

Reduction of IMRT beam complexity through the use of beam modulation penalties in the objective function

Martha M. Matuszak

Department of Radiation Oncology and Department of Nuclear Engineering and Radiological Sciences, University of Michigan, Ann Arbor, Michigan 48109

Edward W. Larsen

Department of Nuclear Engineering and Radiological Sciences, University of Michigan, Ann Arbor, Michigan 48109

Benedick A. Fraass

Department of Radiation Oncology, University of Michigan, Ann Arbor, Michigan 48109

(Received 6 June 2006; revised 21 November 2006; accepted for publication 21 November 2006; published 18 January 2007)

Inverse planned intensity modulated radiation therapy (IMRT) has become commonplace in treatment centers across the world. Due to the implications of beam complexity on treatment planning, delivery, and quality assurance, several methods have been proposed to reduce the complexity. These methods include beamlet intensity restrictions, smoothing procedures, and direct aperture optimization. Many of these methods typically sacrifice target coverage and/or normal tissue sparing in return for increased beam smoothness and delivery efficiency. In the present work, we penalize beam modulation in the inverse planning cost function to reduce beam complexity and increase delivery efficiency, while maintaining dosimetric quality. Three modulation penalties were tested: two that penalized deviation from Savitzky-Golay filtered versions of the optimized beams, and one that penalized the plan intensity map variation (a measure of overall beam modulation). The modulation penalties were applied at varying weights in a weighted sum objective (or cost) function to investigate their ability to reduce beam complexity while preserving IMRT plan quality. The behavior of the penalties was characterized on a CT phantom, and then clinical optimization comparisons were performed in the brain, prostate, and head/neck. Comparisons were made between (i) plans with a baseline cost function (ii) plans with a baseline cost function employing maximum beamlet intensity limits, and (iii) plans with each of the modulation penalties added to the baseline cost function. Plan analysis was based upon dose-volume histograms, relevant dose metrics, beam modulation, and monitor units required for step and shoot delivery. Each of the techniques yielded improvements over a baseline cost function in terms of MU reduction. In most cases, this was achieved with minimal change to the plan DVHs and metrics. In all cases, an acceptable plan was reached with each of the methods while reducing MU substantially. Each individual method has merit as a tool for reducing IMRT beam complexity and could be easily applied in the clinic to improve overall inverse plan quality. However, the penalty based upon the plan intensity map variation consistently produced the most delivery-efficient plans with the fewest computations. © 2007 American Association of Physicists in Medicine.

[DOI: [10.1118/1.2409749](https://doi.org/10.1118/1.2409749)]

Key words: IMRT, optimization, beam complexity

I. INTRODUCTION

The evolution of intensity modulated radiation therapy (IMRT) into a common radiotherapy treatment modality has brought into focus not only its potential advantages, such as improved target conformality, easier planning of simultaneous boost treatments, and reductions in overall normal tissue toxicity, but also potential disadvantages. With the inherent flexibility in planning afforded by the complex inverse planning optimization process comes certain limitations to efficient IMRT planning and delivery. The inverse planning process itself is subject to noise and artifacts that may yield undesirable high-intensity peaks and small beamlet-to-

beamlet fluctuations across a given IMRT field. These artifacts can complicate the IMRT planning process and prevent accurate and efficient deliveries.

Previous efforts to reduce IMRT beam complexity include the application of beamlet restrictions, the use of beam smoothing procedures, and the direct optimization of multi-leaf collimator segments. The use of beamlet restrictions, such as applying maximum intensity limits to beamlets during optimization, can reduce modulation and increase delivery efficiency as long as the limits are not so strict as to interfere with the dosimetric plan objectives.¹ Applying intensity limits allows the cost function to work independently of any outside criteria, although intensity limits may only partly reduce the consequences of the beam complexity

problem by removing large intensity spikes. The large increases in the MU requirements for IMRT plans (compared to a conventional conformal plan) can be a consequence of high-intensity peaks within the optimized beams. The use of maximum intensity limits can remove these peaks and be beneficial for many cases, but it does not prevent smaller fluctuations in modulation between neighboring beamlets that can also increase MU. In these cases, the discretization of intensity levels (also done to varying degrees in leaf sequencing procedures) during optimization may be useful, assuming that an acceptable plan can be found in the reduced solution space.^{2,3}

Smoothing procedures can be applied either inside or outside the optimization loop. Smoothing IMRT beams outside the optimization process may either produce limited results or require reoptimization due to the plan degradation that occurs during postoptimization smoothing.⁴⁻⁶ Inside the optimization loop, smoothing can be done by (i) including smoothness criteria inside the objective function or (ii) smoothing beams after each cycle or iteration.⁷ The latter method is easier to implement, but it suffers from the same difficulty as applying interventions postoptimization: when any procedure is applied outside the cost function, the impact of that procedure on the cost function cannot be weighed according to the dosimetric consequences. Procedures in this category include smoothing after each iteration, applying maximum beamlet intensity limits, and discretizing the number of allowable beamlet intensity levels. These methods generally cannot distinguish between desirable and undesirable gradients and are likely to degrade plans in areas where steep gradients are necessary.^{5,7} Each of the above methods, while superior to methods that are applied outside the optimization loop, may still interfere with the ability of the optimization system to meet the inverse planning objectives. In previous studies, this has resulted in inferior target coverage.⁷

Implemented carefully, methods that penalize modulation as a part of the cost function can effectively improve delivery efficiency, while taking into account the dosimetric tradeoffs to be made with IMRT plan objectives. Spirou *et al.* have presented a comparison of smoothing, using a Savitzky-Golay filter along the direction of the MLC leaf travel, inside and outside the cost function. They found that smoothing inside the cost function was superior in terms of producing sharper dose gradients, better dose homogeneity, and better critical organ sparing, especially for more complex cases.⁷ However, again, care must be taken to design the objective function in such a way that the smoothing process or modulation penalty does not dominate the solution and interfere with normal tissue dose limits or target coverage.⁸

Another approach to reduce beam modulation is to directly optimize the shapes and weights of the actual MLC segments that will be used in the delivery of the IMRT plan.⁹⁻¹⁴ This approach removes the need for leaf sequencing and eliminates errors introduced by sequencing approximations. Although there has not been a comprehensive work comparing direct aperture optimization (DAO) with fluence

map optimization (FMO), several studies point to the conclusion that DAO and FMO perform similarly in terms of dosimetric quality and MU in simple cases, such as breast,⁹ but FMO may still produce better dosimetric results in normal tissues for more complicated cases, such as head and neck.^{10,11} The required MU and segments can be substantially less for DAO in these cases, justifying a slight loss in clinical quality. It should be noted that improved leaf sequencing algorithms have been shown to close the delivery efficiency gap between DAO and FMO.¹⁵

As we further our knowledge of IMRT, it has become clear that the value of an IMRT plan should not be judged solely on the resulting DVHs and dose metrics, but also on the efficiency of delivery and other potential problems associated with the plan, such as increased time and effort needed for quality assurance. The work presented here examines several IMRT beam modulation penalties that can be applied as part of the inverse planning cost function, to determine which is best-suited to preserve plan quality regarding (i) target and normal tissue DVHs and dose metrics and (ii) reduced modulation and improved delivery efficiency.

II. METHODS

To promote beam smoothing as part of the inverse planning objective function, three beam modulation penalties were implemented into our in-house treatment planning and optimization system (UMPlan/UMOpt).¹⁶⁻²² UMOpt is a versatile software package designed for 3-D beamlet-based inverse IMRT optimization that can make use of specialized cost function components, or “costlets.” In this work, we study three new costlets that aim to reduce IMRT beam complexity when included in an inverse plan objective function. The three costlets do the following: (i) minimize deviation from a 1-D Savitzky-Golay filtered beam where filtering is done only in the direction of the MLC travel, (ii) minimize deviation from a 2-D Savitzky-Golay filtered beam, or (iii) minimize a measure of the total plan modulation called the quadratic plan intensity map variation (PIMV_q).¹ In order to perform an investigation of these costlets, each was used in the optimization of an IMRT plan for a CT phantom. After testing in the phantom, nine clinical cases (three each for brain, prostate, and head/neck) were optimized using each of the three modulation costlets (added to the baseline cost function) and compared to the plan generated using just the baseline cost function. Comparison was also made to plans employing the same baseline cost function with maximum beamlet intensity limits. We briefly describe each of these methods below.

