

Personalized dose selection in radiation therapy using statistical models for toxicity and efficacy with dose and biomarkers as covariates

Matthew J. Schipper,^{a,b,*†} Jeremy M. G. Taylor,^a Randy TenHaken,^b Martha M. Matuzak,^b Feng-Ming Kong^c and Theodore S. Lawrence^b

Selection of dose for cancer patients treated with radiation therapy (RT) must balance the increased efficacy with the increased toxicity associated with higher dose. Historically, a single dose has been selected for a population of patients (e.g., all stage III non-small cell lung cancer). However, the availability of new biologic markers for toxicity and efficacy allows the possibility of selecting a more personalized dose. We consider the use of statistical models for toxicity and efficacy as a function of RT dose and biomarkers to select an optimal dose for an individual patient, defined as the dose that maximizes the probability of efficacy minus the sum of weighted toxicity probabilities. This function can be shown to be equal to the expected value of the utility derived from a particular family of bivariate outcome utility matrices. We show that if dose is linearly related to the probability of toxicity and efficacy, then any marker that only acts additively with dose cannot improve efficacy, without also increasing toxicity. Using a dataset of lung cancer patients treated with RT, we illustrate this approach and compare it to non-marker-based dose selection. Because typical metrics used in evaluating new markers (e.g., area under the ROC curve) do not directly address the ability of a marker to improve efficacy at a fixed probability of toxicity, we utilize a simulation study to assess the effects of marker-based dose selection on toxicity and efficacy outcomes. Copyright © 2014 John Wiley & Sons, Ltd.

Keywords: dose finding; phase I; biomarkers; radiation therapy; utilities

1. Introduction

The dose of radiation therapy (RT) used in clinical practice is typically fixed for a given type of cancer and does not vary across subjects. Recommended dose values are either based on historical practice and retrospective toxicity analysis or a formal phase I trial in which radiation dose to the tumor was escalated until it corresponded to a target probability of toxicity [1]. However, it is recognized that there is variability between patients in their sensitivity to radiation. Various biomarkers of this sensitivity have been proposed. These include biologic markers [2], genetic markers [3–5], and standard clinical factors (e.g., smoking as protective for lung toxicity). Some of these markers are measured at baseline (prior to RT start) while others are measured mid-treatment, that is, measured after RT starts but prior to completion of RT. We have hypothesized that mid-treatment markers that measure change in response to a moderate dose of RT will be most predictive of an individual's sensitivity to RT. Regardless of the type and timing (assuming it is prior to completion of treatment), the availability of these markers makes possible a more personalized dose selection.

^aDepartment of Biostatistics, University of Michigan, Ann Arbor, MI, U.S.A.

^bDepartment of Radiation Oncology, University of Michigan, Ann Arbor, MI, U.S.A.

^cDepartment of Radiation Oncology, Georgia Regents University, Augusta, GA, U.S.A.

*Correspondence to: Matthew J. Schipper, Department of Radiation Oncology, University of Michigan, 1215 Washington Heights, Rm 2531, Ann Arbor, MI 48109, U.S.A.

†E-mail: mjschipp@umich.edu

While dose selection is often the intended use, clinical papers on new markers typically focus on measures of the marker's discriminatory ability, such as area under the curve (AUC), which, as we will show, do not directly address the usefulness of the marker for dose selection purposes. When these papers do discuss how the markers might be used for dose selection, they generally do so only on the basis of toxicity, partly because dose selection in RT has historically been limited to giving the highest dose possible subject to toxicity considerations. There have, however, been a few approaches, which explicitly incorporated efficacy considerations in addition to toxicity. One of the earliest was based on the concept of uncomplicated control defined as local tumor control (at a defined timepoint) and no serious complications [6, 7]. One weakness of such an approach is the implicit assumption that toxicity and tumor progression are equally undesirable, an assumption that is seldom reasonable and avoided in our proposed utility-based approach. Also, in contrast to our proposed approach, the notion of uncomplicated control does not allow incorporation of multiple toxicities of varying type and severity. Recently, Jang *et al.* proposed 'Quality-Adjusted Time Without Symptoms or Toxicity' or Q-TWiST, as a single metric with which to compare treatments based on combined efficacy and toxicity endpoints [8]. Q-TWiST was defined as a weighted linear combination of the time spent in various health states. This is an appealing approach but requires detailed patient data on the duration of toxicity.

