University of Michigan School of Public Health
The University of Michigan Department of Biostatistics Working Paper Series

Year 2012 *Paper* 92

The Bayesian Continual Reassessment Method Using a Mixture-of-Uniforms Prior

Thomas M. Braun*

*University of Michigan, tombraun@umich.edu

This working paper is hosted by The Berkeley Electronic Press (bepress) and may not be commercially reproduced without the permission of the copyright holder.

http://biostats.bepress.com/umichbiostat/paper92

Copyright ©2012 by the author.

The Bayesian Continual Reassessment Method Using a Mixture-of-Uniforms Prior

Thomas M. Braun

Abstract

Traditionally, the Bayesian formulation of the Continual Reassessment Method (CRM) is implemented with a one-parameter model describing the association of dose with the probability of dose-limiting toxicity (DLT). Determination of the appropriate value of the prior variance is often done via simulation over a grid search of possible values until suitable operating characteristics are found. However, it is under-appreciated that the operating characteristics for a given value of the prior variance vary by the "skeleton," which is the vector of a priori probabilities of DLT for each dose. The skeleton implicitly leads to a set of indifference intervals, with one interval for each dose, that contain values of the model parameter that support each dose being the MTD. To remove the need of selecting a value for the prior variance, we propose placing a uniform distribution over each of the indifference intervals, making the prior distribution for the model parameter a mixture of uniform distributions. Via simulation, we compare the operating characteristics of the CRM using a traditional continuous prior to using the mixture-of-uniforms prior in a variety of settings and show that the mixture-of-uniforms prior leads to operating characteristics that are less sensitive to the chosen skeleton.

The Bayesian Continual Reassessment Method Using a Mixture-of-Uniforms Prior

Thomas M. Braun
University of Michigan School of Public Health

May 17, 2012

Abstract

Traditionally, the Bayesian formulation of the Continual Reassessment Method (CRM) is implemented with a one-parameter model describing the association of dose with the probability of dose-limiting toxicity (DLT). Determination of the appropriate value of the prior variance is often done via simulation over a grid search of possible values until suitable operating characteristics are found. However, it is under-appreciated that the operating characteristics for a given value of the prior variance vary by the \skeleton," which is the vector of a priori probabilities of DLT for each dose. The skeleton implicitly leads to a set of indi erence intervals, with one interval for each dose, that contain values of the model parameter that support each dose being the MTD. To remove the need of selecting a value for the prior variance, we propose placing a uniform distribution over each of the indi erence intervals, making the prior distribution for the model parameter a mixture of uniform distributions. Via simulation, we compare the operating characteristics of the CRM using a traditional continuous prior to using the mixture-of-uniforms prior in a variety of settings and show that the mixture-of-uniforms prior leads to operating characteristics that are less sensitive to the chosen skeleton.

Keywords : Dose- nding; Dose-escalation; Phase I trial; Variance calibration; Indifference intervals

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

1 Background and Signi cance

Phase I clinical trials are typically the rst studies in humans of agents that have shown promise for the treatment of cancer in pre-clinical animal studies. The goal of most Phase I clinical trials is conceptually simple: to determine which doses of the agent can be given to patients before an unacceptable fraction, $\theta < 1$, of patients begins to experience dose-limiting toxicities (DLTs). The largest dose studied that maintains a DLT fraction closest to is de ned as the maximum tolerated dose (MTD). DLTs are very generally de ned as any Grade 3 (severe) or Grade 4 (life-threatening) non-hematologic side-e ect or any Grade 4 hematologic side-e ect, as de ned by Common Toxicity Criterion (CTC) set forth by the National Cancer Institute (NCI). Commonly used values of vary in the interval $(0.20, 0.35)$, depending upon the nature of the speci c DLTs de ned in the trial.

Given a pre-determined set of doses under study, the rst patient, or cohort of patients, are usually assigned to the lowest dose. Assuming the number of observed DLTs remains low, subsequent patients will be progressively assigned, or escalated, to higher doses. As the trial progresses, eventually a dose will be assigned that results in DLTs occurring at a rate felt to be unacceptable. At this point, the trial could be stopped, with the MTD identi ed at one of the other assigned doses, or subsequent patients could be de-escalated, or assigned to previously studied doses, with the possibility of later re-escalation in future patients, with hopes of collecting more information as to where the MTD is located.

