Int. J. Radiation Oncology Biol. Phys., Vol. 68, No. 4, pp. 1178–1189, 2007

0360-3016/07/$–see front matter

ELSEVIER

doi:10.1016/j.ijrobp.2007.02.051

PHYSICS CONTRIBUTION

A MULTIPLAN TREATMENT-PLANNING FRAMEWORK: A PARADIGM SHIFT FOR INTENSITY-MODULATED RADIOTHERAPY

ROBERT R. MEYER, PH.D.,* HAO H. ZHANG, M.S.,† LAURA GOADRICH, M.S.,*
DARYL P. NAZARETH, PH.D.,‡ LEYUAN SHI, PH.D.,† AND WARREN D. D'SOUZA, PH.D.‡

*Computer Sciences Department and †Industrial and Systems Engineering Department, University of Wisconsin, Madison, WI; and ‡Department of Radiation Oncology, University of Maryland School of Medicine, Baltimore, MD

Purpose: To describe a multiplan intensity-modulated radiotherapy (IMRT) planning framework, and to describe a decision support system (DSS) for ranking multiple plans and modeling the planning surface.

Methods and Materials: One hundred twenty-five plans were generated sequentially for a head-and-neck case and a pelvic case by varying the dose–volume constraints on each of the organs at risk (OARs). A DSS was used to rank plans according to dose–volume histogram (DVH) values, as well as equivalent uniform dose (EUD) values. Two methods for ranking treatment plans were evaluated: composite criteria and pre-emptive selection. The planning surface determined by the results was modeled using quadratic functions.

Results: The DSS provided an easy-to-use interface for the comparison of multiple plan features. Plan ranking resulted in the identification of one to three "optimal" plans. The planning surface models had good predictive capability with respect to both DVH values and EUD values and generally, errors of <6%. Models generated by minimizing the maximum relative error had significantly lower relative errors than models obtained by minimizing the sum of squared errors. Using the quadratic model, plan properties for one OAR were determined as a function of the other OAR constraint settings. The modeled plan surface can then be used to understand the interdependence of competing planning objectives.

Conclusion: The DSS can be used to aid the planner in the selection of the most desirable plan. The collection of quadratic models constructed from the plan data to predict DVH and EUD values generally showed excellent agreement with the actual plan values. © 2007 Elsevier Inc.

Multiplan framework, IMRT planning, Composite criteria, Pre-emptive selection, Quadratic modeling.

INTRODUCTION

Radiation treatment planning requires consideration of two competing objectives: maximizing the amount of radiation delivered to the planning target volume (PTV) and minimizing the amount of radiation delivered to all other tissues. It is difficult to avoid depositing some fraction of the dose to organs at risk (OARs), so planners must determine that a treatment plan will not only deliver an appropriate dose to the tumor, but also that the dose delivered to critical and normal structures is at an acceptably low level. The tradeoff between the above factors leads to consideration of multi-criteria objective techniques.

The utility of multi-criteria approaches with respect to radiation treatment planning has been the subject of previous studies. Recent articles in this area include those by Gopal and Starkschall (1), which presents various two-dimensional plan comparison visualizations; Rosen *et al*. (2), which deals with conformal radiotherapy and emphasizes graphically guided adjustments to base plans; Zhang *et al*. (3), which focuses on sensitivity analysis issues and, as in the earlier work by Yu (4) and by Xing *et al*. (5), develops mechanisms for multiparametric adjustments that seek to construct a sequence of treatment plans that converge to a prespecified set of goals; Romeijn *et al*. (6), which shows the mathematical equivalence of Pareto surfaces based on alternative metrics; and Craft *et al*. (7), which discusses invariance properties of Pareto surfaces. These earlier studies have either been primarily theoretically oriented or have tried to limit parameters in such a way that two-dimensional graphs would be the basis for plan comparison. Consequently, issues of clinically acceptable computing times and user-friendly interfaces, as addressed in this research, that allow plan-wise comparisons based on multiple properties have not been considered.

Reprint requests to: Warren D. D'Souza, Ph.D., Department of Radiation Oncology, University of Maryland School of Medicine, 22 South Greene St., Baltimore, MD 21201. Tel: (410) 328-7159; Fax: (410) 328-2618; E-mail: wdsou001@umaryland.edu

This work was supported in part by National Science Foundation grants DMI-0355567 and DMI-0400294.

Conflict of interest: none.

Received Oct 30, 2006, and in revised form Feb 19, 2007. Accepted for publication Feb 24, 2007.

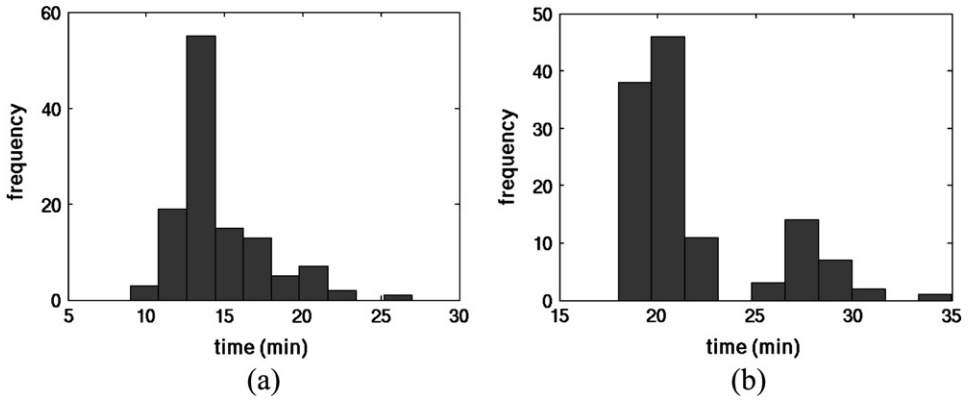

Fig. 1. Planning-time distribution for the multi-plan framework in the (a) head-and-neck case, and (b) pelvis case.

Our goal in this work was to lay the foundation for a paradigm shift for treatment planning. Instead of the current inefficient and time-consuming process of using numerous *plan-evaluate-modify* cycles in a trial-and-error process, we propose the initial generation of multiple plans for a given case. This *multi-plan* approach will produce large amounts of data that will be presented to the clinician through a user-friendly interface that will facilitate the selection of the most appropriate plan from the many alternatives. Even within commercial planning systems there exists a need for decision support tools to evaluate rival treatment plans (*e.g.*, two plans may be evaluated simultaneously in Pinnacle[3] [Philips Medical Systems, OH], and four plans may be ranked by decreasing PTV coverage or OAR protection in the BrainSCAN [BrainLAB, Feldkirchen, Germany] planning system), and currently the available tools for plan comparison are limited at best. Further, the interrelationships between various OARs during plan optimization are not known *a priori*. Often one can improve the dose distribution to one OAR but at the expense of another. In summary, it is difficult to predict the effect of changes in dose–volume constraints corresponding with one OAR on the dose–volume properties of another OAR.

In this work, we address these issues with the use of graphically oriented plan comparators. In addition, we describe a new approach to modeling the planning surface that will assist in understanding the relationships between the OARs and, if necessary, the exploration of treatment alternatives if no acceptable plan is generated within the initial set of plans. In this study, we focus on intensity-modulated radiation therapy (IMRT) of two test cases: (*1*) head-and-neck, and (*2*) pelvis.

METHODS AND MATERIALS

Multi-plan framework

We have developed a framework in which multiple plans can be generated serially for a given case without planner intervention (between plans) using a commercial planning system (Pinnacle[3], version 7.4f), scripting tools available in the planning system, and a user-developed code that generates an automated planning script. For a given case, the user may specify the target structure and the set of OARs, the set of beam angles, the prescription dose, and the set of constraints. For each case, 125 plans (seven-field) were generated by varying the dose–volume constraints on the OARs. Plan optimization was performed using the algorithm in the Pinnacle[3] system and on a SunBlade 2000 workstation with two 900-MHz processors. The distribution of planning times is shown in Fig. 1.

