Superiorization as a novel strategy for linearly constrained inverse radiotherapy treatment planning

Florian Barkmann

E-mail: flobarkmann@gmail.com
Department of Medical Physics in Radiation Oncology, German Cancer Research
Center – DKFZ, Im Neuenheimer Feld 280, 69120 Heidelberg, Germany
Heidelberg Institute for Radiation Oncology – HIRO, Im Neuenheimer Feld 280,
69120 Heidelberg, Germany

Yair Censor

E-mail: yair@math.haifa.ac.il
Department of Mathematics, University of Haifa, Mt. Carmel, 3498838 Haifa, Israel

Niklas Wahl

E-mail: n.wahl@dkfz.de
Department of Medical Physics in Radiation Oncology, German Cancer Research
Center – DKFZ, Im Neuenheimer Feld 280, 69120 Heidelberg, Germany
Heidelberg Institute for Radiation Oncology – HIRO, Im Neuenheimer Feld 280,
69120 Heidelberg, Germany

Abstract.
Objective: We apply the superiorization methodology to the intensity-modulated
radiation therapy (IMRT) treatment planning problem. In superiorization, linear
voxel dose inequality constraints are the fundamental modeling tool within which a
feasibility-seeking projection algorithm will seek a feasible point. This algorithm is
then perturbed with gradient descent steps to reduce a nonlinear objective function.
Approach: Within the open-source inverse planning toolkit matRad, we implement
a prototypical algorithmic framework for superiorization using the well-established
Agmon, Motzkin, and Schoenberg (AMS) feasibility-seeking projection algorithm and
common nonlinear dose optimization objective functions. Based on this prototype,
we apply superiorization to intensity-modulated radiation therapy treatment planning
and compare its performance with feasibility-seeking and nonlinear constrained
optimization. For these comparisons, we use the TG119 water phantom and a head-
and-neck patient of the CORT dataset.
Main Results: Bare feasibility-seeking with AMS confirms previous studies, showing
it can find solutions that are nearly equivalent to those found by the established piece-
wise least-squares optimization approach. The superiorization prototype solved the
linearly constrained planning problem with similar performance to that of a general-
purpose nonlinear constrained optimizer while showing smooth convergence in both
constraint proximity and objective function reduction.
Significance: Superiorization is a useful alternative to constrained optimization in
radiotherapy inverse treatment planning. Future extensions with other approaches to

feasibility-seeking, e. g., with dose-volume constraints and more sophisticated perturbations, may unlock its full potential for high-performant inverse treatment planning.

Date: 2022-07-26

1. Introduction

Numerical optimization methods lie at the heart of state-of-the-art inverse treatment planning for intensity-modulated radiation therapy (IMRT) (Bortfeld and Thieke, 2006; Censor, 2003; Censor and Unkelbach, 2012). Usually, a clinical prescription of the treatment goals forms the input to a nonlinear multi-criteria optimization (MCO) problem with or without additional constraints, depending on the desired patient dose distribution.

During the translation of the clinical goals into an MCO problem, one distinguishes between objectives, i. e., soft goals that compete with each other, and hard constraints designed to ensure, for example, maximal tolerance doses in an organ-at-risk (OAR) and minimal dosage of the target. This versatile approach enables the treatment planner to employ arbitrary combinations of suitable (convex) nonlinear objective functions along with any choice of constraints on the voxels' doses.

This mathematical modeling allows numerical optimization of the fluence of beam elements (beamlets) using a pre-computed normalized dose mapping (Wu and Mohan, 2000). The resulting constrained nonlinear optimization problem is frequently solved by applying an extended (quasi-)Newton approach with sequential quadratic programming (SQP) or interior-point methods (Bazaraa et al., 2006; Bortfeld and Thieke, 2006; Breedveld et al., 2017; Fogliata et al., 2007; Luenberger and Ye, 2008; Wieser et al., 2017; Wu and Mohan, 2000). Recent work has substantially extended the capabilities of inverse planning through multi-criteria Pareto optimization with subsequent exploration of the Pareto surface or stochastic/robust optimization (Küfer et al., 2005; Thieke et al., 2007; Unkelbach et al., 2018).

Computational difficulties may arise in the constrained nonlinear optimization approach. Successful optimizers for nonlinear constrained optimization transform the constrained problem into an unconstrained problem using, for example, barrier functions (in the case of interior-point methods, e. g., Breedveld et al., 2017; Wächter and Biegler, 2006) and the method of Lagrange multipliers in combination with slack variables (Breedveld et al., 2017; Nocedal and Wright, 1999; Wächter and Biegler, 2006). This creates a computational burden when the number of constraints increases. Handling many constraints, for example, linear inequalities for some or all individual voxel dose bounds, can inflate the computational effort because each constraint requires a Lagrange multiplier and an additional slack variable. Possible "workarounds" include minimax-optimization in combination with auxiliary variables or usage of continuous

and differentiable maximum approximations like the LogSumExp and softmax functions (Wieser et al., 2017).

Taking a step back, however, to the starting days of treatment planning research, shows that one does not necessarily need to use an *optimization* approach to solve the purely linearly constrained IMRT problem but could use feasibility-seeking projection algorithms (Censor et al., 1988; Powlis et al., 1989).

In the context of IMRT, such feasibility-seeking translates to seeking a feasible solution that will obey the prescribed lower and upper dose bounds on doses in voxels. If no feasible solution is found, these algorithms find a proximal solution, similar to the piece-wise least-squares approach. Even though they have seen further development over the last decades (Censor, 1999, 2003; Cho and Marks, 2000) and, more recently, also extension to dose-volume constraints (Brooke et al., 2020; Gadoue et al., 2022; Penfold et al., 2017), numerical optimizers have been the preferred choice in the field due to their abilities to handle the nonlinear objective functions, e. g., (generalized) equivalent uniform dose (EUD), which are often desired when prescribing treatment goals.

The work presented here now combines nonlinear objective functions as used in optimization with feasibility-seeking within linearly constraining dose bounds by applying the superiorization method (SM). To do so, the SM uses a *superiorized version of the basic algorithm*, the latter being a user-chosen iterative feasibility-seeking algorithm, which is perturbed by interlacing reduction (not necessarily minimization) steps of the chosen (nonlinear) objective function. This practically steers the iterates of the feasibility-seeking algorithm to a feasible solution with a smaller objective function value.

As a consequence, the SM provides a flexible framework for applications and has attracted interest recently: It has demonstrated its effectiveness for image reconstruction in single-energy computed tomography (CT) (Guenter et al., 2022; Herman et al., 2012), dual-energy CT (Yang et al., 2017) and, more recently, in proton CT (Penfold and Censor, 2015; Schultze et al., 2020), by reducing total variation (TV) during image reconstruction. The SM has also been successfully applied to diverse other fields of applications, such as tomographic imaging spectrometry (Han et al., 2021) or signal recovery (Pakkaranang et al., 2020). For inverse planning, the initial theoretical applicability of superiorization using TV as an objective function has been demonstrated in Davidi et al. (2015).

Building on this preliminary work, we developed, tuned, and evaluated a prototypical superiorization solver for radiotherapy treatment planning problems. To maximize reproducibility and re-usability of our work, our superiorization approach is implemented into the validated open-source radiation therapy treatment planning toolkit matRad (Wieser et al., 2017).

This work is structured as follows: In section 2, we describe the approaches and present the specific version of the SM that we use along with the feasibility-seeking algorithm embedded in it. Section 3 includes our computational work results. Finally, in section 4, we discuss the potential of SM with possible future developments and

conclude our work in section 5.