A variety of Phase I trial designs exist regarding how doses should be escalated and de-escalated during the trial. Using the terms given by Tourneau et al. (2009) and others, existing trial designs are either \rule-based" or \model-based." Rule-based designs determine the dose assignment for a patient or cohort of patients from the number of DLTs observed in one or more of the most recently enrolled patients. The primary motivation of rule-based designs is their simplicity for non-statistical investigators: once the rule is in place, no mathematical computation beyond the proportion of observed DLTs is needed. Although there are many rule-based designs such as the biased-coin design (BCD) (Stylianou and

COBRA
Collection of Biostatistics
Research Archive

http://biostats.bepress.com/umichbiostat/paper92

Flournoy, 2002) and the k-in-a-row design (Oron and Ho , 2009), the most prevalent rule-based design in actual oncology Phase I trials is the 3+3 method; see Lin and Shih (2001) for a more general discussion. In the 3+3 method, patients are enrolled in cohorts of three, with each member of a cohort assigned to the same dose. If no patients in a cohort experience DLT, the next cohort is assigned to the next higher dose. If two or three patients in a cohort experience DLT, the trial is terminated, with the MTD de ned at one dose below the assignment of the current cohort. If one member of the cohort experiences DLT, another cohort of three patients is assigned to the same dose. If no additional DLTs occur, the next cohort is assigned to the next higher dose. However, if at least one additional DLT occurs, the trial is terminated with the MTD de ned at one dose below the assignment of the current cohort. Although the 3+3 design is widely used because it is easy to administer, Ahn (1998), Storer (1989), and Kornet al. (1994), among others, have shown that the 3+3 design tends to both identify the MTD at, and assign a large percentage of patients to, doses below the MTD.

The genesis of model-based Phase I trial designs was publication of the Continual Re-assessment Method (CRM) of O'Quigley, Pepe, and Fisher (1990), which was later modi- ed by Faries (1994) and Goodmanet al. (1995). The CRM assumes that the probability of DLT is related to dose d through a parametric function f (; d) that is monotonic in dose. Two of the most commonly-used models are a one-parameter logistic regression model f (d;) = exp(3 + d)=[1 + exp(3 + d)] and the so-called \empiric" modelf (; d) = d . A prior distribution with mean and variance ^{2} is placed on and the posterior mean of is sequentially updated as each enrolled subject or cohort of subjects is observed for the occurrence of DLT. The posterior mean of is then plugged back intof (; d) to compute an updated probability of DLT for each dose. The dose with updated DLT probability closest to is then assigned to the next cohort of patients. Once all patients have been observed for the possibilty of DLT, a nal updated DLT probability for each dose is computed, and the MTD is de ned as the dose whose updated probability is closest to. Although the CRM

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

Hosted by The Berkeley Electronic Press

was rst proposed using Bayesian methods, a maximum-likelihood version of the CRM for estimation of was later proposed by O'Quigley and Shen (1996). An excellent tutorial on the CRM was published by Garrett-Mayer (2006) and a recent survey of the CRM and later extensions was published by O'Quigley and Conaway (2010).

One of the challenges of the one-parameter CRM model is determining an appropriate value for ^{2}. It is well-known that a very large, usually termed \non-informative," value of ^{2}, which is espoused with Bayesian analyses of large datasets, leads to poor performance of the CRM at the beginning of the trial when few patients have been enrolled and little data is available for computation of the posterior distribution of . Thus, it is better to use a slightly informative value for ^{2} so that the the prior distribution drives the computation of the posterior distribution of at the beginning of the study but allows for the data to drive the computation of the posterior distribution of near the end of the study.

However, the de nition of \slightly informative" is subjective and usually has to be determined through simulations over a range of possible values of ^{2} in a variety of settings to determine which value or range of values for ^{2} leads to the best operating characteristics of the CRM over all settings. One must also bear in mind that the value chosen for ^{2} is dependent upon the skeleton adopted for the CRM, so that if a di erent skeleton were used, the operating characteristics when using ^{2} would change with the new skeleton. Real examples of this phenomenon can be found in Moller (1995) and Neuenschwander et al. (2008), in which the supposed \aggressive" nature of the CRM can be fully explained by the fact that the value of ^{2} used by each was too small for the skeleton each adopted.

