Inverse plan optimization accounting for random geometric uncertainties with a multiple instance geometry approximation (MIGA)

D. L. McShan,[a] M. L. Kessler, K. Vineberg, and B. A. Fraass
Department of Radiation Oncology, University of Michigan, Ann Arbor, Michigan 48109-0010

(Received 15 September 2005; revised 12 January 2006; accepted for publication 7 March 2006; published 28 April 2006)

Radiotherapy treatment plans that are optimized to be highly conformal based on a static patient geometry can be degraded by setup errors and/or intratreatment motion, particularly for IMRT plans. To achieve improved plans in the face of geometrical uncertainties, direct simulation of multiple instances of the patient anatomy (to account for setup and/or motion uncertainties) is used within the inverse planning process. This multiple instance geometry approximation (MIGA) method uses two or more instances of the patient anatomy and optimizes a single beam arrangement for all instances concurrently. Each anatomical instance can represent expected extremes or a weighted distribution of geometries. The current implementation supports mapping between instances that include distortions, but this report is limited to the use of rigid body translations/rotations. For inverse planning, the method uses beamlet dose calculations for each instance, with the resulting doses combined using a weighted sum of the results for the multiple instances. Beamlet intensities are then optimized using the inverse planning system based on the cost for the composite dose distribution. MIGA can simulate various types of geometrical uncertainties, including random setup error and intratreatment motion. A limited number of instances are necessary to simulate Gaussian-distributed errors. IMRT plans optimized using MIGA show significantly less degradation in the face of geometrical errors, and are robust to the expected (simulated) motions. Results for a complex head/neck plan involving multiple target volumes and numerous normal structures are significantly improved when the MIGA method of inverse planning is used. Inverse planning using MIGA can lead to significant improvements over the use of simple PTV volume expansions for inclusion of geometrical uncertainties into inverse planning, since it can account for the correlated motions of the entire anatomical representation. The optimized plan results reflect the differing patient geometry situations which can be important near the surface or heterogeneities. For certain clinical situations, the MIGA optimization approach can correct for a significant part of the degradation of the plan caused by the setup uncertainties. [DOI: 10.1118/1.2191016]

Key words: treatment planning, optimization, geometrical uncertainty

I. INTRODUCTION

The goal of radiotherapy treatment planning (RTP) for conformal therapy is to devise a plan that will accurately deliver a tumorcidal dose of radiation to the clinical target volume (CTV) while avoiding or reducing the dose to nearby anatomy. While treatment planning typically is based on a single volumetric imaging study obtained for planning, many workers have reported the occurrence of localization errors in clinical practice,[1–8] while others have confirmed that incorrectly localized patients suffer decreases in a positive outcome.[9–11] According to the International Commission on Radiation Units Report 50 (ICRU 50) recommendations, these expected geometrical uncertainties should be accounted for through the creation of a planning target volume (PTV) that is defined by adding a margin around the CTV.[12] The PTV then becomes the target volume used for planning the high dose region. This recommendation ensures that the CTV will have a high probability of being treated to the desired dose over the course of treatment. Although normal tissue geometrical uncertainties are frequently ignored, uncertainty in the normal tissue position is sometimes accommodated through the use of an expanded region around the organ at risk (OAR). Using the ICRU 62 recommendations, the expanded region is called a planning organ at risk volume (PRV).[13]

Using these methods to account for geometrical uncertainties has worked well for conformal planning, and careful quantitation of these geometric uncertainties has become critical for the successful implementation of 3D treatment planning,[14–19] provided the PTV and organs at risk are well separated. However, when these normal and target tissues are near each other or overlapping, then the absolute behavior of these volumes is problematic, since it does not allow the specification of the tradeoffs between sparing normal tissue and giving the desired dose to the targets. These tradeoffs are crucial, especially when inverse planning driven by optimization techniques is to be used. A second inherent problem (with multiple planning volumes) is that there is no notion of correlated motion of these structures. Often, particularly for neighboring volumes, the motions of nearby structures are well correlated (i.e., the CTV and the OAR move together),

and the use of expanded volumes (PTV and PRV) does not take into account this correlation. These issues become increasingly important for modern conformal therapy, where IMRT is used to create highly conformal dose distributions, particularly for situations where there are concavities in the shape of the target volume(s), or there are normal tissues very close to the target(s). The highly nonuniform dose distributions that are often obtained with IMRT plans lead to another consideration for geometrical uncertainties: it is possible that an IMRT plan may be much more sensitive to geometrical uncertainties than plans developed using flat radiation fields, particularly where intra-treatment motion is involved.[20–25]

There have been a number of studies into the effects of geometric errors on dose distributions. Lujan *et al.* presented the use of a convolution technique that involves taking the three-dimensional dose calculation results for a given beam arrangement and convolving it with a kernel that models the positional probability distribution.[26] That work included two different uncertainty distributions; a Gaussian-shaped distribution for random setup uncertainties as well as a kernel modeling respiratory motion effects on tumors in the liver. The dose-convolution method has also been described by Craig *et al.*[27,28] and others.[29] The accuracy of such approximations has been examined by McCarter using direct simulation of a limited number of fractions,[30] rather than convolving the uncertainty distribution with the dose distribution, Chetty *et al.* showed similar results with a Monte Carlo dose calculation algorithm using a convolution of the uncertainty distribution with the beam fluence.[31] Organ motion has been studied by Booth and Zavgorodni, who utilized Monte Carlo simulations to examine the dosimetric consequences of organ motion.[32] All of the above approaches include the effects of the geometric uncertainties *a priori* in the calculations themselves. Such an incorporation leads naturally to considerations of planning methods aimed at direct coverage of the CTV.[33–35] However, implementation of this very concept within the inverse planning process must be considered for planning that incorporates use of IMRT.

Although there are many reports of the effect of geometrical uncertainties on dose distributions delivered to the patient, there have been only a few reports attempting to fold geometric uncertainties into the inverse planning process. Löf, *et al.* presented a comprehensive theoretical analysis of setup and motion uncertainty issues, and theoretically described the incorporation of this analysis into their optimization methodology.[36] Birkner *et al.* has reported on a more practical method of accounting for motion issues for optimization based on an "effective dose" generated by a sampling of dose lookups for individual voxels or "subvolumes" based on their spatial displacements for a collection of geometry instances.[37] This approach is similar to the MIGA technique, except that in the Birkner work the dose is not recomputed for each geometry instance. Li *et al.* have also demonstrated a technique for incorporating organ motion into an optimization system using a convolution technique to obtain a "motion smoothed" dose distribution.[38]

Optimization of a proposed radiotherapy treatment plan must take into account the effects of geometrical uncertainties, because any planned dose distribution is degraded by various setup errors and intratreatment organ motion that occurs over the course of many treatments. To design a treatment plan that is robust against (or tolerant of) these uncertainties, we make use of a multiple instance geometry approximation (MIGA) that uses a few representative models of the patient geometry to simulate motions that are likely to occur during delivery of multiple treatments. In the present work, random geometrical uncertainties are modeled using a limited number of representative geometries, and only rigid body transforms are used (though future work will exploit the capability of the model to handle distorted datasets and anatomy). For IMRT plans, explicit beamlet dose calculations are done for each instance of the geometry and the resulting calculations are then summed using the expected population density weights, following, in general, the framework proposed by Lof.[36] In this paper we describe details of the MIGA implementation, initial tests to demonstrate the functionality, and then illustrate the method using the complete analysis of a complex IMRT plan for the head/neck.

II. MATERIALS AND METHODS

Implementation of the multiple instance of geometry approximation (MIGA) methods (described in this paper) into our existing treatment planning and optimization system has required relatively few changes. The system consists of UMPlan, the in-house developed planning system that has been used in our clinic for the past 20 years;[39–43] a convolution/superposition dose calculation (based on work by Mackie[44]) module that had been modified to track the dose according to a predefined fluence grid (beamlets)[45] or to multiple segments;[46] and an inverse planning and optimization system UMOpt[47,48] used for beamlet and multisegment IMRT optimization. The MIGA implementation described below includes (1) mechanisms for describing motion experiments, (2) support for multiple geometry instances, (3) methods used for the definition of beam parameters, (4) the ability to perform and track beamlet dose calculations for each geometry instance, and (5) the optimization technique utilizing the MIGA anatomies. Tests on phantom cases and a clinical head/neck case are also described.