Test cases

To illustrate the multi-plan paradigm, we consider two complex clinical cases: (*1*) head-and-neck, and (*2*) pelvis involving concave-shaped PTVs with three OARs. Figure 2 shows axial images

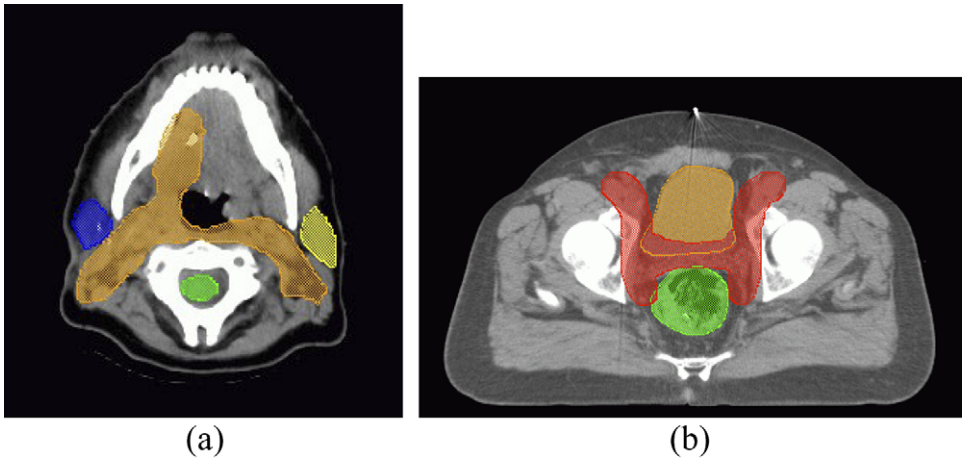

Fig. 2. Axial slices of the (a) head-and-neck case showing the planning target volume (PTV), parotids, and spinal cord, and (b) pelvic case showing the PTV, bladder, and rectum (the small bowel was the third organ at risk).

of both cases. We experimented with five different dose–volume histogram (DVH) constraints for each OAR, thereby resulting in 125 plans. The rationale for the variation in critical structure constraints was to obtain a range of dose–volume values for one critical structure as a function of another.

Head-and-neck case. The head-and-neck case consisted of a 169.2-cm^3 tumor. The PTV was generated by adding a 1-cm margin to the clinical target volume. The target dose specification was as follows: 90% or more of the PTV was required to receive at least 50.4 Gy, and no more than 15% of the PTV was allowed to receive doses in excess of 55.4 Gy. The OARs were the left and right parotids and the spinal cord. The constraints on the parotids were such that ≤33% of the organs were allowed to receive more than 20–32 Gy in 3-Gy increments. The maximum dose constraint on the cord was varied from 33 to 45 Gy in 3-Gy increments.

Pelvic case. The pelvic case consisted of a 733.0-cm^3 target that involves the prostate and nodal volumes. The target dose specification was as follows: 95% or more of the PTV was required to receive at least 45 Gy, and no more than 10% of the PTV was allowed to receive doses in excess of 49.5 Gy. The OARs were the small bowel, rectum, bladder, and femoral heads. For the small bowel, ≤25% was allowed to exceed doses ranging from 19 to 31 Gy in 3-Gy increments. For the rectum and bladder, ≤35% was allowed to exceed doses ranging from 24 to 36 Gy in 3-Gy increments.

Graphic user interface

To be clinically useful, the multi-plan paradigm requires a decision support system (DSS) infrastructure to evaluate and rank treatment plans. We have developed graphic user interface (GUI) decision-support tools to support the multi-plan framework in selecting a plan from alternative plans. The prototype GUI has been implemented in MATLAB (Mathworks, Natick, MA).

The GUI design allows the user to input up to 15 different plans for comparison in groups of 5 at a time. Once the user has selected the plans to compare, the GUI allows selection of structures for plan comparison and the target percentage of organ volume to be used for comparison. Dose values are displayed for the selected percentage of OAR volume for each plan. The user then has the option to sort the 5 plans in order of best to worst. This sort can be done with respect to each individual selected property using the ''Full-Sort'' option (or structure-wise plan ranking [SPR]), which allows a different ordering of plans for each (dose or dose–volume) property. Alternatively, a particular plan ordering resulting from the sort with respect to one property can be carried over to the other properties so that all of the properties for a given plan will occur in the same row of the GUI. This option is referred to as the ''Single-Column-Sort'' option (or plan-wise single-criterion ranking [PSR]).

Note that the two sorting options enable the spreadsheet format to present to the user data corresponding to horizontal line ''slices'' through two different graphic representations of the set of plans. The Full-Sort (SPR) option emphasizes (with its columns) the rankings of the plans with respect to individual organs at specified volumes, and, therefore, represents a horizontal line slice through a representation such as Fig. 3. The Single-Column-Sort (PSR) option emphasizes (with its rows) the selected set of organ dose properties for each of the individual plans (and, therefore, presents a user-selected subset of information from a figure such as Fig. 4); and, in addition, adds plan ranking information based on treatment quality with respect to a selected organ.

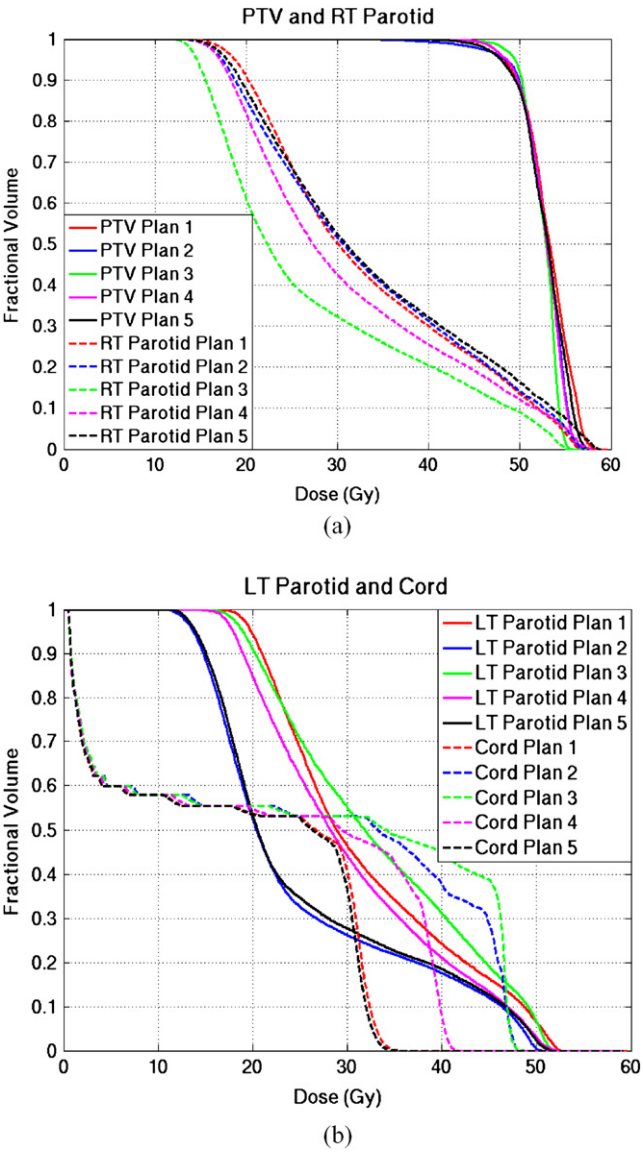

Fig. 3. Dose–volume histograms (DVHs) for the five head-and-neck plans described in Tables 2 and 3. The DVHs are grouped by structures: (a) planning target volume (PTV) and right parotid, and (b) left parotid and cord. All plans were normalized so that ≥90% of the PTV received at least 50.4 Gy.