A more systematic determination of ^{2} was recently developed by Lee and Cheung (2011), who de ned a \least-informative" value of ^{2} that is speci c to: (a) the form of f (d;) used, (b) the a priori DLT probabilities for each dose, known as the \skeleton," and (c) the targeted DLT rate . Knowledge of these three values allows us to divide the positive real line into non-overlapping indi erence intervals (Cheung and Chappell, 2002; Lee and Cheung, 2009), one for each dose d_j; $j = 1; 2; : : : ; J$ that de nes a range of values for for which dose d_j

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

http://biostats.bepress.com/umichbiostat/paper92

would be the MTD. The width of interval j de nes an a priori probability that dose d_j is the MTD and the value of ^{2} that leads to a uniform distribution across the doses is selected as the least-informative variance, denoted $^2_{LI}$.

Our methods choose to use the indi erence intervals in a slightly di erent fashion. Once each indi erence interval has been identi ed, we assume a priori that conditional on being in each interval, has a uniform distribution over that interval, so that the marginal prior distribution for is a mixture of uniform distributions. Through use of this prior, we avoid having to specify the prior variance ^{2}. And given that di erent skeletons require di erent values of ^{2} to obtain equal operating characteristics for the CRM, we expect that our mixture-of-uniforms prior will reduce the variation in the performance of the CRM among di erent skeletons. The Bayesian model averaging CRM (BMA-CRM) of Yin and Yuan (2009a) and the recent work of Daimon et al. (2011) were both motivated by the desire to reduce the impact of skeleton choice on the performance of the CRM. However, both methods are much more computationally intensive than our approach, which we denote the uniform mixture CRM (UM-CRM).

We emphasize that we do not necessarily expect the UM-CRM to better identify the MTD or assign more patients to the MTD than the traditional CRM among all skeletons, although it is likely that the UM-CRM will perform better than the traditional CRM with some skeletons (and vice versa). We explain the speci cs of indi erence intervals and the details of our methods in Section 2, and we examine the performance of our methods via simulation in Section 3. Concluding remarks can be found in Section 4.

2 Proposed Methods

We have a Phase I trial planning to enroll a maximum of N patients, each of whom will be assigned to one of J increasing doses. Each dose $j = 1; 2; : : : J$, has a probability of DLT, p_j, with a corresponding skeleton value p_j supplied by investigators; the targeted DLT rate is . We assume that the p_j can be modeled with a smooth function of dose $f(D_j;)$, where

Hosted by The Berkeley Electronic Press

The Bayesian Continual Reassessment Method Using a Mixture-of-Uniforms Prior

Thomas M. Braun
University of Michigan School of Public Health

May 17, 2012

Abstract

Traditionally, the Bayesian formulation of the Continual Reassessment Method (CRM) is implemented with a one-parameter model describing the association of dose with the probability of dose-limiting toxicity (DLT). Determination of the appropriate value of the prior variance is often done via simulation over a grid search of possible values until suitable operating characteristics are found. However, it is under-appreciated that the operating characteristics for a given value of the prior variance vary by the \skeleton," which is the vector of a priori probabilities of DLT for each dose. The skeleton implicitly leads to a set of indi erence intervals, with one interval for each dose, that contain values of the model parameter that support each dose being the MTD. To remove the need of selecting a value for the prior variance, we propose placing a uniform distribution over each of the indi erence intervals, making the prior distribution for the model parameter a mixture of uniform distributions. Via simulation, we compare the operating characteristics of the CRM using a traditional continuous prior to using the mixture-of-uniforms prior in a variety of settings and show that the mixture-of-uniforms prior leads to operating characteristics that are less sensitive to the chosen skeleton.

Keywords : Dose- nding; Dose-escalation; Phase I trial; Variance calibration; Indifference intervals

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

http://biostat.bepress.com/cobra/ptm/paper93
Hosted by The Berkeley Electronic Press

The Bayesian Continual Reassessment Method Using a Mixture-of-Uniforms Prior

Thomas M. Braun
University of Michigan School of Public Health

May 17, 2012

Abstract

Traditionally, the Bayesian formulation of the Continual Reassessment Method (CRM) is implemented with a one-parameter model describing the association of dose with the probability of dose-limiting toxicity (DLT). Determination of the appropriate value of the prior variance is often done via simulation over a grid search of possible values until suitable operating characteristics are found. However, it is under-appreciated that the operating characteristics for a given value of the prior variance vary by the \skeleton," which is the vector of a priori probabilities of DLT for each dose. The skeleton implicitly leads to a set of indi erence intervals, with one interval for each dose, that contain values of the model parameter that support each dose being the MTD. To remove the need of selecting a value for the prior variance, we propose placing a uniform distribution over each of the indi erence intervals, making the prior distribution for the model parameter a mixture of uniform distributions. Via simulation, we compare the operating characteristics of the CRM using a traditional continuous prior to using the mixture-of-uniforms prior in a variety of settings and show that the mixture-of-uniforms prior leads to operating characteristics that are less sensitive to the chosen skeleton.

