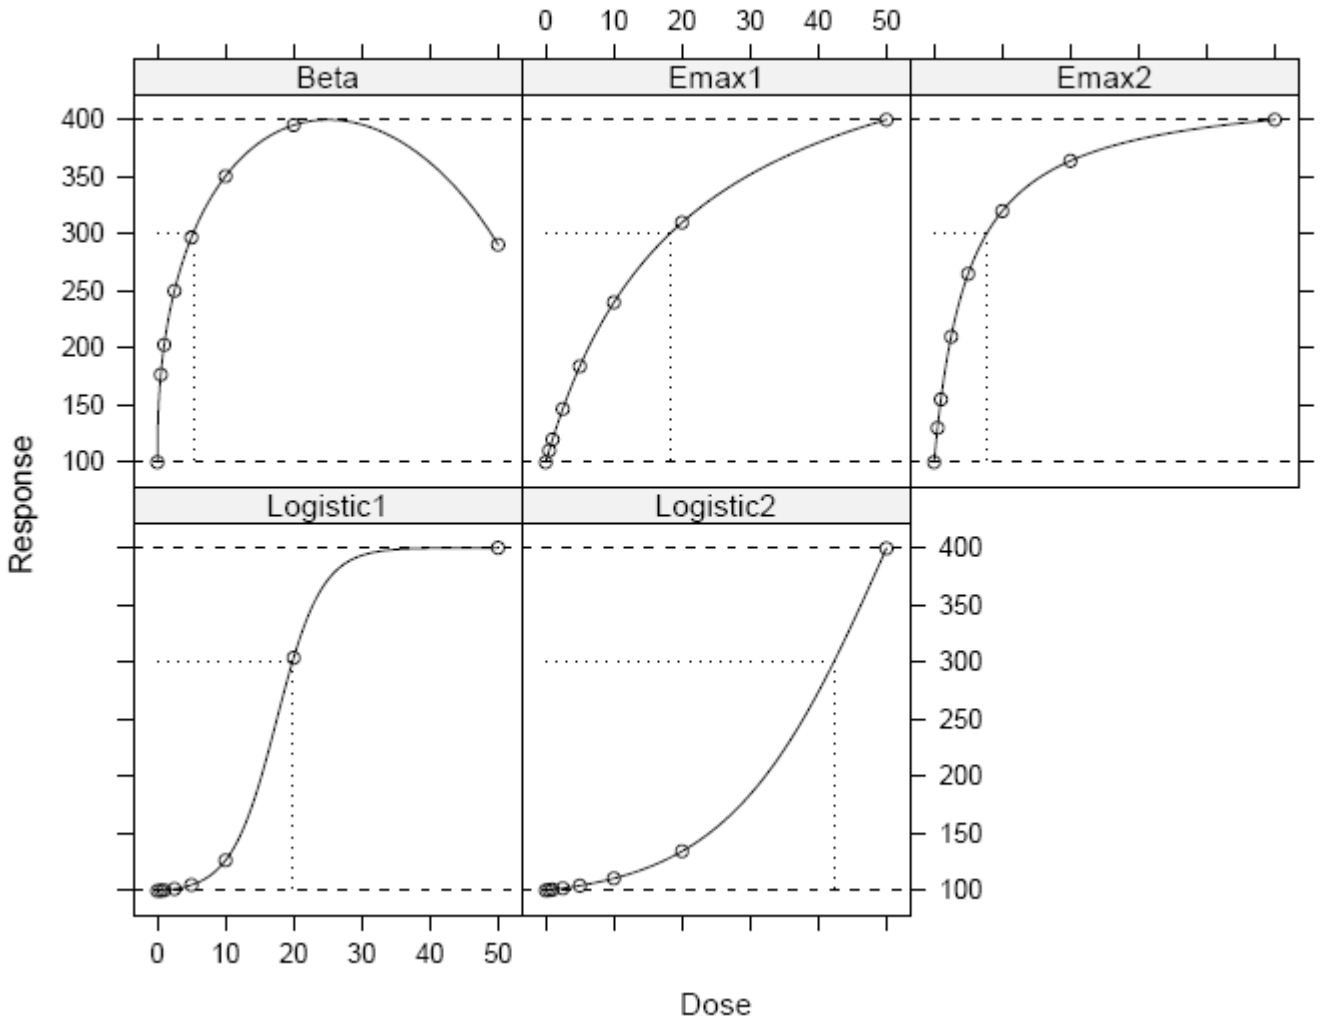

Simulation Study

- Start design: balanced allocations on the doses available
- Interim analyses chosen equally spaced in time
- Candidate Model Set:
  - Two logistic, two emax and one beta model shape (linear model not included, only a simulation scenario)
- Choice of "design" prior:
  - Uniform prior for model weights
  - Prior for $(\theta_{0m}, \theta_{1m})'$ weakly informative (variance 100000)
  - Scaled beta prior for $\boldsymbol{\theta}_m^0$ (mode: parameter guess; $S = \alpha + \beta = 3$)
- At the end of the trial use MCP-Mod procedure (based on AIC)

