Adaptive diffusion smoothing: A diffusion-based method to reduce IMRT field complexity

Martha M. Matuszak^{a)}

*Departments of Radiation Oncology and 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*

Kyung-Wook Jee, Daniel L. McShan, and Benedick A. Fraass

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

(Received 6 September 2007; revised 10 December 2007; accepted for publication 8 February 2008; published 20 March 2008)

Inverse-planned intensity modulated radiation therapy (IMRT) is often able to achieve complex treatment planning goals that are unattainable with forward three-dimensional (3D) conformal planning. However, the common use of IMRT has introduced several new challenges. The potentially high degree of modulation in IMRT beams risks the loss of some advantages of 3D planning, such as excellent target coverage and high delivery efficiency. Previous attempts to reduce beam complexity by smoothing often result in plan degradation because the smoothing algorithm cannot distinguish between areas of desirable and undesirable modulation. The purpose of this work is to introduce and evaluate adaptive diffusion smoothing (ADS), a novel procedure designed to preferentially reduce IMRT beam complexity. In this method, a discrete diffusion equation is used to smooth IMRT beams using diffusion coefficients, automatically defined for each beamlet, that dictate the degree of smoothing allowed for each beamlet. This yields a method that can distinguish between areas of desirable and undesirable modulation. The ADS method has been incorporated into our optimization system as a weighted cost function penalty, with two diffusion coefficient definitions designed to promote: (1) uniform smoothing everywhere or (2) smoothing based on cost function gradients with respect to the plan beamlet intensities. The ADS method (with both coefficient types) has been tested in a phantom and in two clinical examples (prostate and head/neck). Both types of diffusion coefficients produce plans with reduced modulation and minimal dosimetric impact, but the cost function gradient-based coefficients show more potential for reducing beam modulation without affecting dosimetric plan quality. In summary, adaptive diffusion smoothing is a promising tool for ensuring that only the necessary amount of beam modulation is used, promoting more efficient and accurate IMRT planning, QA, and delivery. © 2008 American Association of Physicists in Medicine. [DOI: [10.1118/1.2889703](https://doi.org/10.1118/1.2889703)]

Key words: IMRT optimization, beam complexity, diffusion smoothing

I. INTRODUCTION

Inverse planned intensity modulated radiation therapy (IMRT) is now a common mode of treatment in radiation oncology centers worldwide. As with many new technologies that show potentially large advantages over previous methods, IMRT was quickly thrust into clinical use before some of its potential disadvantages could be evaluated. Because of this, many centers are now performing retrospective studies and working to suppress difficulties seen in IMRT treatments. One such difficulty, which is intricately tied to the roots of IMRT, is the complexity (high modulation) of IMRT fluence patterns. Looking back at more conventional conformal treatments, it is distressing to see that IMRT has sacrificed, sometimes largely, the good delivery efficiency and geometrically robust target coverage of three-dimensional (3D) conformal therapy. To achieve a balance between all important aspects of radiation therapy treatment,

we seek a more ideal IMRT planning scheme, in which the benefits of optimal normal tissue avoidance are obtained without losing the excellent target coverage and high delivery efficiency achieved with conventional planning. To achieve this balance, it is necessary to intelligently reduce the high degree of modulation often present in IMRT intensity patterns.

Considerable effort has been spent trying to reduce the complexity of IMRT fluence patterns while preserving the advantages gained by employing IMRT. Methods that have been considered include beamlet intensity restrictions,¹ direct optimization of delivery,²⁻⁴ and smoothing of intensity modulated beams either during or after optimization.⁵⁻¹³ As the parameters in each of these methods are relaxed, the unconstrained beamlet solution is approached. The goal is to restrain the parameters enough so that the solution is an acceptable compromise between the ideal beamlet solution and a more efficient conventional 3D conformal solution. Unfor-

unately, these previously described smoothing methods may not yield optimally smoothed plans because the methods cannot distinguish between parts of the beam that should and should not be smoothed.

So far, the most promising smoothing methods are those that penalize or account for modulation inside the inverse planning cost function. When smoothing occurs during optimization, given that the smoothing procedure has a relatively low priority, some complexity reduction is usually possible before causing degradation to the plan. Previously, we have shown that significant increases in efficiency are possible while preserving the full dosimetric quality of the original (unsmoothed) IMRT plan.⁹ However, the difficulty remains that the beam is either penalized for being modulated, which is the fundamental feature of IMRT in the first place, or the beam is penalized for not conforming to a certain filtered version of itself, which may not be ideal for the specific beam in question. The amount of smoothing possible is always limited by the eventual tradeoffs that are made with target or normal tissue coverage, leading us to believe that a more adaptive smoothing procedure that can distinguish between areas that should or should not be smoothed could produce superior results.

A spatially adaptive smoothing method was previously investigated by Llacer *et al.* for a simple two-dimensional (2D) test case, and was shown to compromise PTV coverage to a lesser degree than other conventional filtering techniques.⁵ In their work, beamlets around the PTV were manually selected to receive less filtering in order to preserve the intensity near the target. In the current work, we present a new adaptive diffusion smoothing (ADS) method for preferential smoothing in 3D inverse planning. The ADS method, which makes simple use of the common diffusion equation, allows for preferential smoothing by using variable and automatically defined diffusion coefficients. This method has the potential to distinguish between important and nonimportant areas of modulation, facilitating smarter tradeoffs between the cost function and smoothing criteria. While other methods can be used to reduce MU and modulation in IMRT, adaptive diffusion smoothing is unique because of the ability to customize the diffusion coefficients. In fact, diffusion smoothing schemes have been applied in a variety of other medical physics-related applications, such as image processing^{14–17} and the smoothing of Monte Carlo-derived dose distributions.¹⁸ Broser *et al.* used a diffusion filter to enhance the signal-to-noise ratio in microscopy by using local structural data to control the amount of filtering applied to the image.¹⁶ Miao *et al.* used a diffusion filter to denoise Monte Carlo dose distributions based on the local statistical noise.¹⁸

In the following, we (i) describe the ADS method for controlling IMRT beam complexity, (ii) characterize it using a test phantom optimization case, and (iii) demonstrate its utility and potential in clinical IMRT cases.

II. METHODS

II.A. Adaptive diffusion smoothing formulation and implementation

The time-dependent diffusion equation, given in Eq. (1), is found in many areas of science and engineering. This equation describes the propagation of a material from areas of high concentration to areas of low concentration according to a spatially variant diffusion coefficient, $D(x, y)$, which depends on the local properties of the medium (this equation is cast in 2D for application to the beamlet intensity situation)

$$\begin{aligned} \frac{\partial \phi}{\partial t}(x, y, t) &= \nabla \cdot D(x, y) \nabla \phi(x, y, t) \\ &= \frac{\partial}{\partial x} D(x, y) \frac{\partial \phi}{\partial x}(x, y, t) \\ &\quad + \frac{\partial}{\partial y} D(x, y) \frac{\partial \phi}{\partial y}(x, y, t). \end{aligned} \quad (1)$$

By replacing the density of diffusing material, $\phi(x, y, t)$, with beamlet intensity, $I(x, y, t)$, we can make an analogy from the diffusion of a material to the smoothing of an IMRT beam. Using this idea, we have designed a procedure, named “adaptive diffusion smoothing” or ADS, in which an intensity modulated beamlet grid is smoothed using Eq. (1) with a diffusion coefficient $D(x, y)$ that is automatically defined for each beamlet. This allows the smoothing characteristics to adapt to each individual plan. The diffusion coefficients can be derived from any case related parameter(s), allowing for spatially variant smoothing.