Keywords : Dose- nding; Dose-escalation; Phase I trial; Variance calibration; Indifference intervals

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

http://biostat.bepress.com/cobra/ptm/paper93
Hosted by The Berkeley Electronic Press

The Bayesian Continual Reassessment Method Using a Mixture-of-Uniforms Prior

Thomas M. Braun
University of Michigan School of Public Health

May 17, 2012

Abstract

Traditionally, the Bayesian formulation of the Continual Reassessment Method (CRM) is implemented with a one-parameter model describing the association of dose with the probability of dose-limiting toxicity (DLT). Determination of the appropriate value of the prior variance is often done via simulation over a grid search of possible values until suitable operating characteristics are found. However, it is under-appreciated that the operating characteristics for a given value of the prior variance vary by the \skeleton," which is the vector of a priori probabilities of DLT for each dose. The skeleton implicitly leads to a set of indi erence intervals, with one interval for each dose, that contain values of the model parameter that support each dose being the MTD. To remove the need of selecting a value for the prior variance, we propose placing a uniform distribution over each of the indi erence intervals, making the prior distribution for the model parameter a mixture of uniform distributions. Via simulation, we compare the operating characteristics of the CRM using a traditional continuous prior to using the mixture-of-uniforms prior in a variety of settings and show that the mixture-of-uniforms prior leads to operating characteristics that are less sensitive to the chosen skeleton.

Keywords : Dose- nding; Dose-escalation; Phase I trial; Variance calibration; Indifference intervals

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

http://biostat.bepress.com/cobra/ptm/paper93
Hosted by The Berkeley Electronic Press

The Bayesian Continual Reassessment Method Using a Mixture-of-Uniforms Prior

Thomas M. Braun
University of Michigan School of Public Health

May 17, 2012

Abstract

Traditionally, the Bayesian formulation of the Continual Reassessment Method (CRM) is implemented with a one-parameter model describing the association of dose with the probability of dose-limiting toxicity (DLT). Determination of the appropriate value of the prior variance is often done via simulation over a grid search of possible values until suitable operating characteristics are found. However, it is under-appreciated that the operating characteristics for a given value of the prior variance vary by the \skeleton," which is the vector of a priori probabilities of DLT for each dose. The skeleton implicitly leads to a set of indi erence intervals, with one interval for each dose, that contain values of the model parameter that support each dose being the MTD. To remove the need of selecting a value for the prior variance, we propose placing a uniform distribution over each of the indi erence intervals, making the prior distribution for the model parameter a mixture of uniform distributions. Via simulation, we compare the operating characteristics of the CRM using a traditional continuous prior to using the mixture-of-uniforms prior in a variety of settings and show that the mixture-of-uniforms prior leads to operating characteristics that are less sensitive to the chosen skeleton.

Keywords : Dose- nding; Dose-escalation; Phase I trial; Variance calibration; Indifference intervals

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

http://biostat.bepress.com/cobra/ptm/paper93
Hosted by The Berkeley Electronic Press

The Bayesian Continual Reassessment Method Using a Mixture-of-Uniforms Prior

Thomas M. Braun
University of Michigan School of Public Health

May 17, 2012

Abstract

Traditionally, the Bayesian formulation of the Continual Reassessment Method (CRM) is implemented with a one-parameter model describing the association of dose with the probability of dose-limiting toxicity (DLT). Determination of the appropriate value of the prior variance is often done via simulation over a grid search of possible values until suitable operating characteristics are found. However, it is under-appreciated that the operating characteristics for a given value of the prior variance vary by the \skeleton," which is the vector of a priori probabilities of DLT for each dose. The skeleton implicitly leads to a set of indi erence intervals, with one interval for each dose, that contain values of the model parameter that support each dose being the MTD. To remove the need of selecting a value for the prior variance, we propose placing a uniform distribution over each of the indi erence intervals, making the prior distribution for the model parameter a mixture of uniform distributions. Via simulation, we compare the operating characteristics of the CRM using a traditional continuous prior to using the mixture-of-uniforms prior in a variety of settings and show that the mixture-of-uniforms prior leads to operating characteristics that are less sensitive to the chosen skeleton.

