Improving IMRT-plan quality with MLC leaf position refinement post plan optimization

Ying Niu

*Department of Physics, The George Washington University, Washington, DC 20052 and
Department of Radiation Oncology, University of Maryland School of Medicine, Baltimore, Maryland 21201*

Guowei Zhang

Radiation Medicine, Georgetown University Hospital, Washington, DC 20007

Barry L. Berman and William C. Parke

Department of Physics, The George Washington University, Washington, DC 20052

Byongyong Yi and Cedric X. Yu^{a)}

Department of Radiation Oncology, University of Maryland School of Medicine, Baltimore, Maryland 21201

(Received 29 February 2012; revised 29 June 2012; accepted for publication 2 July 2012;
published 27 July 2012)

Purpose: In intensity-modulated radiation therapy (IMRT) planning, reducing the pencil-beam size may lead to a significant improvement in dose conformity, but also increase the time needed for the dose calculation and plan optimization. The authors develop and evaluate a postoptimization refinement (POpR) method, which makes fine adjustments to the multileaf collimator (MLC) leaf positions after plan optimization, enhancing the spatial precision and improving the plan quality without a significant impact on the computational burden.

Methods: The authors' POpR method is implemented using a commercial treatment planning system based on direct aperture optimization. After an IMRT plan is optimized using pencil beams with regular pencil-beam step size, a greedy search is conducted by looping through all of the involved MLC leaves to see if moving the MLC leaf in or out by half of a pencil-beam step size will improve the objective function value. The half-sized pencil beams, which are used for updating dose distribution in the greedy search, are derived from the existing full-sized pencil beams without need for further pencil-beam dose calculations. A benchmark phantom case and a head-and-neck (HN) case are studied for testing the authors' POpR method.

Results: Using a benchmark phantom and a HN case, the authors have verified that their POpR method can be an efficient technique in the IMRT planning process. Effectiveness of POpR is confirmed by noting significant improvements in objective function values. Dosimetric benefits of POpR are comparable to those of using a finer pencil-beam size from the optimization start, but with far less computation and time.

Conclusions: The POpR is a feasible and practical method to significantly improve IMRT-plan quality without compromising the planning efficiency. © 2012 American Association of Physicists in Medicine. [<http://dx.doi.org/10.1118/1.4737518>]

Key words: IMRT, plan-quality refinement

I. INTRODUCTION

Intensity-modulated radiation therapy (IMRT) has been routinely used to shape the high dose volume to a tumor while minimizing exposure to surrounding critical structures.¹ In the past decade, many researchers have devoted themselves to improve the IMRT planning to achieve better treatment quality. Evidently, any room for improving an IMRT plan depends strongly on the degrees of freedom provided by the machines and the related facilities.² Multileaf collimator (MLC) is a critical piece of the equipment which makes an effective delivery of intensity modulated beams possible. During the last two decades, MLCs have been continually refined by linear accelerator vendors to accommodate the development of IMRT. Newer MLCs are designed with narrower leaf widths and fewer movement restrictions to allow more accurate and versatile field shaping.^{3,4}

In IMRT treatment planning, a virtual grid is applied to divide the beam's eye view (BEV) of the planning target volume (PTV) into a series of small rectangular "bixels," each traversed by an imaginary pencil beam of radiation. For plan optimization and dose calculation, the doses deposited in the patient by each of the small "pencil" beams defined by the bixels are precalculated and stored. With precalculated pencil beams, the composite dose distribution from all the involved bixels, as well as for any modifications such as changing the weight of a bixel or adding/removing a bixel, can be quickly computed. For IMRT planning by either optimizing the bixel weights or directly optimizing aperture shapes and weights, the pencil-beam calculation is normally the first step. As such, the grid that defines the bixels determines the resolution of the final intensity map and the boundary roughness of the apertures.

The bixel size sets an artificial constraint in plan optimization by forcing MLC leaves, which can move continuously, to move in discrete steps. Using a high-performance dose-calculation engine to calculate pencil beams, a finer bixel size during optimization might allow the treatment planning system to achieve a better plan quality.⁵ The effect on plan quality caused by the pencil-beam step size, defined as the bixel dimension in the direction of MLC-leaf travel, was studied by Zhang *et al.*⁶ They found that the final objective function value could be decreased significantly when smaller pencil-beam step size was used, leading to significant improvements in dose conformity. These improvements are more pronounced when the structures to be spared are close to the target and a high dose gradient is desired.