To describe the ADS method, we consider the following time-dependent diffusion problem:

$$\begin{aligned} \frac{\partial I_s}{\partial t}(x, y, t) &= \nabla \cdot D(x, y) \nabla I_s(x, y, t) \\ \text{for } 0 < x < X, \quad 0 < y < Y, \end{aligned} \quad (2)$$

with boundary conditions

$$\begin{aligned} \frac{\partial I_s}{\partial x}(0, y, t) &= \frac{\partial I_s}{\partial x}(X, y, t) = 0 \quad \text{for } 0 < y < Y, \\ \frac{\partial I_s}{\partial y}(x, 0, t) &= \frac{\partial I_s}{\partial y}(x, Y, t) = 0 \quad \text{for } 0 < x < X, \end{aligned} \quad (3a)$$

and initial condition

$$I_s(x, y, 0) = I_0(x, y). \quad (3b)$$

Here $I_0(x, y)$ is the unsmoothed intensity map, $I_s(x, y, t)$ is the smoothed intensity map at time t , and $D(x, y) \geq 0$ is the prescribed diffusion coefficient (discussed in detail below). For $t=0$, $I_s(x, y, 0)=I_0(x, y)$ and no smoothing has taken place. For $D(x, y) > 0$, as t increases, $I_s(x, y, t)$ becomes increasingly smoothed. Also, integrating Eq. (2) over $0 \leq x \leq X$ and $0 \leq y \leq Y$ and using the boundary conditions expressed in Eq. (3a), we easily obtain for all $t \geq 0$,

$$\int_0^Y \int_0^X I_s(x, y, t) dx dy = \int_0^Y \int_0^X I_0(x, y) dx dy. \quad (4)$$

Thus, the total intensity of the unsmoothed beam I_0 is automatically preserved for all $t > 0$.

To estimate the smoothed beam at time $\Delta t > 0$, we integrate Eq. (2) over $0 \leq t \leq \Delta t$ to obtain

$$I_s(x, y, \Delta t) = I_0(x, y) + \nabla \cdot D \nabla \int_0^{\Delta t} I_s(x, y, t) dt. \quad (5)$$

To evaluate the integral term, we consider the implicit and explicit time-differencing approximations. Using the implicit approximation, we obtain the following integral term:

$$I_{s,ij} = \frac{I_{o,ij} + \frac{w}{h^2}(D_{i+1/2,j}I_{s,i+1,j} + D_{i-1/2,j}I_{s,i-1,j} + D_{i,j+1/2}I_{s,i,j+1} + D_{i,j-1/2}I_{s,i,j-1})}{1 + \frac{w}{h^2}(D_{i+1/2,j} + D_{i-1/2,j} + D_{i,j+1/2} + D_{i,j-1/2})}, \quad \begin{aligned} D_{i+1/2,j} &= \frac{2D_{ij}D_{i+1,j}}{D_{ij} + D_{i+1,j}}, 0 \text{ on boundaries,} \\ D_{i,j+1/2} &= \frac{2D_{ij}D_{i,j+1}}{D_{ij} + D_{i,j+1}}, 0 \text{ on boundaries.} \end{aligned} \quad (9)$$

Here h is the dimension of one side of the square beamlet (this solution could be easily altered to account for rectangular beamlets). The diffusion coefficients in Eq. (9) are also defined for each beamlet interface, dictating the amount of smoothing between each beamlet at every side. With this discretization, the smoothed beamlet intensity at “ ij ” depends on the smoothed intensities of its neighboring beamlets, making an iteration scheme necessary to solve for $I_{s,ij}$.

On the other hand, using the explicit approximation, Eq. (7), in Eq. (5), we obtain

$$I_s(x, y) = I_0(x, y) + \nabla \cdot w D \nabla I_0(x, y), \text{ where } w = \Delta t. \quad (10)$$

Applying the same cell-centered discretization to I_o , rather than I_s , we obtain the following discretization of the diffusion problem with explicit time differencing [the D values are the same as Eq. (9)]:

$$\begin{aligned} I_{s,ij} &= I_{o,ij} \left[1 - \frac{w}{h^2}(D_{i+1/2,j} + D_{i-1/2,j} + D_{i,j+1/2} + D_{i,j-1/2}) \right] \\ &+ \frac{w}{h^2}(D_{i+1/2,j}I_{o,i+1,j} + D_{i-1/2,j}I_{o,i-1,j} + D_{i,j+1/2}I_{o,i,j+1} \\ &+ D_{i,j-1/2}I_{o,i,j-1}). \end{aligned} \quad (11)$$

Now the smoothed intensities are based explicitly on the original intensities, making the smoothed intensity pattern much easier to calculate. Also, this solution can easily be applied inside of the optimization procedure as a penalized

$$\int_0^{\Delta t} I_s(x, y, t) dt \approx \int_0^{\Delta t} I_s(x, y, \Delta t) dt = \Delta t I_s(x, y, \Delta t) = \Delta t I_s(x, y). \quad (6)$$

Using the explicit approximation, we obtain instead

$$\int_0^{\Delta t} I_s(x, y, t) dt \approx \int_0^{\Delta t} I_s(x, y, 0) dt = \Delta t I_0(x, y). \quad (7)$$

Introducing the implicit approximation, Eq. (6), into Eq. (5), we get

$$-\nabla \cdot w D(x, y) \nabla I_s(x, y) + I_s(x, y) = I_0(x, y), \quad (8)$$

where $w = \Delta t$. Then, using the classic cell-centered spatial discretization¹⁹ of the diffusion operator $(-\nabla \cdot D \nabla)$, we obtain the following discretization of the diffusion problem:

cost function component to be used with a gradient-based optimization scheme. While it would be possible to use Eq. (9) inside the cost function, the cost function gradient with respect to the original intensities cannot be calculated, making it necessary to employ a nongradient based (and likely much slower) optimization algorithm, such as simulated annealing. For these reasons, we employ the explicit solution to the diffusion equation in the remainder of this work.

The explicit diffusion smoothing scheme has several interesting features (many of which are shared with the implicit solution), including:

- (1) If I_o = uniform intensity, then $I_s = I_o$. Thus, ADS does not alter a flat field.
- (2) If $w/h^2(D_{i+1/2,j} + D_{i-1/2,j} + D_{i,j+1/2} + D_{i,j-1/2}) \leq 1$ for all ij , then $I_{o,\min} \leq I_s(x, y) \leq I_{o,\max}$. Thus, if the first inequality is true, the maximum and minimum intensities of the smoothed beam will lie between the maximum and minimum intensities of the unsmoothed beam. (For the implicit solution, this property is always true.) This property ensures that as long as the original beamlet intensity is positive, the smoothed beamlet intensity will also be positive. In addition, this property gives us a guideline for choosing the absolute values of the time step w and diffusion coefficients. We have chosen to limit D to the range between 0 and 10 for all beamlets, and thus to satisfy the above inequality, and to provide the possibility for a relatively high degree of smoothing,

we chose $w=0.02h^2$, where h is the beamlet dimension (usually 0.5 or 1.0 cm) for the remainder of this work. Limiting D between 0 and 10 allows us to keep w fixed while still providing a wide range of smoothing capabilities.

