

Optimization of spatiotemporally fractionated radiotherapy treatments with bounds on the achievable benefit

Melissa R. Gaddy

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

Sercan Yıldız

Statistical and Applied Mathematical Sciences Institute, Research Triangle Park, NC
27709-4006, USA

Jan Unkelbach

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

Dávid Papp

E-mail: dpapp@ncsu.edu

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

Abstract. Spatiotemporal fractionation schemes, that is, treatments delivering different dose distributions in different fractions, can potentially lower treatment side effects without compromising tumor control. This can be achieved by hypofractionating parts of the tumor while delivering approximately uniformly fractionated doses to the surrounding tissue. Plan optimization for such treatments is based on biologically effective dose (BED); however, this leads to computationally challenging nonconvex optimization problems. Optimization methods that are in current use yield only locally optimal solutions, and it has hitherto been unclear whether these plans are close to the global optimum. We present an optimization framework to compute rigorous bounds on the maximum achievable normal tissue BED reduction for spatiotemporal plans.

The approach is demonstrated on liver tumors, where the primary goal is to reduce mean liver BED without compromising any other treatment objective. The BED-based treatment plan optimization problems are formulated as quadratically constrained quadratic programming (QCQP) problems. First, a conventional, uniformly fractionated reference plan is computed using convex optimization. Then, a second, nonconvex, QCQP model is solved to local optimality to compute a spatiotemporally fractionated plan that minimizes mean liver BED, subject to the constraints that the plan is no worse than the reference plan with respect to all other planning goals. Finally, we derive a convex relaxation of the second model in the form

of a semidefinite programming (SDP) problem, which provides a rigorous lower bound on the lowest achievable mean liver BED.

The method is presented on 5 cases with distinct geometries. The computed spatiotemporal plans achieve 12-35 percent mean liver BED reduction over the optimal uniformly fractionated plans. This reduction corresponds to 79-97 percent of the gap between the mean liver BED of the uniform reference plans and our lower bounds on the lowest achievable mean liver BED. The results indicate that spatiotemporal treatments can achieve substantial reductions in normal tissue dose and BED, and that local optimization techniques provide high-quality plans that are close to realizing the maximum potential normal tissue dose reduction.

PACS numbers: 87.55.de, 87.55.kd

1. Introduction

Most radiotherapy treatments are fractionated, meaning that delivery of the total dose is split into multiple treatments delivered over several days or weeks. This is motivated by the *fractionation effect*: the clinical observation that radiation-induced damage to cells is lower if the same physical dose is delivered over multiple fractions, owing to cells' ability to recover from sublethal radiation damage. The most widely used mathematical model of the fractionation effect is the biologically effective dose (BED) model (Fowler 2010). According to this model, the biologically effective dose for a treatment with N fractions, each delivering dose d , is given by

$$b = Nd \left(1 + \frac{d}{\alpha/\beta} \right), \quad (1)$$

where α/β is a tissue-specific parameter. For a fixed total physical dose Nd , the BED is minimized if the dose is split evenly into many fractions. This standard fractionation regimen is desired in normal tissues. On the other hand, for a fixed total physical dose, the BED is maximized if all dose is delivered in few fractions (called *hypofractionation*), which is desired in the tumor. Considering this inherent trade-off of fractionation decisions, it would seem ideal to simultaneously achieve hypofractionation in the tumor while splitting the dose to normal tissues evenly into many fractions. While this may appear unattainable at first glance, this goal can be achieved at least approximately using *spatiotemporal fractionation schemes*.

1.1. Spatiotemporal fractionation schemes

Spatiotemporal fractionation schemes deliver different dose distributions in different fractions in an attempt to minimize BED in healthy tissue and maximize BED in the tumor by hypofractionating parts of the tumor while delivering approximately identical doses to the surrounding tissue. The clinical rationale for spatiotemporal fractionation to optimally exploit the fractionation effect was first proposed in the context of proton radiotherapy (Unkelbach, Zeng & Engelsman 2013, Unkelbach & Papp 2015), where the potential benefit comes from the fact that the dose in the entrance region of a proton beam is largely independent of the beam's range, which provides some flexibility to modify the dose in the tumor without equally affecting the dose in the entrance region. More recently it has also been shown that spatiotemporal fractionation may provide a therapeutic advantage in arc therapy delivered with conventional photon beams (Unkelbach 2015). In this case, arc therapy plans can be created in such a way that each fraction delivers high single-fraction doses to complementary parts of the target volume while creating a similar dose bath in the surrounding normal tissue. This was demonstrated for fractionated radiosurgery treatments of large arteriovenous malformations (Unkelbach et al. 2016).

The optimization of spatiotemporally fractionated treatments is far more challenging than conventional IMRT/IMPT optimization. For most commonly used

objective functions, the fluence map optimization problem in conventional IMRT planning is a large-scale convex optimization problem that can be solved to global optimality using well-established gradient-based (e.g., quasi-Newton) optimization methods (Bortfeld 2006). Because spatiotemporal planning is based on BED rather than physical dose and the fluence maps for different fractions are distinct, the optimization models for optimal spatiotemporal fractionation are inherently nonconvex (see Section 2.2). It is common to use heuristics and local optimization methods for nonconvex models in RT planning, but the global optimality of the resulting plans is rarely discussed. In a recent work, Ajdari & Ghatta (2016) studied a treatment planning problem of similar mathematical structure. They proposed a model predictive control approach which computes a treatment plan iteratively by optimizing the dose distributions of remaining fractions after each fraction, with the assumption that the same dose is delivered in every fraction thereafter. This approach requires only the solution of convex optimization problems; however, it does not promote the type of solutions that we seek in spatiotemporal fractionation: plans that deliver high doses to complementary parts of the tumor in distinct fractions. Furthermore, the model predictive control approach comes with no guarantees that the computed plans are anywhere close to globally optimal.

1.2. The contribution of this paper

So far there have been no computational tools to bound the maximum achievable benefit from spatiotemporal fractionation. Prior works (Unkelbach & Papp 2015, Unkelbach et al. 2016) use local optimization, and it has been unclear if better optimization algorithms may yield substantially better solutions. Our approach combines the local optimization of the nonconvex treatment planning model for spatiotemporal fractionation with the solution of a convex optimization problem that provides a rigorous bound on the maximum achievable benefit from spatiotemporal fractionation. We approach this by formulating the spatiotemporal treatment planning problem as a nonconvex quadratically constrained quadratic programming (QCQP) problem. We then derive a convex relaxation of the QCQP problem in the form of a semidefinite programming (SDP) problem. The SDP relaxation provides a nontrivial lower bound; in particular, this bound is tighter than what can be achieved by replacing each nonconvex constraint in the QCQP model with its convex relaxation. We test our method on two-dimensional slices of 5 liver tumors that represent a variety of patient geometries. Comparing the quality of the locally optimal solutions against the SDP relaxation bounds, we find that the local optimal solutions computed for these cases are indeed nearly globally optimal. While there is no conceptual difficulty in applying the same approach to three-dimensional cases, and in particular the local optimal solutions can be computed using the same methods and software that we used in this paper, the SDP relaxations of the three-dimensional cases cannot be solved in a reasonable amount of time with available off-the-shelf software.

1.3. Relation to prior works

Prior research has addressed the problem of optimizing fractionation decisions based on the BED model (Mizuta, Takao, Date, Kishimoto, Sutherland, Onimaru & Shirato 2012, Unkelbach, Craft, Salari, Ramakrishnan & Bortfeld 2013, Keller et al. 2013, Gay et al. 2013, Saberian et al. 2016, Saberian et al. 2015, Mizuta, Date, Takao, Kishimoto, Sutherland, Onimaru & Shirato 2012). These works aim at maximizing the tumor BED subject to BED constraints to the normal tissue. It was shown that the optimal number of fractions depends not only on the α/β ratios of tumor and normal tissues, but also on the dose distribution. These works, however, all assume uniform fractionation, where the same dose distribution is delivered in all fractions, and only the number of fractions needs to be optimized.

The novelty in the idea of spatiotemporal fractionation lies in the fact that there is a potential advantage of delivering distinct dose distributions in different fractions, purely motivated by the basic fractionation effect as described by the standard BED model. There are several extensions of the BED model that describe higher order biological effects such as incomplete repair of radiation damage between fractions, repopulation of tumors over the course of treatment, accelerated repopulation effects, the effect of chemotherapeutic agents, and reoxygenation of hypoxic tumors (Hall & Giaccia 2012). It was found that some of these models give rise to more complex fractionation schemes, i.e. varying doses per fraction (Bertuzzi et al. 2013, Bortfeld et al. 2015, Salari et al. 2015, Wein et al. 2000, Yang & Xing 2005). However, the role of such models to guide fractionation decisions in clinical practice has been limited. Instead, spatiotemporal fractionation as described in this paper is purely based on the basic fractionation effect, whose existence and clinical relevance is undoubted.

Another approach in which different dose distributions may arise in each fraction is *adaptive radiotherapy* (ART) (Yan et al. 1997, Lu et al. 2008, Kim et al. 2012). This technique uses feedback information, such as changes in the patient anatomy, to modify the treatment plan during the course of a fractionated treatment, e.g. replanning to compensate for tumor shrinkage in lung or head and neck cancer (Sonke & Belderbos 2010, Wu et al. 2009). Instead, spatiotemporal fractionation shows that there is a benefit of delivering distinct dose distributions in different fractions, even in the absence of any changes of the patient over time.

