IMRT optimization: Variability of solutions and its radiobiological impact

Maurizio Mattia,[a)] Paolo Del Giudice,[b)] and Barbara Caccia[c)]
Istituto Superiore di Sanità, Physics Laboratory, Viale Regina Elena 299, 0161 Roma, Italy

(Received 11 September 2003; revised 2 January 2004; accepted for publication 13 February 2004)

We aim at (1) defining and measuring a ''complexity'' index for the optimization process of an intensity modulated radiation therapy treatment plan (IMRT TP), (2) devising an efficient approximate optimization strategy, and (3) evaluating the impact of the complexity of the optimization process on the radiobiological quality of the treatment. In this work, for a prostate therapy case, the IMRT TP optimization problem has been formulated in terms of dose-volume constraints. The cost function has been minimized in order to achieve the optimal solution, by means of an iterative procedure, which is repeated for many initial modulation profiles, and for each of them the final optimal solution is recorded. To explore the complexity of the space of such solutions we have chosen to minimize the cost function with an algorithm that is unable to avoid local minima. The size of the (sub)optimal solutions distribution is taken as an indicator of the complexity of the optimization problem. The impact of the estimated complexity on the probability of success of the therapy is evaluated using radiobiological indicators (Poissonian TCP model [S. Webb and A. E. Nahum, Phys. Med. Biol. **38**(6), 653–666 (1993)] and NTCP relative seriality model [Kallman *et al.*, Int. J. Radiat. Biol. **62**(2), 249–262 (1992)]). We find in the examined prostate case a nontrivial distribution of local minima, which has symmetry properties allowing a good estimate of near-optimal solutions with a moderate computational load. We finally demonstrate that reducing the *a priori* uncertainty in the optimal solution results in a significant improvement of the probability of success of the TP, based on TCP and NTCP estimates. © *2004 American Association of Physicists in Medicine.* [DOI: 10.1118/1.1695650]

PROOF COPY 020405MPH

INTRODUCTION

Evidence is accumulating that intensity-modulated radiation therapy (IMRT) can produce excellent conformal treatment plans even for the most complex target volumes.[1,2] The capability to conform the dose delivery to the tumoral target, preserving the normal tissue, allows us to escalate the therapeutic dose to the tumor with a significant reduction of late effects, increasing the probability of the success of the therapy with respect to the conventional radiotherapy. The determination the best intensity modulation is usually based on heuristic methods. To quantify how much a particular intensity modulation is close to the optimal choice, a cost function is defined that represents the ''error'' associated with the treatment plan obtained. The best beam modulation is found by minimizing such cost function with respect to the parameters defining the treatment plan, through a suitable optimization algorithm. In the general case one must consider the possibility that several local minima may exist for the cost function, associated with suboptimal solutions.

Several issues related to local minima in IMRT treatment planning have been addressed in the recent literature.[1,3] The conditions for the existence of local minima have been investigated; in particular it has been recognized that the choice of dose-volume constraints implies local minima in the cost functions[4,5] as opposed to the quadratic cost function,[6] possibly including maximum dose constraints.[7]

Stochastic algorithms have been proposed to find good solutions in presence of local minima,[8,9] while their relevance have been addressed in terms of the corresponding distribution of values for the cost function: an agreement emerged on the smallness of the effect of local minima, suggesting that adopting a simple deterministic optimization strategy would not greatly affect the quality of the resulting treatment plan.[5,7,10,11]

The above results are not easily converted in a quantitative assessment of the quality of the treatment plan: in particular similar values of the cost functions can correspond to different beam modulation profiles,[7,10] and the question presents itself as to how much those differ in the resulting radiobiological quality (TCP/NTCP). These issues have only recently started being addressed.

In this paper we extend the ''configuration space analysis'' [11–13] to characterize the distribution of suboptimal solutions in the modulation profiles space, and we map those to the corresponding distribution of TCP and NTCP value in order to get a radiobiological score. Such a quantitative physical and radiobiological characterization of suboptimal solutions can help to identify the appropriate compromise between computational resources employed and the probability of finding a treatment plan of prescribed quality. In particular, such an approach allows us to figure out an effective and robust strategy, with respect to the initial conditions and minimization convergence criteria, to find an acceptable solution to the optimization problem. Some recent works[7,12,14] with similar aims to ours use different methods, with partially overlapping results, and in the Conclusions the most significative links with our findings will be discussed.

PROOF COPY 020405MPH

Preliminary results of the present work have been reported in Ref. 15.

METHODS AND MATERIALS

The optimization process aims at defining the "best" treatment plan in terms of the intensity modulation $\mathbf{x}$ of all pencil beams, which are the degrees of freedom of the problem. In the case at hand, an analytical approach poses formidable difficulties (see, however, Ref. 16), and one is often forced to resort to numerical strategies for the approximate solution of the problem.

In such cases, it is useful to visualize the optimization process as the "motion" of a point representing the modulation profile, on the surface describing a "cost function" $F(\mathbf{x})$, whose value at each point (each modulation profile $\mathbf{x}$) encodes the "penalty" for being nonoptimal. The best modulation profile(s) are then identified with the minimum (or minima) of F. The space in which F is defined has as many dimensions as there are pencil beams; it is therefore a high-dimensional space, something to be kept in mind when exercising intuition on the process.

The optimization algorithm defines suitable "forces" which drive the trajectory of the representative point towards the minima of F.

We will use the above "landscape" metaphor throughout as a conceptual context for both the description of the complexity measure and the search for the optimization strategy.

There is no unique definition of complexity, and we adopt the following empirical criterium. We first choose a set of modulation profiles as initial conditions of the optimization process, covering a reasonable range in terms of fluence; for each initial condition an iterative algorithm to be defined shortly finds an "optimal" solution, and the distribution of the optimal solutions found for all the initial conditions is characterized: the broader this distribution, the higher the complexity.

Optimization algorithm

The algorithm has a deterministic and a stochastic component. An initial condition is chosen, $\mathbf{x}_0$, for the modulation of all the pencil beams. Then the following steps are iterated: (1) One pencil beam is chosen at random. (2) A trial increase of fixed size Δ in its modulation is performed. (3) The new value of the cost function F is calculated for this new $\mathbf{x}$ configuration. (4) If the trial change results in a decrease of F the change is accepted and defines the new value of the corresponding modulation; if not, a trial decrease of the same amount is tried, and the same check on F is performed. (5) A new pencil beam is chosen at random, and the preceding steps are repeated. (6) The rate of accepted changes is monitored, and if it drops below a chosen threshold, the size of the trial changes is halved ($\Delta \rightarrow \Delta/2$). We associate the succession of Δ values with corresponding 'stages' of the optimization process. (7) The algorithm stops when Δ falls below a chosen threshold.