In this paper, we focus on how to use expanded toxicity models, in conjunction with efficacy models, to select a dose of RT that is optimal for an individual patient. To simplify calculations, we focus on the estimated survival value at a fixed timepoint as our efficacy endpoint. We also focus on the binary toxicity case but show how multiple binary or ordinal toxicity variables could be utilized. The first approach we consider, called the 'isotoxic' approach, selects a dose that corresponds to an acceptable probability of toxicity for a given patient. The second approach, which we call the 'utility' approach, combines the toxicity model and the efficacy model via a 'utility' function [9] given by the probability of efficacy minus the weighted sum of the probability of various toxicities and selects the dose that maximizes this utility for an individual patient. Using a database of lung cancer patients treated with RT, we build several models for toxicity and efficacy as a function of dose and various markers and use a simulation to illustrate these approaches.

An important consideration when working with RT dose is that RT is distinct from chemotherapy and most other drug settings in that the dose to the tumor is typically different from the dose to the surrounding normal tissue or organ. In a typical treatment plan, a series of radiation beams (typically two to nine), each entering the body at a different location, are targeted at the tumor site. In 3D conformal RT, each beam may be shaped, and the dose of radiation may vary in a linear fashion across the beam. A newer technique, Intensity Modulated RT, splits each beam into a grid of 'beamlets', and the dose of each beamlet can be independently controlled [10]. Although the tumor is the focal point and thus receives the highest dose of radiation, all of the normal tissue in the path of one or more beams also receives some dose of radiation. Taking lung cancer as an example, it is the dose to the tumor (often referred to as the prescription dose) that drives local tumor control, whereas it is the dose to the normal lung tissue (often calculated as mean dose to normal lung tissue) that drives lung toxicity. For a given patient, higher tumor dose generally implies higher mean lung dose. However, because of heterogeneity in tumor size and location, the correlation between tumor dose and mean lung dose, across patients, is often weak.

A RT treatment plan consists of the location and daily (per 'fraction') dose associated with a series of beams along with plan for timing and number of fractions. Specification of a full treatment plan requires choice of number of beams, direction of each beam, and dose for each beamlet within a beam. In this paper, we more narrowly focus on the choice of a single number (tumor dose D or normal tissue dose d) that in turn fully prescribes the treatment plan. In practice, this means that we assume the fractionation scheme, the location of beams, and the relative intensity of each beam are all fixed. This 'plan' can then be scaled up or down, but, importantly, the ratio of D/d is fixed. Throughout the paper, we will use ' r ' to denote the ratio D/d . While r will vary across subjects, its value can easily be calculated for a given patient/plan combination. In the Simulation Study section, we calculate the observed r for each patient, based on the actual treatment plan they were treated with. For prospective use, one could come up with an initial treatment plan to deliver a tumor dose at the high end of possible values and then use the observed r associated with that plan when selecting the final dose (which would just scale down or less likely up the initial plan).

2. Methods

2.1. Notation

Let D_i denote tumor dose, d_i denote dose to normal tissue, r_i the ratio D_i/d_i , and T_i and E_i , respectively, denote binary toxicity and efficacy outcomes with $\pi_{T_i} = P(T_i = 1)$ and $\pi_{E_i} = P(E_i = 1)$, for patient i . We assume availability of models $\pi_{T_i} = f^{-1}(d_i, M_i, \beta)$ and $\pi_{E_i} = g^{-1}(D_i, G_i, \alpha)$, where M and G denote scalars or vectors of marker values, β and α are parameter vectors, and $f()$ and $g()$ are link functions.

2.2. Dose selection strategies

We describe two methods for incorporating toxicity and efficacy markers in dose selection. The first is based on the expanded model for toxicity as a function of M and d and solves $\pi_{T_i} = p^* = f^{-1}(d_i, M_i, \beta)$ for d_i , where p^* is a target probability of toxicity. Note that the assumption of a fixed r , which implies the selection of either D or d , uniquely defines the treatment plan. If no markers for toxicity are available, this approach reduces to selection of the normal tissue dose corresponding to the target rate of toxicity. We refer to this as the ‘isotoxic’ approach because it treats every patient at the same estimated probability of toxicity and ignores the dose-efficacy model.