A. Standard IMRT with a baseline cost function

The standard IMRT optimization plans (the baseline for the comparisons) were simply obtained by optimization of the beamlet intensities according to the minimization of a baseline cost function, without any beam modulation penalties. No restrictions were placed on the allowed values of the beamlet intensities.

B. Maximum intensity limits

IMRT optimization was performed with the baseline cost function as in Sec. II A, but the maximum beamlet intensity in the plan was limited to a predetermined value proportional to the prescription dose divided by the number of beams.¹ This method, discussed in more detail in previous work, was shown to be a simple and effective way to reduce plan MU while reaching all of the plan objectives.¹ We have included it in this analysis to determine whether the more advanced methods of reducing modulation shown below are more advantageous.

C. Beam modulation penalty based on 1-D Savitzky-Golay filtering (SG1D)

The same baseline cost function was used as in the previous two methods, but a modulation penalty based on the one used successfully by Spirou *et al.*⁷ was included as a part of the cost function. In addition to the costs calculated for the target(s) and normal tissues, the following costlet was added into the objective function:

$$C = w \times \sum_{i=0}^B (b_i - s_i)^2. \quad (1)$$

Here, w is the weight assigned to the smoothing importance, B is the number of beamlets in the plan, b_i is the value of the i^{th} beamlet, and s_i is the value of the i^{th} beamlet after the smoothing operation. The smoothed value, s_i , is calculated using a Savitzky-Golay filter.²³ This filter performs a least-squares fit to the elements within a smoothing window to a given degree polynomial (a second order polynomial is used here). The smoothing window is a five row vector of beamlets centered at b_i in the direction of the MLC leaf motion.

D. Beam modulation penalty based on 2-D Savitzky-Golay filtering (SG2D)

Optimization was performed as in Sec. II C, but the smoothed beamlet values, s_i , were calculated using a 5×5 beamlet grid centered at b_i as the smoothing window. This was done to promote modulation reduction in all directions, not just in the direction parallel to MLC travel.

E. Beam modulation penalty based on PIMV_q

Optimization was performed as in the previous two methods, but the beam modulation penalty was based upon a different measure of beam modulation. In addition to the costlets included in the baseline cost function, the following costlet was added to the total objective function calculation:

$$C = w \times \text{PIMV}_q. \quad (2)$$

Here, w is the weight assigned to the smoothing importance and PIMV_q is a quadratic version of the plan intensity map variation (PIMV), described previously.¹ The PIMV provides a measure of the beam modulation across an intensity modulated beam, and PIMV_q was defined to simplify the calculation

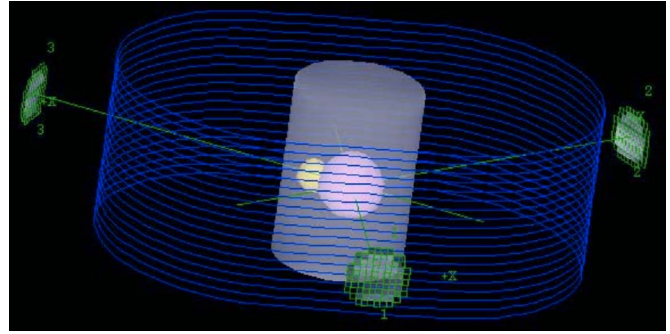


FIG. 1. Illustration of the CT phantom geometry and beam placement. The PTV is the center sphere, “Small” is the off-center sphere, and “Cylinder” is the cylindrical structure. The external contours were not used in the cost function to save calculation time. Note: some figures may appear in color only in the electronic version.

of the gradient during the optimization procedure. The latter quantity is defined as

$$\text{PIMV}_q = \sum_{n=1}^{N_b} \left(\sum_{j=1}^{J-1} \sum_{k=1}^{K-1} \left[(b_{jk} - b_{j,k+1})^2 + (b_{jk} - b_{j+1,k})^2 + \frac{1}{2}(b_{jk} - b_{j+1,k+1})^2 + \frac{1}{2}(b_{jk} - b_{j+1,k-1})^2 \right] \right), \quad (3)$$

where N_b is the number of beams in a plan, J is the maximum number of beamlets in the direction parallel to the motion of the multileaf collimator (MLC), K is the maximum number of beamlets in the direction perpendicular to the motion of the MLC, and b_{jk} is the intensity of the beamlet at the (j,k) grid position. The four terms in Eq. (3) allow all differences between each neighboring beamlet, including the diagonal beamlets, to be taken into account.

The weights used in Eqs. (1) and (2) were varied to demonstrate the potential impact of the beam modulation penalties on the optimal intensity modulated beams as the importance of smoothness was increased.

F. Treatment planning and analysis

As mentioned previously, each of the modulation penalties was validated on a CT phantom and then tested on nine clinical cases—three each for brain, prostate, and head/neck. These plans were then compared to plans employing a baseline cost function or maximum intensity limits. One optimization run was done for the baseline cost function and then several optimization trials were run for each of the other four techniques to characterize plans at a variety of maximum intensity limits and with a range of weights applied to the modulation penalties. The geometry and beam arrangement for the simple CT phantom are shown in Fig. 1, and the high priority planning objectives for the phantom and clinical cases are shown in Table I. In addition, each plan included a lower priority objective of minimizing overall dose to all of the normal structures and uninvolved normal tissues. These objectives were given a low weight so as to not interfere with

TABLE I. Inverse planning objectives for the CT phantom and clinical cases.

CT phantom inverse plan objectives	
Structure	Objectives
PTV	≥ 60 Gy; ≤ 66 Gy
PTV+1 cm	≥ 50 Gy; ≤ 66 Gy; $\leq 20\% \geq 60$ Gy
Small	$\leq 15\% \geq 25$ Gy
Small+1 cm	≤ 66 Gy; Mean ≤ 20 Gy
Cylinder	≤ 70 Gy
Brain inverse plan objectives	
Structure	Objective
PTV1	66 Gy (min 95%, max 105%, and 1% volume up to 110%)
PTV2	60 Gy (min 95%, max 105% of PTV1 with PTV1 coverage priority)
Optic nerves	≤ 60 Gy
Optic chiasm	≤ 60 Gy
Brainstem	≤ 65 Gy
Prostate inverse plan objectives	
Structure	Objectives
Prostate+3 mm	Mean = 75.85 Gy $\pm 3\%$; ≤ 0.5 cc down to 93%; ≤ 0.5 cc up to 115%
Rectum	$\leq 15\% \geq 80$ Gy; $\leq 25\% \geq 75$ Gy; $\leq 35\% \geq 70$ Gy; $\leq 50\% \geq 65$ Gy
Bladder	$\leq 15\% \geq 80$ Gy; $\leq 25\% \geq 75$ Gy; $\leq 35\% \geq 70$ Gy; $\leq 50\% \geq 65$ Gy
Femora	Mean ≤ 50 Gy; $\leq 10\% \geq 52$ Gy
Penile bulb	Mean ≤ 52.5 Gy; $\leq 15\% \geq 70$ Gy
Uninvolved tissue	Max $\leq 100\%$ Rx Dose; Mean ≤ 52.5 Gy; $\leq 15\% \geq 70$ Gy
Head/neck inverse plan objectives	
Structure	Objectives
PTV	70 Gy $\pm 5\%$
Nodal boost PTV	70 Gy $\pm 5\%$
High risk nodal PTV	66 Gy $\pm 5\%$
Low risk nodal PTV	60 Gy $\pm 5\%$
Spinal cord	≤ 45 Gy
Spinal cord+5 mm	≤ 50 Gy
Brainstem	≤ 54 Gy
Contralateral parotid	Mean ≤ 26 Gy
Ipsilateral parotid	Mean ≤ 26 Gy (if possible, otherwise minimize)
Mandible	≤ 70 Gy
Submandibulars	Minimize dose
Oral cavity	≤ 70 Gy

meeting the highly weighted objectives in Table I. The clinical plan objectives were chosen based on in-house IMRT protocols.

Brain cases were optimized with four or five beams originally placed by a dosimetrist, prostate plans were optimized with nine equally spaced axial beams, and head/neck cases were optimized using seven equally spaced axial beams. All cases were planned for a 6 MV linear accelerator (Varian Medical Systems, 21EX) with 120 leaf MLC (0.5 and 1.0 cm

leaf widths). Dose calculations for the inverse planning system were performed by a convolution/superposition algorithm derived from the work of Mackie *et al.*²⁴ All cases were optimized using a quasi-Newton-based search strategy. Plans were initialized with several different starting intensity patterns to ensure that the optimized plans did not represent local minima of the cost function.