Reducing the pencil-beam step size increases the pencil-beam number. Consequently, the pencil-beam calculation time, as well as the required disk space and computer memory, correspondingly increases. Also, the optimization time increases because the search space is enlarged. Even with the latest computers, the added computational burden can be practically prohibitive.

An effective method of improving plan quality and planning efficiency was the coarse-to-fine scheme for the planning of intensity-modulated arc therapy (IMAT).^{7,8} This strategy starts from a suboptimal IMAT plan which is achieved by stratifying the fluence profiles. The MLC leaf positions are subsequently refined using an efficient downhill search.

Here, we introduce a method to refine the MLC leaf positions after an initial plan optimization with a coarse pencil-beam step size. After the optimization, the artificial constraint on the pencil-beam step size is released, and MLC leaves are allowed to move to positions between adjacent gridlines to reduce effectively the step size. We find that, by a special refinement technique to determine new MLC positions after the plan optimization, a significantly better plan quality is achieved with little impact on the overall computational burden.

II. METHODS AND MATERIALS

II.A. Direct aperture optimization (DAO) and simulation environment

Our refinement method, described in Sec. II.B, is tested under the scheme of direct aperture optimization (DAO) (Ref. 9) although it can be applied to other plan optimization schemes. In DAO, the weight of apertures and the positions of MLC leaves are variables in the optimization process. Other treatment parameters are predetermined, such as the orientation of beams, the number of apertures per beam, etc. Compared with the scheme of fluence-map optimization, DAO does not need a leaf-sequencer to convert the optimized intensity map to deliverable MLC apertures. The positions of MLC leaves as well as the aperture weightings are directly optimized by an iterative search algorithm.

Prowess Panther^{TM} (Prowess Inc., Concord, CA), a commercial DAO-based inverse planning system employing the simulated annealing (SA) search algorithm, is used in this

study. During the optimization process, a tentative change is accepted when the objective value decreases. If the objective value increases, the change is accepted with a probability determined by a SA cooling schedule, thereby allowing the optimization to escape from being trapped in a local but not global minimum. To terminate the SA process, a user can set the maximal tolerated number of the consecutive unaccepted changes. Alternatively, the user can set the maximal number of accepted iterations, or manually cease the search when the intended objectives are met.

For a change of MLC-leaf position during iteration, a random travel distance is sampled from a Gaussian distribution. The sampled travel distance is always rounded to be an integral multiple of pencil-beam step size so that the leaf end is at a bixel grid line. As the accepted SA iterations increasing, the variance of sampling decreases and tends toward unity.⁷ Thus, at a late stage of SA process, the tentative travel distances of MLC leaves are often less than one pencil-beam step size, causing no leaf position changes. However, if the use of pencil beams with a finer step size is allowed at this juncture, a further refining search is enabled, and better aperture shapes can be formed.

In this study, the Prowess Panther^{TM} system is performed on a 2.21 GHz PC with dual core AMD CPU. Dose distributions are calculated using the convolution/superposition method with voxels of size $3 \times 3 \times 3$ mm³. The pencil-beam dose distributions are individually calculated for each position in the fluence map, accounting for inhomogeneous media. The Elekta MLC (with 1 cm wide leaves) is commissioned to the planning system.

II.B. Postoptimization refinement (POpR) of leaf positions

The goal of this study is to see if it is possible to refine the leaf positions of all the beam apertures with minimal increase of computational burden after the SA process. The scheme we devised is to take a plan optimized with a coarse pencil-beam step size, and add a postoptimization refinement (POpR) process by relocating MLC leaves. We aim to see if allowing the leaves to move either forward or backward by half a step size would improve the plan quality, and if it would, how the resultant plan of POpR compares with the plan optimized using a smaller pencil-beam step size from the start of SA process. In this study, the objective function and penalties used in the POpR are kept the same as those used in the SA process.

The key for this scheme to work is being able to judge quickly the plan quality based on the resulting dose distributions. To evaluate the dose distribution after a MLC leaf makes a move of half step size, we implemented a method to approximate the dose distributions of a half pencil beam. The approximation method is described in Sec. II.C.

II.C. Computing the dose distributions of half pencil beams

The refinement by using half pencil beams happens at the aperture's boundary [Fig. 1]. Therefore, the number of needed

FIG. 1. The illustration of BEV and the grid for the division of pencil beams. An OAR overlaps with the PTV. If the pencil beams with half step size are provided, a better aperture shape (i.e., the dashed boundary) can be formed based on the optimized configuration of SA (i.e., the solid boundary). Only half pencil beams at the boundary of the aperture need to be calculated for refinement.

half pencil beams for POpR is much less than the number of pencil beams would be required for the SA optimization using the half step size from the start.

