Published in final edited form as:

J Biopharm Stat. 2019 ; 29(4): 625–634. doi:10.1080/10543406.2019.1633654.

Rapid enrollment design for finding the optimal dose in immunotherapy trials with ordered groups

Xiaoqiang Xue^a, Matthew C Foster^b, Anastasia Ivanova^c

^aDecision Sciences, Data Science Safety and Regulatory, IQVIA Inc., Durham, NC, USA

^bLineberger Comprehensive Cancer Center, University of North Carolina, Chapel Hill, NC, USA

^cDepartment of Biostatistics, UNC at Chapel Hill, Chapel Hill, NC, USA

Abstract

In immunotherapy dose-finding trials, the optimal dose is usually defined based on both toxicity and response because the relationship between toxicity and response is different than that seen with cytotoxic antineoplastic therapies. In immunotherapy trials, toxicity and response often require a longer follow-up time compared to trials with cytotoxic agents. The rapid enrollment design has been proposed for dose-finding trials to find the maximum-tolerated dose where the follow-up for toxicity is long and it is desirable to assign a patient to a dose of a new therapy as soon as the patient is enrolled. We extend the rapid enrollment design to immunotherapy trials to find the optimal dose. We further describe how to use the design in immunotherapy trials with ordered groups where efficacy and safety considerations dictate running dose-finding trials in each group separately as efficacy and toxicity at the same dose can vary across groups. The estimation of the optimal dose in each of the groups can be improved in many, but not all, cases by using the monotonicity of toxicity and response among groups.

Keywords

Bayesian isotonic transformation; rapid enrollment design; RED; optimal safe dose; ordered groups; phase 1/2 trials

1. Introduction

Since most oncology therapies studied in the past were cytotoxic, the dose that could be administered to patients was limited by toxicity. A phase 1 study in oncology usually estimates the maximum-tolerated dose (MTD) and then therapeutic efficacy of the MTD is evaluated in a subsequent phase 2 trial. The MTD is usually defined as a dose with a probability of dose-limiting toxicity (DLT) close to target toxicity probability Γ , typically 0.20 or 0.30. Statistical designs for dose-finding methods in oncology have been recently reviewed by Tighiouart et al. (2014) and Sverdlov et al. (2014). In many immunotherapy trials, little or no toxicity has been observed (Weber et al. 2015). To avoid dose escalation to

very high doses without seeing any toxicity in a dose-finding trial of an immunotherapeutic agent, one can collect information on therapeutic response and escalate the dose until sufficient response is observed. In this case, the target dose in a dose-finding trial is defined based on both toxicity and efficacy. There are a number of ways to define the target dose based on toxicity and efficacy (Gooley et al., 1994; Braun 2002; Ivanova 2003; Thall and Cook 2004; Thall and Russell 1998). A number of dose-finding methods have been proposed to estimate the target dose defined based on toxicity and efficacy (Murtaugh and Fisher 1990; Gooley et al., 1994; Li, Durham, and Flournoy, 1995; Chen et al. 2015; Sato et al. 2016). The optimal dose in immunotherapies trials needs to be defined based on an efficacy-toxicity trade-off. Adoptive cell immunotherapies have unique toxicities such as cytokine release syndrome (CRS). Low-grade CRS is tolerable and associated with remission attainment, whereas more severe CRS and neurologic complications are life-threatening (Lee et al., 2015; Neelapu et al., 2018). As the dose of immunotherapy is being increased, the likelihood of attaining a remission increases and so does the likelihood of high-grade CRS. Therefore, we define the optimal dose as the dose that maximizes the probability of success defined as remission without dose-limiting toxicity. As in Ivanova (2003), we add an additional constraint to the definition of the optimal dose that the probability of DLT at the optimal dose should be less than Γ .

In some dose-finding trials, the patient population consists of two or more subpopulations with possibly different susceptibility to toxicity. Examples include adults and children, or subpopulations with different genotypes. Innocenti et al. (2004) suggested that UGT1A1 genotype predicts the occurrence of severe neutropenia during irinotecan therapy. Patients with the TA indel 7 of 7 genotype had much higher risk of developing grade four neutropenia than other patients (Innocenti et al. 2004). Trials in subpopulations with possibly different susceptibility to toxicity are referred to as trials in ordered groups. In immunotherapy trials, CRS and neurologic toxicity are better tolerated in certain populations. For example, it is known that children tolerate CRS and neurologic toxicity better than adults. O'Quigley and Paoletti (2003) proposed a modification of the continual reassessment method (CRM) for trials with ordered groups. Ivanova and Wang (2006) proposed an up-and-down design for ordered groups where the order among the groups is modeled using isotonic regression. Patients in different subgroups, in addition to having differential toxicity, can have different probabilities of therapeutic response at a given dose. The problem we are considering is the ordered group problem where both toxicity and response can be monotone with dose and group.