The algorithm is in part deterministic, in that the motion of the representative point is strictly downhill, and it is never allowed to climb barriers, so that if a local minimum is encountered the optimization will end up trapped there (to be compared, for example, to the "simulated annealing" stochastic optimization). This is intentionally used as an instrumental ingredient in the present study, since we want to *count* the local minima, not to escape from them. But the algorithm also has a stochastic component, since at each point the chosen downhill motion is in a random direction (to be compared to the gradient descent and the like, in which the direction of steepest slope determines the direction of motion). The latter stochasticity is beneficial in improving the ability to explore wider regions of the configuration space. On the other hand, adopting a gradient descent strategy would be prohibitive at the computational level for the intended analysis: for DVH-based cost functions the gradient would have to be computed numerically, and the related computational load should be multiplied by the high number of optimizations needed to explore the F surface and estimate its complexity.

The optimization procedure is repeated for many initial conditions, and the configurations reached at the end of each stage are recorded.

We remark that the value of Δ at each optimization stage defines a characteristic scale which is an upper bound on the "resolution" with which the optimization process is able to probe the F surface. Optimization histories starting with high values of Δ will be called *coarse* in the following, while those starting with small Δ will be termed *refined*. Typically the distribution of final pencil beam intensities is peaked on $\sim$0.4 for the target tumor and $\sim$0 for OAR; at the end of the optimization process $\Delta = 1/2^6 \simeq 0.0156$, giving a $\sim$4% relative uncertainty in the high beams.

Cost function

The cost function F implements "soft" constraints in terms of dose-volume histograms $D_r(y,\mathbf{x})$ for the organ-at-risk (OAR) and the target (see Ref. 17). $D_r(y,\mathbf{x})$ is the fraction of volume in region-of-interest (ROI) r absorbing a dose greater than y when the modulation profile is $\mathbf{x}$.

By "soft" constraints we mean that a penalty is assigned (as it is customary in such problems) for violating the constraints. The kth constraint is defined as in Ref. 5, in terms of maximal and minimal dose ($d_{rk}^{(\max)}$ and $d^{(\min)}$) to be delivered to a given fraction of the volume ($v_{rk}^{(\max)}$ and $v^{(\min)}$) of each ROI r. A minimal dose is defined only for the target.

The explicit expression for F is given by the sum of all the penalties due to excess dose to the different ROIs, and the penalty due to a defect of dose delivered to the target:

$$F(\mathbf{x}) = \sum_{r=1}^{R} F_r(\mathbf{x}) + w^{(\min)} \int_0^{d^{(\min)}} (y - d^{(\min)})^2 \times (D_{(\text{target})}(y,\mathbf{x}) - v^{(\min)})^2 \, dy,$$

where

PROOF COPY 020405MPH

PROOF COPY 020405MPH

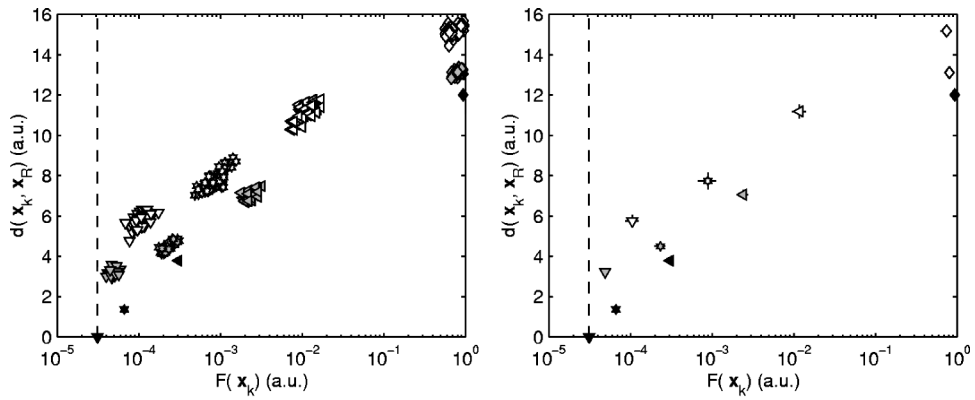

FIG. 1. Optimization plans at different stages. Left panel: the distributions of the optimized modulations at different stages. Each point in the plot represents the distance $d(\mathbf{x}_k,\mathbf{x}_R)$ and the cost $F(\mathbf{x}_k)$ of the modulation $\mathbf{x}_k$ resulting from the kth stage of the optimization processes. Different symbols correspond to different stages: (◇)—stage 0, initial condition; (◁)—stage 1, $\Delta=0.5$; (*)—stage 3, $\Delta=0.125$; (▽)—stage 6, $\Delta=1/2^6$. The implemented algorithm is identified by the colors of the symbols: white—single minimization trajectories are considered; gray—the modulations are averaged 3 by 3 at each stage; black—all the modulations are averaged at each step. The reference modulation $\mathbf{x}_R$ is the average over all the $\mathbf{x}_k$ in the last stage (black ▽). In the case shown $\mathbf{x}_k$ provides the plan with the lowest value of the cost function, the vertical dashed line in both plots. Right panel: the centroids of the distributions shown on the left are plotted. The vertical and horizontal bars are the widths (the standard deviation) of the clouds along the $d(\mathbf{x}_k,\mathbf{x}_R)$ and the F axes, respectively. Data used in this figure are only for illustrative purpose.

$$F_r(\mathbf{x})=\sum_{k=1}^{C} w_{rk}^{(\max)} \int_{d_{rk}^{(\max)}}^{D_r^{-1}(v_{rk}^{(\max)},\mathbf{x})} (y-d_{rk}^{(\max)})^2 \times (D_r(y,\mathbf{x})-v_{rk}^{(\max)})^2\, dy$$

and $D_r^{-1}(v_{rk}^{(\max)},\mathbf{x})$ is the dose corresponding to the maximal volume $v_{rk}^{(\max)}$. The different constraints are assigned different "weights" (importance factors) $w_{rk}^{(\max)}$ which might account, for example, for different clinical priorities. F_r is nonzero for all the ROIs r violating the dose-volume constraints $(D_r^{-1}(v_{rk}^{(\max)},\mathbf{x})\geqslant d_{rk}^{(\max)})$.

