DESIGN OF PHASE I TRIALS

Authors: Alexandros Xyrafas, Martin Bigler
Version 1.0, 27.06.2016

Table of contents

Abbreviations

ATC	Accelerated Titration Design
AUC	Area Under the Curve
BLRM	Bayesian Logistic Regression Model
CRM	Continual Reassessment Method
CTCAE	Common Terminology Criteria for Adverse Events
DL	Dose Level
DLT	Dose-limiting Toxicity
EWOC	Escalation with Overdose Control
mCRM	Modified Continual Reassessment Method
MTD	Maximum Tolerated Dose
mTPI	Modified Toxicity Probability Interval (method)
NOAEL	No Observed Adverse Effect Level
PD	Pharmacodynamics
PGDE	Pharmacologically Guided Dose Escalation
PIPE	Product of Independent Beta Probabilities Dose Escalation (method)
PK	Pharmacokinetics
RP2D	Recommended Phase II Dose
SAKK	Schweizerische Arbeitsgemeinschaft für Klinische Krebsforschung – Swiss Group for Clinical Cancer Research
SAKK CC	SAKK Coordinating Center
TITE-CRM	Time-to-event CRM
TITE-EWOC	Time-to-event EWOC
TPI	Toxicity Probability Interval (method)

Introduction

Phase I clinical trials are an essential step in the development of new anticancer treatment, as some of the agents are often given to humans for the first time (first-in-human). The main goal of these trials is to identify the maximum tolerated dose (MTD), establish the recommended phase II dose (RP2D) and thus schedule the subsequent phase II trials of the proposed new treatments or treatment combinations. Whereas in most other cases phase I trials are performed using healthy volunteers, in oncology usually patients which do not have many treatment options left. The guiding principle for dose escalation in phase I trials in oncology therefore is to avoid exposing too many patients to sub-therapeutic doses while preserving safety and maintaining rapid accrual. In the current document we review some of the most prominent dose escalation methods for phase I trials, including rule- as well as model-based dose methods that have been developed in order to evaluate new anticancer agents. Merits and limitations of all described designs can be found in an overview table in Appendix 1. Furthermore, we review some other general points which are of interest in phase I trials.

Design options

Rule-based designs

The rule-based designs ("up-and-down" designs) assign patients to dose levels (DLs) according to pre-specified rules based on actual observations of target events (e.g., the dose-limiting toxicity (DLT)) from the clinical data.

Examples:

- 3+3 design
- Accelerated Titration Designs (ATD)
- Modified Toxicity Probability Interval (mTPI) method
- "Rolling 6" design
- Pharmacologically Guided Dose Escalation method

Model-based designs

The model-based designs assign patients to DLs and define the RP2D based on the estimation of the target toxicity level by a model depicting the dose – toxicity relationship. The idea is to use statistical models that actively seek a DL that produces a pre-specified probability of DLT by using toxicity data from all enrolled patients to compute a more precise dose – toxicity curve. This method can be conveniently carried out using Bayesian models.

Examples:

- Continual Reassessment Method (CRM)
- Bayesian Logistic Regression Model (BLRM)
- Escalation with Overdose Control

At SAKK, model-based designs or the mTPI method are generally the preferred options for larger phase I trials.

Short description of most important designs

Rule-based designs

3+3 design

The 3+3 design [1] is the design which is used most widely in phase I trials in oncology [2]. Patients are treated in cohorts of three. The first cohort is treated at a starting dose that is considered to be safe, based on extrapolation from animal toxicological data (determination of NOAEL: no observed adverse effect level), and the subsequent cohorts are treated at increasing DLs that have been defined in advance. Historically, dose escalation follows a modified Fibonacci sequence, in which the dose increments become smaller as the dose increases (dose increase by 100%, 67%, 50%, 40%, and 33% for all the rest), though many variations of this sequence are used [3]. The purpose for this was to allow more aggressive dose escalation for the initial DLs to treat fewer patients at sub-therapeutic doses.

The escalation procedure has the following steps:

- If one patient among the three of the first cohort experiences a DLT, then an additional cohort of 3 patients is treated at the same DL.
- If no patient among 3 or 1 among 6 patients experiences a DLT and the DL above has not been tested yet, then the dose is escalated.
- If ≤1 among 6 patients experiences a DLT and the DL above has been tested already, then the trial is stopped and the current DL is considered the maximum tolerated dose (MTD).
- If two or more patients experience a DLT and less than 6 patients have been treated at the DL below, then the dose is de-escalated.
- If two or more patients experience a DLT and 6 patients have already been treated at the DL below, then the trial is stopped and the DL below is considered the MTD.

See Figure 1 for a schematic description of the procedure.

Merits:

- Easy to implement, well understood without need of statistician
- Familiar to investigators and ethics committees
- Safe
- Many modifications exist with cohorts of different numbers: 2+2, 2+4, 3+1+1 (“best-of-five”, see Appendix 1), “Rolling 6 design” (see Appendix 1), 7+7, etc.

Limitations:

- Ignores history other than that of the previous cohort
- Performs same action under different situations (escalate in case of 0 DLT in 3 patients or 1 DLT in 6 patients; de-escalate in case of 2 DLTs in 3 patients or 3DLTs in 3 patients or ≥2 DLTs in 6 patients)
- Not possible to re-escalate
- Does not allow the investigator to change the targeted toxicity easily
- Fixed cohort size (3 or 6)
- The procedure is unnecessarily slow, leading to treatment of excessive numbers of patients at DLs less likely to be efficacious.

- Low probability of selecting the true MTD [5]
- High variability in MTD estimates [6]
- Fixed doses, not flexible

