Robust spatiotemporal fractionation schemes in the presence of patient setup uncertainty

Melissa R. Gaddy

Department of Mathematics, North Carolina State University, Raleigh, NC 27695-8205, USA

Jan Unkelbach

Department of Radiation Oncology, University Hospital Zürich, Zürich, CH 8091, Switzerland

Dávid Papp*

Department of Mathematics, North Carolina State University, Raleigh, NC 27695-8205, USA

(Dated: May 9, 2019)

Purpose. Spatiotemporal fractionation schemes for photon radiotherapy have recently arisen as a promising technique for healthy tissue sparing. Because spatiotemporally fractionated treatments have a characteristic pattern of delivering high doses to different parts of the tumor in each fraction, uncertainty in patient positioning is an even more pressing concern than in conventional uniform fractionation. Until now, such concerns in patient setup uncertainty have not been addressed in the context of spatiotemporal fractionation.

Methods. A stochastic optimization model is used to incorporate patient setup uncertainty to optimize spatiotemporally fractionated plans using expected penalties for deviations from prescription values. First, a robust uniform reference plan is optimized with a stochastic optimization model. Then, a spatiotemporal plan is optimized with a constrained stochastic optimization model that minimizes a primary clinical objective and constrains the spatiotemporal plan to be at least as good as the uniform reference plan with respect to all other objectives. A discrete probability distribution is defined to characterize the random setup error occurring in each fraction. For the optimization of uniform plans, the expected penalties are computed exactly by exploiting the symmetry of the fractions, and for the spatiotemporal plans, quasi-Monte Carlo sampling is used to approximate the expectation.

Results. Using five clinical liver cases it is demonstrated that spatiotemporally fractionated treatment plans maintain the same robust tumor coverage as a stochastic uniform reference plan and exhibit a reduction in the expected mean BED of the uninvolved liver. This is observed for a spectrum of probability distributions of random setup errors with shifts in the patient position of up to 5 mm from the nominal position. For probability distributions with small variance in the patient position, the spatiotemporal plans exhibit an 8-30% reduction in expected mean BED in the healthy liver tissue for shifts up to 2.5 mm and reductions of 5-25% for shifts up to 5 mm.

Conclusions. In the presence of patient setup uncertainty, spatiotemporally fractionated treatment plans exhibit the same robust tumor coverage as their uniformly fractionated counterparts and still retain the benefit in sparing healthy tissues.

* dpapp@ncsu.edu

I. INTRODUCTION

Spatiotemporal fractionation schemes have recently arisen as a promising treatment technique for photon radiotherapy for healthy tissue sparing [1, 2]. Unlike conventional fractionation, spatiotemporally fractionated treatments alter the dose distribution in each fraction in an attempt to optimally exploit the fractionation effect and approximate an ideal treatment that would simultaneously hypofractionate the tumor while delivering a conventionally fractionated dose to healthy tissues. Spatiotemporal fractionation schemes achieve this by delivering high single-fraction doses to different parts of the tumor in each fraction while avoiding a similar localized hypofractionation in the healthy tissue. Thereby, a lower cumulative physical dose is required compared to conventional treatments, which translates into a net reduction of biological dose in the normal tissue.

Previous work has demonstrated the efficacy of spatiotemporal fractionation and its potential to reduce the treatment impact on healthy tissues. It was initially proposed for proton therapy treatments [3, 4], with the rationale that the dose in the entrance region of a proton beam is approximately independent of the beam's range, and therefore the proximal and distal parts of the tumor can be irradiated to higher doses in separate fractions without similarly hypofractionating the entrance region. Subsequently, it was shown that healthy tissue sparing can also be achieved with spatiotemporal fractionation in photon radiotherapy. Specifically, the benefit for photon therapy has been demonstrated in stereotactic body radiotherapy of liver tumors in cases where the uninvolved liver is the main dose-limiting organ [5] and in radiosurgery for large cerebral arteriovenous malformations [2]. The rationale, however, is somewhat different than for proton treatments. In this case, using a sufficiently large number of beams (or arc therapy delivery), treatment plans can be created in such a way that the fractions deliver high single-fraction doses to complementary parts of the target volume while creating a similar dose bath in the surrounding normal tissue. The regions that receive dose in each fraction need not have an identifying geometric feature and are not predetermined; instead, they are identified by the treatment optimization algorithm.

A pressing question that was raised but not addressed in earlier studies on spatiotemporal planning is the impact of dose delivery uncertainty. In all prior work, the benefit from spatiotemporally fractionated photon plans is a result of their characteristic pattern of delivering high single-fraction doses to small parts of the tumor. The resulting sharp dose gradients need to be meticulously aligned in different fractions in order to avoid potentially compromising tumor control. As the regions receiving high dose in each fraction are not predetermined, but are the result of the treatment plan optimization, this issue cannot be addressed using margins around the treated volumes. In this work, we study the impact of random setup uncertainty on spatiotemporal treatments. We utilize stochastic optimization techniques to directly incorporate setup uncertainty in the treatment planning optimization problem. When this uncertainty is properly accounted for, spatiotemporal photon plans can be computed that achieve substantial normal tissue sparing compared to conventional plans while being as robust against random setup uncertainty as their conventionally fractionated counterparts. We briefly discuss other sources of uncertainty in Section IV.

The problem of patient setup uncertainty is not new in the context of IMRT and IMPT. For a recent, comprehensive review of the subject, the reader is referred to [6]. Previous work utilizes stochastic optimization techniques to handle uncertainties in radiotherapy treatments by constraining or optimizing the expected values of quantities such as quadratic penalty functions [7, 8], expected values of TCP-based objective functions [9], or objective functions based on an approximation of the expected cumulative dose [10]. Another stochastic optimization approach is to use a mean

and variance constraint for each voxel to ensure robust target coverage [11]. In the current work we use a scenario-based model with expectations of penalty functions to account for the uncertainty in patient positioning, and use, when necessary, scenario sampling to approximate the expected values. An important difference between spatiotemporal and conventional fractionation lies in the importance of random errors, i.e. setup errors that are different from fraction to fraction. In conventional fractionation, random errors have a minor impact compared to systematic errors, i.e. a setup error that is the same in each fraction. Since spatiotemporal fractionation delivers a distinct dose distribution in each fraction, random errors are both more important and computationally more challenging to account for.

A. Fractionation and the BED model

The most widely used mathematical model of the fractionation effect is the biologically effective dose (BED) model [12], which is based on the linear-quadratic model of a tissue's response to radiation over the course of multiple fractions. The *cumulative BED* delivered to a voxel v in a treatment with T fractions is defined as

$$b_v = \sum_{t=1}^{T} \left(d_{vt} + \frac{d_{vt}^2}{(\alpha/\beta)_v} \right), \tag{1}$$

where d_{vt} is the physical dose delivered to voxel v in fraction t, and $(\alpha/\beta)_v$ is a tissue-specific parameter; see, e.g., [13]. As usual in the conventional fluence map optimization problem of IMRT, the delivered physical dose d_{vt} is a linear function of the fluence map x_t delivered in fraction t. Formally, letting d_t be the vector of dose values $(d_{vt})_{v\in V}$, we have $d_t = Dx_t$, with the dose-influence matrix $D \in \mathbb{R}^{|V|\times B}$, where V is the set of patient voxels and B is the number of beamlets. Similarly, the vector of cumulative BED values is denoted $(b_v)_{v\in V}$.

The BED model postulates that treatments with identical cumulative BED are iso-effective in the tumor and iso-toxic in the healthy tissue. Clinical prescriptions specifying a number of (uniform) fractions and physical doses can thus be converted to cumulative BED prescriptions and thereby generalized to spatiotemporal fractionation schemes. The optimization models in this work utilize standard piecewise quadratic penalty functions to penalize the violation of clinical constraints. To penalize BED that falls below a prescription value b_v^{lo} or exceeds a prescription value b_v^{hi} in voxel v, the penalty functions are $(b_v^{\text{lo}} - b_v)_+^2$ and $(b_v - b_v^{\text{hi}})_+^2$, respectively, where $(y)_+$ denotes the positive part $\max\{0, y\}$. Similar expressions can be used to penalize, for example, excessive mean BED in an organ. Spatiotemporal planning models using a BED-based variant of the generalized equivalent uniform dose (gEUD) can also be considered (see, e.g., [14]), but will not be used in this paper.