To calculate these POpR-involved half pencil beams, it is not unreasonable to assume that the half pencil beam traverses the same anatomy as the whole pencil beam. As shown in Fig. 2(a), if we further ignore the slight difference in divergence between the two abutting half pencil beams that makes the full pencil beam, the dose distributions of them will have the same functional form but just have a lateral shift from each other.

With these assumptions, we can quickly derive the dose distribution of the half pencil beam from the full-sized pencil beam with a shift-and-subtract operation. As shown in Fig. 2, we sum two abutting full-sized pencil beams [Fig. 2(b)], which are prepared before optimization. Then, we shift the summed dose distribution to the right with a distance of the half pencil

FIG. 2. (a) The single pencil beam with a step size of $|\vec{S}|$ can be split into two abutting half pencil beams which can be approximately transformed by shifting $|\vec{S}|/2$ reciprocally. (b)–(d) The dose distribution of the half pencil beam with a step size of $|\vec{S}|/2$ can be obtained by a shift-and-subtract operation.

FIG. 3. The dose profiles of pencil beams on the beam-axis plane at different depth d : 5 mm step size (squares) vs 2.5 mm step size (asterisks). The half pencil beams are calculated by using the prepared full-sized pencil beams with the dose distribution in a homogeneous water phantom. The voxels for storing dose are $3 \times 3 \times 3 \text{ mm}^3$ in size, and the maximum value of all voxels is normalized to 1.

beam step size [Fig. 2(c)]. The dose distribution of the half pencil beam can be easily obtained as the difference between these two broad fields [i.e., the positive part in Fig. 2(d)]. This shift-and-subtract operation can be performed with negligible computer time. In principle, the pencil beam of an arbitrary step size can be generated by shifting with the distance of corresponding step size.

Figure 3 gives the dose distribution resulting from a shift-and-subtraction operation for two half pencil beams with 2.5 mm step size versus the dose distribution of a full-sized pencil beam of 5 mm step size. The generated half pencil-beam dose

profiles not only represent accurately the central line of primary fluences and the depth-dose variation, but also predict the penumbra region correctly. The correct penumbra is very important for the refinement process at the aperture boundary, where the alignment of dose gradient is highly critical to the treatment plan.

II.D. Greedy search

A greedy search is implemented for the POpR. One by one, each MLC leaf is moved forward or backward by half of the original step size. The corresponding dose and the objective function value are evaluated after each move. For each move, if the objective value decreases, the new position will be accepted, whereas the move in the opposite direction will not be tested. If the objective value increases (i.e., the plan worsens), the movement of half a pencil-beam step size in the opposite direction will be tested. If moving the leaf either forward or backward does not improve the plan quality, the leaf position is not changed. After one greedy search is completed for all involved MLC leaves, a new aperture is obtained. The leaves that are not moved will be tested again in the next iteration. No moveable MLC leaf exists after a few iterations. During this greedy-search process, the mechanical constraints are always checked to ensure deliverability.

II.E. Method of testing the POpR

A benchmark phantom case is studied for testing the effectiveness and efficiency of the POpR method. In addition, a head-and-neck (HN) case is studied to verify benefits of the POpR method in clinical planning practice. Multiple planning trials with different pencil-beam step sizes are conducted for both phantom and patient studies. Here, we use an abbreviation POpR-x to denote a trial of SA-optimization process using x mm step size from the start, which is followed by a greedy-search process with tentative moves by half of original pencil-beam step size (i.e., $x/2$ mm). An abbreviation Op-x denotes a trial of SA optimization using x mm step size from the start without an extra greedy-search refinement.

II.E.1. Benchmark phantom study

An IMRT benchmark phantom is used for the simulation study. A cross section of the phantom is shown in Fig. 4, in which a C-shaped PTV partially wraps around a circular organ of risk (OAR). The planning goal is to deliver 100% of prescription dose to 95% of the PTV. No more than 15% of the OAR should receive more than 60% of the prescribed dose. To control hot spots, the maximum dose within the PTV should be no more than 120%, and no more than 10% of healthy tissue volume (HTV) should receive more than 100% of the prescribed dose.