A DVH-based cost function allows a more robust and less case-dependent way to implement dose prescriptions.[5]

The computation of the absorbed dose is performed using a home-made treatment planning system, CARO,[18] based on a semi-empirical model developed at the Medical Physics Department of the "Regina Elena" Institute in Rome, Italy. The algorithm uses contours of the ROIs assuming different homogeneous densities for each volume in the TP. In practice, during the optimization process, for different beam modulations $\mathbf{x}$, the dose is computed as $\mathbf{Kx}$ where $\mathbf{K}$ is a preevaluated kernel assuming the dose absorption to be well approximated by a linear process.

Results are reported for a prostate case treated with a 10 MV photon beam using five fields (8×12 cm) performed with a 120-leaves dynamic multileaf collimator (0.5 cm in the central 20 cm of the field). The dose is computed on a grid whose voxel dimension is $0.42\times 0.25\times 0.5\ \mathrm{cm}^3$.

Tools for analysis

For each one of the initial conditions (modulation vector $\mathbf{x}_0$) the value of the cost function is recorded at different stages of the optimization process, as shown in Fig. 1. After completing the optimization cycle for many initial conditions a "reference" modulation $\mathbf{x}_R$ is defined either as the one corresponding to the smallest final F value reached, or as the mean over the final modulations reached by all the optimization trajectories. The kth stage of each optimization trajectory is then represented as a point in a plane with coordinates $F(\mathbf{x}_k)$ and $d(\mathbf{x}_k,\mathbf{x}_R)$, where $d(\mathbf{x}_k,\mathbf{x}_R)$ is the Euclidean distance between the reference modulation and the modulation $\mathbf{x}_k$ at stage k. As a result each optimization stage is represented by a "cloud" (the distribution of the end-points of the minimization trajectories; see Fig. 1—left), whose centroid (barycenter) approaches an asymptotic position in successive stages, with a corresponding monotonic decrease of the cost function F (see Fig. 1—right).

The vertical and horizontal bars are the standard deviations of the cloud along the distance $d(\mathbf{x}_k,\mathbf{x}_R)$ and the cost function $F(\mathbf{x}_k)$ directions.

Furthermore we test a "nonlocal" generalization of the algorithm in which for each cloud the modulation vectors are grouped and averaged n by n ($n=3$ in the figure). Such heuristics will be shown to provide an effective strategy to get a good estimate of the optimal solution. A limit case is the one in which one averages all the optimization histories at each stage (see black symbols in Fig. 1).

The initial modulations will be either chosen randomly around a reference modulation vector $\mathbf{x}_0=1$ or kept fixed as specified later.

Radiobiological assessment

Results of the optimization process have been evaluated in terms of a radiobiological measure of their quality, using tumor control probability (TCP) and normal tissue complication probability (NTCP) indexes.

The TCP is the probability of completely eradicating all clonogens from the local tumor site. The Poisson TCP model is the one most commonly used in radiotherapy, particularly in the linear-quadratic (LQ) form where the mean number of surviving clonogenic cells $N_s(V,D)$ in a volume V for an initial density ρ_0 and number $N_0=\rho_0 V$ after delivering a

PROOF COPY 020405MPH

PROOF COPY 020405MPH

uniform dose D in n fractions is given by the expression: $N_s(V,D)=\rho_0 V \exp(-\alpha D-\beta D^2/n)$. Accounting for the possible inhomogeneity of the dose distribution, we compute the TCP as the probability that no clonogenic cells survive in any of the subvolumes V_i, from the Poisson statistics:[19,20]

$$\mathrm{TCP}=\prod_i e^{-N_s(V_i,D_i)}.$$

The patient-dependent radio-sensitivity α is taken as a random variable from a normal distribution with mean $\langle\alpha\rangle=0.26\ \mathrm{Gy}^{-1}$ and standard deviation $\sigma_\alpha=0.06\ \mathrm{Gy}^{-1}$.[21] In the Results we test the impact of σ_α on the quality of the optimal treatment plan found. We do not consider other sources of variability either because they are negligible, giving an higher order contribution, or because they lie outside the main focus of the work. The α/β ratio is set to 8.33 Gy;[21] the appropriate α/β value is still debated, and recently lower values have been reported, closer to the best estimate for normal tissues ($\alpha/\beta\in[1.5, 6]$ Gy).[22–24] The dose per fraction has been taken as 2 Gy, the prostate carcinoma being a slowly proliferating tumor, and the density of clonogenic cells $\rho_0=5\times10^6\ \mathrm{cm}^{-3}$.

To evaluate the damage to the organs at risk we have adopted the relative seriality model[25] for calculating the NTCP. In such a model the volume effect is considered as a combination of both serial and parallel subunit organization. If we consider a heterogeneous dose distribution, the probability of damage to the whole organ is given by

$$\mathrm{NTCP}(\mathbf{x})=\left[1-\exp\left(\int_0^\infty \log(1-P(y)^s)\rho(y,\mathbf{x})\,dy\right)\right]^{1/s},$$

where

$$\log P(y)=-\exp[\mathrm{e}\gamma(1-y/D_{50})]\log 2,$$

$\rho(y,\mathbf{x})\,dy$ is the volume fraction of the OAR under study which absorbs a dose in the range $(y,y+dy)$ when the beam modulation is $\mathbf{x}$, s is the relative seriality of the OAR, γ is the "maximum normalized gradient" and D_{50} is the 50% response dose. The parameters are as follows: for the rectum $\gamma=2.2$, $D_{50}=80.0$ Gy, and $s=2.0$,[26] while for the bladder $\gamma=3.0$, $D_{50}=80.3$ Gy, and $s=1.3$.[25]

In Fig. 2 we illustrate how the above radiobiological indexes have been used in this work to analyze the results of the optimization process. The distribution of NTCP versus TCP values obtained for the optimal treatment plans is shown. A radiobiological effectiveness estimate is defined as the percentage of "success cases," enclosed in the bottom right corner, for which TCP>90% and NTCP<10%.

For the chosen constraints on the DVHs, the optimization process allowed a dose at the isocenter as high as 99 Gy, for which we get the best results in terms of percentage of success cases.