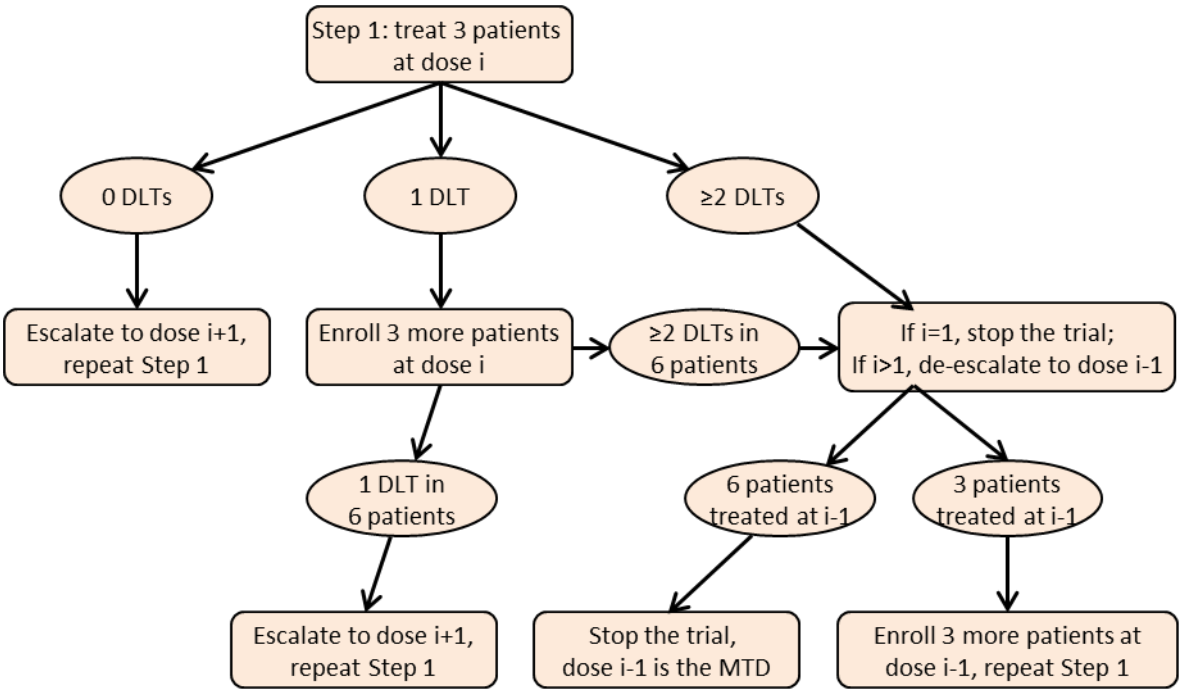

Figure 1. Schema of the 3+3 design, adapted from Ji and colleagues [4].

Accelerated Titration Designs (ATD)

Simon and colleagues [7] proposed three modifications to the classic 3+3 design to accelerate dose finding, considering as standard 40% dose-step increments. The dose-escalation/de-escalation rules are based on definitions of DLT and of "moderate" toxicity and may be protocol specific.

1. "Design 2" in the reference: Only one patient is treated per DL, until one patient has a DLT during the first cycle or two patients have a moderate toxicity (CTCAE grade 2) during the first cycle. Then, the design switches to the 3+3 design.
2. "Design 3" in the reference: Also only one patient is treated per DL but a more rapid dose escalation is performed by using double-dose steps during this stage. The single-patient-cohort stage terminates after one patient has a DLT during the first cycle or two patients have moderate toxicity (CTCAE grade 2) during the first cycle. Then, the design switches to the 3+3 design.
3. "Design 4" in the reference: Again only one patient is treated per DL and a rapid dose escalation is performed by using double-dose steps during this stage. The single-patient-cohort stage terminates after one patient has a DLT at any time during therapy or two patients have moderate toxicity (CTCAE grade 2) during any cycle of treatment. Then, the design switches to the 3+3 design.

The two cases grade 2 toxicity are used for practical reasons, since it is often difficult to determine whether a grade 2 toxicity is treatment related in a heterogeneous population of very ill patients.

The ideas of the accelerated titration designs could also be applied using model-based designs (see below).

Merits:

- Easy to implement, well understood without the need of a statistician

- Partially familiar to investigators and ethics committees
- Faster than 3+3, as escalation is rapid during the first stage of the escalation procedure
- Reduces the total number of patients treated in the trial
- Reduces the number of patients treated at sub-therapeutic doses
- May incorporate intra-patient dose-escalation
- Model-based designs following the initial single-patient cohorts can be incorporated to create hybrid methods

Limitations:

- Similar limitations as 3+3
- More aggressive, therefore may be associated with more risk
- Due to the conservativeness of the investigators, it has been observed that they are often used with an initial dose set much more conservative than would be done for the 3+3 design, reducing their effectiveness
- Few PK/PD data

Modified Toxicity Probability Interval (mTPI) method

Ji and colleagues [9][10] proposed an algorithm-based dose finding method based on toxicity probability intervals (TPI). According to this method, before the start of the trial, a target toxicity interval has to be specified (e.g. 0.20–0.30) and then a specific model is assumed for the toxicity probabilities at each DL. Based on this model, a number *n* of patients treated at a DL, and the observed number of DLTs *d* in these patients one of the following three decisions can be taken:

- Escalate to the next higher DL (E)
- Stay at the current DL (S)
- De-escalate to the next lower DL (D)

Furthermore, if based on the observed DLT rate it is likely that the current dose is highly toxic, the dose is de-escalated and re-escalation to that and higher DLs is not possible anymore (DU).

All these decisions can be calculated before the start of the trial and presented in a two-way table, as shown in Figure 2. This makes the method easy to conduct.

The trial is run until a pre-defined sample size is reached or a pre-defined number of patients are treated at a DL. The MTD is the DL which has the highest probability that the DLT rate lies within the pre-specified target toxicity interval.

Technical details can be found in Appendix 2.

Merits:

- Hybrid method which uses a simple Bayesian model-based method that can easily be summarized into a set of “3+3”-like rules before the trial starts
- Simple implementation (much software available: Excel macro, R code, EAST, etc.)
- Transparent to the clinical investigators (procedure and rules are pre-defined and presented in a table (see Figure 2) at the start of the trial and then used throughout the trial)
- Has good statistical properties [4]:

- Outperforms the traditional 3+3 design both in terms of mean number of patients treated above MTD and the percentage of correct MTD selection
- Can target any user-defined toxicity probability and performs comparably to model-based methods such as the CRM (see below)

Limitations:

- Ignores history other than that of the previous cohort
- A monotone dose-toxicity curve is necessary at the end of the study (isotonic regression)

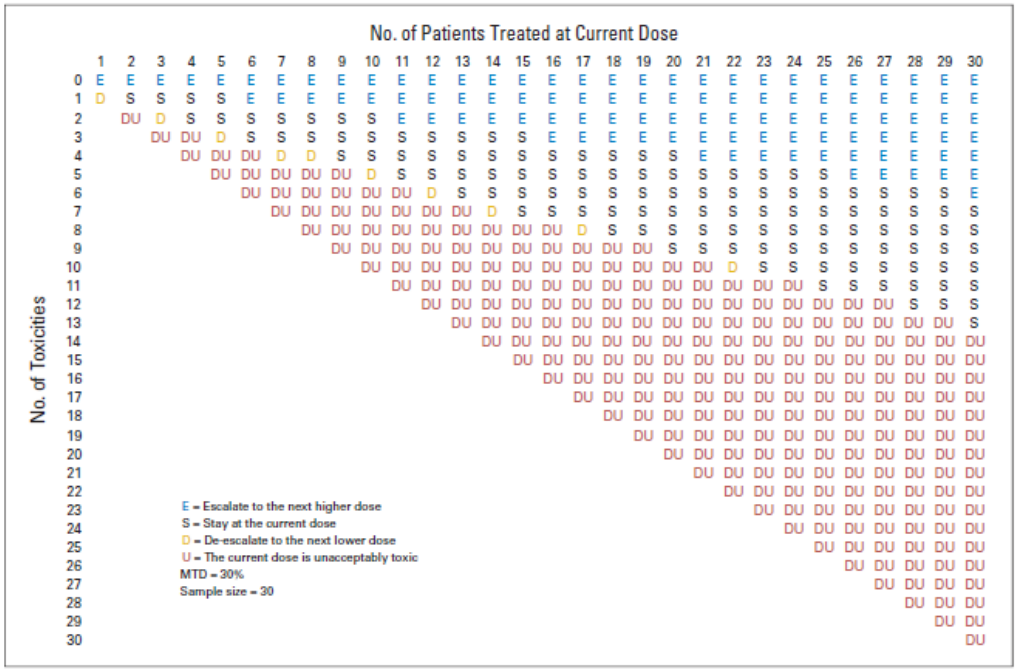

Figure 2. Dose-finding spreadsheet of the modified toxicity probability interval (mTPI) method. The letters in different colors are computed based on the decision rules under the mTPI method and represent different dose-finding actions. In addition to actions de-escalate the dose (D), stay at the same dose (S), and escalate the dose (E), the table includes action unacceptable toxicity (U), which is defined as the execution of the dose-exclusion rule in mTPI.