2. Materials & Methods

This work compares three approaches to model the treatment planning problem in IMRT: (i) the *nonlinear constrained minimization approach* of minimizing an objective function subject to constraints, (ii) the *feasibility-seeking approach* searching for a feasible solution adhering to constraints without considering any objective functions to minimize, and finally, (iii) the *superiorization approach,* which perturbs the feasibility-seeking algorithm to reduce (not necessarily minimize) an objective function. Before introducing these approaches, we briefly recap the discretization of the inverse treatment planning problem.

2.1. Discretization of the inverse treatment planning problem

Computerized inverse treatment planning usually relies on a spatial discretization of the particle fluence, the patient anatomy, and, consequently, the radiation dose.

The patient is represented by a three-dimensional voxelized grid (image) with n voxels numbered $i = 1, 2, \ldots, n$. Based on this image, Q volume of interests (VOIs) S_q, $q = 1, 2, \ldots, Q$ are segmented. This allows us to represent the dose as a vector $\boldsymbol{d} = (d_i)_{i=1}^n$, whose i-th component is the radiation dose deposited within the i-th voxel. For each of the segmentations S_q, we can then easily identify its dosage by finding d_i for all $i \in S_q$.

The radiation fluence is represented as a vector intensities $\boldsymbol{x} = (x_j)_{j=1}^m$, whose j-th component is the intensity of the j-th beamlet. The dose deposition a_i^j for a unit intensity of beamlet j to voxel i can then be precomputed and stored in the *dose influence matrix* $A = (a_i^j)_{i=1, j=1}^{n,m}$, mapping $\boldsymbol{x}$ to $\boldsymbol{d}$ via $\boldsymbol{d} = A\boldsymbol{x}$.

2.2. The constrained minimization approach

In the optimization approach to IMRT treatment planning, the clinically prescribed aims are represented by various (commonly differentiable) objective functions which map the vector of beamlet intensities to the positive real numbers Wu and Mohan, 2000.

For our purposes, we limit ourselves to objective functions $f_p \colon \mathbb{R}^n \to [0, \infty)$, $p = 1, 2, \ldots, P$, operating on the radiation dose $\boldsymbol{d}$ as surrogates for clinical, dose-based goals.

A comprehensive, exemplary list of such common objective functions can be found in Wieser et al. (2017, Table 1) and, for the reader's convenience, also in Appendix A below. These objective functions, which depend on the dose, are related to the intensities $\boldsymbol{x}$ via $\boldsymbol{d} = A\boldsymbol{x}$, which is computed at each iterate/change of $\boldsymbol{x}$ during optimization.

Wishing to fulfill or decide between multiple clinical goals, the resulting multi-objective optimization problem may be scalarized using a weighted sum of several different individual objective functions for the various VOIs S_q. This approach, first introduced for least-squares (as introduced by Bortfeld et al., 1990), can today explore

a plethora of objective functions (Wieser et al., 2017; Wu and Mohan, 2000) while also satisfying hard constraints (Breedveld et al., 2017; Wieser et al., 2017):

$$
\begin{aligned}
\boldsymbol{x}^* = \quad & \arg\min_{\boldsymbol{x}} \sum_{p=1}^{P} w_p f_p(\boldsymbol{d}(\boldsymbol{x})) \\
\text{such that} \quad & c_t^L \leq c_t(\boldsymbol{d}(\boldsymbol{x})) \leq c_t^U \qquad t = 1, 2, \ldots, T, \\
& \boldsymbol{x} \geq 0
\end{aligned}
\tag{1}
$$

Here $w_p \geq 0$, for all $p = 1, 2, \ldots, P$, are user-specified weights reflecting relative importance, f_p are user-chosen individual objective functions, $\boldsymbol{x}$ is the beamlet radiation intensities vector (which is physically bound to the nonnegative real orthant), and c_t are user-chosen individual constraints with lower and upper bounds c_t^L and c_t^U, respectively. While the constraints c_t can, in principle, be nonlinear constraints, we focus here on *linear inequality constraints* representing upper and lower dose prescription bounds.

The inverse planning problem (1), solved with numerical optimization techniques, is commonly used today across treatment modalities (among others Alber and Reemtsen, 2007; Bortfeld et al., 1990; Breedveld et al., 2017; Wieser et al., 2017; Wu and Mohan, 2000). SQP or interior-point methods with a (quasi-)Newton approach are often used to solve the resulting constrained optimization problems (Bazaraa et al., 2006; Bortfeld and Thieke, 2006; Breedveld et al., 2017; Carlsson and Forsgren, 2006; Fogliata et al., 2007; Luenberger and Ye, 2008; Wieser et al., 2017; Wu and Mohan, 2000).

2.3. The feasibility-seeking approach

Our work suggests to not formulate the inverse planning as a constrained optimization problem but as a feasibility-seeking problem for a given set of constraints. The feasibility-seeking approach has been suggested before in literature (see, e. g., Censor et al., 2006, and references therein).

To solve the treatment planning problem with bare feasibility-seeking, we model our dose prescriptions as a system of linear inequalities, which leads to a full discretization of the problem. In general, for every voxel, we define a lower and upper bound of dose and seek a solution, i. e., an intensities vector that fulfills these prescriptions.

The bare feasibility-seeking approach is thus limited to treatment planning problems formulated as a system of linear inequalities and is thereby not as flexible as the constrained optimization approach. Yet, since its formulation is the backbone of the SM allowing to extend feasibility-seeking with soft goals, it will be outlined below using the notation from sections 2.1 and 2.2.

With $\boldsymbol{d}(\boldsymbol{x}) = A\boldsymbol{x}$, the beamlet radiation intensities vector $\boldsymbol{x}$ now has to be recovered from a system of linear inequalities of the form

$$
c_i^L \leq \sum_{j=1}^{m} a_i^j x_j \leq c_i^U, \quad i = 1, 2, \ldots, n\,.
\tag{2}
$$

In principle, individual lower and upper bounds c_i^L and c_i^U can be chosen for each voxel i. Since prescriptions are usually grouped per VOI S_q, the system can be rewritten as:

$$\text{For all } q = 1, 2, \ldots, Q : \ell_q \leq \sum_{j=1}^{m} a_i^j x_j^* \leq u_q \text{ for all } i \in S_q, \tag{3}$$

with ℓ_q and u_q representing the lower and upper dose bounds per VOI S_q, respectively. Since it does not make sense to prescribe positive lower bounds to OARs, these are generally chosen to be equal to zero.

Geometrically, depending on which structure S_q a voxel i belongs to, each physical dose constraint set C_i in each voxel $i = 1, 2, \ldots, n$, is a *hyperslab* (i.e., an intersection of two half-spaces) in the m-dimensional Euclidean vector space $\mathbb{R}^m$.

Aiming at satisfaction of all physical dose constraints along with the nonnegativity constraints is, thus, the following *linear feasibility problem* (which is a special case of the *convex feasibility problem*, see, e. g., Bauschke and Borwein, 1996; Censor and Zaknoon, 2018):

$$\text{Find an } \boldsymbol{x}^* \in W := \{\boldsymbol{x} \in \mathbb{R}^m \mid \text{for all, } q = 1, 2, \ldots, Q, \ \ell_q \leq \sum_{j=1}^{m} a_i^j x_j \leq u_q,$$
$$\text{for all } i \in S_q, \text{ and } \boldsymbol{x} \geq 0\} \tag{4}$$

Such feasibility-seeking problems can typically be solved by a variety of efficient projection methods, whose main advantage, which makes them successful in real-world applications, is computational (see, e.g., Censor and Unkelbach, 2012).

