

A Statistical Software Approach to Bayesian Response-Adaptive Design for Dose Finding

Dawn Woodard and Michael O'Connell
Insightful Corporation
moconnell@insightful.com

Response Adaptive Design – Poster Overview

There are clear advantages of response adaptive design in identifying dose(s) from early phase clinical trials to carry forward in to later phases. These include avoiding giving unsafe and non-efficacious doses to subjects in trials, productivity / timeline gains, and avoidance of later stage failure due to incorrect dose selection in early stage trials.

In this poster, we describe a general statistical software framework for Bayesian response-adaptive design dose finding. The software is based on S+FlexBayes, a new S-PLUS package for flexible Bayesian analysis. The poster first describes S+FlexBayes and then the response-adaptive design methods in context. Trial design with S+SeqTrial, the S-PLUS group sequential trial design package, is covered in passing.

S+FlexBayes

S+flexBayes provides native MCMC functionality for hierarchical generalized linear models and seamless connections with external engines such as WinBUGS and OpenBUGS for additional Bayes model fitting.

The hierarchical generalized linear models include Poisson, Binomial and Gaussian and there are a wide variety of priors available on the usual GLM hyper-parameters e.g. normal, t and mixtures for the (linear) means and inv Chisquare, Inv Wishart, point mass, noninformative and uniform shrinkage for the (linear) variances / covariances. S+FlexBayes enables stratification on data clusters and specification of covariates in link functions i.e. on the tiers of the hierarchy.

S+FlexBayes enables ready fitting of a variety of mixed effects models with Markov Chain Monte Carlo methods including Metropolis-Hastings.

There is a rich inbuilt toolset for pre-processing, convergence diagnosis, summarizing / plotting the posterior distributions, and validation of the results.

S+FlexBayes Example – Orthodont data

The Orthodont data set provided with S-PLUS consist of measurements of the distance from the pituitary gland to the pterygomaxillary fissure. These measurements were taken each two years from age eight to fourteen on 16 male and 11 female children (Potthoff and Roy, 1964). An analysis of these data carried out by Pinheiro and Bates (2001) using

nlme() in S-PLUS found clear differences in the growth pattern between female and male children.

Our hierarchical linear model setup in S+FlexBayes allows us to model the variation in the growth rate due to Sex in the second level model:

$$\beta_{kj} = \alpha_{k0} + \alpha_{k1}Sex + v_{kj}$$

where $k=0,1$ correspond to the subject intercept and slope respectively. This is a more “natural” way to model the growth rate as sex-dependent. The first level model is

$$distance_{ij} = \beta_{0j} + \beta_{1j}age_{ij} + \varepsilon_{ij}$$

Note that by including a second level model there is no need to include fixed effects into the first level. The fixed effects are taken over by the population model corresponding to the posterior means of (β_{0j}, β_{1j})

S+FlexBayes Analysis

```
> beta.prior = bayes.normal(zero, identity)
> alpha.prior = bayes.nonInformative()
> error.prior = bayes.nonInfoPower(-1)
> betaCov.prior = bayes.invWishart(df= 2, scale= diag(2)/2)
```

You need to put all these priors in a single prior object

```
> ortho.prior = bhlm.prior(error.var= error.prior,
random.coef=beta.prior, level2.coef= alpha.prior, random.var=
betaCov.prior, common.error.var= 2)
```

The last term in the list, `common.error.var = 2`, specifies that all the groups share the same variance, i.e. $\sigma_j^2 = \sigma^2$, for $j=1, \dots, 23$.