A second approach, and one which we view as preferable, combines the models for T and E and selects the dose, which maximizes the utility function $U()$:

$$U(d_i, r_i, M_i, G_i) = P(E_i = 1 | D_i = r_i * d_i, G_i) - \theta * P(T_i = 1 | d_i, M_i) \quad (1)$$

Note that D_i can be replaced with $r_i * d_i$ so that Eq. (1) is a function of only a single dose term, d_i . In many instances, there are multiple types of toxicity. For example, in lung cancer, esophageal and heart toxicity are also of concern and are commonly considered in treatment planning. The expression in Eq. (1) is then easily generalized to

$$U(d_i, r_i, M_i, G_i) = P(E_i = 1 | D_i = r_i * d_i, G_i) - \sum_j \theta_j * P(T_{ij} = 1 | d_{ij}, M_i) \quad (2)$$

where normal tissue or organ is indexed by j . We refer to this as the ‘utility’ approach. The last term in Eq. (2) is similar to several recently proposed metrics for phase I trials, including ‘total toxicity burden’ [11]. We note that in both approaches, additional covariates such as smoking status or baseline lung function could also be included in the model for toxicity. In practice, use of this approach requires one to choose a value for θ . This value could be elicited from clinicians, or it could be selected by individual patients. All else being equal, as θ increases, Eq. (1) will be maximized by smaller d . Thus, θ could also be viewed as a tuning parameter, with its value chosen to result in an acceptable marginal rate of toxicity for a population of patients, calculated as

$$P(T = 1) = 1/n \sum_i P(T_i = 1 | \hat{d}_i(\theta, r_i, M_i))$$

where \hat{d}_i maximizes Eq. (1) and is thus a function of θ , r_i , and M_i . We note that maximizing $\sum_i U_i(d_i, r_i, M_i)$ over d_i and choosing the value of θ such that $P(T = 1) = p^*$ is analogous to a Lagrange multiplier approach. The utility function approach is more general than the isotoxic approach, because the isotoxic approach forces $P(T_i = 1) = p^*$ for all i , whereas the utility function method only forces the average of $P(T_i = 1)$ to equal p^* .

In addition to controlling the average rate of toxicity in the utility function approach, it may be desirable to also limit an individual patient’s probability of toxicity. This can be accomplished by incorporating an additional individual level constraint in the maximization of $U(d_i)$. For example, $d_i < d_{max}$ or $P(T_i = 1) < c$. The latter approach has been used in dose finding designs to determine the set of ‘acceptable’ doses from which an optimal dose is then selected on the basis of toxicity and efficacy considerations [12].

Efficacy and the efficacy/toxicity trade-off are increasingly being considered in dose finding trial designs [13]. We describe two specific examples here and note the similarity to the utility approach described earlier. The first is a proposed dose finding trial design based on bivariate toxicity and efficacy outcomes [14]. Briefly, their method assumes ordinal toxicity and efficacy outcomes and the utility (or desirability) associated with each possible bivariate outcome, which is elicited from the clinician. An example for binary toxicity and efficacy is given in the matrix in the succeeding text.

	E=0	E=1
T=1	U_{01}	U_{11}
T=0	U_{00}	U_{10}

Given models for T and E as a function of dose, one can calculate the expected utility for any dose value by a weighted sum as

$$E(U|d, r, M, G) = \sum_{j,k} U_{jk} * P(E = j, T = k|d, r, M, G)$$

If we assume that, conditional on dose and marker values, T and E are independent, a little algebra shows that any utility matrix of the following form has expected value equal to Eq. (1).

	E=0	E=1
T=1	$-\theta$	$1-\theta$
T=0	0	1

The authors in [14] fixed $U_1 = 100$ and $U_{01} = 0$, with the interior values set by clinicians. If both E and T are binary, this allows choice of two values. Use of Eq. (1) reduces this to choice of a single parameter by forcing $U_{10} - U_{00} = U_{11} - U_{01}$, that is, the increase in utility or benefit associated with positive efficacy is the same regardless of whether toxicity is experienced or not.

A second example of dose finding based on both efficacy and toxicity is given in [15], who define the desirability of a dose in terms of $\pi_E(P(E = 1))$ and $\pi_T(P(T = 1))$ as $1 - x$, where x is given by

$$\left(\frac{1 - \pi_E}{1 - \pi_E^*}\right)^p + \left(\frac{\pi_T}{\pi_T^*}\right)^p = x^p \tag{3}$$

where π_E^* and π_T^* are, respectively, elicited from the clinician as the smallest efficacy rate they would accept if the treatment was associated with no toxicity and the highest toxicity rate they would accept if the treatment had 100% efficacy. When $p = 1$, it can be shown that

$$c * (1 - x) = \pi_E - \theta * \pi_T$$

where c is a constant not depending on dose and $\theta = \frac{\pi_E^* - 1}{\pi_T^*}$. Thus, in the case where $p = 1$, selecting dose to maximize $(1 - x)$ from Eq. (3) is equivalent to maximizing Eq. (1).