- (3) $\sum_{\text{all beamlets}} I_o = \sum_{\text{all beamlets}} I_s$ or the smoothed beam preserves the total intensity of the unsmoothed beam. [The discretization scheme preserves Eq. (4).] This property ensures that the smoothed plan will not be drastically different from the original plan and that the contributing beams will be consistent. In some cases, this feature could be considered restrictive, but in the iterative ADS penalty smoothing scheme discussed below, this property is beneficial because it does not interfere by causing a shift in the total intensity from one beam to another.
- (4) Because of the properties of the diffusion operator, the ADS process preferentially suppresses high frequency components of I_o . This is an important feature, because, unlike other smoothing techniques such as polynomial fitting, ADS ensures that the smoothed beam is actually a less modulated version of the original beam. No unwanted high frequency components can arise as artifacts of the ADS smoothing process.
- (5) If $D_{ij}=0$ in beamlet ij , then $I_{s,ij}=I_{o,ij}$; if D_{ij} is large, $I_{s,ij}$ is strongly smoothed. This gives us a basis on which to define D according to the amount of smoothing desired in a certain beamlet.
- (6) The explicitly smoothed $I_{s,ij}$ is determined by the original plan I_o only at the ij beamlet and its four nearest neighbors. The implicit solution smoothes over the entire beam. Thus, in addition to being simpler to calculate, the explicit solution also leads to more local smoothing, which is most likely more desirable for IMRT beams. However, if a more globally smoothed beam is desired, exploration of the implicit method would certainly be warranted.

We believe that the full power of the diffusion smoothing procedure lies in the definition of the diffusion coefficient, D . D can be defined in a multitude of ways, the only constraint being $D_{ij} \geq 0$ for all beamlets. We know from the above properties that if $D_{ij}=0$ for a certain beamlet, then the smoothed intensity of that beamlet will equal the original intensity. Conversely, if D_{ij} is large for a beamlet, then there will be a large amount of smoothing between the original and smoothed beamlet intensities. This gives the user a high degree of control in the amount of smoothing applied over the field and allows for spatially variable amounts of smoothing.

Smoothing procedures that are applied as a cost function penalty appear to have the fewest drawbacks in terms of plan degradation compared to those applied outside the cost function. Therefore, we chose to characterize adaptive diffusion smoothing for use inside the cost function. From the previous paragraph, we can specify D so that modulation is penalized strongly or weakly in given areas of the beam. For example, if we would like to keep sharp edges at the target boundary, we can choose D to have low or zero components near the edges of the beamlet intensity map. Similarly, if we

do not wish to penalize a large intensity gradient over a target/normal tissue interface, we can choose D to be small over those beamlets. Because there are practically no restrictions on the definition of the coefficient values, there are many parameters that can be used to prescribe D . As a first characterization of the ADS method, we will consider two logical ways to define the diffusion coefficients.

Uniform: In IMRT cases with simple geometries, there is generally a large solution space of acceptable plans, and the cost function is not sensitive to small changes in the beamlet intensity patterns. In these situations, the simplest way to define D is as a uniform value across the field. Although this definition of D does not exploit the full potential of defining individual diffusion coefficients (and thus does not distinguish between desirable and undesirable modulations), it provides a good starting point to test the diffusion smoothing scheme and assess whether customized diffusion smoothing coefficients are necessary for simple plans.

Gradient: The ultimate goal of smoothing an IMRT intensity pattern should be to maximally smooth the field with the minimum change in cost function value, which translates to the minimum negative effect on the inverse plan dose prescription. In a gradient-based optimization method, the partial derivatives of the cost function (CF) with respect to each of the beamlet intensities (I_{ij}), $\partial \text{CF} / \partial I_{ij}$, must be calculated at every iteration. Each of these partial derivatives describes how important each beamlet value is to minimizing the cost function value. Beamlets with large absolute values of the gradient have the property that altering those beamlets would have a large effect on the cost function, and vice versa. We have observed that these gradients, at convergence, can vary by several orders of magnitude, with a small percentage of the gradients having very high values, many having very small values and a few having moderate values. With this information, we can define D to be a function of these gradient values to achieve little to no smoothing in high gradient beamlets, moderate smoothing in moderate gradient beamlets, and a high degree of smoothing in low gradient beamlets. We have studied a variety of possible formulations for D , including making D_{ij} inversely proportional to $|\partial \text{CF} / \partial I_{ij}|^n$. In order to achieve the desired distribution of D , we have concluded that the following function is a robust and tunable formulation for D that can be used in most, if not all, IMRT cases:

$$D_{ij} = \frac{10}{1 + a \left(\frac{|\partial \text{CF} / \partial I_{ij}|}{s} \right)^n}. \quad (12)$$

Here s is a gradient scaling factor which we set equal to the median value of $|\partial \text{CF} / \partial I_{ij}|$. (The median value was chosen over the mean because of the tendency for some gradients to have extremely high values and bias the mean toward the higher end.) The parameters a and n are tunable parameters that allow D to be further customized depending on the individual case. It is likely that similar plans will have similar optimal values of a and n . The parameter a serves to shift the

FIG. 1. (a) An optimized brain IMRT beam, (b) the diffusion smoothed beams when using uniform coefficients, (c) absolute value of the cost function gradients with respect to the beamlets, $|\partial CF / \partial I_{ij}|$, (d) adaptive diffusion smoothing coefficients that are defined as a function of (c), and (e) diffusion smoothed beams when using the gradient-based diffusion coefficients in (d). Beams are shown in the beam's eye view in relation to the PTV, brainstem, optic nerves, and optic chiasm. Blue represents low intensity beamlets (gradients, or diffusion coefficients) while red represents high intensity.

transition from high to low D and n controls the steepness of the falloff from high to low D .

In theory, Eq. (12) implies that the highest amount of smoothing will occur in beamlets that have a small effect on the dose prescription goals, and only minor smoothing will occur in beamlets that have a large effect on the cost function. Thus, the use of a variable diffusion coefficient defined in this way yields a method that can distinguish between desirable and undesirable modulations.