To specify the likelihood (error) structure of the model, set

```
> ortho.likelihood = bhlm.likelihood(type= "normal")
```

To specify the number of chains and the way the parameters are to be initialized in each chain, create a sampler list

```
> ortho.sampler = bhlm.sampler(nBurnin = 2000,
nIter = 10000, nThin = 10, nChains= 3, init.point="prior")
```

Now you are ready to launch the fitting routine

```
> ortho.bhlm = bhlm(random.formula= distance ~ I(age- 11),
level2.formula= ~Sex, group.formula= ~Subject, data= Orthodont,
prior= ortho.prior, likelihood= ortho.likelihood, sampler =
ortho.sampler)
```

You can look at the results by typing

```
> ortho.bhlm
```

```
*****
** Results from chain 1 **
```

```

(b, a)
Random Effects:
numeric matrix: 54 rows, 3 columns.
              mean stdev Bayes Factor
(Intercept):M16 23.2479 0.6191 0.000000
I(age - 11):M16  0.6110 0.2504 0.012938
(Intercept):M05 23.2339 0.6255 0.000000
I(age - 11):M05  0.8093 0.2587 0.001449
(Intercept):M02 23.5518 0.6298 0.000000
I(age - 11):M02  0.7731 0.2501 0.002920
.....
Random effects Variance parameter [ sigma ]:
              mean stdev Bayes Factor
random:sigma:1.1 3.5511 1.2057 2.5331e-006
random:sigma:1.2 0.0537 0.2176 6.8939e-001
random:sigma:2.1 0.0537 0.2176 6.8939e-001
random:sigma:2.2 0.2071 0.0688 0.0000e+000

Level 2 (random coefficients regression) Coefficients:
              mean stdev Bayes Factor
(Int):(Int)      23.8083 0.4007 0.00000
Sex:(Int)        -1.1301 0.3977 0.00344
(Int):I(age-11)  0.6290 0.1024 0.00000
Sex:I(age - 11) -0.1517 0.1030 0.08415
Error Variance:
mean stdev Bayes Factor
Measurement Error [ sigma ] 1.315814 0.1254685 0
.....

```

It is simple to visualize the (level 2) parameter posterior densities:

```

> plot( orthodont.bhlm@model[[1]], which = "univariate", region =
"hpd", level = 0.95 )

```

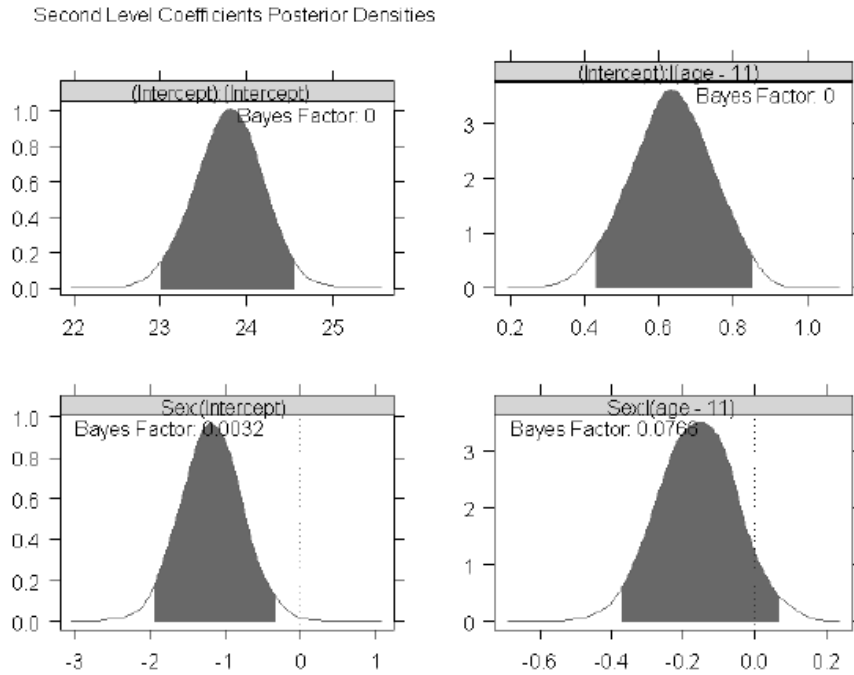


Figure 1. Posterior densities of parameters from the hierarchical linear model fit to the orthodont data using S+FlexBayes

Bayes Response-Adaptive Design

Our response-adaptive design is based on the EffTox method of Thall and Cook (2004) and involves fitting a Bayesian hierarchical model to the probability of the outcomes “toxicity”, “efficacy”, and “no effect” after each small cohort of subjects.