Model-based designs

Continual Reassessment Method (CRM) including modifications

The CRM is described in detail by O'Quigley and colleagues [11][12][13].

The novelty of the CRM was determining the dose for the next patient on the basis of the toxicity outcomes for all patients previously treated on the trial using a mathematical model for the association between dose and toxicity. This dose-toxicity curve is fitted to the data and each patient is assigned to the dose most likely to be associated with the target toxicity level, designated as MTD [14][15].

The original CRM algorithm is as follows:

1. A *target probability of DLT* (denoted ϑ) has to be specified before starting the trial, e.g. ϑ=0.30. Furthermore, a prior assumption on the probability of a DLT depending on the dose has to be made using a one-parameter curve (prior dose-toxicity curve).
2. One patient is treated at that dose for which the absolute difference of the estimated probability of DLT and the target probability of DLT is minimal.

3. The one-parameter toxicity curve is re-estimated using all data of all patients treated so far (posterior dose-toxicity curve). A graphical example of prior and posterior dose-toxicity curves can be found in Figure 3
4. Steps 2–3 are repeated until some pre-specified stopping criteria are met (e.g. pre-defined total sample size reached).

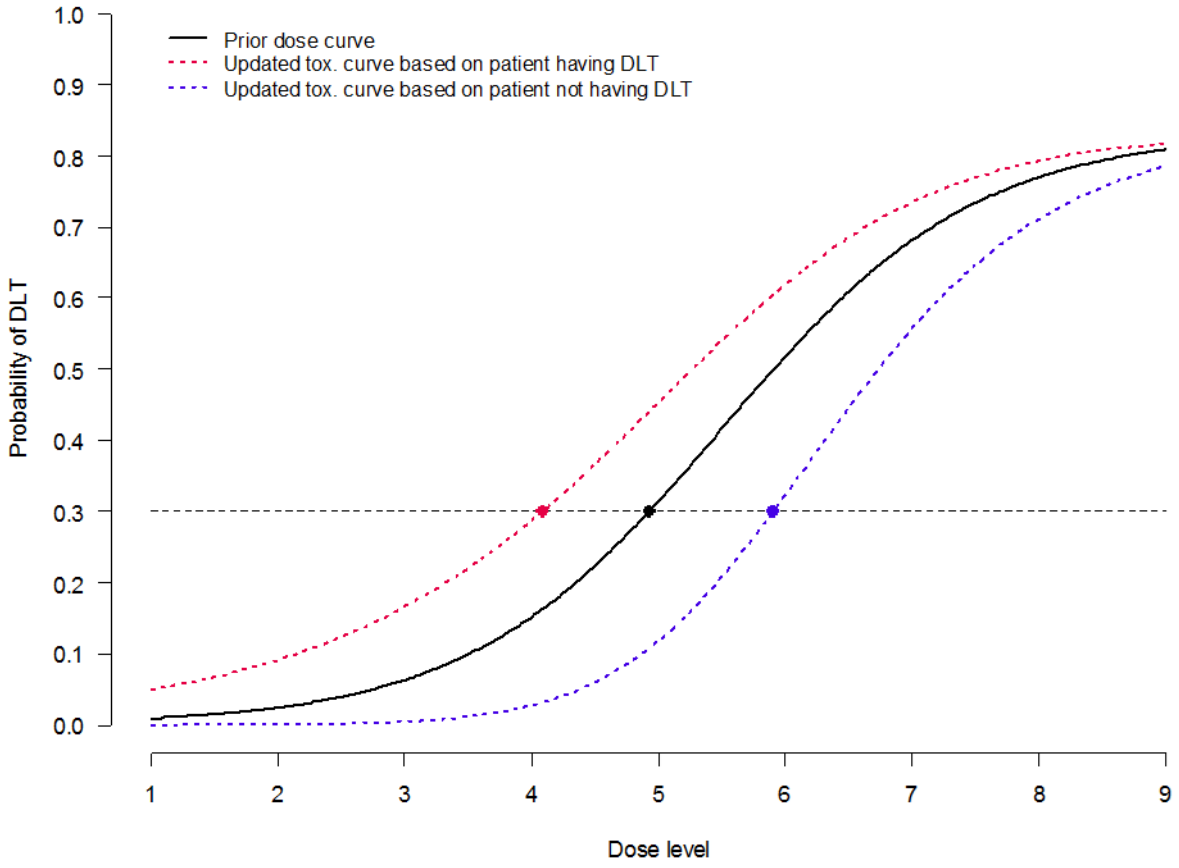

Figure 3. Prior and posterior dose-toxicity curves updating the estimate of the model parameters with one patient per cohort. With a desired DLT rate of 0.30, the DL for the first patient is DL5. If the first patient experiences a DLT at DL5, then the dose is de-escalated to DL4. If the first patient does not experience a DLT at DL5, then the dose is escalated to DL6.

There has been much criticism on the traditional CRM design regarding safety concerns such as the large dose escalation steps that were allowed based on little information or the starting DL based on experts' initial estimations, which could potentially put patients in danger.

Since then, there have been many adaptations of the original CRM algorithm.

Goodman and colleagues [6] retained the same general ideas of the CRM, with some modifications, calling it his proposed model "modified CRM" (mCRM). The general idea is to begin the study like a "3+3" (or a similar variant) and then switch to a CRM-type escalation procedure once the first DLT is observed. As such, the mCRM requires a set of pre-specified doses in which to escalate. The mCRM has some other key differences from the original version:

- DLs for escalation are pre-defined
- Starting dose is the lowest considered pre-defined DL

- Patients are treated in cohorts of 2 or 3.
- Any given dose escalation cannot increase by more than 1 DL, although dose de-escalations can be large

Below, there is a list of some of the many adaptations to the original CRM algorithm that can be found in the literature:

- *Extended CRM*: Two-stage design starting with 3+3 (or variants) until one DLT is observed and then switching to the original CRM using all the information of stage 1 in order to estimate the dose-toxicity curve [16].
- *Restricted CRM*: Revised version of extended CRM preventing escalations from exceeding one DL [16].
- *Practical CRM*: Uses preclinical information about toxicities to select doses that produce low (10%) and high (90%) rate of DLT and then estimate the DTC that fits these 2 points [17].
- *Time-to-event CRM (TITE CRM), see Appendix 1*: Uses a time-to-event approach with a weight function (relevant to the follow-up time of the patient to DLT) in order to incorporate time-to-toxicity of the patients in escalation decisions during the trial [18].