The GUI also includes the capability of using equivalent uniform dose (EUD) values for plan comparison. Equivalent uniform dose was calculated by:

$$EUD = \left(\frac{1}{N}\sum_{i=1}^{N} D_i^a\right)^{1/a},$$

where N is the total number of voxels corresponding to a given structure, D_i is the dose to the i^{th} voxel, and a is a structure-specific parameter that describes the dose–volume effect. Table 1 summarizes the values for a that were used in the EUD calculations for the head-and-neck and pelvic cases. We used values of a that have been previously published by Wu *et al.* (8). The value of a for the small bowel was obtained using previously derived parameters (9) obtained from Emami *et al.* (10). The GUI also includes a feature that allows the user to display DVHs for selected plans simultaneously.

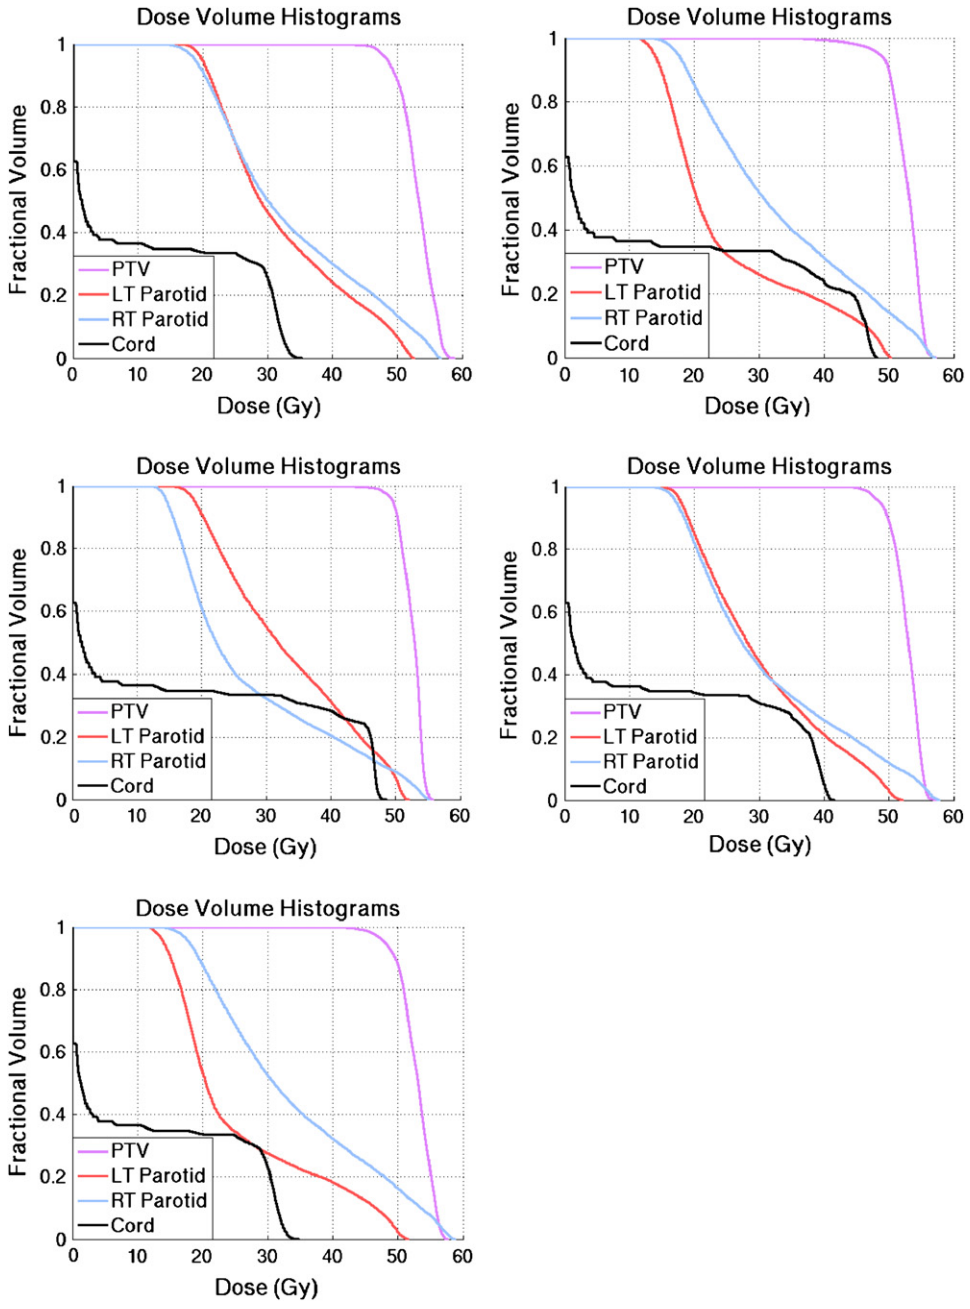

Fig. 4. Competing objectives in intensity-modulated radiotherapy planning as seen in dose–volume histograms of the planning target volume (PTV) and organs at risk (OARs) for five sample head-and-neck plans with differing input constraints. Note that no plan outperforms the other when one considers all the OARs.

Plan ranking scheme

When multiple plans are generated, it may often be the case that no one single plan outperforms all others when all OARs are considered. We describe two approaches to aid in the selection of the treatment plan: (*1*) composite criteria, and (*2*) pre-emptive selection.

Composite criteria. The plans presented in the GUI (described above) can be ranked using a composite multi-structure criterion C of the form $C = \sum \alpha_i f_i\,(p_i - t_i)$, where the α_i are a set of user-specified weights, f_i are linear, piecewise-linear, or quadratic, p_i are plan properties, and the t_i are user-specified goals for those properties. The lower the value of C, the more preferable the plan. C represents a distance metric that measures the distance from a given plan to an ideal plan (as characterized by the set of goals t_i). It should be noted that distance measures in higher dimensions have been the subject of research (Aggarwal and Charu [11]), because they can be problematic in terms of nearest neighbor selection and are also subject to individual preferences. The work here emphasizes the latter issue because it allows the user to employ whatever distance measure is deemed most appropriate. In our implementation, we used a linear formulation for $f_i\,(p_i - t_i)$, therefore, $f_i\,(p_i - t_i) = (p_i - t_i)$. One could also choose to use a piecewise-linear or quadratic form for f_i to give greater emphasis to OAR overdose. We used $\alpha_i =$ in this work, but other structure-specific weights can be specified by the user. The selection of α_i is strongly

Table 1. Equivalent uniform dose parameters for the head-and-neck and pelvic cases

Case	a
Head-and-neck	
Planning target volume	−8.0
Left parotid	5.0
Right parotid	5.0
Cord	7.4
Pelvic	
Planning target volume	−10.0
Bladder	6.0
Rectum	6.0
Small bowel	7.7

influenced by clinician preferences and outside the scope of this work. We illustrate this approach with an example in the Results section (Plan ranking: Composite criteria).

Pre-emptive selection. The general approach of pre-emptive multi-criteria optimization is to prioritize goals and then use these priorities to sequentially determine and apply a series of additional constraints that successively reduce the size of the set of acceptable plans. One approach to this process is to solve a sequence of optimization problems generated by successively adding constraints based on the prioritized goals. We assume that the columns of the SPR view are arranged in order of decreasing structure importance. The results in the first column therefore reflect an ordering of the plans with respect to the highest-priority goal, and inspection of this column leads to the selection of a threshold value for that goal, beyond which a plan is no longer considered acceptable. In the next iteration, the second column is scanned for those plans that satisfied the constraint corresponding to the first goal, proceeding downward in the column until a selected threshold for the second goal is reached. This step eliminates some of the plans that satisfied the initial goal. Successive iterations are then performed with additional goal constraints for successive columns until one plan (or a small number of plans) remains. This approach is illustrated with an example in the Results section (Plan ranking: Pre-emptive selection).