2. Mathematical model for optimal spatiotemporal fractionation

2.1. Uniform and nonuniform fractionation using the BED model

The natural generalization of the BED model in (1) to fractionated treatments delivering different doses in different fractions is to define the (cumulative) BED as

$$b = \sum_{t=1}^N \left(d_t + \frac{d_t^2}{(\alpha/\beta)} \right), \quad (2)$$

where d_1, \dots, d_N are the doses delivered (in any order) in fractions 1 through N . IMRT planning using BED can be performed analogously to conventional IMRT optimization, using similar objective functions and constraints in the treatment planning optimization model, but substituting BED in place of physical dose (Unkelbach & Papp 2015). We define x_1, \dots, x_N to be the vectors of beamlet weights delivered in fractions 1 through N , and V to be the set of volume elements (voxels) used for dose calculation during the optimization. Using the cumulative BED from (2), we obtain the nonuniform fractionation problem:

$$\begin{aligned}
& \min_{x,d,b} F(b) \\
& \text{s.t.} \quad b_v = \sum_{t=1}^N (d_{vt} + \frac{d_{vt}^2}{(\alpha/\beta)_v}) \quad \forall v \in V \\
& \quad Dx_t = d_t \quad t = 1, \dots, N \\
& \quad x_t \geq 0 \quad t = 1, \dots, N,
\end{aligned} \tag{3}$$

where D is the usual dose-influence matrix, and the objective function F represents the desired clinical goals of target coverage, conformity, and organ sparing. Defining I as the index set of the objectives, we write the objective function as the sum

$$F(b) = \sum_{i \in I} w_i F_i(b), \tag{4}$$

where each term F_i represents a single clinical objective (such as deviation from a prescribed lower or upper bound on the minimum, maximum, or mean BED of a structure), and each positive weight w_i represents the relative importance of a clinical goal.

BED-based optimization can also be performed analogously for conventional, uniformly fractionated treatments. Throughout this work, we use a *uniform reference plan* as a benchmark to evaluate the benefit of nonuniform fractionation. We obtain the uniform reference plan by solving (3) with the additional constraint that $x_1 = \dots = x_N$. Eliminating the redundant variables, the uniform reference plan is the optimal solution of the following problem:

$$\begin{aligned}
& \min_{x,d,b} F(b) \\
& \text{s.t.} \quad b_v = Nd_v \left(1 + \frac{d_v}{(\alpha/\beta)_v} \right) \quad \forall v \in V \\
& \quad Dx = d \\
& \quad x \geq 0.
\end{aligned} \tag{5}$$

Figure 4 shows an example of a reference plan and a nonuniform 5-fraction treatment of a patient harboring a large liver tumor. The reference plan (Figure 1c) was computed by solving the BED-based treatment plan optimization model (5). Figure 1a shows the 5 dose distributions of a nonuniformly fractionated plan. Both plans were

computed by optimizing the same objectives described below in Section 3 to represent identical clinical goals. The reference plan yields a conventional treatment that irradiates the tumor to the prescribed dose in each fraction. The nonuniformly fractionated plan delivers high single-fraction doses to parts of the tumor while delivering a similar low-dose bath to the surrounding tissue in each fraction. The high single-fraction doses in the tumor allow for a reduction in the total physical dose delivered (Figure 4(b)). Since the dose to the surrounding normal tissue is approximately uniformly fractionated, this also yields a reduction of BED in the normal tissue.




Figure 1. Dose distributions for Case 1, a large central lesion within the liver. Also shown are the contours of the heart, the esophagus, and the spinal cord; these organs are not dose-limiting. (a) Physical dose distributions in each of the five fractions show that the nonuniformly fractionated treatment hypofractionates different parts of the tumor. (b) Total physical dose delivered throughout the nonuniformly fractionated treatment. (c) Physical dose distribution of the uniformly fractionated reference plan. (d) DEQ5 of the nonuniformly fractionated plan, which is the uniform plan that is isoeffective in delivering the same BED as the nonuniformly fractionated plan. (e) The difference between the physical dose in the uniform plan and the DEQ5 for the nonuniform plan, or (c) minus (d). This shows that the spatiotemporal plans reduce dose in the healthy liver and in the entrance region of the beams that expose the liver the most. All numerical quantities shown are in [Gy].

2.2. Nonconvexity of nonuniform fractionation

It is well-known that conventional fluence map optimization is a convex optimization problem when the clinical goals are modeled using a convex objective function F of the

physical dose. Despite the apparent nonconvexity introduced by the quadratic equality constraints in the uniform fractionation problem (5), this model is convex when the physical doses d are restricted to clinically relevant values and F is a piecewise quadratic penalty function similar to the dose-based objective functions (Unkelbach & Papp 2015). This is the case for the objectives used in this paper. Hence, the uniform reference plan can be computed using the gradient-based local optimization methods commonly used in IMRT plan optimization.

The nonuniform fractionation problem (3), however, is nonconvex.[‡] This nonconvexity is an inherent characteristic of the problem that cannot be eliminated by reformulating the model, as the following argument shows: by definition, permuting the fractions in an optimal treatment plan leads to another optimal plan, yet the average of these $N!$ treatment plans will be a plan with identical fractions, which in general will not be optimal. Thus, gradient-based optimization methods used to solve the uniform model can only yield locally optimal solutions for the nonuniform model. Figure 2 illustrates this problem, showing that substantially different locally optimal spatiotemporal plans may exist for the same case.

2.3. The constrained nonuniform model

A drawback of formulation (3)–(4) is that the optimal solution does not realize the maximum benefit of spatiotemporal fractionation over conventional fractionation in any given objective. Instead, the benefit is distributed among the terms of the objective function; in other words, the benefit is realized as a combination of smaller improvements with respect to the different clinical goals. In the liver cases we study in Section 4, the primary gain is expected to be in lowering BED to the liver without compromising target coverage and conformity. This suggests an alternative formulation in which we minimize only one of the clinical objectives of (4), say F_1 , while constraining the solution to be at least as good as the uniform reference plan with respect to each other objective F_i ($i \in I, i \neq 1$).

Let b^* be the BED distribution delivered by the uniform reference plan, which is obtained by solving (5), and let F_1 be the penalty for the mean BED in the liver exceeding zero. We modify (3) to include the constraints that the solution must be as good as the uniform reference plan with respect to all objectives besides F_1 , and we

[‡] The nonconvexity of the formulation comes from the composition of the piecewise quadratic penalty function penalizing the underdose of the target and the function defining the BED. As a general rule, the composition of a convex nonincreasing function (such as our penalty function) with a convex quadratic function (such as the BED) can be nonconvex, unlike the composition of a convex nondecreasing function (such as the functions used to penalize overdose) and the BED, which is always convex.

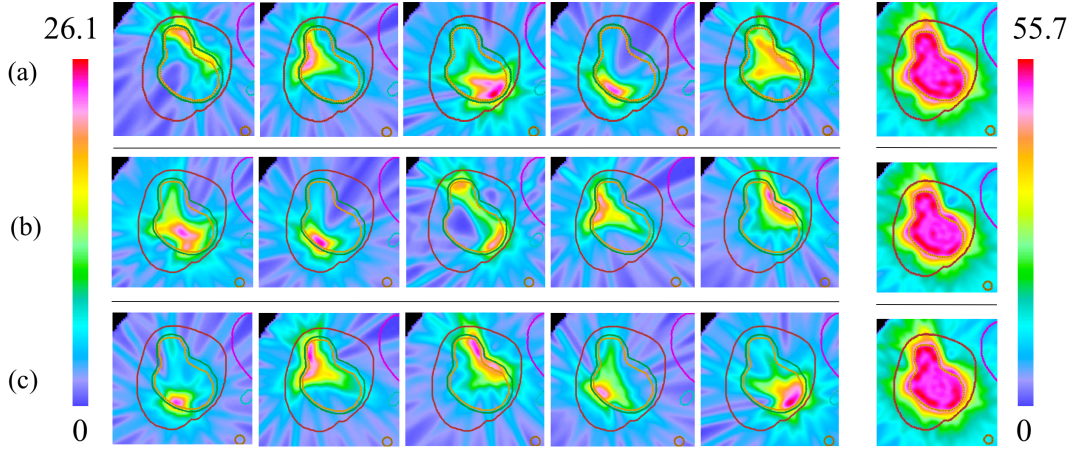


Figure 2. Rows (a), (b), and (c) are physical dose distributions from three additional locally optimal 5-fraction nonuniform treatments for Case 1, which is the same clinical liver case displayed in Figure 4. They are all locally optimal solutions of the model (3). The first five panels of each row are the dose distributions in the five nonuniform fractions, and the last panel on each row is the equivalent dose DEQ5 (see Section 4.2). The solutions exhibit the same pattern: different subregions of the tumor receive a high single-fraction dose in different fractions. Note that the emergent “partitions” in the optimized plans are a result of the optimization. This pattern supports the rationale that the benefit of spatiotemporal fractionation is a result of hypofractionating parts of the tumor while maintaining a consistent low dose in the surrounding tissue. The difference in the hypofractionated regions in each solution also demonstrates that several qualitatively different locally optimal spatiotemporal treatments may exist for the same case. All numerical quantities shown are in [Gy].

obtain the *constrained* nonuniform spatiotemporal fractionation problem:

$$\begin{aligned}
 & \min_{x,d,b} F_1(b) \\
 & \text{s.t.} \quad F_i(b) \leq F_i(b^*) \quad i \in I, i \neq 1 \\
 & \quad b_v = \sum_{t=1}^N (d_{vt} + \frac{d_{vt}^2}{(\alpha/\beta)_v}) \quad \forall v \in V \\
 & \quad Dx_t = d_t \quad t = 1, \dots, N \\
 & \quad x_t \geq 0 \quad t = 1, \dots, N.
 \end{aligned} \tag{6}$$

The first set of constraints ensures that the improvement in the objective F_1 is not at the cost of sacrificing the other clinical objectives; the computed spatiotemporal plan is either preferable or identical to the uniform reference plan with respect to every objective.