B. Modeling dose delivery uncertainty

Random setup errors, along with other sources of uncertainty in dose delivery, can be modeled as uncertainty in the dose-influence matrix D. We use the notation Ω for the set of possible setup errors, or *scenarios*, and D^ω for the realization of D in scenario $\omega \in \Omega$. Under random setup uncertainty, the patient positioning error is potentially different in each fraction, and it is convenient to think of each scenario as a T-dimensional vector $\omega = (\omega_1, \ldots, \omega_T)$. The associated probability function is denoted by P. For simplicity, in this paper we will work with finitely many discrete scenarios, that is, $\Omega \subset \mathbb{R}^T$ is a finite set, and each scenario ω has a positive probability $P(\omega)$.

With the dose-influence matrix being a random variable, the physical dose and the cumulative BED in each voxel also become random variables. The cumulative BED for voxel v in scenario ω is

$$b_v^\omega = \sum_{t=1}^{T} \left(d_{vt}^{\omega_t} + \frac{(d_{vt}^{\omega_t})^2}{(\alpha/\beta)_v} \right). \tag{2}$$

As a result, the plan quality that we aim to optimize is also random. Following [15] and many others, we quantify plan quality by the expected value of an appropriately chosen penalty function that penalizes deviation from prescribed BED values, such as the piecewise quadratic penalty functions in Section I A above. For example, to ensure that the cumulative BED delivered to a particular voxel remains above a prescribed value b_v^{lo} with sufficiently high probability, we may minimize (or constrain) the quantity

$$\mathbb{E}_P \left[(b_v^{\text{lo}} - b_v)_+^2 \right]. \tag{3}$$

To lighten the notation, for the rest of the paper, we shall drop the subscript P from the expected values, as the probability distribution is always the same.

II. MATERIALS AND METHODS

A. Stochastic optimization of uniformly fractionated treatments

To ensure a fair comparison between spatiotemporal fractionation and uniform fractionation, uniformly fractionated treatment plans are optimized using the same BED-based prescriptions and the same stochastic model of random setup uncertainty as the spatiotemporal plans. The benefit of spatiotemporal fractionation can be evaluated by comparing the spatiotemporal plans to their robust uniformly fractionated counterparts. In our optimization models, all clinical objectives are implemented using standard piecewise quadratic penalty functions. Structures that are sufficiently distant from the target volumes are likely to be nearly or entirely unaffected by random setup uncertainty; it is sufficient to consider only the nominal scenario in the objectives involving these structures. The objectives concerning target structures and nearby volumes are implemented using expected piecewise quadratic penalties.

Let I^+ and I^- denote the index sets of the clinical objectives associated with structures that are assigned an expected penalty for the BED exceeding the prescription amount b^{hi} or BED falling below a prescription amount b^{lo}, respectively. Let I^m be the index set for the clinical objectives penalizing mean BED exceeding m^{hi}. Similarly, let $\bar{I}^+$, $\bar{I}^-$, and $\bar{I}^m$ denote the index sets of the clinical objectives that involve nominal penalty values. Let the set I be the union of the six aforementioned index sets. Let V_i be the set of voxels associated with the objective $i \in I$, with $V = \cup_{i \in I} V_i$ denoting the set of all patient voxels considered in the optimization model. Lastly, let $\bar{\omega}$ denote the nominal scenario.

Piecewise quadratic penalty functions $F_i(\cdot)$ are used to mathematically define the clinical objectives, either for the

nominal scenario $\bar{\omega}$ or the expected value over all scenarios:

$$F_i(b) = \begin{cases} \sum\limits_{v \in V_i} \left(b_v^{\bar{\omega}} - b_{iv}^{\mathrm{hi}}\right)_+^2 & \forall i \in \bar{I}^+ \\ \sum\limits_{v \in V_i} \left(b_{iv}^{\mathrm{lo}} - b_v^{\bar{\omega}}\right)_+^2 & \forall i \in \bar{I}^- \\ \left(\frac{1}{|V_i|} \sum\limits_{v \in V_i} b_v^{\bar{\omega}} - m_i^{\mathrm{hi}}\right)_+^2 & \forall i \in \bar{I}^m \\ \mathbb{E}\left[\sum\limits_{v \in V_i} \left(b_v - b_{iv}^{\mathrm{hi}}\right)_+^2\right] & \forall i \in I^+ \\ \mathbb{E}\left[\sum\limits_{v \in V_i} \left(b_{iv}^{\mathrm{lo}} - b_v\right)_+^2\right] & \forall i \in I^- \\ \mathbb{E}\left[\left(\frac{1}{|V_i|} \sum\limits_{v \in V_i} b_v - m_i^{\mathrm{hi}}\right)_+^2\right] & \forall i \in I^m. \end{cases}$$

Using this notation, the uniformly fractionated reference plans are obtained by solving optimization problems of the form

$$\begin{aligned} \min_{x,d,b} \quad & \sum_{i \in I} q_i F_i(b) \\ \text{s.t.} \quad & b_v^\omega = \sum_{t=1}^{T} \left(d_{vt}^{\omega_t} + \frac{(d_{vt}^{\omega_t})^2}{(\alpha/\beta)_v}\right) && \forall\, v \in V, \quad \omega = (\omega_1, \ldots, \omega_T) \in \Omega \\ & d_t^\omega = D^{\omega_t} x && \forall\, \omega \in \Omega, \quad t = 1, \ldots, T \\ & x \geq 0, \end{aligned} \tag{4}$$

with the penalty weights $q_i \geq 0$ reflecting the relative importance of each clinical objective. The variables d_{vt}^ω and b_v^ω can be eliminated using the equality constraints; this yields an optimization problem with the fluence x as the only variable, and the only constraint is the nonnegativity of the fluence.

In a similar manner to [4], we verify that the BED-based optimization model for uniformly fractionated treatment plans is convex in the domain of "reasonable" doses d for all typical α/β values and number of fractions T. Because the BED is a monotone increasing function of the physical dose and the piecewise quadratic penalties are convex functions of the BED, the only potentially problematic term of the objective function is the underdose penalty whose deterministic form is $f(b_v) = q\left(b_v^{\mathrm{lo}} - b_v\right)_+^2 = q\left(b_v^{\mathrm{lo}} - T\left(d_v + \frac{d_v^2}{\alpha/\beta}\right)\right)_+^2$ for a single voxel v. This is not a convex function of the physical dose. However, it can be seen that $\frac{\partial^2 f}{\partial d_v^2} > 0$ for physical dose values d_v that satisfy

$$d_v > -\frac{\alpha/\beta}{2} + \frac{1}{\sqrt{3}} \sqrt{\frac{(\alpha/\beta) b_v^{\mathrm{lo}}}{T} + \frac{(\alpha/\beta)^2}{4}}. \tag{5}$$

For the parameter values used in the current work ($b_v^{\mathrm{lo}} = 100$ Gy, $\alpha/\beta = 10$ for the tumor, and $T = 5$, as detailed later in Section II D), the above expression yields a lower bound of 3.66 Gy on the physical dose in order for the uniform model to be convex, which is not a binding constraint in the tumor, where the prescription BED b_v^{lo} and number of fractions T correspond to $d_v = 10$ Gy of physical dose per fraction. More generally, we may compare the lower bound in (5) with the prescribed per-fraction dose d_v^{lo} corresponding to a BED of b_v^{lo} delivered in T fractions, which is given

by the formula

$$d_v^{\text{lo}} = -\frac{\alpha/\beta}{2} + \sqrt{\frac{(\alpha/\beta)b_v^{\text{lo}}}{T} + \frac{(\alpha/\beta)^2}{4}}. \tag{6}$$

This reveals that regardless of the parameters b_v^{lo}, α/β, and T, the underdose penalty is convex for per-fraction dose values satisfying

$$d_v > d_v^{\text{lo}}/\sqrt{3}, \tag{7}$$

which is reasonably required from every acceptable uniformly fractionated plan.