Modeling the planning surface

Often, the relationship between the PTV–OAR and OAR–OAR constraints is not well understood during the optimization process. The data available from the generation of multiple plans provide the basis for modeling the planning surface. The goal of modeling the planning surface is to gain an understanding of the changes induced on the dose–volume values of one OAR by changing the constraints on another OAR. In addition, if the planning surface can be adequately modeled, then insights into the potential benefits (or the absence of such benefits) that might be obtained from the generation of additional plans will be facilitated through the use of such a collection of models derived from the properties of the initial plan set. The user can specify treatment properties for an OAR (these could be dose–volume value, maximum dose, or EUD) to be modeled using linear or quadratic representations. In this work, we construct models for a variety of OAR properties. The input space corresponds to constraints for each of the OARs, and the output space that is modeled corresponds to the achieved OAR properties, which include maximum dose, dose–volume values, and EUD.

A quadratic model of a PTV or OAR property p_j would have the following algebraic form:

$$p_j \sim k_j + l_j u + u' M_j u,$$

where j indexes the plan properties, u is a vector of variables corresponding to the constraint settings that were varied to generate the plan set, and k_j (a scalar), l_j (a vector), and M_j (a symmetric matrix) are fitting parameters. This is a second-order Taylor approximation of a multivariate function representing the property as a function of the constraint settings. The fitting parameters were computed using linear programming optimization to minimize the maximum relative error of the fit of the p_j values of the plans. We minimized the maximum relative error instead of minimizing the sum of squared errors because the former limits the largest error in comparison with minimizing *average* squared-error. (Note that a linear model corresponding to a first-order Taylor approximation has only the $k_j + l_j u$ terms.) Thus, the quadratic model includes terms of the form mu_k^2 and $\mu u_k u_l$, where m and μ correspond to entries of the matrix M_j. These terms yield a second-order model that takes into account the combined effects of the interaction of pairs of constraint settings (each quadratic model thus has a total of 10 terms: a constant term k_j, a sum of thee linear terms corresponding to the inner product $l_j u$, and a sum of three quadratic and three bilinear terms corresponding to the term $u' M_j u$). We have also experimented with a cubic model, but we found that this led to only a marginal improvement in fitting quality relative to the quadratic model, so we do not consider further extension of the model below. The user could apply this model to predict the value of the property of interest p_j for any setting of the independent variables u. Those settings that provided desirable predicted property values could then be input to the treatment-planning software to generate the resulting treatment plan, which finally could be evaluated and compared with previous plans using our decision-support tools to assess the results of this second round of plan generation.

The problem of determining the model parameters that minimize the maximum relative error is stated as an unconstrained optimization problem:

$$\text{minimize } (\max_i \{|p_j^i - k_j + l_j u^i + u^{i\prime} M_j u^i| / p_j^i\}),$$

where i denotes a plan index, and the optimization is performed with respect to the set of variables consisting of the unknown model parameters k_j, l_j, and M_j. This problem can be solved by applying standard linear programming software (such as CPLEX) to the equivalent problem:

$$\text{minimize } y$$
$$\text{such that } -y \le (p_j^i - k_j + l_j u^i + u^{i\prime} M_j u^i)/p_j^i \le y \text{ for i} = 1, \ \ldots \ , \ \text{N} \ ,$$

where N is the number of plans (125 in our experiments). Linear programming problems of this size are solved in a fraction of a second with current software. Following the above formulation, the functional form for a plan property can be expressed as $P \sim \rho + \alpha.lp + \beta.c + \chi.rp + \delta.lp^2 + \tau.c^2 + \varepsilon.rp^2 + \phi.c.lp + \gamma.c.rp + \eta.lp.rp$, where P is the output plan property, ρ, α, β, χ, δ, τ, ε, ϕ, γ, and η are the coefficients in the quadratic model and lp, rp, and c are the input plan constraints corresponding to three OARs.

RESULTS

GUI for plan evaluation

In our results, we only consider results for the OARs because the plans were normalized to achieve the same PTV coverage. We illustrate the ''multi-objective'' dilemma often

Table 2. Input constraints and actual plan values for doses to 33% of parotid volumes and cord maximum in five plans that illustrate the improvement of some organs at risk over others in any given plan (shown in Fig. 3)

	Input constraints (Gy)			Plan values (Gy)		
Plan no.	33% of left parotid	33% of right parotid	Cord maximum	33% of left parotid	33% of right parotid	Cord maximum
1	32	32	33	40.2	42.2	35.6
2	20	32	45	27.9	42.4	48.6
3	32	20	45	40.3	34.3	48.0
4	26	26	45	37.4	39.2	48.1
5	20	32	33	31.4	43.8	38.3

facing the clinician in IMRT planning, along with some of the features of the GUI using five plans for a head-and-neck case corresponding to the input constraints listed in Table 2. Table 3 shows the PTV volume receiving at least 110% of the prescription dose and EUD for the PTV and OARs. Figure 3 shows an overlay of DVHs by structure corresponding to the five plans, whereas Fig. 4 shows the same DVHs on a plan-by-plan basis. It is clear that no plan outperforms the other when considering all OARs, whereas the DVHs for the PTV are nearly the same for all plans because PTV dose was normalized across plans.

We illustrate how a GUI that was developed to support a multi-plan framework can be used to sort and rank plans. Figure 5a shows sample spreadsheets from the GUI for five head-and-neck plans. This screenshot shows the SPR or Full-Sort option based on physical dose properties. Plan 2 results in the lowest dose to 33% of the left parotid, Plan 3 results in the lowest dose to the 33% of the right parotid, and Plan 5 produces the lowest maximum dose to the spinal cord. Figure 5b shows the GUI spreadsheet in Single-Column Sort or PSR mode. We describe how to quantitatively evaluate these spreadsheets below. Figure 6 shows the GUI spreadsheet based on EUDs for the five head-and-neck plans (Figs. 3 and 4).

Plan ranking

Composite criteria. In the head-and-neck case, we set all $\alpha_i = 1$, and we set the targets of the parotid and cord dose as follows: 37.5 Gy for the maximum cord dose, 36 Gy to 33% of the left parotid, and 39 Gy to 33% of the right parotid. These values correspond to the plan targets t_i described in the Methods section (Plan ranking scheme: Composite criteria). If we consider the weighted sum objective for the five plans shown in Fig. 5b, we note that Plan 5 would rank highest with the smallest composite score value. If desired, the user can select the SPR view, which can be used to validate the selection of Plan 5, provided that its behavior with respect to other columns was considered acceptable. Alternatively, the user can employ the PSR view sorted by the composite column to display the other properties of Plan 5 in Row 1. The same methodology can be used to rank all 125 plans in the head-and-neck and pelvic cases, with the ''optimal'' plan being the one corresponding to the plan with the smallest composite objective. The plan targets for the pelvic case are 33 Gy to 35% of the rectum, 35 Gy to 25% of the bowel, and 40 Gy to 35% of the bladder. Similar criteria can be used to sort plans by EUD in Fig. 6.

