Comparison of a Simple Bayesian Decision-Theoretic Design for Dose-Toxicity Trials with Traditional 3 + 3 Method

Lynn Eudey, Shenghua (Kelly) Fan, Tianyao Lu

Department of Statistics and Biostatistics, California State University East Bay

Abstract

Phase I dose-finding trials using a new drug are conducted to study the safety, explore the pharmacokinetics and pharmacodynamics and to identify the maximum tolerated dose (MTD) of the new drug. MTD is defined as a specific percentile of the dose-toxicity curve. Due to ethical considerations the trials use adaptive designs, increasing the dose in patients based on patient results at all doses.

The traditional 3 + 3 design is often used but is inefficient. Fan, Lu and Wang (2012, Statistics in Medicine) propose a simple and flexible Bayesian decision-theoretic design. Their method uses a working model with conjugate priors which produce analytic posterior distributions for monotonic dose-toxicity curves. Moreover, a utility function can be selected to reflect the interest of the trial. In this poster we present simulation results comparing their method with the traditional 3 + 3 method. The possible problem due to prior misspecification is also investigated.

Key words: Bayesian adaptive-design, dose-finding trial, decision theory

Introduction

Phase I studies are typically the first exposure of a new drug or therapy in humans. Often therapies known to be effective for one indication are tried in a different indication. The main purpose of a Phase I trial is to study the safety, pharmacokinetics, and pharmacodynamics of the new drug. In addition, a primary objective of a Phase I study in drug-therapeutic clinical trials is to estimate the maximum tolerated dose (MTD) of the new drug. The definition of MTD is the dose at which r-% of the patients suffer dose-limiting toxicity (DLT). Often this percentile is the 33rd percentile.

In some studies the "patients" undergoing Phase I trials are healthy volunteers, frequently from the pharma/biotech company developing the drug. In oncology trials the patients can be cancer patients who have undergone every possible currently marketed therapy to no avail. In all cases, especially when the new therapy has known uncomfortable side-effects (nausea, muscle weakness, pain, alopecia) the approach to estimation of the MTD should use as few patients as possible.

Kelley and Venook (2013) point out that in some patient populations there is a high rate of background adverse events necessitating Phase I trials in the patient population. This is especially true if there are no known reliable surrogate endpoints. It is therefore critical to have an efficient phase I design for combined agents to study the toxicity profile of

combined agents even if they are well established agents. Bayesian methods are practical for well-established agents as we have "prior information" of the toxicity profiles of agents.

Methods

Traditional 3+3 Design

The traditional 3 + 3 design is frequently used for Phase I trials to determine the MTD (when defined as the 33rd percentile). This study often has 5 – 6 doses predetermined for experimentation and treatment is assigned to cohorts of three patients. The first cohort is assigned to the lowest dose and if no toxicities occur, then another three patients are assigned to the next highest dose. If a dose has one patient (of the three, denoted by 1/3) experience a toxicity response, then that same dose is given to the next cohort. If a dose has two or three patients experience toxicity (denoted by 2+/3) then this dose exceeds the MTD and the dose is lowered. If less than two cohorts have experienced this lower dose, then the next cohort is assigned to the lower dose. The MTD is defined as the highest dose with at most one-third of the patients experiencing toxicity out of six patients (two cohorts).
(See Berry, Carlin, Lee, Muller (2011).)

Using an algorithm adapted from that of Berry, Carlin, Lee, Muller (2011), a R program was written to simulate the use of 3+3 design using six increasing doses. Below are three grids showing the use of the design.