Follow-up time for toxicity or response can be rather long in a trial with an immunotherapy agent. While cytotoxic chemotherapy often results in early, predictable toxicities such as myelosuppression, immune-related adverse events are less predictable. For instance, monoclonal antibodies that interrupt immune checkpoint pathways can cause autoimmune phenomena that are delayed months from the start of therapy, and may or may not be associated with tumor response (Naidoo et al. 2015). A number of methods have been developed for dose-finding trials with delayed outcome. These methods allow using information from all patients, not only patients who completed their follow-up. Cheung and Chappell (2000) proposed an extension of the CRM to trials with long follow-up for toxicity, the time-to-event CRM (TITE-CRM). Salter et al. (2015) described the TITE-CRM for

ordered groups. Ivanova et al. (2016) proposed a more conservative approach than the one used in TITE-CRM to handle potential future toxicity in patients still in follow-up. The rapid enrollment design (RED) (Ivanova et al. 2016) is a Bayesian dose-finding design that imputes potential toxicities from patients still in follow-up and integrates with the decision rule regarding what dose to assign to the next patient.

In this paper, we extend the RED to the problem of finding the optimal dose and show how to apply the new design in dose-finding trials with ordered groups. This work was motivated by an ongoing dose-finding trial of autologous chimeric antigen receptor-modified T cells (CAR-T cells) in patients with acute lymphoblastic leukemia. In trials with CAR-T cells, a molecular safety switch can be triggered by a small molecule. Since little or no toxicity was experienced by subjects in other trials of this small molecule (Zhou et al. 2015), in our trial the target dose of the small molecule was defined as the dose where the probability of therapeutic response (improvement in cytokine release syndrome) without toxicity is maximized among doses with the probability of toxicity of at most 0.25. Patients are enrolled in two separated cohorts: pediatric (from 3 to 18 years old) and adult (>18 years old). Pediatric patients show higher tolerance for the CAR-T cells as compared to adult patients, hence we assume that the probability of unacceptable CAR-T cell toxicity in the adult population is higher than or equal to the probability of such toxicity in pediatric patients treated with the same CAR-T cell dose.

In this paper, we describe the estimation procedure and give examples in Section 2. The design is described in Section 3. In Section 4, we show how to temporarily impute unobserved outcomes to mitigate the uncertainty from patients still in follow-up when a new assignment is made. Simulations are presented in Section 5 and conclusions in Section 6.

2. Estimating under order restrictions

2.1. Notation

Each patient is followed for toxicity and response for a fixed period of time. Toxicities and responses observed outside the observation window do not count. With regard to toxicity (T) and therapeutic response (R), there are four possible outcomes: no response and no toxicity, T^-R^- , response and no toxicity, T^-R^+ , no response and toxicity, T^+R^- , and response and toxicity, T^+R^+ . Let groups $1, 2, \dots, I$ be the sub-populations with different susceptibility to toxicity. For a subject in group i , $i = 1, \dots, I$, assigned to dose k , $k = 1, 2, \dots, K$, denote $\pi_{ik1} = \Pr(T^-R^-)$, $\pi_{ik2} = \Pr(T^-R^+)$, $\pi_{ik3} = \Pr(T^+R^-)$ and $\pi_{ik4} = \Pr(T^+R^+)$. Observations from different subjects are independent. Let $\{y_{ik1}, y_{ik2}, y_{ik3}, y_{ik4}\}$ be the number of patients at dose k , $k = 1, 2, \dots, K$, and in group i , $i = 1, 2, \dots, I$, out of $n_{ik} = \sum_{l=1}^4 y_{ikl}$ patients with outcomes T^-R^- , T^-R^+ , T^+R^- , and T^+R^+ , correspondingly. Further denote the number of patients with positive therapeutic response by $y'_{ik2} = y_{ik2} + y_{ik4}$, and the number of patients with toxicity by $y'_{ik3} = y_{ik3} + y_{ik4}$. Let $\pi'_{ik2} = \pi_{ik2} + \pi_{ik4}$ be the probability of response, and $\pi'_{ik3} = \pi_{ik3} + \pi_{ik4}$ be the probability of toxicity. The vector of counts $(y_{ik1}, y_{ik2}, y_{ik3}, y_{ik4})$ follows a multinomial distribution with parameters n_{ik} and $(\pi_{ik1}, \pi_{ik2}, \pi_{ik3}, \pi_{ik4})$.

2.2. Isotonic estimation of toxicity and response probabilities in a trial with ordered groups

Since $y_{il} = (y_{1l}, y_{2l}, \dots, y_{Kl})$ follows a multinomial distribution the likelihood is.

$$L[y_{1l}, y_{2l}, \dots, y_{Il} | \pi_{1l}, \pi_{2l}, \dots, \pi_{Il}] = \prod_{i=1}^I \prod_{k=1}^K \prod_{l=1}^4 \{\pi_{ikl}\}^{y_{ikl}}.$$

We use a Bayesian approach and assume $\text{Dirichlet}(\alpha_{k1}, \alpha_{k2}, \alpha_{k3}, \alpha_{k4})$ prior on $(\pi_{ik1}, \pi_{ik2}, \pi_{ik3}, \pi_{ik4})$. The Dirichlet is the conjugate prior for multinomial distribution with posterior distribution computed as.

$$(\pi_{ik1}, \pi_{ik2}, \pi_{ik3}, \pi_{ik4} | y_{ik1}, y_{ik2}, y_{ik3}, y_{ik4}) \sim \text{Dirichlet}(\alpha_{k1} + y_{ik1}, \alpha_{k2} + y_{ik2}, \alpha_{k3} + y_{ik3}, \alpha_{k4} + y_{ik4}).$$