Various pencil-beam step sizes are initialized for the planning trials. Two trials (i.e., POpR-10 and POpR-5) for testing the method are conducted. The trials (i.e., Op-10, Op-5, and Op-2) with only SA-optimization processes are conducted for performance comparison. Because fractional-millimeter pencil-beam step size is not allowed in Prowess Panther^{TM}

FIG. 4. The C-shaped benchmark phantom. The phantom is a cylinder, 10 cm in diameter and 10 cm in length. The OAR is a central cylinder, 1.6 cm in diameter and 8 cm in length. The PTV, which has the same length with OAR, is a 2.5 cm wide half-annulus that surrounds the OAR over 180°. The rest of volume, which does not belong to PTV and OAR, is defined as HTV.

system, 2 mm of pencil-beam step size is selected instead of 2.5 mm. For each planning trial of Op-x or POpR-x, the same set of constraints and objectives is used. For each trial, seven 6 MV photon beams equally spaced in gantry angle are initialized for DAO with seven apertures per beam direction.

II.E.2. HN case study

An oropharyngeal cancer case is used to test our POpR method in clinical practice. As shown in Fig. 5, the PTV including a primary target is of a concave shape, which partially wraps around the spinal cord, the brain stem, and bilateral neck nodes. The PTV is 545.9 cm³, and the length is 20.7 cm in the superior-inferior (SI) direction extending from the cochleae to the clavicles. Besides the brainstem and the cervical spine, the OARs also include the parotid glands on both sides, which overlap with PTV by 7.4% on the right side and 5.6% on the left.

FIG. 5. The PTV and OARs delineations on a set of planning CT for the first course of IMRT treatment. (a) A transverse plane presents PTV, parotids, and spinal cord. A 0.3 cm margin is added to the cord. (b) A 3D view of all of involved volumes and organs. A skin volume is defined to have 0.3 cm depth from the patient surface, which is not shown here.

TABLE I. Dose specifications at University of Maryland Medical Center for the HN case.

Target coverage	Acceptance criteria
PTV = 59.4 Gy	$V_{95} \geq 95\%$ PTV volume $V_{105} \leq 5\%$ PTV volume
Normal tissue constraint	
Spinal cord	$D_{\max} < 42$ Gy
Brain stem	$D_{\max} < 45$ Gy
Cochlea	$D_{\max} < 42$ Gy
Parotid glands	$D_{\text{mean}} \leq 25$ Gy

^aNote: V_{95} = the volume receives at least 95% of the PTV prescription dose; $D_{\max}$ = the maximal dose; D_{mean} = the mean dose.

Plans are optimized to meet the dose specification shown in Table I. The planning goal is to deliver 59.4 Gy in 33 fractions to the entire PTV. The doses to the spinal cord and brain stem are limited to 42 Gy–45 Gy, respectively, so as to leaving room for subsequent boost to the primary target. The close proximity of the PTV to the skin and the overlaps between the PTV and the parotid glands create a challenging case with conflicting objectives.

Initial pencil-beam step sizes of 10 mm and 5 mm are used, as these are common choices in clinical settings. Two SA-optimization trials (i.e., Op-10 and Op-5) and two refinement trials (i.e., POpR-10 and POpR-5) are conducted using the same set of constraints and objectives based on the dose specification. Seven 6 MV photon beams are used for each trial. The beam angles are 0° , 51° , 103° , 160° , 200° , 255° , 305° , which are chosen to best target the PTV volume. Due to the MLC limitation on the maximal travel distance across the centerline, the fields are split when BEVs are large. The total number of initialized apertures is 154 (i.e., 22 for each beam).

III. RESULTS

III.A. The benchmark phantom study results

III.A.1. Efficiency and effectiveness of POpR

Table II shows the comparisons of the pencil-beam preparing time, the SA search time, the successful SA iterations, and the objective values. We easily find better plan qualities resulting from finer pencil-beam step sizes, which accords with

the conclusions of Zhang *et al.*⁶ To ensure valid comparison of efficiency and effectiveness, multiple runs were conducted to ensure that the cooling schedule is adequate and all trials converge to the minimum.

As shown in Table II, about twice the optimization time is spent (e.g., 5–6 h for Op-5 vs 2–3 h for Op-10) to obtain an improvement in objective value using finer pencil-beam step size rather than a coarse one. However, the time needed for our refinement (POpR), including the time to generate the half-sized pencil beams, is only 1–3 min, which is negligible compared with the time consumed in the SA-optimization processes. The refinement, for which 1200–1500 leaves are involved, results in around 35% of the leaves being relocated during the fine adjustment. The entire time of refinement is dominated by the time needed for generating half-sized pencil beams.

Compared with the Op-10 trial, the refinement process of the POpR-10 trial significantly lowers the average objective values by 57% of the unit objective value. In the POpR-5 trial, the average objective values decrease by 21% compared to the Op-5 trial. Comparing the POpR trials using the coarse step sizes with the Op-x trials using finer pencil-beam step sizes (i.e., POpR-10 vs Op-5 and POpR-5 vs Op-2), we only find 10%–15% differences in average objective values.