Simulation Results: 1) MED estimate

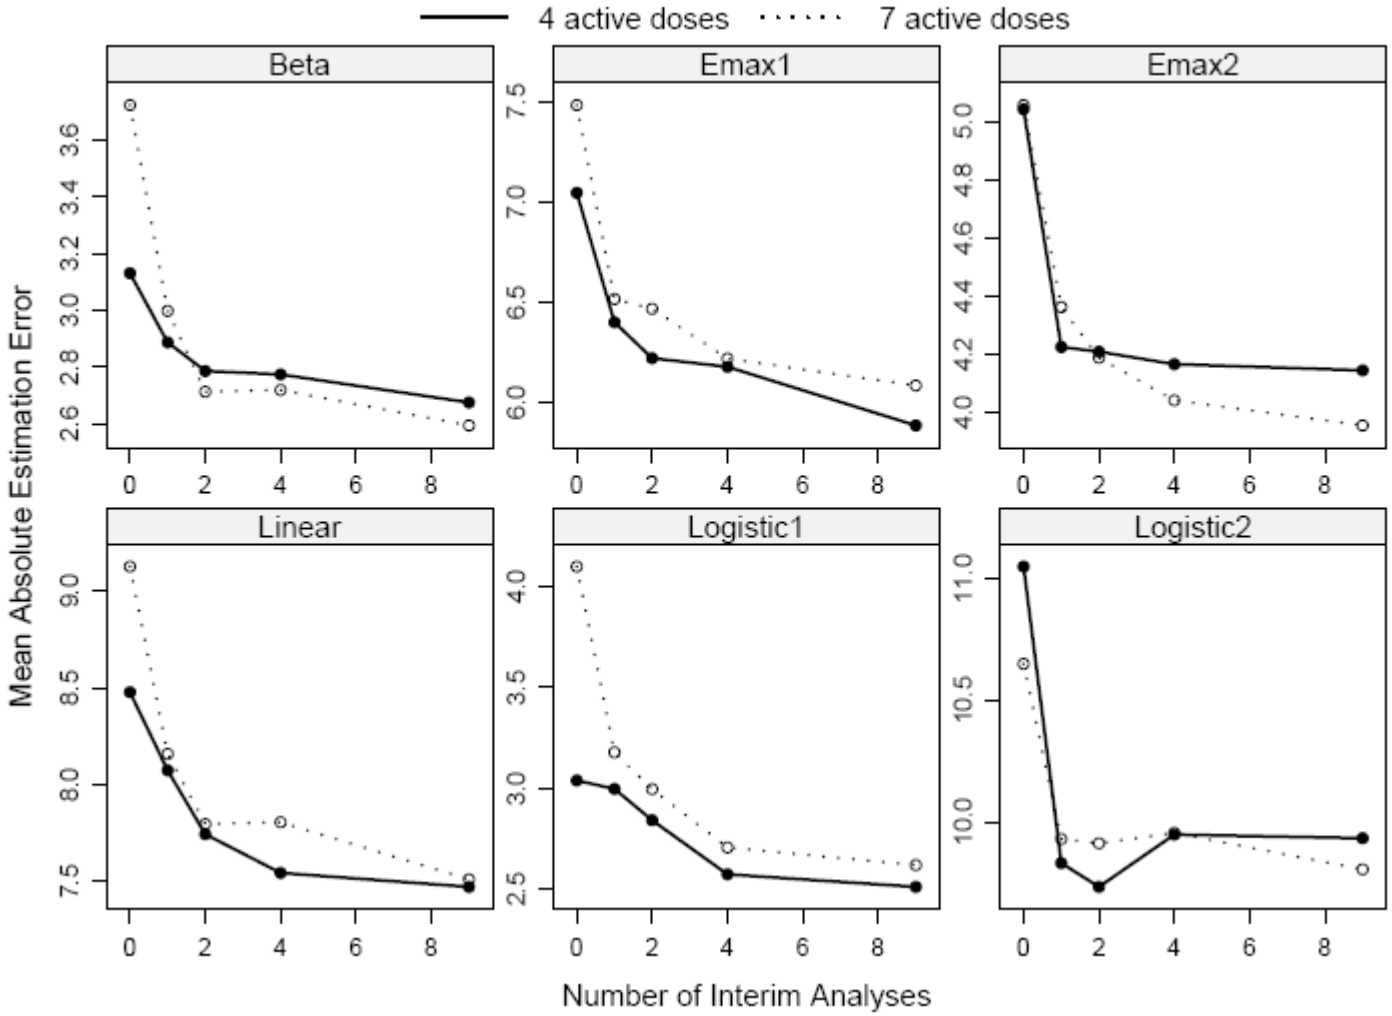

Simulation Results: 2) Patient allocations

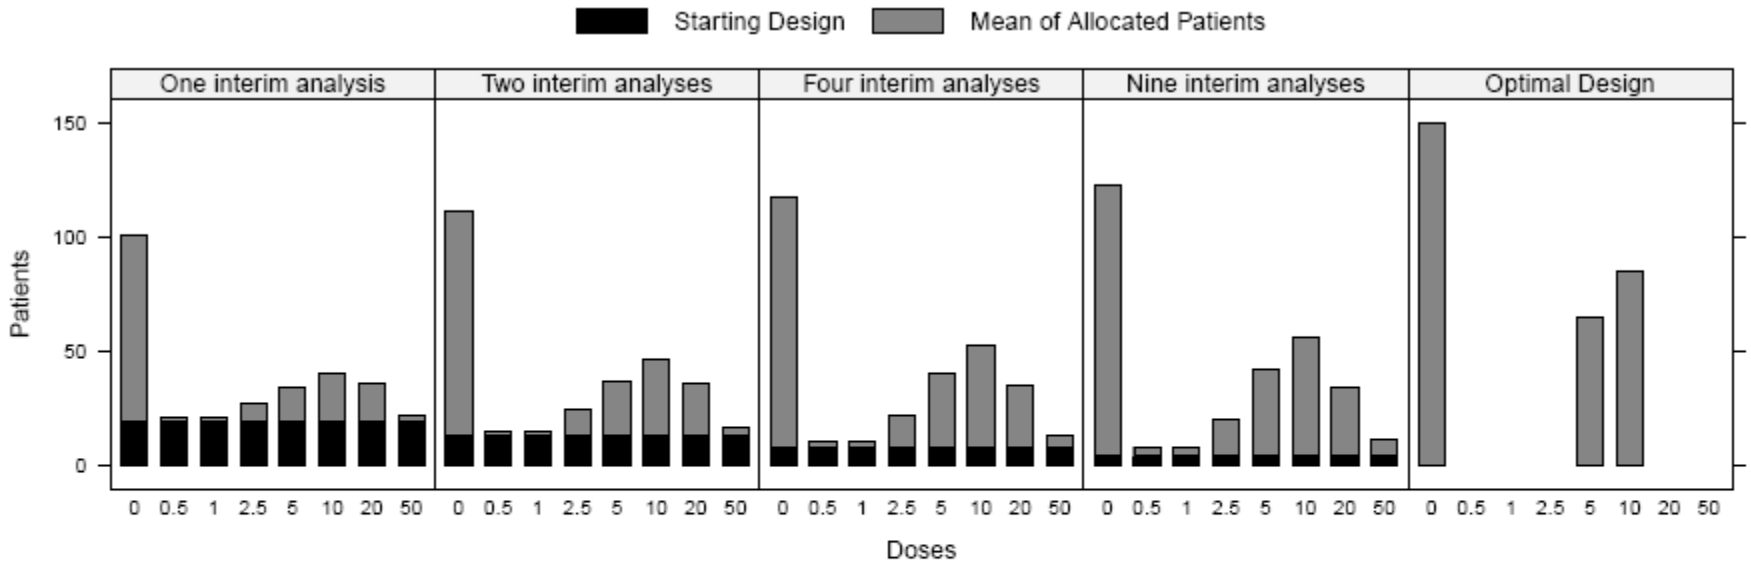

Patient allocations for Emax2 model (true MED = 7.7), right panel MED-optimal design for Emax2 model

Additional Simulation Results

Additional simulations results

- How many additional patients would be needed, for fixed design to achieve the same MED estimation precision as the adaptive design?
  For Emax2 model: 200 additional patients
- Results under increased σ: Qualitatively similar results, relative benefit of adaptive design decreases slightly
- Sensitivity analysis with respect to moderate mis-specification of prior distribution: Largely identical results

Simulation Conclusions

- Adaptation can improve a design, but there are no guarantees.
- The benefit of adaptation depends on different aspects.