After optimization, leaf sequencing for static MLC (SMLC) delivery was performed with an in-house-developed

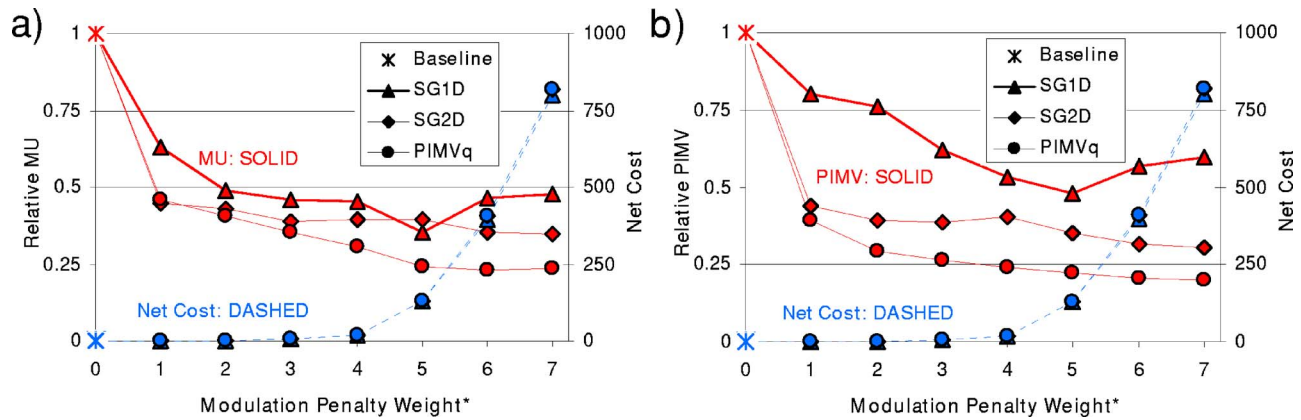


FIG. 2. Optimization results for the CT phantom when including modulation penalties in the cost function. Net cost and relative SMLC are shown as a function of increasing modulation penalty weight in (a), and relative plan intensity map variation and net cost are shown in (b).

leaf sequencer based on the method reported by Bortfeld *et al.*²⁵ Delivery sequences allow up to 250 segments per beam, with the goal of achieving a correspondence between planned and delivered intensities of 1%.

Optimization results for each of the different techniques were compared using dose metrics, DVHs, field modulation and complexity, and efficiency of SMLC delivery. The results from the modulation penalty validation in the CT phantom are presented first and then all methods are compared in the three clinical sites. Our analysis of these results attempted to determine which method produced the highest quality IMRT plans, taking into account the target and normal tissue planning objectives and other important factors such as overall beam complexity and delivery efficiency.

III. RESULTS

A. Modulation penalties in the phantom

The CT phantom case (Fig. 1) was optimized with the baseline objective function described in Table I, and then optimization was performed including each of the modulation penalties with a number of different weights. Figure 2(a) shows the net cost (objective function value less the modulation penalty) and relative SMLC MU as a function of increasing modulation penalty weight (this is a qualitative weight for plotting purposes, not the actual weight used in the optimization, as that varied for each case and method) when including each of the three modulation penalties in the cost function. The results from the baseline cost function are also shown. With increasing modulation penalty weights, we observe large decreases in MU along with a rising trend in the net cost. The net cost increases because, as the modulation penalty increases, the plan deviates further from the baseline optimized plan. To minimize the modulation penalty, tradeoffs are usually made with the lower priority objective of minimizing total dose to the normal tissues. This is indicated by the slowly increasing net cost at the lower penalty weights. As the weight increases further, we observe a steeper increase in the net cost as the modulation penalties begin to dominate the optimization and various high priority objectives are violated. In a clinical case, one would likely

not choose a smoothing penalty weight that interfered with the high priority objectives. The extent to which the modulation penalty would be allowed to trade off with lower priority objectives would have to take into account the increase in normal tissue dose versus the reduction in MU and beam complexity.

To examine the relationship between the monitor units and beam modulation, we also calculated the total plan intensity map variation [equal to a nonquadratic version of Eq. (3), which is a measure of plan modulation] in Fig. 2(b). The relative plan intensity map variation varied from 0.20 to 0.80, indicating the possibility for large reductions in beam complexity with the use of the modulation penalties. We also saw that the use of the PIMV_q penalty produced plans that had the lowest amount of modulation, followed by use of the SG2D penalty, and then the SG1D penalty. In addition, because the SG1D plans promote smoothing only in the direction of the MLC travel, we saw only a slight increase in the MU as compared to the SG2D plans. This means that the SG1D plans retained the most modulation while still seeing a reduction in MU. At higher SG1D penalties, the beams started to deviate from the trend of decreasing modulation and MU. Therefore, this penalty may be less useful at high weights where nonconformance to the filtered beams is highly penalized.

As shown in Fig. 2, the penalty weights were chosen so that the optimized plans would have almost identical objective function values. This required several iterations of choosing the modulation penalty weights, but it allowed for fair plan comparisons. Dose-volume histograms and the intensity modulated beams for the cases optimized at relatively low penalty weights (a and b) and high penalty weights (c and d) are shown in Fig. 3. For the lower penalty weights, which correspond to a weight of 3 in Fig. 2, the normal tissue DVHs were nearly identical to those in the baseline plan. The target coverage is not compromised, but does become more homogeneous as the fields become smoother. In Fig. 3(b), large differences between the beam intensity patterns for each of the modulation penalized plans and the baseline plan can be seen. For the SG1D penalty, the optimal

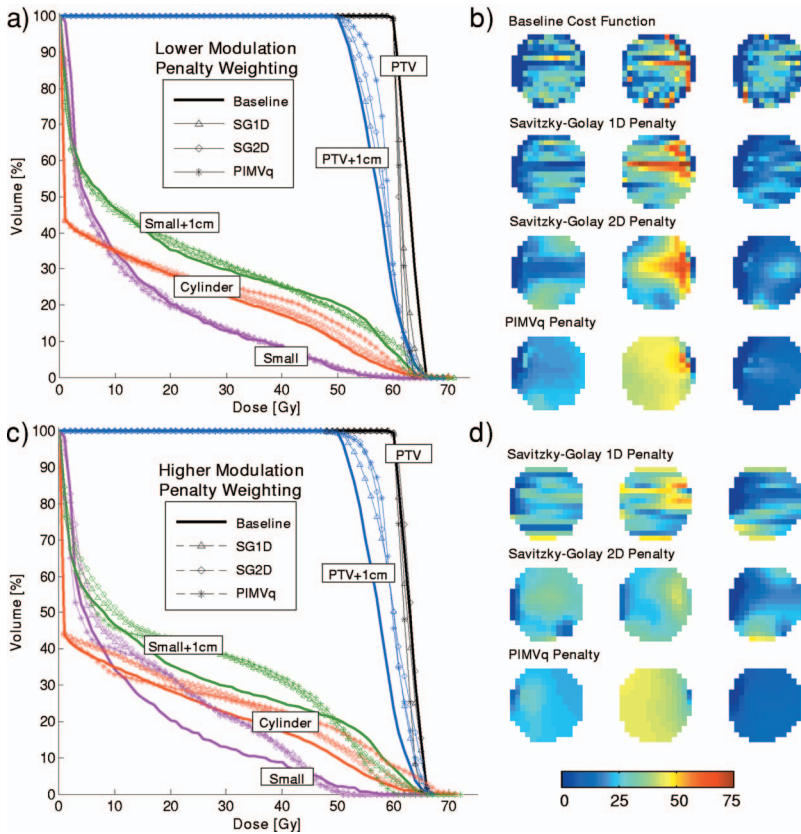


FIG. 3. (a) Optimized DVHs and (b) intensity modulated beams when using a relatively low modulation penalty weight (corresponds to a weight of 3 in Fig. 2) in the phantom case. (c) Optimized DVHs and (d) intensity modulated beams when using a relatively higher modulation penalty weight (corresponds to a weight of 7 in Fig. 2) in the same case. Results using the baseline cost function are shown for comparison.

beams approach a steplike function in the direction perpendicular to the MLC motion. This happens because the filtering is done only in the direction of leaf travel. As Spirou *et al.* suggest, we chose a five beamlet smoothing window, since using a larger smoothing window resulted in an exaggeration of this step function observation, with each row of beamlets having an almost identical weight. Conversely, we found that the use of a smaller window resulted in minimal smoothing and minimal reduction in MU as compared to the baseline plan. The use of the SG2D penalty produced more highly smoothed versions of the original plan with large reductions in (i) the amount of the modulation, (ii) the number of high intensity beamlets, and (iii) the size of beamlet-to-beamlet variations. In the plan optimized with the PIMV_q penalty, the fields were relatively uniform, with the second beam delivering the majority of the intensity. In these beams, the modulation was high only in areas near the overlap of the PTV and the “Small” structure.