III.A.2. Dosimetric benefits from POpR

For each Op-x trial, the best plan from one of the ten runs is selected to show dosimetric results. The selected Op-x plan is used to generate a refined POpR-x plan. Figure 6 shows iso-dose curves for the Op-10 plan and the corresponding POpR-10 plan. The comparisons of dose volume histograms (DVHs) are presented in Fig. 7. Detailed dose analysis of all Op-x and POpR-x plans is shown in Table III. All the dose distributions for comparison are scaled to ensure that 95% of PTV receives 100% prescription dose. (For geometrically challenging cases in planning, the PTV coverage will be insufficient if the prescription dose is set to D_{50} of PTV as recommended by ICRU.¹⁰ In this paper, the prescribing is contingent on the specific clinical requirement.)

The POpR makes better homogeneity of PTV dose, which is found in Table III based on the decrease of near-maximum (D_2) dose and the increase of the near-minimum (D_{98})

TABLE II. The performances of planning trials with or without a greed-search refinement.

Trial	Pencil-beam calculating time	Optimization time (h)	Successful SA iterations	Objective values (10 runs)
Op-10 ^a	15 min	2–3	~30 000	1.87 ± 0.06
POpR-10 ^b	15 min	2–3	~30 000	1.30 ± 0.06
Op-5	30 min	5–6	~50 000	1.20 ± 0.06^c
POpR-5	30 min	5–6	~50 000	0.99 ± 0.05
Op-2	1 h 15 min	8–10	~100 000	0.84 ± 0.05

^aOp-10 denotes the SA-optimization trial without a refinement process using 10 mm pencil-beam step size.

^bPOpR-10 denotes the post-optimization refinement trial with 10 mm initial pencil-beam step size.

^cAmong all the runs of Op-5, the run with minimal objective value is chosen to be the unit objective value to normalize the objective values of all the runs of the different trials of POpR-x and Op-x.

FIG. 6. Isodose lines of the central transverse slices for the plans of Op-10 and POpR-10. The isodose lines from inner to outer positions represent 116%, 100%, 80%, 60%, 40%, and 30% of the prescription dose.

dose after refinement. (To avoid the reporting inaccuracy due to a single computation point, the near-maximum and near-minimum doses, which are recommended by ICRU,¹⁰ are used as indicators instead of the maximal and minimal dose.) The homogeneity index (HI), which is defined as $(D_2 - D_{98})/D_{50}$, is also presented in Table III, indicating the improvement of PTV coverage after refinement. Compared with the Op-*x* plans, the POpR-*x* plans perform better in controlling hot spots of PTV, which is observed through the shrinking area of 116% isodose curve, as well as the values of V_{116} and V_{120} . The improvements in OAR sparing due to the refinement is indicated by the changes of dose gradient, which can be found when comparing the isodose lines of plans in Fig. 6. The refining benefits for OAR sparing are also shown as the decreases of D_{mean} , $D_{near-max}$, and V_{60} . After refining, the hot area of HTV becomes smaller according to the move of 100% isodose curve and the decrease of V_{100} . The dosimetric performances of POpR-10 and POpR-5 even reach or approach the performances of Op-5 and Op-2, respectively, which are demonstrated by the HI of PTV, the mean dose to the OAR, etc.

III.B. The HN case study

III.B.1. Efficiency of POpR in the HN case

The planning times and objective values of the SA-optimization trials and the greedy-search refinement trials are shown in Table IV. The objective values decrease on average by 17% when we refine the pencil-beam step size from 10 mm to 5 mm, resulting in nearly equal objective values as obtained using 5 mm step size from the start. Refining plans from 5 mm step size to 2.5 mm step size yields 4% average decrease in the objective values. Only 5 min are needed for the greedy-search process of each POpR-*x* trials. Compared with the performances of the Op-*x* trials at the late stage of SA processes (i.e., <1% decrease in objective value during 15 min), the performances of POpR-*x* trials show an advantage in saving planning time.

III.B.2. Dosimetric benefits from POpR in the HN case

Table V shows the dosimetric performances of the Op-*x* plans and the corresponding POpR-*x* plans with various plan-

FIG. 7. The DVH comparison of the resulted plans of Op-10 (dotted-dashed lines), POpR-10 (solid lines), and Op-5 (dashed lines). The plan of POpR-10 is a refined result based on the optimized segments of Op-10.

TABLE III. Dosimetric data analysis of the optimized plans of Op-10, Op-5, Op-2 and the refined plans of POpR-10, POpR-5 for the phantom study.