In the stochastic model, the expected value of the underdose penalty is the probability-weighted linear combination of deterministic quadratic underdose penalties; thus, a sufficient condition for the convexity of the stochastic model (4) is that the per-fraction dose values satisfy (7) in every scenario.

B. Stochastic optimization model for spatiotemporal treatments

It is straightforward to formulate an optimization model analogous to (4) to compute an optimal robust spatiotemporal plan, with only two small changes to the model (4). First, the fluence map in each fraction t needs to be a separate vector of decision variables x_t, and then the corresponding random physical dose d_t^ω is computed as $d_t^\omega = D^{\omega_t} x_t$. One way to compare spatiotemporal treatment plans with uniform reference plans would be to optimize both plans for the same objective function (i.e. the same set of objective weights q_i). However, improvements in terms of quadratic penalty functions are difficult to interpret. In the application to liver tumors in Section II D, we instead quantify the benefit of spatiotemporal fractionation via the mean BED reduction in the healthy liver for a fixed BED delivered to the tumor. To that end, we minimize the objective function corresponding to the primary clinical objective (the mean liver BED) while constraining the remaining penalty function values to be at least as good as the penalty values in the uniform reference plan. This ensures that the benefit from spatiotemporal fractionation is measured in terms of a single, interpretable, scalar quantity representing the primary clinical objective, and that the measured benefit does not come at the cost of sacrificing other clinical objectives.

Next, we provide the mathematical details for this formulation. Without loss of generality, let V_1 denote the voxel set associated with the primary clinical objective and let F_1 be the corresponding penalty function. Let b^* denote the BED distribution of the uniformly fractionated reference plan, with $F_i(b^*)$ the value of the i-th penalty function evaluated with the uniform plan. For all clinical objectives aside from the primary clinical objective, we constrain the spatiotemporally fractionated treatment plan to have a smaller or equal penalty function value than the uniform reference plan. Then the optimization model for spatiotemporally fractionated plans can be written as follows:

$$\begin{aligned}
\min_{x,d,b} \quad & F_1(b) & & \\
\text{s.t.} \quad & F_i(b) \le F_i(b^*) & & \forall i \in I,\ i \neq 1 \\
& b_v^\omega = \sum_{t=1}^{T} \left(d_{vt}^{\omega_t} + \frac{(d_{vt}^{\omega_t})^2}{(\alpha/\beta)_v} \right) & & \forall v \in V, \quad \omega = (\omega_1, \ldots, \omega_T) \in \Omega \\
& d_t^\omega = D^{\omega_t} x_t & & \forall \omega \in \Omega, \quad t = 1, \ldots, T \\
& x_t \ge 0 & & t = 1, \ldots, T.
\end{aligned} \tag{8}$$

Unlike the optimization model for uniformly fractionated treatments, the spatiotemporal optimization model is nonconvex, and the solutions to this model can only be certified to be locally optimal. For spatiotemporal planning without uncertainty, it has been shown that in spite of the nonconvexity, locally optimal solutions achieve nearly all of the potential benefit of healthy tissue sparing with spatiotemporal fractionation [14, 16].

C. Modeling random patient setup uncertainty

We model uncertainty in patient positioning by considering scenarios in which the patient is shifted slightly from the nominal position.

1. *Probability distributions of patient position*

First, let us consider a one-dimensional probability distribution of the patient's position when the setup uncertainty is restricted to one axis of motion. For computational convenience, we assume that the patient is shifted an integer number of voxels on either side of the nominal position. Let p_s be the probability that the patient's position is s voxels away from the nominal position. (Thus, p_0 is the probability of the patient being in the nominal position.) We assume that the probabilities p_s satisfy

$$p_s = \gamma^{|s|} p_0 \quad s \in S \tag{9}$$

for some parameter $\gamma \in [0, 1]$ and a finite set $S \subset \mathbb{Z}$. In particular, the patient positioning error is assumed to be symmetric and bounded, with its mode in the nominal position. The parameter γ allows us to consider a spectrum of probability distributions. As γ increases, the probabilities of larger errors increase. When $\gamma = 0$, the probability of the nominal scenario is 1, and all other shifts have probability zero. At the other extreme, $\gamma = 1$, the nominal scenario and all shifts have an equal probability of occurring. In all of our experiments, $S \subset \{-2, -1, 0, 1, 2\}$.

We assume that the patient positioning errors along each axis are independent. Considering $|S|$ scenarios for the error along each of the n axes, the number of scenarios for the patient position in one fraction is $|S|^n$. In a treatment with T fractions, the number of scenarios is $|S|^{n \times T}$. For example, in some of our computational experiments, a five-fraction treatment plan is computed that allows the patient position to be shifted up to two voxels from the nominal position in either direction, in two dimensions. The number of distinct scenarios in these experiments is $5^{2 \times 5} \approx 10^7$.

2. *Computing the expected penalty values*

Because we use a finitely-supported probability distribution, the expected values of the penalty functions can be computed exactly as a finite sum over the scenarios. However, the large number of scenarios in the spatiotemporal models becomes prohibitive during optimization, as the time required to compute the penalty function values and their gradients is proportional to the number of scenarios. This can be mitigated by scenario sampling. Because the dimension of the uncertainty, nT, is small (in our experiments, $n = 2$ and $T = 5$), we use quasi-Monte Carlo (QMC) integration, rather than Monte Carlo, to sample the scenarios and approximate expected penalty values. QMC integration methods, which use evenly distributed pseudo-random samples from the domain of integration, are known

to have a better rate of convergence than Monte Carlo when the dimension of uncertainty (the number of integration variables) is small. In particular, we use a standard rank-one QMC lattice rule with a random shift [17].

We note that for uniformly fractionated plans, scenario sampling is not necessary to evaluate the expected penalty values. In a uniformly fractionated treatment, every permutation of the T shifts in a scenario $\omega = (\omega_1, \ldots, \omega_T)$ will yield the same cumulative BED because the same fluence map is delivered in each fraction, and the cumulative BED is not dependent on the order of the fractions. Thus, after a penalty is evaluated for a particular sequence of T shifts, its weight can be adjusted to account for all permutations of that shift sequence. Formally, let $\hat{\Omega} \subset \Omega$ be the set of scenarios $\hat{\omega} = (\hat{\omega}_1, \ldots, \hat{\omega}_T)$ such that every scenario $\omega = (\omega_1, \ldots, \omega_T) \in \Omega$ is a permutation of exactly one scenario in $\hat{\Omega}$; furthermore, for each $\hat{\omega} \in \hat{\Omega}$, let $N(\hat{\omega})$ be the number of scenarios in Ω that are a permutation of $\hat{\omega}$, and let $P(\omega)$ be the probability of scenario ω. Then the expectation of the penalty function F satisfies

$$\mathbb{E}\left[F(b_v)\right] = \sum_{\omega \in \Omega} P(\omega) F(b_v^{\omega}) = \sum_{\hat{\omega} \in \hat{\Omega}} N(\hat{\omega}) P(\hat{\omega}) F(b_v^{\hat{\omega}}). \tag{10}$$

Even with T as little as 5, $\hat{\Omega}$ is substantially smaller than Ω, which vastly reduces the number of function evaluations required to compute the exact expected value. For example, in our experiments with 2-dimensional patients and patient positioning errors up to 2 voxels in either direction along both axes, the number of scenarios can be reduced from $5^{10} \approx 10^7$ to 118 755. This eliminates the need for scenario sampling in the optimization models for uniformly fractionated treatments because the expected penalty value over all scenarios can be exactly computed in each iteration of the optimization algorithm.

D. Numerical experiments with clinical liver cases

We computed uniformly fractionated and spatiotemporally fractionated treatment plans for five two-dimensional slices of liver tumors that represent a variety of patient geometries. The same cases were also used in [16]. In this work, we considered patient setup errors in the two-dimensional transverse plane; we did not consider superior-inferior motion of the patient that would move the voxels out of the transverse slice.

1. Clinical liver cases and prescriptions

In Cases 1, 2, and 3, the lesion is centrally located within the liver, which is the primary dose-limiting organ. Case 1 has a large lesion within the liver, Case 2 has a small lesion, and Case 3 contains two separate lesions. In Case 4, the liver abuts the chest wall, and in Case 5, the chest wall, kidney, and GI tract are near the tumor. The patient geometries can be found in Figure 1.