To use the information on monotonicity of response and toxicity with dose within each group and among groups, we use the Bayesian isotonic transformation (Dunson and Neelon 2003), to project from $R^{3 \times I \times K}$ into the restricted space Ω using a minimal distance mapping. The number of dimensions in $R^{3 \times I \times K}$, $3 \times I \times K$, is the number of groups, I , multiplied by the number of doses in each group, K , and multiplied by 3, the number of degrees of freedom in the 2×2 table as $\pi_{ik1} + \pi_{ik2} + \pi_{ik3} + \pi_{ik4} = 1$. The restricted space $\Omega \subset R^{3 \times I \times K}$ is a set of vectors of multinomial probabilities such that

- (a) $\pi'_{i13} \leq \pi'_{i23} \leq \dots \leq \pi'_{iK3}$, $i = 1, 2, \dots, I$;
- (b) $\pi'_{i12} \leq \pi'_{i22} \leq \dots \leq \pi'_{iK2}$, $i = 1, 2, \dots, I$;
- (c) $\pi'_{1k3} \geq \pi'_{2k3} \geq \dots \geq \pi'_{Ik3}$, $k = 1, 2, \dots, K$;
- (d) $\pi'_{1k2} \leq \pi'_{2k2} \leq \dots \leq \pi'_{Ik2}$, $k = 1, 2, \dots, K$.

Here restriction a) reflects the assumption of monotonicity of toxicity with dose in each group; restriction b) reflects the monotonicity of response with dose in each group; restriction c) indicates that, for a given dose, toxicity is monotone with group number; and restriction d) indicates that, for a given dose, response is monotone with the group number. In all problems known to us the direction of the order a) is as in (1). In some cases monotonicity of response, condition b), cannot be assumed. In many cases, however, we can assume that the probability of response is increasing and then decreasing, that is, follows an umbrella order with unknown peak. Our method can be modified to accommodate the umbrella order. The direction of the order in c) and d), increasing or decreasing, can vary depending on the problem. That is, both toxicity and response can be increasing with group number, toxicity can be increasing and response decreasing, or nothing can be said about the order of response probabilities in d).

Following Dunson and Neelon (2003), the Bayesian isotonic transformation is obtained by drawing samples from the posterior distribution of unconstrained parameter vectors $(\pi^*_{ik1}, \pi^*_{ik2}, \pi^*_{ik3}, \pi^*_{ik4})$, $k = 1, 2, \dots, K$, $i = 1, 2, \dots, I$. To map this set of unconstrained vectors

to a constrained set, we first calculate the number of each of the four outcomes at each dose level k within each group i that corresponds to the draw. The numbers are computed as $\tilde{y}_{ikl} = \left(\sum_{l=1}^4 y_{ikl} + \sum_{l=1}^4 \alpha_l \right) \times \pi_{ikl}^*$. Treating the counts $\tilde{y}_{ikl}$ as observed data we obtain constrained maximum likelihood estimates of the parameter vector by maximizing the likelihood $\hat{\Pi} = \sup \prod_{i=1}^I \prod_{k=1}^K \prod_{l=1}^4 (\pi_{ikl})^{\tilde{y}_{ikl}}$ where the supremum is taking over all vectors $(\pi_{ik1}, \pi_{ik2}, \pi_{ik3}), k = 1, 2, \dots, K, i = 1, 2, \dots, I$, that satisfy restrictions a), b), c) and d) in (1). These estimates serve as a projection of the initial draw to the constrained set of parameter vectors.

2.3. Examples

We illustrate the computation described in Section 2.2 by two examples. We maximize the likelihood by using an optimization routine `constrOptim` in R software (R Core Team 2017). The first example is a hypothetical example with a single group and 4 dose levels after 21 patients have been assigned. Table 1 displays the observed counts in the four cells of the 2×2 table for each dose and observed proportions or responses and toxicities. We present posterior means of transformed variables under the order constraints on the probability of toxicity $\pi'_{13} \quad \pi'_{23} \quad \pi'_{33} \quad \pi'_{43}$ and the probability of therapeutic response $\pi'_{12} \quad \pi'_{22} \quad \pi'_{32} \quad \pi'_{42}$. We used a Dirichlet prior with parameter vector (0.5, 0.5, 0.5, 0.5).

An example with two ordered groups is displayed in Table 2, where 42 subjects, 21 in group 1 and 21 in group 2, have been assigned to 3 dose levels. The Bayesian isotonic transformation estimates are calculated under the assumption a)-d) in (1). That is, both toxicity and therapeutic response are monotone with dose within each group, and, at each dose, toxicity is higher in group 1 compared to group 2, and that response in group 1 is lower compared to group 2.