In the software that we currently use at the SAKK CC, the following modifications/options from the original CRM are available:

- The starting dose could be the lowest, 2nd lowest or 3rd lowest dose available
- Different stopping rules are available

By default, no previously untested doses are skipped while escalating and dose escalation is not allowed in case the previous subject experienced a DLT.

Merits:

- Few patients are treated at low, ineffective doses
- Information from all treated patients are taken into consideration for the recommended dose

Limitations:

- Not easily understood by non-statisticians
- It is imperative that a statistician is involved in CRM trial design and CRM trial conduct
- Needs prior specifications
- Trial length may be too long
- Potential difficulties in acceptance from institutional review boards and regulatory agencies

Bayesian Logistic Regression Model (BLRM)

The two-parameter BLRM parameterizes the probability of toxicity using the logistic transformation. The ability to combine data from patients treated at different DLs is a valuable feature of the regression model. This effectively enables us to "borrow" statistical strength, which can result in more accurate predictions of the future outcomes for patients who are given specific doses of the drug. Compared to the one-parameter CRM model, the logistic model imparts greater flexibility and more accurately depicts the dose-toxicity curve that characterizes the treatment under evaluation.

This method, using a fully Bayesian approach, proposes to classify the probability of DLT into 4 regions of the posterior distribution: Under-dosing, Targeted toxicity, Excessive toxicity and Unacceptable toxicity, while the intervals for excessive and unacceptable toxicity are sometimes combined for overdose control (see Figure 4).

More details of prior tuning are described by Neuenschwander and colleagues [19].

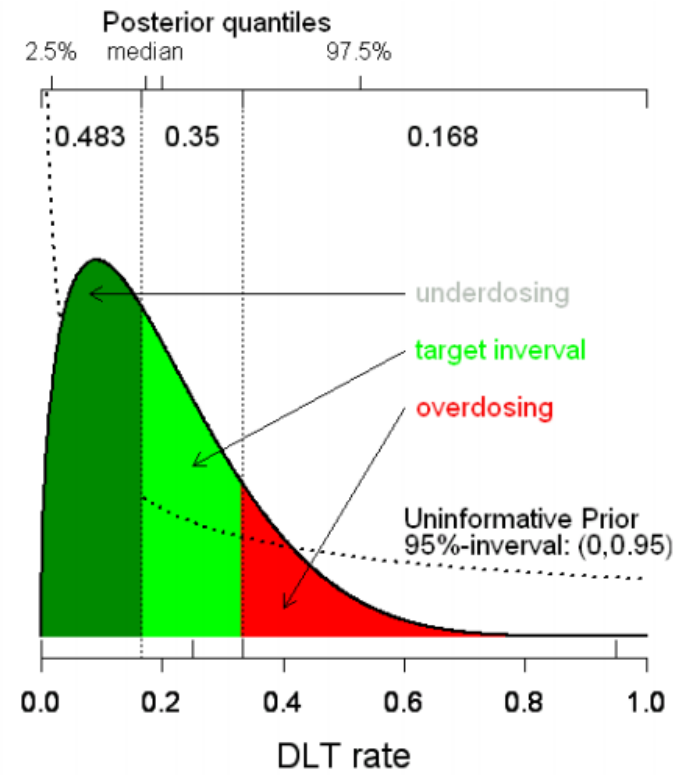

Figure 4. Example of categorization of posterior distribution into regions in the BLRM method (excessive and unacceptable toxicity are combined).

A usual classification of the posterior distribution is:

Under-dosing:	$\vartheta \in (0, 0.17]$	
Targeted toxicity:	$\vartheta \in (0.17, 0.33]$	
Excessive toxicity:	$\vartheta \in (0.33, 0.6]$	
Unacceptable toxicity:	$\vartheta \in (0.6, 1]$,	where ϑ refers to the probability of toxicity.

For the selection of the next candidate dose, the posterior probabilities of each region are examined and a dose is recommended based on a pre-specified acceptance of risks and benefits. A common recommendation rule is to maximize the probability of toxicity within the targeted toxicity category, while keeping the probability of excessive and unacceptable toxicity under 25%.

Regarding the dose selection method, given that we deal with a fully Bayesian approach, a formal loss function (see below) should be used, where the optimal decision is the one that minimizes the corresponding Bayes risk (posterior expected loss).

A simple loss function could be defined such as:

$$\mathcal{R}(\theta, dose) = \begin{cases} l_1{=}1 \text{ if under-dosing} \\ l_2{=}0 \text{ if targeted toxicity} \\ l_3{=}1 \text{ if excessive toxicity} \\ l_4{=}2 \text{ if unacceptable toxicity} \end{cases}$$

An alternative dose selection method uses escalation with overdose control, selecting the dose that maximizes the posterior probability of targeted toxicity for all doses where the posterior probability of overdosing (either excessive or unacceptable toxicity) is less than the user specific threshold.

Merits:

- Two-parameter logistic model estimates better the dose-toxicity curve
- All available information from previous patients is used to guide dose assignment for the next cohort
- Next recommended dose is the one which has the highest probability of having an acceptable and desired toxicity probability
- Overdose control is taken into consideration

Limitations:

- Similar limitations as of the CRM.

Other considerations

Late toxicities

In treating cancer patients, late toxicities may arise depending on the nature of the therapy and disease (e.g. there might be complications of normal tissues some months after radiation therapy or severe chronic graft-versus-host disease after bone marrow transplantation). With the emergence of novel monoclonal antibodies, such as immunotherapies that target the human immune system to fight tumor cells, clinical researchers also face challenges of late-onset or cumulative toxicities due to long pharmacokinetic half-life and a potential delayed effect on the disease. This issue becomes more prominent for phase I trials of combination therapies [20][21], as combining two or more agents of different mechanisms may result in unexpected cumulative toxicities.

Basing dose finding on DLTs scored within a fixed window (e.g. 1 treatment cycle) is a practical compromise motivated by algorithms that choose doses sequentially for successive patients based on previous patients' doses and outcomes. Conventional methods consider a patient to be fully evaluated if he/she experiences toxicity within this fixed window (the so called DLT window) or is followed for this length of time without toxicity. A problem with such methods is that they may cause investigators to treat an undesirably large number of patients at toxic doses before late-onset toxicities are first observed [22]. As a result, a trial may fail to identify the MTD or select a dose that appears to be safe in the pre-specified fixed window for DLT evaluation but cannot be tolerated by patients beyond it [23].

A safer approach would be to use a very large DLT window and wait for each patient's outcome before choosing the next patient's dose. However, in most settings this would introduce substantial inconvenience from excessive trial durations [24].

To conclude, we should probably bear in mind the recommendations given by Postel-Vinay and colleagues [22], according to which, the recommended dose for further studies should incorporate all available information – notably toxicities observed after cycle 1 and intolerable clinical grade 2 toxicities.