We use both Proportional Odds (Thall and Russell, 1998) and Continuation Ratio (Thall and Cook, 2004) models to infer probabilities of toxicity and efficacy at each dose:

Proportional Odds model:

$$\text{logit}(\pi_T) = \mu + \beta * \text{dose}$$

$$\text{logit}(\pi_{T \vee E}) = \mu + \alpha + \beta * \text{dose}$$

Continuation Ratio model:

$$\text{logit}(\pi_T) = \mu_T + \beta_T * \text{dose}$$

$$\text{logit}(\pi_{E|\sim T}) = \mu_{E|\sim T} + \beta_{E|\sim T} * \text{dose}$$

The dose for the next cohort is based on an adaptation criterion e.g. Thall and Russell (1998) select from the doses with low prob. of toxicity, the dose with the largest prob. of having high efficacy. Thall and Cook (2004) choose the dose with the highest expected “desirability” (a function of both efficacy and toxicity).

For the Thall and Russell (1998) adaptation:

Unacceptable doses have e.g.

$$\Pr(\pi_T > .1) > .95 \quad \text{or} \quad \Pr(\pi_E < .5) > .95$$

Out of the acceptable doses, choose the one that has the highest $\Pr(\pi_E > .5)$

The workflow sequence is as follows:

1. Define a set of candidate designs
 - set cohort size and max #cohorts
 - define toxicity and efficacy probability thresholds
 - set up the outer procedure for adapting the trial to the observed toxicity and efficacy probability posterior distributions
 - set up the inner models for toxicity and efficacy
2. Simulate response data from a variety of potential response models e.g. Emax, quadratic etc. as considered by Pinheiro et al. (2006).
3. Fit the toxicity and efficacy models in the designs to the simulated data
 - obtain the samples and return an object of class posterior
 - convert posterior parameter estimates to probabilities of toxicity and efficacy
 - use the posterior samples for parameter inference
4. Summarize the model results graphically for each design and simulation combination
5. Compare candidate designs across response simulation models and decide on a design for the trial e.g. in terms of robustness to simulated response models.

Bayes Response-Adaptive Design – Simulation

Step 1. Define the designs

```
#Define the cohorts
```

```
cohSize <- 3 # the cohort size  
nCohorts <- 10 # the max number of cohorts
```

```
#Define the acceptable tox / eff probabilities
```

```
#Thall and Russell (1998) dose-finding method
```

```
lowestEff <- 0.5  
highestTox <- 0.1
```

```
#Define the outer procedure thresholds
```

```
pEff <- 0.1
```

```
pTox <- 0.1
```

```
#If more than 90% of the posterior efficacy probability distribution is less than 0.5, reject the dose
```

```
#If more than 90% of the posterior toxicity probability distribution is greater than 0.1, reject the dose
```

```
#Define the inner models for efficacy and toxicity
```

```
Proportional Odds model:
```

$$\text{logit}(\pi_T) = \mu + \beta * \text{dose}$$

$$\text{logit}(\pi_T \vee E) = \mu + \alpha + \beta * \text{dose}$$

Continuation Ratio model:

$$\text{logit}(\pi_T) = \mu_T + \beta_T * \text{dose}$$

$$\text{logit}(\pi_{E|\sim T}) = \mu_{E|\sim T} + \beta_{E|\sim T} * \text{dose}$$

Step 2. Models for Efficacy and Toxicity (Truth)

Simulate Efficacy Probability(Dose) from Proportional Odds and Continuation Ratio models. Simulate toxicity as logistic model