3. Mathematical model for bounding the maximum achievable benefit

In this section, we formulate the convex optimization model for bounding the maximum achievable benefit from spatiotemporal fractionation. To that end, we write problem

(3)–(4) as a quadratic optimization problem with quadratic constraints. We assume that F is a weighted sum (4) of penalty functions that penalize deviations from prescribed BED values. Let V_i be the set of voxels involved in clinical objective $i \in I$. Each F_i is a piecewise quadratic penalty function that penalizes either BED above a prescribed threshold b_v^{hi} in voxels $v \in V_i$, or BED below a prescribed threshold b_v^{lo} in voxels $v \in V_i$, or mean BED above a prescribed mean BED m_i^{hi} in a structure i . Let I^+ , I^- , and I^m be the index sets of objectives of the first, second, and third type respectively. (Thus, the index set I in (4) is the union of the three sets I^+ , I^- , and I^m .) Letting $(y)_+$ denote the positive-part function $\max(y, 0)$, we express the clinical goal $i \in I$ with the penalty function

$$F_i(b) = \begin{cases} \sum_{v \in V_i} (b_v - b_{iv}^{hi})_+^2 & \text{if } i \in I^+, \\ \sum_{v \in V_i} (b_{iv}^{lo} - b_v)_+^2 & \text{if } i \in I^-, \\ \left(\frac{1}{|V_i|} \sum_{v \in V_i} b_v - m_i^{hi} \right)_+^2 & \text{if } i \in I^m. \end{cases} \quad (7)$$

The overall objective function F is the weighted sum (4).

The prescription BED values b_{iv}^{hi} and b_{iv}^{lo} can be derived from clinical dose prescriptions of fractionated treatments that employ similar physical doses per fraction as expected to be used in the spatiotemporally fractionated treatments. Typically one would use the same threshold b^{hi} or b^{lo} in every voxel v of the same structure V_i , but we opted for the above, more general, formulation to allow for distance-dependent thresholds often used to improve the conformity of the dose distributions.

Noting that $(y)_+$ is the smallest number z satisfying the inequalities $z \geq y$ and $z \geq 0$, we introduce the auxiliary optimization variables p_{iv} for $i \in I^+$ and q_{iv} for $i \in I^-$, whose values (when minimized) are equal to the quantities $(b_v - b_{iv}^{hi})_+$ and $(b_{iv}^{lo} - b_v)_+$ for voxel v . Similarly, we introduce the variables r_i for $i \in I^m$ for the mean-BED penalties $(\frac{1}{|V_i|} \sum_{v \in V_i} b_v - m_i^{hi})_+$.

Finally, we eliminate the variables b_v from the problem by replacing them with $\sum_{t=1}^N (d_{vt} + \frac{d_{vt}^2}{(\alpha/\beta)_v})$. We arrive at the following quadratically constrained quadratic

programming (QCQP) formulation of the nonuniform fractionation problem:

$$\begin{aligned}
\min_{x,d,p,q,r} \quad & \sum_{i \in I^+} \sum_{v \in V_i} w_i p_{iv}^2 + \sum_{i \in I^-} \sum_{v \in V_i} w_i q_{iv}^2 + \sum_{i \in I^m} w_i r_i^2 \\
\text{s.t.} \quad & p_{iv} \geq -b_{iv}^{hi} + \sum_{t=1}^N (d_{vt} + \frac{d_{vt}^2}{(\alpha/\beta)_v}) \quad \forall i \in I^+, \forall v \in V_i \\
& p_{iv} \geq 0 \quad \forall i \in I^+, \forall v \in V_i \\
& q_{iv} \geq b_{iv}^{lo} - \sum_{t=1}^N (d_{vt} + \frac{d_{vt}^2}{(\alpha/\beta)_v}) \quad \forall i \in I^-, \forall v \in V_i \\
& q_{iv} \geq 0 \quad \forall i \in I^-, \forall v \in V_i \\
& r_i \geq \frac{1}{|V_i|} \sum_{v \in V_i} \left(\sum_{t=1}^N d_{vt} + \frac{d_{vt}^2}{(\alpha/\beta)_v} \right) - m_i^{hi} \quad \forall i \in I^m \\
& r_i \geq 0 \quad \forall i \in I^m \\
& Dx_t = d_t \quad \forall t = 1, \dots, N \\
& x_t \geq 0 \quad \forall t = 1, \dots, N.
\end{aligned} \tag{8}$$

As the physical dose d_{vt} is a linear function of the beamlet weights x_t , the d_{vt} variables can also be eliminated from the formulation, and all the inequality constraints can be seen as quadratic inequalities in the primary decision variables x_1, \dots, x_N . We introduce the matrix $X_t = x_t x_t^T$ and consider each matrix element as an auxiliary decision variable. The quadratic inequalities in (8) can then be written as linear inequalities in x_t and X_t . For instance, on the right-hand side of the first set of inequalities we substitute

$$-b_{iv}^{hi} + \sum_{t=1}^N (d_{vt} + \frac{d_{vt}^2}{(\alpha/\beta)_v}) = \sum_{t=1}^N \left\langle \begin{pmatrix} 1 & x_t^T \\ x_t & X_t \end{pmatrix}, \begin{pmatrix} -\frac{b_{iv}^{hi}}{D^T e_v} & \frac{e_v^T D}{2} \\ \frac{D^T e_v}{2} & \frac{1}{(\alpha/\beta)_v} D^T e_v e_v^T D \end{pmatrix} \right\rangle,$$

where $\langle A, B \rangle = \sum_{i,j} A_{ij} B_{ij}$ is the component-wise (or Frobenius) inner product of matrices, and e_v is the characteristic vector with a 1 in the v -th position and zeros elsewhere. (That is, $e_v^T D$ is simply the v -th row of the D matrix.) The resulting optimization model has a convex quadratic objective function and only linear constraints, aside from the nonconvex quadratic equations $X_t = x_t x_t^T$ for $t = 1, \dots, N$. We obtain a convex relaxation of this problem by replacing each of these equations with the weaker convex constraint $X_t - x_t x_t^T \succcurlyeq 0$ (meaning that the difference $X_t - x_t x_t^T$ is positive semidefinite). The latter constraint is indeed convex, as it is equivalent to the linear matrix inequality $\begin{pmatrix} 1 & x_t^T \\ x_t & X_t \end{pmatrix} \succcurlyeq 0$.

Since in the original nonconvex model we have $X_t = x_t x_t^T$ and $x_t \geq 0$, it is clear that each X_t is also component-wise nonnegative. Hence, we can add the inequalities $X_t \geq 0$ for $t = 1, \dots, N$ to the convex relaxation. (This component-wise inequality should not be confused with the linear matrix inequality $X_t \succcurlyeq 0$.) Note that even though the inequality $X_t \geq 0$ is redundant in the original model, $X_t \geq 0$ is not a

redundant constraint in the convex relaxation; adding it to the optimization problem tightens the bound.

A further simplification is possible. Since the relaxation is convex and symmetric in the fractions, we can assume without loss of generality that at the optimum we have $x_1 = \dots = x_N$ and $X_1 = \dots = X_N$. Thus, we can eliminate the variables corresponding to the different fractions and use only a single variable x and X in place of each x_t and X_t . (This shows that our convex relaxation does not distinguish between the uniformly and nonuniformly fractionated models, although the bound does depend on the number of fractions.)

Finally, using the shorthand $C_v = \frac{1}{(\alpha/\beta)_v} D^T e_v e_v^T D$, we arrive at the following convex relaxation of (8):

$$\begin{aligned}
\min_{x, X, p, q, r} \quad & \sum_{i \in I^+} \sum_{v \in V_i} w_i p_{iv}^2 + \sum_{i \in I^-} \sum_{v \in V_i} w_i q_{iv}^2 + \sum_{i \in I^m} w_i r_i^2 \\
\text{s.t.} \quad & p_{iv} \geq N \left\langle \begin{pmatrix} 1 & x^T \\ x & X \end{pmatrix}, \begin{pmatrix} -\frac{b_{iv}^{hi}}{N} & \frac{e_v^T D}{2} \\ \frac{D^T e_v}{2} & C_v \end{pmatrix} \right\rangle & \forall i \in I^+, \forall v \in V_i \\
& p_{iv} \geq 0 & \forall i \in I^+, \forall v \in V_i \\
& q_{iv} \geq -N \left\langle \begin{pmatrix} 1 & x^T \\ x & X \end{pmatrix}, \begin{pmatrix} -\frac{b_{iv}^{lo}}{N} & \frac{e_v^T D}{2} \\ \frac{D^T e_v}{2} & C_v \end{pmatrix} \right\rangle & \forall i \in I^-, \forall v \in V_i \\
& q_{iv} \geq 0 & \forall i \in I^-, \forall v \in V_i \\
& r_i \geq \frac{N}{|V_i|} \sum_{v \in V_i} \left\langle \begin{pmatrix} 1 & x^T \\ x & X \end{pmatrix}, \begin{pmatrix} -\frac{m_i^{hi}}{N} & \frac{e_v^T D}{2} \\ \frac{D^T e_v}{2} & C_v \end{pmatrix} \right\rangle & \forall i \in I^m \\
& r_i \geq 0 & \forall i \in I^m \\
& \begin{pmatrix} 1 & x^T \\ x & X \end{pmatrix} \succcurlyeq 0, \quad x \geq 0, \quad X \geq 0.
\end{aligned} \tag{9}$$

The optimal objective function value of this problem is a lower bound for the global minimum of the spatiotemporal fractionation problem (8).

