Bayesian Model Averaging Continual Reassessment Method BMA-CRM

Guosheng Yin and Ying Yuan

August 26, 2009

This document provides the statistical background for the Bayesian model averaging continual reassessment method (BMA-CRM). The BMA-CRM is a Bayesian model-based phase I clinical trial design. The primary goal of the BMA-CRM is to identify the maximum tolerated dose (MTD) of a new drug, which is typically defined as the dose with a dose-limiting toxicity (DLT) probability that is closest to the target toxicity rate. We assume that the DLT is recorded as a binary outcome, and monotonically increases with respect to the dose level.

1 CRM

We first introduce the continual reassessment method (CRM) of O'Quigley, Pepe and Fisher (1990). The CRM is a model-based dose-finding approach that uses a single unknown parameter to link the true toxicity probabilities with the prespecified toxicity probabilities corresponding to the prior mean toxicity probability set.

More specifically, let $(d_1, \ldots, d_J)$ denote a set of J prespecified doses of the new drug under investigation. Let ϕ be the target toxicity rate specified by investigators. The CRM, assumes a working dose-toxicity model, such as

$$\text{pr}(\text{toxicity at } d_j) = \pi_j(\alpha) = p_j^{\exp(\alpha)} \tag{1.1}$$

for $j = 1, \ldots, J$, where p_j is a prespecified constant, and α is an unknown parameter. A

normal prior distribution $N(0, \sigma^2)$ is often assigned to α,

$$f(\alpha) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{\alpha^2}{2\sigma^2}\right),$$

with $\sigma^2 = 2$.

To determine the constants p_j's in the probability model (1.1), the user is often asked to first specify the prior mean probability of toxicity, say $\{s_1, \ldots, s_J\}$, at each of the J dose levels under consideration. These prior mean toxicity probabilities should satisfy the monotonic toxicity order

$$0 < s_1 < s_2 < \ldots < s_J < 1.$$

Then, the values of the p_j's are computed through

$$E\{p_j^{\exp(\alpha)}\} = \int p_j^{\exp(\alpha)} f(\alpha) d\alpha = s_j. \tag{1.2}$$

That is, we solve for the values of the p_j's so that the prior mean probability of toxicity at each dose is the elicited value by investigators.

During the trial, the unknown parameter α is continuously updated based on the accumulating data in order to identify the dose with a given target toxicity level. Each new cohort of patients is assigned to the dose with an estimated toxicity probability closest to the prespecified target thus far.

Suppose that among n_j patients treated at dose level j, y_j patients have experienced DLT. Let D denote the observed data, $D = \{(n_j, y_j), j = 1, \ldots, J\}$. Based on the binomial distribution for the toxicity outcome, the likelihood function is given by

$$L(D|\alpha) \propto \prod_{j=1}^{J} \{p_j^{\exp(\alpha)}\}^{y_j} \{1 - p_j^{\exp(\alpha)}\}^{n_j - y_j}.$$

Following Bayes' theorem, we can estimate the toxicity probabilities using the corresponding posterior means of $\pi_j(\alpha)$,

$$\hat{\pi}_j = \int p_j^{\exp(\alpha)} \frac{L(D|\alpha) f(\alpha)}{\int L(D|\alpha) f(\alpha) d\alpha} d\alpha.$$

After updating the posterior estimates of the toxicity probabilities at all of the doses considered, the recommended dose for the next cohort of patients is the one that has a toxicity probability closest to the target ϕ. Thus, a new cohort of patients is assigned to dose level j^* such that

$$j^* = \text{argmin}_{j\in(1,\ldots,J)}|\hat{\pi}_j - \phi|,$$

with a restriction that untried doses cannot be skipped. The trial continues until the exhaustion of the total sample size, and then the dose with a posterior toxicity probability closest to ϕ is selected as the MTD.

For safety, a stopping rule is implemented to protect against the possibility that even the lowest dose is too toxic. If the posterior probability that the lowest dose is more toxic than the target is higher than a prespecified threshold, say 0.9, then the trial will stop.

2 BMA-CRM

A major issue associated with the CRM is the arbitrariness in the prespecification of the prior mean toxicity probabilities $\{s_1, \ldots, s_J\}$. Due to lack of toxicity information on a new drug, investigators may have quite different opinions on the prior mean toxicity probabilities.

To avoid subjectivity in the specification of $\{s_1, \ldots, s_J\}$, the BMA-CRM prespecifies multiple sets of mean toxicity probabilities. Each set of prior mean probabilities of toxicity corresponds to a CRM model of the form (1.1). During the trial, conditional on the observed data, these different models usually yield different estimates of the toxicity probabilities $(\hat{\pi}_1, \ldots, \hat{\pi}_J)$. Some of them may be close to the true values, while others may not, depending on how well the models fit the accumulating data. To accommodate the uncertainty in the specification of these CRMs, we take a Bayesian model averaging (BMA) approach to average the $\hat{\pi}_j$'s across the CRM models to obtain the BMA estimate of the toxicity probability for dose level j. In other words, we incorporate the uncertainty in the prespecification of the

toxicity probabilities into the estimation procedure such that the potential estimation bias caused by a misspecification of the s_j's can be averaged out.