They commonly can handle very large-size problems of dimensions beyond which other, more sophisticated currently available, methods start to stutter or cease to be efficient. This is because the building blocks of a projection algorithm are the projections onto the given individual sets. These projections are actually easy to perform, particularly in linear cases such as hyperplanes, half-spaces, or hyperslabs.

For the purpose of this paper, we define such an iterative feasibility-seeking algorithm via an algorithmic operator $\mathcal{A} : \mathbb{R}^m \to \mathbb{R}^m$,

$$\boldsymbol{x}^0 \in \mathbb{R}^m, \ \boldsymbol{x}^{k+1} = \mathcal{A}(\boldsymbol{x}^k), \ k = 1, 2, \ldots, \tag{5}$$

whose task is to (asymptotically) find a point in W.

The algorithmic structures of projection algorithms are sequential, simultaneous, or in-between, such as in the block-iterative projection (BIP) methods (see, e. g., Davidi et al., 2009; Gordon and Gordon, 2005, and references therein) or in the more recent string-averaging projection (SAP) methods (see, e. g., Bargetz et al., 2018; Censor and Zaslavski, 2014, and references therein). An advantage of projection methods is that they work with the initial, raw data and do not require transformation of, or other operations on, the sets describing the problem.

For our prototype used here in conjunction with the SM described in the sequel, we rely on the well-established Agmon, Motzkin, and Schoenberg (AMS) relaxation method

for linear inequalities (Agmon, 1954; Motzkin and Schoenberg, 1954). Implemented sequentially and modified for handling the bounds $\boldsymbol{x} \geq 0$, it is outlined in algorithm 1. We denote $\ell := (\ell_q)_{q=1}^Q$ and $u := (u_q)_{q=1}^Q$.

Algorithm 1 The AMS Sequential Relaxation Method's algorithmic operator $\mathcal{A}^{\mathrm{AMS}}$

1: **function** $\mathcal{A}^{\mathrm{AMS}}(\boldsymbol{x}, A, u, \ell, \lambda, \nu)$

2:　　$I = \mathrm{CS}(n)$ 　　　　　　　　　　　　　　*select control sequence*

3:　　**for all** $i \in I$ **do**

4:　　　　**if** $\langle \boldsymbol{a}^i, \boldsymbol{x} \rangle > u_q$ **then** 　　　　　*u_q for structure containing i-th voxel*

5:　　　　　　$\boldsymbol{x} \leftarrow \boldsymbol{x} - \lambda \nu_i \dfrac{\langle \boldsymbol{a}^i, \boldsymbol{x} \rangle - u_q}{||\boldsymbol{a}^i||_2^2} \boldsymbol{a}^i$

6:　　　　**else if** $\langle \boldsymbol{a}^i, \boldsymbol{x} \rangle < \ell_q$ **then** 　　　　*ℓ_q for structure containing i-th voxel*

7:　　　　　　$\boldsymbol{x} \leftarrow \boldsymbol{x} - \lambda \nu_i \dfrac{\ell_q - \langle \boldsymbol{a}^i, \boldsymbol{x} \rangle}{||\boldsymbol{a}^i||_2^2} \boldsymbol{a}^i$

8:　　　　**else** 　　　　　　　　　　　　　　　　*Do nothing*

9:　　　　**end if**

10:　　**end for**

11:　　**for** $j = 1, 2, \ldots m$ **do** 　　　　　　　*Ensure nonnegativity of $\boldsymbol{x}$*

12:　　　　$x_j \leftarrow \begin{cases} x_j, & \text{for } x_j \geq 0 \\ 0, & \text{for } x_j < 0 \end{cases}$

13:　　**end for**

14:　　**return** $\boldsymbol{x}$

15: **end function**

During an iteration, algorithm 1 iterates over all rows of the dose matrix A and handles sequentially the right-hand side and the left-hand side of individual constraints from eq. (3). The *control sequence* (CS) (Censor and Zenios, 1998, Definition 5.1.1) determines the order of iterating through the matrix rows/constraints. When a corresponding voxel dose inequality is violated, the algorithm performs geometrically a projection of the current point $\boldsymbol{x}$ onto the violated half-space with a user-chosen relaxation parameter $0 < \lambda \leq 2$. The original AMS algorithm is modified in algorithm 1 to allow the relaxation for each voxel i to be weighted with ν_i and by performing projections onto the nonnegative orthant of $\mathbb{R}^m$ (in steps 11–13) to return only nonnegative intensities $\boldsymbol{x}$. The vector $\boldsymbol{a}^i = (a_i^j)_{j=1}^m$ is the i-th row of the dose matrix A and is the normal vector to the half-space represented by that row and $||\boldsymbol{a}^i||_2^2$ is its square Euclidean norm.

2.4. The superiorization method and algorithm

The SM is built upon application of a feasibility-seeking approach (section 2.3) to the constraints in the second and third lines of eq. (1). But instead of handling the constrained minimization problem of eq. (1) with a full-fledged algorithm for constrained minimization, the SM interlaces into the feasibility-seeking iterative process (i. e., the "the basic algorithm") steps that reduce locally in each iteration the objective function value.

Accordingly, the SM does not aim at finding a constrained minimum of the combined objective function $f(\boldsymbol{x}) = \sum_{p=1}^{P} w_p f_p(\boldsymbol{x})$ of eq. (1) over the constraints. It rather strives to find a feasible point that satisfies the constraints and has a reduced – not necessarily minimal – value of f.

In the following, we give a brief and focused introduction to SM. A more detailed explanation and review can be found in, e.g., Censor et al. (2021, Section II) and references therein (see also Censor, 2015, 2017; Censor et al., 2014; Davidi et al., 2009; Herman, 2014; Herman et al., 2012; Schultze et al., 2020).

In general, the SM is intended for *constrained function reduction problems* of the following form (Censor, 2017, Problem 1):

Problem 1 (The constrained function reduction problem of the SM)
Let W be a given set (such as in eq. (4)) and let $f : \mathbb{R}^m \to \mathbb{R}$ be an objective function (such as in eq. (1)). Let $\mathcal{A}$ from eq. (5) be an algorithmic operator that defines an iterative "basic algorithm" for feasibility-seeking of a point in W. Find a vector $\boldsymbol{x}^* \in W$ whose function value is lesser (but not necessarily minimal) than that of a point in W that would have been reached by applying the basic algorithm alone.

The SM approaches this question by investigating the *perturbation resilience* (Censor, 2015, Definitions 4 and 9) of $\mathcal{A}$, and then proactively using such perturbations, to locally reduce the values f of the iterates, in order to steer the iterative sequence generated by algorithm $\mathcal{A}$ to a solution with smaller objective function value. The structure of the superiorization algorithm implemented here is given by algorithm 2 with explanations here and in section 2.4.1.

Except for the initialization in steps 1–3, algorithm 2 consists of the perturbations phase (steps 5–19) and the feasibility-seeking phase (steps 20–23).

In the perturbation phase, the objective function f is reduced using gradient descent steps. The step-size β of these gradient updates is calculated by α^s where α is a fixed user-chosen constant, called *kernel*, $0 < \alpha < 1$ so that the resulting step-sizes are nonnegative and form a summable series. The power s is incremented by one until the objective function value of the newly acquired point is smaller or equal to the objective function value of the point with which the current perturbations phase was started.

The parameter N determines how many perturbations are executed before applying the next full sweep of the feasibility-seeking phase. The basic algorithm 1 with algorithmic operator $\mathcal{A}^{\mathrm{AMS}}$, used throughout this work, is indeed "perturbation resilient" (Censor and Zaslavski, 2013).