Note that to obtain (9), we linearized every quadratic constraint in (8), even though all of them were convex except for those involving q_{iv} . It can be seen that this way (9) yields a tighter bound than what could be obtained by simply replacing the concave quadratic constraints with a convex relaxation and keeping the convex quadratic constraints intact. This is because unlike the linearizations of the concave quadratics (which are relaxations), the linearizations of the convex quadratics are tighter than the original constraints. Our derivation of (9) shows that the convex model as a whole is indeed a relaxation of (8) despite the fact that this cannot be seen when we compare the models constraint by constraint.

3.1. Bounding the maximum benefit in a given objective

The convex relaxation of the constrained nonuniform model (6) can be derived analogously to how we obtained the relaxation (9) of the model (3). For simplicity of notation, we assume (consistently with the application to liver tumors) that the mean dose objective $F_1(b) = (\frac{1}{|V_1|} \sum_{v \in V_1} b_v - m_1^{hi})_+^2$ is the primary objective. Then the convex optimization model bounding the minimum value of $F_1(b)$ from below is the following:

$$\begin{aligned}
& \min_{x, X, p, q, r} && r_1 \\
& \text{s.t.} && \sum_{v \in V_i} p_{iv}^2 \leq F_i(b^*) && \forall i \in I^+ \\
& && \sum_{v \in V_i} q_{iv}^2 \leq F_i(b^*) && \forall i \in I^- \\
& && r_i \leq F_i(b^*) && \forall i \in I^m, i \neq 1 \\
& && p_{iv} \geq N \left\langle \begin{pmatrix} 1 & x^T \\ x & X \end{pmatrix}, \begin{pmatrix} -\frac{b_{iv}^{hi}}{N} & \frac{e_v^T D}{2} \\ \frac{D^T e_v}{2} & C_v \end{pmatrix} \right\rangle && \forall i \in I^+, \forall v \in V_i \\
& && p_{iv} \geq 0 && \forall i \in I^+, \forall v \in V_i \\
& && q_{iv} \geq -N \left\langle \begin{pmatrix} 1 & x^T \\ x & X \end{pmatrix}, \begin{pmatrix} -\frac{b_{iv}^{lo}}{N} & \frac{e_v^T D}{2} \\ \frac{D^T e_v}{2} & C_v \end{pmatrix} \right\rangle && \forall i \in I^-, \forall v \in V_i \\
& && q_{iv} \geq 0 && \forall i \in I^-, \forall v \in V_i \\
& && r_i \geq \frac{N}{|V_i|} \sum_{v \in V_i} \left\langle \begin{pmatrix} 1 & x^T \\ x & X \end{pmatrix}, \begin{pmatrix} -\frac{m_i^{hi}}{N} & \frac{e_v^T D}{2} \\ \frac{D^T e_v}{2} & C_v \end{pmatrix} \right\rangle && \forall i \in I^m \\
& && r_i \geq 0 && \forall i \in I^m \\
& && \begin{pmatrix} 1 & x^T \\ x & X \end{pmatrix} \succcurlyeq 0, \quad x \geq 0, \quad X \geq 0.
\end{aligned} \tag{10}$$

3.2. Solution methods

Although the optimization problems (9) and (10) are convex, their solution is not straightforward using the gradient-based algorithms commonly used in treatment planning optimization, such as the projected gradient descent method and its variants. This is because computing the violation of the linear matrix inequality constraints $\begin{pmatrix} 1 & x^T \\ x & X \end{pmatrix} \succcurlyeq 0$ and projecting on the set of points satisfying this inequality require an expensive matrix factorization in each step of the method. Optimization problems involving this type of constraints are called *semidefinite programs* (Vandenberghe & Boyd 1996, Boyd & Vandenberghe 2004). They have been extensively studied in the numerical optimization literature, and several reliable and efficient algorithms are available for their solution.

Formally, a semidefinite program is a convex optimization problem with a linear objective function and with constraints that are all either linear inequalities or *linear*

matrix inequalities of the form $A(x) \succcurlyeq 0$, where $A(\cdot)$ is an affine function that maps the optimization variables x to the space of real symmetric matrices. The problems (9) and (10) have convex quadratic objectives and constraints, rather than only linear ones, but it can be shown that convex quadratic inequalities can be equivalently written as linear matrix inequalities (Vandenberghe & Boyd 1996); hence their inclusion is without loss of generality.

Large-scale semidefinite programs are routinely solved using Newton-type methods implemented in widely available optimization software such as SeDuMi (Sturm 1999) and MOSEK (MOSEK ApS 2017). In our experiments we used MOSEK to solve the semidefinite programs (9) and (10).

4. Application to liver tumors

We examined the benefit of spatiotemporal fractionation on two-dimensional slices of five clinical liver cases with distinct geometries. The first three cases feature centrally-located lesions within the liver; Case 1 has a large lesion, Case 2 has a small lesion, and Case 3 has two separate lesions within the liver. In each of these cases the liver is the main dose-limiting organ. In Case 4, the tumor abuts the chest wall, and Case 5 is a challenging geometry where both the chest wall and the bowel are dose-limiting and need to be included in the treatment plan optimization model.

4.1. Experimental setup

4.1.1. The BED prescription and α/β ratios. Five-fraction treatments were planned for all cases. To derive the upper and lower BED thresholds b^{hi} and b^{lo} , typical prescription doses and normal tissue constraints for 5-fraction liver SBRT were converted into BED values using an α/β ratio of 10 in the tumor and 4 in all normal tissues. For example, the prescription lower bounds b^{lo} for the GTV and PTV were chosen to be 100 Gy and 72 Gy, respectively, which correspond to 50 Gy and 40 Gy of physical dose delivered in 5 fractions assuming uniform fractionation. The complete list of objectives and constraints is as follows:

- A BED of 100 Gy is prescribed to the GTV.
- A BED of 72 Gy is prescribed to the PTV.
- BED exceeding 115.5 Gy BED in the GTV is penalized.
- BED exceeding 100 Gy BED in the PTV is penalized.
- The plan is to be conformal: a linear BED falloff is aimed for in a 3cm margin around the PTV.
- The mean BED in the liver excluding the GTV is minimized.
- In Cases 4 and 5, BED exceeding 96.2 Gy in the chest wall is penalized. (This corresponds to 35 Gy physical dose in 5 fractions.)
- In Case 5, BED exceeding 75 Gy in the affected sections of the GI tract is penalized.

Each of these objectives is implemented via a penalty function shown in Equation (7).

4.1.2. Computing the uniformly fractionated reference plan. The BED-based fluence map optimization problem (5) was solved to obtain optimal beamlet weights for the 5-fraction uniform reference plan, which is a high-quality plan that delivers the same dose distribution in each of its fractions. Dose-influence matrices were calculated for 21 equispaced coplanar beams using the Quadrant Infinite Beam (QIB) dose calculation algorithm implemented in CERR version 5.2 (Deasy et al. 2003). These 21-beam IMRT plans represent the quality achievable by high-quality arc therapy plans (Papp & Unkelbach 2014). In this work, we limited the computations to two-dimensional slices of the patient voxels and optimized beamlet weights for a single row of the 1x1 cm beamlet grid. As discussed in Section 2.2, although the uniform plan optimization problem is defined using a combination of convex and nonconvex constraints, when the beamlet weights x and physical doses d are restricted to clinically relevant values, the problem is convex, and the global optimal solution can be computed using local optimization algorithms (Unkelbach & Papp 2015). In our experiments, the optimization of the reference plans was performed with Matlab using the L-BFGS-B solver (Zhu et al. 2011). The runtimes to find the (globally) optimal solutions were less than 10 seconds on a standard desktop computer for all of the cases.

4.1.3. Computing the spatiotemporally fractionated plans. After computing the uniform reference plan, we computed nonuniformly fractionated treatment plans in which the beamlet weights and corresponding dose distributions are not the same in each of the five fractions. As described in Section 2.3, we solve (6) to minimize the mean BED in the liver, subject to the constraints that the solution must be at least as good as the uniform reference plan computed in Section 4.1.2 with respect to all other objectives. A perturbation of the fluence maps of the uniform reference plan was used as an initial solution for the optimization. § These nonconvex problems were solved to local optimality using Matlab’s built-in optimizer `fmincon` (MathWorks 2016), which utilizes an interior-point algorithm to find local solutions to a constrained nonlinear program. Runtimes for computing a locally optimal solution of (6) ranged from 15 minutes to 4 hours.

4.2. Results

Figures 4 and 3 show a graphical comparison of the uniform reference plan with locally optimal nonuniformly fractionated solutions for Cases 1 and 2; the graphical results for the remaining three cases can be found in the supplement. Table 1 summarizes the computed bounds and benefits observed in all five cases. In each case, the nonuniform

§ Beamlet weights were multiplied by independent random factors drawn uniformly from the interval $[0, 2]$. The quality of the computed local optimal solutions did not appear to be sensitive to the specific way the plans were perturbed.