In each of the five cases, an α/β ratio of 10 was used for the target structure and an α/β ratio of 4 was used for all healthy tissues. Five-fraction treatments were optimized to be consistent with common clinical practice. The optimization models incorporated the following clinical objectives:

- Minimize the expected mean BED in the uninvolved liver, defined as the liver minus the GTV.
- Penalize the expected shortfall from 100 Gy BED in the GTV. This corresponds to 50 Gy of physical dose delivered in 5 fractions.

- Minimize the nominal BED in the unclassified tissue.
- For Cases 4 & 5: Penalize excess of 96.25 Gy BED in the chest wall in the nominal scenario. This is equivalent to 35 Gy of physical dose in 5 fractions.
- For Cases 4 & 5: A conformal plan was desired, with a prescribed linear falloff from 175 Gy BED to 15 Gy BED in 3 cm in the nominal scenario. This is equivalent to a falloff from 50 Gy to 10 Gy of physical dose in 3 cm.
- For Case 5: Minimize the nominal BED in the kidneys.

BED excess and shortfall were penalized quadratically (from a threshold of 0 Gy if not specified otherwise), as detailed in Section II A.

The nominal dose-influence matrix D was obtained using the Quadrant Infinite Beam (QIB) dose calculation algorithm implemented in CERR version 5.2 [18]. In this work we considered two dose grids: one that was downsampled from the CT resolution by a factor of two, and one that matched the full CT resolution. For all five cases, the two-dimensional slices contain approximately 12,000 voxels in the downsampled cases and approximately 50,000 voxels in the full-resolution cases. Recall that the patient setup error is modeled as error in the dose-influence matrix, so each scenario $\omega \in \Omega$ has its corresponding matrix D^ω. Instead of repeating the dose calculation for every possible position of the patient to explicitly find each D^ω, we adopt the static dose cloud assumption, which is commonly used in treatment planning optimization [6, 19]. Under this assumption, the position of the patient does not affect the dose distribution, and if a shift of a patient occurs, the nominal dose distribution is delivered to the shifted patient. This assumption improves the computational efficiency: each D^ω can be assembled from the rows of the nominal D with no additional dose calculation. The static dose cloud assumption may not be applicable for superficial tumors or for tumors near interfaces of tissues with starkly different densities, but the assumption is reasonable for these clinical liver cases because the treatment area is composed of mostly homogeneous tissue. In cases with a tumor in the superior part of the liver close to the lung, it may be most practical to adopt the static dose cloud approximation for the treatment plan optimization and employ a more accurate dose calculation in the final stages of optimization and plan evaluation.

2. Probability distributions of patient shifts

In this work we considered a spectrum of probability distributions on three different supports. On the downsampled dose grid, we considered probability distributions supported on 1- and 2-voxel shifts, and the voxel sizes for these experiments were 2.54 mm × 2.54 mm for the first four cases and 2.18 mm × 2.18 mm for Case 5. With other treatment sites in mind that have smaller setup uncertainties, we also ran a set of experiments without downsampling and using the full CT resolution with probability distributions supported on 1-voxel shifts. These experiments used a voxel size of 1.27 mm × 1.27 mm for Cases 1–4 and 1.09 mm × 1.09 mm for Case 5. In all of the experiments, we assumed that the range and probability distribution of the random setup error is the same along the two axes of motion.

As described in Section II C 1, the parameter γ is used to adjust the probabilities of each point in the support of the distribution. Values of γ that correspond to several values for the variance σ^2 of the shifts are chosen to illustrate the

benefit of spatiotemporal fractionation across a spectrum of probability distributions. (The variance is measured in units of voxels instead of mm.) In the distributions of 1-voxel shifts, the variance ranges from 0 to a maximum value of $2/3$, and for the 2-voxel distributions, the variance ranges from 0 to 2. This corresponds to nominal optimization when the variance is 0 to a uniform probability distribution of the shifts when the variance is maximized.

3. Optimizing robust uniformly fractionated reference plans

Uniformly fractionated treatment plans with 21 equispaced coplanar beams were optimized to approximate a high-quality VMAT plan delivering the modulated fluence maps over multiple small arc sectors [20]. We optimized the beamlet weights of a single row of a 1 cm × 1 cm grid of beamlets. Each of the five cases had a few hundred beamlets. Uniformly fractionated treatment plans were optimized for six values of σ^2 for 2-voxel shifts and five values of σ^2 for 1-voxel shifts. The optimization problem (4) in Section II A was solved to global optimality using the solver L-BFGS-B [21]. As detailed in Section II C 2, the quasi-Monte Carlo approximations for the expected GTV penalty are not necessary because the exact penalty values can be computed by taking advantage of the symmetry of the fractions. For probability distributions supported on 1-voxel shifts, the exact penalties can be calculated with 1287 scenarios, and the optimization runtimes were on the order of a few minutes. For probability distributions supported on 2-voxel shifts, the calculation of the exact penalty values required 118 755 scenarios, which yielded runtimes of up to 28 hours.

The uniformly fractionated treatment plans served as high-quality reference plans that were used as a point of comparison for the spatiotemporally fractionated plans optimized in the next section. To ensure that the uniformly fractionated plans were sufficiently robust, the penalty weights q_i in the objective function of (4) were adjusted so that in each of the cases, there was a 95% probability that 95% of the GTV received the prescribed 100 Gy BED.

4. Optimizing robust spatiotemporal plans

After the uniformly fractionated reference plans were computed, the optimization problem (8) in Section II B was solved to compute robust spatiotemporal treatment plans. Because of the large number of scenarios in the spatiotemporal models, we used quasi-Monte Carlo sampling to approximate the expected GTV penalty, as described in Section II C 2. Locally optimal spatiotemporal plans were obtained by solving the optimization problem (8) using an in-house implementation of the augmented Lagrangian method, which is an iterative method for solving nonlinearly constrained nonlinear optimization problems with a primal-dual approach; for details, the reader is referred to [22, Sec. 6.4]. We used a scenario sample of 2048 lattice points for each iteration of the augmented Lagrangian algorithm and a large sample of 524 288 lattice points to evaluate the constraint violation for the multiplier updates. Runtimes ranged from a few hours with the 1-voxel shift distributions to 43 hours for the 2-voxel shift distributions with the largest value of gamma.

5. Computing equieffective dose

For an interpretable comparison between spatiotemporally fractionated treatment plans and their uniformly fractionated reference plans, we compute the equivalent dose DEQ5 of the expected BED. The DEQ5 is the physical dose distribution that if delivered in 5 uniform fractions would yield the same BED as the expected BED for each voxel; i.e. the DEQ5 is the quantity satisfying

$$\mathbb{E}\left[5d\left(1+\frac{d}{\alpha/\beta}\right)\right] = \mathrm{DEQ5} + \frac{(\mathrm{DEQ5})^2/5}{\alpha/\beta}$$

for uniformly fractionated plans and

$$\mathbb{E}\left[\sum_{t=1}^{5}\left(d_t+\frac{d_t^2}{\alpha/\beta}\right)\right] = \mathrm{DEQ5} + \frac{(\mathrm{DEQ5})^2/5}{\alpha/\beta}$$

for spatiotemporal plans.

III. RESULTS

The benefit of spatiotemporal fractionation is measured by the reduction in expected mean liver BED in the healthy liver tissue from the uniform reference plan. By construction, the spatiotemporal plans have the same plan quality as the uniform reference plan with respect to all other clinical objectives. Tables I, II, and III summarize the mean liver BED reductions in spatiotemporal plans for all of the cases and probability distributions in our numerical experiments. Table I contains the reductions for probability distributions supported on shifts up to 2 voxels (approximately 5 mm) in each axis, and Tables II and III display the reductions in mean liver BED of spatiotemporal plans for probability distributions supported on 1-voxel shifts for downsampled (approximately 2.5 mm) and full-resolution (approximately 1.25 mm) dose grids. For reference, the tables also report the mean liver BED reductions that could be obtained in the case that no uncertainty is present.