As an illustration of each of the above diffusion coefficient definitions and the ADS method itself, Fig. 1(a) shows a standard optimized intensity modulated beam from a brain cancer treatment example with the PTV and several critical structures outlined. Figure 1(b) shows the diffusion smoothed version of that beam using uniform diffusion coefficients across the entire field. The beam was “diffusion smoothed” over five iterations to accentuate the smoothing for illustration purposes. Figures 1(c)–1(e) demonstrate the gradient-based diffusion smoothing process. Figure 1(c) shows the $|\partial CF / \partial I_{ij}|$ values, which are highest in areas of the beam that project onto the critical normal tissues. Figure 1(d) shows the gradient-based diffusion coefficients that are calculated using Eq. (12) with $a=1$ and $n=2$. Finally, Fig. 1(e) shows the diffusion smoothed beam calculated with the diffusion coefficients shown in Fig. 1(d). This beam is also shown after five diffusion smoothing calculations, to highlight the differences between the two methods. We see that when D is uniform, the degree of smoothing is the same over the entire field, smoothing out the modulation near the organs at risk. However, when D is defined using the cost function gradients with respect to the beamlet intensities, the least amount of smoothing occurs in areas where modulation is necessary to meet the plan objectives. In this case, these areas occur around in the overlap regions of the PTV, brainstem, and optic structures. This is shown by the profiles in Fig. 2, which were taken in the direction of travel (shown by the arrow in Fig. 1) for one leaf pair. The base line beam intensity profile has very high and low intensities while the ADS profiles are much smoother. The ADS-gradient profile

preserves the intensity gradients near the brainstem and chiasm overlap regions, while the ADS-uniform profile has a smaller slope. It should be noted that the smoothing that occurs with the gradient-based diffusion coefficients is not always intuitive, as in this beam, since the gradient values depend not only on the normal structure locations but also on the interplay between each of the beams.

Due to the fact that even small changes in the beamlet weights can cause large changes in the cost function, smoothing is usually most successful when applied inside the cost function. Thus, for the remainder of this work, the adaptive diffusion smoothing procedure is used inside the IMRT inverse cost function. To do this, we first calculate the diffusion smoothed beamlet intensities and then calculate the deviation between the original beamlet intensities and the smoothed intensities. This deviation is then penalized as a

FIG. 2. Beam intensity profiles for the base line and ADS plans in the direction of leaf travel for one leaf pair in Fig. 1. The leaf pair is shown by the white arrow in Fig. 1. The beam area that projects onto the brainstem and chiasm is shown by the black arrows.

TABLE I. Test case plan objectives and PDU scale for comparisons to base line plan.

Structure	Goal	Approximate loss in plan quality (relative to base line plan) equal to a PDU of 1.0
PTV	% vol > 57 Gy = 100	0.1 Gy increase to 100% (same as 1 Gy to 1%)
	% vol > 63 Gy = 0	0.1 Gy increase to 100%
OAR1	% vol > 35 Gy = 15	0.1 Gy increase to 85%
	% vol > 63 Gy = 0	0.5 Gy increase to 100%
	Minimize mean dose	1 Gy increase
OAR2	% vol > 20 Gy = 0	0.1 Gy to 100%
	% vol > 63 Gy = 0	0.5 Gy to 100%
	Minimize mean dose	1 Gy increase
Normal tissue	% vol > 63 Gy = 1	1 Gy increase to 99%
	Minimize mean dose	2 Gy increase

part of a weighted sum cost function with weight, p . Specifically, the following adaptive diffusion smoothing penalty is added to the total cost:

$$\text{ADS penalty} = p \times \sum_{\text{all } ij} (I_{o,ij} - I_{s,ij})^2. \quad (13)$$

Here $I_{s,ij}$ is calculated using Eq. (11), and D and w must abide by property 2. As stated previously, D was chosen to fall between 0 and 10 and $w = 0.02h^2$, where h is the beamlet dimension. The individual D values are variables between 0 and 10, depending on the type of adaptive diffusion smoothing coefficients chosen.

II.B. Characterization in CT phantom and clinical examples

To characterize the method and coefficient choices, we applied adaptive diffusion smoothing to a test case with a central spherical target surrounded by two normal structures. The plan consisted of three 6 MV beams with 0.5 cm by 0.5 cm beamlets covering the PTV. The base line inverse plan objectives are shown in Table I. The point density in the structures was adequate to properly sample the region and minimize any point-based artifacts in the beamlet intensities. Beamlet intensity optimization was first performed using the base line cost function with UMOpt,²⁰ our in-house IMRT optimization software package. UMOpt supports a variety of different cost function components, or “costlets”²¹ and allows the implementation of new costlets, such as the adaptive diffusion smoothing penalty. After optimization with the base line cost function, the ADS penalty was added to the cost function at a given weight and the plan was reoptimized. This was repeated with increasing ADS penalty weights to study the consequences of increasing the importance of the ADS smoothness costlet. In these studies, the weight of the ADS penalty was systematically varied to analyze the range of plans and tradeoffs possible for both the uniform and gradient-based diffusion coefficients.

In addition to the comparisons made between base line IMRT and ADS plans, we also optimized plans using the PIMV_q method,⁹ which penalizes the quadratic plan intensity map variation defined by

$$\text{PIMV}_q = \sum_{n=1}^{N_b} \left(\sum_{j=1}^{J-1} \sum_{k=1}^{K-1} \left[(b_{jk} - b_{j,k+1})^2 + (b_{jk} - b_{j+1,k})^2 + \frac{1}{2}(b_{jk} - b_{j+1,k+1})^2 + \frac{1}{2}(b_{jk} - b_{j+1,k-1})^2 \right] \right). \quad (14)$$

Here 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 PIMV_q penalty has been shown to be a simple, yet viable smoothing costlet, and we apply it here to determine whether there is an advantage to using the more sophisticated ADS scheme.

After characterizing the ADS procedure and penalty in the simple test case, we employed the same scheme to clinical examples in the prostate and head/neck. Both cases had seven equispaced 6 MV beams; the base line planning objectives are shown in Tables II and III and reflect in-house IMRT protocols. The normal tissue goals for the prostate are conservatively based on several published toxicity recommendations from RTOG 9406.^{22,23} 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 originally based on the work of Mackie *et al.*,²⁴ but optimized for beamlet calculations. As stated previously, treatment planning was performed with our in-house 3D treatment planning and IMRT optimization software packages, UMPlan and UMOpt.^{20,21,25} Leaf sequencing for static MLC (SMLC) delivery was performed with an in-house-developed leaf se-

TABLE II. Prostate plan objectives and PDU scale for comparisons to base line plan.

Structure	Goal	Approximate loss in plan quality (relative to base line plan) equal to a PDU of 1.0
Prostate +5 mm	% vol > 78 Gy=100	0.1 Gy decrease to 100% (same as 1 Gy to 1%)
	% vol > 88 Gy=0	0.1 Gy increase to 100%
Rectum	% vol > 45 Gy=20	0.1 Gy increase to 80%
	% vol > 85 Gy=0	0.1 Gy increase to 100%
	Minimize mean dose	1 Gy increase
Bladder	<30% gets >45 Gy	0.1 Gy increase to 70%
	% vol > 85 Gy=0	0.1 Gy increase to 100%
	Minimize mean dose	1 Gy increase
Penile bulb	Mean <35	1 Gy increase
	% vol > 85 Gy=0	0.1 Gy increase to 100%
	Minimize mean dose	1 Gy increase
Femurs	Mean <30	1 Gy increase
	<10% gets >40	1 Gy increase to 90%
	% vol > 45 Gy=0	1 Gy increase to 100%
	Minimize mean dose	1 Gy increase
Normal tissue	Maximum <88 Gy	2 Gy
	Minimize mean dose	2 Gy increase

quencer based on a method published by Bortfeld *et al.*²⁶ Delivery sequences allow up to 250 segments per beam, with a goal of achieving a correspondence between planned and delivered intensities of 1%. The ADS penalty was evaluated by comparing the results of the ADS penalized plans to the base line IMRT plan and the PIMV_q plan for each case.