A. Motion experiments

For analysis of any single plan with respect to geometric uncertainties, or for optimization of that plan in the face of geometric uncertainties, we have defined the infrastructure to allow the creation of "motion experiments." Each motion experiment describes one particular way of including geometrical uncertainties into the planning, dose calculation, and/or optimization process. For example, the user might specify a multiple-instance MIGA experiment consisting of seven different instances, in order to compare this "experiment" to the result of a dose-convolution version of the plan, or against the static version of the plan.

To allow the user to specify the multiple geometry instances to be used for MIGA, the motion experiment information is described using a simple text file using keyword-value pairs to specify (1) the experiment description, (2) the geometry instance selection (datasets as described below), (3) the relative weights (population density weights), (4) the motion experiment method. One can choose to use (1) the MIGA approach, (2) a convolution of a position probability distribution for random setup uncertainties, or (3) a breathing motion kernel (as described by Lujan[49] to evaluate the effects of respiratory motion).

B. Patient geometry instances

The UMPlan treatment planning system has from its initial design supported the use of multiple anatomical datasets.[39] This feature has been generally used to allow the integration of data from multiple imaging sources,[40,50] but is quite suitable for the description of multiple patient anatomical descriptions (geometry instances) that are used in the MIGA implementation. A 4×4 transformation matrix is used to describe the transformations of each dataset relative to a common reference system (the treatment reference system).[39] One useful enhancement to the existing system allows the definition of duplicate datasets that can then be used to implement rigid body translations/rotations, which is the way that the multiple MIGA datasets are currently used. For each of these duplicate datasets, all anatomical information is duplicated, including associated contours, surfaces, and other information. Once the duplicated dataset is defined, the desired transformations (translation and rotation) operations can be defined using the dataset registration (or deregistration, in this case) options in the planning system.

Within the MIGA infrastructure, it is also possible to handle non-rigid body transformations between datasets. In this case, the individual datasets truly represent different geometry instances that are derived from different imaging studies (e.g., CT scan sets taken on different days or taken at different phases of the respiratory cycle). In these cases, there may still be some rigid body transformation(s) that will place the multiple imaging studies into basic alignment, however, a more complex mapping between different datasets may be used. An example of this mapping might be the use of thin plate splines with control points to map the distortions involved in the lung due to respiration.[51] Since various techniques for this distortion mapping are currently being investigated, the motion experiment infrastructure uses a secondary file that describes the mapping between datasets using a series of corresponding control points sufficient to describe a warping transformation (e.g., using thin plate splines[25] for interpolation). Further study of MIGA with distorted instances will be performed once the distortion mapping research currently underway is completed.

When performing beamlet optimization, the patient anatomical description(s) included in each dataset (multislice contours and derived surfaces) are used to form a voxel-based (volumetric) description of individual regions (or combinations of regions). These regions are then sampled (using random subsampling of a high-resolution voxel grid) to form a list of geometrical points to be used for dose calculation and evaluation. For MIGA motion experiments, the same set of points is applied to the different geometry instances (datasets) using the geometrical mapping transforms.

An important issue for MIGA (or any method handling geometric uncertainties) is the appropriate definition of the target volume(s). Normally, clinicians define the gross tumor volume(s) (GTVs) and/or clinical target volume(s) (CTVs) that describe the regions to be irradiated. Conventionally, one would add a margin to this volume to form a planning target volume (PTV) that then represents a volume that will ensure adequate coverage of the CTV, despite the geometrical uncertainties. As we are now attempting to account for this motion more explicitly, the target volume that should be used for planning with MIGA is not the conventional PTV volume, but something more like the CTV. How to accurately define this CTV, and which errors and uncertainties are to be contained within the CTV, are beyond the scope of this work. Here, we will take the CTV(s) to be the volume(s) to which we wish to deliver the treatment prescription(s), while including geometrical uncertainties due to random setup errors and/or intratreatment motion using the MIGA concept.

C. Plan definition and beam setup

Typically IMRT planning begins with a preselected set of beams that are placed around the patient (1) uniformly in an axial plane, or (2) with directions determined interactively based on traditional "forward planning" experience. Each beam's aperture(s) is typically formed so that the PTV(s) are covered. Using a Beam's Eye View (BEV) display is the best way to ensure geometrically that there will be adequate target coverage and that normal tissues irradiation will be minimized. In UMPlan, it is possible to display within the BEV not only the nominal reference anatomy, but also the anatomical information from the other shifted or warped datasets. This allows the planner to enlarge jaw and MLC boundaries to account for the expected (as modeled by the MIGA datasets) motions of the clinical target volume(s). For IMRT optimization, a beamlet grid is defined (resolution typically 1×1 cm or 0.5×0.5 cm (defined at the source to axis distance). These grid elements are initially set to match the rectangular jaw shape but those beamlets that fall outside of the BEV target projection are normally defined to receive an intensity of 0 and are excluded from the beamlet intensity optimization.

D. Dose calculation methods

The calculation method used for normal (non-MIGA) beamlet IMRT planning utilizes a convolution superposition (CVSP) algorithm.[44] The algorithm has been modified to track the dose to each point per unit fluence through each beamlet.[45] In effect, the dose contribution of each beamlet (a rectangular cross section of the primary beam fluence or pencil beam) is calculated. Note that the convolution calculation propagates dose through dose spread arrays that are corrected for density changes between the primary interaction site and

the dose deposition sites. For inverse planning purposes, the dose is computed to the set of points defined within each region of interest, and so the dose per point per unit beamlet intensity is computed for each beamlet.

For the multiple geometry instances used here, the dose calculation procedure is exactly the same, except that the calculations are repeated for each of the different geometry instances. The geometrical shifts between the different datasets affect both the density grid used for heterogeneity corrections and the positions of the calculation points relative to the beam. For each dataset, the same beamlet description is used, and the calculations are again computed to the same (but shifted or warped) set of points for each region. The following equation is used to determine the effective unit beamlet dose per calculation point:

$$\langle D(j,k)\rangle = \sum_{i=1}^{n} W_i \times D(i,j,k), \tag{1}$$

where the effective dose $\langle D(j,k)\rangle$ represents the unit intensity motion-weighted beamlet dose for point j and beamlet k. This dose represents the weighted sum of the beamlet doses $D(i,j,k)$ for each geometrical instances i. The weight W_i for each instance represents the integral probability for a particular geometrical instance normalized so that $\Sigma_{i=1}^{n} W_i = 1.0$.

E. Optimization

For normal (non-MIGA) optimization, individual beamlet doses are retrieved, and the weights of each beamlet are optimized using cost function scoring that is based on the sum of the dose to each point from all beamlets. The UMOpt system uses a flexible cost function methodology that allows a wide range of functions to be used to assign a cost for dose to individual regions.[47] The costs (an individual region can have multiple costs or "costlets") for each of the different regions are combined to form a total cost for the plan. Simulated annealing, quasi-Newton, or pre-emptive lexicographic ordering search engines are used to drive the weight adjustments based on the changes in the total cost (or energy) of the system.

For MIGA optimization, the calculational results from each of the datasets are retrieved, and the beamlet doses (per unit intensity) for each point are summed using the probability weights from the motion experiment file. Different probability weightings (specified in different experiment control files) can be assessed without having to recalculate the point doses. The resulting point doses, which are the motion-averaged dose to each point per unit intensity, are then used to determine the beamlet weights that result in the optimal cost function result, just as is done in the static planning case. Dose volume histograms (DVHs) are calculated using the dose points contained in each structure, just as in the non-MIGA case. The DVH analyzes the dose delivered to each point over the patient treatment course as approximated by the MIGA weighted multiple instances.