Pre-emptive selection. Suppose that the cord maximum dose is assigned the highest priority, and a cord maximum dose constraint of 37.5 Gy is added. Using the SPR view (Fig. 5a), note that the feasible plans with respect to this constraint are Plans 1 and 5. Next, let us assume that the left parotid dose (to 33% of its volume) is assigned second priority, with the requirement that the parotid dose not be greater than 36 Gy. From the first SPR column in Fig. 5a, we observe that Plans 1 and 5 are still feasible. Finally, the right parotid dose (to 33% of its volume) is assigned the next-highest priority, and an additional constraint is imposed such that the right parotid dose not be greater than 39 Gy. Using the SPR view again, we note that only Plan 1 is feasible. The above description provides an illustrative (but not unique) example of successive selection of dose values resulting in plan ranking. Note that the above example is not to be taken as a recipe for optimal selection of dose values for pre-emptive selection.

The above ranking scheme can be extended to all 125 plans generated. Using the same priority assignment as discussed above, the addition of the cord maximum dose constraint

Table 3. PTV hotspot* and EUD values for the PTV and OARs in the five plans shown in Fig. 3

Plan no.	PTV hotspot (%)	PTV EUD (Gy)	Left parotid EUD (Gy)	Right parotid EUD (Gy)	Cord EUD (Gy)
1	12.4	51.2	40.0	42.3	26.7
2	8.9	50.6	35.5	42.7	38.2
3	0.2	50.7	39.1	39.0	38.3
4	<0.1	50.7	37.8	40.8	38.0
5	20.0	51.1	36.8	43.9	27.0

Abbreviations: PTV = planning target volume; EUD = equivalent uniform dose; OAR = organ at risk.
* Percentage volume of PTV receiving at least 110% of prescription dose.

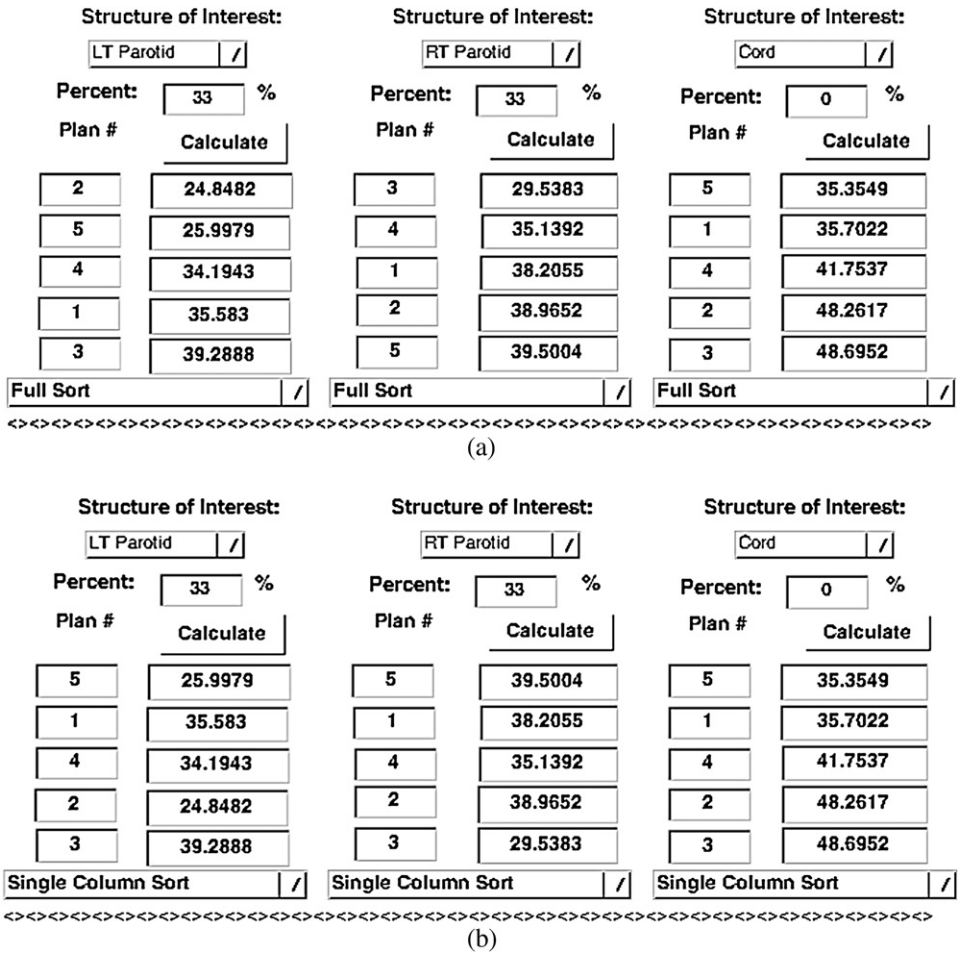

Fig. 5. (a) Graphic user interface (GUI) using ''Full-Sort'' or structure-wise plan ranking to rank plans in terms of diminishing quality for selected volumes for each of the three organs at risk (OARs) for the five sample plans in the head-and-neck case. (b) GUI using ''Single-Column-Sort'' or plan-wise single-criterion ranking to rank plans in terms of diminishing quality for the cord, along with the properties for the other two OARs for that plan ranking.

results in 7 feasible plans. The addition of the left parotid dose constraint results in only 5 of those 7 plans being feasible. Finally, the addition of the right parotid dose constraint results in only 3 plans being feasible. At this point the user can use additional criteria (such as the PTV hotspot or spatial dose distribution) to select the ''optimal'' treatment plan from the selection of 3 feasible plans. To rank all 125 plans in the pelvic case, the rectum dose (to 35% of its volume) was

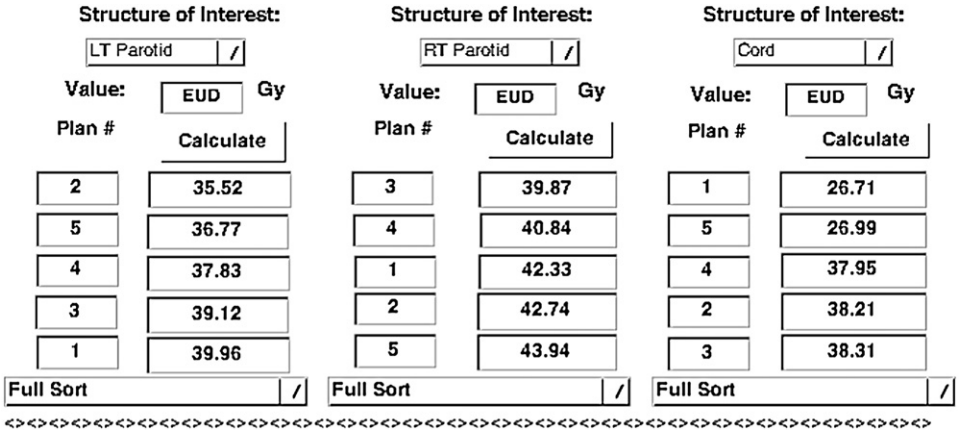

Fig. 6. Graphic user interface using ''Full-Sort'' to rank plans in terms of diminishing equivalent uniform dose quality for each of the three organs at risk in the head-and-neck case for the five sample plans in Tables 2 and 3.