Plan comparisons (without and with varying amounts of ADS smoothing) are described by examining dose-volume histograms, relevant dose metrics, IMRT beam complexity, and delivery efficiency (MU required). However, comparisons of different plans, especially when somewhat different optimization schemes are used, can be difficult, especially since values of the total cost function or individual costlets do not have any specific clinical relevance that can be used to compare the importance of the tradeoffs that are used to achieve the final “optimal” plans. Therefore, we describe here a method for (1) choosing the objective function weights in a clinically relevant manner and (2) judging the quality of inverse plan compromises (or tradeoffs). We call this method “plan degradation units,” or PDU. The goal of the PDU construct is to describe a consistent unit of “tradeoff,” since different kinds of compromises are typically made among the many goals involved in a clinical inverse treatment plan. To facilitate these comparisons, we develop a PDU scale for each case along with the base line objectives in Tables I–III. This scale is developed along with the design of the base line cost function: we assign a concession or sacrifice value for each cost function goal which corresponds to 1 PDU. Each concession is meant to correspond to a consistent level of plan degradation. For example, we could as-

sign a PDU of 1.0 to be a minimum PTV dose of 59.9 Gy instead of 60 Gy, while also defining 1 PDU as 99% of the PTV receiving 60 Gy and 1% of the PTV receiving 59 Gy. For a less important objective, a PDU of 1.0 may correspond to larger dose concession, such as allowing the unspecified mean normal tissue dose to increase by 2 Gy. Once we have designed the PDU scale, we work backward to assign the objective function weights so that the final total cost of the objective function is equal to the PDU value. The total cost in a weighted-sum cost function is given by, $C = \sum w_i f_i$, where w is the weight of the objective f . For one component of the cost function, the cost contribution is simply $w_i f_i$. To normalize all of the objectives to their unit PDU values, we set the individual cost components, $C_i = 1$, and use the objective concession levels given in the table to calculate w . For example, the minimum PTV dose concession in the previous example that was equal to 1 PDU was 0.1 Gy. Therefore, in a quadratic cost function, the weight is equal to 1 divided by 0.1, or 100. Inspection of the PDU scales shown in Tables I–III shows that we have chosen PDU scales that are conservative. In other words, the PDU value is meant to be very sensitive to small changes in the dosimetric goals. This was done purposefully, to show that a significant amount of smoothing is possible in IMRT plans without sacrificing clinical plan quality. In Tables I–III the plan quality metric sacrifices (relative to the achieved base line plan objectives in Gy that correspond to a PDU of 1.0) that were used to assign the cost function weights and judge the final plan quality are given.

TABLE III. Head/neck plan objectives and PDU scale for comparisons to base line plan.

Structure	Goal	Approximate loss in plan quality (relative to base line plan) equal to a PDU of 1.0
PTV70	% vol > 69.3 Gy = 100 % vol > 77 Gy = 0	0.15 Gy decrease to 100% (same as 1 Gy to 1%) 0.2 Gy increase to 100%
PTV64	% vol > 63.4 Gy = 100 % vol > 70.4 Gy = 0	0.15 Gy decrease to 100% 0.2 Gy increase to 100%
PTV60	% vol > 59.4 Gy = 100 % vol > 66 Gy = 0	0.15 Gy decrease to 100% 0.2 Gy increase to 100%
Cord	Maximum < 50 Gy Minimize mean dose	0.1 Gy increase 1 Gy increase
Brainstem	Maximum < 54 Gy Minimize mean dose	0.1 Gy increase 1 Gy increase
Mandible	Maximum < 73.5 Gy Minimize mean dose	0.2 Gy increase 1 Gy increase
Parotids	Maximum < 77 Gy Mean < 26 Gy Minimize mean dose	0.2 Gy increase 0.25 Gy increase 1 Gy increase
Esophagus larynx	Maximum < 50 Gy Minimize mean dose Maximum < 77 Gy	1 Gy increase 0.5 Gy increase 1 Gy increase
Oral cavity	Mean < 49 Gy Minimize mean dose	0.5 Gy increase 1 Gy increase
Normal tissue	Maximum < 73.5 Gy Minimize mean dose	1 Gy increase 2 Gy increase

III. RESULTS

III.A. Phantom study

The adaptive diffusion smoothing procedure was implemented into our 3D treatment optimization system infrastructure as a penalty or costlet to be used inside the cost function of an inverse IMRT plan. A test case was used to test the ADS implementation and characterize the two different methods for setting diffusion coefficients. A base line plan was established for this case by optimizing beamlet weights according to the minimization of the cost functions in Table I. Then, to study the impact of including the ADS penalty at varying weights in the cost function, we reoptimized the base line plan while systematically increasing the ADS penalty weights.

Results of the phantom comparisons are shown in Fig. 3. Figure 3(a) shows the simple three field beam arrangement and anatomy. Figure 3(b) illustrates the tradeoff between MU and dosimetric plan quality as the modulation penalty weights are increased for the PIMV_q, ADS-uniform, and ADS-gradient smoothing penalties. This figure plots the relative MU required to deliver the plan as a function of plan

degradation (PDUs) with respect to the base line plan. We would like to note that the PDU value is equal to the objective function value. Table I contains the PDU scale for the phantom case. A PDU of 1 corresponds to a loss in plan quality relative to the base line plan and could be equivalent to, for example, a 0.1 Gy decrease in the minimum dose to the target, or a 2 Gy increase in the mean dose to the normal tissue, or a lesser combination of the two. As stated previously, we have used these values to assign the objective weights inside the objective function and purposefully made the PDU scale sensitive to changes in plan quality, to show that a large reduction in MU is possible with small losses in plan quality. In Figs. 3(b) and 3(c), the plans optimized with each different method, to similar relative MU values, are compared [denoted by the circled plans in Fig. 3(b)]. Figure 3(c) shows the dose-volume histograms (DVHs) for each of these plans compared to the base line plan, and Fig. 3(d) illustrates the effect of the various smoothing penalties on the beamlet intensity maps. For plans requiring approximately the same MU, the ADS-gradient penalty yielded the highest dosimetric quality plan. With a PDU of 2.0, the dif-

FIG. 3. (a) The phantom geometry and beam arrangement, (b) relative MU (to the base line plan) as a function of plan degradation units, (c) dose-volume histograms for the circled plans in (b), and (d) the corresponding intensity modulated beams in the beam's eye view for the phantom case 1 for standard IMRT and the three different modulation penalties. The outline of the two OARs is shown by the dotted lines. Red corresponds to high intensity beamlets and blue to low intensity beamlets.

ference between the base line plan and ADS-gradient plan is very small, demonstrating the conservative PDU scale. From the beam's eye view in Fig. 3(d), we see a large reduction in overall modulation when using each of the smoothing methods, although the complexity reduction is slightly different for each technique.