	Case description	Mean liver BED [Gy]				Gap closed	Sparing factor
		conventional	spatiotemporal	reduction	lower bnd.		
1	Large central lesion	84.54	75.87	12.75%	73.38	77.69%	0.6663
2	Small lesion	26.14	19.47	34.26%	18.58	88.23%	0.3830
3	Two small lesions	59.54	50.24	18.51%	48.03	80.80%	0.5879
4	Lesion abutting ch. wall	47.51	38.65	22.92%	37.65	89.86%	0.5289
5	Lesion abutting GI tract	88.67	77.38	14.59%	77.02	96.91%	0.7028

Table 1. Summary of mean liver BED reductions from spatiotemporal fractionation, lower bounds for mean liver BED, and sparing factors for each of the five cases. The “gap closed” values provide a measure of how close the local optimal solutions are to achieving the lower bound on the mean liver BED; see Eq. (11) for the definition. The sparing factor is a value that determines the dependence of the optimal fractionation schedule of a fixed dose distribution on the patient geometry (see Section 4.2.2). The remarkable benefit seen in Case 2 agrees with the fact that the sparing factor is substantially lower in this case than in the other cases.

plan is constrained to maintain the same target coverage as the uniform reference plan, yet the mean liver BED in the nonuniform plans is substantially lower than in the uniform plans (see Table 1). Case 1 displays the smallest reduction in the mean liver BED with an approximately 13% improvement, and Case 2 has the largest reduction of approximately 34% in the mean liver BED. Note that these improvements are achieved without sacrificing any other clinical objective, as by definition, the spatiotemporal plans computed by solving (6) are at least as good as the uniform reference plan with respect to every objective. In particular, the spatiotemporal plans are as conformal as the corresponding uniform reference plans.

The effectiveness of a spatiotemporal treatment can be quantified by its equivalent dose DEQ5, defined via

$$\text{BED} = \sum_{t=1}^5 \left(d_t + \frac{d_t^2}{\alpha/\beta} \right) = \text{DEQ5} + \frac{(\text{DEQ5})^2/5}{\alpha/\beta},$$

which yields

$$\text{DEQ5} = 5 \left(-\frac{\alpha/\beta}{2} + \sqrt{\left(\frac{\alpha/\beta}{2} \right)^2 + \frac{\alpha/\beta}{5} \sum_{t=1}^5 \left(d_t + \frac{d_t^2}{\alpha/\beta} \right)} \right).$$

The DEQ5 is the cumulative dose distribution of a uniformly fractionated 5-fraction treatment that achieves the same BED as the given spatiotemporal treatment. (Note the difference between DEQ5 and the also commonly used *equieffective dose*, EQX, which is the total physical dose that needs to be delivered in a uniform treatment with a dose per fraction of X Gy to achieve the same BED.) Using DEQ5 has the advantage that the spatiotemporal plan can be directly compared to the physical dose distribution of the uniform reference plan. The comparison is shown in panels (c)–(e) of Figures 4

and 3 for Cases 1 and 2, and in the Supplement for the remaining cases. By computing the difference between the two plans, we observe that the nonuniformly fractionated plans maintain the same tumor BED as the uniformly fractionated plans while reducing the mean dose and BED in healthy liver tissue. The same can also be demonstrated by comparing the DVH curves of the physical dose and DEQ5 of the two plans. For Case 1, these curves are included in the Supplement (Figure 7). The DVH curves show that there is substantial dose reduction in the liver (both the DEQ5 and physical dose curves shift to the left), and that this is a consequence of lower physical dose (but nearly identical DEQ5) delivered to the tumor.

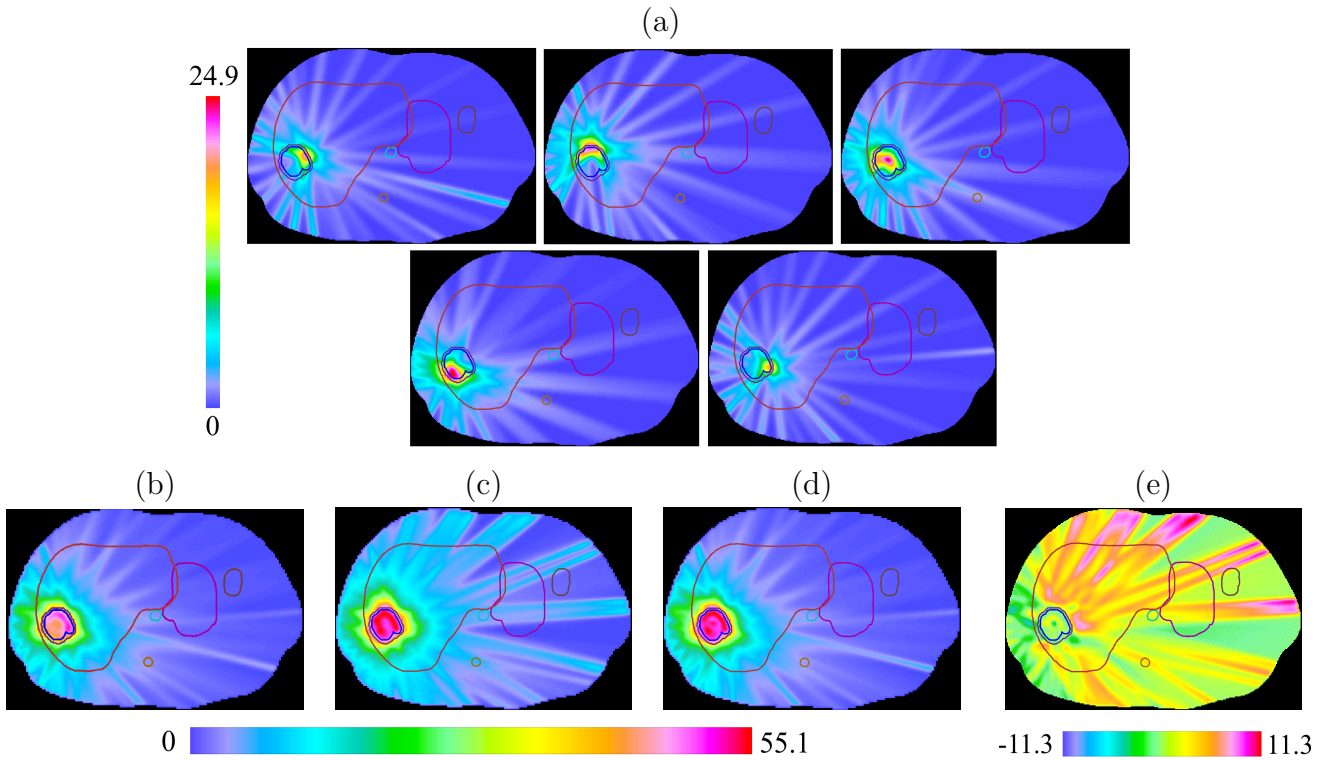


Figure 3. Dose distributions for Case 2, a small lesion within the liver. (a) Physical dose distributions in a 5-fraction nonuniformly fractionated treatment. Although the tumor is small, the spatiotemporal plan still hypofractionates different subregions of the target. (b) Total physical dose delivered by the five nonuniform fractions. (c) Physical dose distribution of the uniform reference plan. (d) DEQ5 of the nonuniformly fractionated plan. As in Case 1, the nonuniformly fractionated plan provides similar target coverage as the uniform reference plan, with a more conformal dose distribution. (e) The plot of (c) minus (d) shows that the majority of the BED reduction is in the healthy liver and in the entrance regions of the beams exposing the liver the most. All numerical quantities shown are in [Gy].

4.2.1. Bounds on the maximum achievable liver BED reduction. Because the nonuniform fractionation problem is nonconvex, the nonuniformly fractionated solutions presented in this paper are only certified to be locally optimal. To gauge the quality of

these solutions, for each case we computed a lower bound on the minimum liver BED, as described in Section 3.1. These lower bounds, also shown in Table 1, are close to the mean liver BED values achieved by the locally optimal nonuniformly fractionated solutions. For example, in Case 1, spatiotemporal fractionation reduced the mean liver BED to 75.9 Gy, which is a reduction of 8.7 Gy from the uniform reference plan. The lower bound from the SDP relaxation is 73.4 Gy, which means that no treatment plan can achieve a reduction of more than 11.2 Gy from the uniform reference plan. It is important to note that the value of 73.4 Gy is only a lower bound, and we are not able to guarantee the existence of a nonuniformly fractionated solution that achieves this value. The mean liver BED value for the true globally optimal plan may be anywhere between 75.9 and 73.4 Gy.

We compared the mean liver BED reduction in the locally optimal nonuniformly fractionated plan with our bound on the maximum possible reduction by computing the fraction of the gap between the mean liver BED in the uniform reference plan and the lower bound from the SDP relaxation that the nonuniformly fractionated plan closes. Formally, if m_{ref} and m_{sp} are the mean liver BEDs achieved by the uniform reference plan and the locally optimal spatiotemporal plan, respectively, and m_{SDP} is the lower bound obtained from the SDP relaxation, then the fraction of the *gap closed* by the spatiotemporal plan is the ratio

$$\frac{m_{\text{ref}} - m_{\text{sp}}}{m_{\text{ref}} - m_{\text{SDP}}}. \quad (11)$$

In particular, a ratio of 100% is the best possible value, which is attained if the spatiotemporal plan matches the SDP lower bound, and 0% is the worst possible value, attained when the uniform plan cannot be improved upon using spatiotemporal fractionation. Note that this ratio being close to 100% does not only speak directly for the quality of the spatiotemporal plan; it also demonstrates the sharpness of the lower bound. For example, in Case 1, the nonuniform plan's reduction of 8.7 Gy is 78% of the upper bound on the maximum possible BED reduction, which is 11.2 Gy. The nonuniformly fractionated plans closed 78-97% of the gap between the mean liver BED of the uniform reference plan and our lower bound on the lowest achievable mean liver BED.