RESULTS

To characterize the complexity of the IMRT plan optimization problem we take in sequence the following steps:

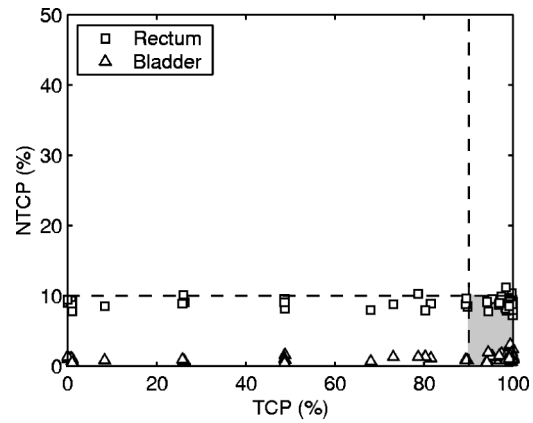

FIG. 2. Distributions of NTCP versus TCP for each OAR. Each symbol corresponds to the treatment plan resulting from a given optimization stage: the abscissa is the TCP for the plan and the ordinate is the NTCP for the considered OAR. For a plan there are as many points as the OARs taken into account. In the prostate case under study the OARs are the bladder (△) and the rectum (□): To each cloud in the left plot of Fig. 1 there correspond two distributions of symbols. A "success case" is a plan having a TCP above a threshold (90%), and a NTCP for all OARs below a limit value (10%). The percentage of success, the fraction of plans in the gray region, is a radiobiological measure of the quality of the optimization results at a given stage.

(1) We perform a large number of optimizations starting from a wide distribution of random initial conditions, and characterize the distribution of final solutions.
(2) We next estimate to what extent the distribution of final configurations is reduced by adopting the same initial condition for many optimization histories.
(3) Finally, from the gained insight into the shape of the cost function, we devise an approximate strategy to quickly infer the location of low-lying minima.

Inferring the distribution of suboptimal solutions

We now illustrate, using the representation described in the section "*Tools for analysis*," how the solutions found at successive stages of the optimization process are distributed in the configuration space of the modulation profiles. We choose a large number (91 in the examples shown) of randomly chosen initial conditions for the optimization process, and we record the sequence of configurations generated by each optimization history.

We then define as our "reference configuration" the average of the modulation profiles among the end-point configurations reached by the optimization histories. Such configuration mean consistently turned out to have a lower value of the cost function (i.e., the final configurations are systematically "less optimal" than their average, with the exception of the initial conditions). Figure 3 illustrates for the above reference configuration the DVHs for the tumor target and the considered OARs (rectum, bladder, and the body); for the same configuration, in Fig. 4, a sample dose distribution is shown across a section 1.5 cm from the isocenter, from which a high degree of the dose conformation of the tumor volume is apparent.

The centroids of the analyzed distributions are shown in Fig. 5(a). It is seen that in successive stages the distance

PROOF COPY 020405MPH

PROOF COPY 020405MPH

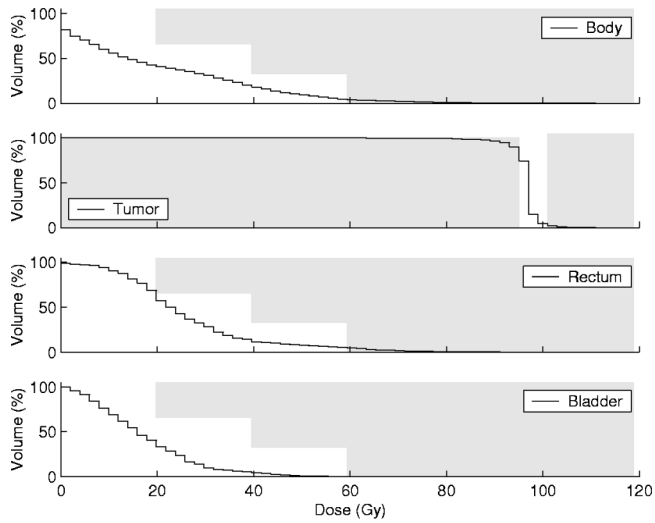

FIG. 3. DVHs for the tumor target and the considered OARs (rectum, bladder and the body) for the reference beam modulations. The shaded regions are the regions subject to penalty according to the dose-volume constraints derived from Ref. 17.

between the cloud of solutions and the reference configuration decreases linearly with decreasing values of the logarithm of the cost function. The black symbols which correspond to the averages at successive stages have consistently values of the cost function distinctly lower than the members of the cloud. The picture emerges of the end-points at each stage being symmetrically distributed around lower configurations, at a moderately fluctuating distance from them; the centroid of the clouds slowly drifts from a stage to the next. Whatever the underlying geometry of the cost function surface, the message of this first stage of analysis is that a single deterministic optimization trajectory would, at best, converge very slowly to a low-lying solution with high computational load, or be stopped by the intrinsic roughness of the surface.

Figures 5(b) and 5(c), characterize the clouds at successive stages and their centroids in terms of TCP, NTCP, and the related probability of success of the treatment plan. Panel (b) illustrates a sensitivity analysis of the probability of success with respect to the α radio-sensitivity parameter. It is seen that for radio-resistant cases even at the end of the optimization process there are high chances of a poor result of the plan. On the other hand, panel (c) shows that already at intermediate stages of optimization the estimated TCP and NTCP imply a success, for the modulations corresponding to the centroids of the clouds.

Figure 6 provides an indication of the computational effort needed to complete the analysis of Fig. 5, in terms of the average number of steps (computations of the cost function) needed to complete an optimization history. To compute the centroids of the successive clouds would therefore imply multiplying the number in Fig. 6 for each stage by the number of histories.

Dependence on the initial conditions

We now examine the implication of the stochastic component of the chosen optimization algorithm. The Δ parameter provides an effective measure of the ''resolution power'' with which the algorithm is able to sense the irregularities of the cost function surface. At any point of the F surface reached by the algorithm, the direction of the gradient of F determines the allowed directions (those for which the cost function decreases), as the unit vectors parallel to the coordinate axis, forming an acute angle with the gradient (remember that at each step the algorithm chooses only one pencil beam to be changed, which amounts to a move in the direction of one of the coordinate axes). On average, the