		Op-10	POpR-10	Op-5	POpR-5	Op-2
PTV	$D_{\text{near-max}}$ (%)	125.9	119.7	120.1	119.5	118.3
	$D_{\text{near-min}}$ (%)	90.3	97.0	92.8	94.4	96.0
	HI	0.314	0.204	0.243	0.225	0.200
	V_{116} (%)	26.0	11.2	17.8	16.1	10.7
	V_{120} (%)	9.8	2.3	1.5	2.3	0.2
OAR	D_{mean} (%)	45.7	43.6	42.9	41.3	41.1
	$D_{\text{near-max}}$ (%)	75.4	74.2	74.3	69.3	68.4
	V_{60} (%)	19.9	17.0	16.0	12.8	10.4
HTV	V_{100} (%)	12.1	7.6	7.3	6.5	6.3

^aNote: D_{mean} = the mean dose; $D_{\text{near-max}}$ = D_2 , the minimal dose received by 2% of most irradiated volume; $D_{\text{near-min}}$ = D_{98} , the minimal dose received by 98% of most irradiated volume; HI = $(D_2 - D_{98})/D_{50}$. All of doses are presented as the percentage of PTV prescription dose.

quality indicators. The plans of Op-10 and POpR-10 are also compared in DVHs as shown in Fig. 8. We normalize doses of plans to make 96% of PTV receive at least 95% prescription dose (i.e., $V_{95} = 96\%$).

Comparing the plans of POpR-10 and Op-10, a better PTV coverage is obtained after refinement, which is indicated by the decrease of HI. The hot spots of PTV also are reduced according to the decreases of V_{103} and V_{105} . The mean doses to the spinal cord and brain stem decrease by 0.7 Gy and 0.6 Gy, respectively. The near-maximum dose received by the brain stem decreases by 1.7 Gy, and the near-maximum dose to the spinal cord decreases by 0.4 Gy. The near-maximum cochlea doses are not improved in general, because the acceptance criterion concerning the cochlea has been well met before refinement. The V_{95} of skin decreases by 0.4%. Satisfactory mean doses to the parotids (i.e., <25 Gy) are not achieved in both plans of POpR-10 and Op-10. However, the refinement decreases the mean doses to the two parotids by 0.6 Gy and 0.5 Gy, respectively.

Comparing the plans of POpR-5 and Op-5, small improvements of refinement are found in the hot-spot area (e.g., V_{103}) and mean dose to parotids as presented in Table V, while the values of other quality indicators are left unchanged in general.

IV. DISCUSSION

The POpR method helps the optimizer efficiently approach the optimum which was previously unreachable without using

TABLE V. Dose-data comparison of the SA-optimization trails (Op-x) and the postoptimization refinement trials (POpR-x) for the HN case.

		Op-10/POpR-10	Op-5/POpR-5
PTV	$D_{\text{near-min}}$ (Gy)	54.4/54.2	54.0/54.0
	$D_{\text{near-max}}$ (Gy)	63.3/62.4	62.1/62.1
	HI	0.146/0.135	0.134/0.134
	V_{105} (%)	6.7/2.6	2.4/2.0
	V_{103} (%)	37.2/24.2	27.4/25.6
Spinal cord	$D_{\text{near-max}}$ (Gy)	39.3/38.9	39.7/39.6
	D_{mean} (Gy)	30.7/30.0	31.3/31.2
Brain stem	$D_{\text{near-max}}$ (Gy)	37.6/35.9	37.8/37.9
	D_{mean} (Gy)	25.8/25.2	26.8/26.7
Right parotid	D_{mean} (Gy)	27.4/26.8	28.9/28.8
Left parotid	D_{mean} (Gy)	26.2/25.7	26.3/26.0
Right cochlea	$D_{\text{near-max}}$ (Gy)	22.8/24.2	25.5/25.8
Left cochlea	$D_{\text{near-max}}$ (Gy)	26.1/25.1	25.5/25.2
Skin	D_{mean} (Gy)	20.4/20.3	20.7/20.7
	$D_{\text{near-max}}$ (Gy)	57.0/56.1	56.8/56.8
	V_{95} (%)	2.3/1.9	2.1/2.1

finer pencil-beam step sizes from the start of the SA process. In both the phantom and HN studies, the refined plans show improvements in both PTV homogeneity and OAR sparing compared to the plans generated with coarse pencil-beam step sizes, and the qualities of the refined plans approach those optimized with finer step sizes from the start. As compared with the use of finer pencil beams, POpR achieves similar improvements with negligible time relative to the time needed for the initial plan optimization rather than doubling the computation time.