- One aspect: Quality of the starting design:
  Bad starting design $\rightarrow$ Possibility to improve.
  Good starting design $\rightarrow$ Not much improvement possible.
- Another aspect: Are the dose levels properly chosen?
- Choice of the design criterion is important. Optimizing efficiency in dose estimation can decrease efficiency in other aspects.

Conclusions

- Main idea of our approach: Combine optimal design (MED optimal designs) and Bayesian ideas (shrinkage and posterior model probabilities)
- Account for inherent model uncertainty by using an averaged design criterion.
- Approach is not "fully" Bayesian, but computationally very efficient:
  Uses low dimensional quadrature and does not rely on MCMC (important for simulation based evaluation of a design).

References

- Bornkamp, B., Pinheiro, J. C. and Bretz, F. (2009). MCPMod - An R Package for the Design and Analysis of Dose-Finding Studies. *Journal of Statistical Software*, **29**(7), 1–23.
- Bretz, F., Pinheiro, J. C., and Branson, M. (2005). Combining multiple comparisons and modeling techniques in dose-response studies. *Biometrics*, **61**(3), 738–748.
- Dette H., Bretz F., Pepelyshev A., Pinheiro J.C. (2008). Optimal Designs for Dose Finding Studies. *Journal of the American Statistical Association*, **103**, 1225–1237.
- Pinheiro, J. C., Bretz, F., and Branson, M. (2006). Analysis of dose-response studies: Modeling approaches. In: *Dose Finding in Drug Development*. Ting, N. (ed.), Springer, New York, pp. 146–171.
- Pinheiro, J.C., Bornkamp, B., and Bretz, F. (2006). Design and analysis of dose finding studies combining multiple comparisons and modeling procedures. *Journal of Biopharmaceutical Statistics*, **16**(5), 639–656.