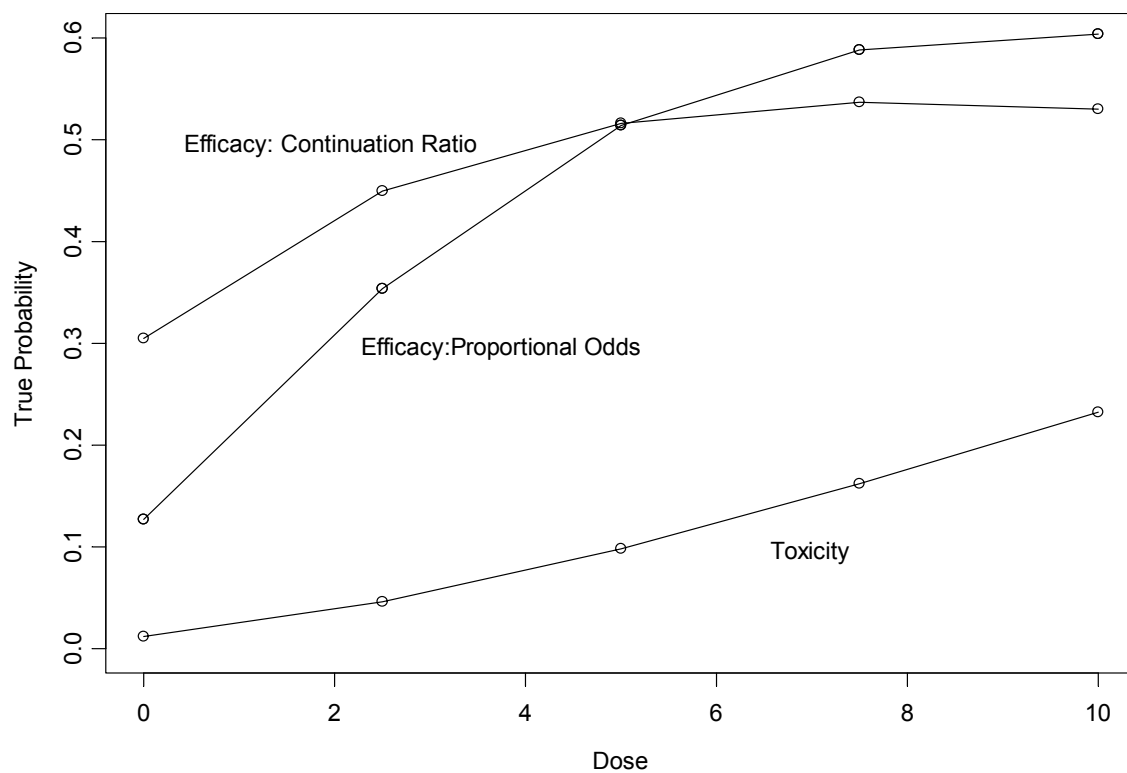


Figure 2. Simulated Efficacy and Toxicity (truth)

Step 3. Run the Simulation

```
# Initialize the patient doses and outcomes
patientDose <- numeric(0)
outcome <- numeric(0)
# Start at the minimum dose
currentDoseInd <- 1
# Initialize the cohort index
coh <- 1
# Run simulation, adaptively assigning dose
```

```

# to patients
# After each cohort, plot point/interval
# estimates of probability of efficacy and
# toxicity at each dose.

chooseNextDoseTR(patientDose, outcome, doseValues, doseTrans,
currentDoseInd, lowestEff, highestTox, pEff, pTox, coefPriors)
{
# Choose the next dose; Thall and Russell (1998)
# Run model, either the continuation ratio model
# or the proportional odds model
data <- list(patientDose = patientDose,
doseLevels = doseTrans, outcome = outcome)
if(model == "CR")
postSamples <- runEffToxCRModel(data = data,
coefPriors = coefPriors)
else if(model == "PO")
postSamples <- runEffToxPOModel(data = data,
coefPriors = coefPriors)
...
runEffToxCRModel(data, coefPriors) {
# Continuation ratio; Thall and Cook (2004),
# Fit toxicity and cond'l efficacy separately
# Assign independ normal priors for fixed coeffts
toxPrior <- bpm.prior(fixed.coef =
bayes.normal(mean = coefPriors$toxCoefMean,
cov=diag(coefPriors$toxCoefSD^2)))
...
# Specify the sampler control parameters
sampler <- bpm.sampler(nBurnin = 1000, nThin =
1, nSamples = 10000, nChains = 1)
# Fit the tox and eff-given-no-toxicity models
postTox <- bbm(fixed.formula = tox ~
patientDoseLevel, trials.formula = ~ trials,
data = toxModelData, prior = toxPrior,
sampler = sampler, engine = "OpenBUGS")
...
# Get the posterior samples from the models
betaTox <- getSamples(postTox)
...
for(j in (1:nDoses)) {
# Calculate the probability of toxicity at dose j
ptoxDoseSamp[, j] <- betaTox[, 1] +
betaTox[, 2] * data$doseLevels[j]
ptoxDoseSamp[, j] <- exp(ptoxDoseSamp[, j])/
(1 + exp(ptoxDoseSamp[, j]))
...
}
postSamples <- cbind(ptoxDoseSamp,
peffDoseSamp)
# Return Pr(tox) and Pr(eff) distn samples for
# each dose