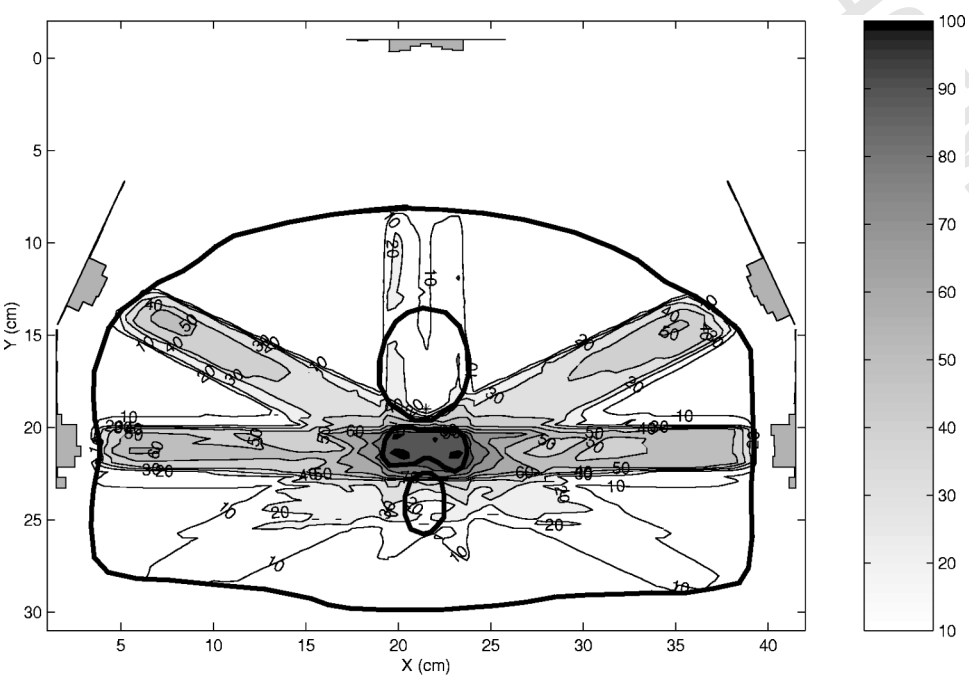

FIG. 4. Dose distribution across a section 1.5 cm from the isocenter for the reference beam modulations. The dose levels are coded according to the gray-scale color bar on the right. The modulation profiles and orientations of the five beams are also shown. The asterisk marks the isocenter. The thick curves enclose the body, the OARs (bladder and rectum) and the target: the latter is that containing the higher values of absorbed dose.

PROOF COPY 020405MPH

PROOF COPY 020405MPH

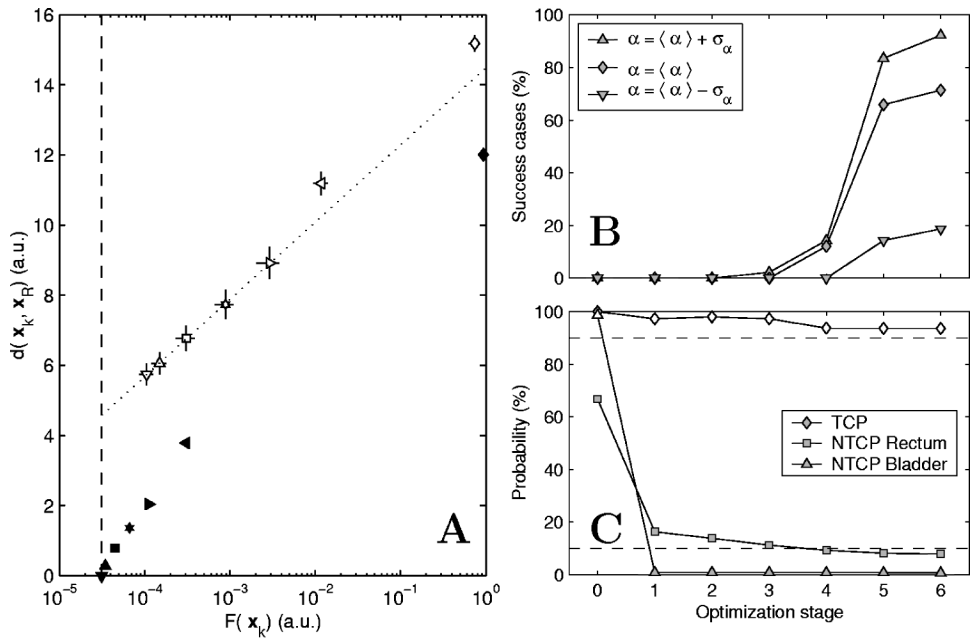

FIG. 5. A measure of the complexity of the optimization process. All the optimization histories start from random initial modulations $\mathbf{x}_0$ as in Fig. 1. (a) The centroids of the distributions (white symbols) and the averages of the modulations (black symbols) are reported at different stages: (◇)—stage 0, initial conditions; (◁)—stage 1, $\Delta=0.5$; (▷)—stage 2, $\Delta=0.25$; (*)—stage 3, $\Delta=0.125$; (□)—stage 4, $\Delta=0.0625$; (△)—stage 5, $\Delta=1/2^5$; (▽)—stage 6, $\Delta=1/2^6$. Dotted line is a linear fit of the centroids of the last five optimization stages: $y=p_1 \log x+p_2$, where $p_1=0.9552\pm0.0013$ and $p_2=14.48\pm0.031$. The asymptotic distance of the cloud centroid from the reference modulation (the intersection between the dashed and dotted lines) is $y_{\min}=4.572\pm0.031$. (b) Dependence of the radiobiological quality of the optimized plans on the tumor radiosensitivity: percentage of success cases from the distributions shown in (a) for different α ($\langle\alpha\rangle=0.26\ \mathrm{Gy}^{-1}$ and $\sigma_\alpha=0.06\ \mathrm{Gy}^{-1}$) across the optimization stages. (c) TCP, bladder NTCP and rectum NTCP (see the legend) for the average modulations [black symbols in panel (a)]: For the TCP model we used the worst $\alpha=\langle\alpha\rangle-\sigma_\alpha=0.20\ \mathrm{Gy}^{-1}$. The number of optimization histories is 100.

representative point on the cost function surface will move in the direction of the average over the above allowed unit vectors (in two dimensions, that would be the direction bisecting the quadrant containing the gradient vector). Given an initial value of Δ, small enough that the gradient of the cost function surface does not vary significantly over distances of order Δ, the average direction just defined will not fluctuate much, and the algorithm becomes essentially (locally) deterministic.

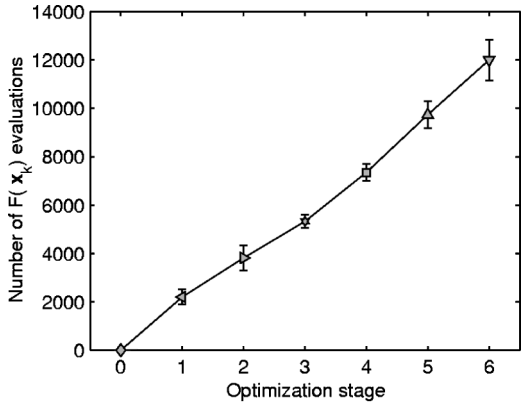

FIG. 6. Number of evaluations of the cost function needed to complete the different optimization stages. The averages (symbols) and the standard deviations (error bars) of the needed steps are computed over the 91 minimizations of Fig. 5.