2.3. Evaluation of potential dose selection markers

Reports on new biomarkers for toxicity or efficacy typically report odds ratios and measures of discriminatory ability such as AUC, with associated p -values from tests of no discriminatory ability or no increase in AUC relative to a dose alone model. These metrics are useful and relevant if the new marker is to be used in a purely predictive fashion (e.g., counsel patients regarding risk of toxicity). Often, however, the goal is to use the new marker to select a dose in such a way that overall outcomes are improved. More precisely, there is an expectation that marker-based dose selection will result in either (i) increased π_E at the same π_T or (ii) decreased π_T at the same π_E . However, neither of these events may occur, even when using a marker that substantially increases AUC over the AUC for dose alone. As an illustration, consider the hypothetical lung cancer scenario depicted in Figure 1. Suppose a biomarker is discovered, which enables us to identify two groups of patients: one with a uniformly higher and the other with a lower risk of toxicity. So rather than predicting toxicity using only mean lung dose (d , solid line), we can now more accurately predict toxicity using the new biomarker (dashed lines). For simplicity of visual presentation, assume that there are 100 patients to be treated, all with the same tumor size and location (so that the ratio d/D is the same for all patients). Using the dose only toxicity model, we could treat all patients at a tumor dose of 70 Gy (equivalent to $d = 15$ Gy), and the toxicity and efficacy rates would be 25% and 60%. With the new biomarker, we could instead treat each patient at their ‘personalized’ dose

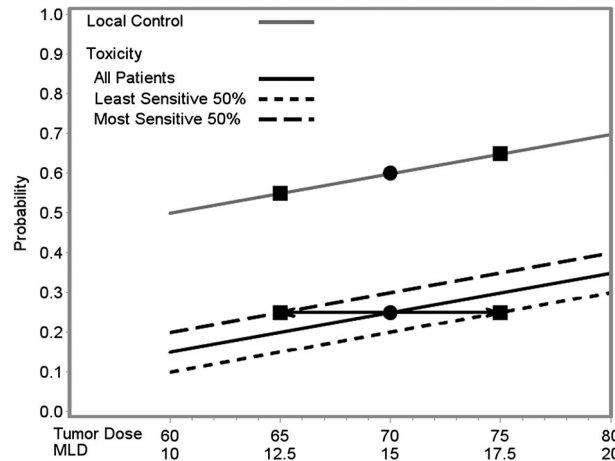


Figure 1. Illustration of result that when probabilities of toxicity and efficacy (local control) are linearly related to dose, a prognostic marker for toxicity cannot be used to improve efficacy without also increasing toxicity.

corresponding to a toxicity probability of 25% (65 or 75 Gy for more or less sensitive patients). With this approach, the overall toxicity rate is still 25%, but the overall efficacy rate is now the average of 55% (for patients treated at 65 Gy) and 65% (for patients treated at 75 Gy), equal to 60%, which is the same as that achieved by ignoring the biomarker. To see that this result also holds in the more realistic setting where r varies across patients, suppose

$$P(E_i = 1 | D_i = r_i d_i) = \alpha_0 + \alpha_1 r_i d_i$$

and

$$P[T_i = 1 | d_i, M_i] = \beta_{01} I(M_i = 1) + \beta_{02} I(M_i = 2) + \beta_1 d_i$$

where M denotes marker group 1 or 2, which are of the same size so that $P(M_i = 1) = P(M_i = 2) = 0.5$ (although this is not necessary for the result to hold). Then, the probability of toxicity conditional only on dose is given by

$$P(d_i | T_i = 1) = \beta_0 + \beta_1 d_i$$

where $\beta_0 = .5 * (\beta_{01} + \beta_{02})$

We can calculate the marginal (across the population of patients) probabilities of T and E following two different dose selection strategies.

Strategy 1: Treat all patients at $d = (p^* - \beta_0) / \beta_1$

Strategy 2: If $M_i = 1$, treat at $d = (p^* - \beta_{01}) / \beta_1$, else if $M_i = 2$, treat at $d_i = (p^* - \beta_{02}) / \beta_1$

It is readily seen that under both strategies, the expected population toxicity rate = p^* . Under strategy 1, d_i is fixed and constant across patients, but r_i varies between patients. Thus, the average probability of efficacy is given by integrating the subject specific probabilities over the distribution of r .