4.2.2. The optimal number of fractions. In order to fairly evaluate the benefit that comes from spatiotemporal fractionation, we need to consider the gain to be had from changing the fractionation schedule without altering the physical dose distribution. In particular, some cases might already benefit from hypofractionation alone, in which case the improvement we see in the spatiotemporal plans could be a combination of the benefit of allowing for greater hypofractionation in the tumor and the benefit of allowing for different dose distributions in different fractions. We computed the optimal uniformly fractionated treatment plans with up to five fractions by solving the uniform model (5) for $N = 1, \dots, 5$. In all of the cases, the five-fraction uniform plans had the lowest mean liver BED among all computed uniform plans. Hence, Table 1 compares the

5-fraction spatiotemporal plans against the best uniform plans with up to five fractions.

This agrees with previous work on the dependence of optimal fractionation schedules on the patient geometry and dose distribution. When the goal is to minimize the mean dose to a single dose-limiting parallel organ such as the liver, the optimal fractionation schedule is a function of the *effective sparing factor* $\bar{\delta}$ which was introduced independently by Unkelbach, Craft, Salari, Ramakrishnan & Bortfeld (2013) and Keller et al. (2013). It was shown that if $\bar{\delta} > \frac{(\alpha/\beta)_N}{(\alpha/\beta)_T} = 0.4$, then increasing the number of fractions is optimal, and if $\bar{\delta} < 0.4$, then lowering the number of fractions is optimal. In Case 2, the effective sparing factor is approximately 0.4, while in the remaining four cases, the sparing factor is well above 0.4. This provides an explanation for the observation that the benefit of spatiotemporal fractionation was largest for Case 2; the treatment quality of a uniformly fractionated 5-fraction treatment and a single-fraction treatment is approximately equal. Therefore, achieving a benefit through spatiotemporal fractionation relies to a lesser extent on achieving near-uniform fractionation in the normal liver.

5. Discussion

Results. The results indicate that spatiotemporal treatments can achieve substantial reductions in both BED and physical dose to the liver and other normal tissue. The approximately 15-20% mean liver BED reduction is consistent with the benefit of spatiotemporal fractionation observed in previous work for brain lesions (Unkelbach et al. 2016). Our work further shows that local optimization techniques provide high-quality plans that are close to realizing the maximum potential normal tissue dose reduction.

The results in this paper are limited to the optimization of treatment plans for two-dimensional slices of clinical liver cases in order to demonstrate the concepts introduced in the paper, but in a practical setting the computations would have to be carried out for the entire patient. There is no additional difficulty in computing locally optimal nonuniform solutions for three-dimensional cases with the same algorithms that we used in our study (Unkelbach et al. 2016). Therefore, the main challenge in extending this study to three-dimensional cases is the computational complexity of solving the SDP relaxations in Section 3. In the nonuniform fractionation optimization problem there is at least one variable for each beamlet in the fluence maps and at least one constraint for every voxel associated with a quadratic penalty function. In the SDP relaxation, with the introduction of the matrix variables X , the number of variables is roughly squared. SDP problems of this size cannot be solved for three-dimensional cases in a reasonable amount of time with the available off-the-shelf software. Even with the two-dimensional case, the solver took up to 7 hours to solve the SDP relaxation. It is an interesting direction for future research to devise customized algorithms for the solution of these large-scale SDP problems that arise for the three-dimensional cases. For now, the results obtained for the two-dimensional slices suggest that the locally

optimal solutions computed for the nonuniform fractionation problem are close to global optimality.

Static-beam IMRT versus VMAT. As mentioned in Section 4, the treatment plans in this work have a higher number of beams than a conventional IMRT plan to approximate a VMAT plan, where we expect nonuniform fractionation to display the most benefit. The nonuniformly fractionated plans that are most effective in lowering BED and total physical dose are those in which the tumor is hypofractionated while the surrounding tissue is fractionated. In other words, these plans deliver a high single-fraction dose to parts of the tumor during each fraction and a consistent lower dose to the liver and other healthy tissue throughout all of the fractions. VMAT treatments are particularly suitable for delivering such a nonuniformly fractionated plan, thanks to their characteristic low-dose bath delivered to the healthy tissue.

Uncertainty. Further research is needed to investigate the impact of various sources of uncertainty on spatiotemporally fractionated treatments. We anticipate that the effect of small changes in the radiobiological parameters is negligible, as they only slightly affect what fractionated treatments are isoeffective. Interfractional patient motion and soft tissue deformation is a greater concern for spatiotemporal treatments in which each fraction delivers precariously aligned dose distributions with sharp dose gradients in the interior of the target structure. Future research will consider including additional safeguards such as bounds on the dose gradients as well as stochastic and robust optimization methods to explicitly incorporate uncertainty into the planning. We also note that spatiotemporal fractionation can be applied to treatments with fewer distinct fractions that are delivered a number of times each; for example, a 15-fraction scheme can be designed (entirely analogously to our plans) that consist of only 3 distinct fractions that are delivered 5 times each. The repeated delivery of identical fractions may mitigate somewhat the effects of random interfraction motion. Finally, other technologies, currently under development, may also make the safe delivery of spatiotemporal liver SBRT treatments feasible in the future. These include MLC and couch tracking, and image-guided radiotherapy using an MR-linac.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. DMS-1719828. Additionally, this material was based upon work partially supported by the National Science Foundation under Grant DMS-1127914 to the Statistical and Applied Mathematical Sciences Institute. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Disclosures

The authors have no conflicts of interest to disclose.

References

- Ajdari, A. & Gbate, A. (2016). A model predictive control approach for discovering nonstationary fluence-maps in cancer radiotherapy fractionation, *in* T. Roeder, P. Frazier, R. Szechtman, E. Zhou, T. Huschka & S. Chick (eds), *Proceedings of the 2016 Winter Simulation Conference*, IEEE, pp. 2065–2075.
- Bertuzzi, A., Bruni, C., Papa, F. & Sinisgalli, C. (2013). Optimal solution for a cancer radiotherapy problem, *Journal of Mathematical Biology* **66**(1-2): 311–349.
- Bortfeld, T. (2006). IMRT: a review and preview, *Physics in Medicine and Biology* **51**(13): R363.
URL: <http://stacks.iop.org/0031-9155/51/i=13/a=R21>
- Bortfeld, T., Ramakrishnan, J., Tsitsiklis, J. & Unkelbach, J. (2015). Optimization of radiation therapy fractionation schedules in the presence of tumor repopulation, *INFORMS Journal on Computing* **27**(4): 788–803.
URL: <https://doi.org/10.1287/ijoc.2015.0659>
- Boyd, S. P. & Vandenberghe, L. (2004). *Convex Optimization*, Cambridge University Press.
- Deasy, J. O., Blanco, A. I. & Clark, V. H. (2003). CERR: A computational environment for radiotherapy research, *Medical Physics* **30**: 979–985. <https://doi.org/10.1118/1.1568978>.
- Fowler, J. F. (2010). 21 years of biologically effective dose, *British Journal of Radiology* **83**(991): 554–568.
URL: <http://dx.doi.org/10.1259/bjr/31372149>
- Gay, H. A., Jin, J.-Y., Chang, A. J. & Haken, R. K. T. (2013). Utility of normal tissue-to-tumor α/β ratio when evaluating isodoses of isoeffective radiation therapy treatment plans, *International Journal of Radiation Oncology*Biophysics* **85**(1): e81–e87.
URL: <http://dx.doi.org/10.1016/j.ijrobp.2012.09.021>
- Hall, E. & Giaccia, A. (2012). *Radiobiology for the radiologist*, Lippincott Williams & Wilkins.
- Keller, H., Hope, A., Meier, G. & Davison, M. (2013). A novel dose–volume metric for optimizing therapeutic ratio through fractionation: Retrospective analysis of lung cancer treatments, *Medical Physics* **40**(8): 084101.
- Kim, M., Gbate, A. & Phillips, M. H. (2012). A stochastic control formalism for dynamic biologically conformal radiation therapy, *European Journal of Operational Research* **219**(3): 541–556.
URL: <http://dx.doi.org/10.1016/j.ejor.2011.10.039>
- Lu, W., Chen, M., Chen, Q., Ruchala, K. & Olivera, G. (2008). Adaptive fractionation therapy: I. Basic concept and strategy, *Physics in Medicine and Biology* **53**(19): 5495–5511.
URL: <https://doi.org/10.1088/0031-9155/53/19/015>
- MathWorks (2016). MATLAB optimization toolbox documentation. <https://www.mathworks.com/help/optim/ug/fmincon.html/>.
- Mizuta, M., Date, H., Takao, S., Kishimoto, N., Sutherland, K. L., Onimaru, R. & Shirato, H. (2012). Graphical representation of the effects on tumor and OAR for determining the appropriate fractionation regimen in radiation therapy planning, *Medical Physics* **39**(11): 6791–6795.
URL: <http://dx.doi.org/10.1118/1.4757580>
- Mizuta, M., Takao, S., Date, H., Kishimoto, N., Sutherland, K. L., Onimaru, R. & Shirato, H. (2012). A mathematical study to select fractionation regimen based on physical dose distribution and the linear-quadratic model, *International Journal of Radiation Oncology*Biophysics* **84**(3): 829–833.
- MOSEK ApS (2017). MOSEK Optimization Suite release 8.0.0.94. <http://http://docs.mosek.com/8.0/intro.pdf>.