Intra-patient dose escalation

The likelihood that patients will receive sub-therapeutic doses of the treatment in the initial stages of a phase I trial remains an ethical consideration. An approach that could enhance the proportion of patients who are exposed to a potentially therapeutic DL would be to allow escalation within individual patients (intra-patient dose escalation), given that the patient has experienced no or minimal toxicity in the DL he/she initially was assigned to [25].

Although the idea of intra-patient dose escalation is appealing, it is not widely applied for both theoretical and practical reasons. Intra-patient dose escalation may mask the cumulative effects of treatment or, at the very least, would make them harder to differentiate from chronic or delayed toxic effects. However, regardless of the trial design used, chronic, delayed, or cumulative toxic effects are generally not well captured by most phase I trials because most patients with advanced cancers do not remain on study for extended periods of time. Furthermore, it can be difficult to present and interpret results of trials that allow intra-patient dose escalations because a single patient may contribute data for several DLs [2]. It is important, thus, that a minimum number of newly recruited patients be treated at each DL in order to evaluate the tolerability of the DLs in those individuals, rather than having risk confusion by combining the intra-patient escalated and the new patients. In practice, intra-patient dose escalation will often be irrelevant because information about the safety of the drug is required prior to dose escalation, by which time many phase I patients will have progressed. Arguably, as the actual chance of benefit (in terms of response rate) on a phase I is usually quite small, the added benefit for an individual patient to be accrued by escalation to one or two DLs becomes even smaller. As patients suitable for intra-patient dose escalation are stable or benefiting from the trial treatment prior to dose escalation, the incremental gains to such patients may be small compared to the risk of increased toxicity.

Combination therapies

Combination therapies are quite often handled the same way as single agent therapies. DLs are prespecified, and the doses are escalated and de-escalated according to the observed DLTs. Typically, a set of predetermined DL combinations are explored based on the single-agent MTD or other preclinical data demonstrating synergy. The dose of one agent under investigation is escalated while the dose of the second agent remains constant until a tolerable combination DL is achieved. Often all possible combination levels cannot be feasibly explored.

Several issues surrounding the design and identification of the MTD of a combination regimen need careful consideration. Ideally, an understanding of the underlying biologic rationale for the combination would be available. For example, are the toxicity profiles of the agents overlapping, or additive? Is the efficacy of the two agents' additive, complementary, or synergistic?

Despite the increased testing of such combination treatments in oncology, few designs for dose escalation of two or more agents have been proposed [26][27].

Among the most prominent design proposals are the BLRM for combination therapies, which was developed by Neuenschwander and colleagues [28] and works like the single-agent BLRM, but with 5 parameters in total (2 per agent plus one parameter reflecting the interaction between the two agents). Another proposal is a method called “product of independent beta probabilities dose escalation design for dual-agent phase I trials” (PIPE), developed by Mander and Sweeting [29].

Phase IB trials

Phase Ib trials are usually trials of the combination of investigational agents with standard anticancer agents or with other therapies (e.g. radiation therapy). These combination trials are often started after initial trials of the investigational agent have shown evidence of tolerability and some evidence of single-agent activity. Single agent studies may be unnecessary in situations where the new investigational agent acts as a cytotoxicity potentiator, lacking expected inherent anticancer activity. Investigators should provide a strong rationale for the drug combination, supported by preclinical studies, for proposed trials. New agents may be combined if there is sound evidence to support this study design. These combination investigational agent studies are evaluated on a case-by-case basis.

Conducted in patients diagnosed with the disease, or a condition for which the study drug is intended, phase Ib trials try to demonstrate some biomarker, surrogate, or possibly clinical outcome that could be considered for "proof of concept." Proof of concept in these kind of trials typically confirms the hypothesis that the current prediction of biomarker, or outcome benefit is compatible with the mechanism of action.

These trials are usually small with only 2 to 4 DLs to be tested. In these cases, the differences among different dose escalation methods are generally small. Nevertheless, simulations performed by the Statistics Unit at the SAKK CC and presented at the joint meeting of the International Biometric Society Austro-Swiss and Italian Regions 2015 have shown that even in these cases model-based designs should be preferred.

De-escalation

Instead of using a traditional dose escalation design, dose de-escalation might be applied in the situation where the drug under investigation has been used in other indications, the safety profile is well-known, and the investigators feel confident regarding the true MTD.

In these cases, a de-escalation works similar to an escalation design, but the starting dose is the highest one (instead of the lowest one). In case of (unexpected) high toxicity, the dose is adjusted to lower DLs.

Cohort size

For trials that enroll sequential cohorts with dose-escalation between cohorts, the choice of cohort size should consider the amount of risk that is acceptable in the study population. Larger cohorts might be necessary to provide reasonable assurance of safety before escalating the dose of a product intended to treat a disease that is less serious and for which the tolerance for accepting risk might be lower. Smaller cohorts might be adequate for a product that is intended to treat a serious or life-threatening disease where a greater potential benefit may justify a higher risk. Standardized protocol designs, such as the 3+3 design, are often used for dose escalation of oncology products. However, the cohort size in such a design might not be appropriate for other therapeutic areas where there is less tolerance of risk, and a larger cohort might be needed to provide a greater assurance of safety

prior to dose escalation. In addition, other study objectives, such as assessments of tolerability, feasibility, and pharmacologic activity may influence choice of cohort size.

For products, whose manufacturing capacity is limited, a practical limit on cohort size might be placed, particularly in the early clinical development. The prevalence of the proposed study population may also limit the cohort size. When considering the limitations due to manufacturing capacity and prevalence of the study population, sponsors should select a cohort size that is feasible, but still adequate to meet the study objectives [30].

Regarding how many patients should be allowed to be treated at a time it should be noted that this is mainly an issue of trial logistics and does not depend on the choice of design. All of the described phase I designs allow to recruit from one up to three patients in parallel (or even more in some designs).

Expansion cohorts

Recent reports show that phase I trials are increasingly using dose-expansion cohorts in order to better characterize the toxicity profiles of experimental agents or trial disease-specific cohorts. After the dose-escalation phase, an additional number of patients are treated at the MTD, who comprise the dose expansion cohort. The RP2D, defined as the DL recommended for future trials, is based on review of all the data and safety considerations, which could also include data obtained during the expansion phase.

While the primary aim of the expansion cohort remains to further evaluate the safety profile of an experimental agent and thus confirm the MTD, the broader aim varies from exploring pharmacokinetic (PK), pharmacodynamics (PD), efficacy, or even some biomarker-related endpoints.

The dose expansion cohorts have a large diversity, as it can be more than one, at more than one DL, with drug combinations, or even following a different treatment schedule. Often the dose expansion cohorts have different eligibility criteria in order to study a specific targeted population within a certain disease type, allowing some preliminary efficacy assessments. In some cases even numerous dose expansion cohorts with different disease entities are treated in parallel, so as to get some information in which indications one should focus in the subsequent phase II.

The simplest method to use the safety data of a dose-expansion cohort in the RP2D evaluation is by performing a retrospective analysis of the overall safety data. After having treated all patients in the dose-expansion cohort, the RP2D could be estimated based on both the MTD of the dose-escalation part and the safety data of the dose-expansion part. More complex designs allow changing the dose during the dose-expansion part adaptively, based on cumulating toxicity data [31].