F. Test phantoms and experiments

For the initial testing of the MIGA concepts, a cylindrical phantom 30 cm diameter was defined with a centered target sphere of 6 cm diameter, both with unit density. A normal cord-like object of 1 cm diameter and 10 cm length was placed near the target. Multiple datasets were then defined to delineate a series of shifts along various axes to document the correct behavior of the various software features involved in the MIGA method. For the tests described here, 13 datasets were defined with offsets in a direction parallel to the centers of the target and cord structures with offsets of 0, and ±2, 4, 6, 8, 10, and 15 mm. Three nonopposing beams were then used for the initial testing, including testing of various fields using beamlet grids of 2×2 mm, 5×5 mm, and 10×10 mm. The beamlets were defined to cover the full range of the motion of the target. Figure 1 shows the geometry, beam arrangement, and beamlet patterns for the 5×5 mm plan.

To demonstrate the results obtained with the MIGA-based optimization, and to assess the sensitivity of the results to the number of instances needed to approximate a Gaussian-distributed uncertainty distribution of setup errors, three different population density distributions, each represented by a Gaussian with a full width half max (FWHM) of 4, 8, and 12 mm in a direction perpendicular to the beam direction were evaluated. For each Gaussian distribution, we optimized initially identical plans, based on using 1, 3, or all 13 instances. For the three instance experiment, the nominal dataset was included in the mix using a weight of 0.5 for the unshifted dataset and 0.25 for the shifted datasets. The 13 datasets were weighted according to the Gaussian probabilities for each of the displacements. Figure 2 illustrates the weights used for the different number of instances used for the 12 mm FWHM experiments.

Each plan was optimized using each of the geometry experiments. The cost functions used for these tests had as goals a target dose of 100 with a tolerance of ±2% with a powered dose difference penalty (the plan was penalized if the dose dropped below 98% or above 102% and the penalty was the dose difference below or above the tolerance window, to the power of six) summed and averaged for all points within the region. To reduce the dose outside of the target region, the surrounding region was assigned a penalty using a threshold function that penalized any doses above zero using the dose difference. The cordlike structure was penalized with a threshold function starting at a relative dose level of 60% and the difference above was penalized to a power of 5. All plans were then evaluated on the 13 instance dataset that more accurately represents the full Gaussian distribution (i.e., the "truth" for this experiment).

G. Clinical example

To demonstrate the use of the MIGA technique in a realistic clinical situation, a MIGA optimization for a clinical head/neck IMRT plan is compared to a static IMRT plan. The planning protocol followed for this comparison is used for a

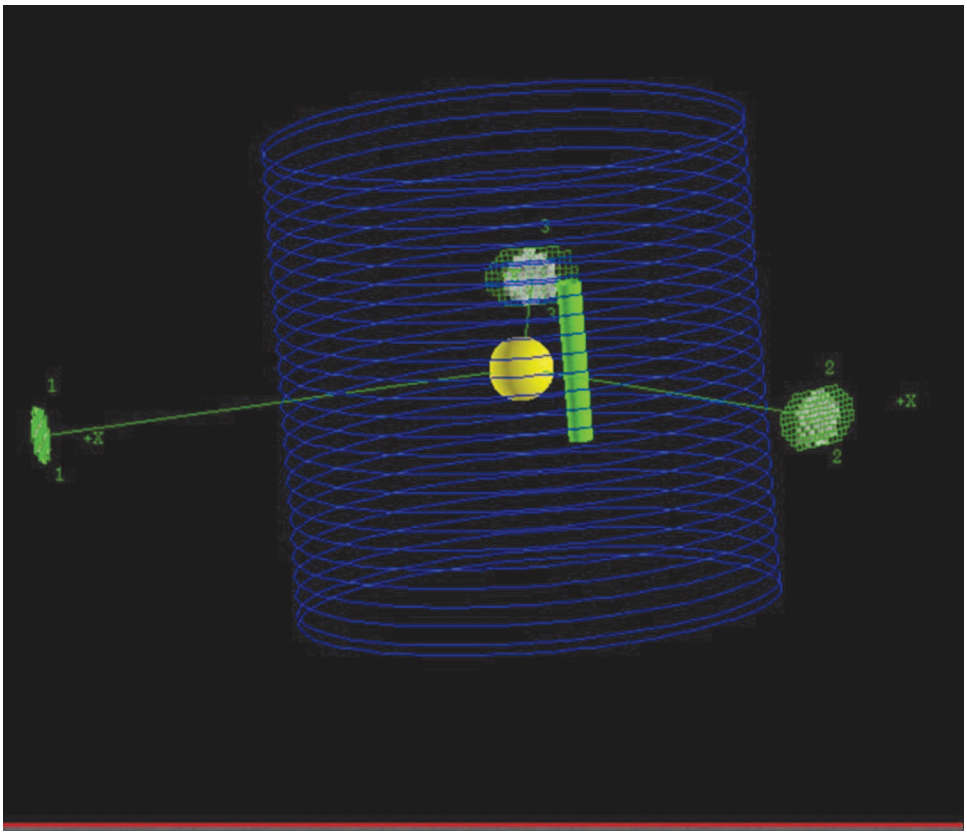

FIG. 1. Test phantom with three non-opposed IMRT beams. Beamlets per beam are constrained to the potential motion displacement, which is ±15 mm in the direction between the target (yellow) and the cord (green).

clinical IMRT study[53,54] and is similar to the requirements of the RTOG H0022 Head/Neck IMRT study.[55] For this comparison, target volume doses of 66, 60, and 54 Gy were prescribed for the high risk volumes, high risk nodes, and contralateral nodal volumes, respectively, as shown in Fig. 3. The cost functions used for the two optimization trials are listed in Table I, and require a uniform (±5%) dose to each of the targets, dose <45 Gy to the spinal cord, attempt to reduce the mean dose to the parotids below 26 Gy, and minimize all other normal tissue doses. For this simple example,

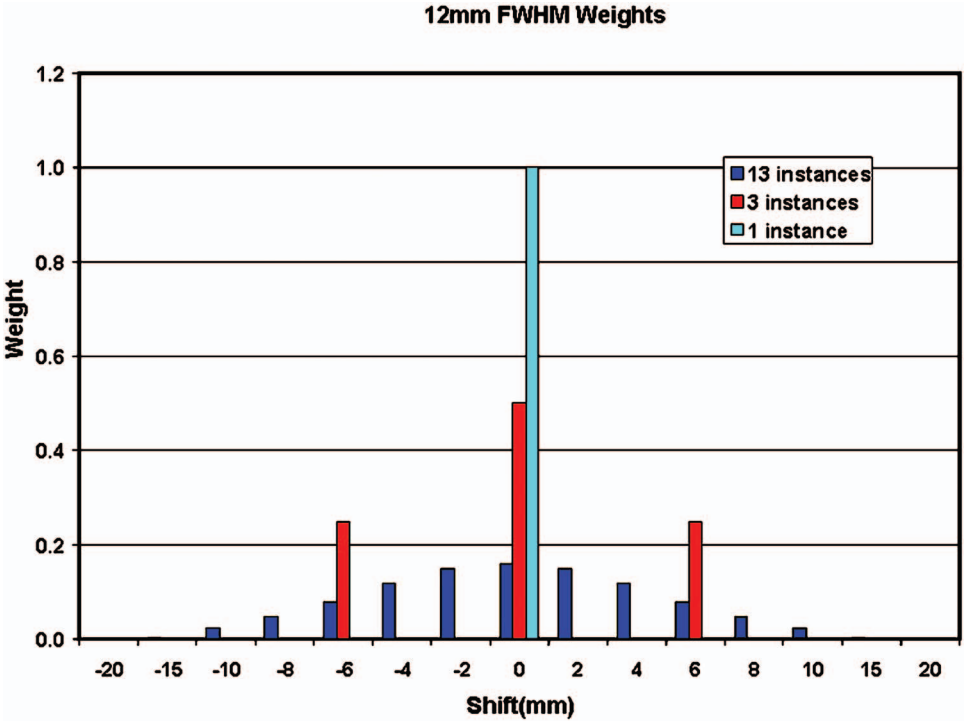

FIG. 2. Dataset weights for 1- (default), 3- $(0, \pm 6$ mm), and 13- $(0, \pm 2, \pm 4, \pm 6, \pm 8, \pm 10, \pm 15$ mm) instance distributions for the 12 mm FWHM uncertainty distribution.