3. The rapid enrollment design to find the optimal dose

Consider the problem of finding the optimal dose in the case of no stratification in the population, that is, a single group problem. We define the optimal dose as the dose where the probability of response without toxicity is maximized, subject to the restriction on the marginal probability of DLT. Ivanova (2003) noted that in plausible scenarios this goal can be replaced with the goal of finding a dose where the probability of toxicity is equal to the probability of no response and no toxicity. Therefore, we will be working with three outcomes that can be obtained from the 2×2 table in Section 2. For ease of notation we drop a group index in this section, and use $\pi_{k1}, \pi_{k2}, \pi_{k3}, \pi'_{k3}$, to denote the probabilities of no toxicity and no response, no toxicity and response, and toxicity at dose k , correspondingly. Let $\hat{\pi}_{k1}, \hat{\pi}_{k2}, \hat{\pi}'_{k3}$ be corresponding posterior means obtained under order restrictions as described in Section 2. Ivanova (2003) replaced the objective of maximizing the probability of response and no toxicity as a function of dose with the objective of finding the root of the equation $\pi'_{k3} = \pi_{k1}$. Our dose escalation algorithm is based on looking for a dose q such that, $\pi'_{q3} = \pi_{q1}$. Because $\pi'_{q3} - \pi_{q1}$ is monotone with dose, it is much easier to solve the equation, $\pi'_{q3} - \pi_{q1} = 0$, than maximize the probability of response and no toxicity as a function of dose (Ivanova 2003). Additionally, we are imposing a safety

constraint that the probability of toxicity at the optimal dose is less than the target level Γ , $\pi'_{q3} < \Gamma$. In immunotherapy trials, it can happen that there is sufficient efficacy and not much toxicity at most doses. To avoid escalation to higher doses in this case, we impose the sufficient efficacy condition and check if we already reached a dose with sufficient efficacy $\pi_{q2} > \gamma$. This prevents unnecessary escalation to higher doses.

The proposed design to find the optimal dose can be viewed as an extension of the rapid enrollment design (Ivanova et al. 2016). Below we describe the new rapid enrollment design for optimal dose (REDO) applied to a single group. Let $\{1, \dots, k\}$ be the set of doses with at least one patient assigned.

1. *Initial escalation.* Do not change the dose unless at least m patients are assigned to the dose and their toxicity outcomes are observed (except when the safety rule is invoked). We recommend using cohorts of $m=3$. Assigning at least three patients prevents making a decision based on insufficient information.
2. If $\hat{\pi}_{k1} > \hat{\pi}'_{k3}$ and at least m patients have been assigned to dose level k , the next patient is assigned to dose level $k+1$ (or k if $k=K$).
3. If $\hat{\pi}_{k1} \leq \hat{\pi}'_{k3}$ and if there is a dose s such that $\hat{\pi}_{s1} = \hat{\pi}'_{s3}$, the next patient is assigned to dose level s . Otherwise, let $s, s+1$ be such that $\hat{\pi}_{s1} > \hat{\pi}'_{s3}$ and $\hat{\pi}_{(s+1)1} < \hat{\pi}'_{(s+1)3}$. That is, given the data, the optimal dose is somewhere between dose levels s and $s+1$. Let probability γ_s reflect how close the difference $\pi'_{s3} - \pi_{s1}$ is to 0, $\gamma_s = \Pr\{-e < \pi'_{s3} - \pi_{s1} < e\}$. The next patient is assigned to the dose s^* , where $s^* = \operatorname{argmax}(\gamma_s, \gamma_{s+1})$, that is, to the dose corresponding to the higher value of γ_s and γ_{s+1} .
4. *Safety rule.* Do not assign patients to the dose with $\Pr\{\pi'_{k3} > \Gamma\} > 0.9$, $k=1, 2, \dots, K$. If $\Pr\{\pi'_{k3} > \Gamma\} > 0.9$, the trial is stopped because the lowest dose is too toxic.
5. *Sufficient efficacy rule.* For $s^* \geq 2$, if $\Pr\{\pi_{(s^*-1)2} > \gamma\} > 0.7$ and $\Pr\{\pi_{s^*2} > \gamma\} > 0.7$, that is, both doses s^*-1 and s^* have sufficient efficacy, assigned the next patient to dose s^*-1 . Otherwise, assign to s^* .
6. *Estimation of the optimal dose at the end of the trial.* Let N be the total sample size in the trial divided by the number of doses. That is, N is the average number of subjects assigned to a dose. At the end of the trial compare proportions of successes at the doses with at least N subjects assigned. The dose with the highest estimated probability of success and tolerable toxicity rate according to the safety rule 5) is declared the optimal dose. No dose is selected if the lowest dose is deemed too toxic by the safety rule 5).

To apply this design in the case of ordered groups, e.g. children and adults, a separate dose-escalation as described above is run in children and in adults. At each step, all data, from children and adults are used to estimate the quantities in the 2×2 matrix for each dose and group as described in Section 2.