There are also possibilities to include efficacy criteria in the estimation of the RP2D, by, e.g., using sequential model-based dose-finding algorithms that aim to recommend a dose optimal in both safety and efficacy [32].

Iasonos and O'Quigley [31] suggest that the size of the dose-expansion cohort should be at least 50% of the pre-expansion sample size, or a minimum of 12 to 15 patients, to obtain some meaningful preliminary evidence for efficacy. Boonstra and colleagues [33] conducted a simulation study and

conclude that a 10 to 20 patient dose expansion cohort substantially increases the probability of selecting the true MTD, if the dose-finding portion enrolls approximately 12 to 18 patients.

At SAKK we strongly reccomend to perform a (larger) dose-finding trial (using a model based design or mTPI) without a dose-expansion part instead of a smaller dose-finding (using, e.g., a 3+3 design) part with subsequent dose-expansion. This has the advantage of a more reliable estimate of the true MTD, based on all accrued patients, allowing dose changing using the same algorithm during the whole trial, when necessary.

Glossary

This section explains some of the most important terms which are often used in the context of phase I clinical trials.

Term	Definition
Cohort	Group of patients treated at a dose level.
Dose-limiting toxicity (DLT)	Toxic effects that are presumably related to the treatment that are considered unacceptable (because of their severity and/or irreversibility) and that limit further dose escalation. DLTs are defined before beginning the trial and are protocol specific. They are typically defined based on toxic effects seen in the first cycle and specified using a standardized grading criteria, for example, Common Terminology Criteria for Adverse Events (CTCAE).
Dose toxicity curve	It is a curve that reflects the relationship between treatment dose and probability of toxicity of the treatment.
Maximum tolerated dose (MTD)	Phase I trials that use model-based methods: the dose that produces the target toxicity level. Phase I trials conducted in the United States: the highest dose level at which ≤33% of patients experience DLT. Phase I trials conducted in Europe and Japan: the lowest dose level at which ≥33% of patients experience DLT (a misnomer in the sense that the MTD is actually not a tolerable dose). In SAKK we use the definition of MTD as in the United States.
Pharmacodynamics	Pharmacologic effects of the drug on the body (e.g., nadir neutrophil or platelet count, non-hematologic toxicity, molecular correlates, imaging endpoints).
Pharmacokinetics	Pharmacologic effects of the body on the drug (i.e., the time course of drug absorption, distribution, metabolism, and excretion).
Recommended phase II dose (RP2D)	Phase I trials with a toxicity endpoint that are conducted in the United States: the MTD. Phase I trials with a toxicity endpoint that are conducted in Europe and Japan: one dose level below the MTD. In SAKK we use the definition of RP2D as in the United States.
Target toxicity level	The maximum probability of DLT that is considered acceptable in the trial. The target toxicity level in phase I trials is typically between 20% and 33%.

References

[1] Storer BE (1989). Design and analysis of phase I clinical trials. *Biometrics*, 45(3): 925–37.

[2] Le Tourneau C, Lee JJ, Siu LL (2009). Dose Escalation Methods in Phase I Cancer Clinical Trials. *J Natl Cancer Inst*, 101: 708–20.

[3] Omura AG (2003). Modified Fibonacci search. *J Clin Oncol*. 21(16): 3177–80.

[4] Ji Y, Wang SJ (2013). Modified toxicity probability interval design: a safer and more reliable method than the 3+3 design for practical phase I trials. *J Clin Oncol*, 31(14): 1785–91.

[5] Thall PF and Lee SJ (2003). Practical model-based dose-finding in phase I clinical trials: Methods based on toxicity. *Int J Gynecol Cancer*, 13, 251–61.

[6] Goodman SN, Zahurak ML, Piantadosi S (1995). Some Practical Improvements in the Continual Reassessment Method for Phase I Studies. *Stat Med*, 14: 1149–61.

[7] Simon R, Freidlin B, Rubinstein L, Arbuck SG, Collins J, Christian MC (1997). Accelerated Titration Designs for Phase I Clinical Trials in Oncology. *J Natl Cancer Inst*, 89(15): 1138–47.

[8] Collins JM (2000). Innovations in Phase Trial Design: Where Do We Go Next? *Clin Cancer Res*, 6: 3801–2.

[9] Ji Y, Li Y, Bekele BN (2007). Dose finding in phase I clinical trials based on toxicity probability intervals. *Clin Trials,* 4: 235–44.

[10] Ji Y, Liu P, Li Y, Bekele BN (2010). A modified toxicity probability interval method for dose-finding trials. *Clin Trials,* 7: 653–63.

[11] O'Quigley J, Pepe M, Fisher L (1990). Continual Reassessment Method: A Practical Design for Phase 1 Clinical Trials in Cancer. *Biometrics,* 46(1): 33–48.

[12] O'Quigley J, Chevret S (1991). Methods for dose finding studies in cancer clinical trials: a review and results of a Monte Carlo study. *Stat Med*, 10: 1647–64.

[13] O'Quigley J, Shen LZ (1996). Continual Reassessment Method: A Likelihood Approach. *Biometrics,* 52: 673–84.

[14] Garrett-Mayer E (1996). The continual Reassessment Method for Dose-Finding Studies: A Tutorial. *Clin Trials*, 3: 57–71.

[15] Ivy SP, Siu LL, Garrett-Meyer E, Rubinstein L (2010). Approaches to Phase 1 Clinical Trial Design Focused on Safety, Efficiency, and Selected Patient Populations: A report from the Clinical Trial Design Task Force of the National Cancer Institute Investigational Drug Steering Committee. *Clin Cancer Res*, 16(6): 1726–36.

[16] Moller S (1995). An Extension of the Continual Reassessment Methods Using a Preliminary Up-and-Down Design in a Dose Finding Study in Cancer Patients, in Order to Investigate a Greater Range of Doses. *Stat Med*, 14: 911–22.

[17] Piantadosi S, Fisher JD, Grossman S (1998). Practical Implementation of a Modified Continual Reassessment Method for Dose-Finding trials. *Cancer Chemother Pharmacol*, 41: 429–36.

[18] Cheung YK, Chappell R (2000). Sequential Designs for Phase I Clinical Trials with Late Onset Toxicities. *Biometrics*, 56(4): 1177–82.

[19] Neuenschwander B, Branson M, Gsponer T (2008). Critical Aspects of the Bayesian Approach to Phase I Cancer Trials. *Stat Med*, 27: 2420–39.

[20] Muler J, McGinn C, Normolle D, Lawrence T, Brown D, Hejna G, Zalupski M (2004). Phase I Trial Using a Time to Event Continual Reassessment Strategy for Dose Escalation of Cisplatin Combined with Gemcitabine and Radiation Therapy in Pancreatic Cancer. *J Clin Oncol*, 22: 238–43.

[21] Huang B, Kuan PF (2014). Time-to-Event Continual Reassessment Method Incorporating Treatment Cycle Information with Application to an Oncology Phase I Trial. *Biom J*, 56(6): 933–46.