Keywords : Dose- nding; Dose-escalation; Phase I trial; Variance calibration; Indifference intervals

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

http://biostat.bepress.com/cobra/ptm/paper93
Hosted by The Berkeley Electronic Press

The Bayesian Continual Reassessment Method Using a Mixture-of-Uniforms Prior

Thomas M. Braun
University of Michigan School of Public Health

May 17, 2012

Abstract

Traditionally, the Bayesian formulation of the Continual Reassessment Method (CRM) is implemented with a one-parameter model describing the association of dose with the probability of dose-limiting toxicity (DLT). Determination of the appropriate value of the prior variance is often done via simulation over a grid search of possible values until suitable operating characteristics are found. However, it is under-appreciated that the operating characteristics for a given value of the prior variance vary by the \skeleton," which is the vector of a priori probabilities of DLT for each dose. The skeleton implicitly leads to a set of indi erence intervals, with one interval for each dose, that contain values of the model parameter that support each dose being the MTD. To remove the need of selecting a value for the prior variance, we propose placing a uniform distribution over each of the indi erence intervals, making the prior distribution for the model parameter a mixture of uniform distributions. Via simulation, we compare the operating characteristics of the CRM using a traditional continuous prior to using the mixture-of-uniforms prior in a variety of settings and show that the mixture-of-uniforms prior leads to operating characteristics that are less sensitive to the chosen skeleton.

Keywords : Dose- nding; Dose-escalation; Phase I trial; Variance calibration; Indifference intervals

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

http://biostat.bepress.com/cobra/ptm/paper93
Hosted by The Berkeley Electronic Press

The Bayesian Continual Reassessment Method Using a Mixture-of-Uniforms Prior

Thomas M. Braun
University of Michigan School of Public Health

May 17, 2012

Abstract

Traditionally, the Bayesian formulation of the Continual Reassessment Method (CRM) is implemented with a one-parameter model describing the association of dose with the probability of dose-limiting toxicity (DLT). Determination of the appropriate value of the prior variance is often done via simulation over a grid search of possible values until suitable operating characteristics are found. However, it is under-appreciated that the operating characteristics for a given value of the prior variance vary by the \skeleton," which is the vector of a priori probabilities of DLT for each dose. The skeleton implicitly leads to a set of indi erence intervals, with one interval for each dose, that contain values of the model parameter that support each dose being the MTD. To remove the need of selecting a value for the prior variance, we propose placing a uniform distribution over each of the indi erence intervals, making the prior distribution for the model parameter a mixture of uniform distributions. Via simulation, we compare the operating characteristics of the CRM using a traditional continuous prior to using the mixture-of-uniforms prior in a variety of settings and show that the mixture-of-uniforms prior leads to operating characteristics that are less sensitive to the chosen skeleton.

Keywords : Dose- nding; Dose-escalation; Phase I trial; Variance calibration; Indifference intervals

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

http://biostat.bepress.com/cobra/ptm/paper93
Hosted by The Berkeley Electronic Press

The Bayesian Continual Reassessment Method Using a Mixture-of-Uniforms Prior

Thomas M. Braun
University of Michigan School of Public Health

May 17, 2012

Abstract

Traditionally, the Bayesian formulation of the Continual Reassessment Method (CRM) is implemented with a one-parameter model describing the association of dose with the probability of dose-limiting toxicity (DLT). Determination of the appropriate value of the prior variance is often done via simulation over a grid search of possible values until suitable operating characteristics are found. However, it is under-appreciated that the operating characteristics for a given value of the prior variance vary by the \skeleton," which is the vector of a priori probabilities of DLT for each dose. The skeleton implicitly leads to a set of indi erence intervals, with one interval for each dose, that contain values of the model parameter that support each dose being the MTD. To remove the need of selecting a value for the prior variance, we propose placing a uniform distribution over each of the indi erence intervals, making the prior distribution for the model parameter a mixture of uniform distributions. Via simulation, we compare the operating characteristics of the CRM using a traditional continuous prior to using the mixture-of-uniforms prior in a variety of settings and show that the mixture-of-uniforms prior leads to operating characteristics that are less sensitive to the chosen skeleton.