$$\begin{aligned} P(E = 1) &= 1/n \sum_n P\left(E_i = 1 \mid d_i = \frac{p^* - \beta_0}{\beta_1}, r_i\right) \\ &= \alpha_0 + \alpha_1 \frac{p^* - \beta_0}{\beta_1} \bar{r} \end{aligned}$$

Under strategy 2, d varies across subjects according to M . Similar algebra to the preceding text yields

$$\begin{aligned} P(E = 1) &= P(M = 1) * \left(\alpha_0 + \alpha_1 \frac{p^* - \beta_{01}}{\beta_1}\right) * \bar{r} + P(M = 2) * \left(\alpha_0 + \alpha_1 \frac{p^* - \beta_{02}}{\beta_1}\right) * \bar{r} \\ &= \alpha_0 + \alpha_1 \frac{p^* - \beta_0}{\beta_1} \bar{r} \end{aligned}$$

Thus, despite the use of a marker that allows better prediction of toxicity, strategy 2 results in the same marginal probability of efficacy as strategy 1, which is based only on a dose. This result depends on the linear dose-outcome relationships and the additive effect of the marker. Another crucial assumption is that there is no interaction between dose and marker. This result would not hold in the presence of an interaction or when the true relation between the probability of toxicity and dose is not linear.

3. Simulation study

The aforementioned work was motivated by our work with a database of outcomes and biomarkers for lung cancer patients treated with RT. Using these data, we fit models for efficacy and toxicity as a function of dose with or without various biomarkers. To illustrate the issues associated with marker-based dose selection in RT, we conducted a simulation study.

3.1. Toxicity models

The primary dose limiting toxicity for these patients is lung toxicity (pneumonitis grade 2 or higher), which we model as a function of mean lung dose (d) and cytokine levels measured from blood draws pre and during RT. Details of five logistic regression toxicity models that fit to these data are given in Table I. Maximum likelihood was used to estimate model parameters. We use the level of the cytokine IL8 measured at baseline and the cytokines $TGF\beta$ and $TGF\alpha$, each measured mid-treatment. Higher baseline levels of IL8 appear to be protective, whereas increased values of the other cytokines mid-treatment appear to increase the risk of toxicity. For illustration purposes, we utilize $TGF\beta$ as a continuous predictor and $TGF\alpha$ as a binary predictor with values classified as above or below the median. Figure 2 shows the

Model	Toxicity equation (on logit scale)	AUC	Local progression free survival at 2 years (%)		
			Isotoxic	Utility without efficacy markers	Utility with efficacy markers
1	$-5.16 + 0.2 * d$	0.72	40.0	44.1	48.4
2	$-3.66 + 0.22 * d - 0.85 * \log(IL8)$	0.78	42.2	45.8	50.2
3	$-3.72 + 0.19 * d - 0.79 * \log(IL8) + 0.03 * TGF\beta * d$	0.83	44.8	47.0	51.0
4	$-5.34 + 0.17 * d + 0.07 * d * I(TGF\alpha > 3)$	0.76	44.2	45.1	52.1
5*	$-5.6 + 0.12 * d + 0.16 * d * I(TGF\alpha > 3)$	NA	54.0	54.5	57.6

Toxicity rate is equal to 15% in every case.

*Hypothetical model.

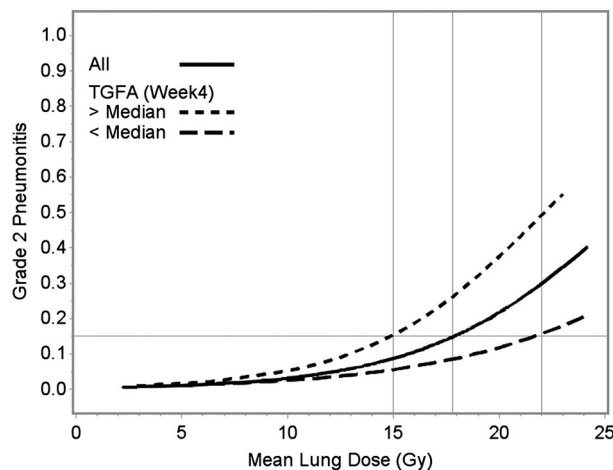


Figure 2. Estimated probability of lung toxicity for all patients and by mid-treatment transforming growth factor alpha (TGFA) levels.

fitted models 1 and 4. Models 3 and 4 have dose * marker interactions with no main effect for marker. While we have also fit models using these markers as main effects, here we present only the marker * dose interaction models for illustration purposes. We do not address the very important issues of multiplicity or model validation. Finally, we note that model 5 is hypothetical and was included merely to illustrate the effect of a stronger dose * marker interaction.