[22] Postel-Vinay S, Collette L, Paoletti X, Rizzo E, Massard C, Olmos D, Fowst C, Levy B, Mancini P, Lacombe D, Ivy P, Seymour L, Le Tourneau C, Siu LL, Kaye SB, Verweij J, Soria JC (2014). Towards new methods for the determination of dose limiting toxicities and the assessment of the recommended dose for further studies of molecularly targeted agents – Dose-Limiting Toxicity and Toxicity Assessment Recommendation Group for Early Trials of Targeted therapies, an European Organisation for Research and Treatment of Cancer-led study. *Eur J Cancer*, 50: 2040–9.

[23] Millward M, Underhill C, Lobb S, McBurnie J, Meech S, Gomez-Navarro J, Marshall M, Huang B, Mather C (2013). Phase I Study of Tremelimumab plus PF-3512676 in Patients with Melanoma or Advanced Solid Tumors. *Br J Cancer*, 108: 1998–2004.

[24] Bekele B, Ji Y, Shen Y, Thall P (2008). Monitoring Late-Onset Toxicities in Phase I Trials Using Predictive Risks. *Biostatistics*. 9: 442–57.

[25] Eisenhauer EA, Twelves C, Buyse M (2006). Phase I Cancer Clinical Trials: A Practical Guide. *Oxford University Press*, New York.

[26] Thall PF, Millikan RE, Mueller P, Lee SJ (2003). Dose-finding with two agents in phase I oncology trials. *Biometrics*, 59: 487–96.

[27] Mandrekar SJ (2014). Dose-Finding Trial Designs for Combination Therapies in Oncology. *J Clin Oncol*, 32(2): 65–7.

[28] Neuenschwander B, Matano A, Tang Z, Roychoudhury S, Wandel S, Bailey S (2015). A Bayesian Industry Approach to Phase I Combination Trials in Oncology. *Statistical Methods in Drug Combination Studies*: 95–135.

[29] Mander AP, Sweeting MJ (2015). A product of independent beta probabilities dose escalation design for dual-agent phase I trials. *Stat Med,* 34(8): 1261–76.

[30] Considerations for the Design of Early-Phase Clinical Trials of Cellular and Gene Therapy Products. Guidance for industry. FDA, 2015.

[31] Iansonos A, O'Quigley J (2013). Design Considerations for Dose-Expansion Cohorts in Phase I Trials. *J Clin Oncol*, 31(31): 4014–22.

[32] Ivanova A (2003). A new dose finding design for bivariate outcomes. *Biometrics*, 59(4): 1001–7.

[33] Boonstra PS, Shen J, Taylor JGM, Braun TM, Griffith KA, Daignault S, Kalemkerian GP, Lawrence TS, Schipper MJ (2015). A Statistical Evaluation of Dose Expansion Cohorts in Phase I Clinical Trials. *J Natl Cancer Inst*, 107(3).

[34] Storer BE (2001). An evaluation of phase I clinical trial designs in the continuous dose-response setting. *Stat med*, 20(16): 2399–408.

[35] Skolnik JM, Barrett JS, Jayaraman B, Patel D, Adamson PC (2008). Shortening the Timeline of Pediatric Phase I Trials: The Rolling Six Design. *J Clin Oncol*, 26(2): 190–5.

[36] Babb J, Rogatko A, Zacks S (1998). Cancer Phase I Clinical Trials: Efficient Dose Escalation with Overdose Control. *Stat Med*, 17: 1103–20.

[37] Collins JM, Grieshaber CK, Chabner BA (1990). Pharmacologically guided phase I clinical trials based upon preclinical drug development. *J Natl Cancer Inst*, 82: 1321–6.

[38] Chu PL, Lin Y, Shih WJ (2009). Unifying CRM. EWOC designs for phase I cancer clinical trials. *J Stat Plan Inference*, 139: 1146–63.

[39] Wages NA, Conaway MR, O'Quigley J (2013). Using the time-to-event continual reassessment method in the presence of partial orders. *Stat Med*, 32(1): 131–41.

[40] Mauguen A, Le Deley MC, Zohar S (2011). Dose-finding approach for dose escalation with overdose control considering incomplete observations. *Stat Med*, 30(13): 1584–94.

[41] Yuan Y, Yin G (2011). Robust EM Continual Reassessment Method in Oncology Dose Finding. *J Am Stat Assoc*, 106(495): 818–31.

[42] Liu S, Ning J (2013). A Bayesian Dose-finding Design for Drug Combination Trials with Delayed Toxicities. *Bayesian Anal,* 8(3): 703–22.

Appendix 1: Overview of merits and limitations of different designs

Design	Merits	Limitations
3+3 design	– Easy to implement, well understood without need of statistician – Familiar to investigators and ethics committees – Safe – Many modifications exist with cohorts of different numbers	– Ignores history other than that of the previous cohort – Performs same action under different situations – Not possible to re-escalate – Does not allow the investigator to change the targeted toxicity easily – Fixed cohort size (3 or 6) – The procedure is unnecessarily slow, leading to treatment of excessive numbers of patients at DLs less likely to be efficacious. – Low probability of selecting the true MTD – High variability in MTD estimates – Fixed doses, not flexible
Accelerated Titration Designs	– Easy to implement, well understood without the need of a statistician – Partially familiar to investigators and ethics committees – Faster than 3+3, as escalation is rapid during the first stage of the escalation procedure – Reduces the total number of patients treated in the trial – Reduces the number of patients treated at sub-therapeutic doses – May incorporate intra-patient dose-escalation – Model-based designs following the initial single-patient cohorts can be incorporated to create hybrid methods	– Similar limitations as 3+3 – More aggressive, therefore may be associated with more risk – Due to the conservativeness of the investigators, it has been observed that they are often used with an initial dose set much more conservative than would be done for the 3+3 design, reducing their effectiveness – Few PK/PD data
Modified Toxicity Probability Interval Method	– Can easily be summarized into a set of “3+3”-like rules before the trial starts – Simple implementation (much software available: Excel macro, R code, EAST, etc.) – Transparent to the clinical investigators (procedure and rules are predefined and presented in a table at the start of the trial and then used throughout the trial) – Has good statistical properties: • Outperforms the traditional 3+3 design both in terms of mean number of patients treated above MTD and the percentage of cor-	– Ignores history other than that of the previous cohort – A monotone dose-toxicity curve is necessary at the end of the study (isotonic regression)

	rect MTD selection • Can target any user-defined toxicity probability and performs comparably to model-based methods such as the CRM	
Continual Reassessment Method	– Few patients are treated at low, ineffective doses – Information from all treated patients are taken into consideration for the recommended dose	– Not easily understood by non-statisticians – It is imperative that a statistician is involved in CRM trial design and CRM trial conduct – Needs prior specifications – Trial length may be too long – Potential difficulties in acceptance from institutional review boards and regulatory agencies
Bayesian Logistic Regresssion Model	– Two-parameter logistic model estimates better the dose-toxicity curve – All available information from previous patients is used to guide dose assignment for the next cohort – Next recommended dose is the one which has the highest probability of having an acceptable and desired toxicity probability – Overdose control is taken into consideration	– Similar limitations as of the CRM

Appendix 2: Other dose escalation methods

Best-of-five design

This is an aggressive design appropriate for a target toxicity probability in the range 0.30–0.40 [33]. Starting at k=1, three patients are evaluated at dose k. If no patients have DLT, then the dose is escalated to dose k+1. If all patients have DLT, then the trial is stopped. Otherwise an additional patient is entered at the kth dose for a total of four. If now one of four patients has DLT, then the dose is escalated to dose k+1. If three of four have DLT, then the trial is stopped. Otherwise one last patient is entered at dose k for a total of five. If now two of five patients have DLT, then the dose is escalated to dose k+1, otherwise the trial is stopped.