The dosimetric studies for both the phantom and HN cases clearly demonstrate that the POpR decreases in the objective values, producing real improvements in plan qualities. It is verified that the approximations used applied to half pencil beams provides enough accuracy to ensure the proper dose predictions in the greedy-search process, while saving much time for refinement.

Many factors may affect the performance of POpR in specific cases. Notable benefits of POpR are expected after we optimize a challenging case with competing objectives. For example, the benefits of POpR in improving the PTV homogeneity are not so evident in the HN case compared with those in the phantom case, which accords with the difference of the two cases in how much the objective-value decreases. In the phantom case, the PTV surrounding the OAR with a very small gap causes a competition between the PTV con-

TABLE IV. The performance of SA-optimization trials (Op-10 and Op-5) and greedy-search refinement trials (POpR-10 and POpR-5) for the HN patient.

Trial	Pencil-beam calculating time	Optimization time (h)	Successful SA iterations	Objective values (10 runs)
Op-10	~40 min	1–1.5 h	~10 000	1.80 ± 0.02
POpR-10	~40 min	1–1.5 h	~10 000	1.50 ± 0.03
Op-5	~1 h 20 min	2.5–3 h	~16 000	1.48 ± 0.02
POpR-5	~1 h 20 min	2.5–3 h	~16 000	1.42 ± 0.02

FIG. 8. The DVH comparison of the trials without or with a refinement process for a HN case: Op-10 (dotted-dashed lines) vs POPR-10 (solid lines).

formity and the OAR sparing. Because such competition happens in each of transverse planes, the volume involved in the objective competition and the refinement is large, therefore, causing a notable dosimetric difference before versus after POPR. However, in the HN case, the major objective competition happens at the boundaries or overlaps of the PTV with the parotids, which are small regions. Consequently, the volumes involved in the refinement are small compared with the total PTV. Therefore, the improvement of the PTV homogeneity in the HN case is not as clear as that in the phantom study. In contrast, the benefits of sparing the parotids from the POPR are pronounced due to the small volumes of parotids.

Numerous studies have found that a strong relationship exists between dose to the parotid gland and the salivary function after treatment.^{11–13} Studies have shown a gradual decrease in mean parotid dose produces a gradual decrease in the normal tissue complication probability.¹⁴ Thus, a clinical significance exists when a moderate improvement of parotid sparing is obtained. A consistent benefit from POPR in parotid sparing is expected generally, because the volumes of the parotids are normally small and comparable to the volumes involved the refinement.

In this study, the objective function and penalties used in the POPR are kept the same as those used in the SA process. Thus, the refinement is not biased to any specific regional dose (e.g., the parotid dose). On the other hand, we find a significant decrease in objective value due to the POPR does not always translate into the same degree of plan quality improvement. This occurs because the optimization is driven by minimizing a specific objective function that convolutes a multiobjective problem into a single value. The objective function may not ideally represent all the dosimetric desirabilities. For example, the defined volumes and the corresponding objective function weightings, which are specified before the start of an optimization, may not correspond to the desirabilities of particular regional doses at the end stage of the optimization. Some studies introduced voxel-specific penalties instead of structure-specific penalties to fine-tune the regional doses (e.g., the parotid-PTV boundary in HN

case, the prostate-rectum boundary in prostate case, etc.) to meet clinical requirements.^{15,16} They found that, both tumor hot spots and OAR doses could be controlled by varying the regional voxel-specific penalties after a conventional inverse planning. Others have introduced Pareto optimization strategies that vary the parameters in the objective function within a practical range to create a multidimensional surface, called the Pareto front, rather than a single optimal objective value.¹⁷

For IMAT planning, the optimizer must cope with more beam orientations and more stringent machine constraints rather than the situation of segmental IMRT. Bedford's study has proven a downhill search is efficient for fine-tuning an IMAT plan with coarse segments, while maintaining the deliverability.⁸ Enabling finer pencil beams, the POPR may also bring an extra improvement for an IMAT plan.

V. CONCLUSIONS

We have developed a method (POPR) to refine the MLC leaf positions after an IMRT-plan optimization with a coarser set of leaf positions. A benchmark phantom study and a HN case study consistently show a plan-quality improvement after the refining processes. As compared to the use of finer pencil-beam step size from the start of optimization, the POPR method achieved similar results without doubling the optimization time or expanding the computer-memory requirement. Further study is needed to see the benefits of our method when applied to a planning scheme with variable penalties by regions and for IMAT planning.

ACKNOWLEDGMENT

This work was supported in part by National Institutes of Health (NIH) Grant No. RO1CA133519.