3.2. Efficacy model

Because RT is a local and not systemic treatment, we utilize local progression free survival (LPFS), defined as the minimum time to either local progression of the treated tumor or death, as our measure of efficacy. To estimate the effect of biomarkers and tumor dose on LPFS, we used Cox proportional hazard models because many patients are censored during the first 2 years. Our goal was to use biomarkers to identify patients who would (or would not) benefit from higher RT dose. We used a stepwise procedure in which dose and dose * biomarker interactions were included as potential covariates. For biomarkers, we used expression values of 84 miRNAs, which are detectable in serum. In addition to dose, 23 dose * miRNA interactions were jointly significant at the 0.10 level. From the fitted Cox model, we calculated the dose coefficient for patient i as $\alpha_0 + \sum_j \alpha_j * X_{ij}$, where α and α_j denote the parameter coefficients for dose and the interaction between dose and the j^{th} miRNA term and X_{ij} denotes the value of the j^{th} miRNA for the i^{th} patient. Although this biomarker could be used in a continuous fashion, we chose to group patients into two groups for illustration purposes in the simulation study. Patients whose dose coefficient was in the lowest 1/3 had significantly worse LPFS and no dose effect compared with patients with higher dose coefficient (Figure 3), so we used these two groups. LPFS at 2 years is a commonly utilized endpoint in RT lung trials, and we use it here as our binary efficacy endpoint and estimate it from the fitted Cox model as

$$\hat{S}(2 | D, G) = \hat{S}_0(2)^{\exp(\alpha_1 D * I(G=1) + \alpha_2 D * I(G=2))}$$

where $\hat{S}_0(2)$ is the Nelson–Aalen estimate of baseline survival at 2 years. Dose had a highly significant impact on LPFS in group 1 (estimated hazard ratio = 0.90 per 1 Gy increase in tumor dose ($p < 0.01$) but not in group 2 (estimated hazard ratio = 0.99, $p = 0.78$).

3.3. Simulation methods

The simulation (or virtual clinical trial) followed these steps for each of the five toxicity models and each patient (total of 82) in the database.

- (1) If using Isotoxic approach: Calculate the value d corresponding to 15% probability of toxicity. If $d_i > 25$, then set $d_i = 25$.




Figure 3. Kaplan–Meier estimates of local progression free survival for groups of patients defined by miRNA group and dose.

OR

If using Utility approach: For given θ , calculate d_i , which maximizes the utility given by Eq. (1), using either the dose only efficacy model or the dose and miRNA efficacy model, and adjust θ so that the population toxicity rate is 0.15

- (2) Calculate the corresponding tumor dose $D_i = d_i * r_i$. If $D_i > 90$, then reset D_i to 90 and calculate the corresponding d_i
- (3) Calculate $P(T_i=1 | d_i, M_i)$
- (4) Calculate $P(E_i=1 | D_i, G_i)$

After doing this for all patients, we then calculated the average of $P(T_i = 1)$ and $P(E_i = 1)$ across patients for each model and dose selection method.

3.4. Simulation results

Figure 4 shows utility-dose curves for example patients based on model 4. A comparison of the gray and black lines shows how increasing θ results in selecting a smaller dose. Similarly, the optimal dose is higher for patients with shallow dose-toxicity curves ($TGF\alpha < \text{median}$) than for patients whose toxicity increases more rapidly with dose.

The toxicity models that use the cytokine values in a continuous fashion sometimes ‘recommend’ very high or very low dose values (e.g., $d = 30$ Gy). The rationale for the dose constraints in the simulation is to avoid doses that clinicians would be uncomfortable prescribing, as well as to avoid extrapolating the model predictions to dose values not seen in the dataset used to fit the models. To facilitate comparisons of the different strategies/models and to avoid outcomes where one strategy had better efficacy but worse toxicity than another strategy, we targeted the same level of toxicity with each so that they can be directly compared with respect to efficacy.

Simulation results are given in Table I. Relative to the dose only model 1, use of models 2–4 resulted in improved efficacy without increased toxicity. Although there is no statistical interaction of dose and IL8 in model 2, the increase in π_T (not $\text{logit}(\pi_T)$) varies with IL8. Thus, despite the earlier finding regarding additive markers, model 2 results in improved efficacy. Intuitively, the reason why model 4 allows improvement over model 1 is that the required reduction in d in the sensitive patients is just 2 Gy, while the possible increase in d in the tolerant patients is 4 Gy (Figure 2). A general conclusion from the results is that the ability to improve efficacy at a fixed toxicity is not closely correlated with AUC. For example, the AUC of model 4 is just 0.76, considerably lower than the 0.83 for model 3. Yet both result in approximately the same efficacy improvement because of the stronger marker * dose interaction in the former.