4. Mitigating for delayed toxicity and clinical response

In this section, we propose a method for temporarily imputing toxicity and response for patients still in follow-up. The goal is to have a conservative approach for imputation to use with the design described in Section 3. Let T_1 be the length of follow-up for toxicity, and T_2 be the length of follow-up for response. Both response and toxicity are defined based on a given follow-up time T_1 and T_2 . We would like to assign a new patient to a dose using partial information from patients being followed for toxicity and response. To mitigate the uncertainty regarding toxicity due to patients still in follow-up for toxicity we add temporary toxicities. Temporary toxicities are added to make the design less likely to escalate to higher doses as this is almost always a requirement in a phase 1 trial. As it usually takes longer to observe response than toxicity, here we assume $T_1 < T_2$. The case $T_1 > T_2$ can be handled similarly. Let u be the patient's current follow-up time. For the design in Section 3, we consider three outcomes: response and no toxicity, T^-R^- , no response and toxicity, T^-R^+ , and toxicity T^+ . At the end of follow-up for both efficacy and toxicity patient's outcome will be one of the three possible outcomes (1,0,0), (0,1,0), or (0,0,1). If toxicity has been observed already in the interval $(0, u)$, the patient is assigned the outcome (0,0,1) and we do not add any temporary outcomes. If $T_1 < u < T_2$, no imputation is necessary because the toxicity observation window has closed. If the patient is still in follow-up for toxicity, $u < T_1$, and therapeutic response has been observed already, the temporary outcome we add is $(0, u/T_1, 1 - u/T_1)$. If $u < T_1$, and neither toxicity nor therapeutic response has been observed yet, the temporary outcome to assign to this patient is $(0, u/T_1, 1 - u/T_1)$. The reason we do not add temporary toxicity outcomes to both no response and response categories is because we are choosing the most conservative approach, that is, the approach resulting in the slowest escalation of doses. Since we escalate the dose when the estimated probability of the first outcome, no response and no toxicity, is higher than the probability of the third outcome, toxicity, the temporary outcome resulting in the slowest dose escalation is $(0, u/T_1, 1 - u/T_1)$.

5. Simulation study

Table 3 presents six simulation scenarios. Group 1 probabilities in the first four scenarios are from O'Guiley et al. (2001). As O'Guiley et al. (2001) considered the case of a single group, we set the probabilities in the second group so that the probability of toxicity is the same or lower compared to the first group. We added a scenario with a sudden increase in probability of toxicity to verify that the safety rule provides necessary protection from overdosing of patients. Scenario 6 is a scenario where the lowest dose has sufficient efficacy. This scenario was added to make sure that the sufficient efficacy rule is working and that we are not going to escalate to higher doses. We used $\Gamma = 0.25$ and $\alpha = 0.6$. Dirichlet prior with the parameter vector (0.5,0.5,0.5) was used. Patients were assigned one at a time. The sample size was 20 patients per trial or per group in trials with ordered groups. For each scenario of design, 1000 simulation runs were run. We compared the designs in a set-up when outcome is immediate and when the outcome is delayed. For delayed outcome simulations, the observation window for toxicity was 5 weeks from the start of the therapy, and the observation window for therapeutic response was 10 weeks from start of the therapy. A new patient was enrolled on Monday every other week. For patients with DLT, the time to DLT

was generated following uniform distribution in $(0,35)$, and for patients with therapeutic response and no DLT, the time to clinical response was generated from uniform $(0,70)$.

We compared REDO with Ivanova optimal dose design (IOD) from Ivanova (2003). The IOD is a group up-and-down design to find the dose where the probability of toxicity is equal to the probability of no response and no toxicity. The decision rule for IOD is as follows. Suppose the most recent subject was allocated to level $j \in \{1, \dots, K\}$. Assign the next subject to

- i. dose d_{j-1} , if the most recent subject had toxicity;
- ii. dose d_j , if the most recent subject had response and no toxicity;
- iii. dose d_{j+1} , if the most recent subject had no response and no toxicity.

Table 4 presents the comparison of REDO and IOD. We show results for immediate outcome for IOD and REDO and results for REDO with delayed outcome. When outcomes are delayed we used the imputation method described in Section 4. Overall the new design performs better than IOD. Additionally, the new design can be applied to trials with delayed outcomes, while the IOD cannot accommodate delayed outcomes in a straightforward manner. When outcome is delayed, the REDO with mitigation of toxicities tends to recommend lower doses. This is expected because escalation to higher doses is slower under the conservative toxicity mitigation rule. Since the escalation with delayed outcome is conservative, REDO does not reach the optimal dose when it is the highest dose as in Scenario 4. The planned number of patients in trials with delayed toxicity should be higher than in trials with slow accrual or when toxicity outcome is known quickly. For example, when the sample size is increased from 20 to 30 patients per trial, the proportion of trials that recommend the highest dose in this scenario increases from 0.34 to 0.61. Despite slower escalation, clinicians prefer using this conservative mitigation rule since major focus is usually on safety and preventing rapid escalation to possibly unsafe doses. The safety rule provides a good protection from overdosing of patients in Scenario 5. The sufficient efficacy rule in REDO is effective in Scenario 6 and prevents unnecessary escalation to higher doses. IOD was simulated without the sufficient efficacy rule and therefore it continued to escalate to higher doses.

Table 5 shows the results for the case of two ordered groups with 20 patients assigned within each group. The two groups exchange the information through isotonic estimation of the probabilities of the three outcomes. This results in improved performance in group 1 compared to a single group set-up (Table 4) in all scenarios except scenario 4 (Table 5). In scenario 4, the escalation is slower when ordered groups are considered because the probabilities of toxicity and efficacy in group 1 and group 2 are very similar. One of the reasons to run a trial with ordered groups rather than a single trial is to ensure safety. Combining two subpopulations with different toxicity profiles might lead to overdosing of patients with high susceptibility to toxicity. The requirement to run a trial with ordered groups can, however, lead to a decrease in efficiency of the optimal dose estimation as in scenario 4. If the groups are truly ordered, the estimation is improved. Overall the new design yields good estimation of the optimal dose in both groups.