Example 1: Nine patients with two toxic events

Cohort	**Dose (increasing from left to right)**					
	1	**2**	**3**	**4**	**5**	**6**
1	0/3 (increase)					
2		1/3 (stay)				
3		1/3 (MTD)				

Example 2: Fifteen patients with four toxic events

Cohort	**Dose (increasing from left to right)**					
	1	**2**	**3**	**4**	**5**	**6**
1	0/3 (increase)					
2		1/3 (stay)				
3		0/3 (increase)				
4			1/3 (stay)			
5			2/3 (decrease)			
		2 cohorts already - MTD				

Example 3: Fifteen patients with two toxic events

Cohort	**Dose (increasing from left to right)**					
	1	**2**	**3**	**4**	**5**	**6**
1	0/3 (increase)					
2		0/3 (increase)				
3			0/3 (increase)			
4				2/3 (decrease)		
5			0/3 (MTD)			

A Simple Bayesian Design

Fan, Lu and Wang (2013) propose a simple Bayesian design for Phase I trials. They use a Beta prior, a newly proposed isotonic regression working model, a simple gain function, and update the working data set with each cohort's outcome. The Beta prior (parameters a and b) is updated to a Beta posterior. Each dose, x_i, has probability $p_i = a_i / (a_i + b_i)$ of toxic event on average. Stopping rules based current information can be incorporated. Here, we stop if 1) the lowest dose is too toxic, 2) we reach the maximum number of patients.

The design has seven components and only one basic assumption, that probability of toxicity is monotonically increasing. An implied assumption for the working data is that if a patient experiences a toxicity at a low dose then he/she will experience a toxicity at all higher doses. Additionally, if a patient experiences no toxicity at a high dose then he/she will experience no toxicity at lower doses. Thus, the working data set can be updated with all the implied outcomes based on the observed outcomes of the current cohort.

If there are K doses, then up to K posterior distributions are updated to new Beta distributions and the next cohort is assigned a dose based on the maximum expected gain over each of the K doses.

Their design can implement a variety of gain functions and also weight the value of current patients and future patients. For this poster we use their One-Step-Look-Ahead algorithm which considers the next patient as the last patient in the trial (even though it may not be). The gain function is $-|\, p_i - r \,|$, where r is the target toxicity level for MTD (here 0.33). The next dose is the dose that maximizes the expected gain using the updated Beta distributions. Start with the lowest dose, increase doses until a toxicity occurs, then incorporate the working model. The first set of p_i 's are based on prior information about the therapy.

If we let $T = 0$ for no toxic event and $T = 1$ for a toxic event, then after treating a patient at dose x_i , given patient's toxicity t_i , then the updated Beta distribution has parameters ($a_i + t_i$, $b_i + 1 - t_i$). For cohorts with more than 1 patient, update the parameters using the total number of toxicities (and non-toxicities). For example, given a patient is treated with dose 3 and experiences a toxic event, then the working data is updated as shown in the table below.

Example 4: Use of the Bayesian Design. Given a patient is treated at dose 3 and experiences a toxicity.

Dose →	**1**	**2**	**3**	**4**	**5**	**6**
Working Data	No change	No change	+1	+1 Implied by model	+1 Implied by model	+1 Implied by model
Beta parameter, *a*	a_1	a_1	a_1 +1	a_1 +1	a_1 +1	a_1 +1
Beta parameter, *b*	b_1	b_1	b_1	b_1	b_1	b_1

The next dose would be assigned choosing dose i, that gives maximum $E(-|p_i - r\,|)$ over the six updated Beta distributions.

Example 5: Use of the Bayesian Design. Given a patient is treated at dose 3 and experiences no toxicity.

Dose →	**1**	**2**	**3**	**4**	**5**	**6**
Working Data	0 Implied by model	0 Implied by model	0	No change	No change	No change
Beta parameter, *a*	a_1	a_1	a_1	a_1	a_1	a_1
Beta parameter, *b*	b_1+1	b_1+1	b_1+1	b_1	b_1	b_1

The next dose would be assigned choosing dose i, that gives maximum $E(-|p_i - r|)$ over the six updated Beta distributions.

Simulations

Simulations (m = 10,000) were run for two different MTD targets, $r = 0.20$ and $r = 0.30$. The proposed Bayesian design is adaptable to different MTD targets, while the 3 + 3 design was developed for a target around 0.33. The Bayesian design can incorporate several dose selection rules. Here we use one-step-look-ahead which regards each successive patient as the last patient and the gain function $-|p - r|$ (referred to as OSLA(1)). Cohort size for the 3+3 design is fixed at 3. The proposed Bayesian design can use different cohort sizes; here we use a cohort of 1 patient. Additionally, the proposed Bayesian design can incorporate stopping rules based on the working data set or stop at a pre-set maximum sample size n_{max}. Here we stop at 16 patients for OSLA(1).