As expected, the benefit of spatiotemporal fractionation decreases with larger setup uncertainty. However, when the variance in the patient position is small, spatiotemporally fractionated plans maintain a large reduction in expected mean liver BED. With the downsampled dose grid, for the smallest nonzero σ^2, Case 4 displays the largest reduction of 24.9% for 2-voxel shifts, while Case 5 has the smallest reduction of 4.9%. The former corresponds to a DEQ5 reduction from 19.3 to 16.6 Gy, while the latter corresponds to a DEQ5 reduction from 33.0 to 32.1 Gy. The reductions are even larger for spatiotemporal plans with 1-voxel shifts on both the downsampled and the full-resolution dose grid. This suggests that the magnitude of the largest undetected random patient setup error is just as important as the variance of the setup error.

As the variance σ^2 increases, the fractions of the spatiotemporal plans become increasingly uniform. As seen in Figure 2a, the fractional doses are highly modulated for the smaller variances, i.e. each fraction delivers a very high dose to complementary parts of the tumor. As the variance increases, the fractional dose distributions become more uniform, i.e. each fraction delivers a significant dose to most of the target volume. This trend is expected because more uniform dose distributions help maintain robustness against greater misalignment of the fractional dose distributions. This explains the decreasing benefit of spatiotemporal fractionation over uniform reference plans as σ^2

increases. These trends in the dose distributions were observed in all of the five cases; see additional figures in the Supplement.

In Figure 2b, the DEQ5 of the expected BED for Case 3 is plotted for the spatiotemporal plan for three σ^2 values. The robust treatment plans exhibit a "margin" of higher dose around the GTV to maintain robustness despite the fact that there is no explicit margin, such as the PTV, contoured around the GTV. As seen in the figure, the width of the "margin" around the GTV increases as the variance σ^2 of patient shifts increases. The stochastic optimization model adjusts the width of the margin around the GTV depending on the probabilities of patient shifts, and this margin is optimized to exactly maintain robustness but also to avoid unduly irradiating the healthy liver with too large a margin.

Figure 3 shows the evolution of the optimal treatment plans as the variance of patient shifts increases by comparing the DEQ5 of the expected BED for uniform and spatiotemporal plans. As σ^2 increases, the spatiotemporal plan becomes more similar to the uniformly fractionated plan, and the benefit in healthy tissue sparing decreases. This trend can also be observed in the dose-volume histogram (DVH) curves in Figure 4. In this figure, we observe that as σ^2 increases, the curves for the uniform and spatiotemporal plans move closer to one another.

A key feature of the robust spatiotemporal plans is that the reduction in mean BED in the healthy liver tissue occurs simultaneously with equally robust coverage of the tumor. In all of the cases presented in this work, the spatiotemporal plan meets the same robustness criterion as its uniform reference plan: there is at least a 95% probability that 95% of the GTV receives 100 Gy BED. In most cases we found that even though the spatiotemporal plan and the uniform plan have equal expected GTV penalties, the spatiotemporal plans generally have a higher probability that 95% of the GTV receives the prescription 100 Gy BED. In addition, in most cases the spatiotemporal plan has a higher expected mean GTV BED than the uniformly fractionated reference plan.

The similar robustness between the two fractionation schemes can be seen in Figure 5, which displays a DVH plot for the DEQ5 values comparing the uniform reference plan and the spatiotemporally fractionated treatment plan for Case 3. Figure 5 displays a "cloud" of DVH curves that shows the DVH curve for every scenario in a sample with 4096 scenarios. The spatiotemporal plan exhibits more variability in the tumor, as evidenced through a greater spread of DVH curves, but we notice that the variability occurs in the region of higher DEQ5. The spatiotemporal plan maintains the same robust GTV coverage as the uniform reference plan in the critical region around DEQ5 of 50 Gy. DVH plots similar to Figure 4 and Figure 5 for the remaining cases can be found in the Supplement.

IV. DISCUSSION

Benefit and potential applications of spatiotemporal fractionation. Spatiotemporal fractionation improves the ratio of tumor versus normal tissue BED by achieving partial hypofractionation in the tumor along with more uniform fractionation in normal tissues. Consequently, potential clinical applications of spatiotemporal fractionation are tumors that are eligible for hypofractionation, i.e. lesions treated with SBRT and SRS. In contrast, treatment sites that rely on fractionation to protect dose-limiting normal tissues within the target volume, such as most head and neck cancers or glioblastoma, are not suited. In addition, spatiotemporal fractionation mostly improves on the mean BED to the surrounding normal tissues rather than the maximum BED adjacent to the target volume. Based on these considerations, large liver tumors where the prescription dose is limited by the mean dose to the non-involved

liver represent one of the most promising clinical applications of spatiotemporal fractionation - except for the problem of motion uncertainty, which is particularly pressing for liver tumors.

Prior to this work, the potential benefit of spatiotemporal fractionation was shown, but the problem of uncertainty was previously unaddressed. In this work we have demonstrated that spatiotemporal fractionation schemes retain a portion of their benefit in healthy tissue sparing even in the presence of random patient setup uncertainty. As expected, the benefit of spatiotemporal fractionation decreases with increasing setup uncertainty. However, even with patient positioning errors of up to 5 mm along both axes, spatiotemporal fractionation schemes yielded a reduction in the expected mean liver BED over uniformly fractionated plans in each case, without any compromise in robustness or in the other clinical objectives. For smaller setup errors of 1-2 mm, which are applicable for intracranial lesions, a large portion of the benefit remains. Therefore, brain lesions treated with SRS such as large AVMs, benign brain tumors, and large cerebral metastases represent promising applications of spatiotemporal fractionation.

Practical treatment planning. Spatiotemporal plans computed in this study were obtained following a two-step procedure: the computation of a BED-based uniformly fractionated reference plan followed by the optimization of the spatiotemporal plan using model (8). The primary limitation of this procedure for practical treatment planning is the computation time required for optimization, in particular the time required to compute the optimal solution to model (8). This is a constrained optimization model with nonlinear nonconvex constraints, which we were only able to solve in reasonable time after sufficient simplifications to the model (two-dimensional geometry, coarse MLC resolution, static dose cloud assumption). In this study, the two-step approach was necessary to carefully quantify the benefit of spatiotemporal fractionation over conventional IMRT. However, in a clinical setting, where it is not necessary to compare the two fractionation schemes, spatiotemporal plans could be computed by simply minimizing a weighted sum of (BED-based) piecewise quadratic penalty functions subject to no constraints other than the nonnegativity of the fluence, similarly to model (4). These optimization problems can be solved much faster; in fact, our augmented Lagrangian algorithm to solve model (8) relies on iteratively solving a large number of bound-constrained optimization problems. The primary aim of this work is to establish that robust spatiotemporal treatments can maintain a benefit in healthy tissue sparing. Future studies evaluating the feasibility of incorporating spatiotemporal planning in the clinical workflow may focus on more efficient computational methods for spatiotemporal planning under uncertainty using the bound-constrained formulation, which should allow the computation of spatiotemporal treatment plans in a more realistic setting (three-dimensional geometry, high-resolution beamlet and dose grid).

Measures of robustness. In our optimization models, the expected value of the piecewise quadratic penalty of underdose in the GTV is used as the primary measure of the robustness of a treatment plan. The constraints in the spatiotemporal optimization model are active at the solution, which means that the expected penalty values for the GTV are equal in the uniform and spatiotemporal treatment plans. Thus, with respect to the expected GTV penalty robustness measure, the uniformly fractionated treatment plans and the spatiotemporal plans are equally robust. In practice, robustness of a treatment plan is often assessed with a DVH-like criterion, such as the probability that 95% of the GTV receives the prescription BED [23, 24]. An exact optimization model with a DVH-based robustness measure would require a very large mixed-integer program with binary variables for all of the GTV voxels in every scenario, which is computationally intractable. In our models we used the expected GTV penalty to have a tractable measure of robustness, but after the optimization we also evaluated the robustness of the optimal solutions using DVH criteria. We have found that in nearly every experiment, the spatiotemporal plans are more robust than the

uniform reference plans with respect to this robustness measure. Thus, the healthy tissue sparing could potentially be even higher than reported if the spatiotemporal plans and the uniform plans were equally robust with respect to the DVH-based robustness criterion.