```

```
return(posterior(postSamples))
}
```

Step 4. Summarize results for each design and simulation combination

```
...
# Plot the intervals for piEff and piTox
effToxIntervalPlot(piEffEst = piEffEst,
  piToxEst = piToxEst, piEffInterval =
  piEffInterval, piToxInterval =
  piToxInterval, lowestEff = lowestEff,
  highestTox = highestTox,
  doseValues = doseValues)
# need to have at least one acceptable dose
# to continue; or stop at max #cohorts
return(chosenDoseInd)
}
```

Step 5. Compare candidate designs across simulation models; decide on design for trial

We compared two efficacy/toxicity models (PO and CR) fit to data generated from same two models (PO and CR) under the dose-adaptive design defined in step 1. All priors, e.g. means and s.d.'s of dose-response parameters were chosen to be diffuse. We did 100 simulations and tracked the number of times that each simulation resulted in identifying the correct dose.

The CR model performed better than the PO model no matter whether the data were generated from the CR or the PO model, as shown in Table 1.

		True Model	
		Proportional Odds ($\mu = -2.5$, $\alpha = 2.5$, $\beta = 2$)	Continuation Ratio ($\mu_T = -2.5$, $\beta_T = 2$, $\mu_{E T} = 0.15$, $\beta_{E T} = 1$)
Fit Model	Proportional Odds	41 / 100	50 / 100
	Continuation Ratio	46 / 100	60 / 100

Table 1. Comparison of two efficacy/toxicity models (PO and CR) fit to data generated from same two models (PO and CR) under the dose-adaptive design defined in Step 1.