Algorithm 2 The superiorized version of the feasibility-seeking basic algorithm $\mathcal{A} = \mathcal{A}^{\mathrm{AMS}}$

1: $k \leftarrow 0$

2: $\boldsymbol{x}^k \leftarrow \boldsymbol{x}^0$

3: $s \leftarrow -1$

4: **while** stopping rule not met **do**

5:      $t \leftarrow 0$ **start of perturbation phase**

6:      $\boldsymbol{x}^{k,t} \leftarrow \boldsymbol{x}^k$

7:      **while** $t < N$ **do**   **apply N function reductions**

8:          $loop \leftarrow \texttt{true}$

9:          **while** $loop$ **do**

10:             $s \leftarrow s + 1$

11:             $\beta \leftarrow \alpha^s$ **Step size adaptation**

12:             $\boldsymbol{z} \leftarrow \boldsymbol{x}^{k,t} - \beta \nabla f(\boldsymbol{x}^{k,t})$ **Function reduction step**

13:             **if** $f(\boldsymbol{z}) \leq f(\boldsymbol{x}^{k,t})$ **then**   **Function reduction check**

14:                $t \leftarrow t + 1$

15:                $\boldsymbol{x}^{k,t} \leftarrow \boldsymbol{z}$

16:                $loop \leftarrow \texttt{false}$

17:             **end if**

18:          **end while**

19:      **end while**

20:      $\nu_k \leftarrow \eta^k \nu$ **start of feasibility-seeking phase**

21:      $\boldsymbol{x}^{k+1} \leftarrow \mathcal{A}^{\mathrm{AMS}}(\boldsymbol{x}^{k,t}, A, u, \ell, \lambda, \nu_k)$

22:      $k \leftarrow k + 1$

23: **end while**

24: **return** $\boldsymbol{x}^k$

The superiorization approach has the advantage of letting the user choose any task-specific algorithmic operator $\mathcal{A}$ that will be computationally efficient, independently of the perturbation phase, as long as perturbation resilience is preserved.

For our IMRT treatment planning problem using voxel dose constraints as introduced in eqs. (2) to (4), $\mathcal{A}$ can be – besides the chosen AMS algorithm – any of the wide variety of feasibility-seeking algorithms (see, e.g., Bauschke and Borwein, 1996; Bauschke and Koch, 2013; Cegielski, 2012; Censor and Cegielski, 2015; Censor and Zenios, 1998; Censor et al., 2012).

The principles of the SM have been presented and studied in previous publications (consult, e.g. Censor, 2015; Herman, 2014; Herman et al., 2012), but, to the best of our knowledge, this is the first work applying the SM to a treatment planning problem with an objective function of the general form $f(\boldsymbol{x}) := \sum_{p=1}^{P} w_p f_p(\boldsymbol{x})$ from eq. (1).

2.4.1. Modifications of the prototypical superiorization algorithm To control the initial step-size, we "warm start" the algorithm with larger kernel powers s within the first iteration, which substantially improves the algorithm's run-time. For our purposes, we chose an initial increment of $s \leftarrow s + 25$.

In the feasibility-seeking phase, instead of weighting all projections onto the half-spaces equally via the relaxation parameters, each projection can also be given an individual weight $0 < \nu_i < 1$ representing the "importance" of the i-th inequality constraint (i. e., voxel).

Further, as shown in step 20 of algorithm 2, weights can be reduced after each iteration to improve stability. Similar to how the step-sizes are reduced in the perturbation phase, we utilize another kernel $0 < \eta < 1$ and use its powers η^k to reduce the weights in step 20 by incrementing k after each feasibility-seeking sweep. The new weights are then calculated by $\eta^k \cdot \nu$, where ν are the initial weights.

Finally, we integrate four different control sequences to iterate through the rows of A. Apart from following the cyclic order according to voxel indices, we experimented with a random order and with sequences choosing rows with increasing or decreasing weights ν_i.

2.4.2. Stopping criteria The algorithm was terminated after a given maximal number of iterations was reached or after a certain time limit was exceeded, or when the stopping criterion formulated below was met. The default number of maximum iterations was 500 and the default wall-clock duration was set to $50\,\mathrm{min}$.

The stopping criterion that we used consists of two parts, both of which must be met for three consecutive iterations for the algorithm to stop. The first part of the stopping criterion is that the relative change of the objective function f defined by

$$\left| \frac{|f(\boldsymbol{x}^{k+1}) - f(\boldsymbol{x}^k)|}{\max\{1, f(\boldsymbol{x}^k)\}} \right| \tag{6}$$

becomes smaller than 10^{-4}.

For the second part of the stopping criterion, we define the square of the weighted L_2-norm of the constraints violations by‡

$$V(\boldsymbol{x}) := \frac{1}{n} \sum_{i=1}^{n} \frac{(\ell_q - \langle \boldsymbol{a}^i, \boldsymbol{x} \rangle)_+^2 + (\langle \boldsymbol{a}^i, \boldsymbol{x} \rangle - u_q)_+^2}{||\boldsymbol{a}^i||_2^2}, \tag{7}$$

where ℓ_q and u_q depend on which structure the i-th voxel belongs to. This second part of the stopping rule is met if the relative change of V defined by

$$\left| \frac{|V(\boldsymbol{x}^{k+1}) - V(\boldsymbol{x}^k)|}{\max\{1, V(\boldsymbol{x}^k)\}} \right| \tag{8}$$

is smaller than 10^{-3}.

All tolerances of the stopping criteria can be customized and also set to a negative number to turn off single stopping criteria or early stopping altogether.

---

‡ For any real number r we use: $(r)_+ := \max\{0, r\}$.

2.5. Implementation

The superiorization prototype described above was implemented in the open-source cross-platform "matRad" (Bangert et al., 2020; Cisternas et al., 2015; Wieser et al., 2017), which is a multi-modality radiation dose calculation and treatment planning toolkit written in Matlab. The implementation is publicly available on the matRad GitHub repository on a research branch.§

The implementation in matRad facilitates comparison against plans generated on the same datasets with a nonlinear optimizer, as matRad implements a number of common objective functions used in treatment planning (compare to Appendix A and Wieser et al. (2017, Table I)). While matRad provides interfaces to both the open-source Interior Point OPTimizer (IPOPT) (Wächter and Biegler, 2006) as well as to Matlab's built-in interior-point algorithm from `fmincon`, only the first was used for our comparisons.

matRad performs all computations in a fully-discretized model with a voxel grid. The "dose matrix" A is stored as a compressed sparse column matrix computed for all analyses using matRad's singular value decomposed pencil-beam algorithm (Bortfeld et al., 1993) validated against a clinical implementation (Wieser et al., 2017).

3. Results

3.1. Proof-of-work: Phantom plan

To demonstrate the applicability of superiorization to the IMRT treatment planning problem, we first evaluate a small example using the horseshoe phantom defined in the AAPM TG119 Report (Ezzell et al., 2009). The phantom is part of the CORT dataset (Craft et al., 2014) and consequently available with matRad.

We created an equidistantly spaced 5-field IMRT photon plan with $5\,\text{mm} \times 5\,\text{mm}$ beamlet doses (resulting in 1918 pencil-beams and a corresponding sparse dose influence matrix with 9.3×10^7 non-zero entries in 3.5×10^6 voxels).

With this setup, we generated treatment plans using three different approaches: (i) constrained minimization with IPOPT, (ii) the AMS algorithm for feasibility-seeking only, and (iii) the SM with the AMS algorithm. Different combinations of nonlinear objective functions and linear inequality constraints on dose were evaluated and compared across these approaches.

For analysis, we use dose-volume histograms (DVHs) and axial dose (difference) slices, as well as the evolution plots of the objective function values and the constraint violations.