Figure 7 repeats the analysis of the previous section for a very small initial value of Δ. This would correspond to effectively removing much of the residual stochasticity in the optimization history, and in the limit of a totally deterministic trajectory the end-point cloud would contract to a point (the nearest available local minimum). In Fig. 7 the reference configuration is again the average modulation of the final cloud. This case with small initial Δ can be viewed as one in which the (smaller) stochastic component makes the different trajectories emerging from the common zero initial condition to ''diffuse'' to a lesser extent.

In this more deterministic case the optimization trajectory ends up nearer to the common reference configuration; so, in this case there seems to be a ''good'' direction leading from the null initial condition to low-lying minima, and changing frequently the direction of motion on F worsens the result. In other words, from the chosen initial condition, with all pencil beams zero, there appears to be a groove available towards low-lying minima, such that the best strategy is to deviate as little as possible from the groove. In fact, the chosen initial condition is special in that it selects a wide subspace of the pencil beams which should ultimately be zero because they point towards an OAR, and the optimization process has to monotonically increase the modulation for the pencil beams hitting the target, until the required constraints are satisfied. The latter process would be visualized as sliding along the groove, while any move erroneously trying to increase the

PROOF COPY 020405MPH

PROOF COPY 020405MPH

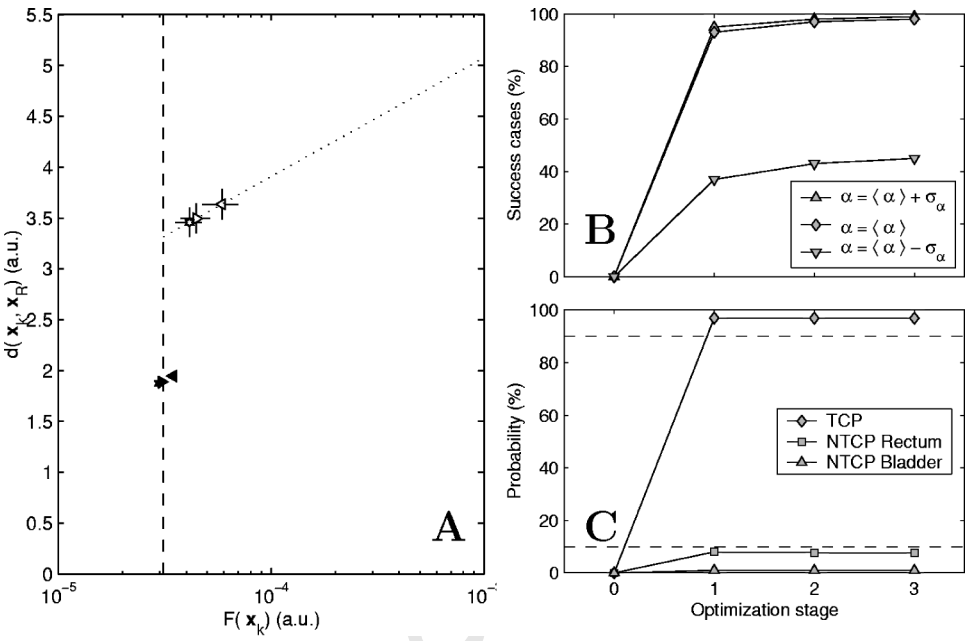

FIG. 7. Distributions of "refined" optimized histories at different stages starting from the same initial condition ($\mathbf{x}_0=0$). (a) The centroids of the distributions (white symbols) and the mean of the modulations (black symbols) are reported: (◇)—stage 0, initial condition (not visible for the chosen x-scale); (◁)—stage 1, $\Delta=0.0625$; (▷)—stage 2, $\Delta=1/2^5$; (*)—stage 3, $\Delta=1/2^6$. Dotted line—linear fit as in Fig. 5(a) where $p_1=0.5082\pm0.0056$, $p_2=8.587\pm0.055$ and $y_{\min}=3.3150\pm0.0002$. (b) Dependence of the radiobiological quality of the optimized plans on the tumor radiosensitivity, as in Fig. 5(b). (c) TCP, bladder NTCP, and rectum NTCP versus the optimization stage as in Fig. 5(c). One hundred optimization histories are used.

modulation of the other pencil beams would amount to trying to climb the walls.

As has been reported by various authors, an informed choice of the initial condition could greatly reduce the complexity of the optimization problem. In order to clarify this point a fluence field based on the shadow projected by the tumor along the beam directions has been analyzed. The obtained results are very similar to the case of a null starting modulation (data not shown): also in this case a subset of pencil beams are "frozen" in the good starting position, and the optimization process wanders in a narrow subspace (the above groove) where only the remaining degrees of freedom have to be adjusted. For this reason it is possible that both the null initial condition and the uniformly modulated pencil beam on the tumor shadow are in fact a good initial condition in most cases.

Panels (b) and (c) in Fig. 7 illustrate the radiobiological assessment for the present case. The probability of success [panel (b)] is improved with respect to the random initial conditions (Fig. 5): high probability of success is reached more quickly, and the sensitivity α is lowered. The TCP and NTCP estimates clearly show the implication of the peculiar initial condition: both start with a zero value, as they should given that no dose is delivered, to be contrasted with the case of random initial conditions, for which both TCP and NTCP start near 100%.

A heuristic approach to reducing the uncertainty of the solution

In the preceding analysis we consistently observed that, at a given stage of the optimization process, the average modulation over many sampled histories constitute a better modulation choice compared to all the single ones. This observation suggests a strategy for obtaining good approximate solutions with a reduced computational effort, compared to stochastic approaches such as the simulated annealing. Indeed, one is led to try performing a small number of optimizations, averaging the end-point modulations and using the average as the tentative solution of the optimization problem. Figure 8 illustrates to what extent such a strategy works in the case at hand.

For the same optimization trajectories of Fig. 5 (white and black symbols), at each stage we also plot in panel (a) the configurations resulting from averaging two by two (light gray symbols) and four by four (dark gray symbols) all modulation vectors. It is seen from panel (b) [(c)] that the distribution of the success probability in this case is comparable to that of Fig. 5, so that picking at random any two (four) light (dark) gray points would provide a good solution even in the worst case.

So, when no clue on a good initial condition is available, it is tempting to assume that a good strategy would be to perform a few optimizations starting from random initial conditions, and adopting as the solution the average of the final set of modulations found.