In Figs. 3(c) and 3(d), the DVHs and beams for plans optimized at higher modulation penalties (corresponding to a weight of 7 in Fig. 2) are shown. In these optimization trials, we saw an increase in the normal tissue dose while the target coverage remained fairly consistent due to the design of the cost function. However, there were slight violations of the minimum dose objective in the PTV expansion and the dose-volume objective in the “Small” structure. As a result of the increasing modulation penalty weight, the beam complexity was further decreased compared to Fig. 3(b), although the peak locations in the SG1D beams were inconsistent with the original plan. The SG2D beams were smoothly varying and

the intensity was well distributed throughout the three beams. The PIMV_q beams were similar to the beams produced at the lower penalty weight, although (i) they were much more uniform, (ii) beam 2 was slightly more intense, and (iii) the modulation was further reduced. Looking at the large differences in the beam intensity patterns, it appears that the phantom case has a large solution space, i.e., there are many different beam combinations that lead to very similar DVH results.

For this CT phantom example, the PIMV_q penalty achieved the same total objective function values with lower degrees of beam complexity and MU than the SG penalties. However, tradeoffs are still being made within the dose objectives, and each of the modulation penalties causes slightly different tradeoffs to be made in the cost function. Thus, each of the penalties may have merit in different geometries and in situations where planners may have different tradeoff preferences. We hypothesize that the SG1D and SG2D penalties will be useful in cases that require high intensity gradients across a field, and that the PIMV_q penalty will be most applicable because it generally provides the same objective value for the least amount of beam complexity. These ideas will recur in the discussion of the results from the brain, prostate, and head/neck trials.

B. Clinical studies

Nine clinical cases were optimized using each of the three modulation penalties in the objective function at varying weights. The results were compared to plans using the base-




FIG. 4. Plots illustrating relative MU versus net cost tradeoffs in the brain cases when using each of the modulation penalties compared to the baseline cost function and maximum intensity limits. All values are relative to the baseline plan.

line objective functions without or with maximum beamlet intensity limits. The inclusion of the modulation penalties at reasonable weights did not compromise the clinical planning objectives and generally resulted in smoother, less complex intensity patterns that could be delivered with significant reductions in monitor units compared to the baseline IMRT plans. This gain in delivery efficiency was sometimes at the expense of increased dose to the normal tissues. Therefore, in addition to the results presenting the maximum MU reductions while still meeting the clinical objectives, we have included the maximum MU reduction while maintaining all normal tissue mean doses achieved with the baseline cost function. This was done by adding mean dose costlets along with the modulation penalties to the baseline cost function at the mean dose levels achieved by the baseline plan optimization. These results will be discussed for each individual treatment site.

1. Brain

In the three brain cases tested, each beam complexity reduction technique produced plans that met the high priority dose objectives and could be delivered with a significant MU reduction compared to the baseline IMRT plan. Figure 4

shows all of the optimization runs plotted as relative MU versus net cost. In each plan shown, all of the high priority objectives in Table I are met. The increase in delivery efficiency seen in the brain is usually gained as a result of a tradeoff between the modulation penalty (or maximum beamlet intensity) and the low-priority objective of minimizing overall dose to the normal tissues. This can be seen in Fig. 5(a), which shows, for *brain3*, optimized DVHs with the baseline cost function, with maximum intensity limits, and with each of the modulation penalties applied. The plans shown are at moderate modulation penalty weights and are very similar in net cost. The intensity modulated beams from each of the techniques are shown in Fig. 5(b). A reduction in overall beam modulation from the baseline cost function to the other techniques can be seen, and the MU reductions are shown in parentheses. The same general trend in relative MU was observed in the other brain cases with the PIMV_q modulation penalty, consistently producing the largest reduction in MU at similar objective function values, followed by maximum intensity limits, the SG2D penalty, and then the SG1D penalty (Fig. 4). The maximum possible MU reductions for each method compared to the baseline plan are shown in Table II for each of the brain cases. The maximum possible

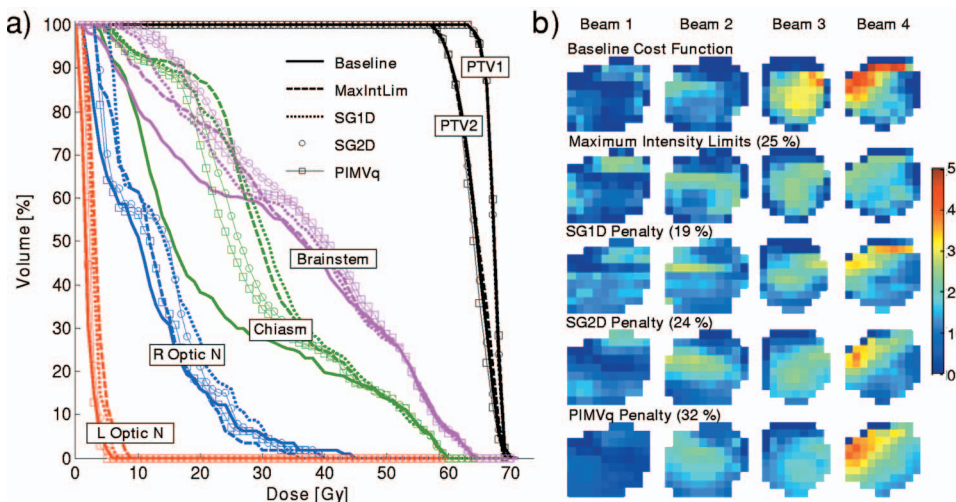


FIG. 5. (a) *Brain3* DVHs shown optimized with each of the different techniques for reducing modulation and (b) corresponding *Brain3* beams optimized with each of the beam modulation reduction methods. SMLC reductions are in parentheses.

TABLE II. Maximum MU reductions possible in the brain using each optimization technique while (i) meeting high priority plan objectives and (ii) maintaining mean normal tissue doses from the baseline plan.

Optimization technique	Brain1		Brain2		Brain3	
	(i) only	(i) and (ii)	(i) only	(i) and (ii)	(i) only	(i) and (ii)
Maximum intensity limits	49.0%	9.4%	27.3%	2.5%	31.1%	4.7%
SG1D filter penalty	35.8%	3.9%	10.1%	0.0%	22.4%	13.5%
SG2D filter penalty	49.2%	12.7%	18.1%	3.6%	27.5%	12.9%
PIMV _q penalty	56.7%	16.4%	25.8%	12.8%	39.1%	13.1%

MU reduction is defined as the maximum MU reduction possible while still meeting all of the high priority clinical planning objectives in Table I. In practice, these plans may not be acceptable to the physician because there can be significant mean dose increases in the normal tissues. Therefore, we have also included the more modest maximum MU reductions possible when maintaining the mean normal tissue doses from the baseline plans. Realistic MU reductions for clinical practice may lie between these two values and would be based on the individual plan and physician tradeoff preferences. The smaller decreases in MU for *brain2* may be explained by a lack of modulation in the fields for the original cost function, as the plan was relatively simple without involvement of the optic nerves. The SG penalties may not have resulted in a large MU reduction because the original beams already closely resembled the filtered beam. Since the PIMV_q plans were penalized based on the total modulation, they have the greatest MU decrease, along with the relatively flat maximum intensity limited plans.