3.1.1. *General usability of the AMS feasibility-seeking projection algorithm* We first validate that our implemented projection algorithm AMS is capable of finding comparable treatment plans to those found by established optimization algorithms

§ `https://github.com/e0404/matRad/tree/research/superiorization`

when applied to a straightforward piece-wise least-squares objective function for the unconstrained minimization of residuals.

The setup prescribes 60 Gy to the C-shaped target. To achieve this prescription, we bound the dose in the target by (60 ± 1) Gy. To the two OARs, "Core" and "Body", upper bounds (a. k. a. tolerance doses) are prescribed, resulting in the parameters given in table 1.

For nonlinear minimization with IPOPT, the tolerance doses serve as parameters for respective penalized piece-wise least-squares objective functions while for AMS the tolerances directly translate into linear inequalities and the weights proportionally increase the relaxation parameters.

Table 1: Dose inequalities / prescriptions and penalty weights used for minimization and for AMS feasibility-seeking.

VOI	w_p	tolerance / inequality constraint
Target	1000	$59\,\text{Gy} < \boldsymbol{d} < 61\,\text{Gy}$
Core	100	$\boldsymbol{d} < 20\,\text{Gy}$
Body	30	$\boldsymbol{d} < 30\,\text{Gy}$

(a) axial dose slice, IPOPT

(b) axial dose slice, AMS

(c) dose difference, IPOPT – AMS

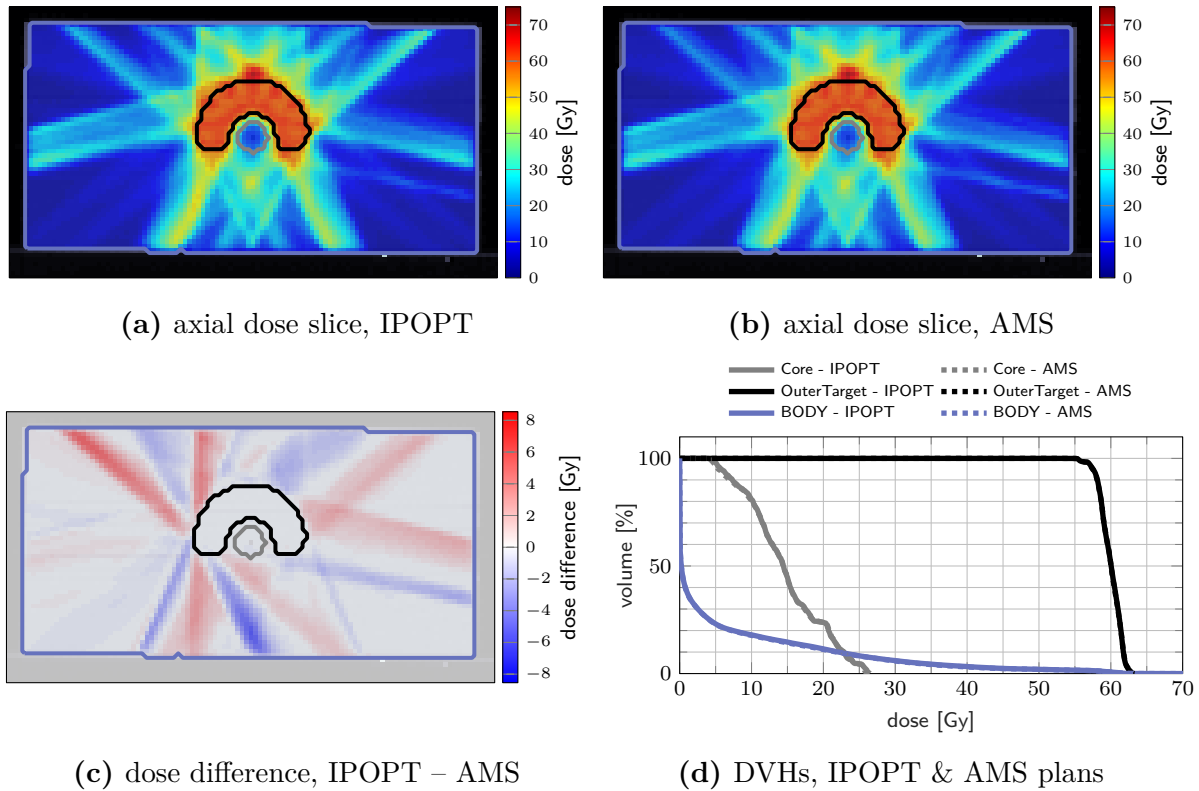**(d)** DVHs, IPOPT & AMS plans

Figure 1: Comparison of treatment plans obtained by nonlinear minimization with IPOPT and by feasibility-seeking with AMS, using the tolerances from table 1.

Figure 1 confirms that feasibility-seeking with weighted AMS is able to find dose distributions of similar quality as conventional nonlinear unconstrained minimization of a piece-wise least-squares objective function. While resulting in different intensity-modulation patterns, nearly congruent DVHs are observed.

A crude performance analysis though measures substantially longer run-times for the AMS approach (about five times slower than unconstrained minimization). This difference is mainly driven by the fact that AMS does sequential iteration through the matrix rows in each sweep.

This investigated scenario is, however, not intended to display any performance advantages of the AMS algorithm, but only to validate its behavior and confirm the long-known ability of such feasibility-seeking algorithms to yield acceptable treatment plans (Censor et al., 1988; Powlis et al., 1989).

3.1.2. Inverse planning with superiorization Using the same phantom and irradiation geometry as in section section 3.1.1, the feasibility problem used in 3.1.1 was modified to enforce some hard linear inequality constraints while minimizing an objective function. When the constraints are feasible, superiorization using AMS as the basic algorithm will find a feasible point while perturbing the iterates of the feasibility-seeking algorithm towards smaller (not necessarily minimal) function values with objective function reduction steps.

As reference, nonlinear constrained minimization with IPOPT with a logistic maximum approximation for minimum/maximum (compare Wieser et al., 2017, Table 1), was used. Three prescription scenarios were investigated: (I) linear inequalities on the target ($59\,\mathrm{Gy} < \boldsymbol{d} < 61\,\mathrm{Gy}$), (II) additional linear inequalities on the "Core" structure ($\boldsymbol{d} < 30\,\mathrm{Gy}$), and (III) only linear inequalities on the "Core" ($\boldsymbol{d} < 30\,\mathrm{Gy}$). The parameters are detailed in table 2.

Table 2: Dose inequality constraints, objective functions, and penalty weights used separately for constrained minimization and for superiorization. The Roman numerals in parentheses for the inequality constraints describe their usage in the plans, respectively. The functions in the right-hand column stem from Wieser et al. (2017, Table 1) and are identified here in Appendix A below.