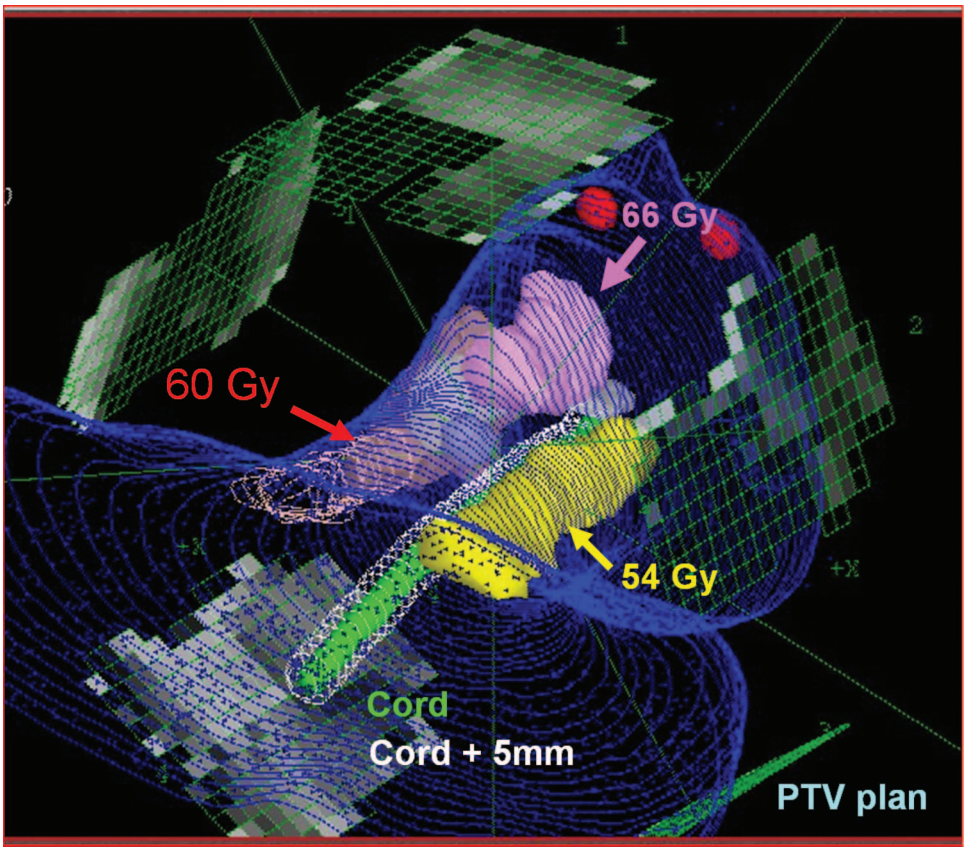

FIG. 3. Clinical example of a five axial field head/neck IMRT plan: Targets (pink, red, yellow) get doses of 66, 60, ande 54 Gy. Cord (green), Cord +5 mm (white), larynx (white), parotids (yellow), eyes (red), and oral cavity (gray) are also shown.

five equispaced axial beams of 1×1 cm beamlets were used for the static and MIGA IMRT plans (Fig. 3).

The MIGA dataset was constructed by replicating the original CT dataset six times and shifting the datasets in each of the $\pm x$, y, and z directions. The full MIGA anatomy thus contained 7 datasets (the original with weight 0.5, and the 6 shifted datasets, each with a weight of 0.0833), as shown schematically in Fig. 4. The shifts used for each direction were determined by the setup uncertainties measured for similar head/neck patients in our clinic,[56] and in this case were ±0.42 cm (ant-post), ±0.75 cm (sup/inf), and ±0.45 cm (left/right). While the MIGA plan was optimized to give the appropriate dose to the CTV's drawn by the physician for each of the targets, the static plan was optimized to deliver the desired dose to a planning target volume (PTV) created by anisotropic 3-D expansions of the individual CTVs using margins of 0.42 cm (ant/post), 0.75 cm (sup/inf), and 0.45 cm (lat), which is our normal approach to account for this setup error distribution. Since the PTV created for some CTVs extended outside the surface of the patient (a common problem for inverse planning in the head/neck), modified PTVs were created by editing the PTVs back to the patient surface ("PTV-s"), or 5 mm inside the patient surface ("PTV-5"). For this simple clinical example, the "MIGA," "PTV," "PTV-s," and "PTV-5" plans were all evaluated on the seven-instance MIGA dataset. Results of the dose volume histograms obtained for these static and MIGA plans, using the seven-instance MIGA dataset, are analyzed.

TABLE I. Prioritized goals for head and neck parotid sparing protocol.

Cord	Max < 45 Gy
Cord+5 mm	Max < 50 Gy
Primary GTV, boost CTVs	66 Gy
Ipsilateral nodes	60 Gy
Contralateral and retropharyngeal nodes	54 Gy
Target homogeneity	±5%
Parotiods	Mean dose < 26 Gy
All normal tissue	Minimize

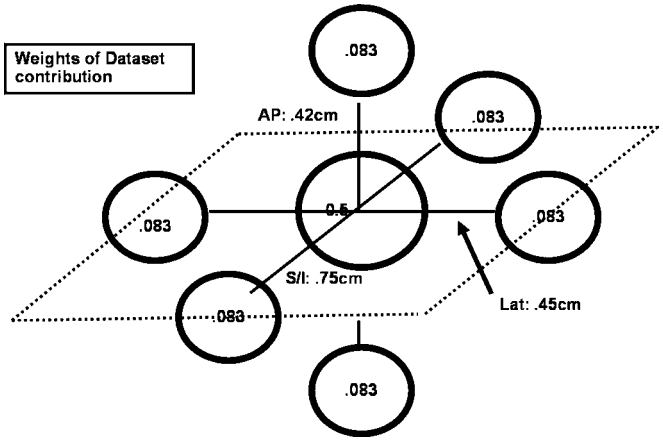

FIG. 4. The 7-instance MIGA dataset, with weights 0.5 (center) and 0.083 (outside edges), and shifts (from the nominal dataset) of ±0.42 cm (ant/post), ±0.75 cm (sup/inf), ±0.45 cm (left/right).

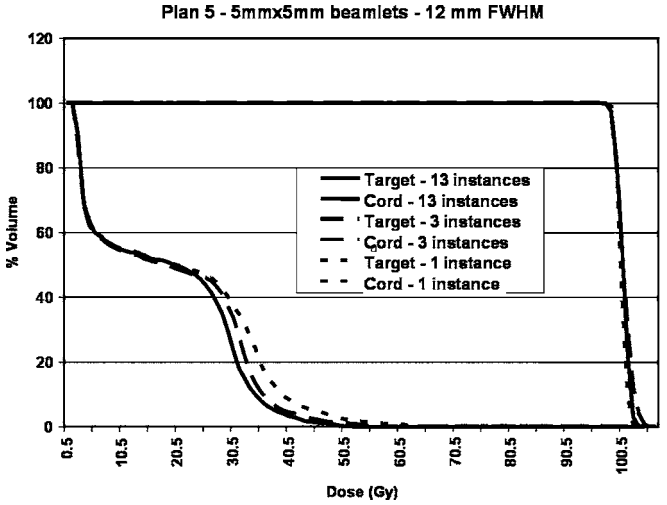

FIG. 5. Dose volume histogram for test phantom case with 3 nonopposing beams optimized with 1- (no motion), 3-, and 15-instance datasets (evaluated using the 15 Gaussian datasets) for a motion compromised geometry with FWHM of 8 mm.

III. RESULTS

A. Test geometries

The results of the simple test geometry experiments [Fig. 2 shows the geometry setup for the nominal dataset (no shift)] are illustrated in Fig. 5. The range of beamlets for the lateral fields account for the maximum expected excursion of ±1.5 cm for the experiments analyzed. As one would expect, optimizing with the multiple instance geometry approximation results in the beamlet weights being increased along the direction of motion (between the two instances) to ensure dose coverage of the target. Figure 5 demonstrates the dose volume histogram (DVH) results for an anticipated random uncertainty in the positioning along the direction between the target and the cord-like object. The results are shown for the static (1-instance) versus 3-instance, and 13-instance datasets, assuming a Gaussian distribution with a FWHM of 12 mm. For the static case, a PTV was used and was expanded by ±6 mm in the direction of the anticipated motion. This comparison shows how closely the static and 3-instance MIGA plans agree with the 13-instance MIGA plan, which in this case represents the best that could be achieved given the particular random motion assumption. Figure 5 demonstrates that the single instance example achieves adequate target coverage through the use of the PTV definition. However, the cord DVH for the PTV plan shows a higher dose than resulted with any of the MIGA optimized plans. The 3-instance DVH matches closely with the 13-instance DVH, indicating that a limited number of instances can be used to approximate the assumed Gaussian motion.