2. Prostate

In the three prostate cases tested, each technique produced plans that met the high priority dose volume objectives and demonstrated significant reductions in MU. Relative MU are plotted versus net cost in Fig. 6 for all of the prostate cases tested. In *prostate1*, the MU reductions were largest for plans using maximum intensity limits or the PIMV_q penalty. In *prostate2* and *prostate3*, the PIMV_q plans consistently required the fewest MU at a similar net cost and demonstrated

the most predictable and consistent behavior at higher weights compared to the Savitzky-Golay filter-based modulation penalties. All techniques, however, produced acceptable plans with large MU reductions. The main tradeoff for improved delivery efficiency was, again, a slight increase in overall dose to the normal tissues. This can be seen in the *prostate1* DVH comparisons for all methods in Fig. 7(a). All plans shown met the objectives in Table I and have similar values of net cost. The DVHs for the plans with modulation penalties are all very similar, while the maximum intensity limited plan had the highest dose to the rectum along with reduced dose to the femora. This demonstrates the availability of plan tradeoffs possible to achieve the same objective function value. Four of the nine beams from each of the optimized plans are shown in Fig. 7(b). Here we see that the maximum intensity limited beams look very different from the rest of the beams, supporting the large shape differences that were observed in the plan DVHs. The reduction in overall beam complexity for the non-baseline plans is apparent in the intensity maps, and the MU percent reductions compared to the baseline cost function are shown in parentheses. To further illustrate the tradeoff between beam smoothing and normal tissue mean dose, Fig. 8 shows mean doses plotted as a function of relative SMLC MU for each of the normal structures in each prostate case. The PTV mean and minimum doses are also shown. This demonstrates that the minimum dose to the target stays constant as the modulation penalties increase. However, the mean doses to the normal tissues increase as the MU decrease. This is more prominent




FIG. 6. Plots illustrating relative MU versus net cost tradeoffs in the prostate cases when using each of the modulation penalties compared to a baseline cost function and maximum intensity limits. All values are relative to the baseline plan.

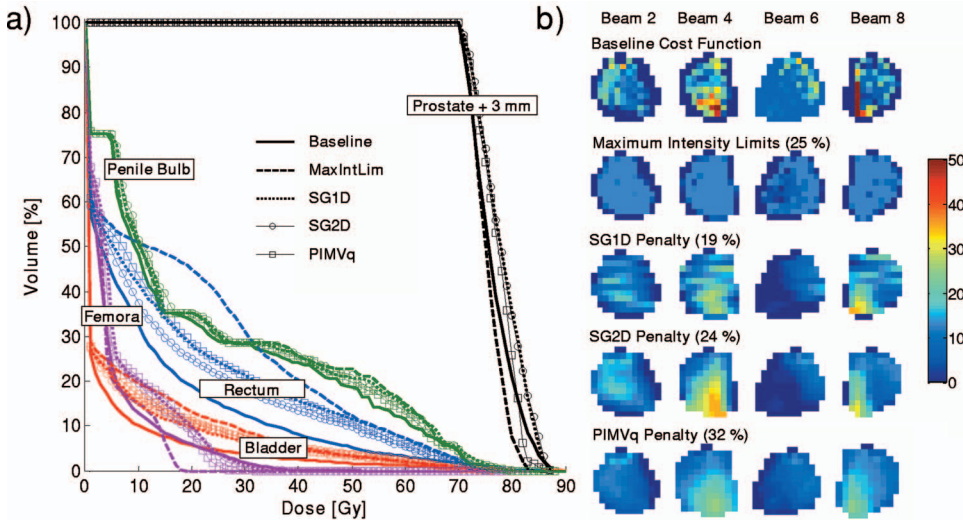


FIG. 7. (a) *Prostate1* DVHs shown optimized with each of the beam modulation reduction methods and (b) corresponding *Prostate1* beams 2, 4, 6, and 8 out of 9 optimized with each of the beam modulation reduction methods. SMLC reductions are in parentheses.

in structures that are close to the target, such as the rectum and bladder. The plots for *prostate1* and *prostate2* also show the inconsistent behavior of the SG penalties at higher weights.

Table III shows the maximum MU reductions possible when applying each of the modulation reduction methods. The first column shows the maximum MU reduction possible while still meeting all of the high priority plan objectives, but allowing for tradeoffs with the lower priority objective of

minimizing overall dose to the normal tissues. The second column shows the MU reduction possible while meeting the high priority objectives and maintaining the mean normal tissue doses achieved in the baseline plan. We have included the DVHs for the *prostate1* plans that correspond to the latter in Fig. 9. Here, we see a slight change in the shape of the rectum DVH, but the rest of the normal tissue and target DVHs are unchanged. Thus, MU reductions occur on the order of 18.9%–26.6% when applying the PIMV_q penalty in

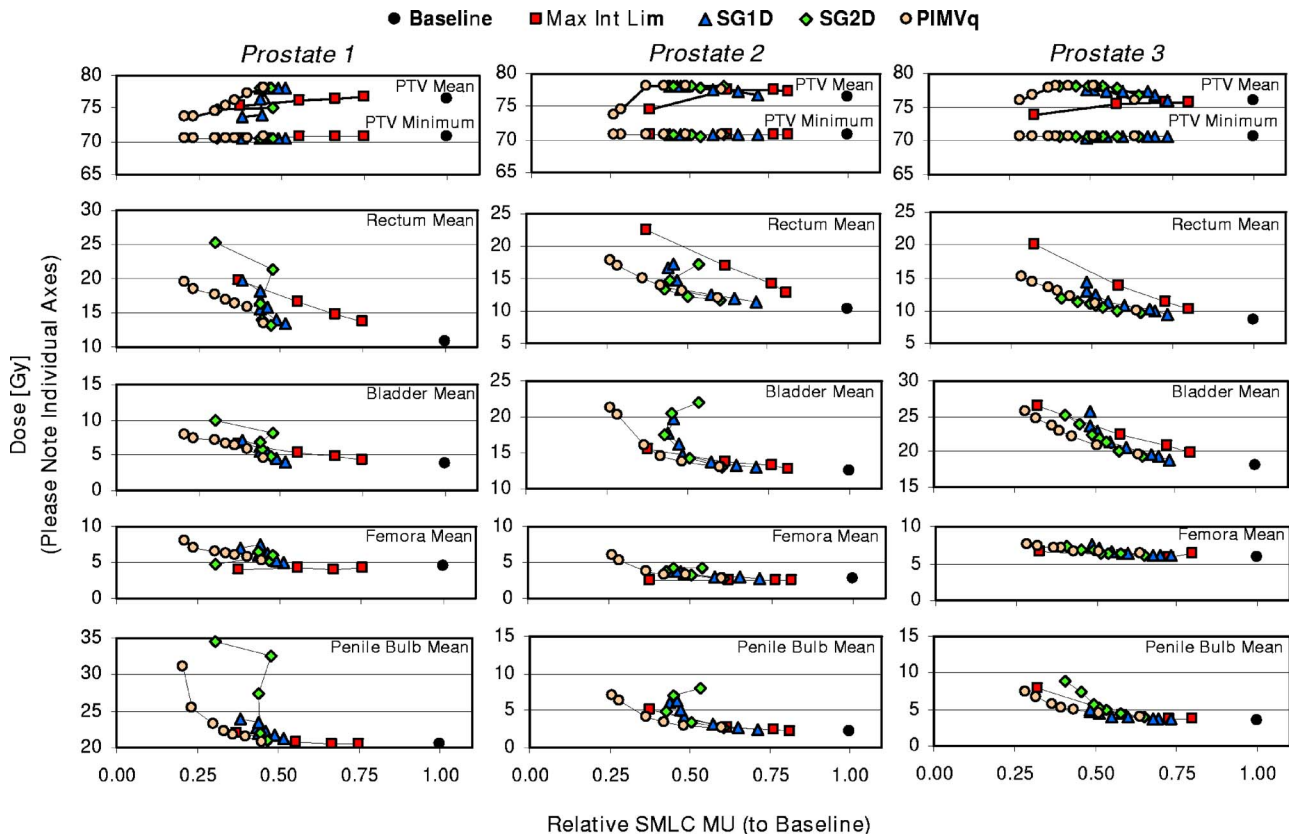


FIG. 8. Doses in each prostate case plotted as a function of SMLC MU when using each of the optimization techniques. Trends in the PTVminimum and mean dose are shown along with normal tissue mean doses.

TABLE III. Maximum MU reductions possible in the prostate using each optimization technique while (i) meeting high priority plan objectives and (ii) maintaining mean normal tissue doses from the baseline plan.