- Papp, D. & Unkelbach, J. (2014). Direct leaf trajectory optimization for volumetric modulated arc therapy planning with sliding window delivery, *Medical Physics* **41**(1): 011701. <http://dx.doi.org/10.1118/1.4835435>.
URL: <http://scitation.aip.org/content/aapm/journal/medphys/41/1/10.1118/1.4835435>
- Saberian, F., Ghate, A. & Kim, M. (2015). A two-variable linear program solves the standard linear-quadratic formulation of the fractionation problem in cancer radiotherapy, *Operations Research Letters* **43**(3): 254–258.
URL: <http://doi.org/10.1016/j.orl.2015.02.005>
- Saberian, F., Ghate, A. & Kim, M. (2016). Optimal fractionation in radiotherapy with multiple normal tissues, *Mathematical Medicine and Biology* **33**(2): 211–252.
URL: <http://dx.doi.org/10.1093/imammb/dqv015>
- Salari, E., Unkelbach, J. & Bortfeld, T. (2015). A mathematical programming approach to the fractionation problem in chemoradiotherapy, *IIE Transactions on Healthcare Systems Engineering* **5**(2): 55–73.
- Sonke, J. & Belderbos, J. (2010). Adaptive radiotherapy for lung cancer, *Seminars in Radiation Oncology* **20**(2): 94–106.
- Sturm, J. F. (1999). Using SeDuMi 1.02, a Matlab toolbox for optimization over symmetric cones, *Optimization Methods and Software* **11–12**(1–4): 625–653. See also <http://sedumi.ie.lehigh.edu/>.
- Unkelbach, J. (2015). Non-uniform spatiotemporal fractionation schemes in photon radiotherapy, in D. A. Jaffray (ed.), *World Congress on Medical Physics and Biomedical Engineering, June 7-12, 2015, Toronto, Canada*, Springer International Publishing, pp. 401–404. http://dx.doi.org/10.1007/978-3-319-19387-8_97.
- Unkelbach, J., Bussi re, M. R., Chapman, P. H., Loeffler, J. S. & Shih, H. A. (2016). Spatiotemporal fractionation schemes for irradiating large cerebral arteriovenous malformations, *International Journal of Radiation Oncology*Biophysics* **95**(3): 1067 – 1074.
URL: <http://www.sciencedirect.com/science/article/pii/S0360301616001152>
- Unkelbach, J., Craft, D., Salari, E., Ramakrishnan, J. & Bortfeld, T. (2013). The dependence of optimal fractionation schemes on the spatial dose distribution, *Physics in Medicine and Biology* **58**(1): 159–167. <https://dx.doi.org/10.1088/0031-9155/58/1/159>.
- Unkelbach, J. & Papp, D. (2015). The emergence of nonuniform spatiotemporal fractionation schemes within the standard BED model, *Medical Physics* **42**(5): 2234–2241. <https://doi.org/10.1118/1.4916684>.
- Unkelbach, J., Zeng, C. & Engelsman, M. (2013). Simultaneous optimization of dose distributions and fractionation schemes in particle radiotherapy, *Medical Physics* **40**(9): 091702.
- Vandenbergh, L. & Boyd, S. P. (1996). Semidefinite programming, *SIAM Review* **38**(1): 49–95. <http://dx.doi.org/10.1137/1038003>.
- Wein, L. M., Cohen, J. E. & Wu, J. T. (2000). Dynamic optimization of a linear-quadratic model with incomplete repair and volume-dependent sensitivity and repopulation, *International Journal of Radiation Oncology*Biophysics* **47**(4): 1073–83.
- Wu, Q., Chi, Y., Chen, P. Y., Krauss, D. J., Yan, D. & Martinez, A. (2009). Adaptive replanning strategies accounting for shrinkage in head and neck IMRT, *International Journal of Radiation Oncology*Biophysics* **75**(3): 924–932.
- Yan, D., Vicini, F., Wong, J. & Martinez, A. (1997). Adaptive radiation therapy, *Physics in Medicine and Biology* **42**(1): 123–132.
URL: <https://doi.org/10.1088/0031-9155/42/1/008>
- Yang, Y. & Xing, L. (2005). Optimization of radiotherapy dose-time fractionation with consideration of tumor specific biology, *Medical Physics* **32**(12): 2567–79.
- Zhu, C., Byrd, R., Nosedal, J. & Morales, J. L. (2011). L-BFGS-B: Software for large-scale bound-constrained optimization, version 3.0. <http://users.iems.northwestern.edu/~nosedal/lbfgsb.html>.

Supplementary Figures

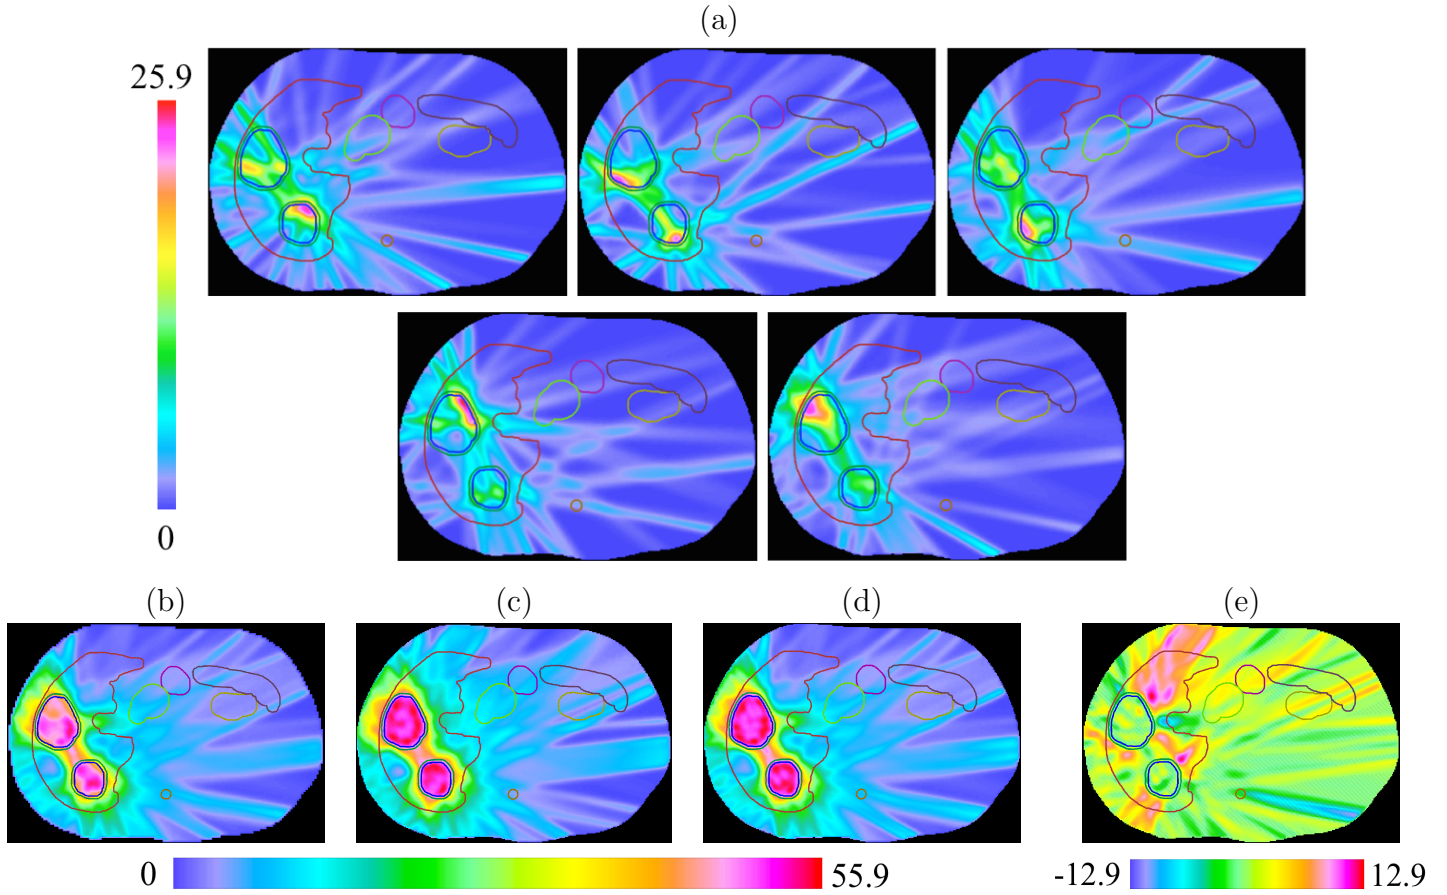


Figure 4. Nonuniformly-fractionated dose distributions for Case 3, which contains two lesions within the liver. (a) Physical dose distributions in each of the five fractions show that the nonuniformly fractionated treatment hypofractionates different parts of the tumor. (b) Total physical dose delivered throughout the nonuniformly fractionated treatment. (c) Physical dose distribution of the uniformly fractionated reference plan. (d) DEQ5 of the nonuniformly fractionated plan, which is the uniform plan that is isoeffective in delivering the same BED as the nonuniformly fractionated plan. (e) The difference between the physical dose in the uniform plan and the DEQ5 for the nonuniform plan, or (c) minus (d). This shows that the benefit of nonuniform fractionation is mostly in reduced dose in the healthy liver and in the entrance region of those beams that expose the liver the most. All numerical quantities shown are in [Gy].