Other sources of uncertainty. In this work, we focus on random setup uncertainty. Setup errors that are different between fractions yield a misalignment of BED contributions from different fractions, and hence represent the main concern in spatiotemporal fractionation. In addition, random errors could not be accounted for via margins because the dose distributions per fraction are determined during the optimization and not beforehand. In future work, systematic errors should also be included in robust treatment plan optimization. However, systematic errors lead to a shift (rather than a degradation) of the cumulative BED distribution. Therefore, systematic errors affect spatiotemporal treatments in the same way as conventional treatments, and could, in principle, be mitigated via PTV margins. For the application to liver tumors, intrafraction motion of the tumor due to respiration represents an additional uncertainty. However, modern treatment devices such as MR-linacs, which allow for gated treatment based on real-time MR imaging, may be used to mitigate such uncertainties.

V. CONCLUSION

We have demonstrated that in the presence of random patient setup uncertainty, robust spatiotemporally fractionated treatments can be optimized using a stochastic optimization model. These treatments maintain the same robust tumor coverage as their uniformly fractionated counterparts, while achieving a benefit in healthy tissue sparing.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. DMS-1719828.

DISCLOSURES

The authors have no conflicts of interest to disclose.

REFERENCES

[1] J. Unkelbach, Non-uniform spatiotemporal fractionation schemes in photon radiotherapy, in *World Congress on Medical Physics and Biomedical Engineering, June 7-12, 2015, Toronto, Canada*, edited by D. A. Jaffray, pages 401–404, Springer International Publishing, 2015, http://dx.doi.org/10.1007/978-3-319-19387-8_97.

[2] J. Unkelbach, M. R. Bussière, P. H. Chapman, J. S. Loeffler, and H. A. Shih, Spatiotemporal Fractionation Schemes for Irradiating Large Cerebral Arteriovenous Malformations, International Journal of Radiation Oncology*Biology*Physics **95**, 1067 – 1074 (2016).

[3] J. Unkelbach, C. Zeng, and M. Engelsman, Simultaneous optimization of dose distributions and fractionation schemes in particle radiotherapy, Medical Physics **40**, 091702 (2013).

[4] J. Unkelbach and D. Papp, The emergence of nonuniform spatiotemporal fractionation schemes within the standard BED model, Medical Physics **42**, 2234–2241 (2015), https://doi.org/10.1118/1.4916684.

[5] J. Unkelbach, D. Papp, M. R. Gaddy, N. Andratschke, T. Hong, and M. Guckenberger, Spatiotemporal fractionation schemes for liver stereotactic body radiotherapy, Radiotherapy and Oncology **125**, 357–364 (2017).

[6] J. Unkelbach, M. Alber, M. Bangert, R. Bokrantz, T. C. Chan, J. O. Deasy, A. Fredriksson, B. L. Gorissen, M. van Herk, W. Liu, H. Mahmoudzadeh, O. Nohadani, J. V. Siebers, M. Witte, and H. Xu, Robust radiotherapy planning, Physics in Medicine & Biology **63**, 22TR02 (28pp) (2018).

[7] J. Unkelbach and U. Oelfke, Incorporating organ movements in IMRT treatment planning for prostate cancer: Minimizing uncertainties in the inverse planning process, Medical Physics **32**, 2471–2483 (2005).

[8] R. Bohoslavsky, M. G. Witte, T. M. Janssen, and M. van Herk, Probabilistic objective functions for margin-less IMRT planning, Physics in Medicine & Biology **58**, 3563–3580 (2013).

[9] M. G. Witte, J. van der Geer, C. Schneider, J. V. Lebesque, M. Alber, and M. van Herk, IMRT optimization including random and systematic geometric errors based on the expectation of TCP and NTCP, Medical Physics **34**, 3544–3555 (2007).

[10] D. Fontanarosa, H. P. van der Laan, M. Witte, G. Shakirin, E. Roelofs, J. A. Langendijk, P. Lambin, and M. van Herk, An in silico comparison between margin-based and probabilistic target-planning approaches in head and neck cancer patients, Radiotherapy and Oncology **109**, 430–436 (2013).

[11] M. Chu, Y. Zinchenko, S. G. Henderson, and M. B. Sharpe, Robust optimization for intensity modulated radiation therapy treatment planning under uncertainty, Physics in Medicine and Biology **50**, 5463–5477 (2005).

[12] J. F. Fowler, 21 years of biologically effective dose, British Journal of Radiology **83**, 554–568 (2010).

[13] A. Niemierko and M. Goitein, Modeling of normal tissue response to radiation: The critical volume model, International Journal of Radiation Oncology, Biology, Physics **25**, 135–145 (1992).

[14] A. Adibi and E. Salari, Spatiotemporal radiotherapy planning using a global optimization approach, Physics in Medicine & Biology **63**, 035040 (2018).

[15] J. Unkelbach and U. Oelfke, Inclusion of organ movements in IMRT treatment planning via inverse planning based on probability distributions, Physics in Medicine & Biology **49**, 4005–4029 (2004).

[16] M. R. Gaddy, S. Yildiz, J. Unkelbach, and D. Papp, Optimization of spatiotemporally fractionated radiotherapy treatments with bounds on the achievable benefit, Physics in Medicine & Biology **63**, 015036 (2018).

[17] P. L'Ecuyer and C. Lemieux, Variance Reduction via Lattice Rules, Management Science **46**, 1214–1235 (2000).

[18] J. O. Deasy, A. I. Blanco, and V. H. Clark, CERR: A computational environment for radiotherapy research, Medical Physics **30**, 979–985 (2003), https://doi.org/10.1118/1.1568978.

[19] M. Sharma, E. Weiss, and J. V. Siebers, Dose deformation-invariance in adaptive prostate radiation therapy: Implication for treatment simulations, Radiotherapy and Oncology **105**, 207–213 (2012).

[20] D. Papp and J. Unkelbach, Direct leaf trajectory optimization for volumetric modulated arc therapy planning with sliding window delivery, Medical Physics **41**, 011701 (2014), http://dx.doi.org/10.1118/1.4835435.

[21] C. Zhu, R. H. Byrd, and J. Nocedal, L-BFGS-B, FORTRAN routines for large scale bound constrained optimization, ACM Transactions on Mathematical Software **23**, 550–560 (1997).

[22] A. Ruszczyński, *Nonlinear Optimization*, Princeton University Press, Princeton, NJ, 2005.

[23] M. van Herk, Errors and margins in radiotherapy, Seminars in Radiation Oncology **14**, 52–64 (2004).

[24] J. J. Gordon, N. Sayah, E. Weiss, and J. V. Siebers, Coverage optimized planning: Probabilistic treatment planning based on dose coverage histogram criteria, Medical Physics **37**, 550–563 (2010).

		σ^2					
	Case description	0	$\frac{1}{4}$	$\frac{1}{2}$	1	$\frac{3}{2}$	2
1	Large central lesion	26.8	8.3	5.5	4.1	3.0	2.7
2	Small lesion	32.4	13.6	12.3	8.8	7.9	6.9
3	Two small lesions	28.6	11.6	9.4	7.5	6.0	4.6
4	Lesion abutting ch. wall	50.1	24.9	19.2	14.1	11.6	10.7
5	Lesion abutting GI tract	23.6	4.9	3.0	1.1	1.1	1.0

TABLE I. Percent reductions in expected mean liver BED of spatiotemporal plans over uniformly fractionated plans for probability distributions supported on 2-voxel shifts in either axis of the transverse plane. The variance σ^2 of patient shifts ranges from 0, where the nominal position has probability 1, to 2, where there is an equal probability of shifting to the 25 points around the nominal position. The voxel sizes are 2.54 mm × 2.54 mm for Cases 1-4 and 2.18 mm × 2.18 mm for Case 5.

		σ^2				
	Case description	0	$\frac{1}{6}$	$\frac{1}{3}$	$\frac{1}{2}$	$\frac{2}{3}$
1	Large central lesion	26.8	10.1	8.3	6.8	6.1
2	Small lesion	32.4	16.6	16.4	15.8	15.1
3	Two small lesions	28.6	14.5	12.9	12.0	11.2
4	Lesion abutting ch. wall	50.1	29.9	25.7	23.5	21.0
5	Lesion abutting GI tract	23.6	7.7	5.4	4.2	3.2

TABLE II. Percent reductions in expected mean liver BED of spatiotemporal plans over uniformly fractionated plans for probability distributions supported on 1-voxel shifts in either axis of the transverse plane. The variance σ^2 of patient shifts ranges from 0, where the nominal position has probability 1, to $2/3$, where there is an equal probability of shifting to the 9 points around the nominal position. The voxel sizes are 2.54 mm × 2.54 mm for Cases 1-4 and 2.18 mm × 2.18 mm for Case 5.