B. Clinical head/neck example results

Dose volume histograms (DVHs) for the targets and normal tissues for this example case are shown in Figs. 6(a), 6(b), and 6(c). Target volume comparisons are shown in Fig. 6(a), which shows the PTV plan and MIGA plan doses to the

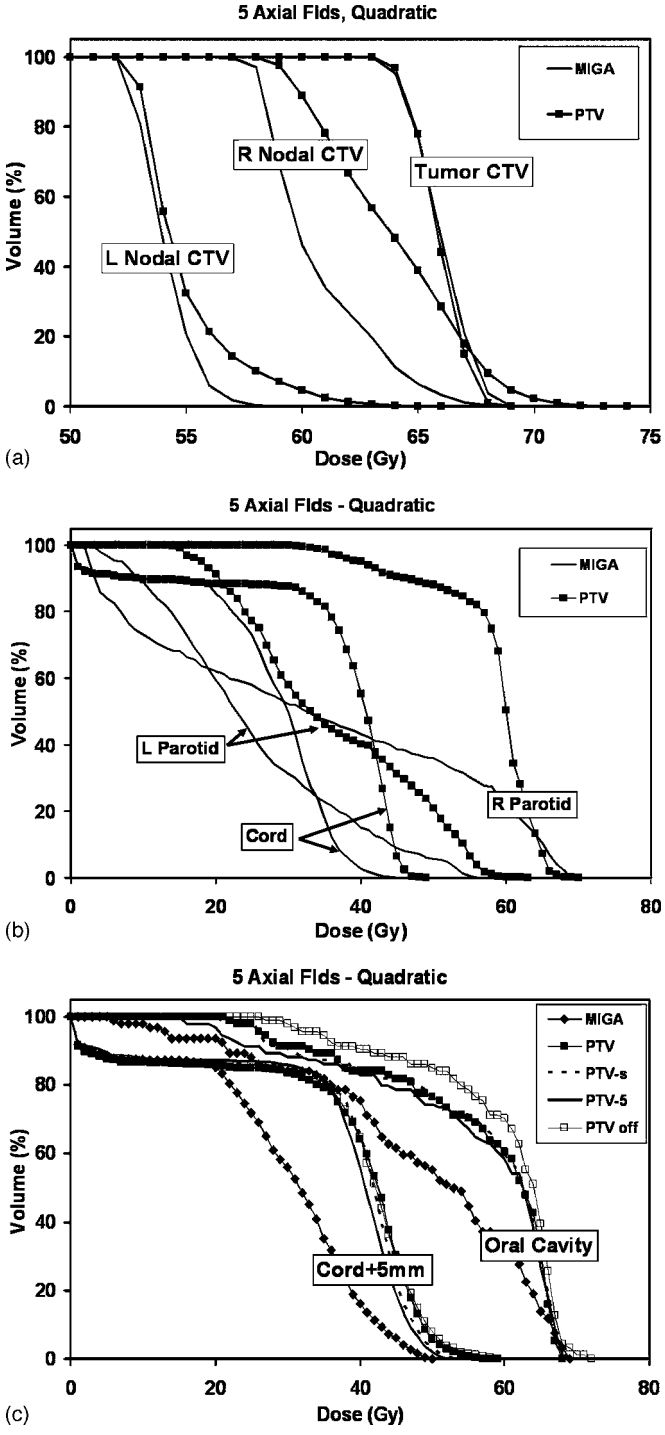

FIG. 6. (a) DVHs for the three targets (CTV66, CTV60, and CTV 54). (b) DVHs for normal tissues (R and L parotid and spinal cord). (c) DVHs for normal tissues (cord+5 mm, and oral cavity). Results also shown for PTV-s for surface limited PTV and PTV-5 for PTV restricted to be 5 mm inside surface.

CTVs (CTV66, CTV60, CTV54). The CTV doses are compared since the goal of both the MIGA optimization method and the PTV expansion is to assure that the dose delivered to the CTVs is the prescribed dose, even in the case of setup uncertainty. The plots in Fig. 6(a) show that both the PTV and MIGA plans correctly cover all the CTVs with the de-

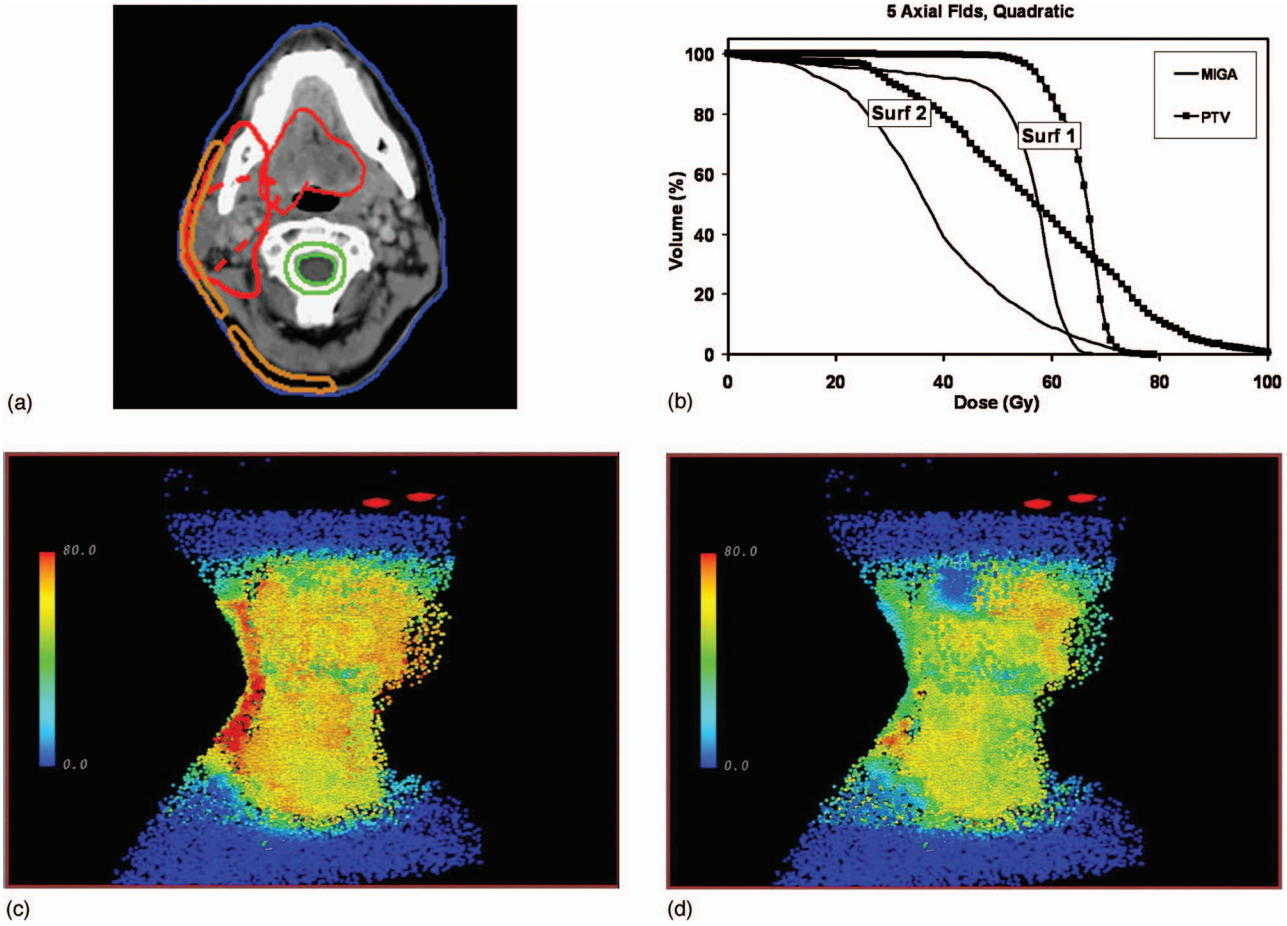

FIG. 7. (a) Two regions highlighting the dose to the skin were delineated—one immediately next to the target volumes, and another not directly next to targets. (b) DVHs for the PTV and MIGA plans to the two surface regions. (c) Lateral view of dose distribution (all calculated points) due to the PTV plan. The color of each point is defined by the dose using the displayed color table. (d) Lateral view of dose distribution due to MIGA plan. Note the decreased dose along the lateral surface of neck, as well as the improved sparing of the parotid.

sired doses, as they should. Note, however, that the MIGA DVHs do a better job of keeping the dose to the lower dose CTVs down closer to their prescribed doses, rather than the PTV plans, which must give the lower dose CTVs more dose to cover the PTVs for the high dose targets (since those PTVs overlap the lower dose CTVs). Figures 6(b) and 6(c) show the DVHs for the normal tissues, illustrating that the MIGA method significantly decreases the dose delivered to all normal tissues when compared to the PTV plans.