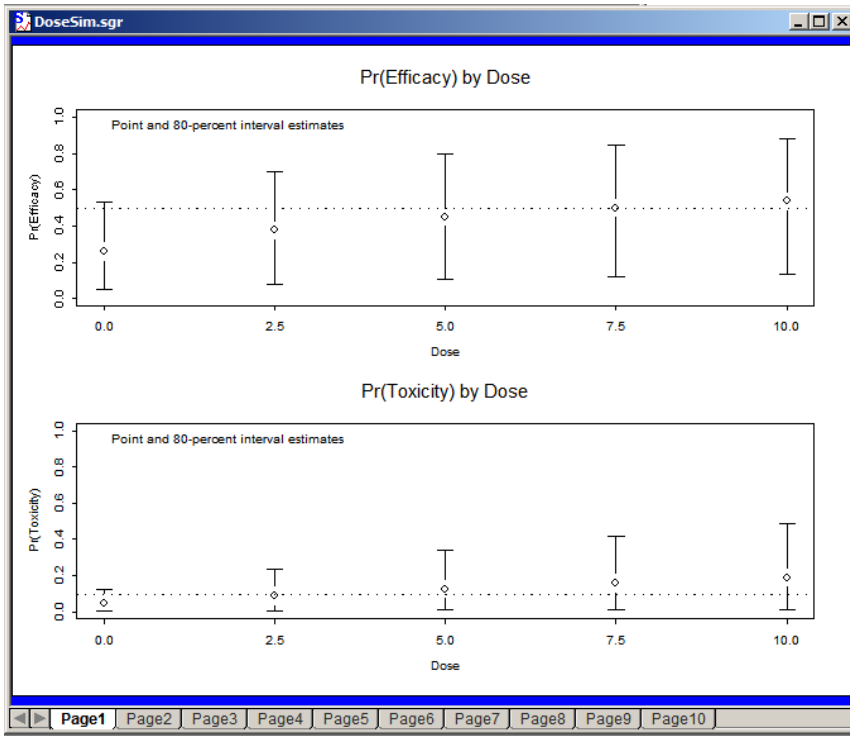


Figure 3. Posterior Efficacy and Toxicity by dose for a trial design using the adaptation method of Thall and Russell (1998). Top: CR model after 1st cohort. Bottom; CR model after 10th cohort.

Related Work

We presented a model where the response was coded as efficacious and/or toxic. If the response is time-to-event, the practical trial design is subject to the complexities of accrual, drop-out and time to event. Figure 4 shows recent work with S+SeqTrial to address this situation.

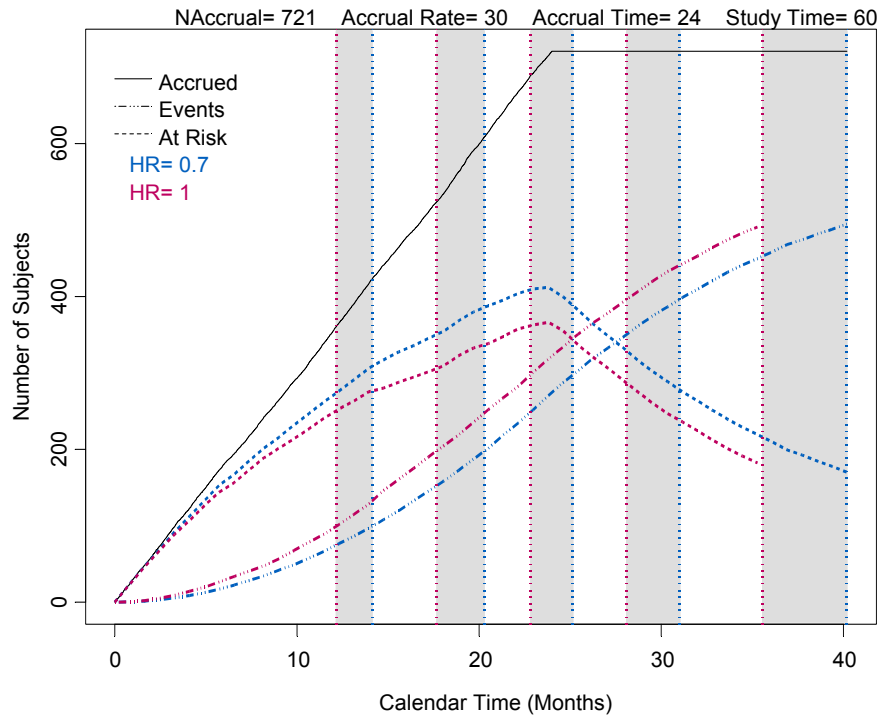


Figure 4. Adaptive-sequential design for time-to-event.

Other Applications of S+FlexBayes

Another key application of Bayes methodology is the assessment of drug safety. Adverse event data are sparse across a large number of potential preferred terms within system organ classes. The inherent issues of sparse (count, presence/absence) data and test multiplicity can be addressed directly in Bayes models. Figure 5 shows some results from such a model, as fashioned after Berry and Berry (2004), fit using S+FlexBayes, and displayed in the interactive Insightful Clinical Review Solution.

