

PubH 7470: STATISTICS FOR TRANSLATIONAL & CLINICAL RESEARCH



PHASE I CLINICAL TRIALS:

CONTINUAL REASSESSMENT METHOD (CRM)

When designing cancer clinical trials for development and evaluation of therapeutic interventions, two special aspects must be taken into consideration: (1) the target population (cancer patients) and (2) the fact that all anti-cancer drugs under investigation are cytotoxic agents.

It is true that newer drugs are safer – and more efficient, but toxicity is still a very serious concern – especially in cancer drug testing.

Potential therapeutic agents for cancer treatment can induce severe safety concern - even at lower dose levels; they can generate severe toxicities than most of pharmaceutical agents for treatment of other diseases. Possible adverse effects may be irreversible, even fatal. As a result, phase I trials are conducted only on cancer patients - not healthy volunteers.

Patients in cancer clinical trials are those with malignant tumors; most cancers are life-threatening and the disease process is usually irreversible. Patients in phase I trials are mostly **terminal cancer patients who have failed all standard therapies** and for whom the new anti-tumor agent being tested may be the **last hope**

At this stage, the investigator is facing a classic, fundamental **dilemma**: it's a conflict of scientific versus ethical intent.

We put some patients at risk for their own **benefits and benefits of others**; and it's an unavoidable conflict because these patients have failed under all standard therapies.

We need to reconcile the risks of toxicity to patients with the potential benefit to these same patients and make an efficient design to use no more patients than necessary.

Thus, the phase I cancer trial may be viewed as a problem in optimization: maximizing the dose-toxicity evaluation, while minimizing the number of patients treated .

The primary scientific objective of the evaluation of new chemotherapeutic agents in cancer patients in phase I trials is to employ an efficient dose-finding design to reach “the maximum dose with an acceptable and manageable safety profile” for use in subsequent phase II trials. The most commonly used design is the “standard design”, or some of its variations – including the fast-track when toxicity is less severe.

“STANDARD DESIGN”

- Start at lowest dose, enroll groups of 3 pts
 - (i) Move up if none of first 3 have toxicity;
 - (ii) If two or three patients have dose-limiting toxicity (DLT), stop.
 - (iii) If one of 3 has toxicity, enroll 3 more at same dose and move up if none of the second cohort have toxicity. Otherwise, stop
- After moving up, repeat the process at new dose

MAXIMUM TOLERATED DOSE

- If dose escalation not possible, i.e. being stop at a given dose, this dose is considered as “above MTD”.
- When a dose is judged as above the MTD, the **next lower dose** is often declared the MTD (Optional: six patients have been treated with no more than 1 showing DLT; that is dose de-escalation is an option).

ST DESIGN: COMMON CRITIQUES

- Patients enter early are likely treated sub-optimally; may be we need to move up faster (against the principle of “good medicine”!)
- Only few patients left when MTD reached, not enough to estimate MTD’s toxicity rate (against the principle of “good statistics”!)

The standard design is not robust; expected rate of the selected MTD is strongly influenced by the doses used. If the trial is such that there are many dose levels below the MTD then the standard design will choose a dose far too low with greater probability than if there are fewer dose levels below the MTD.

Standard Design is **SAFE**, i.e. few patients are exposed to and died because of toxicities.

However, “**safe**” does not necessarily mean “**good**”; if not given enough medication, the patient would be killed by the cancer/disease.

In addition:

- (1) It does not target a particular toxicity rate associated with MTD.
- (2) It does not make use of all available toxicity data; escalation rule depends solely of toxicity outcomes of the current dose.

However, back to the very basic aim of “reaching the maximum dose with an acceptable and manageable safety profile”; what is considered “acceptable”. How high is high? How do you know if you have no target ; and this is even more serious because the result is not robust – actually few investigators never knew what is the toxicity rate of their resulting MTDs.

STATISTICAL FORMULATION

The MTD could be statistically interpreted as some “**percentile**” of a tolerance distribution or dose-response curve in terms of the presence or absence of the DLT. In other words, the MTD is the dose at which a specified proportion of patients, say, θ , experiencing DLT. Storer (1997) indicated that the value of θ is usually in the range from .1 to .4. In the previous lecture, we used .3-.4 or 30% to 40%. There is no magic number, it depends on the severity of the side effects and if they are treatable.