Keywords : Dose- nding; Dose-escalation; Phase I trial; Variance calibration; Indifference intervals

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

http://biostat.bepress.com/cobra/ptm/paper93
Hosted by The Berkeley Electronic Press

The Bayesian Continual Reassessment Method Using a Mixture-of-Uniforms Prior

Thomas M. Braun
University of Michigan School of Public Health

May 17, 2012

Abstract

Traditionally, the Bayesian formulation of the Continual Reassessment Method (CRM) is implemented with a one-parameter model describing the association of dose with the probability of dose-limiting toxicity (DLT). Determination of the appropriate value of the prior variance is often done via simulation over a grid search of possible values until suitable operating characteristics are found. However, it is under-appreciated that the operating characteristics for a given value of the prior variance vary by the \skeleton," which is the vector of a priori probabilities of DLT for each dose. The skeleton implicitly leads to a set of indi erence intervals, with one interval for each dose, that contain values of the model parameter that support each dose being the MTD. To remove the need of selecting a value for the prior variance, we propose placing a uniform distribution over each of the indi erence intervals, making the prior distribution for the model parameter a mixture of uniform distributions. Via simulation, we compare the operating characteristics of the CRM using a traditional continuous prior to using the mixture-of-uniforms prior in a variety of settings and show that the mixture-of-uniforms prior leads to operating characteristics that are less sensitive to the chosen skeleton.

Keywords : Dose- nding; Dose-escalation; Phase I trial; Variance calibration; Indifference intervals

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

http://biostat.bepress.com/cobra/ptm/paper93
Hosted by The Berkeley Electronic Press

The Bayesian Continual Reassessment Method Using a Mixture-of-Uniforms Prior

Thomas M. Braun
University of Michigan School of Public Health

May 17, 2012

Abstract

Traditionally, the Bayesian formulation of the Continual Reassessment Method (CRM) is implemented with a one-parameter model describing the association of dose with the probability of dose-limiting toxicity (DLT). Determination of the appropriate value of the prior variance is often done via simulation over a grid search of possible values until suitable operating characteristics are found. However, it is under-appreciated that the operating characteristics for a given value of the prior variance vary by the \skeleton," which is the vector of a priori probabilities of DLT for each dose. The skeleton implicitly leads to a set of indi erence intervals, with one interval for each dose, that contain values of the model parameter that support each dose being the MTD. To remove the need of selecting a value for the prior variance, we propose placing a uniform distribution over each of the indi erence intervals, making the prior distribution for the model parameter a mixture of uniform distributions. Via simulation, we compare the operating characteristics of the CRM using a traditional continuous prior to using the mixture-of-uniforms prior in a variety of settings and show that the mixture-of-uniforms prior leads to operating characteristics that are less sensitive to the chosen skeleton.

Keywords : Dose- nding; Dose-escalation; Phase I trial; Variance calibration; Indifference intervals

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

http://biostat.bepress.com/cobra/ptm/paper93
Hosted by The Berkeley Electronic Press

The Bayesian Continual Reassessment Method Using a Mixture-of-Uniforms Prior

Thomas M. Braun
University of Michigan School of Public Health

May 17, 2012

Abstract

Traditionally, the Bayesian formulation of the Continual Reassessment Method (CRM) is implemented with a one-parameter model describing the association of dose with the probability of dose-limiting toxicity (DLT). Determination of the appropriate value of the prior variance is often done via simulation over a grid search of possible values until suitable operating characteristics are found. However, it is under-appreciated that the operating characteristics for a given value of the prior variance vary by the \skeleton," which is the vector of a priori probabilities of DLT for each dose. The skeleton implicitly leads to a set of indi erence intervals, with one interval for each dose, that contain values of the model parameter that support each dose being the MTD. To remove the need of selecting a value for the prior variance, we propose placing a uniform distribution over each of the indi erence intervals, making the prior distribution for the model parameter a mixture of uniform distributions. Via simulation, we compare the operating characteristics of the CRM using a traditional continuous prior to using the mixture-of-uniforms prior in a variety of settings and show that the mixture-of-uniforms prior leads to operating characteristics that are less sensitive to the chosen skeleton.