One of the regions that is significantly affected by setup error is the buildup region and skin, especially where beams tangentially irradiate the skin. One finds that there are often large intensities in the beamlets that tangentially intersect the patient surface. In contrast to many of the usual tricks used to handle target regions very near the surface, including the addition of virtual bolus, resetting the depth of buildup region calculation points to the depth of d_{max} and others, MIGA plans can do accurate calculations to the buildup region points, since correct geometry is used for the dose calculations for each dataset instance in the MIGA plans. MIGA plans, since they take into account the potential setup differences, do not have the same high beamlet fluences at the patient surface. This is shown by looking at the DVHs for two 3-D structures [Fig. 7(a)] defined to document the dose to the buildup region (1) for Surf-1, the region near the target volumes; and (2) for Surf-2, a surface region away from the target volumes. The DVHs for these two structures [Fig. 7(b)] show that while the dose near the surface from the PTV plan might be as large as 100 Gy (when evaluated over the MIGA 7-instance dataset), the MIGA plan can keep these doses lower, as desired to minimize patient skin reactions. The overall improvement in dose distribution, especially on the lateral aspect of the neck, can be seen by comparing Figs. 7(c) and 7(d), which show the dose to each point calculated within the patient, for the PTV and MIGA plans respectively. The MIGA plan clearly shows lower doses along the lateral aspect of the neck, and also clearly shows decreased dose to the parotid.

One more interesting point about the MIGA plan solutions is illustrated by Fig. 8. The MIGA solutions are not simply a blurred out version of the PTV plan. Rather, the optimization finds a different solution when the MIGA method is used. This is illustrated by comparing the fluence distributions for beam 1 of the PTV and MIGA plans, where the fluences are quite different. In this (anecdotal) example, the type of inten-

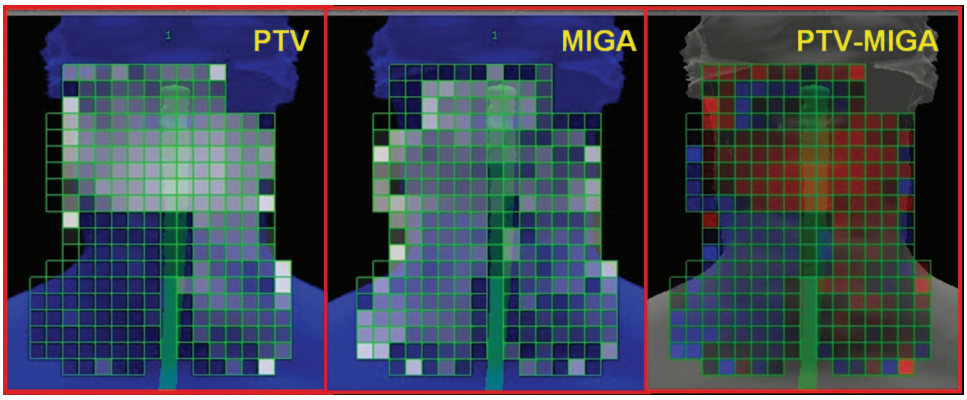

FIG. 8. Example intensity distributions for Beam 1: PTV plan, MIGA plan, difference (PTV plan—MIGA plan).

sity patterns chosen by the optimization were clearly different, an effect which could benefit from further study.

IV. DISCUSSION

In recent years, a great deal of effort has been spent delineating the scope of various kinds of geometrical errors in radiation therapy treatment planning and delivery. It is clear that clinical 3-D conformal therapy treatment, with its dose distributions tightly conformed to the target volume(s), may be highly sensitive to both systematic and random setup errors, as well as to intratreatment motion. The use of IMRT treatment delivery, with highly modulated intensity distributions within each radiation field, may be even more sensitive to these geometric errors. Therefore, an understanding of the various types of patient motion, and optimization of the plan so that the degradation of the planned treatment due to such motions is minimized, is critical for radiotherapy treatment planning.

Fractionated radiation treatment involves patient localization errors (setup errors) introduced by the imprecision of day to day patient setup, physiological changes, and also by intratreatment motion throughout the course of each fraction. The only accommodation that the typical treatment planning process makes for these "motions" is the use of simple margins around the target volume (e.g., the PTV designed to assure that the CTV is always included in the high dose region, even as the various geometric motions occur). As mentioned before, this simple margin-based approach runs into difficulties whenever the expanded volumes [either PTV or PRV (Planning Risk Volume)] overlap, since any overlap of the volumes makes it impossible to fulfill the goal of the expanded volumes without requiring some sort of tradeoff to be made. This is particularly crucial when optimization techniques are used, as in the case of inverse planning for IMRT. In inverse planning, the inverse plan's cost function determines how best to deliver the desired dose to the target(s) while at the same time minimizing the dose to normal tissues, rather than simply using the PTV margin to assure that the target volume will be treated, regardless of the implications for surrounding normal tissues. In order to create a treatment plan that is not degraded by the setup errors and motion that occur during treatment, the plan must be created and optimized while incorporating the expected setup error and motion distributions into the plan optimization process.

The MIGA method thus creates a model of the patient geometry that can incorporate the geometrical setup error and intratreatment motion expected for the patient. This description is formed by the use of multiple instances of the patient geometry, weighted to best represent the expected distribution of patient geometries that will be encountered during the treatment course. The instances can be simple duplicates of the original CT dataset, translated to represent the random setup error for the patient, or they can be independent CT datasets acquired at different times, or for different phases of the respiratory cycle, for example. While in this work we have considered only simple translations of duplicate datasets, representing random setup errors, the MIGA infrastructure can be used for much more complex representations of the situation. These more complex representations will be studied in more detail in forthcoming publications.

In the simple phantom experiments, all of the different multiple instance representations showed improved results for the MIGA-based plans over the use of a simple static plan. As expected, the use of a large number of instances provides a more accurate simulation of the potential motions than using a fewer number of instances. However, since the calculation time and computer resources necessary for the calculation depend on the number of instances, it is appropriate to try to find the minimal number of instances which may be representative of the actual expected distribution of geometries. Figure 5 demonstrated that the three-instance dataset (nominal plus two extremes) did a reasonable job of representing the full Gaussian uncertainty distribution, so we have used this kind of MIGA dataset (seven instances when used in 3-D) for our initial clinical examples. However, it is clear that more work is necessary to determine the best representations to be used for clinically relevant problems.

Another issue still to be addressed in future work involves the size of the beamlets and how that interacts with the range of motion involved in the Gaussian uncertainty distribution, including specifically the distance between different instances. For large beamlet grids and small uncertainties, it is expected that it will be difficult for the optimization to respond to the small geometric uncertainties, while small beamlets and larger distances between the instances may lead

to quite different response of the optimization to the incorporation of the multiple instances into the optimization process (including possibly aliasing). A reasonable rule-of-thumb based on stochastic sampling theory might suggest that it will be difficult to correct for motions smaller than half the beamlet grid width. In general, there are many issues related to the small number of instances that we have used in these initial examples, and future work will address a number of these points. One issue of particular interest is the importance of the tails of the uncertainty distribution: how important is the incorporation of an instance that represents the two or three sigma deviations of the distribution?