Optimization technique	Prostate1		Prostate2		Prostate3	
	(i) only	(i) and (ii)	(i) only	(i) and (ii)	(i) only	(i) and (ii)
Maximum intensity limits	62.6%	20.9%	62.4%	12.3%	67.6%	12.1%
SG1D filter penalty	61.8%	12.3%	56.1%	2.2%	51.6%	8.8%
SG2D filter penalty	69.6%	26.3%	57.2%	16.2%	59.4%	21.9%
PIMV _q penalty	79.2%	26.6%	71.5%	18.9%	71.7%	26.4%

the prostate with essentially no loss in clinical quality compared to the baseline plans. These MU reductions are actually quite similar to the reductions seen in what appear to be “worse” plans in Fig. 7. This demonstrates one of the disadvantages of using a weighted sum cost function and shows why one must be explicit in the definition of the objectives. If dose to the normal tissues is weakly penalized, as it usually is in a weighted sum cost function so as to not interfere with the more important objectives, the plan will have little incentive to reduce that dose. In Fig. 9, where we include costlets that limit the normal tissue mean doses to those we know were already achievable in the baseline plan, those limits were able to be adhered to while still reducing the MU. We would also like to note that the considerable MU decreases observed in the prostate compared to the brain are likely a function of both baseline plan complexity and the reduced beamlet size (0.5 cm by 0.5 cm as compared to 1 cm by 1 cm) used in the prostate plans.

3. Head/neck

The head/neck body site included the highest number of targets and normal structures in the cost function. The cases

tested had large primary PTVs and nodal volumes, requiring large treatment fields that encompassed substantial (if not all) volumes of many important normal structures. Consequently, there was little room to improve plans based upon the addition of a beam modulation penalty while not interfering with the high priority dose objectives. Therefore, the MU reductions observed (Fig. 10) for the three head/neck plans were smaller and plateaued more quickly than those in the prostate. Still, it was possible to reduce modulation and MU with little loss in plan quality according to the DVHs and dose metrics. As in the other body sites, target coverage was not compromised, due to the design of the cost function. Plans became unacceptable as the modulation penalty grew too large and, in these cases, the first thing to be violated was generally the sparing of one parotid or the spinal cord maximum dose. The PIMV_q method penalizes all modulation, no matter the location; therefore PIMV_q penalized plans have the largest mean spinal cord doses. While a small change may not be clinically important, larger increases in mean cord dose, such as those obtained with higher penalty weights, may warrant a change in the cost function or use of a smoothing mechanism that would not penalize modulation near the spinal cord (to be reported in a separate publication). All other mean doses and DVHs are very consistent throughout the optimization runs, with mean doses varying, in most cases, by less than 1 Gy. The PIMV_q penalty produced the highest MU reductions, followed by the maximum intensity limits, SG2D, and SG1D penalties, with the SG1D penalty producing the most inconsistent MU reductions (see Fig. 10). The highest MU reductions possible while still meeting all of the high priority objectives are shown in Table IV. As in the other sites, it was also possible to achieve substantial MU reductions while maintaining the mean normal tissue doses achieved in the baseline plan (Table IV).

HN2 had additional boost targets within each nodal volume, making it difficult to meet all of the objectives. Many plans had difficulties in satisfying the target dose homogeneity requirements due to the number of overlapping targets. Because of this, the limiting factor in the optimization was not the normal tissue dose. This also explains why the MU reductions were similar when adding in the mean normal tissue dose objectives in Table IV. The performance of the maximum intensity limited plans was poor in this case, demonstrating the need for the higher intensities to reach the homogeneity goals in the targets. The PIMV_q penalty, at higher weights, successfully reduced MU with almost no

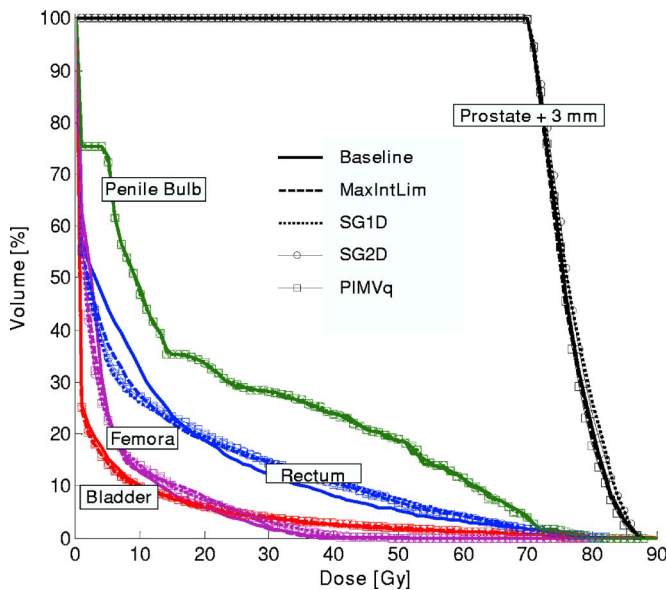


FIG. 9. Prostate1 DVHs shown optimized with each of the beam modulation reduction methods and maximum intensity limits when including mean dose costlets for the normal tissues equal to the mean doses achieved by the baseline plan.




FIG. 10. Plots illustrating relative MU versus net cost tradeoffs in the head/neck cases when using each of the modulation penalties compared to the baseline cost function and maximum intensity limits. All values are relative to the baseline plan.

change in plan DVHs or metrics. In this case there was only a 1.7 Gy increase in mean dose to spinal cord at the highest $PIMV_q$ penalty. This plan represented a 52% reduction in MU as compared to the baseline cost function.

Achieving a mean dose of less than 26 Gy in the left parotid (the only one that could be spared) appeared to be the limiting factor in the optimization in *HN3*. This case shows that sometimes limiting the maximum intensity is not desirable. Figure 11(b) shows that higher intensities were necessary in areas near the left parotid (the lower right of beam 7, for example) to create gradients in the targets to compensate for reducing the parotid dose. In the maximum intensity limited plan, this was not possible and those plans became inferior. The remaining techniques were able to reduce the modulation while still meeting the plan objectives. The fact that the smoothing costlets were still utilizable in this complex geometry demonstrates one advantage of including modulation penalties as weighted costlets versus applying beamlet restrictions in an absolute sense. The $PIMV_q$ penalized plan again had the highest MU reductions, followed by the SG2D and SG1D filter penalties. The latter techniques had slightly reduced normal tissue mean doses compared to the $PIMV_q$ penalty, which can be seen in the DVHs in Fig. 11(a).

IV. DISCUSSION

A high degree of modulation in intensity modulated radiotherapy beams can lead to large increases in treatment time and monitor units (as compared to conventional 3DCRT

plans), and planning and quality assurance time. None of these increases in time and effort is desirable. In this work, we have investigated several techniques for reducing beam modulation, with the goal of not significantly altering IMRT plan quality or the ability of the optimizer to reach the prescribed inverse plan objectives. Each of the modulation penalties studied produced a fairly continuous range of plans and results at varying weights of the modulation penalty. At higher weights, there was a greater tradeoff between the low priority objectives (typically overall normal tissue dose minimization) and a greater decrease in MU. We have also demonstrated that MU reductions on the order of 13%–52% are possible in clinical cases with these complexity-reduction methods, while maintaining the normal tissue mean doses achieved with a baseline cost function. This provides the user with a choice of the degree of normal tissue dose increase acceptable for a certain increase in plan efficiency. Each of the modulation penalties could be a powerful tool for a dosimetrist, physicist, or physician to manipulate an IMRT plan on a case by case basis, depending on the specific goals of the plan.

Applying maximum intensity limits to beamlet plans is a simple solution to reduce MU in plans that do not require high degrees of freedom to meet the plan objectives. Thus, this method was successful in both the brain and prostate body sites. However, in the more complicated head/neck cases, limitations were encountered when certain objectives were on the border of being violated. In these cases, the

TABLE IV. Maximum MU reductions possible in the head/neck using each optimization technique while (i) meeting high priority plan objectives and (ii) maintaining mean normal tissue doses from the baseline plan.