		σ^2				
	Case description	0	$\frac{1}{6}$	$\frac{1}{3}$	$\frac{1}{2}$	$\frac{2}{3}$
1	Large central lesion	25.6	17.5	15.0	13.2	12.4
2	Small lesion	34.4	28.5	26.0	24.6	23.6
3	Two small lesions	28.6	22.8	20.7	19.4	18.6
4	Lesion abutting ch. wall	50.0	43.0	39.8	37.6	35.9
5	Lesion abutting GI tract	25.5	17.6	14.4	12.7	11.8

TABLE III. Percent reductions in expected mean liver BED of spatiotemporal plans over uniformly fractionated plans for probability distributions supported on 1-voxel shifts in either axis of the transverse plane with the full-resolution dose grid. The variance σ^2 of patient shifts ranges from 0, where the nominal position has probability 1, to $2/3$, where there is an equal probability of shifting to the 9 points around the nominal position. The voxel sizes are 1.27 mm × 1.27 mm for Cases 1-4 and 1.09 mm × 1.09 mm for Case 5.

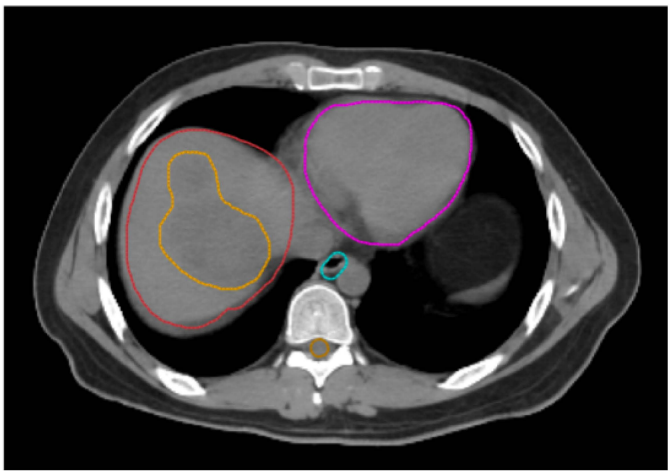

Case 1: Large central lesion

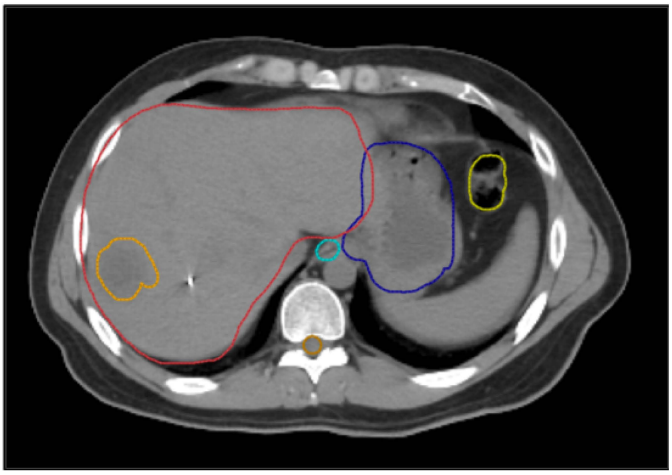

Case 2: Small lesion

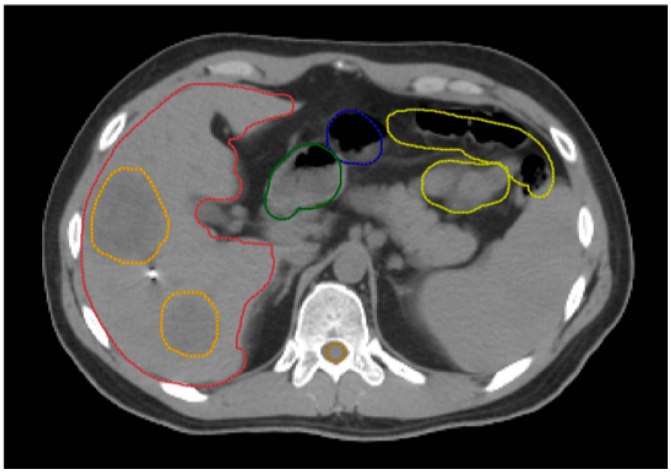

Case 3: Two small lesions

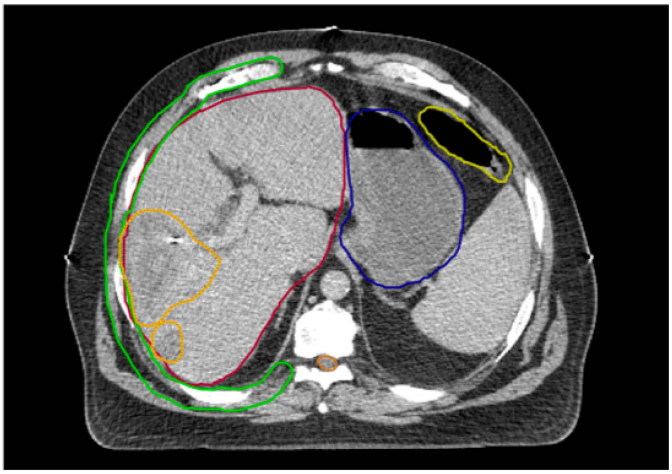

Case 4: Lesion abutting chest wall

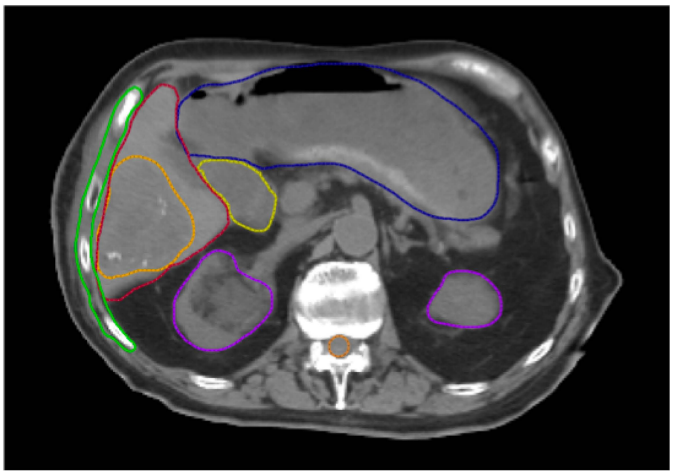

Case 5: Lesion abutting GI tract

FIG. 1. Patient geometries for each of the five cases. The structures contoured on the CT scan are the liver (red), GTV (orange), chest wall (light green), large and small bowel (yellow), stomach (dark blue), esophagus (light blue), kidneys (purple), spinal cord (light brown), heart (magenta), and duodenum (dark green).