^{a)} Author to whom correspondence should be addressed. Electronic mail: cyu002@umaryland.edu

¹L. K. Mell, A. K. Mehrotra, and A. J. Mundt, "Intensity-modulated radiation therapy use in the U.S., 2004," *Cancer* **104**(6), 1296–1303 (2005).

²C. X. Yu, C. J. Amies, and M. Svatos, "Planning and delivery of intensity-modulated radiation therapy," *Med. Phys.* **35**(12), 5233–5241 (2008).

³J. B. Fiveash, H. Murshed, J. Duan, M. Hyatt, J. Caranto, J. A. Bonner, and R. A. Popple, "Effect of multileaf collimator leaf width on physical dose distributions in the treatment of CNS and head and neck neoplasms with intensity modulated radiation therapy," *Med. Phys.* **29**, 1116–1119 (2002).

⁴H. D. Kubo, R. B. Wilder, and C. T. E. Pappas, "Impact of collimator leaf width on stereotactic radiosurgery and 3D conformal radiotherapy treatment plans," *Int. J. Radiat. Oncol., Biol., Phys.* **44**, 937–945 (1999).

⁵T. Bortfeld, U. Oelfke, and S. Nill, "What is the optimum leaf width of a multileaf collimator?," *Med. Phys.* **27**, 2494–2502 (2000).

⁶G. Zhang, Z. Jiang, D. Shepard, M. Earl, and C. Yu, "Effect of beam-let step-size on IMRT plan quality," *Med. Phys.* **32**(11), 3448–3454 (2005).

⁷K. Otto, "Volumetric modulated arc therapy: IMRT in a single gantry arc," *Med. Phys.* **35**, 310–317 (2008).

⁸J. L. Bedford, "Treatment planning for volumetric modulated arc therapy," *Med. Phys.* **36**, 5128–5138 (2009).

- ⁹D. M. Shepard, M. A. Earl, X. A. Li, S. Naqvi, and C. Yu, "Direct aperture optimization: A turnkey solution for step-and-shoot IMRT," *Med. Phys.* **29**, 1007–1018 (2002).
- ¹⁰ICRU Report 83, "Prescribing, recording, and reporting intensity-modulated radiation therapy," *Journal of the International Commission on Radiation Units and Measurements* (Oxford University Press, Oxford, UK, 2010), Vol. 10(1).
- ¹¹A. Eisbruch, H. M. Kim, J. E. Terrell, L. H. Marsh, L. A. Dawson, and J. A. Ship, "Xerostomia and its predictors following parotid-sparing irradiation of head-and-neck cancer," *Int. J. Radiat. Oncol., Biol., Phys.* **50**, 695–704 (2001).
- ¹²M. W. Münter, C. P. Karger, S. G. Hoffner, H. Hof, C. Thilman, V. Rudat, S. Nill, M. Wannenmacher, and J. Debus, "Evaluation of salivary gland function after treatment of head-and-neck tumors with intensity-modulated radiotherapy by quantitative pertechnetate scintigraphy," *Int. J. Radiat. Oncol., Biol., Phys.* **58**, 175–184 (2004).
- ¹³A. I. Blanco, K. S. Chao, I. El Naga, G. E. Franklin, K. Zakarian, M. Vicic, and J. O. Deasy, "Dose-volume modeling of salivary function in patients with head-and-neck cancer receiving radiotherapy," *Int. J. Radiat. Oncol., Biol., Phys.* **62**, 1055–1069 (2005).
- ¹⁴T. Dijkema, C. P. J. Raaijmakers, R. K. Ten Haken, J. M. Roesink, P. M. Braam, A. C. Houweling, M. A. Moerland, A. Eisbruch, and C. H. J. Terhaard, "Parotid gland function after radiotherapy: The combined Michigan and Utrecht experience," *Int. J. Radiat. Oncol., Biol., Phys.* **78**, 449–453 (2010).
- ¹⁵C. Cotrutz and L. Xing, "IMRT dose shaping with regionally variable penalty scheme," *Med. Phys.* **30**, 544–541 (2003).
- ¹⁶C. Holdsworth, M. Kim, J. Liao, and M. Phillips, "The use of a multiobjective evolutionary algorithm to increase flexibility in the search for better IMRT plans," *Med. Phys.* **39**, 2261–2274 (2012).
- ¹⁷D. Craft, T. Halabi, and T. Bortfeld, "Exploration of tradeoffs in intensity-modulated radiotherapy," *Phys. Med. Biol.* **50**, 5857–5868 (2005).