It's kind of strange but in employing the Standard Design, investigators set no goal for θ , the proportion of patients experiencing DLT (that's why we studied to see what it was in a previous lecture). The result is implicit that with the stopping of "2-out-of-6" in Standard or "3+3" design; it's usually 30%-40% - but, again, the design is not robust.

A STATISTICAL MODEL

Let Y be the binary response such that $Y=1$ denote the occurrence of a pre-defined DLT and $\{d_i ; i= 1,2,\dots,n\}$ be a set of fixed dose levels used in a phase I trial. Let

$$p(\mathbf{x}) = \Pr[Y=1|\mathbf{x}] \ \&$$

$$\text{logit } [p(\mathbf{x})] = \log \{p(\mathbf{x})/[1-p(\mathbf{x})]\};$$

The relationship between Y and dose level d could be described by the logistic model:

$$\text{logit } [p(\mathbf{x})] = \alpha + \beta\mathbf{x},$$

where x is usually the log of the dose d ; or x is d

In general, one can consider $p(x) = F(\alpha + \beta x)$ where $F(\cdot)$ is some cumulative distribution function (cdf). For example, in addition to the logistic, we may consider the “probit” model. Some even consider one-parameter family - such as the “hyperbolic tangent” or “exponential”.

SUPPORTS FOR LOGISTIC MODEL

The fit and the origin of the linear logistic model could be easily traced as follows. When a dose y of an agent is applied to a pharmacological system, the fractions f_a and f_u of the system affected and unaffected satisfy the so-called “median effect principle” (Chou, 1976):

$$\frac{f_a}{f_u} = \left\{ \frac{y}{ED_{50}} \right\}^m$$

where ED_{50} is the “median effective dose” and “ m ” is a Hill-type coefficient; $m = 1$ for first-degree or Michaelis-Menten system. The median effect principle has been investigated much very thoroughly in pharmacology.

If we set “ $\pi = f_a$ ”, the median effect principle and the logistic regression model are completely identical with a slope $\beta_1 = m$.

MAXIMUM TOLERATED DOSE (MTD)

Let $\mathbf{l}_\theta = \text{logit}(\theta)$, then the MTD is defined by $\mathbf{x}_\theta = \log(\text{MTD}) = (\mathbf{l}_\theta - \alpha) / \beta$ (it's likely not be one of the doses used in a phase I trial under standard design).

$$\underline{\text{Model:}} \quad \text{logit} [p(\mathbf{x})] = \alpha + \beta \mathbf{x}$$

Note that if we have some estimates of the toxicity rates associated with the dose levels used in a phase I standard design, we could estimate to toxicity rate of the resulting MTD and compare to the θ of choice. Of course, the exact rates are not known, but clinicians should have some estimates; otherwise, it would be difficult for the investigator to “justify” the selected dose levels.

The process **could** consist of these steps:

(i) choose the “maximum tolerated level” θ ,

(ii) choose a design and calculate its MTD’s expected toxicity rate r_0 , and

(iii) compare the calculated expected rate r_0 to the selected level θ ; then

(iv) Do it again if needed: trial by error

Besides many elements of arbitrariness (choosing the level θ for the problem and estimating the rates r_i 's for the planned doses); the basic problem is, according to Storer (1989, 1993), the standard dose escalation design “frequently failed to provide a convergent estimate of MTD” (so, even if we know what we want, i.e. θ , **the standard design might not get us there**). The alternative is a newer design, “continual reassessment method”.

The primary objective of phase I trials is to find the maximum dose, called MTD, with an acceptable and manageable safety profile for use in subsequent phase II trials. But that's the investigator's objective, not the patient's. Patients in phase I trials are mostly terminal cancer patients for which the new anti-tumor agent being tested may be the last hope. Designs, such as **“standard design”**, **do not serve them - at least not ethically.**

According to “principle of “good medicine”, the patient should be treated with the best treatment the doctor knows. Patients enter early to a Phase I trial with Standard Design are likely treated sub-optimally; they receive a treatment level that the attending physician knows to be inferior. Some of these patients would likely die before any other therapy can be attempted. The newer design, the “continual reassessment method (CRM)” is an attempt to correct that by **giving each patient a better chance of a favorable response.**

In addition to the attempt to treat each patient more ethically, the CRM also updates the information of “the dose-response relationship” as observations on DLT become available and then to use this information to concentrate the next step of the trial around the dose that might correspond to the anticipated target toxicity level. It does so using a Bayesian framework, even though it has been argued that the CRM could be explained by likelihood approach.