Not surprisingly, for each of the models considered, the utility approach (even without using the efficacy markers) outperforms the isotoxic approach. By allowing the probability of toxicity to vary across subjects, the method ‘spends its toxicity’ in a manner that results in the largest possible efficacy. By restricting every subject to be treated at a fixed π_T , the isotoxic method exposes some patients to a dose that, relative to a lower dose, may increase π_T with only a small corresponding increase in π_E , whereas




Figure 4. Utility values for example patients when $r = (D/d) = 6$

other subjects that might experience relatively large increases in π_E if π_T were allowed to be higher, have their dose limited. Conversely, in the utility-based approach, the only subjects exposed to higher π_T are those with a larger than average increase in π_E . In addition, we have shown that the utility function approach can target any desired level of toxicity for a population of patients through selection of an appropriate value of θ . Thus, prior to obtaining marker measurements, each patient's expected probability of toxicity and efficacy is equal to the population mean values (e.g., those given in Table I), which are higher for efficacy and the same for toxicity. These are the relevant values to consider for an individual patient deciding whether to pursue a dose selection scheme based on biomarkers. Incorporating the efficacy marker allowed further gains in outcome as seen in Table I. Use of toxicity model 4 with the efficacy marker allowed an increase in LPFS from 40% to 52% without increasing toxicity.

3.5. Correlation between toxicity and efficacy markers

Table II gives the relation between miRNA group and TGF α group for patients in our data. The two variables are not highly correlated (odds ratio = 1.29, Fisher's exact test $p = 0.64$).

Intuitively, a high degree of positive correlation would limit the potential gains of the marker-based dose selection strategies discussed in this paper. The reason is that the same patients whose LPFS would improve the most with high dose would also experience the largest increase in toxicity. Similarly, negative correlation would seem to allow the largest improvements. Patients with high efficacy sensitivity but low toxicity sensitivity would be natural candidates for high dose, whereas patients with low efficacy sensitivity but high toxicity sensitivity would naturally receive lower dose. To investigate this empirically, we repeated the simulation study after changing miRNA group status for some patients. To induce positive correlation, we took the 10 patients who were TGF α sensitive but miRNA not sensitive and changed their miRNA group to sensitive. To keep the four marginal totals constant, we also took 10 of the 31 patients who were TGF α not sensitive but miRNA sensitive and changed their miRNA group to not sensitive. We repeated this in the opposite direction to induce negative correlation. To remove any effect of $r(D/d)$, which varies between patients, we also set $r = 5$ for all patients. The results in Table III confirmed our intuition with the negative correlation case allowing the greatest local control improvements.

4. Discussion

A natural test of the ability of these dose selection methods to improve outcomes would be a clinical trial in which patients were randomized between marker-based and standard (fixed) dose selection. At the same time, clinical use of these methods would ideally be based on validated models linking the probabilities of toxicity and efficacy with dose and marker values. To build such models requires data on patients treated over a range of dose values. If one were confident in the form of the models but less confident regarding the actual parameter values, a Bayesian approach could be taken in which the parameter values were updated throughout the trial based on accrued data. If the marker(s) were used to

Table II. Frequency table of patient's toxicity marker (TGF α) group and efficacy marker (miRNA signature) group.

	TGF α < 3 (not sensitive)	TGF α > 3 (sensitive)	
miRNA group = 1 (not sensitive)	16	10	26
miRNA group = 2 (sensitive)	31	25	56
	47	35	

Table III. Effect of correlation between efficacy and toxicity markers on ability to improve local control.

Correlation between normal tissue and tumor dose sensitivity	Odds ratio	Local progression free survival at 2 years		
		Isotoxic	Utility without efficacy markers	Utility With efficacy markers
Observed in data	1.29	53	57	59
Positive	Infinite	50	54	58
Negative	0	61	62	62

identify a small number of groups of patients with distinct dose-toxicity or dose-efficacy relationships, a more standard dose finding trial could be performed seeking to identify, for each marker defined group, the dose that maximized the utility.