Keywords : Dose- nding; Dose-escalation; Phase I trial; Variance calibration; Indifference intervals

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

http://biostat.bepress.com/cobra/ptm/paper93
Hosted by The Berkeley Electronic Press

The Bayesian Continual Reassessment Method Using a Mixture-of-Uniforms Prior

Thomas M. Braun
University of Michigan School of Public Health

May 17, 2012

Abstract

Traditionally, the Bayesian formulation of the Continual Reassessment Method (CRM) is implemented with a one-parameter model describing the association of dose with the probability of dose-limiting toxicity (DLT). Determination of the appropriate value of the prior variance is often done via simulation over a grid search of possible values until suitable operating characteristics are found. However, it is under-appreciated that the operating characteristics for a given value of the prior variance vary by the \skeleton," which is the vector of a priori probabilities of DLT for each dose. The skeleton implicitly leads to a set of indi erence intervals, with one interval for each dose, that contain values of the model parameter that support each dose being the MTD. To remove the need of selecting a value for the prior variance, we propose placing a uniform distribution over each of the indi erence intervals, making the prior distribution for the model parameter a mixture of uniform distributions. Via simulation, we compare the operating characteristics of the CRM using a traditional continuous prior to using the mixture-of-uniforms prior in a variety of settings and show that the mixture-of-uniforms prior leads to operating characteristics that are less sensitive to the chosen skeleton.

Keywords : Dose- nding; Dose-escalation; Phase I trial; Variance calibration; Indifference intervals

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

http://biostat.bepress.com/cobra/ptm/paper93
Hosted by The Berkeley Electronic Press

The Bayesian Continual Reassessment Method Using a Mixture-of-Uniforms Prior

Thomas M. Braun
University of Michigan School of Public Health

May 17, 2012

Abstract

Traditionally, the Bayesian formulation of the Continual Reassessment Method (CRM) is implemented with a one-parameter model describing the association of dose with the probability of dose-limiting toxicity (DLT). Determination of the appropriate value of the prior variance is often done via simulation over a grid search of possible values until suitable operating characteristics are found. However, it is under-appreciated that the operating characteristics for a given value of the prior variance vary by the \skeleton," which is the vector of a priori probabilities of DLT for each dose. The skeleton implicitly leads to a set of indi erence intervals, with one interval for each dose, that contain values of the model parameter that support each dose being the MTD. To remove the need of selecting a value for the prior variance, we propose placing a uniform distribution over each of the indi erence intervals, making the prior distribution for the model parameter a mixture of uniform distributions. Via simulation, we compare the operating characteristics of the CRM using a traditional continuous prior to using the mixture-of-uniforms prior in a variety of settings and show that the mixture-of-uniforms prior leads to operating characteristics that are less sensitive to the chosen skeleton.

Keywords : Dose- nding; Dose-escalation; Phase I trial; Variance calibration; Indifference intervals

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

http://biostat.bepress.com/cobra/ptm/paper93
Hosted by The Berkeley Electronic Press

The Bayesian Continual Reassessment Method Using a Mixture-of-Uniforms Prior

Thomas M. Braun
University of Michigan School of Public Health

May 17, 2012

Abstract

Traditionally, the Bayesian formulation of the Continual Reassessment Method (CRM) is implemented with a one-parameter model describing the association of dose with the probability of dose-limiting toxicity (DLT). Determination of the appropriate value of the prior variance is often done via simulation over a grid search of possible values until suitable operating characteristics are found. However, it is under-appreciated that the operating characteristics for a given value of the prior variance vary by the \skeleton," which is the vector of a priori probabilities of DLT for each dose. The skeleton implicitly leads to a set of indi erence intervals, with one interval for each dose, that contain values of the model parameter that support each dose being the MTD. To remove the need of selecting a value for the prior variance, we propose placing a uniform distribution over each of the indi erence intervals, making the prior distribution for the model parameter a mixture of uniform distributions. Via simulation, we compare the operating characteristics of the CRM using a traditional continuous prior to using the mixture-of-uniforms prior in a variety of settings and show that the mixture-of-uniforms prior leads to operating characteristics that are less sensitive to the chosen skeleton.