6. Conclusions

We extended the rapid enrollment design to dose-finding trials with the goal of finding the optimal dose. The new rapid enrollment design to find the optimal dose, REDO, performs better than the design in Ivanova (2003). In many immunotherapy trials, the follow-up for toxicity and response is long and hence decisions about what dose to assign need to be made based on partial data. We proposed a mitigation rule that imputes potential toxicities from patients still in follow-up for toxicity. Such a rule plays an important role in the beginning of the trial when no information is available at a dose. We took a conservative approach to mitigate uncertainty of unobserved toxicities as this is often the choice of phase 1 investigators. We describe how to apply the new design in the case where there are two or more subpopulations with different toxicity and/or response profiles. Separate dose-finding trials are usually run as the maximum-tolerated and optimal doses might be different across subpopulations. If one of the subpopulations is known to be more susceptible to toxicity, ethical considerations dictate to take the known order of toxicity rates among the subpopulations into account and model the toxicity jointly across subpopulations. Joint modeling improves the accuracy of estimation resulting in more precise dosing of patients in the trial. We do not recommend using joint modeling if the order of subpopulations is not known as it can lead to bias of the toxicity and response rate estimates. The sample size in a dose-finding trial is usually not large enough to estimate the order between subpopulations with sufficient confidence, leaving the investigators to rely on the information available before the trial.

Acknowledgments

The authors thank the two guest editors for overseeing the review process for this special issue and anonymous reviewers for tremendously helpful comments.

Funding

This work was supported by the NIH [P01 CA142538].

References

- Braun TM 2002 The bivariate continual reassessment method: Extending the CRM to phase I trials of two competing outcomes. *Controlled Clinical Trials* 23 (3):240–256. [PubMed: 12057877]
- Chen Z, Yuan Y, Li Z, et al. 2015 Dose escalation with over-dose and under-dose controls in phase I/II clinical trials. *Contemporary Clinical Trials* 43:133–141. doi:10.1016/j.cct.2015.05.014. [PubMed: 26012358]
- Cheung YK, and Chappell R. 2000 Sequential designs for phase I clinical trials with late-onset toxicities. *Biometrics* 56:1177–1182. [PubMed: 11129476]
- Dunson DB, and Neelon B. 2003 Bayesian inference on order-constrained parameters in generalized linear models. *Biometrics* 59:286–295. [PubMed: 12926713]
- Gooley TA, Martin PJ, Fisher LD, and Pettinger M. 1994 Simulation as a design tool for phase I/II clinical trials: An example from bone marrow transplantation. *Controlled Clinical Trials* 15:450–462. doi:10.1016/0197-2456(94)90003-5. [PubMed: 7851107]
- Innocenti F, Undevia SD, Iyer L, Chen PX, Das S, Kocherginsky M, Karrison T, Janisch L, Ramirez J, Rudin CM, et al. 2004 Genetic variants in the UDP-glucuronosyltransferase 1A1 gene predict the risk of severe neutropenia of irinotecan. *Journal of Clinical Oncology* 22 (8):1382–1388. doi: 10.1200/JCO.2004.07.173. [PubMed: 15007088]