We have limited our attention to situations in which a patient's initial dose is selected according to the model and then possibly 'scaled' up or down mid-treatment. We have not considered the situation in which patients are re-planned in such a way that the spatial distribution of dose is changed part way through treatment. However, we expect that methods to use mid-treatment imaging to redistribute dose to less sensitive normal tissue, or to more aggressive tumor regions, could be utilized 'on top' of the proposed methods to allow greater improvements in outcome. Study of this is more difficult as it requires the ability to generate new treatment plans, which requires patient images, dose maps and as a result is much more computationally intensive. Nevertheless, such an approach seems likely to allow additional gains over the scaling approach.

Acknowledgements

This research was supported by the grant PO1 CA59827. Authors wish to thank reviewers for helpful comments and suggestions.

References

1. Emami B, Perez C, Herskovich A, Phase Hederman M. Phase I/II study of treatment of locally advanced (T3T4) non-oat cell lung cancer with high dose radiotherapy (rapid fractionation): Radiation therapy oncology group study. *International Journal of Radiation Oncology Biology Physics* 1988; **15**:1021–1025.
2. Stenmark MH, Cai XW, Shedden K, Hayman JA, Yuan S, Ritter T, Ten Haken RK, Lawrence TS, Kong FM. Combining physical and biologic parameters to predict radiation-induced lung toxicity in patients with non-small-cell lung cancer treated with definitive radiation therapy. *International Journal of Radiation Oncology Biology Physics* 2012; **84**(2): e217–e222.
3. Yuan ST, Ellingrod VL, Schipper M, Stringer KA, Cai X, Hayman JA, Yu J, Lawrence TS, Kong FM. Genetic variations in TGFbeta1, tPA, and ACE and radiation-induced thoracic toxicities in patients with non-small-cell lung cancer. *Journal of Thoracic Oncology* 2013; **8**(2):208–213.
4. Wang W, Shedden K, Ten Haken R, Hayman J, Arenberg D, Lawrence T, Kong FS. Single nucleotide polymorphisms in DNA repair genes may be associated with survival in patients with non-small cell lung cancer treated with definitive radiation therapy. *International Journal of Radiation Oncology Biology Physics* 2012; **84**(3S):S69.
5. Yuan X, Liao Z, Liu Z, Wang LE, Tucker SL, Mao L, Wang XS, Martel M, Komaki R, Cox JD, Milas L, Wei Q. Single nucleotide polymorphism at rs1982073:T869C of the TGF{beta}1 gene is associated with the risk of radiation pneumonitis in patients with non-small-cell lung cancer treated with definitive radiotherapy. *Journal of Clinical Oncology* 2009; **27**(20):3370–3378.
6. Ågren A, Brahme A, Turesson T. Optimization of uncomplicated control for head and neck tumors. *International Journal of Radiation Oncology Biology Physics* 1990; **19**(4):1077–1085.
7. Emami B. Three-dimensional conformal radiation therapy in bronchogenic carcinoma. *Seminars in Radiation Oncology* 1996; **6**(2):92–97.
8. Jang RW, Le Maître A, Ding K, Winton T, Bezjak A, Seymour L, Shepherd FA, Leighl NB. Quality-adjusted time without symptoms or toxicity analysis of adjuvant chemotherapy in non-small-cell lung cancer: an analysis of the National Cancer Institute of Canada Clinical Trials Group JBR.10 trial. *Journal of Clinical Oncology* 2009; **27**(26):4268–4273.
9. Berger JO. *Statistical Decision Theory and Bayesian Analysis*. Springer-Verlag: New York, 1985.
10. Kessler M, McShan D, Epelman M, Vineberg K, Eisbruch A, Lawrence T, Costlets Fraass B. A generalized approach to cost functions for automated optimization of IMRT treatment plans. *Optimization and Engineering* 2005; **6**:421–448.
11. Bekele NB, Thall FP. Dose-finding based on multiple toxicities in a soft tissue Sarcoma trial. *Journal of the American Statistical Association* 2004; **99**:26–35.
12. Cai C, Yuan Y, Ji Y. A bayesian phase I/II design for oncology clinical trials of combining biological agents. *Journal of the Royal Statistical Society: Series C* 2014; **63**:159–173.
13. Zohar S, Chevret S. Recent developments in adaptive designs for phase I/II dose finding studies. *Journal of Biopharmaceutical Statistics* 2007; **17**:1071–1083.
14. Thall PF, Nguyen HQ. Adaptive randomization to improve utility-based dose-finding with bivariate ordinal outcomes. *Journal of Biopharmaceutical Statistics* 2012; **22**:785–801.
15. Berry S, Carlin BP, Lee JJ, Muller P. *Bayesian Adaptive Methods for Clinical Trials*. CRC Press: Boca Raton, 2001.