Results

The tables below show the results of the simulations. For Table 1, Table 2, and Table 3 the targeted dose has a probability of toxicity of 0.20 ($r = 0.20$). This is followed by graphs showing the results for those simulations. Another simulation (below) shows the results for a targeted dose with a probability of toxicity of 0.30 ($r = 0.30$).

Table 1: Target $r = 0.20$. Probabilities of toxicities for doses 1 – 6.

Scenario\Dose	**1**	**2**	**3**	**4**	**5**	**6**
High Toxicity	0.15	**0.20**	0.38	0.52	0.70	0.80
Medium Toxicity	0.01	0.02	0.09	**0.20**	0.40	0.58
Low Toxicity	0.001	0.002	0.003	0.01	0.07	**0.20**

Table 2: Proportion of patients assigned to each dose using scenarios of Table 1 over 10,000 simulations. Modes in bold.

Scenario	Design	Dose Allocation					
		1	**2**	**3**	**4**	**5**	**6**
High Toxicity	**OSLA(1)**	0.40	**0.50**	0.06	0.03	0.01	0.00
	3 + 3	0.061	0.32	**0.42**	0.17	0.026	0.001
Medium Toxicity	**OSLA(1)**	0.06	0.07	0.25	**0.54**	0.04	0.04
	3 + 3	0.01	0.02	0.11	0.30	**0.41**	0.15
Low Toxicity	**OSLA(1)**	0.06	0.06	0.06	0.06	0.19	**0.56**
	3 + 3	0.00001	0.0001	0.0014	0.028	0.25	**0.72**

Table 3: Dose identified as MTD using scenarios of Table 1 over 10,000 simulations. Modes in bold.

Scenario	Design	MTD Determination					
		1	**2**	**3**	**4**	**5**	**6**
High Toxicity	**OSLA(1)**	0.38	**0.60**	0.02	0	0	0
	3 + 3	0.20	**0.49**	0.27	0.047	0.0022	0
Medium Toxicity	**OSLA(1)**	0	0.01	0.25	**0.74**	0	0
	3 + 3	0.0018	0.033	0.22	**0.46**	0.24	0.045
Low Toxicity	**OSLA(1)**	0	0	0	0	0.19	**0.81**
	3 + 3*	0	0	0.0001	0.0029	0.023	**0.14**

*MTD not determined by 3+3 design and low toxicity scenario for 83.4% of the 10,000 simulations.

Figure 1: Target $r = 0.20$, scenario of Table 1, number of patients used for 3+3 design (for OSLA(1) fixed number of 16 patients).

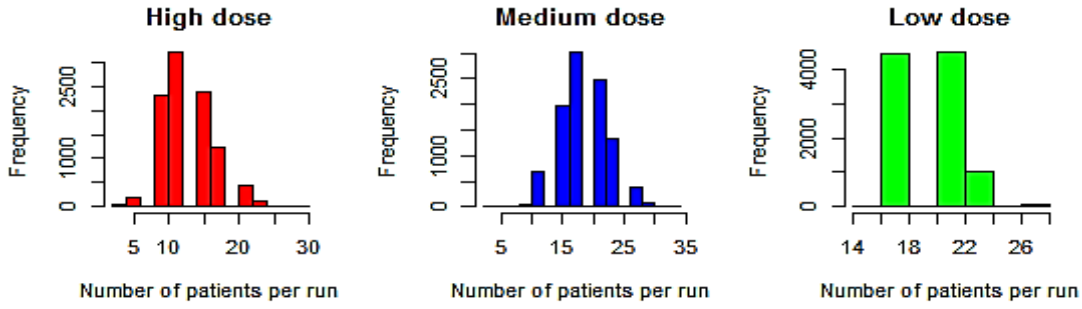

Ave = 14, s = 3.78 Ave = 19, s = 3.89 Ave = 20, s = 2.02

Figure 2: Target $r = 0.20$, scenario of Table 1, number of toxicities for 3+3 design.