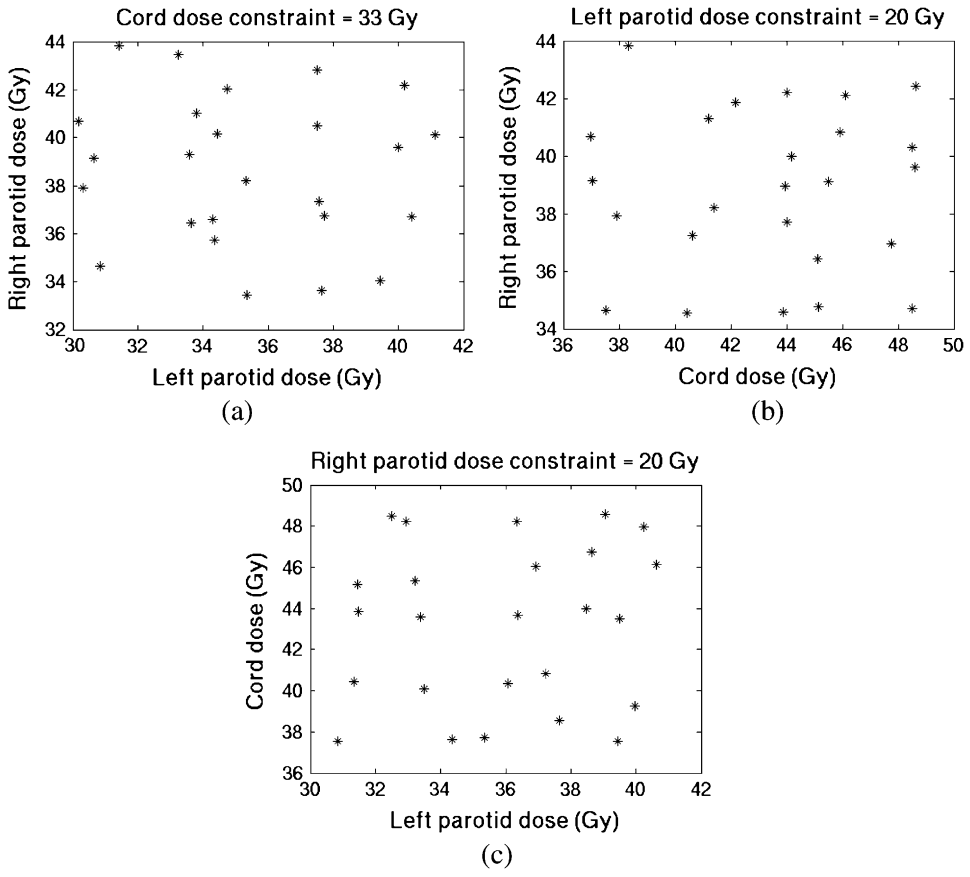

Fig. 7. Space of plans corresponding to the minimum values of input constraints for the (a) cord, (b) left parotid, and (c) right parotid. Each figure corresponds to 25 plans for the minimum input organ-at-risk constraint.

assigned the highest priority. The addition of the constraint that the rectum dose not be greater than 33 Gy results in 23 of the initial 125 plans being feasible. If the bowel dose (to 25% of its volume) is assigned the next-highest priority, and a constraint requiring it to be no greater than 35 Gy is added, this results in 7 of the 23 plans still being feasible. If a final constraint requiring that bladder dose (to 35% of its volume) not be greater than 40 Gy, this leaves only 2 feasible plans. Again, additional criteria such as the PTV ''hotspot'' or spatial distribution can be used to select the ''optimal'' plan. Similar criteria can be used to rank plans by EUD as well (Fig. 6).

Modeling the planning surface

The difficulty in understanding the planning surface is illustrated by trying to identify the space corresponding to the plan properties for minimum values of OAR input constraints. The space of plans corresponding to the minimum values of plan properties is different for each of the three minimum constraint values (Fig. 7).

Although the PTV hotspot (volume of PTV receiving 110% of the prescription dose) was modeled, the data were considered too random to fit any mathematic model. Therefore, we only consider modeling the OAR properties. Table 4 summarizes the relative error for the head-and-neck case produced by modeling the plan properties using quadratic and linear models. For the head-and-neck case, the quadratic model generally gave satisfactory results (<6% difference) for the cord and left and right parotids DVH values and the parotid EUD values. The cord EUD data, however, resulted in a poor fit to the data. The quadratic model results in smaller fit errors compared with the linear model.

Figure 8 illustrates some details of these results for the quadratic fit to the dose received by 33% of the right parotid and

Table 4. Summary of maximum relative error (as percentages) for head-and-neck case produced by modeling the plan properties using quadratic and linear models. The plan properties considered are the maximum cord dose, the dose received by 33% of the parotids, and the organ at risk equivalent uniform doses (EUDs)

	Cord (maximum dose)	Left parotid (33% volume)	Right parotid (33% volume)	Cord EUD	Left parotid EUD	Right parotid EUD
Quadratic	2.6	4.8	2.8	13.4	5.4	4.2
Linear	3.6	5.8	3.1	17.9	5.5	5.0

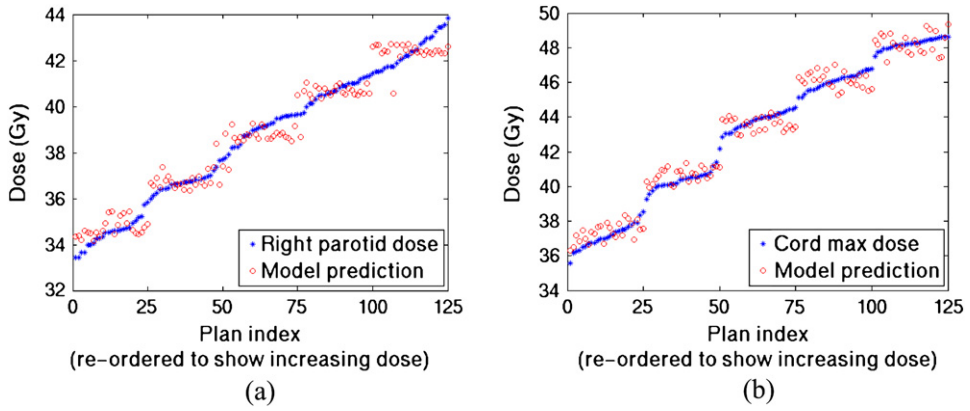

Fig. 8. Results of quadratic models obtained by minimizing maximum relative error between predicted plan property and actual plan property for (a) dose received by 33% of right parotid, and (b) maximum cord dose.

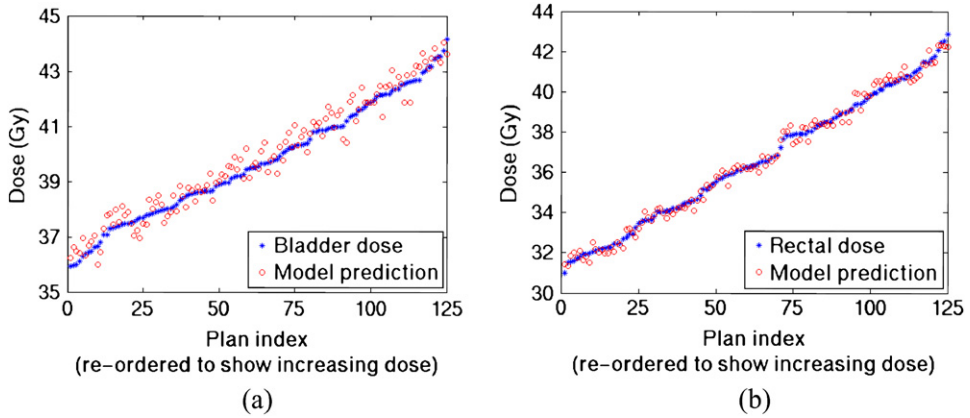

Fig. 9. Results of quadratic models obtained by minimizing maximum relative error between predicted plan property and actual plan property for (a) dose received by 35% of the bladder, and (b) dose received by 35% of the rectum. Note that the maximum relative error of the model prediction is $<1.5\%$ in (b) and that the predicted doses "track" the actual doses.