The CRM is very attractive and has fostered a heated debate or debates which last for more than a decade. There are many variations of the CRM, we'll describe here a scheme based on a specific prior; the principle and the process are the same if another model is selected.

Step 1: Choose the “maximum tolerated level” θ , the toxicity rate at the recommended dose level or MTD’s (say, $\theta=.33$ or whatever); this is a basic difference with standard design (SD).

Step 2: Choose a fixed number of patients to be enrolled; usually $n = 19-24$; this is another difference with SD (where the number of patients needed is variable).

Step 3: The CRM uses binary response (DLT or not); Let Y be the binary response such that $Y=1$ denote the occurrence of a pre-defined DLT. Let $p(x) = \Pr[Y=1|x]$ and

$$\text{logit } [p(x)] = \log \{p(x)/[1-p(x)]\}$$

The next step is to choose a statistical model representing the relationship between Y and dose level; for example, it could be described by the logistic (or probit) model :

$$\text{logit } [p(x)] = \alpha + \beta x$$

where x is the log of the dose d ; or x is dose d .

Step 4: Use the baseline

response/toxicity, adverse-effect rate (dose = 0) to calculate and fix the “intercept” α .

Step 5: Under the Bayesian framework, choose a prior distribution for the “slope” β ; for example, “unit exponential” - one with probability density function $g(\beta) = \exp(-\beta)$

Step 6: From the model: $\text{logit} [p(x); \beta] = \alpha + \beta x$, with β placed at “the prior mean” and set $p(x)$ equal to the target rate θ , solve for dose x . This is **dose for the first patient**, a dose determined to reflect the current belief of the investigator/doctor as the dose level that produces the probability of DLT closest to the target rate θ - the “maximum dose with an acceptable and manageable safety”. This step fits the “principle of good medicine”! - the patient is treated at the MTD.

Step 7: After the first patient's toxicity/adverse-effect result becomes available, the “posterior distribution” of β is calculated and the posterior mean of β is substituted in $\text{logit} [p(x); \beta] = \alpha + \beta x$.

The next patient is treated at the dose level x whose probability $p(x)$ is the target rate θ (with calculated posterior mean of β). **This step is repeated in subsequent patients** every time toxicity/adverse-effect result becomes available and the posterior distribution of β is re-calculated.

There are more than one ways to calculate the “posterior mean”, one of which can proceed as follows – without going through the posterior distribution:

From the model “logit $[p(x_i; \beta)] = \alpha + \beta x_i$ ”, the (Bernoulli) likelihood at dose i is:

$$L(\beta) = p(x_i; \beta)^{\delta_i} [1-p(x_i; \beta)]^{1-\delta_i}$$

And the mean β of is calculated from:

$$E(\beta) = \frac{\int \beta L(\beta) g(\beta) d\beta}{\int L(\beta) g(\beta) d\beta}$$

Finally the MTD is estimated as the dose level for the hypothetical $(n+1)$ th patient; n has been pre-determined, usually 19-24.

The original CRM, proposed by O'Quigley et al (1990), drew some critiques and/or some strong opposition. Korn et al. (1994), Goodman et al. (1995), and Ahn (1998) pointed out the followings. Corrections have been proposed by those authors plus Thall et al (1999), Zohar and Chevret (2001), Storer (2001), among others.

First, the CRM might start the trial with an initial dose far above the “customary” lowest dose that is often one-tenth the LD10 in mice. This possibility makes many clinicians and **regulatory agencies** (e.g. FDA) reluctant to implement the CRM. After all, this is the first trial in human, little is known about the dose range - except results from animal studies. Some might go higher at the first dose, but not more than one-third the LD10 in mice. **Some proposed that the trial always starts with the lowest dose as the dose for the first patient; CRM would start with the second dose/patient.**

Secondly, there is a possibility that dose could be escalated for more than one dose level at a time (traditionally, as in standard design, doses are equally-spaced on the log scale or following a modified Fibonacci sequence with increases of 100, 67, 50, 40, and 33% for fifth and subsequent doses). Moller (1995) gave an example showing that the **first dose could be escalated to the top level when the first patient has no LTD.**

To overcome this problem, some proposed that one could pre-determined a set of doses to be used in the trial just as under the standard design. From the model, $\text{logit}[p(x); \beta] = \alpha + \beta x$, and use the prior mean of β to solve for doses with (prior) probabilities 5%, 10%, 15% etc... Start the first patient at the lowest dose (as previously proposed) and the magnitude of dose escalation is limited/imputed to one dose level only between two consecutive patients.