CONCLUSIONS

We have demonstrated that, for the case of a prostate IMRT therapy optimization, the multiplicity and distribution of suboptimal solutions can be estimated, and the impact of such distribution on the clinical quality of the treatment can be evaluated. As a byproduct of the analysis, we have also devised a way for quickly estimating near-optimal solutions.

The presence and relevance of local minima in TP optimization problems is a much debated subject,[1,3] and the qualitative indication is that, when beam orientations are not involved in the optimization procedure, local minima are present in general, but do not pose a real problem.[7,11] Other authors, however, have long applied stochastic approaches like the simulated annealing to TP optimization problems, known to efficiently deal with the local minima problem.

PROOF COPY 020405MPH

PROOF COPY 020405MPH

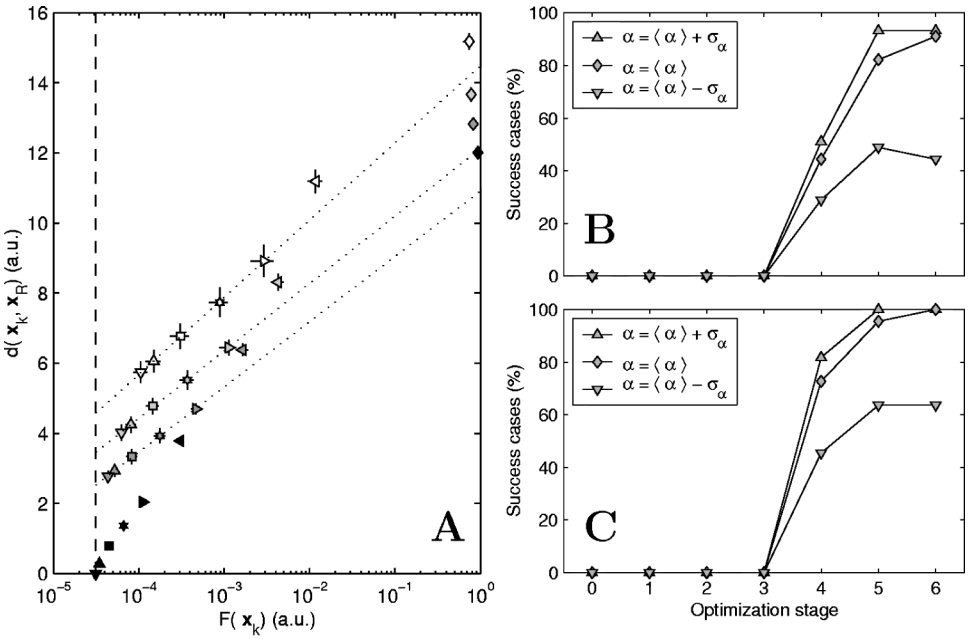

FIG. 8. Distributions of "coarse" optimized histories at different stages starting from random initial conditions as in Fig. 5 when modulations are grouped 2 by 2 (light gray symbols) and 4 by 4 (dark gray symbols). Other symbols and optimization stages as in Fig. 5. (a) Dotted line on light gray symbols—linear fit over the distribution centroids of the histories grouped 2 by 2 on the last five optimization stages: $y=p_1 \log x+p_2$, where $p_1=0.835\pm0.018$, $p_2=12.12\pm0.15$, and $y_{\min}=3.460\pm0.018$; Dotted line on dark gray symbols—linear fit over the 4 by 4 centroids of the last five optimization stages where $p_1=0.808\pm0.018$, $p_2=10.91\pm0.17$, and $y_{\min}=2.524\pm0.014$. (b) and (c) Susceptibility to the tumor radiosensitivity of the plans quality versus the optimization stages, respectively, for histories grouped 2 by 2 and 4 by 4 [see Fig. 5(b) for details].

Probing complexity of the cost function using a large number of random initial conditions for the optimization process is an established strategy in the wider domain of optimization theory.[12] The analysis of the results in a suited bidimensional space (distance from a reference modulation *versus* cost function) allowed us to introduce a complexity index for the IMRT plans. In agreement with other authors, we found a high degree of degeneracy in the solutions for the studied prostate case:[7,11,27] the local minima have very similar cost function values, even if the corresponding beam modulations can be quite different. In addition, our work proves that neglecting such differences, and then the presence of local minima, significantly affect the radiobiological quality of the TP, if a reasonable variability of the model parameters is considered. After we completed the present work, Ref. 14 report a qualitatively similar approach to evaluate the radiobiological impact of the optimization process in presence of local minima; the authors compare errors due to being trapped in a local minima to those due to a wrong choice of the convergence criterium. The context of the analysis in Ref. 14 is quite different from ours, once neither IMRT nor DVH-based cost functions are taken into account.

On the other hand we also found that the computationally expensive option of simulating annealing and similar approaches is not necessarily a mandatory choice. Indeed, we have proved that a much faster option is available in the examined case, due to specific symmetry properties of the problem. It remains to be investigated, of course, how general such features are. In particular, it should be checked that the end-point modulations found by the optimization process are isotropically distributed around low-lying minima, possibly unreachable either because of other local minima in between, or simply because it would take an enormous time to reach them.

In view of an assessment of the potential of the method, it will be relevant of course to check the extent to which the method and the results extend to other pathologies and are robust against all the sources of variability (variable anatomy among different patients, the same patient in time, etc.). Planned improvements include using a more accurate dose delivery algorithm like a Montecarlo simulation.

If proven to hold in other cases, the method could serve as a basis for a pathology-based evaluation of the need for refined and computationally expensive optimization procedures *vs* the gain in terms of the radiobiological quality of the treatment.

Another relevant issue relates to the actual precision (and reproducibility) with which a prescribed IMRT modulation profile can be implemented by the multi-leaf system,[28,29] since this sets the scale below which any differences in suboptimal solutions are ineffective.

ACKNOWLEDGMENTS

We are indebted to M. Benassi and S. Carpino for making the CARO software available to us. Many discussions with M. Benassi and S. Marzi were very helpful at various stages of the work. We are grateful to A. E. Nahum for a careful and critical reading of the manuscript.