For the initial examples presented in this work, the assumption was made that the uncertainty distributions are Gaussian. While this may be reasonable for most situations, it is easy to think of situations where it is not the appropriate choice. Breathing causes longer dwell times at the peaks of inspirations and expiration cycles. A patient laying in a cradle may have a bimodal state depending on comfort or other factors (clothing, cradle design,…). The MIGA method may simulate any kind of distribution—it is certainly not limited to Gaussians. If these distributions are known, then they could potentially be approximated within MIGA.

The MIGA approach offers a number of advantages over other approximations that use motion blurring of the dose distributions. These blurring techniques do not take into account changes in geometry shape and do not properly account for changes in depth (and distance from source) or proximity to heterogeneities. While still approximate, the MIGA approach at least provides accurate calculations for the individual geometry instances. For deep-seated tumor locations, this advantage may not always be clinically important, but for tumors near the patient's skin or for sites where there is significant motion, the accuracy of the calculated doses could potentially make a crucial difference in the overall clinical relevance of the MIGA-based optimized plan. The accuracy of the MIGA-based planning process can also be improved if more detailed simulations (more instances) of the expected clinical situation are used.

The clinical relevance of a MIGA-based plan optimization is of course dependent on how well the simulated geometrical uncertainty distribution correlates with what actually happens during the patient's treatments. For example, if there is an uncorrected systematic shift of the patient's setup location vis a vis the central point of the MIGA instance distribution, the MIGA plan will not be a good representation of the treatment to be delivered. Of course, this is true for any prospective estimate of the geometrical setup accuracy for a specific patient, and in fact this is the main justification for the current interest in adaptive radiotherapy[57] paradigms that adjust the treatment geometry repeatedly as more individualized data on uncertainties for the particular patient are obtained using daily measurements at each treatment fraction.

While the MIGA approach provides an improvement over dose blurring techniques (to account for motion) by providing explicit dose calculation for a sampling of setup offsets, it does generally make the initial assumption that there is no systematic offset and only models random setup or motion errors. However, with daily quantitative assessment of the actual patient setup, it is possible to use the MIGA methodology to first generate a "dose-to-date" estimate of the treatments that have been given to date. One can then incorporate systematic offsets to skew the instance weights or assume that the offsets are removed, and use MIGA to reoptimize the balance of the treatment. It is clear that care is needed in clearly defining the experiments to which the MIGA method is applied.

Using the multiple instance approximation (MIGA) to simulate the distribution of geometries involved in a patient's treatment course highlights a basic limitation to most current treatment planning and optimization approaches. For most forward and inverse planning, the goals of the treatment are specified using the total dose to be delivered over the patient's treatment course, either as a total dose (e.g., 70 Gy), or a relative dose distribution (e.g., prescribe to the 95% isodose surface). When considering MIGA, or any other method of dealing with setup uncertainty, the use of only the total dose to be delivered is a problem. For example, imagine that a 2-instance MIGA simulation is performed, and that it turns out that the way that the MIGA optimization achieves a uniform dose of 70 Gy to the target volume is to give very unevenly distributed doses to a given voxel (e.g.. 50% higher doses in one instance, and 50% lower dose in the other). The total dose in all voxels may be 70 Gy, but when the patient is in one position, the daily dose to the target is significantly overdosed, while in the other position, it is underdosed. Clearly, to believe that the MIGA optimization is going to be clinically useful in this case, the cost function for the optimization should be modified to either (1) force dose/fraction constraints as well as the total dose costs, or (2) specifically include dose/fraction corrections into the optimization through the use of linear/quadratic model estimates of a bioeffect.[58] Although the MIGA simulation may not be an explicit representation of the actual fractionation of the patient (i.e., one instance for each fraction), the bioeffect modeling still needs to include the fractionation effects in an explicit way. This is a new level of planning complexity that has been investigated very little.[18,59] Clearly this is a consideration when the dose to a particular region will see different levels of dose (and different dose rates) from day to day. While applying the bioeffect modeling (converting dose to BED (biologically effective dose)) is conceptually easy, the incorporation of this technique may dramatically increase the resource requirements and slow the fractionation-corrected optimization process.

V. CONCLUSIONS

A novel optimization technique that incorporates the simulation of setup uncertainty and/or patient intratreatment motion has been developed using a multiple instance of geometry approximation (MIGA) method. The MIGA-based optimization approach simulates the expected distribution of patient geometries (location, anatomy, positioning), and then optimizes a single IMRT plan using a weighted sum of the

behavior over all the simulated geometry instances. While in this first description of the MIGA technique the various instances are all constructed from the initial planning scan set by duplicating the CT scan set, and then translating each new dataset to approximate a Gaussian distribution of random setup errors, the implementation can easily support the use of distinct measured CT scan sets (for example, different parts of the respiratory cycle), as well as other methods of creating geometry instances that demonstrate the distortions that occur due to organ motions, tumor growth, weight loss, etc.

A number of simple examples have demonstrated that it is possible to use a limited number of instances (typically at least three in each direction) while still achieving a good representation of a Gaussian setup uncertainty distribution. This same kind of representation (total of seven instances, in three dimensions), has been used to demonstrate significant improvements in a head and neck IMRT plan over what would typically occur if PTV expansions of the various CTVs were used to compensate for the setup uncertainty. While there are a large number of issues to be investigated further, both to optimize the use of the method, as well as determining the best way to make use of the method for specific clinical protocols or situations, the initial results with the method are quite promising. A number of clinical and methodological studies are underway to further investigate use of the MIGA optimization method described in this work.

ACKNOWLEDGMENTS

We would like to thank many colleagues who have helped with this study, including Anna Samuelsson Back, Robin Marsh, Edie Krueger, Martha Coselmon, Jeff Radawski, and many others. This work has been supported in part by National Institutes of Health Grant No. 2 P01 CA59827.

[a]Corresponding author and reprint requests: Daniel L. McShan, Ph.D. Department of Radiation Oncology, University of Michigan, UH-B2C432, Box 0010, 1500 E. Medical Ctr. Dr. Ann Ar bor, Michigan 48109. Telephone: (734)-936-4290; fax: (734)-936-7859; electronic mail: dlmcshan@umich.edu