Optimization technique	<i>HN1</i>		<i>HN2</i>		<i>HN3</i>	
	(i) only	(i) and (ii)	(i) only	(i) and (ii)	(i) only	(i) and (ii)
Maximum intensity limits	41.8%	26.9%	40.0%	37.1%	22.6%	18.3%
SG1D filter penalty	28.8%	28.8%	45.3%	36.2%	31.4%	25.7%
SG2D filter penalty	36.5%	34.0%	50.8%	50.8%	37.5%	29.5%
$PIMV_q$ penalty	47.4%	35.6%	52.3%	52.3%	38.5%	33.1%

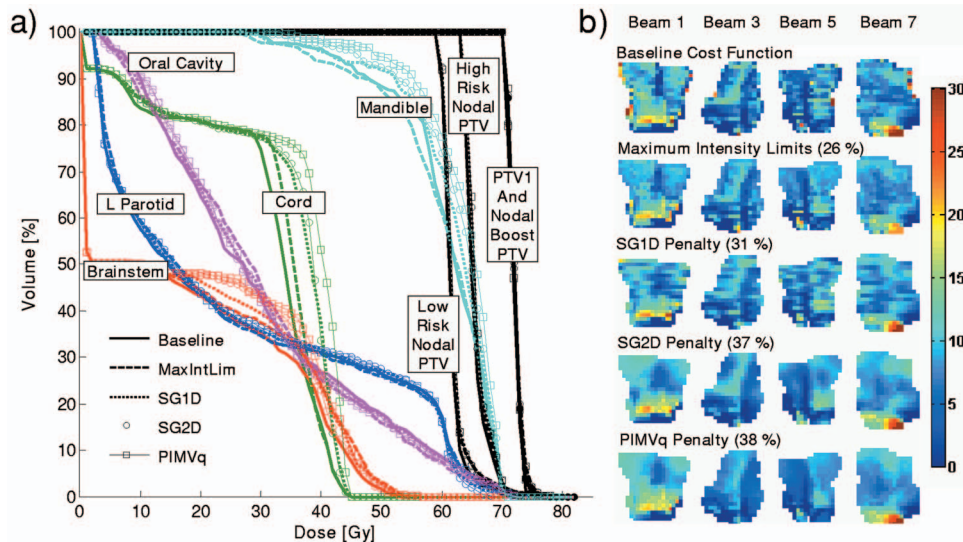


FIG. 11. (a) *HN3* DVHs shown optimized with each of the beam modulation reduction techniques and (b) corresponding *HN3* beams optimized with each of the beam modulation reduction methods. SMLC reductions are in parentheses.

maximum intensity limits interfered with the optimizer's ability to meet the plan objectives, thus requiring a more sophisticated smoothing mechanism.

To allow more degrees of freedom in the quest for reduced beam complexity, the use of terms in the cost function to penalize beam modulation were also investigated. The first penalty was based on a 1-D Savitzky-Golay filter (SG1D penalty) applied only in the direction of the MLC travel, while the second penalty was based on a 2-D filter (SG2D penalty). The third modulation costlet penalized the sum of squared differences between all neighboring beamlets (PIMV_q penalty). Each of these penalties, when applied with a reasonable weight in a weighted sum cost function, successfully reduced plan modulation, while still allowing the optimizer ample freedom to meet the plan objectives. An advantage of using these penalties inside the cost function is that tradeoffs with normal structure doses are controlled by the weight of the modulation penalty. This would be impossible if a smoothing procedure were applied iteratively or postoptimization. The SG1D penalty has been used previously and shown to be effective compared to applying smoothing operations after each optimization iteration.⁷ We have observed that while this method can effectively reduce modulation and MU in most cases, its behavior can be somewhat unpredictable at increased weights of the penalty. Also, it (i) only decreases modulation in one direction, creating a steplike intensity pattern in the direction perpendicular to leaf travel, which is entirely nonintuitive when compared to patient anatomy, (ii) could introduce unwanted errors in QA and delivery, and (iii) could potentially increase a plan's sensitivity to geometric shifts in the direction perpendicular to MLC travel. However, these potentially negative features were useful in cases where smaller structures were present in the field (such as the penile bulb in the prostate and the spinal cord in the head/neck). Still, in the majority of the cases studied here, using the SG2D penalty was superior in terms of overall beam modulation reduction and MU reduction, while producing plans with similar DVHs and dose

metrics. We believe that in most cases, a 2-D modulation penalty should be applied to reduce potential problems with QA and delivery.

Finally, the PIMV_q penalty most consistently reduced MU as the weight was increased. This method eliminates the need for a separate computation of the smoothed filtered beam, as it simply penalizes a direct measure of the overall beam modulation. The behavior of this penalty was very predictable for every plan studied. A potential problem in this method is that it penalizes all modulation, whereas the SG penalties only penalize deviation from a smoothed version of the plan. In the head/neck plans, where the spinal cord frequently lies in the beam, the beam gradient is reduced as a result of the penalty, and the mean dose to the spinal cord is increased. In cases such as this, the Savitzky-Golay methods may be best suited, or changes to the cost function may need to be made when using the PIMV-based penalty. Also, it may be advantageous to examine the use of "smarter" smoothing mechanisms that can distinguish between areas that should or should not be smoothed. Such a method has been investigated by Llacer *et al.*, who showed that for a simple 2-D test case, a space-variant filter can be more effective at smoothing without compromising PTV coverage than other conventional filtering techniques.⁸ We are currently investigating another method based on diffusion principles for 3-D inverse planning that will allow for preferential smoothing based on any desired parameter, such as the beam intensity, gradient, or proximity to organ at risk; these results will be reported separately.²⁶

As mentioned previously, a competing method for improving plan delivery accuracy and efficiency is the optimization of the size and shape of the segments to be used in IMRT delivery. We are currently working on a fair comparison of an in-house developed algorithm, called direct segment optimization (DSO),^{27,28} to fluence map optimization with smoothing and improved leaf sequencing, to determine

which produces the best dosimetric results in complicated cases while aiming to improve overall planning and delivery efficiency.

It is of interest to note that several authors have reported an increase in planning efficiency when incorporating smoothing into the cost function in contrast to smoothing at each iteration.^{7,29} We have also observed this trend and agree that this is most likely due to an increase in curvature of the cost function near the global minimum, reducing the time to fine-tune the plan after convergence.

Another potential advantage to reducing beam complexity is a decrease in delivery time. In fact, we observed delivery time reductions on the order of 25% in the prostate cases studied here, but did not see a significant reduction in delivery time in the other body sites. It may be that the effective leaf travel is not substantially reduced until the beam is almost flat. However, more analysis needs to be done on the delivery time implications of these methods.

One disadvantage of incorporating smoothing into the cost function is a difficulty in analyzing the tradeoffs between the smoothness criteria and the target and normal tissue objectives. However, this problem remains for assessing all weighted-sum cost function plans, and the proposed Pareto surface navigation could be applied to analyze the tradeoffs in question.³⁰ We are also investigating the use of a modulation penalty as a priority stage in a Lexicographic ordering optimization.³¹ Depending on the available solution space, this method could eliminate the need to assess the tradeoffs with the smoothing penalty and other planning objectives. Nevertheless, it currently remains a case-by-case situation to properly incorporate smoothing into a cost function, and this general approach would most likely include iterating through several importance factors for smoothness before reaching the desired solution. The resulting trial and error process may offset the decreased time needed for optimization, but the ultimate increase in plan quality based upon improved QA and delivery efficiency should prove beneficial.

V. CONCLUSIONS

Three different modulation penalties were applied in the inverse planning cost function to reduce the complexity of IMRT beams. After validation in a CT phantom, the penalties were tested in the brain, prostate, and head/neck. Application of each of the penalties in a weighted sum cost function produced plans that met prescribed planning objectives with large decreases in monitor units required for delivery. The large decreases in plan modulation not only improve delivery efficiency, but potentially decrease planning and quality assurance time or difficulty. We believe that each of these methods has merit for improving IMRT plan quality and could be used as a cost function component on a regular basis for the purpose of reducing unnecessary beam modulation. For the greatest benefit in MU reduction, we recommend the use of a 2-D modulation penalty and have found that the plan intensity map variation penalty is best suited for the majority of cases, as it produced the greatest reduction in

monitor units with the least amount of computation needed. However, the Savitzky-Golay 2-D filter penalty may be better at preserving a certain degree of modulation for cases in which critical normal structures lie within the bounds of the beam(s). We expect that the inclusion of these various tools will continue to refine and improve the application and planning of intensity modulated radiation therapy treatments.