[a]Electronic mail: mattia@iss.infn.it
[b]Electronic mail: paolo.delgiudice@iss.infn.it
[c]Author to whom correspondence should be addressed. Electronic mail: caccia@iss.it

[1]C. Nutting, D. P. Dearnaley, and S. Webb, "Intensity modulated radiation therapy: a clinical review," Br. J. Radiol. **73**, 459–469 (2000).
[2]Intensity Modulated Radiation Therapy Collaborative Working Group, "Intensity-modulated radiotherapy: current status and issues of interest," Int. J. Radiat. Oncol., Biol., Phys. **51**(4), 880–914 (2001).
[3]M. Iori, "Methods for physical and radiobiological optimisation in radiotherapy with intensity modulation," Phys. Medica **17**(2), 55–73 (2001).
[4]J. O. Desay, "Multiple local minima in radiotherapy optimization problems with dose-volume constraints," Med. Phys. **24**(7), 1157–1161 (1997).

PROOF COPY 020405MPH

PROOF COPY 020405MPH

[5] S. V. Spirou and C. S. Chui, ''A gradient inverse planning algorithm with dose-volume constraints,'' Med. Phys. **25**(3), 321–333 (1998).
[6] S. Webb, *The Physics of Conformal Radiotherapy* (Institute of Physics Publishing, Bristol, 1997).
[7] J. Llacer, J. O. Deasy, T. R. Bortfeld, T. D. Solberg, and C. Promberger, ''Absence of multiple local minima effects in intensity modulated optimization with dose-volume constraints,'' Phys. Med. Biol. **48**(2), 183–210 (2003).
[8] S. Webb, ''Optimization of conformal radiotherapy dose distributions by simulated annealing,'' Phys. Med. Biol. **34**(10), 1349–1370 (1989).
[9] S. M. Morrill, K. S. Lam, R. G. Lane, M. Langer, and I. I. Rosen, ''Very fast simulated reannealing in radiation therapy treatment plan optimization,'' Int. J. Radiat. Oncol., Biol., Phys. **31**(1), 179–188 (1995).
[10] S. Marzi, M. Mattia, P. Del Giudice, B. Caccia, and M. Benassi, ''Optimization of intensity modulated radiation therapy: assessing the complexity of the problem,'' Ann. Ist. Super. Sanità **37**(2), 225–230 (2001).
[11] Q. Wu and R. Mohan, ''Multiple local minima in IMRT optimization based on dose-volume criteria,'' Med. Phys. **29**, 1514–1527 (2002).
[12] S. A. Solla, G. B. Sorkin, and S. R. White, ''Configuration space analysis for optimization problems,'' in *Disordered Systems and Biological Organization*, volume 20 of *F*, edited by E. Bienenstock *et al.*, NATO ASI (Springer-Verlag, Berlin, 1986), pp. 283–293.
[13] C. G. Rowbottom and S. Webb, ''Configuration space analysis of common cost functions in radiotherapy beam-weight optimization algorithms,'' Phys. Med. Biol. **47**(1), 65–77 (2002).
[14] R. Jeraj, C. Wu, and T. R. Mackie, ''Optimiser convergence and local minima errors and their clinical importance,'' Phys. Med. Biol. **48**(17), 2809–2827 (2003).
[15] B. Caccia, P. Del Giudice, M. Mattia, M. Benassi, and S. Marzi, ''An evaluation of the overall sensitivity of the optimization process in the IMRT technique,'' Radiother. Oncol. **61**(Suppl. 1), S83 (2001).
[16] A. Brahme, *Radiation Therapy Physics* (Springer-Verlag, Berlin, 1995), pp. 209–246.
[17] B. Emami *et al.*, ''Tolerance of normal tissue to therapeutic irradiation,'' Int. J. Radiat. Oncol. **21**(1), 109–122 (1991).
[18] M. Benassi and R. Paoluzi, ''An empirical formula, continuous in field parameters and depth dose for axial dose for co-60 gamma radiation,'' Br. J. Radiol. **45**, 475 (1972).
[19] S. Webb and A. E. Nahum, ''A model for calculating tumour control probability in radiotherapy including the effects of inhomogeneous distributions of dose and clonogenic cell density,'' Phys. Med. Biol. **38**(6), 653–666 (1993).
[20] A. E. Nahum and B. Sanchez-Nieto, ''Tumour control probability modelling: basic principles and applications in treatment planning,'' Phys. Medica **17**(2), 13–23 (2001).
[21] A. E. Nahum, B. Movsas, E. M. Horwitz, C. C. Stobbe, and J. D. Chapman, ''Incorporating clinical measurements of hypoxia into tumor local control modeling of prostate cancer: implications for α/β ratio,'' Int. J. Radiat. Oncol., Biol., Phys. **57**(2), ■■–■■ (2003).
[22] J. Fowler, R. Chappel, and M. Ritter, ''Is α/β for prostate tumors really low?'' Int. J. Radiat. Oncol., Biol., Phys. **50**(4), 1021–1031 (2001).
[23] S. Levegrün, A. Jackson, M. J. Zelefsky, M. K. Skwarchuk, E. S. Venkatram, W. Schlegel, Z. Fuks, S. A. Leibel, and C. C. Ling, ''Fitting tumor control probability models to biopsy outcome after three-dimensional conformal radiation therapy of prostate cancer: pitfalls in deducing radiobiologic parameters for tumors from clinical data,'' Int. J. Radiat. Oncol., Biol., Phys. **51**(4), 1064–1080 (2001).
[24] J. Z. Wang, M. Guerrero, and X. A. Li, ''How low is the α/β ratio for prostate cancer?'' Int. J. Radiat. Oncol., Biol., Phys. **55**(1), 194–203 (2003).
[25] P. Kallman, A. Agren, and A. Brahme, ''Tumour and normal tissue responses to fractionated non-uniform dose delivery,'' Int. J. Radiat. Biol. **62**(2), 249–262 (1992).
[26] L. P. Muren, N. Jesben, A. Gustafsson, and O. Dahl, ''Can doseresponse models predict reliable normal tissue complication probabilities in radical radiotherapy of urinary bladder cancer? the impact of alternative radiation tolerance models and parameters,'' Int. J. Radiat. Oncol., Biol., Phys. **50**(3), 627–637 (2001).
[27] M. Alber, G. Meedt, F. Nüsslin, and R. Reemtsen, ''On the degeneracy of the imrt optimization problem,'' Med. Phys. **29**, 2584–2589 (2002).
[28] L. J. Verhey, ''Issues in optimization for planning of intensity-modulated radiation therapy,'' Semin Radiat. Oncol. **12**(3), 210–218 (2002).
[29] M. Langer, ''Application of coloring theory to reduce intensity modulated radiotherapy dose calculations,'' Med. Phys. **27**(9), 2077–2083 (2000).

PROOF COPY 020405MPH

PROOF COPY 020405MPH