The large amounts of smoothing possible in this case may be indicative of the simple geometry and cost function. In fact, the optimal a and n values in Eq. (12) for this case were 0.1 and 2, respectively. While $n=2$ results in a reasonable fall-off from high D to low D , the low value of a shifts this falloff so that it occurs at a fairly high value of D , meaning that the majority of the beamlets were maximally smoothed.

As expected, the steep fluence gradients and modulation near normal tissue interfaces and overlap regions appear to be preserved to a greater extent in the ADS-gradient beams than in the ADS-uniform beams. On the other hand, the PIMV_q penalty promotes an overall reduction of the intensity variation and high intensities in the fields and leads to the distribution of intensity more evenly across the fields. When applying the truly adaptive gradient-based diffusion coeffi-

cients in ADS, there appears to be a higher degree of smoothing in areas that are not in tissue overlap regions. This is expected, because the cost function gradients with respect to the vector of beamlet intensities ($\partial CF / \partial \mathbf{I}$) are likely to be smaller in these regions. Because of this additional smoothing, the ADS-gradient beams may result in lower MU than the ADS-uniform and PIMV_q penalties, even though the ADS-gradient beams have steeper fluence gradients and modulation near the normal tissue interfaces and overlap regions.

These initial results demonstrate that diffusion smoothing is a promising method for controlling modulation in IMRT fields, with an advantage when using the adaptive gradient-based coefficients.

III.B. Clinical examples

To gain more insight into the merits of each of the two different coefficient definition methods, we tested each method for the more complicated clinical geometries of the prostate and head/neck. These cases were optimized with the

FIG. 4. (a) relative MU (to the base line plan) as a function of plan degradation units, (b) dose-volume histograms for the circled plans in (a), and (c) the corresponding intensity modulated beams in the beam's eye view for the prostate case for base line IMRT and the three different modulation penalties. From left to right is base line, PIMV_q, ADS-uniform, and ADS-gradient. The bladder, rectum, and penile bulb are shown in the beam's eye view.

cost functions shown in Tables II and III, and the ADS penalties (using both diffusion coefficient definition methods). As in the phantom case, the base line cost function was first optimized, and then smoothing penalties were added at increasing weights inside the cost function. The weights were systematically increased to observe the tradeoffs made with the base line goals.

III.B.1. Prostate

Figure 4(a) shows results of the prostate base line optimization along with the potential for reducing MU through use of the ADS and PIMV_q penalties, as a function of plan degradation units. The addition of any of the smoothing penalties reduced the MU substantially, although a greater benefit is observed with the ADS-gradient penalty. The maximum reductions in MU were around 26%–36%, although this would be much higher if more degradation was allowed in the plan objectives. It can be seen that the ADS-gradient method achieves its maximum result with smaller tradeoffs (PDUs) than the other penalties. This is because the ADS-gradient penalty does not highly penalize the important areas of modulation. Therefore, this penalty will be nearly zero for high quality plans, while the other methods can still create large penalties because they are penalizing overall modulation. In Fig. 3 we showed that in plans with similar MU requirements, the ADS-gradient plan had the best dosimetric quality. Conversely, in Fig. 4, we compare plans with similar

dosimetric quality. Figure 4(b) shows the DVHs for the base line plan and the other plans at a PDU level of approximately 6. Because of the number of structures in the cost function, this level of plan degradation is very slight, as seen by inspection of the (nearly identical) DVHs. In Fig. 4(c), we display a typical beam intensity distribution for each optimization method, along with the rectum, bladder, and penile bulb in the beam's eye view. Small differences are observed between the ADS-uniform and ADS-gradient beams: the ADS-gradient beams are more uniform in areas where there are no normal tissue overlap regions and allow for quicker falloff between high and low intensity regions. These features combine throughout all of the beams to produce more smoothing and MU reduction at the same level of plan quality. The smoothing that occurs is a function of the interplay between the cost function and many beamlets from seven directions; therefore, one must be careful not to draw significant conclusions from the display of a single beam. For example, the ADS-uniform beam from this left posterior direction is more intense than the other beams, but this is offset by a smaller contribution from the anterior beam in this case. An inspection of all seven beams demonstrates that the PIMV_q beams display an overall flattening of intensity across the beams with the intensity contributions from each beam becoming more uniform. On the other hand, the ADS-gradient plans show a shift in intensity contributions to the beams that intersect the fewest organs at risk. From this point-of-view,

FIG. 5. (a) Relative SMLC monitor units are shown as a function of the plan degradation base line cost for the head/neck plan. (b) DVHs and (c) a typical intensity modulated beam are shown for the circled plan in (a) for each of the optimization methods. From left to right is base line, ADS-gradient, ADS-uniform, and PIMV_q. Several normal structures are outlined in the beam's eye view.

the ADS-gradient beams are much more intuitive than the PIMV_q beams. There were only slight differences observed in the dose distributions in each of the techniques.

In comparison to the optimal a and n values in Eq. (12) for the phantom compared to this prostate example, the n value remained at 2 while the a value increased to an optimum of 0.5. Again, the relative small a value means that the majority of the beamlets are being maximally smoothed, but not quite to the same degree as the phantom case. A preliminary observation is that the a value is an indication of the difficulty of the geometry and/or cost function.

III.B.2. Head/neck

The inverse plan objectives used in the head/neck example were very strict and closely reflect our current clinical standard. Despite the strict cost function and seemingly small amount of solution space to work with, the application of both the ADS and PIMV_q penalties resulted in a substantial reduction of modulation and MU (Fig. 5). This demonstrates that, even in complicated cases, there may still be a large range of plans that can achieve similar results, and that some of those plans may be more desirable in terms of plan efficiency. The use of the ADS (and PIMV_q) smoothing costlets enabled us to find a more efficient plan without sacrificing the quality achieved with the base line plan. In contrast to the phantom and prostate cases, there was no substantial differ-

ence in MU reduction between the different smoothing costlets [Fig. 5(a)]. The ADS-gradient and PIMV_q penalties showed nearly identical results in terms of MU reduction, and the ADS-uniform penalty was slightly worse. Figure 5(c) demonstrates typical qualitative differences between beams smoothed with each of the different methods. These differences are small, which may point to the fact the head/neck plan solution space with the additional objectives of modulation reduction is fairly small. The ADS beams both appear to do a slightly better job at sparing the larynx in the beam shown, although the DVHs are nearly identical for all methods. The optimal a and n values in Eq. (12) may also indicate the complexity of the case itself. The values found for this case were $a=5$ and $n=0.25$ —quite different results than in the previous two cases. Here the optimal diffusion coefficients falloff very slowly from high to low intensity and are focused on a smaller and lower range of values, due to the combination of the low n and high a values.

To demonstrate the ability of the ADS-gradient penalty to preserve intensity gradients as well as to smooth at high penalty weights, the final (most penalized) ADS-gradient plan is shown in Fig. 6. Figure 6(a) illustrates the difference between the base line and ADS-smoothed plans via a 3D visualization of the two plans with several of the important regions of interest displayed. This ADS-gradient plan is noticeably smoother and can be delivered with 57% fewer

FIG. 6. (a) Head/neck geometry and intensity modulated beams for the base line plan and the last (most smoothed) ADS-gradient plan shown in Fig. 6. The ADS-gradient plan can be delivered with 57% fewer SMLC monitor units with minimal effect on the base line plan quality, as shown in the DVHs in (b).