VOI	w_p	$c(\boldsymbol{d})$	$f(\boldsymbol{d})$
Target	1000	$59\,\mathrm{Gy} < \boldsymbol{d} < 61\,\mathrm{Gy}$ (I & II)	$f_{\mathrm{sqdev}}(\boldsymbol{d}; 60\,\mathrm{Gy})$
Core	100	$\boldsymbol{d} < 30\,\mathrm{Gy}$ (II & III)	$f_{\mathrm{sqdev+}}(\boldsymbol{d}; 20\,\mathrm{Gy})$
Body	30	-	$f_{\mathrm{sqdev+}}(\boldsymbol{d}; 30\,\mathrm{Gy})$

Figure 2 compares dose distributions and DVHs after superiorization and after constrained minimization. The respective evolution of the objective function values and the constraint violations (calculated by the infinity norm over all inequality constraint functions, corresponding to the maximum residual) is exemplarily shown in fig. 3 for
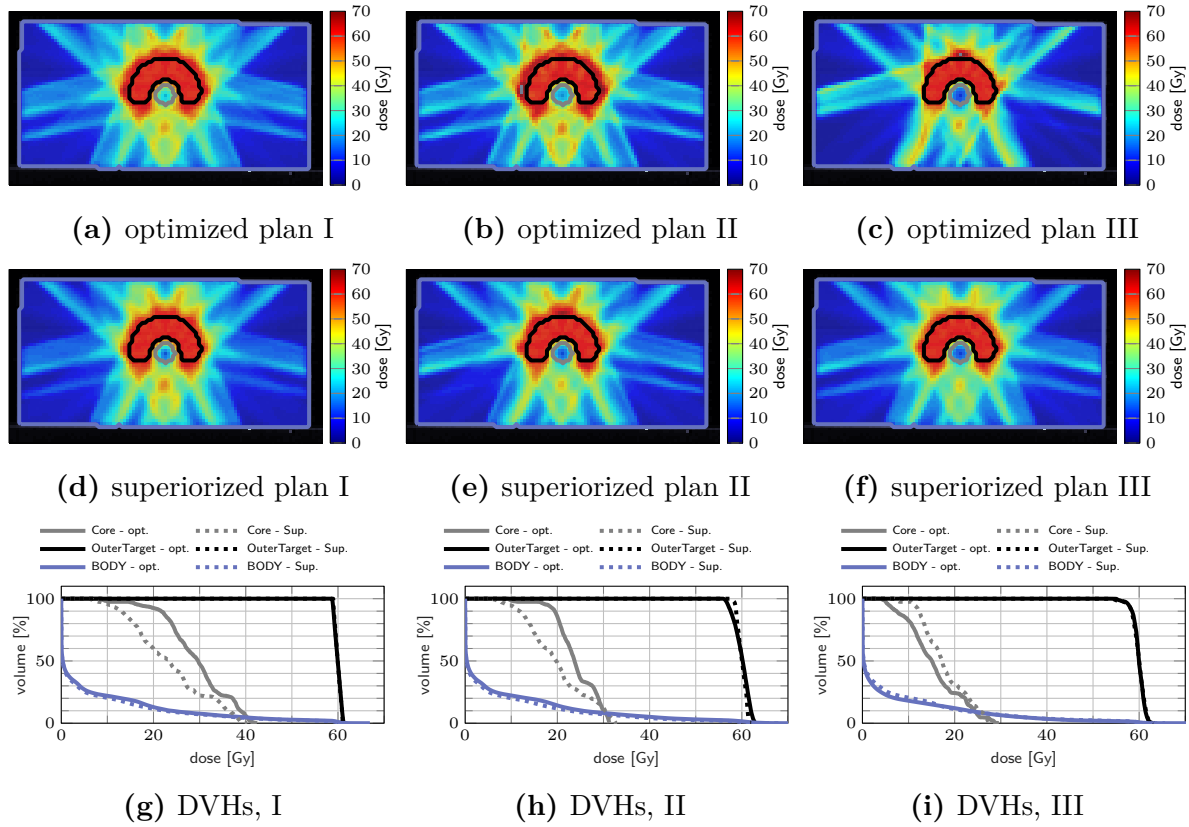
(a) optimized plan I   **(b)** optimized plan II   **(c)** optimized plan III

(d) superiorized plan I   **(e)** superiorized plan II   **(f)** superiorized plan III

(g) DVHs, I   **(h)** DVHs, II   **(i)** DVHs, III

Figure 2: Comparison of treatment plans obtained by superiorization and by constrained minimization. The top row (a)–(c) shows axial dose distribution slices after constrained minimization, the middle row (d)–(f) shows axial dose distribution slices after superiorization. The corresponding DVHs are shown in the bottom row (g)–(i).

plan I.

Comparing plan quality, both plans adhere to the linear inequality constraints when the problem is feasible (which is the case for plans I & III) as seen in the DVHs. In plan I, superiorization appears to reach better OAR sparing with reduced mean and maximum dose, while in plan III constrained minimization achieves better OAR sparing. For plan II, which poses an infeasible problem, both target coverage and mean OAR sparing are improved for superiorization, yet at higher OAR maximum dose than obtained through constrained minimization.

The evolution of the objective function and constraint violation for plan I in fig. 3 exhibits a "typical" behavior of superiorization, seeing a strong decrease in the objective function values within the first iterations, followed by a slower slight increase as the perturbations' step-sizes diminish. Both approaches were stopped after the maximum number of iterations (1000) was reached.

Nearly similar constraint violation is achieved by both methods, while constrained minimization resulted in higher objective function values than superiorization, which can be attributed to the difference in OAR sparing. For all investigated plans I–III,

superiorization showed a much "smoother" evolution of objective function and constraint violation than observed in the constrained minimization approach.
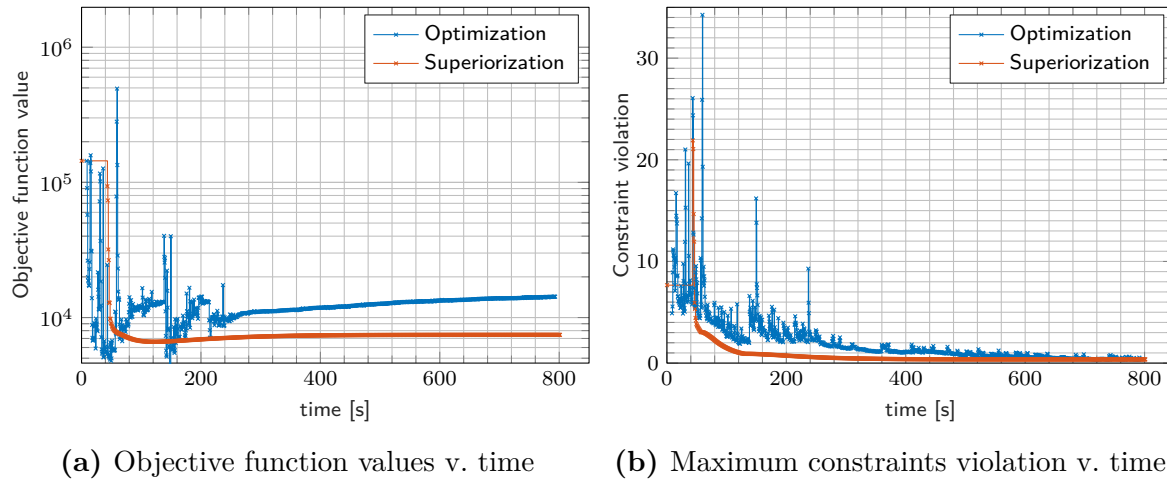

(a) Objective function values v. time   **(b)** Maximum constraints violation v. time

Figure 3: Objective Function values **(a)** and maximum constraint violation **(b)** over time for plan I shown in fig. 2. . Each cross indicates a full iteration.

3.2. Head-and-neck case

To prove the usability of superiorization in a conventional planning setting, we applied the SM to a head-and-neck case with a wider range of available objective functions, i.e., including common DVH-based objectives.

Coverage of the planning target volumes (PTVs) was enforced using voxel inequality constraints. Again, the results of superiorization were compared to those obtained by solving the constrained minimization problem. All objectives and constraints are given in table 3.

Both solvers use the same stopping criteria for the maximum constraint violation (smaller than 0.01 Gy is acceptable) and objective function change of value (smaller than 0.1 % in three consecutive iterations/sweeps).

Figure 4 shows exemplary axial dose slices and the DVHs for the plans generated with constraint minimization and with the SM. Quantitative run-time information and evolution of objective function and constraint violation are provided in fig. 5.