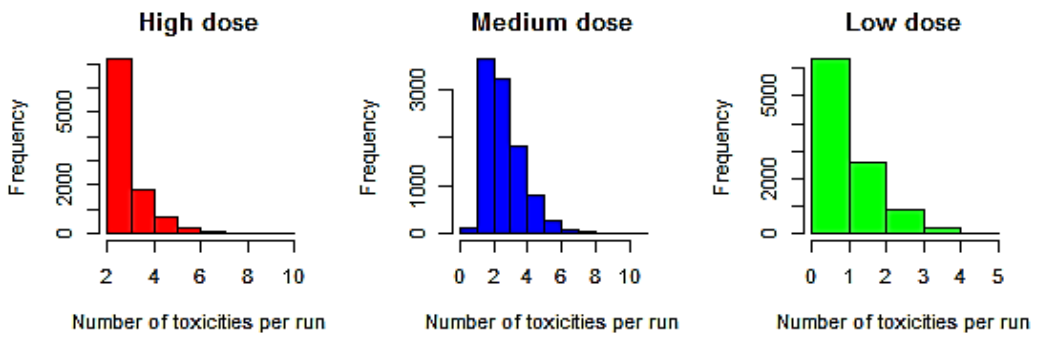

Ave = 4, s = 1.14 Ave = 4, s = 1.21 Ave = 2, s = 1.09

Figure 3: Target $r = 0.20$, scenario of Table 1, number of toxicities for OSLA(1) design.

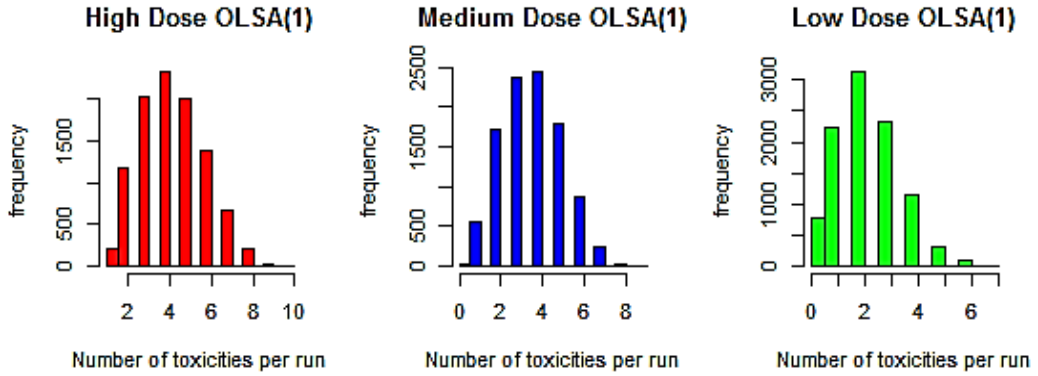

Ave = 4, s = 1.59 Ave = 4, s = 1.45 Ave = 2, s = 1.27

For Table 4, Table 5, and Table 6 the targeted dose has a probability of toxicity of 0.30 ($r = 0.30$). And the graphs (Figures 4 – 6) for Tables 4 – 6 follow.

Table 4: Target $r = 0.30$. Probabilities of toxicities for doses 1 – 6.

Scenario\Dose	**1**	**2**	**3**	**4**	**5**	**6**
High Toxicity	**0.30**	0.53	0.77	0.87	0.95	0.98
Medium Toxicity	0.05	0.10	0.20	**0.30**	0.50	0.70
Low Toxicity	0.02	0.03	0.06	0.10	0.18	**0.30**

Table 5: Proportion of patients assigned to each dose using scenarios of Table 4 over 10,000 simulations. Modes in bold.

		Dose Allocation					
Scenario	**Design**	**1**	**2**	**3**	**4**	**5**	**6**
High Toxicity	**OSLA(1)**	**0.894**	0.080	0.021	0.005	0	0
	3 + 3	**0.540**	0.40	0.06	0	0	0
Medium Toxicity	**OSLA(1)**	0.064	0.14	0.34	**0.37**	0.067	0.15
	3 + 3	0.051	0.13	0.26	**0.30**	0.22	0.039
Low Toxicity	**OSLA(1)**	0.063	0.062	0.063	0.15	0.28	**0.38**
	3 + 3	0.027	0.043	0.050	0.17	0.31	**0.36**

Table 6: Dose identified as MTD using scenarios of Table 4 over 10,000 simulations. Modes in bold.