Let $(M_1, \ldots, M_K)$ be the models corresponding to each set of prior mean toxicity probabilities $\{(s_{11}, \ldots, s_{1J}), \ldots, (s_{K1}, \ldots, s_{KJ})\}$; and correspondingly, based on (1.2), we can solve for $\{(p_{11}, \ldots, p_{1J}), \ldots, (p_{K1}, \ldots, p_{KJ})\}$. Model M_k $(k = 1, \ldots, K)$ in the CRM is given by

$$\pi_{kj}(\alpha_k) = p_{kj}^{\exp(\alpha_k)}, \quad j = 1, \ldots, J.$$

Let $\mathrm{pr}(M_k)$ be the prior probability that model M_k is the true model. If there is no preference a priori for any single model in the CRM case, we can assign equal weights to different CRM models by simply setting $\mathrm{pr}(M_k) = 1/K$. At a certain stage of the trial, based on the observed data $D = \{(n_j, y_j), j = 1, \ldots, J\}$, the likelihood function under model M_k is

$$L(D|\alpha_k, M_k) \propto \prod_{j=1}^{J} \{p_{kj}^{\exp(\alpha_k)}\}^{y_j} \{1 - p_{kj}^{\exp(\alpha_k)}\}^{n_j - y_j}.$$

The posterior model probability for M_k is given by

$$\mathrm{pr}(M_k|D) = \frac{L(D|M_k)\mathrm{pr}(M_k)}{\sum_{i=1}^{K} L(D|M_i)\mathrm{pr}(M_i)}$$

where $L(D|M_k)$ is the marginal likelihood of model M_k,

$$L(D|M_k) = \int L(D|\alpha_k, M_k) f(\alpha_k|M_k) d\alpha_k,$$

α_k is the power parameter in the CRM associated with model M_k, and $f(\alpha_k|M_k)$ is the prior distribution of α_k under model M_k.

The BMA estimate for the toxicity probability at each dose level is given by

$$\bar{\pi}_j = \sum_{k=1}^{K} \hat{\pi}_{kj} \mathrm{pr}(M_k|D), \quad j = 1, \ldots, J, \tag{2.1}$$

where $\hat{\pi}_{kj}$ is the posterior mean of the toxicity probability at dose level j under model M_k, i.e.,

$$\hat{\pi}_{kj} = \int p_{kj}^{\exp(\alpha_k)} \frac{L(D|\alpha_k, M_k) f(\alpha_k|M_k)}{\int L(D|\alpha_k, M_k) f(\alpha_k|M_k) d\alpha_k} d\alpha_k.$$

By assigning $\hat{\pi}_{kj}$ a weight of $\text{pr}(M_k|D)$, the BMA method automatically identifies and favors the best fitting model, thus $\bar{\pi}_j$ is always close to the best estimate. Therefore, the decision of dose escalation or de-escalation in the trial is based upon $\bar{\pi}_j$ as opposed to $\hat{\pi}_{kj}$.

In addition, we implement a safety rule: if the posterior probability that the lowest dose is more toxic than the target is higher than a prespecified threshold, say 0.9, i.e.,

$$\sum_{k=1}^{K} \text{pr}\{\pi_{k1}(\alpha_k) > \phi | M_k, D\}\text{pr}(M_k|D) > 0.9. \tag{2.2}$$

then the trial will stop.

3 BMA-CRM software

The BMA-CRM software provides a user-friendly interface for practitioners to use the BMA-CRM methodology to design phase I clinical trials. The software consists of three main components:

(1) trial design model parameters;

(2) simulation run, which is used to assess the operating characteristics of the BMA-CRM by simulating trials under various practical scenarios;

(3) trial conduct, which is used to carry out the actual trial.

Specifically, the BMA-CRM software requires users to specify the following desing parameters.

- Maximum sample size.
- Cohort size. In practice, patients are often treated in cohorts, and the most common cohort size is three, although users could set the cohort size to one or two.
- Target toxicity probability, i.e., the value of ϕ.

- Safety stopping probability, which is the threshold of the safety rule as described in (2.2). The default value 0.9 is a reasonable value for most of the cases.

- Number of doses (i.e., J).

- Starting dose level. For safety, we recommend starting the trial from the lowest dose level, although users may start the trial at the physician specified dose level. The default starting dose level is 1.

- Number of probability sets (i.e., K). The BMA-CRM method requires prespecifying multiple sets of prior means of the toxicity probabilities. Two to five sets are reasonable, and the default number of probability sets is 3. When $K = 1$, the BMA-CRM reduces to the standard CRM.

- Toxicity probabilities in each probability set, i.e., the values of $K \times J$ prior mean toxicity probabilities $\{(s_{11}, \ldots, s_{1J}), \ldots, (s_{K1}, \ldots, s_{KJ})\}$. We recommend that these sets should be specified to represent different prior opinions on toxicity probabilities of the doses under investigation. For example, the first set may represent an aggressive prior guess by specifying a low dose level as the target dose, whereas the second and third sets stand for neutral and conservative prior opinions, respectively, by specifying an intermediate or high dose level as the target dose. The graph of dose-toxicity curves provided in the software is designed to help users achieve this goal. The basic principle of specifying these toxicity probability sets is that they should be diversified enough to cover potentially possible dose-toxicity profiles. As long as one of the probability sets is close to the true dose-toxicity profile, the BMA-CRM should yield a near-optimal performance.

Simulations under various practical scenarios provide a very useful tool to assess the performance of the design. To conduct a simulation study, users are further required to

specify the number of simulated trials, the seed used to generate random numbers for the simulation, and the true toxicity probability at each dose level. We can simulate, say 1,000 trials, to examine on average, whether the BMA-CRM can select the target dose with a substantially high percentage. Because the underlying true dose-toxicity curve is unknown, we should simulate a variety of scenarios to cover the range of cases in reality.

To use the BMA-CRM in practice, an important assumption is that the toxicity outcome needs to be observed shortly after the initiation of the treatment.

REFERENCES

O'Quigley, J., Pepe, M., and Fisher, L. (1990). Continual reassessment method: A practical design for phase I clinical trials in cancer. *Biometrics* **46**, 33-48.

Yin, G. and Yuan, Y. (2009). Bayesian model averaging continual reassessment method in phase I clinical trials. *Journal of the American Statistical Association*. In press.