Both techniques were able to generate a plan that satisfies the linear inequalities up to the allowed violation threshold. Considering absolute run-time, the plan generated with the SM satisfied the stopping criteria after 400 s, with constrained minimization failing to converge until the maximum number of iterations was reached.

SM spent most of the time in the first sweep/iteration, where it focuses on multiple objective function evaluations to generate a large initial decrease (as already observed above). It continuously decreases the objective function values together with decreasing constraints violation, reaching acceptable constraints violation more slowly than the run with constrained minimization.

Table 3: Dose inequality constraints, objective functions and penalty weights used for optimization and for superiorization on the head-and-neck case. The functions in the right-hand side column are identified here in Appendix A.

VOI	w_p	$c(\boldsymbol{d})$	$f(\boldsymbol{d})$
PTV70	1000	$66.5\,\mathrm{Gy} < \boldsymbol{d} < 77\,\mathrm{Gy}$	$f_{\mathrm{sqdev}}(\boldsymbol{d}; 70\,\mathrm{Gy})$
PTV63	1000		$f_{\mathrm{sqdev}}(\boldsymbol{d}; 63\,\mathrm{Gy})$
PTV63	1000		$f_{\mathrm{minDVH}}(\boldsymbol{d}; 60\,\mathrm{Gy}, 95\,\%)$
Spinal Cord PRV	100	$\boldsymbol{d} < 50\,\mathrm{Gy}$	$f_{\mathrm{sqdev}+}(\boldsymbol{d}; 15\,\mathrm{Gy})$
Parotid L & R	100		$f_{\mathrm{sqdev}+}(\boldsymbol{d}; 10\,\mathrm{Gy})$
Optic Nerve L & R	100		$f_{\mathrm{maxDVH}}(\boldsymbol{d}; 50\,\mathrm{Gy}, 10\,\%)$
Larynx	300		$f_{\mathrm{sqdev}+}(\boldsymbol{d}; 15\,\mathrm{Gy})$
Chiasm	100		$f_{\mathrm{maxDVH}}(\boldsymbol{d}; 50\,\mathrm{Gy}, 10\,\%)$
Cerebellum	100		$f_{\mathrm{sqdev}+}(\boldsymbol{d}; 15\,\mathrm{Gy})$
Brainstem PRV	100	$\boldsymbol{d} < 30\,\mathrm{Gy}$	$f_{\mathrm{sqdev}+}(\boldsymbol{d}; 15\,\mathrm{Gy})$
NT / Body	100		$f_{\mathrm{mean}}(\boldsymbol{d})$

However, using the same stopping criteria, the SM reached a solution with a much lower objective function value (approximately one-third of the value achieved by the constrained minimization plan). This is also visible in the dose slices and DVH, which show more normal tissue/OAR sparing for the SM plan. All results are, naturally, only valid for the experiments we performed. Further work, with varying algorithmic parameters, initialization points, and stopping criteria, is necessary to make more general statements.

4. Discussion

In this work, we applied the novel superiorization method, which solves a system of linear inequalities while reducing a nonlinear objective function, to inverse radiotherapy treatment planning. On a phantom and on a head-and-neck case, we demonstrated that superiorization can produce treatment plans of similar quality to plans generated with constrained minimization.

Superiorization showed a smooth convergence behavior for both objective function reduction and constraint violation decrease, including the "typical" behavior of strong initial objective function reduction with subsequent diminishing objective function reduction – including potential slight increase – while proximity to the feasible set within the dose inequality constraints is achieved.

4.1. The mathematical framework of constrained minimization and of superiorization for treatment planning

At the heart of the superiorization algorithm lies a feasibility-seeking algorithm (in this work, the AMS relaxation method for linear inequalities). This means that
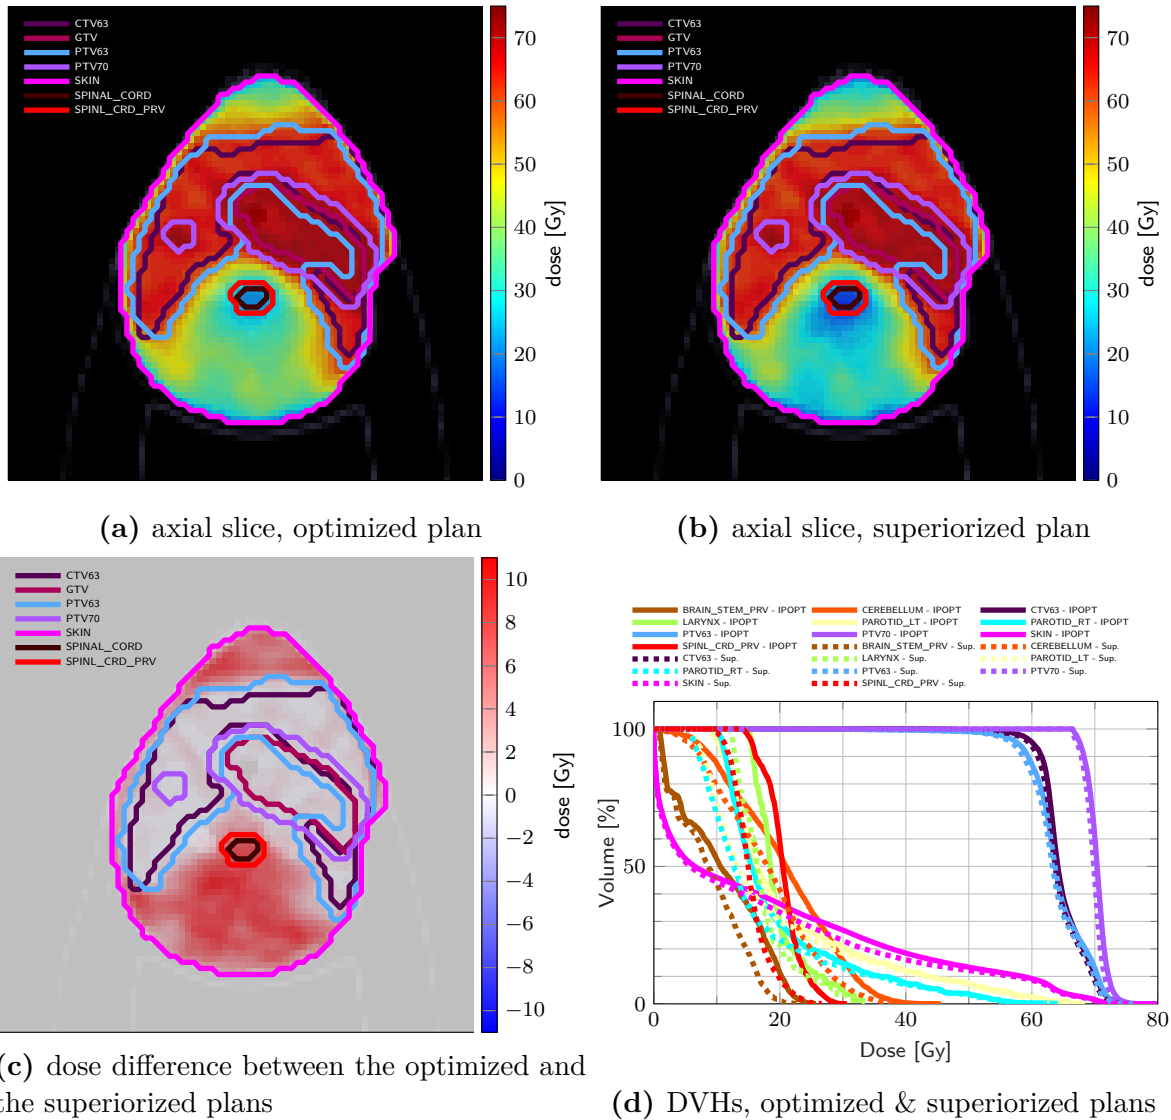
(a) axial slice, optimized plan

(b) axial slice, superiorized plan

(c) dose difference between the optimized and the superiorized plans

(d) DVHs, optimized & superiorized plans

Figure 4: Comparison of treatment plans after constrained minimization with IPOPT and after superiorization (with AMS as the basic algorithm) using the tolerances from table 3.

superiorization handles the treatment planning problem as a feasibility-seeking problem for linear inequality dose constraints that should be fulfilled while reducing (not necessarily minimizing) an objective function along the way.