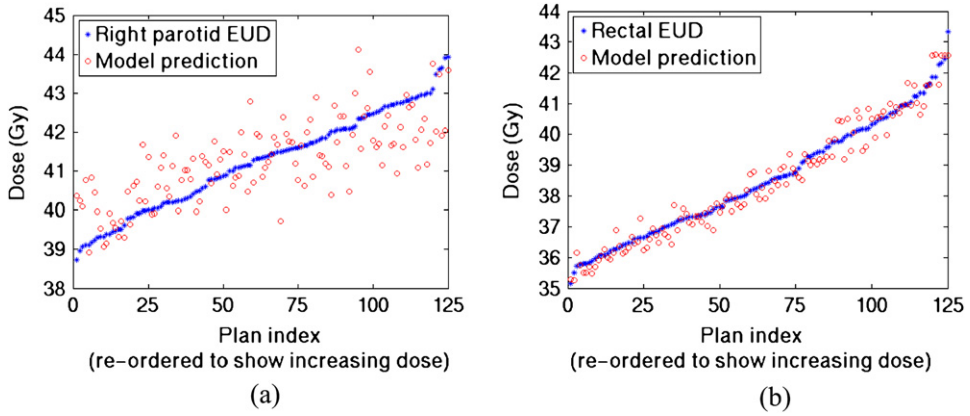

Fig. 10. Results of quadratic models obtained by minimizing maximum relative error between predicted plan property and actual plan property for (a) the right parotid equivalent uniform dose (EUD) in the head-and-neck case, and (b) the rectum EUD in the pelvic case.

the maximum cord dose. (In these figures the plan indices are reordered in each subfigure, so that the displayed delivered doses are increasing. This is done to show the distribution of the delivered doses and to better allow comparisons of delivered doses with the corresponding doses as predicted by the quadratic models. Similar observations apply to Figs. 9 and 10).

Table 5 summarizes the relative error obtained for the linear and quadratic models in the pelvic case. The quadratic model improves the fit error compared with the linear model.

Table 5. Summary of relative error (as percentages) for pelvic case produced by modeling the plan properties using quadratic and linear models

	Bladder (35% volume)	Rectum (35% volume)	Bowel (25% volume)	Bladder EUD	Rectum EUD	Bowel EUD
Quadratic	1.9	1.5	1.8	1.9	1.7	2.2
Linear	2.2	2.1	2.1	2.5	2.4	2.5

The plan properties considered are the maximum cord dose, the dose received by 33% of the parotids, and the organ at risk equivalent uniform doses (EUDs).

Figure 9 shows the excellent agreement of the dose values predicted by the quadratic model with the actual rectum and bladder dose values (to 35% of their volumes) for the 125 plans. The quadratic model again fits the data with lower errors than the linear model. The fit errors for the pelvic case were generally smaller than those in the head-and-neck case. The fit errors were <3% for the pelvic case and, except for the cord EUD, the fit errors were <6% in the head-and-neck case. Figure 10 shows the fits of the quadratic model for the right parotid and rectum EUD values. The fit errors were generally better for the pelvic case compared with the head-and-neck case.

Having demonstrated that the plan properties can be adequately modeled with the aforementioned fit errors, we can proceed to understand the relationship between OARs. Because most IMRT plans typically involve more than two OARs, graphically describing the interrelationship between them can pose a challenge. The quadratic model used to model each plan property can be used to ''fill'' the plan surface with a finer resolution than that presented by the 125 plans. We present two graphic displays by which the plan surface may be explored. One can use one-dimensional representations of a specific plan property corresponding to one OAR as a function of the constraint settings for another OAR for fixed values of the remaining OARs. This is illustrated in Fig. 11. Figure 11a shows the dose to 35% of the rectum as a function of the bowel constraint settings, five different rectum constraint settings, and for a fixed bladder constraint of 27 Gy. Figure 11b shows the rectum dose as a function of the bladder constraint settings, five rectum constraint settings, and for a fixed bowel constraint of 28 Gy. Alternatively, one can use contour plots to display the dose–volume property for one OAR as a function of the input dose constraints for two OARs. (If there were more than three OARs, the constraint settings for all but two of the OARs would have to be fixed to obtain such a projection of the quadratic model.) Figure 12 shows representative contour plots for the dose to 35% of the rectum as a function of the bladder and bowel input dose constraints for a fixed-input rectum dose constraint. One can navigate through these contour plots to determine the optimal region of plan space, and if the optimal plan does not exist in the initial collection of the multi-plan framework generated, then the needed input constraints using both Figs. 11 and 12 can be determined for the desired output plan properties. The achievable phase space points fall inside the projections in Figs. 11 and 12.

DISCUSSION

Although advanced optimization algorithms continue to be developed, the planning surface is not known *a priori*. As a result, IMRT planning continues to be iterative, and the quality of IMRT plans is generally a function of the experience of the planner. The goal in this work was twofold: (*1*) to develop a DSS that will enable the sorting and ranking of treatment plans, and (*2*) to model the plan surface using quadratic models and linear programming optimization. Our results demonstrate that the plan surface can generally be modeled using quadratic models and linear programming techniques with errors of <5%. We also found that the

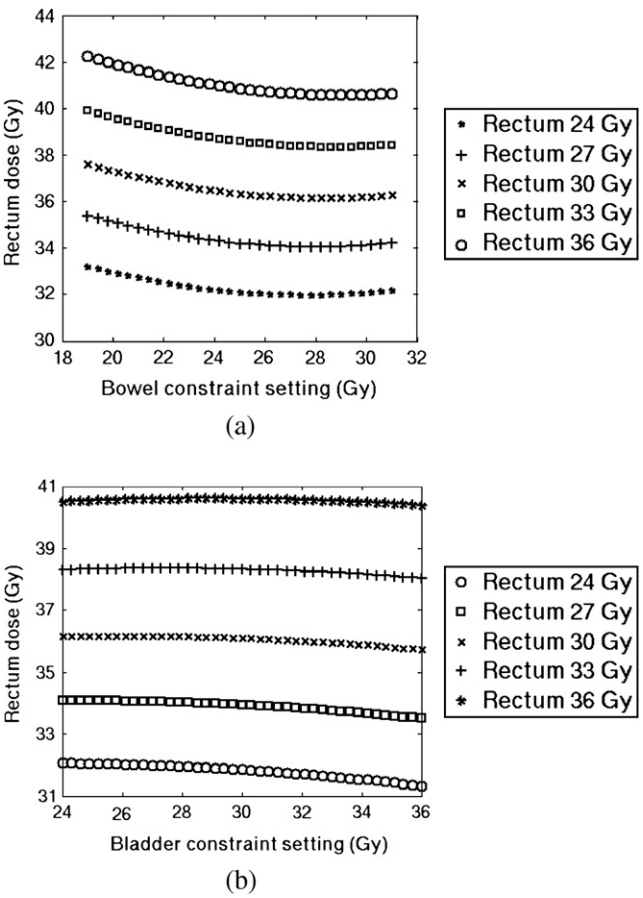

Fig. 11. Dose to 35% of the rectum for five rectum input constraint settings as a function of the (a) bowel (bladder dose input constraint fixed at 27 Gy), and (b) bladder (bowel input dose constraint fixed at 28 Gy) dose-constraint settings.