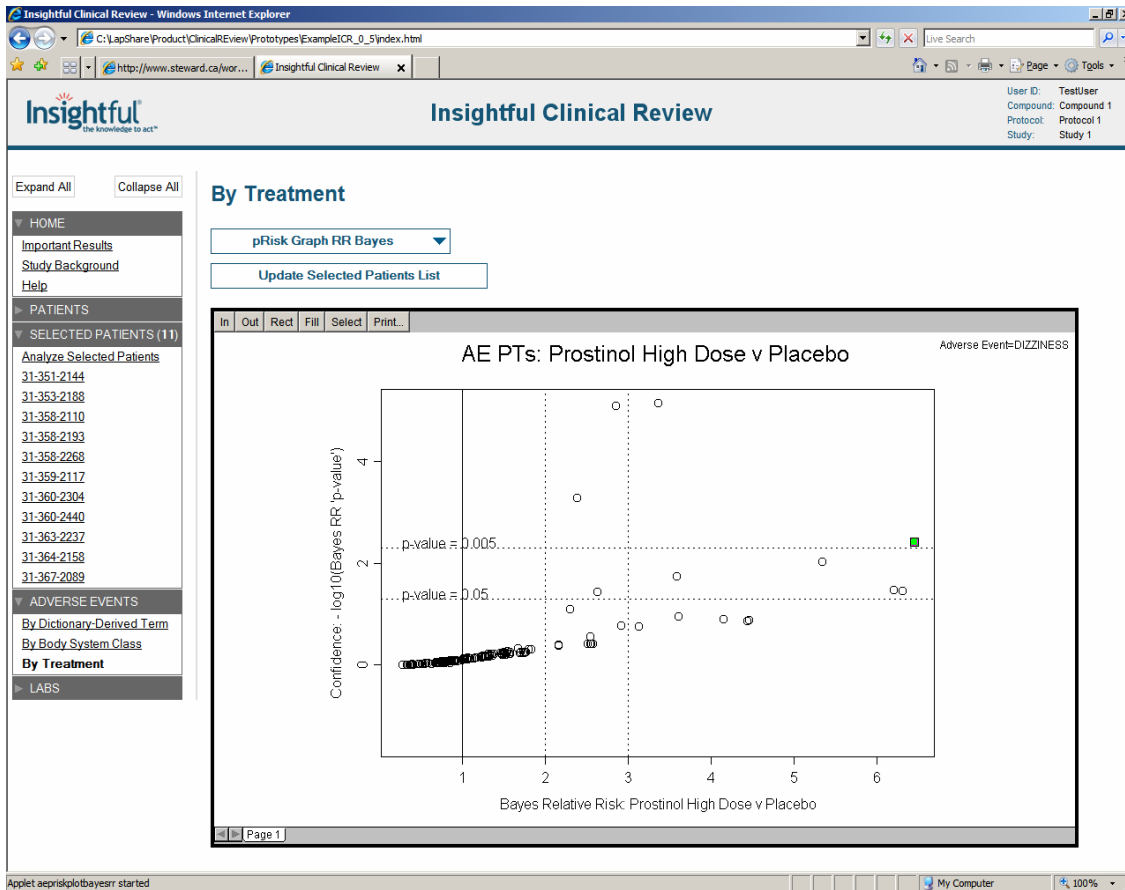


Figure 5. p-Risk plot of Adverse Events inside the Insightful Clinical Review Solution.

References

- O'Quigley, J, Hughes, MD and Fenton, T (2001). Dose-finding designs for HIV studies. *Biometrics* 57, 1018-1029.
- Pinheiro, J., Bornkamp, B., and Bretz, F. (2006). Design and analysis of dose finding studies combining multiple comparisons and modeling procedures. *J Biopharm. Statistics*, 16, 639-656.
- Thall, PF and Cook, JD (2004). Dose-finding based on efficacy-toxicity trade-offs. *Biometrics* 60, 684-693.
- Thall, PF and Russell, KE (1998). A strategy for dose-finding and safety monitoring based on efficacy and adverse outcomes in phase I/II clinical trials. *Biometrics* 54, 251-264.

Contact Information

Michael O'Connell, Ph.D.
 Director, Life Science Solutions
 Insightful Corp
 RTP NC
moconnell@insightful.com
 800-5690123 x454

© Insightful Corporation 2008. All rights reserved.

INSIGHTFUL MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS POSTER. S-PLUS, S+FlexBayes, S+SeqTrial, Insightful Clinical Review, Insightful, and the Insightful logo are registered trademarks of Insightful Corporation in the United States and/or other countries.