- Ivanova A. 2003 A new dose-finding design for bivariate outcomes. *Biometrics* 59:1001–1007. [PubMed: 14969479]
- Ivanova A, and Wang K. 2006 Bivariate isotonic design for dose-finding with ordered groups. *Statistics in Medicine* 25:2018–2026. doi:10.1002/sim.2312. [PubMed: 16220476]
- Ivanova A, Wang Y, and Foster MC. 2016 The rapid enrollment design for phase I clinical trials. *Statistics in Medicine* 35 (15):2516–2524. doi:10.1002/sim.6886. [PubMed: 26833922]
- Lee DW, Kochenderfer JN, Stetler-Stevenson M, Cui YK, Delbrook C, Feldman SA, Fry TJ, Orentas R, Sabatino M, Shah NN, et al. 2015 T cells expressing CD19 chimeric antigen receptors for acute lymphoblastic leukaemia in children and young adults: a phase 1 dose-escalation trial. *Lancet* 385:517–528. doi:10.1016/S0140-6736(14)61403-3. [PubMed: 25319501]
- Li W, Durham SD, and Flournoy N. 1995 An adaptive design for maximization of a contingent binary response In *New developments and applications in experimental design*, ed. Flournoy N, Roserberger WF, and Wong WK, 50–61. Hayward, CA: Institute of Mathematical Statistics.
- Murtaugh PA, and Fisher LD. 1990 Bivariate binary models of efficacy and toxicity in dose-ranging trials. *Communications in Statistics. Theory and Methods*. 19 (6):2003–2020. doi: 10.1080/03610929008830305.
- Naidoo J, Page DB, Li BT, Connell LC, Schindler K, Lacouture ME, Postow MA, and Wolchok JD. 2015 Toxicities of the anti-PD-1 and anti-PD-L1 immune checkpoint antibodies. *Annals of Oncology* 26 (12):2375. [PubMed: 26371282]
- Neelapu SS, Tummala S, Kebriaei P, Wierda W, Gutierrez C, Locke FL, Komanduri KV, Lin Y, Jain N, Daver N, et al. 2018 Chimeric antigen receptor T-cell therapy – Assessment and management of toxicities. *Nature Reviews Clinical Oncology* 15(1):47–62. doi: 10.1038/nrclinonc.2017.148.
- O'Guiley J, Huges MD, and Fenton T. 2001 Dose-finding designs for HIV studies. *Biometrics* 57:1018–1029. [PubMed: 11764240]
- O'Quigley J, and Paoletti X. 2003 Continual reassessment method for ordered groups. *Biometrics* 59:430–440. [PubMed: 12926728]
- R Core Team. 2017 R: A language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing <https://www.R-project.org/>.
- Salter A, O'Quigley J, Cutter GR, and Aban IB. 2015 Two-group time-to event continual reassessment method using likelihood estimation. *Contemporary Clinical Trials* 45:340–345. doi:10.1016/j.cct.2015.09.016. [PubMed: 26409251]
- Sato H, Hirakawa A, and Hamada C. 2016 An adaptive dose-finding method using a change-point model for molecularly targeted agents in phase I trials. *Statistics in Medicine* 35 (23):4093–4109. doi:10.1002/sim.6981. [PubMed: 27221807]
- Sverdlov O, Wong WK, and Ryznik Y. 2014 Adaptive clinical trial designs for phase I cancer studies. *Statistics Surveys* 8:2–44. doi:10.1214/14-SS106.
- Thall PF, and Cook JD. 2004 Dose-finding based on efficacy-toxicity trade-offs. *Biometrics* 60 (3): 684–693. doi:10.1111/j.0006-341X.2004.00218.x. [PubMed: 15339291]
- Thall PF, and Russell KE. 1998 A strategy for dose-finding and safety monitoring based on efficacy and adverse outcomes in phase I/II clinical trials. *Biometrics* 54 (1):251–264. [PubMed: 9544520]
- Tighiouart M, Liu Y, and Rogatko A. 2014 Escalation with overdose control using time to toxicity for cancer phase I clinical trials. *PLoS One* 93: e93070. doi: 10.1371/journal.pone.0093070
- Weber JS, Yang JC, Atkins MB, and Disis ML. 2015 Toxicities of immunotherapy for the practitioner. *Journal of Clinical Oncology* 33:2092–2099. doi:10.1200/JCO.2014.60.0379. [PubMed: 25918278]
- Zhou X, Dotti G, Krance RA, Martinez CA, Naik S, Kamble RT, Durett AG, Dakhova O, Savoldo B, Di Stasi A, et al. 2015 Inducible caspase-9 suicide gene controls adverse effects from alloplete T cells after haploidentical stem cell transplantation. *Blood* 125 (26):4103–4113. doi:10.1182/blood-2015-02-628354. [PubMed: 25977584]

Table 1.

An example of Bayesian isotonic estimation in a single group. Observed counts are presented as the number of outcomes over the total number of patients at the dose.

	Dose	d_1	d_2	d_3	d_4
Observed Counts	T^-R^-	2/3	2/6	2/6	0/6
	T^-R^+	1/3	1/6	2/6	2/6
	T^+R^-	0/3	0/6	1/6	1/6
	T^+R^+	0/3	3/6	1/6	3/6
	T^+	0.00	0.50	0.33	0.67
	R^+	0.33	0.67	0.50	0.83
Bayesian Isotonic Posterior Mean	T^-R^-	0.54	0.41	0.25	0.06
	T^-R^+	0.28	0.19	0.30	0.31
	T^+R^-	0.10	0.06	0.18	0.18
	T^+R^+	0.08	0.34	0.27	0.45
	T^+	0.18	0.40	0.45	0.63
	R^+	0.36	0.53	0.57	0.76

Table 2.

An example of Bayesian isotonic estimation in two groups under constraints in (1). Observed counts are presented as the number of outcomes over the total number of patients at the dose.

	Dose	Group 1			Group 2		
		d ₁	d ₂	d ₃	d ₁	d ₂	d ₃
Observed Counts	T ⁻ R ⁻	2/3	2/6	2/6	1/3	2/6	2/6
	T ⁻ R ⁺	1/3	1/6	2/6	1/3	1/6	2/6
	T ⁺ R ⁻	0/3	0/6	2/6	1/3	1/6	1/6
	T ⁺ R ⁺	0/3	3/6	0/6	0/3	2/6	1/6
	T ⁺	0.00	0.50	0.33	0.33	0.50	0.33
	R ⁺	0.33	0.67	0.33	0.33	0.50	0.50
Bayesian Isotonic Posterior Mean	T ⁻ R ⁻	0.53	0.40	0.24	0.46	0.32	0.26
	T ⁻ R ⁺	0.18	0.14	0.28	0.29	0.26	0.29
	T ⁺ R ⁻	0.18	0.13	0.28	0.18	0.15	0.17
	T ⁺ R ⁺	0.11	0.33	0.20	0.07	0.27	0.28
	T ⁺	0.29	0.46	0.48	0.25	0.42	0.45
	R ⁺	0.29	0.47	0.48	0.36	0.53	0.57

Table 3.

Scenarios with four doses and two ordered groups.