ACKNOWLEDGMENTS

The authors would like to acknowledge the work of Daniel McShan, Ph.D., and Kyung-Wook Jee, Ph.D., for their work on our in-house optimization system. In addition, this work was supported in part by NIH Grant No. P01-CA59827.

- ¹M. M. Coselmon, J. M. Moran, J. Radawski, and B. A. Fraass, "Improving IMRT delivery efficiency by applying intensity limits during inverse planning," *Med. Phys.* **32**, 1234–1245 (2005).
- ²X. Sun, P. Xia, and N. Yu, "Effects of the intensity levels and beam map resolutions on static IMRT plans," *Med. Phys.* **31**, 2402–2411 (2004).
- ³J. Seco, P. M. Evans, and S. Webb, "An optimization algorithm that incorporates IMRT delivery constraints," *Phys. Med. Biol.* **47**, 899–915 (2002).
- ⁴L. Ma, "Smoothing intensity-modulated treatment delivery under hardware constraints," *Med. Phys.* **29**, 2937–2295 (2002).
- ⁵R. Mohan, M. Arnfield, S. Tong, Q. Wu, and J. Siebers, "The impact of fluctuations in intensity patterns on the number of monitor units and the quality and accuracy of intensity modulated radiotherapy," *Med. Phys.* **27**, 1226–1237 (2000).
- ⁶X. Sun and P. Xia, "A new smoothing procedure to reduce delivery segments for static MLC-based IMRT planning," *Med. Phys.* **43**, 2785–2794 (2004).
- ⁷S. Spirou, N. Fournier-Bidoz, J. Yang, C. Chui, and C. Ling, "Smoothing intensity-modulated beam profiles to improve the efficiency of delivery," *Med. Phys.* **28**, 2105–2112 (2001).
- ⁸J. Llacer, N. Agazaryan, T. D. Solberg, and C. Promberger, "Degeneracy, frequency response and filtering in IMRT optimization," *Phys. Med. Biol.* **49**, 2853–2880 (2004).
- ⁹B. van Asselen, M. Schwarz, C. van Vliet-Vroegindeweij, J. V. Lebesque, B. J. Mijneer, and E. M. F. Damen, "Intensity-modulation radiotherapy of breast cancer using direct aperture optimization," *Radiother. Oncol.* **79**, 162–169 (2006).
- ¹⁰D. M. Shepard, M. A. Earl, X. A. Li, S. Naqvi, and C. Yu, "Direct aperture optimization: A turnkey solution for step-and-shoot IMRT," *Med. Phys.* **29**, 1007–1018 (2002).
- ¹¹Y. Chen, D. Michalski, C. Houser, and J. M. Galvin, "A deterministic least-squares algorithm for beam weight optimization in conformal radiotherapy," *Phys. Med. Biol.* **47**, 1647–1658 (2002).
- ¹²W. De Gerssem, F. Claus, C. De Wagter, and W. De Neve, "An anatomy-based beam segmentation tool for intensity-modulated radiation therapy and its application to head-and-neck cancer," *Int. J. Radiat. Oncol., Biol., Phys.* **51**, 849–859 (2001).
- ¹³Y. Chen and J. M. Galvin, "A sequential method for aperture-based IMRT treatment planning," *The 14th International Conference of the use of Computers in Radiotherapy*, Seoul, Korea, 2004.
- ¹⁴H. E. Romeijn, R. K. Ahuja, J. F. Dempsey, and A. Kumar, "A rigorous approach to aperture modulation in IMRT treatment planning," *The 14th International Conference on the use of Computers in Radiotherapy* (2004).
- ¹⁵D. Cao, M. A. Earl, S. Luan, and D. M. Shepard, "Continuous intensity map optimization (CIMO): A novel approach to leaf sequencing in step and shoot IMRT," *Med. Phys.* **33**, 859–867 (2006).
- ¹⁶B. A. Fraass and D. L. McShan, "3-D treatment planning: I. Overview of a clinical planning system," in *The Use of Computers in Radiation Therapy*, edited by I. A. D. Bruinvis, F. H. van der Giessen, H. J. van Kleffens, and F. W. Wittkamper (North-Holland, Elsevier Science, Amsterdam, 1987), pp. 273–276.
- ¹⁷B. A. Fraass, D. L. McShan, and K. J. Weeks, "3-D treatment planning:

- III. Complete beam's-eye-view planning capabilities," in *The Use of Computers in Radiation Therapy*, edited by I. A. D. Bruinvis, F. H. van der Giessen, H. J. van Kleffens, and F. W. Wittkamper (North-Holland, Elsevier Science, Amsterdam, 1987), pp. 193–196.
- ¹⁸B. A. Fraass, D. L. McShan, R. K. TenHaken, and K. M. Hutchins, "3-D treatment planning: V. A Fast 3-D photon calculation model," in *The Use of Computers in Radiation Therapy*, edited by I. A. D. Bruinvis, F. H. van der Giessen, H. J. van Kleffens, and F. W. Wittkamper (North-Holland, Elsevier Science, Amsterdam, 1987), pp. 521–525.
- ¹⁹D. L. McShan and B. A. Fraass, "3-D treatment planning: II. Integration of gray scale images and solid surface graphics," in *The Use of Computers in Radiation Therapy*, edited by I. A. D. Bruinvis, F. H. van der Giessen, H. J. van Kleffens, and F. W. Wittkamper (North-Holland, Elsevier Science, Amsterdam, 1987), pp. 41–44.
- ²⁰D. L. McShan, B. A. Fraass, and A. S. Lichter, "Full integration of the beam's eye view concept into computerized treatment planning," *Int. J. Radiat. Oncol., Biol., Phys.* **18**, 1485–1494 (1990).
- ²¹M. L. Kessler, D. L. McShan, M. Epelman, K. A. Vineberg, A. Eisbruch, T. S. Lawrence, and B. A. Fraass, "Costlets: a generalized approach to cost functions for automated optimization of IMRT treatment plans," *Optim. Eng.* **6**, 421–448 (2005).
- ²²J. H. Kim, N. Dogan, D. L. McShan, and M. L. Kessler, "An AVS-based system for optimization of conformal radiotherapy treatment plans," 1995 International Advanced Visual Systems User and Developer Conference, Boston, MA, 1995, pp. 417–423.
- ²³A. Savitzky and M. J. E. Golay, "Smoothing and differentiation of data by simplified least squares procedures," *Anal. Chem.* **36**, 1627–1639 (1964).
- ²⁴T. R. Mackie, J. W. Scrimger, and J. J. Battista, "A convolution method of calculating dose for 15-MV x rays," *Med. Phys.* **12**, 188–196 (1985).
- ²⁵T. R. Bortfeld, D. L. Kahler, T. J. Waldron, and A. L. Boyer, "X-ray field compensation with multileaf collimators," *Int. J. Radiat. Oncol., Biol., Phys.* **28**, 723–730 (1994).
- ²⁶M. M. Matuszak, E. W. Larsen, and B. A. Fraass, "Adaptive diffusion smoothing: A novel method to control IMRT field complexity based on the diffusion equation," *Med. Phys.* **33**, 2201 (2006) (abstract).
- ²⁷B. A. Fraass, D. L. McShan, and M. L. Kessler, "Dose-base conformal field shaping using automated optimization," in *XIIIth International Conference on the Use of Computers in Radiotherapy*, edited by T. Bortfeld and W. Schlegel (Springer, Heidelberg, Germany, 2000), pp. 32–35.
- ²⁸B. A. Fraass, J. M. Moran, K. A. Vineberg, R. Marsh, and D. L. McShan, "Automated optimization of 3DCRT and multisegment IMRT plans," *ES-TRO*, Amsterdam, the Netherlands, 2004 (abstract).
- ²⁹M. Alber and F. Nusslin, "Intensity modulated photon beams subject to a minimal surface smoothing constraint," *Phys. Med. Biol.* **45**, N49–N52 (2000).
- ³⁰D. Craft, T. Halabi, and T. Bortfeld, "Exploration of tradeoffs in intensity-modulated radiotherapy," *Phys. Med. Biol.* **50**, 5857–5868 (2005).
- ³¹K.-W. Jee, D. L. McShan, and B. A. Fraass, "Intuitive multicriteria IMRT optimization using a lexicographic approach," *Med. Phys.* **31**, 1715 (2004) (abstract).