MU. The loss in dosimetric quality is minimal, which can be seen in the DVH comparison of the two plans in Fig. 6(b). The substantial differences noted in the intensity modulated beams between the two techniques suggest that there may be significant differences in the dose distributions of the two plans. Figure 7 shows dose distributions from two transverse slices for the base line and ADS-gradient plans. The dose distributions in the more superior slice [(a) and (b)] are very similar, and the dose distributions in the more inferior slice [(c) and (d)] are slightly different. The anterior and posterior lateral ADS-gradient beams contribute more dose in this plane than in the base line plan, which can be seen in Fig. 6. In addition, the base line isodose lines are slightly less smooth than the ADS-gradient isodose lines.

IV. DISCUSSION

This work has shown that ADS can significantly reduce unneeded modulation in the intensity distributions for IMRT plans. In the phantom and prostate, we are able to reduce MU by approximately 30%–40% with no loss in the plan quality when using ADS, and MU reductions greater than 40% are attainable with only very small concessions in the base line plan. While all smoothing penalties performed well, the adaptive gradient-based diffusion coefficients in the ADS-gradient penalty were able to reduce MU by around

10% more at higher penalty weights in the phantom and prostate cases. This advantage may be due to the fact that the ADS-gradient penalty more appropriately penalizes the less important modulation more than the more important modulation. Therefore, it can preserve—with minimum penalty—the essential modulation in the plan while smoothing large regions of the beam. In the prostate, this results in more uniform areas in each beam that require fewer MU to deliver. Although a large amount of smoothing was achieved in the head/neck case with the ADS-gradient method, significant differences in MU reduction were not seen in this example compared to the $PIMV_q$ penalty. Superior improvement in delivery efficiency with the more sophisticated ADS-gradient method may not have been observed in this case because the cost function gradients were much higher and fluctuated more than in the other sites, due to the large number of clinical goals, and the comparative difficulty of the cost function. In addition, we expect that some uncertainties in the gradients (due to the point sampling, for example) will exist and may lead to undesirable variation in the diffusion coefficients. This is the subject of ongoing study.

This preliminary evaluation of adaptive diffusion smoothing with spatially variant diffusion coefficients (the ADS-gradient penalty) revealed that it has great potential as a tool to reduce IMRT beam complexity in regions where the com-

FIG. 7. Transverse dose distributions corresponding to the base line [(a) and (c)] and ADS-gradient [(b) and (d)] head/neck plans shown in Fig. 6. PTV70, 64, and 60 are shown in yellow, red, and violet, respectively. The parotids glands are contoured in blue and the spinal cord is contoured in green.

plexity is not necessary to produce a quality plan. While more work is needed to evaluate the ADS method, the results of this study show that the use of the ADS penalty does not have to lead to a reduction in dosimetric plan quality and can significantly reduce modulation and MU. In the phantom and prostate examples tested in this work, the use of uniform smoothing methods such as the ADS-uniform penalty or PIMV_q penalty was adequate to reduce beam complexity. However, use of the gradient-based diffusion coefficients in the ADS-gradient penalty more successfully smoothed in areas where modulation was not essential to meeting the plan objectives. Nonetheless, more work on additional clinical cases and body sites is warranted and underway. Additional case studies will also provide guidance on the optimal a and n values that can be used in Eq. (12) to define the gradient diffusion coefficients. Having these parameters is an advantage because it allows us to further customize the method to individual treatment sites, but it can also be a disadvantage if this customization is required for each individual patient. The geometries shown in this work demonstrated a significant range of optimal a and n values, and tuning these pa-

rameters could be time-consuming if it had to be done for each case. Some preliminary testing has suggested that a and n can be fixed for similar geometries, and future work will aim to provide estimated values for a and n in a variety of treatment sites.

In this article, we have not studied smoothing outside of the plan optimization (e.g., within a leaf sequencer), other than in the initial example (Fig. 1), which was used to demonstrate how the smoothing (by itself) works. However, it would be possible to study the application of an ADS algorithm for postoptimization smoothing within the leaf sequencing step. In this kind of study, we would expect the diffusion coefficient definition to have a much greater impact on the final solution. Thus, for centers employing postoptimization smoothing or filtering inside the leaf sequencing process, gradient-based diffusion smoothing could be an attractive option.

One issue with smoothing IMRT plans is the effect it will have on the geometric sensitivity of the plan. This complicated question can only be fully answered through the simulation of a large number of treatment courses for a variety of sites. However, preliminary testing on the prostate case shows that there may be an advantage to using the ADS-gradient method to improve target coverage in the face of setup uncertainty. More detailed analysis on geometric uncertainty, including breathing and organ motion on several patients, will have to be performed to draw any firm conclusions on the robustness of diffusion smoothing to geometric uncertainty.

One of the most exciting results of this study is the large amount of smoothing possible without affecting the quality of the plan. We plan to retrospectively analyze a series of clinical protocol plans to learn how much smoothing would have been possible with the use of the ADS-gradient penalty. In addition, we are currently using the Lexicographic Ordering²⁷ method to quantify the tradeoffs between the modulation penalties and the plan objectives. This will allow our physicians to make educated choices between smoothing and the plan objectives.

On the other hand, for weighted-sum cost functions, the development of a more quantitative way to choose the objective weights and show comparisons between plans using the new idea of PDU has been successful. Many current optimization algorithms rely on the use of conventional weighted-sum cost functions and require a large number of trial-and-error iterations to choose the proper weights for the individual objectives. The addition of a modulation penalty can affect the other objectives in different ways, and evaluating the overall cost after the addition of the modulation penalty can be difficult. The plan degradation unit scale puts a value on different degrees of plan degradation and makes it more intuitive to assign the objective weights and evaluate the tradeoffs that are made when including a modulation penalty. Instead of simply providing a “cost” with no obvious clinical relevance, the PDU value gives a more reliable gauge of the change in plan quality. We believe that the adoption of this methodology will be very useful for judging and designing cost functions for clinical use.

Further applications of diffusion smoothing in IMRT are the use of specialized diffusion coefficients to manipulate beamlet intensities to reduce plan sensitivity to setup errors, organ motion, and even undesirable delivery artifacts such as tongue and groove underdosage.^{28,29} In future work, we plan to use custom diffusion coefficients for adaptive radiation therapy to ensure that large gradients do not occur in areas that may require corrections in an adaptive scheme. This should make it easier to apply feedback during the treatment course, to make fractional changes in the intensity patterns required to correct or change the dose prescription. Thus, the increase of delivery efficiency may be just one of the possible applications of the ADS method.