Scenario	Outcome	Group 1				Group 2			
		d ₁	d ₂	d ₃	d ₄	d ₁	d ₂	d ₃	d ₄
1	Pr(T ⁻ R ⁻)	0.74	0.13	0.15	0.20	0.77	0.45	0.15	0.10
	Pr(T ⁻ R ⁺)	0.20	0.70	0.60	0.50	0.20	0.50	0.70	0.60
	Pr(T ⁺)	0.06	0.17	0.25	0.30	0.03	0.05	0.15	0.30
2	Pr(T ⁻ R ⁻)	0.17	0.20	0.10	0.10	0.65	0.18	0.15	0.20
	Pr(T ⁻ R ⁺)	0.70	0.50	0.50	0.40	0.30	0.65	0.50	0.40
	Pr(T ⁺)	0.13	0.30	0.40	0.50	0.05	0.17	0.35	0.40
3	Pr(T ⁻ R ⁻)	0.90	0.65	0.15	0.20	0.90	0.68	0.40	0.16
	Pr(T ⁻ R ⁺)	0.10	0.30	0.70	0.50	0.10	0.30	0.50	0.70
	Pr(T ⁺)	0.00	0.05	0.15	0.30	0.00	0.02	0.10	0.14
4	Pr(T ⁻ R ⁻)	0.80	0.70	0.40	0.16	0.80	0.60	0.48	0.20
	Pr(T ⁻ R ⁺)	0.20	0.30	0.50	0.70	0.20	0.40	0.50	0.70
	Pr(T ⁺)	0.00	0.00	0.10	0.14	0.00	0.00	0.02	0.10
5	Pr(T ⁻ R ⁻)	0.79	0.25	0.10	0.05	0.80	0.21	0.10	0.10
	Pr(T ⁻ R ⁺)	0.20	0.70	0.40	0.35	0.20	0.75	0.50	0.42
	Pr(T ⁺)	0.01	0.05	0.50	0.60	0.00	0.04	0.40	0.50
6	Pr(T ⁻ R ⁻)	0.10	0.09	0.09	0.08	0.10	0.10	0.10	0.09
	Pr(T ⁻ R ⁺)	0.90	0.90	0.90	0.90	0.90	0.90	0.90	0.90
	Pr(T ⁺)	0.00	0.01	0.01	0.02	0.00	0.00	0.00	0.01

Table 4.

Comparison of the design to find the optimal dose from Ivanova (2003) (IOD) and the REDO using group 1 scenarios from Table 3. REDO with toxicity mitigation is used when outcome is delayed. The true optimal dose is shown in bold.

Scenario	Dose	d ₁	d ₂	d ₃	d ₄
1	IOD	0.03	0.73	0.18	0.06
	REDO	0.03	0.80	0.13	0.04
	REDO, delayed outcome	0.02	0.93	0.04	0.01
2	IOD	0.75	0.17	0.07	0.01
	REDO	0.80	0.15	0.05	0.00
	REDO, delayed outcome	0.89	0.10	0.01	0.00
3	IOD	0.01	0.05	0.79	0.16
	REDO	0.00	0.04	0.83	0.12
	REDO, delayed outcome	0.01	0.11	0.85	0.03
4	IOD	0.00	0.02	0.23	0.74
	REDO	0.00	0.03	0.18	0.80
	REDO, delayed outcome	0.04	0.11	0.51	0.34
5	IOD	0.01	0.86	0.12	0.01
	REDO	0.01	0.92	0.07	0.00
	REDO, delayed outcome	0.01	0.94	0.05	0.00
6	IOD	0.43	0.37	0.15	0.05
	REDO	0.70	0.17	0.10	0.04
	REDO, delayed outcome	0.75	0.21	0.04	0.00

Table 5.

Proportion of trials each dose was recommended as the optimal dose in the two ordered groups in scenarios from Table 3. REDO with toxicity mitigation is used when outcome is delayed. The true optimal dose is shown in bold.

Scenario	Dose	Group 1				Group 2			
		d ₁	d ₂	d ₃	d ₄	d ₁	d ₂	d ₃	d ₄
1	REDO	0.03	0.88	0.08	0.01	0.00	0.14	0.66	0.20
	REDO, delayed outcome	0.03	0.95	0.02	0.00	0.01	0.26	0.67	0.06
2	REDO	0.83	0.14	0.03	0.00	0.05	0.71	0.19	0.06
	REDO, delayed outcome	0.95	0.05	0.00	0.00	0.09	0.81	0.08	0.01
3	REDO	0.00	0.08	0.89	0.04	0.00	0.01	0.18	0.81
	REDO, delayed outcome	0.02	0.19	0.79	0.00	0.00	0.06	0.39	0.55
4	REDO	0.01	0.05	0.50	0.45	0.00	0.02	0.09	0.90
	REDO, delayed outcome	0.06	0.15	0.71	0.08	0.00	0.06	0.27	0.66
5	REDO	0.01	0.92	0.07	0.00	0.00	0.77	0.20	0.03
	REDO, delayed outcome	0.01	0.97	0.02	0.00	0.00	0.88	0.11	0.00
6	REDO	0.69	0.21	0.08	0.03	0.65	0.15	0.10	0.10
	REDO, delayed outcome	0.84	0.15	0.01	0.00	0.66	0.22	0.10	0.03