		MTD Determination					
Scenario	**Design**	**1**	**2**	**3**	**4**	**5**	**6**
High Toxicity	**OSLA(1)**	**0.96**	0.04	0	0	0	0
	3 + 3	**0.89**	0.11	0	0	0	0
Medium Toxicity	**OSLA(1)**	0	0.066	0.38	**0.51**	0.047	0
	3 + 3	0.060	0.20	**0.34**	0.32	0.071	0.0023
Low Toxicity	**OSLA(1)**	0	0	0.002	0.085	0.37	**0.54**
	3 + 3*	0.010	0.020	0.060	0.16	0.30	**0.45**

*MTD not determined by dose 6 for 3+3 design and low toxicity scenario.

Figure 4: Target r = 0.30, scenario of Table 4, number of patients used for 3+3 design (for OSLA(1) fixed number of 16 patients).

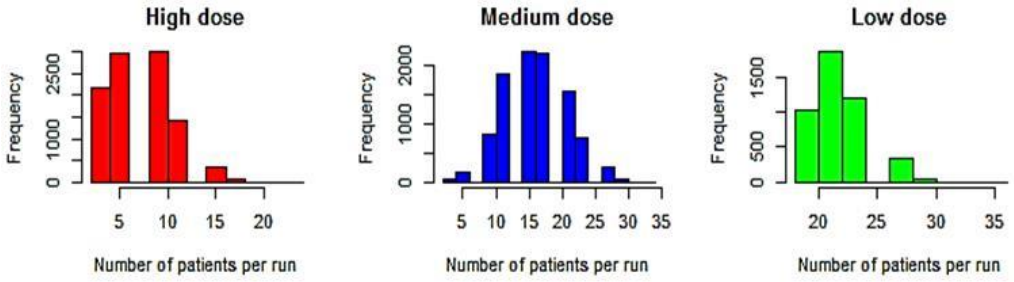

Ave = 8, s = 3.37 Ave = 17, s = 4.89 Ave = 22, s = 2.80

Figure 5: Target $r = 0.30$, scenario of Table 4, number of toxicities for 3+3 design.

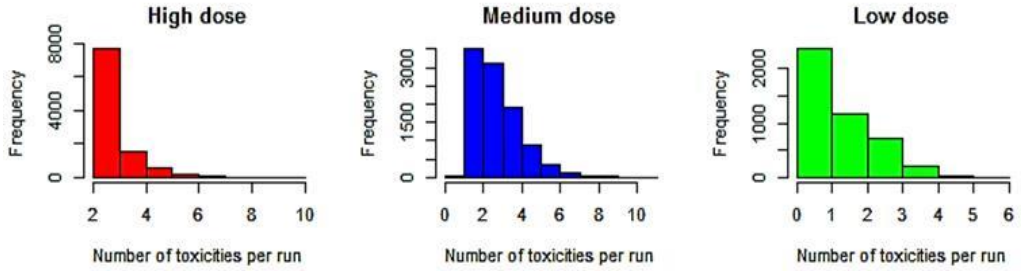

Ave = 3, s = 1.05 Ave = 4, s = 1.23 Ave = 2, s = 1.18

Figure 6: Target $r = 0.30$, scenario of Table 4, number of toxicities for OSLA(1) design.