V. CONCLUSION

The diffusion equation has been used in a procedure that preferentially smooths IMRT plans, using a diffusion coefficient matrix that allows the degree of smoothing to adapt to each individual plan. This procedure was used to define an ADS penalty, applied inside an inverse planning cost function, to promote overall smoothing and monitor unit reduction. Two methods for definition of the diffusion coefficients—to promote uniform smoothing and smoothing based on the beamlet gradients (partial derivatives of the cost function with respect to the beamlet intensity)—were applied and tested on a CT phantom and two clinical examples. Without compromising the base line cost function, MU reductions on the order of 30% and 40% were obtained with the ADS penalties. Compared to the ADS-uniform penalty, the ADS-gradient penalty was better able to preserve intensity gradients and modulation in important areas of the IMRT fields, leading to an advantage in reducing MU in the phantom and prostate cases. This was possible because the gradient-based diffusion coefficients preferentially induce smoothing in the beam where it does not interfere with meeting the dose prescription objectives. All smoothing penalties were equally successful in the head/neck example. Overall, the ADS procedure and penalty is a promising tool for smoothing the unnecessary modulation in IMRT plans which may well have additional important uses due to the possibility of customizing the diffusion coefficients for purposes beyond the reduction of MU.

ACKNOWLEDGMENT

This work was partially supported by NIH Grant No. P01-CA59827.

^{a)}Electronic mail: martha.matuszak@beaumont.edu

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).

²J. L. Bedford and S. Webb, "Constrained segment shapes in direct-aperture optimization for step-and-shoot IMRT," *Med. Phys.* **33**, 944–958 (2006).

³B. A. Fraass and D. L. McShan, "Optimization of conformal and IMRT plans with direct segment optimization," in *XVth International Conference on the Use of Computers in Radiation Therapy*, edited by J.-P. Bissonnette (Novel Digital, Toronto, ON, 2007).

⁴D. M. Shepard, M. A. Earl, X. Li, S. Naqvi, and C. Yu, "Direct aperture optimization: A turnkey solution for step-and-shoot IMRT," *Med. Phys.*

29, 1007–1018 (2002).

⁵J. Llacer, N. Agazaryan, T. D. Solberg, and C. Promberger, "Degeneracy, frequency response and filtering in IMRT optimization," *Phys. Med. Biol.* **49**, 2853–2880 (2004).

⁶J. Markman, D. A. Low, A. W. Beavis, and J. O. Deasy, "Beyond bixels: Generalizing the optimization parameters for intensity modulated radiation therapy," *Med. Phys.* **29**, 2298–2304 (2002).

⁷M. Alber and F. Nusslin, "Intensity modulated photon beams subject to a minimal surface smoothing constraint," *Phys. Med. Biol.* **45**, N49–N52 (2000).

⁸L. Ma, "Smoothing intensity-modulated treatment delivery under hardware constraints," *Med. Phys.* **29**, 2937–2952 (2002).

⁹M. M. Matuszak, E. W. Larsen, and B. A. Fraass, "Reduction of IMRT beam complexity through the use of beam modulation penalties in the objective function," *Med. Phys.* **34**, 507–520 (2007).

¹⁰R. 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).

¹¹S. V. Spirou, N. Fournier-Bidoz, J. Yang, C. S. Chui, and C. C. Ling, "Smoothing intensity-modulated beam profiles to improve the efficiency of delivery," *Med. Phys.* **28**, 2105–2112 (2001).

¹²S. Webb, D. J. Convery, and P. M. Evans, "Inverse planning with constraints to generate smoothed intensity-modulated beams," *Phys. Med. Biol.* **43**, 2785–2794 (1998).

¹³X. Sun and P. Xia, "A new smoothing procedure to reduce delivery segments for static MLC-based IMRT planning," *Med. Phys.* **31**, 1158–1165 (2004).

¹⁴A. A. Samsonov and C. R. Johnson, "Noise-adaptive nonlinear diffusion filtering of MR images with spatially varying noise levels," *Magn. Reson. Med.* **52**, 798–806 (2004).

¹⁵P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion," *IEEE Trans. Pattern Anal. Mach. Intell.* **12**, 629–639 (1990).

¹⁶P. J. Broser, R. Schulte, S. Lang, A. Roth, F. Helmchen, J. Waters, B. Sakmann, and G. Wittum, "Nonlinear anisotropic diffusion filtering of three-dimensional image data from two-photon microscopy," *J. Biomed. Opt.* **9**, 1253–1264 (2004).

¹⁷Z. Ding, J. C. Gore, and A. W. Anderson, "Reduction of noise in diffusion tensor images using anisotropic smoothing," *Magn. Reson. Med.* **53**, 485–490 (2005).

¹⁸B. Miao, R. Jeraj, S. Bao, and T. R. Mackie, "Adaptive anisotropic diffusion filtering of Monte Carlo dose distributions," *Phys. Med. Biol.* **48**, 2767–2781 (2003).

¹⁹R. J. Stamm'ler and M. J. Abbate, *Methods of Steady-State Reactor Physics in Nuclear Design* (Academic, London, 1983).

²⁰J. Kim, N. Dogan, D. L. McShan, and M. L. Kessler, "An AVS-based system for optimization of conformal radiotherapy treatment plans," International Advanced Visual Systems User and Developer Conference, Boston, MA, 1995 (unpublished).

²¹M. L. Kessler, D. L. McShan, M. Epelman, K. 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. M. Michalski, K. Winter, J. A. Purdy, M. Parliament, H. Wong, C. A. Perez, M. Roach, W. Bosch, and J. D. Cox, "Toxicity after three-dimensional radiotherapy for prostate cancer on RTOG 9406 dose Level V," *Int. J. Radiat. Oncol., Biol., Phys.* **62**, 706–713 (2005).

²³M. Roach, K. Winter, J. M. Michalski, J. D. Cox, J. A. Purdy, W. Bosch, X. Lin, and W. S. Shipley, "Penile bulb dose and impotence after three-dimensional conformal radiotherapy for prostate cancer on RTOG 9406: Findings from a prospective, multi-institutional, phase I/II dose-escalation study," *Int. J. Radiat. Oncol., Biol., Phys.* **60**, 1351–1356 (2004).

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).

²⁵B. A. Fraass and D. L. McShan, "3-D treatment planning: I. Overview of a clinical planning system," in *International Conference on the Use of Computers in Radiation Therapy*, edited by I. Bruinvis, F. van der Giesen, H. van Kleffens, and F. Wittkamper (North Holland, Amsterdam, 1987).

²⁶T. R. Bortfeld, D. L. Kahler, and A. Boyer, "X-ray field compensation with multileaf collimators," *Int. J. Radiat. Oncol., Biol., Phys.* **28**, 723–730 (1994).

- ²⁷K.-W. Jee, D. L. McShan, and B. A. Fraass, "Lexicographic ordering: Intuitive multicriteria optimization for IMRT," *Phys. Med. Biol.* **52**, 1845–1861 (2007).
- ²⁸M. M. Matuszak, E. W. Larsen, R. K. Ten Haken, and B. A. Fraass, "Uniform radiobiological targeting with adaptive diffusion smoothing during IMRT optimization," in *XVth International Conference on the Use of Computers in Radiation Therapy*, edited by J.-P. Bissonnette (Novel Digital, Toronto, ON, 2007).
- ²⁹M. M. Matuszak, "Controlling beam complexity in intensity modulated radiation therapy," Ph.D. thesis, University of Michigan, 2007.