If doses with (prior) probabilities 5%, 10%, 15% etc... are pre-determined then each patient is treated at such a dose closest to the calculated dose from the Bayesian CRM process.

As previously mentioned, some investigators are uncomfortable at another feature of the CRM where CRM uses only a cohort of one patient for the dose adjustment. - just as many are with “the fast-track design”. Of course, it can be easily changed by **increasing the size of the cohort to two or three patients.**

The problem is, if all of those modifications are implemented, the resulting design - whatever you call it - would be similar to the standard design (Korn et al, 1994). There seem to be no perfect solution to the very fundamental dilemma, the conflict between scientific efficiency & ethical intent.

The strength of the CRM are still its three properties: (1) it has a well-defined goal of estimating a percentile of the dose-toxicity relationship, (2) it should converge to this percentile with increasing sample sizes, and (3) the accrual is pre-determined. The standard design does not have these characteristics.

In addition, there seems to be no way to overcome the problem that, under CRM and cohorts of size one, the dose for the next patient can be determined only after the result on the DLT for the current patient becomes available. This goes against the desire by most investigators to complete phase I trials as rapidly as possible - not only with the minimum number of patients but also in a minimal amount of time. This is mostly due to the urgent need to identify new active drugs and phase II efficacy trials cannot begin until completion of the phase I trial.

In other words, Korn et al. pointed out two severe deficiencies of the CRM. First, **trials will take too long to complete** - using cohorts of size one - especially when there is no shortage of patient (if one did not have concerns with this, one could accrue one patient at a time to the SD rather than three at a time - where one enrolls the second and third patient without waiting for toxicity result from the first). Second, **some patients - especially the first few - may be treated at dose level higher than the intended MTD.**

A typical problem with Bayesian method is, at early time when very little data available, results/decisions are dominated by the choice of the prior. With a poorly chosen prior, some early patients maybe treated at doses higher than MTD which is defined as “ the highest possible dose with acceptable toxicity”. And dose-severity relationship is not linear; toxicity seen at doses higher than MTD will likely more serious than one seen at MTD.

Of course, as mentioned, an obvious remedy to the first deficiency is to have CRM accrue more than one patient - say, three - at a time to a dose level. This, however, worsen the second deficiency of the CRM, its tendency to treat patients at a high doses.

Regulatory agencies - and protocol review panels and IRBs - may be too rigid; they are more concerns about side effects: (1) patients may die from side effects - some are fatal, but they may also die - and more likely so - from the disease if not treated with enough medication (as mentioned, these are mostly terminal patients; some would likely die before any other therapy can be attempted), and (2) some side effects are not serious or treatable/reversible; in such situation, methods such as CRM should be seriously considered.

The problem for the time being is only the ease of application, or lack of it. There are software in the public domain; for example,

http://biostatistics.mdanderson.org/SoftwareDownload/SingleSoftware.aspx?Software_Id=13

but only for the original algorithm proposed by O'Quigley et al. It would be more efficient if, say, available in SAS with some options for more flexibility.

The next problem is logistic arrangement: Can we afford to have a statistician calculating the next dose for CRM design when we need it? (About 50-100 active early-phase trials are open at any time in any large medical center) . Or how to make it simple enough to make its implementation practical? In addition, **we need immediate access to toxicity results!**

There are no exercises today –

REFERENCES

- Storer (1989); Biometrics 45: 925-937
- **O'Quigley et al (1990); Biometrics 46: 33-48**
- Gatsonis & Greenhouse (1992); Stat Med 11:1377-1389
- **Goodman et al. (1995); Stat Med 14: 1149-1161**
- **Korn et al (1994); Stat Med 13: 1799-1806**
- O'Quigley and Shen (1996); Biometrics 52: 673-684
- Ahn (1998); Stat Med 17: 1537-1549
- Heyd and Carlin (1999); Stat Med 18: 1307-1321
- Zohar and Chevret (2001); Stat Med 20; 2827-2843
- Storer (2001); Stat Med 20: 2399-2408.