Keywords : Dose- nding; Dose-escalation; Phase I trial; Variance calibration; Indifference intervals

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

http://biostat.bepress.com/cobra/ptm/paper93
Hosted by The Berkeley Electronic Press

The Bayesian Continual Reassessment Method Using a Mixture-of-Uniforms Prior

Thomas M. Braun
University of Michigan School of Public Health

May 17, 2012

Abstract

Traditionally, the Bayesian formulation of the Continual Reassessment Method (CRM) is implemented with a one-parameter model describing the association of dose with the probability of dose-limiting toxicity (DLT). Determination of the appropriate value of the prior variance is often done via simulation over a grid search of possible values until suitable operating characteristics are found. However, it is under-appreciated that the operating characteristics for a given value of the prior variance vary by the \skeleton," which is the vector of a priori probabilities of DLT for each dose. The skeleton implicitly leads to a set of indi erence intervals, with one interval for each dose, that contain values of the model parameter that support each dose being the MTD. To remove the need of selecting a value for the prior variance, we propose placing a uniform distribution over each of the indi erence intervals, making the prior distribution for the model parameter a mixture of uniform distributions. Via simulation, we compare the operating characteristics of the CRM using a traditional continuous prior to using the mixture-of-uniforms prior in a variety of settings and show that the mixture-of-uniforms prior leads to operating characteristics that are less sensitive to the chosen skeleton.

Keywords : Dose- nding; Dose-escalation; Phase I trial; Variance calibration; Indifference intervals

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

http://biostat.bepress.com/cobra/ptm/paper93
Hosted by The Berkeley Electronic Press

The Bayesian Continual Reassessment Method Using a Mixture-of-Uniforms Prior

Thomas M. Braun
University of Michigan School of Public Health

May 17, 2012

Abstract

Traditionally, the Bayesian formulation of the Continual Reassessment Method (CRM) is implemented with a one-parameter model describing the association of dose with the probability of dose-limiting toxicity (DLT). Determination of the appropriate value of the prior variance is often done via simulation over a grid search of possible values until suitable operating characteristics are found. However, it is under-appreciated that the operating characteristics for a given value of the prior variance vary by the \skeleton," which is the vector of a priori probabilities of DLT for each dose. The skeleton implicitly leads to a set of indi erence intervals, with one interval for each dose, that contain values of the model parameter that support each dose being the MTD. To remove the need of selecting a value for the prior variance, we propose placing a uniform distribution over each of the indi erence intervals, making the prior distribution for the model parameter a mixture of uniform distributions. Via simulation, we compare the operating characteristics of the CRM using a traditional continuous prior to using the mixture-of-uniforms prior in a variety of settings and show that the mixture-of-uniforms prior leads to operating characteristics that are less sensitive to the chosen skeleton.

Keywords : Dose- nding; Dose-escalation; Phase I trial; Variance calibration; Indifference intervals

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

http://biostat.bepress.com/cobra/ptm/paper93
Hosted by The Berkeley Electronic Press

The Bayesian Continual Reassessment Method Using a Mixture-of-Uniforms Prior

Thomas M. Braun
University of Michigan School of Public Health

May 17, 2012

Abstract

Traditionally, the Bayesian formulation of the Continual Reassessment Method (CRM) is implemented with a one-parameter model describing the association of dose with the probability of dose-limiting toxicity (DLT). Determination of the appropriate value of the prior variance is often done via simulation over a grid search of possible values until suitable operating characteristics are found. However, it is under-appreciated that the operating characteristics for a given value of the prior variance vary by the \skeleton," which is the vector of a priori probabilities of DLT for each dose. The skeleton implicitly leads to a set of indi erence intervals, with one interval for each dose, that contain values of the model parameter that support each dose being the MTD. To remove the need of selecting a value for the prior variance, we propose placing a uniform distribution over each of the indi erence intervals, making the prior distribution for the model parameter a mixture of uniform distributions. Via simulation, we compare the operating characteristics of the CRM using a traditional continuous prior to using the mixture-of-uniforms prior in a variety of settings and show that the mixture-of-uniforms prior leads to operating characteristics that are less sensitive to the chosen skeleton.

Keywords : Dose- nding; Dose-escalation; Phase I trial; Variance calibration; Indifference intervals

COBRA
A BEPRESS REPOSITORY
Collection of Biostatistics
Research Archive

http://biostat.bepress.com/cobra/ptm/paper93
Hosted by The Berkeley Electronic Press