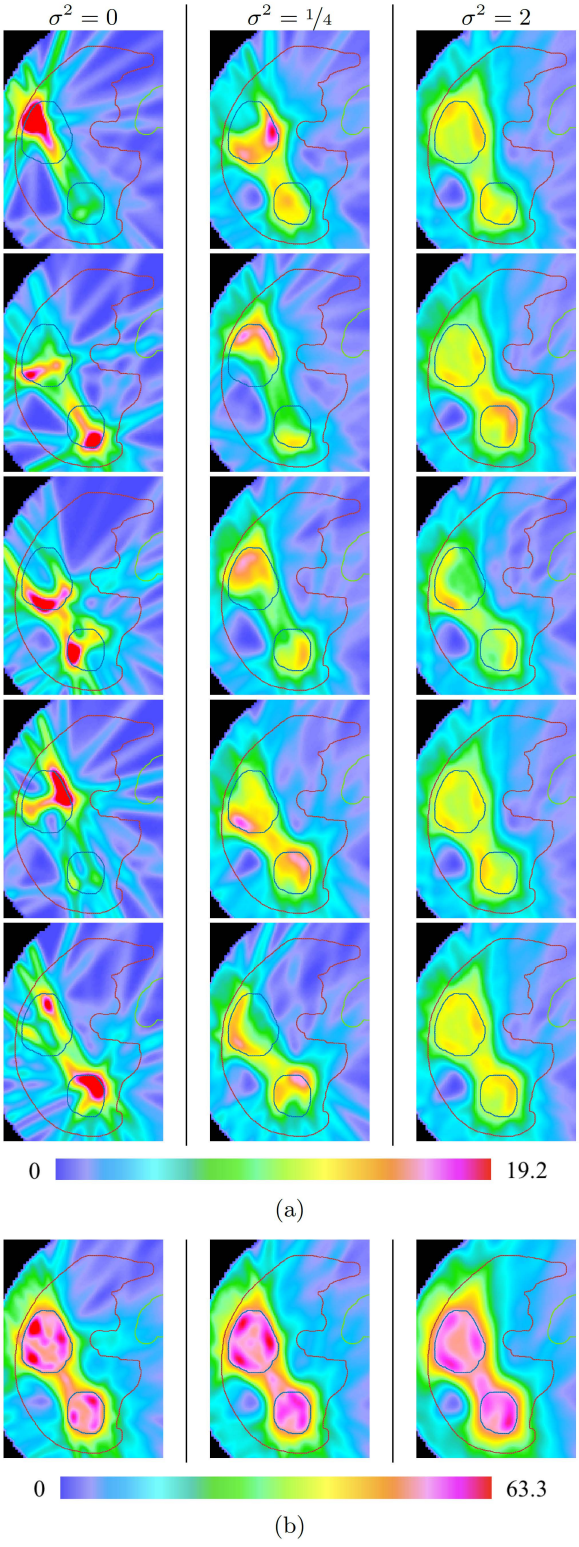

FIG. 2. (a) Physical dose distributions for the five fractions of the optimal spatiotemporal treatment plans for Case 3 with three different values for the variance of patient shifts, plotted in the nominal scenario. Each column contains a spatiotemporal treatment plan that was optimized with a different value of σ^2. As the variance increases, the fractional doses go from highly modulated to more uniform. (b) Distributions of DEQ5 for the expected BED of the plans. All numerical values are shown in Gy.

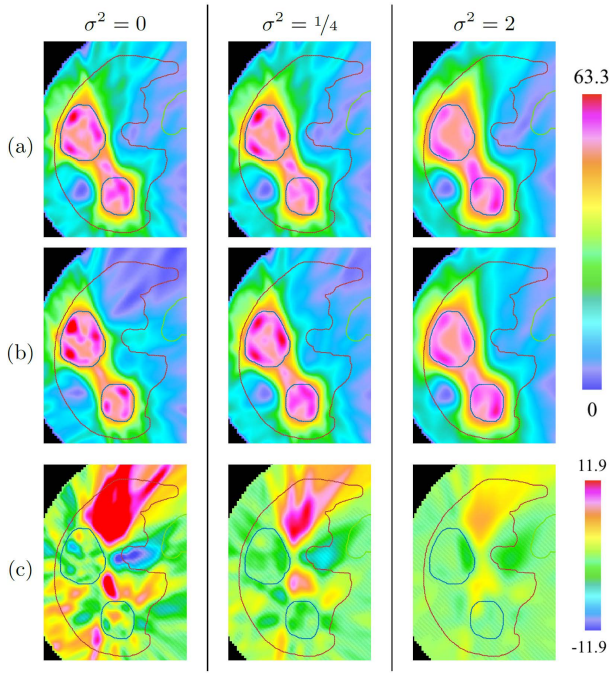

FIG. 3. DEQ5 distributions visualizing the benefit of the robust spatiotemporal plans from Figure 2. (a) Distributions of DEQ5 for the expected BED of the uniformly fractionated plans for Case 3, for three different values of the variance of patient shifts. (b) Distributions of DEQ5 for the expected BED of the spatiotemporal plans. (c) The DEQ5 for the uniform plans minus the DEQ5 of the spatiotemporal plans. All numerical values are shown in Gy.

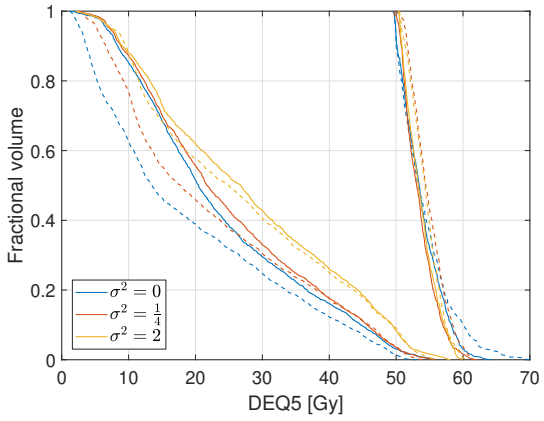

FIG. 4. Dose-volume histogram (DVH) curves for the DEQ5 of expected BED for the GTV and healthy liver in Case 3, with three values of the variance σ^2 of patient shifts. The solid lines are the curves for the uniform plans and the dashed lines are for the spatiotemporal plans. As σ^2 increases, the healthy liver sparing deteriorates and the uniform and spatiotemporal DVH curves become closer together.

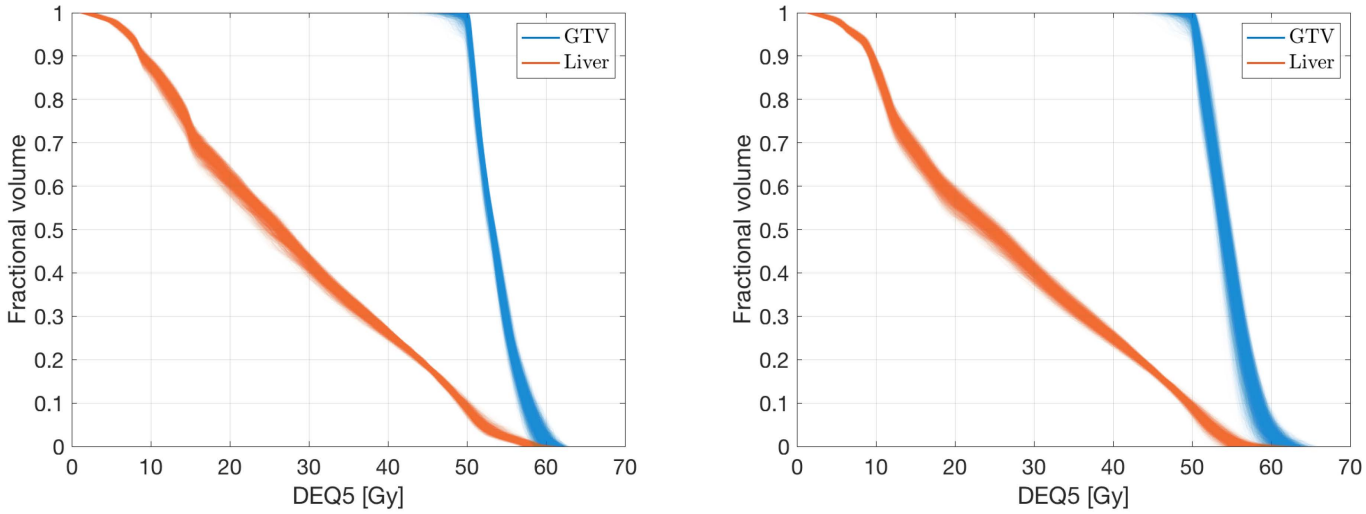

(a) (b)

FIG. 5. A comparison of dose-volume histograms (DVHs) for the DEQ5 of the BED for (a) the uniform reference plan, and (b) the spatiotemporal plan, for Case 3 with variance $\sigma^2 = 2$ of patient shifts up to two voxels from the nominal position. Each curve in the DVH "cloud" is the curve for a single scenario in a sample of 4096 scenarios. The spatiotemporal plan exhibits more variability in the regions of high BED, as evidenced by the greater spread of the curves, but the two treatment plans match in the important shoulder region around a DEQ5 of 50 Gy.