As in the 3+3 design, the MTD is defined as the DL below the dose where the trial was stopped.

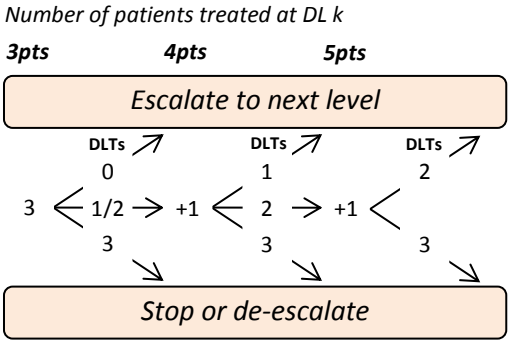

Figure 5. Best-of-five design.

Rolling-6 design

Skolnik and colleagues [35] proposed another modification to the 3+3 design. As a general rule of this design, accrual is suspended only when awaiting data from six patients. Nevertheless, DLTs are observed continuously, and the decision on whether to enroll a new patient onto the current, next highest or next lowest DL is based on the available data at the time of the new patient enrollment. So, DL assignment is based on the number of patients enrolled in the cohort, the number of DLTs observed and also on the number of patients at risk for experiencing a DLT.

Pharmacologically Guided Dose Escalation (PGDE) method

The PGDE method **Error! Reference source not found.** is another variation of the traditional 3+3 design that has not been widely used in clinical practice. This approach assumes that DLTs can be predicted by plasma drug concentrations and that animal models can accurately reflect this relationship in humans.

The PGDE method has two stages. A pre-specified plasma exposure defined by the area under the curve (AUC) for drug concentration as a function of time is extrapolated from preclinical data. Then, pharmacokinetic data are obtained for each patient in real time in order to determine the subsequent DL. As long as the pre-specified plasma exposure is not reached, dose escalation proceeds with one patient per DL and typically at 100% dose increments. When the target AUC is reached or if DLTs occur, dose escalation switches to the traditional 3+3 design with smaller (usually around 40%) dose increments (stage 2).

The PGDE method has not been widely adopted due to practical obstacles (see limitations below). So, despite some encouraging reports, the success of PGDE has been more academic than practical [8].

In clinical practice, the PGDE method has reliably defined the RP2D for some cytotoxic agents such as certain anthracyclines and platinum compounds but has been found to be inappropriate for other classes of cytotoxic agents such as the antifolates, which display a high inter-patient heterogeneity in pharmacokinetics (PK).

Merits:

- Faster dose escalation
- Fewer patients used
- Provides some data on PK inter-patient variability

Limitations:

- Logistic difficulties in obtaining real-time PK results, which are required to determine the safety of the subsequent dose escalation
- Differences between species used for dose estimation and humans may affect utility of the method
- Problems in extrapolating preclinical PK data to phase I studies with different treatment schedules
- Risk of exposing the next patient to a highly toxic dose if the AUC obtained in the preceding patient was atypically low due to inter-patient variability in drug metabolism

Escalation With Overdose Control (EWOC)

In order to overcome some criticism regarding the exposure of patients to high toxic doses when using the mCRM, Babb and colleagues [36] proposed a Bayesian adaptive design called EWOC. This method uses a similar idea to the CRM of treating each patient at the dose estimated to be closest to the MTD, but it places heavier penalties on overdosing than underdosing. The probability of administering a dose that exceeds the MTD for each higher DL is assessed after each patient, with an interdiction of dose escalation if this probability exceeds some critical pre-specified value (e.g. 25% and 5% for the probabilities of overdosing and excessive overdosing, respectively). Chu and colleagues [38] have shown that the continual reassessment method and the EWOC method can be unified in a hybrid model that seems to be able to determine the target dose more expeditiously than the EWOC method and results in smaller overdose proportions than the continual reassessment method.

Time-to-event CRM (TITE-CRM)

The TITE-CRM, originally proposed by Cheung and Chappell [18], is an extension of the original CRM method. Instead of using a binary outcome for toxicity, it uses a time-to-event approach with a weight function to incorporate time-to-toxicity of the patients in escalation decisions during the trial. The weight (from 0 to 1) is an increasing function of the follow-up time of the patient, and if a DLT occurs, the weight becomes 1. Patients with incomplete follow-up will be incorporated into the statistical model for MTD estimation without holding enrollment and thus this method results in a faster enrollment and shorter trial. Since then, variations of TITE-CRM have been developed including the work of Wages and colleagues [39], which extended the TITE-CRM design in the presence of partial ordering for a drug combination trial. Mauguen and colleagues [40] presented a hybrid design (TITE-EWOC) by introducing the time-to-event approach in the Escalation with Overdose Control (EWOC)

method [36]. Yuan and Yin [41] proposed an expectation-maximization CRM approach to handling late-onset toxicity. A recent paper by Liu and colleagues [42] proposed a data augmentation design for delayed toxicity by treating the unobserved toxicities as missing data.

Modified toxicity probability interval method – technical comments

A Beta-binomial model for toxicity probabilities at each DL is assumed. Given a target toxicity rate and toxicity outcomes at any dose, one calculates posterior probabilities of three non-overlapping intervals that partition the [0, 1] interval. These intervals correspond to the decisions of dose escalation, staying at the same dose, or dose de-escalation. The interval with the highest posterior probability triggers the decision for the next cohort of patients.

The algorithm includes a stopping rule which may facilitate early termination of a trial if dose 1 is excessively toxic, and an exclusion rule which prevents escalation to a dose that is likely to be highly toxic. At the end of the trial, the MTD is selected as the dose for which the isotonic-transformed posterior probability of toxicity is closest to the target toxicity probability.

Ji and colleagues [10] proposed a modified version of the TPI procedure (mTPI), based on the unit probability mass statistic (ratio of the probability mass of the interval and the length of the interval). The mTPI inherits many attractive features of the TPI procedure, but it is simpler as it only requires specification of an equivalence interval, within which any dose has a toxicity probability so close to the true MTD that physicians would agree to select them as the estimated MTD.

The trial is run until a pre-defined sample size is reached or a pre-defined number of patients are treated at a DL. Isotonic regression is used to determine the toxicity rates at each DL.

An attractive feature of this method is that all dose-escalation decisions can be pre-calculated and presented in a two-way table (as shown in Figure 2).