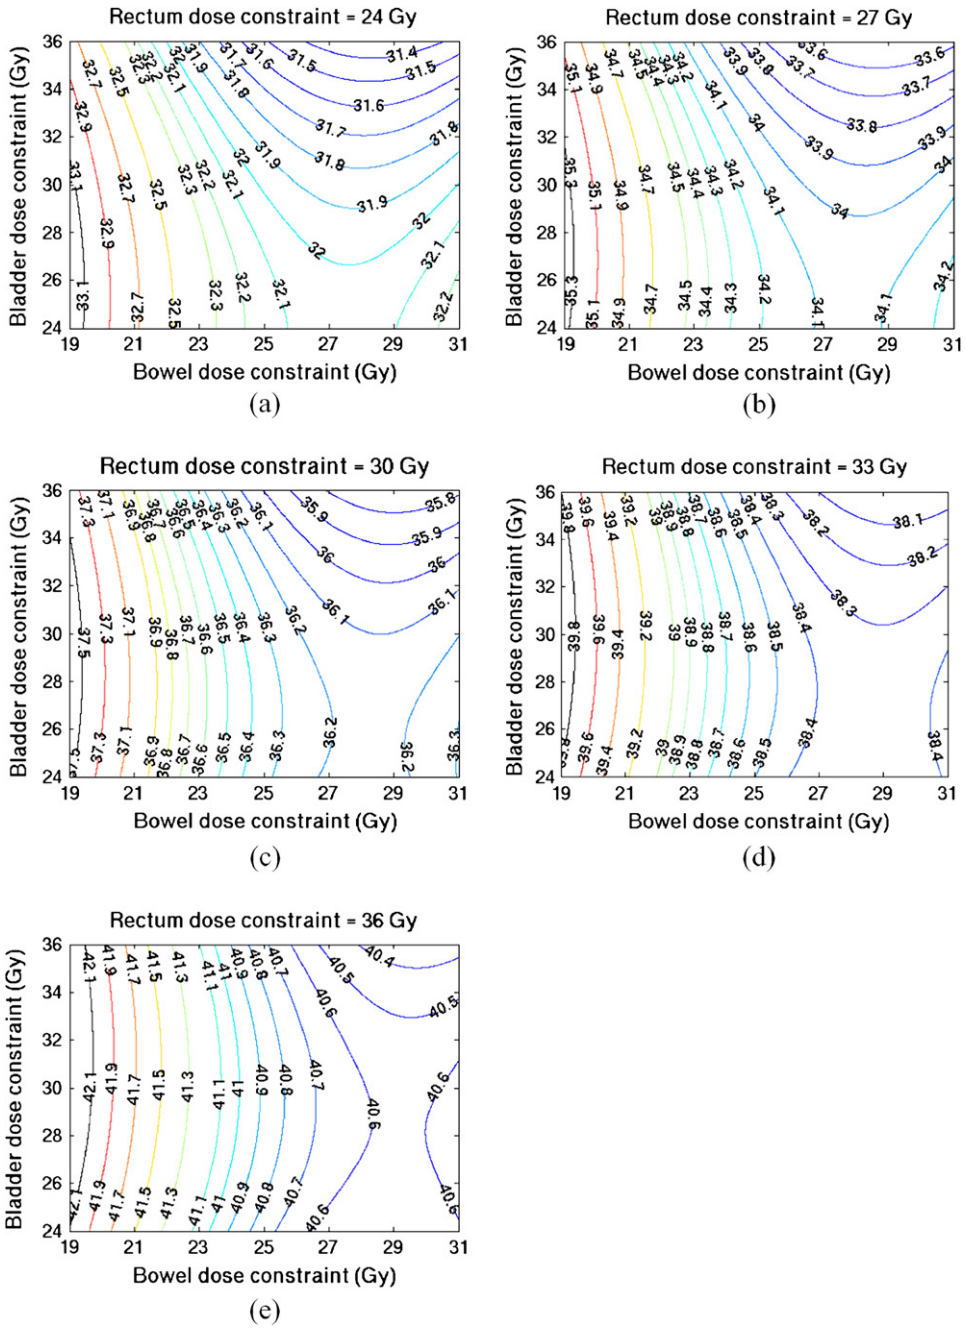

Fig. 12. Contour plots of the rectum dose in the pelvic case as a function of the input dose constraints to 25% of the bowel and 35% of the bladder for input dose constraint of (a) 24 Gy, (b) 27 Gy, (c) 30 Gy, (d) 33 Gy, and (e) 36 Gy to 35% of the rectum.

quadratic model outperformed the linear model. We have restricted ourselves to three OARs and 125 plans per case owing to the limited computing resources. However, the framework can easily be extended to encompass more of the plan surface by increasing the number of plans and/or OARs.

Because of the limitations of computing resources, we believe that the gains to be derived from computerized treatment planning have yet to be fully realized. To overcome the limitation of computing resources, we have started exploring the multi-plan framework in a distributed high-throughput computing environment. Although radiation therapy is delivered on an outpatient basis and during the workday, computing resources remain largely idle during off-peak hours.

This work is a preliminary effort to change the paradigm of treatment planning from a ''plan-and-evaluate'' followed by ''modify-if-necessary'' approach to a multi-plan framework in which the user has the ability to gain insight into patient-specific planning surface behavior. This process is intended to eliminate guesswork on the part of the planner and provide the clinician with a richer data set to interactively determine the available plan options. The tools developed here are a first step in enabling the user to navigate through the

multidimensional ''plan space'' to select the plan that achieves the desired clinical objectives.

In summary, the multi-plan framework can serve two purposes: (*1*) the determination of the best plan from the collection of plans generated using either composite or pre-emptive criteria, and (*2*) the visualization of the dependence of OAR plan properties on the input constraints. During conventional treatment planning for IMRT, it is difficult to determine the quality of rival treatment plans that seem to outperform other plans with respect to one or more but not all OAR plan properties and to determine the dependence of an OAR property on the set of input constraints.

A proposed clinical workflow using the results of this work would involve the generation of a large collection of treatment plans in which the input constraints for the OARs are varied to span a range of values. The resulting plans could then be evaluated, and the most desirable plan from this collection could be selected. Additionally, the collection of plans could be used to model the actual OAR plan properties obtained to aid in understanding the relationship between OAR actual plan properties and the input constraints corresponding to all OARs. By modeling the planning surface using a quadratic formalism and linear programming techniques, one can generate contour plots such as the ones shown in Fig. 12. Such plots will assist the clinician in visualizing the tradeoffs in IMRT planning for which competing objectives exist. If the clinician can identify a region in these contour plots that is most desirable, and a plan corresponding to that region does not exist in the initial collection of plans, one can then use the input OAR constraints corresponding to the desired output OAR plan properties to generate additional plan(s).

REFERENCES

1. Gopal R, Starkschall G. Plan space: Representation of treatment plans in multidimensional space. *Int J Radiat Oncol Biol Phys* 2002;53:1328–1336.
2. Rosen I, Liu HH, Childress N, *et al*. Interactively exploring optimized treatment plans. *Int J Radiat Oncol Biol Phys* 2005; 61:570–582.
3. Zhang X, Wang X, Dong L, *et al*. A sensitivity-guided algorithm for automated determination of IMRT objective function parameters. *Med Phys* 2006;33:2935–2944.
4. Yu Y. Multiobjective decision theory for computational optimization in radiation therapy. *Med Phys* 1997;24:1445–1454.
5. Xing L, Li JG, Donaldson S, *et al*. Optimization of importance factors in inverse planning. *Phys Med Biol* 1999;44:2525–2536.
6. Romeijn HE, Dempsey JF, Li JG. A unifying framework for multi-criteria fluence map optimization models. *Phys Med Biol* 2004;49:1991–2013.
7. Craft D, Halabi T, Bortfeld T. Exploration of tradeoffs in intensity-modulated radiotherapy. *Phys Med Biol* 2005;50: 5857–5868.
8. Wu Q, Mohan R, Niemierko A, *et al*. Optimization of intensity-modulated radiotherapy plans based on equivalent uniform dose. *Int J Radiat Oncol Biol Phys* 2002;52: 224–235.
9. Thieke C. Multicriteria optimization in inverse radiotherapy planning [Ph.D. thesis]. Heidelberg, Germany: Ruprecht-Karls-Universität; 2003. p. 33–37.
10. Emami B, Lyman J, Brown A, *et al*. Tolerance of normal tissue to therapeutic irradiation. *Int J Radiat Oncol Biol Phys* 1991;21: 109–122.
11. Aggarwal, Charu C. Re-designing distance functions and distance based-applications for high dimensional data. *ACM SIGMOD Record* 2001;30:13–18.