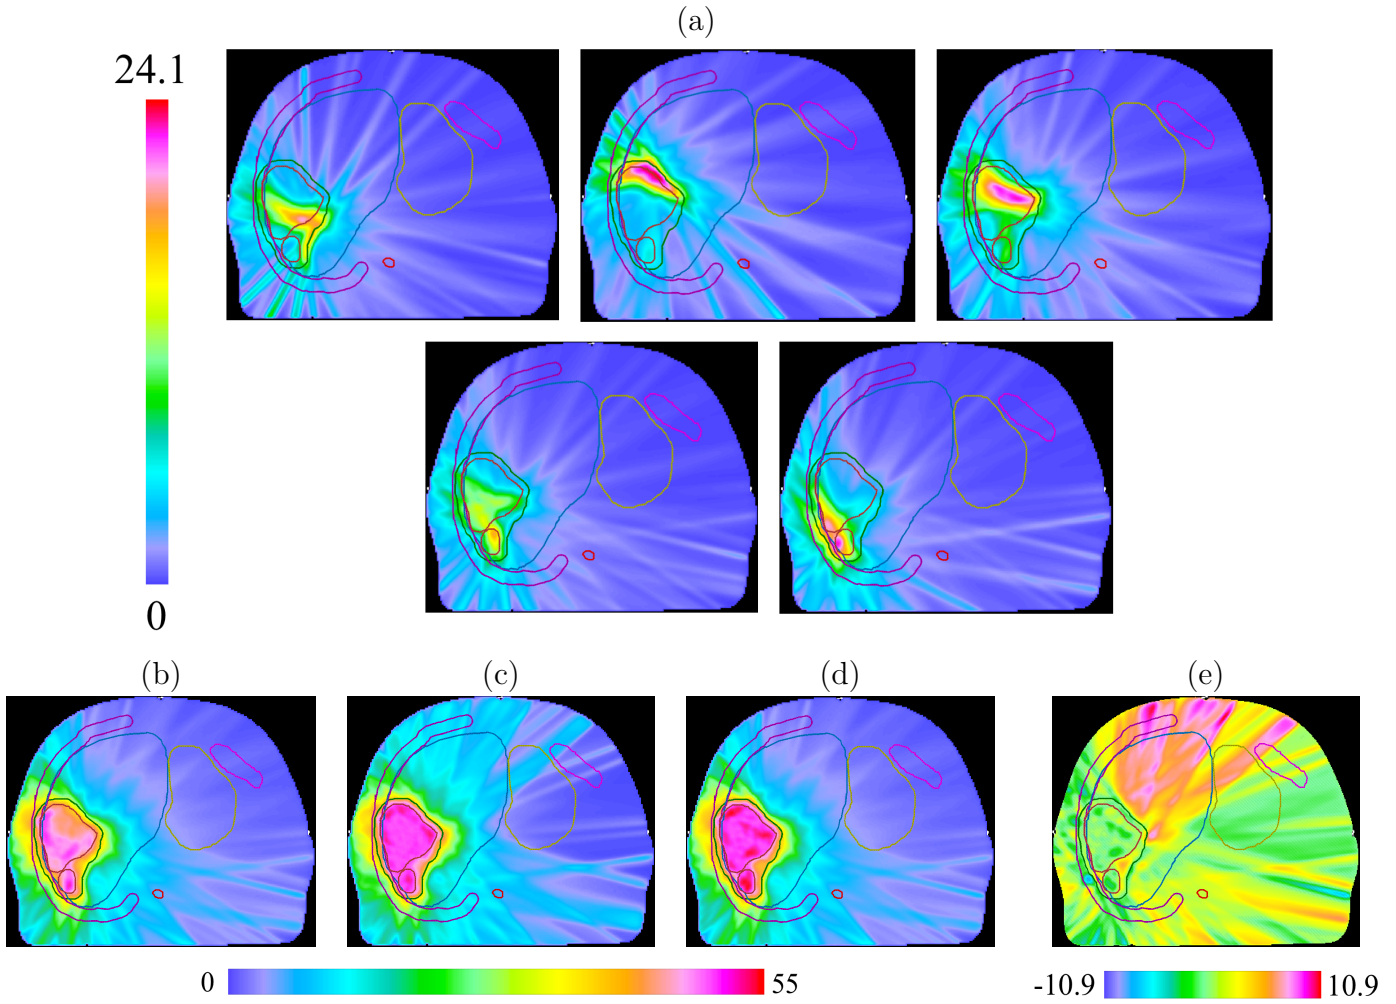


Figure 5. Dose distributions for Case 4, a lesion is abutting the chest wall. (a) Physical dose distributions for each fraction of a nonuniformly fractionated treatment. (b) Total physical dose delivered by the five nonuniform fractions. (c) Physical dose distribution of the uniform reference plan. (d) DEQ5 of the nonuniformly fractionated plan. (e) Plot of (c) minus (d). All numerical quantities shown are in [Gy].

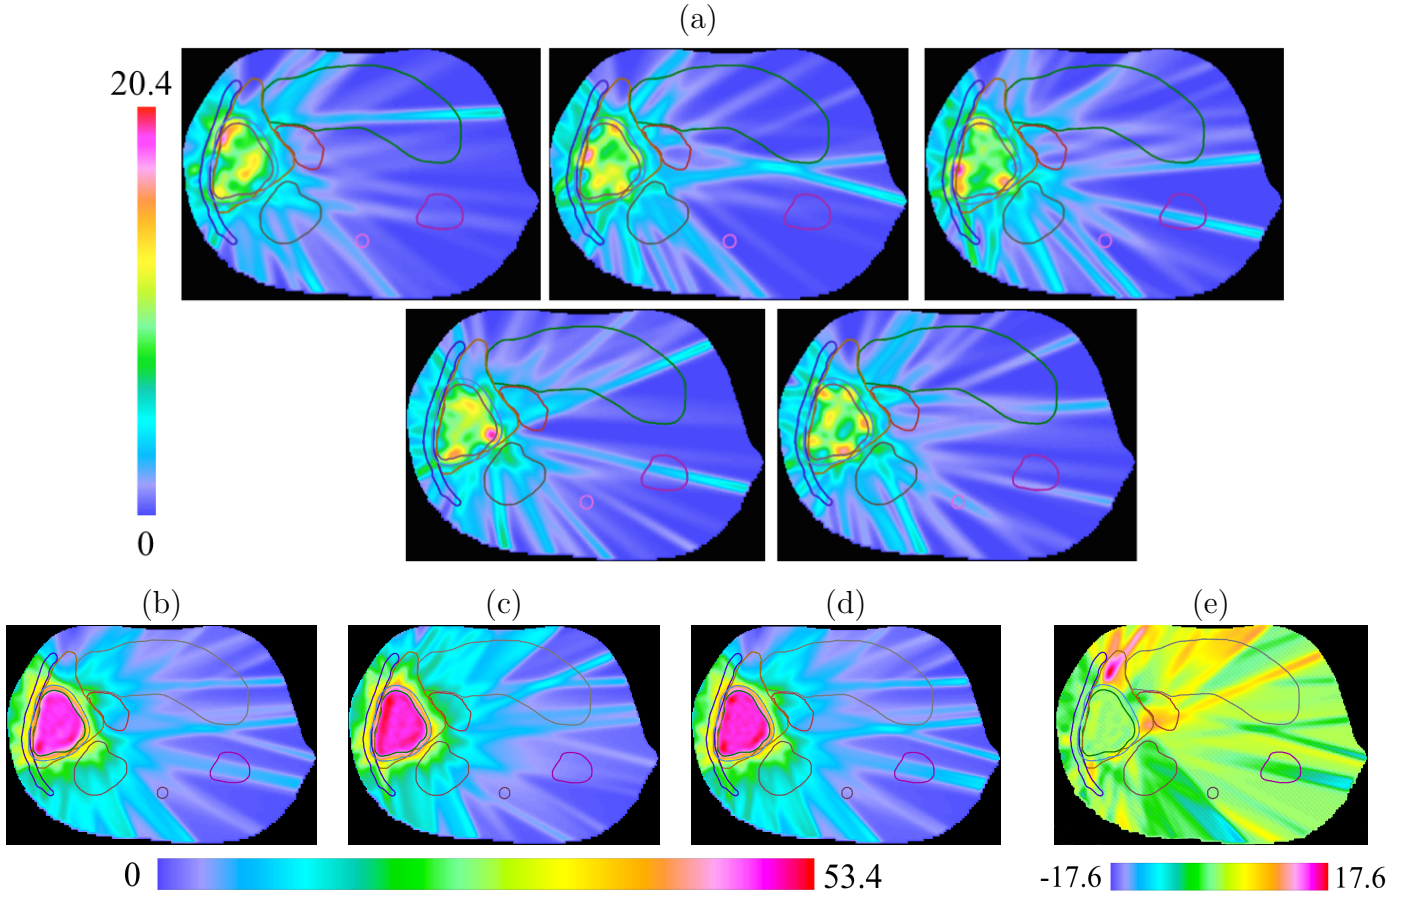




Figure 6. Dose distributions for Case 5, a lesion abutting the GI tract. (a) Physical dose distributions for each fraction of a nonuniformly fractionated treatment. (b) Total physical dose delivered by the five nonuniform fractions. (c) Physical dose distribution of the uniformly fractionated reference plan. (d) DEQ5 of the nonuniformly fractionated plan. (e) Plot of (c) minus (d). All numerical quantities shown are in [Gy].



(a)



(b)

Figure 7. Dose-volume histogram for (a) total physical dose and (b) DEQ5 for various structures in Case 1. The dotted lines are the values from the uniformly fractionated reference plan and the solid lines are from the spatiotemporal plan. These curves indicate that the spatiotemporal plan achieves an overall reduction in physical dose, as observed by the lines shifted to the left in (a). Additionally, in (b) we note that the spatiotemporal plan maintains DEQ5 in the tumor while reducing dose to the healthy liver tissue.