[1]R. W. Byhardt, J. D. Cox, A. Hornburg, and G. Liermann, Int. J. Radiat. Oncol., Biol., Phys. **4**, 881–887 (1978).
[2]L. Coia, J. Chu, R. Larsen, and R. Myerson, Int. J. Radiat. Oncol., Biol., Phys. **12**, 1697–1705 (1986).
[3]W. J. Graham, D. W. Anderson, D. J. Landry, and C. R. Bogardus, Int. J. Radiat. Oncol., Biol., Phys. **7**, 1117–1119 (1981).
[4]S. E. Griffiths, R. G. Pearcey, and J. Thorogood, Int. J. Radiat. Oncol., Biol., Phys. **13**, 1583–1588 (1987).
[5]A. G. Haus and J. E. Marks, Invest. Radiol. **8**, 384–391 (1973).
[6]P. K. I. Kartha, A. Chung-Bin, T. Watchor, and F. R. Hendrickson, Med. Phys. **2**, 331–332 (1975).
[7]P. K. I. Kartha, A. Chung-Bin, T. Wachtor, and F. R. Hendrickson, Int. J. Radiat. Oncol., Biol., Phys. **2**, 797–799 (1977).
[8]I. Rabinowitz, J. Broomberg, M. Goitein, K. McCarthy, and J. Leong, Int. J. Radiat. Oncol., Biol., Phys. **11**, 1857–1867 (1985).
[9]A. L. Boyer, in *Radiation Oncology Physics-1986*, Medical Physics Monograph No. 15, edited by H. R. E. Kereiakes and C. G. B. J. G. Kereiakes (American Institute of Physics, New York, 1987), p. 438.
[10]L. L. Doss, in *Proceedings of Workshop on Quality Control in the Radiotherapy Department; Vol. Cancer and Leukemia Group B*, edited by L. Reinstein, New York, 1979, pp. 57–62.
[11]J. J. Kinzie, G. E. Hanks, C. J. Maclean, and A. Kramer, Cancer **52**, 2223–2226 (1983).
[12]ICRU Report 50: Prescribing, Recording, and Reporting Photon Beam Therapy, edited by ICRU (International Committee on Radiation Units and Measurements, Bethesda, MD, 1993).
[13]ICRU, ICRU Report 62: Prescribing, Recording, and Reporting Pfoton Beam Therapy (supplement to ICRU Report 50), Vol. 62, Bethesda, MD, 1999.
[14]A. Dutreix, Int. J. Radiat. Oncol., Biol., Phys. **13**, 1291–1296 (1987).
[15]M. Goitein, Med. Phys. **12**, 608–612 (1985).
[16]G. Kutcher, G. Mageras, and S. Leibel, Semin. Radiat. Oncol. **5**, 134–145 (1995).
[17]Suit HD, D. Phil, J. Becht, J. Leong, M. Stracher, W. C. Wood, L. Verhey, and M. Goitein, Int. J. Radiat. Oncol., Biol., Phys. **14**, 777–786 (1988).
[18]R. K. Ten Haken, in *ESTRO, 1997*, edited by T. Bortfeld, B. Mijnheer, and A. Nahum (1997), p. 8.
[19]M. M. Urie, M. Goitein, K. Doppke, J. G. Kutcher, T. LoSasso, R. Mohan, J. E. Munzenrider, M. Sontag, and J. W. Wong, Int. J. Radiat. Oncol., Biol., Phys. **21**, 91–107 (1991).
[20]T. Bortfeld, S. B. Jiang, and E. Rietzel, Semin. Radiat. Oncol. **14**, 41–51 (2004).
[21]T. Bortfeld, K. Jokivarsi, M. Goitein, J. Kung, and S. B. Jiang, Phys. Med. Biol. **47**, 2203-20 (2002).
[22]C. S. Chui, E. Yorke, and L. Hong, Med. Phys. **30**, 1736–1746 (2003).
[23]R. George, P. J. Keall, V. R. Kini, S. S. Vedam, J. V. Siebers, Q. Wu, M. H. Lauterbach, D. W. Arthur, and R. Mohan, Med. Phys. **30**, 552–562 (2003).
[24]G. D. Hugo, N. Agazaryan, and T. D. Solberg, Med. Phys. **30**, 1052–1066 (2003).
[25]S. B. Jiang, C. Pope, K. M. Al Jarrah, J. H. Kung, T. Bortfeld, and G. T. Chen, Phys. Med. Biol. **48**, 1773–1784 (2003).
[26]A. E. Lujan, R. K. Ten Haken, E. W. Larsen, and J. M. Balter, Med. Phys. (submitted).
[27]T. Craig, V. Moiseenko, J. Battista, and J. Van Dyk, Int. J. Radiat. Oncol., Biol., Phys. **57**, 833–842 (2003).
[28]T. Craig, J. Battista, and J. Van Dyk, Med. Phys. **30**, 2012–2020 (2003).
[29]J. C. Stroom, H. C. de Boer, H. Huizenga, and A. G. Visser, Int. J. Radiat. Oncol., Biol., Phys. **43**, 905–919 (1999).
[30]S. D. McCarter and W. A. Beckham, Phys. Med. Biol. **45**, 923–931 (2000).
[31]I. J. Chetty, M. Rosu, N. Tyagi, L. H. Marsh, D. L. McShan, J. M. Balter, B. A. Fraass, and R. K. Ten Haken, Med. Phys. **30**, 1776–1780 (2003).
[32]J. T. Booth and S. F. Zavgorodni, Phys. Med. Biol. **46**, 1369–1377 (2001).
[33]A. Bel, M. van Herk, and J. V. Lebesque, Med. Phys. **23**, 1537–1545 (1996).
[34]J. H. Killoran, H. M. Kooy, D. J. Gladstone, F. J. Welte, and C. J. Beard, Int. J. Radiat. Oncol., Biol., Phys. **37**, 213–221 (1997).
[35]J. Leong, Phys. Med. Biol. **32**, 327–334 (1987).
[36]J. Lof, B. K. Lind, and A. Brahme, Phys. Med. Biol. **43**, 1605–1628 (1998).
[37]M. Birkner, D. Yan, M. Alber, J. Liang, and F. Nusslin, Med. Phys. **30**, 2822–2831 (2003).
[38]J. G. Li and L. Xing, Med. Phys. **27**, 1573-8 (2000).
[39]D. L. McShan and B. A. Fraass, 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 (Elsevier Science Publishers BV, North-Holland, New York, 1987), pp. 41–44.
[40]B. A. Fraass and D. L. McShan, 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 (Elsevier Science Publishers BV, North-Holland, New York, 1987), pp. 273–276.
[41]B. A. Fraass, D. L. McShan, R. K. TenHaken, and K. M. Hutchins, 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 (Elsevier Science Publishers BV, North-Holland, 1987), pp. 521–525.
[42]B. A. Fraass, D. L. McShan, and K. J. Weeks, 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 (Elsevier Science Publishers BV, North-Holland, New York, 1987), pp. 193–196.
[43]D. L. McShan, R. K. TenHaken, and B. A. Fraass, 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 (Elsevier Science Publishers BV, North Holland, New York, 1987), pp. 249–252.
[44]T. Mackie, A. Bielajew, D. Rogers, and J. Battista, Phys. Med. Biol. **33**,

1–20 (1988).
[45]D. L. McShan, B. A. Fraass, and M. L. Kessler, Med. Phys. **25**, A149 (1998).
[46]B. A. Fraass, D. L. McShan, and M. L. Kessler, in *Proceedings of the XIIIth International Conference On the Use of Computers In Radiotherapy*, edited by U. o. Heidelberg (T Bortfeld, W Schlegel, Heidelberg, Germany, 2000), pp. 32–35.
[47]M. L. Kessler, D. L. McShan, M. A. Epelman, K. A. Vineberg, A. Eisbruch, T. S. Lawrence, and B. A. Fraass, Opt. Eng. **6**, 425–448 (2005).
[48]J. J. H. Kim, N. Dogan, D. L. McShan, and M. L. Kessler, in *An AVS-based System for Optimization of Conformal Radiotherapy Treatment Plans*, Boston MA, 1995, pp. 417–423.
[49]A. E. Lujan, E. W. Larsen, J. M. Balter, and R. K. Ten Haken, Med. Phys. **26**, 715–720 (1999).
[50]R. K. Ten Haken, A. F. Thornton Jr., H. M. Sandler, M. L. LaVigne, D. J. Quint, B. A. Fraass, M. L. Kessler, and D. L. McShan, Radiother. Oncol. **25**, 121–133 (1992).
[51]K. K. Brock, D. L. McShan, R. K. Ten Haken, S. J. Hollister, L. A. Dawson, and J. M. Balter, Med. Phys. **30**, 290–295 (2003).
[52]F. L. Bookstein, *Morphometric Tools for Landmark Data: Geometry and Biology* (Cambridge University Press, Cambridge, 1991).
[53]A. Eisbruch, J. A. Ship, L. A. Dawson, H. M. Kim, C. R. Bradford, J. E. Terrell, D. B. Chepeha, T. N. Teknos, N. D. Hogikyan, Y. Anzai, L. H. Marsh, R. K. Ten Haken, and G. T. Wolf, World J. Surg. **27**, 832–7 (2003).
[54]A. Eisbruch, Med Dosim **27**, 99–104 (2002).
[55]RTOG, 2001.
[56]R. Marsh, J. Balter, V. L. Evans, and A. Eisbruch, Med. Dosim **22**, 293-7 (1997).
[57]D. Yan, F. Vicini, J. Wong, and A. Martinez, Phys. Med. Biol. **42**, 123–132 (1997).
[58]H. D. Thames Jr., H. R. Withers, L. J. Peters, and G. H. Fletcher, Int. J. Radiat. Oncol., Biol., Phys. **8**, 219–226 (1982).
[59]M. van Herk, M. Witte, J. van der Geer, C. Schneider, and J. V. Lebesque, Int. J. Radiat. Oncol., Biol., Phys. **57**, 1460–1471 (2003).