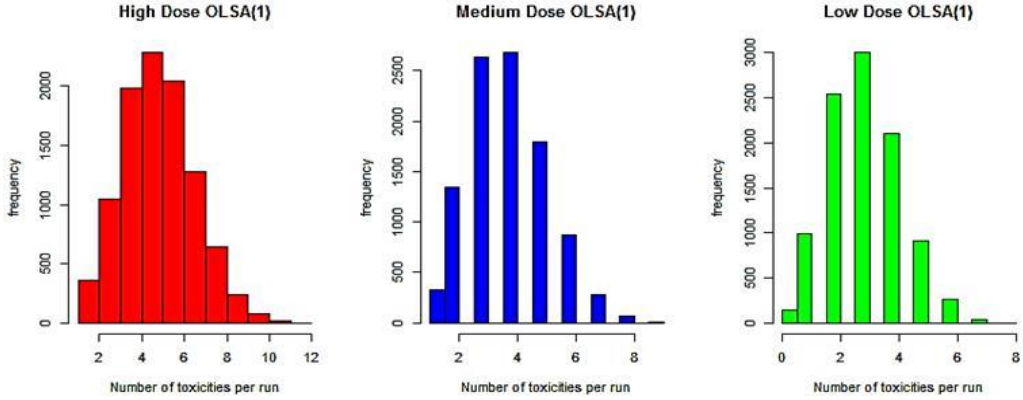

Ave = 6, s = 1.68 Ave = 4, s = 1.39 Ave = 3, s = 1.29

Conclusions

A simple Bayesian design provides a flexible model for determining maximum tolerated dose. The proposed Bayesian design is simple to implement and requires little computation time. Updating the Beta distribution parameters takes one line of R code and provides an analytic posterior Beta distribution. In their article, Fan, Lu, and Wang (2012) compare their results with the continual reassessment method (CRM) and the continual reassessment method likelihood version (CRML). In this study we compared the results of the tradition 3+3 design with the Bayesian approach (OSLA(1)) using a fixed number of 16 patients. To summarize the results, we look at proportion of times the correct dose was selected under the two designs; the number of patients each design used; and the number of toxicities experienced under each design.

In comparing whether the designs selected the correct dose for MTD, the modal choices are shown in Tables 3 and 6 (for the two different scenarios). In all cases, the Bayesian design (OSLA(1)) had the modal frequencies for the correct dose. Also, the OSLA(1)

design selected the correct dose with higher relative frequency than that of the 3+3 design. In addition, in each scenario the 3+3 design shows higher proportion of simulations than OSLA(1) shows with a dose above the target being selected. For example, in Table 3, with medium toxicity the 3+3 design selected doses 5 and 6 nearly 30% of the time, when dose 4 should be the MTD, while OSLA(1) shows 0% above dose 4 being identified as the MTD. This is in part because the target is 0.20 (lower than what 3+3 is designed for). But, the same is true when the target is 0.30 (Tables 4-6); in the low-toxicity scenario the 3+3 design designates the MTD at too high a dose 11% of the time, while OSLA(1) does only 4% of the time. For the medium-toxicity scenario, 3+3 designates too high a dose about 7% of the time, while OSLA(1) does only about 5% of the time.

Quite often the 3+3 design required many more than 16 patients. The proposed Bayesian design can incorporate stopping rules which may result in fewer patients being used than the number used under the 3+3 design. Here, Figures 1 and 4 show that, with the exception of the high dose scenarios, the 3+3 design can frequently go over 16 patients.

The proposed Bayesian design can be modified to accommodate any target probability for MTD (here we used 0.20 and 0.30). In contrast, the 3+3 design targets a fixed probability somewhere between 0.20 and 0.30, with no flexibility for other target probabilities.

In conclusion, the proposed Bayesian design appears to be more efficient and more accurate than the traditional 3+3 design. In addition, the proposed Bayesian design provides more flexibility in selecting a target probability for rate of toxicity for Phase I study and incorporating stopping rules.

References

Berry, S.M., B.P. Carlin, J.J. Lee, and P. Muller, Editors. 2010. *Bayesian Adaptive Methods for Clinical Trials,* Chapman & Hall/CRS Biostatistics Series Boca Raton, FL

Fan, SK, Y. Lu, and Y-G Wang. 2012. A simple Bayesian decision-theoretic design for dose-finding trials. *Statistics in Medicine,* v. 31, 3719 – 3730

Kelley, R.K., and A.P. Venook. 2013. Novel Therapeutics in Hepatocellular Carcinoma: How Can We Make Progress? *Am Soc Clin Oncol Educ Book.* doi: 10.1200/EdBook_AM.2013.33.e137