Constrained optimization algorithms, on the other hand, tackle the same data, i.e., constraints and objective function, as a full-fledged optimization problem. With the IPOPT package, for example, inequality constraints become logarithmic barrier functions and are incorporated as a linear combination into the Lagrangian function, whose minimization then enforces the constraints (Wächter and Biegler, 2006).

When the problem is hardly feasible, finding the right Lagrange multipliers may then dominate the optimization problem in its initial stages. Superiorization with
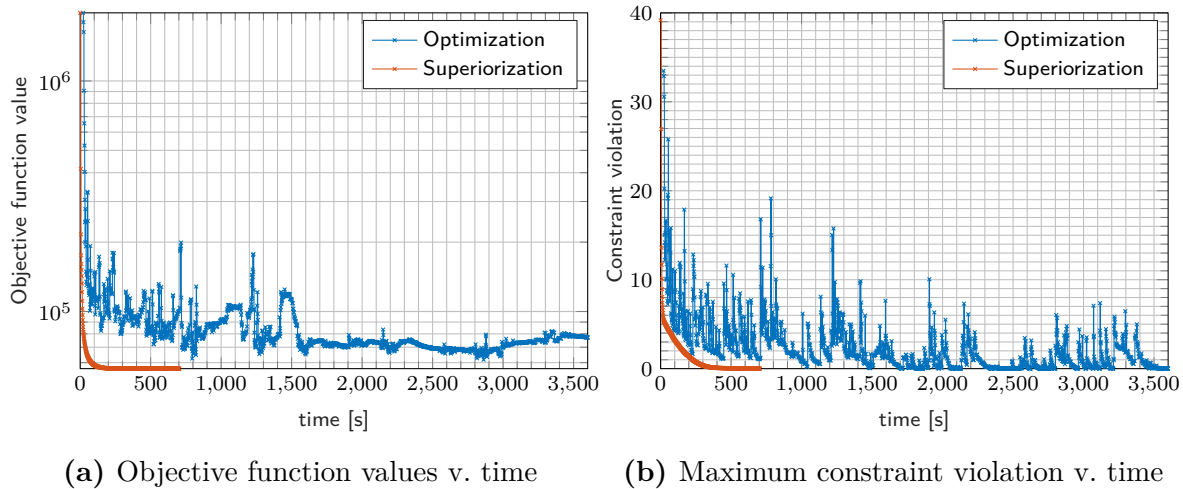
(a) Objective function values v. time   **(b)** Maximum constraint violation v. time

Figure 5: Evolution of objective Function values **(a)** and constraint violation **(b)** with run-time for the plan shown in fig. 4.

a feasibility-seeking projection algorithm will smoothly reduce the proximity to the constraints, even for infeasible constrained problems, while the perturbations in the objective function reduction phase reduce the objective function value.

Our current implementation is, however, specifically geared for linear constraints. Yet other works on feasibility-seeking have shown that other relevant constraints, like, e.g., DVH constraints, can be incorporated into the feasibility-seeking framework, since they can still be interpreted as linear inequalities on a subset (relative volume) of voxels (Brooke et al., 2020; Gadoue et al., 2022; Penfold et al., 2017).

4.2. Comparability of run-time, convergence and stopping criteria

We demonstrated that feasibility-seeking for inverse IMRT treatment planning is practically equivalent to the least-squares approach if similar prescriptions are set. However, obtaining the final solution with feasibility-seeking took more time than with unconstrained minimization with our prototype implementation in Matlab.

Stopping criteria, convergence and run-times are more comparable when considering the constrained minimization vis-à-vis superiorization. Our prototype superiorization algorithm "converged" as fast as the used constrained nonlinear minimization algorithm when using the same objective functions and linear inequalities, exhibiting more smooth progress during the iterations.

Recognizing the limited scope of the experiments presented here, our results about the superiorization method need further work to become well established. For example, the stopping criteria play a substantial role in both optimization and superiorization. Further modification of the respective parameters may lead to earlier or later stopping of either of the algorithms.

Run-time and convergence of a constrained nonlinear optimization algorithm can be influenced by incorporating second derivatives instead of relying on a low-memory

approximation to the quasi-Newton approach. In addition, alternative nonlinear minimum/maximum dose constraint implementations are possible. An advantage of the SM is that such "workarounds" are not necessary.

Despite these limitations, we demonstrated that a straightforward superiorization implementation was able to solve the given treatment planning problem arriving at dosimetrically comparable treatment plans.

4.3. Dosimetric performance

The treatment plans obtained with constrained minimization and with superiorization show some dosimetric differences. For the three different linearly constrained setups on the TG199 phantom, these differences were most pronounced on the OAR, and less pronounced for the target dosage.

In the setups with target dose inequality constraints, superiorization reached better OAR sparing. This may be a result of multiple interacting factors: the strong initial objective function decrease in superiorization pulling down the dose in the OAR, and potential too early stopping of the constrained minimizer.

Further, in the infeasible setting with linear inequality constraints on both target and OAR, superiorization has the advantage that the feasibility-seeking algorithm will still smoothly converge to a proximal point.

The improved OAR sparing did not occur when only using dose inequality constraints on the OAR. However, in this case, the differences in DVHs of the OAR are only substantial below a dose of 20 Gy and, thus, of limited significance, since a piece-wise least-squares objective was used that does not contribute to the objective function at dose values below 20 Gy.

The head-and-neck case also reproduces the better OAR sparing for all evaluated OARs, at slightly reduced target coverage for the non-constrained CTV63 and PTV63. Here, the difference in convergence speed was most significant. Through all cases, the superiorization exhibited the smooth evolution of both objective function value and constraint violation, which in turn suggests robustness against changes in the stopping criteria as well.

These encouraging results show that superiorization can create acceptable and apparently "better" treatment plans. Additional work on more cases or planning benchmarks, with varying tuning parameters of both constrained minimization and superiorization approaches is needed to assess the convergence, run-time, and dosimetric quality of the solutions.

4.4. Outlook

With the proof-of-concept put forward in this work, there are many possible directions to further investigate the application of superiorization algorithms to the radiotherapy inverse treatment planning problem. From the perspective of a treatment planner, one

may focus on enabling further constraints, e.g., DVH-based constraints, that are often used in treatment planning.

Some of these constraints are also representable as modified linear inequalities or convex and non-convex sets and, thus, can efficiently be solved using a feasibility-seeking algorithm. Even nonlinear constraints that are based, for example, on normal-tissue complication probability or equivalent uniform dose could be incorporated in the current definition of the superiorization algorithm if the "basic algorithm" in the feasibility-seeking phase of the SM is replaced by any other perturbation resilient projection method that can handle nonlinear constraints. Such algorithms exist in the literature.

Moreover, superiorization might also be extended to use more complex function reduction steps and inherent criteria. For example, a "true" backtracking line search could be performed, similar to approaches in optimization, since a perturbation resilient "basic algorithm" might be able to handle much more complex function reduction steps.

Considering these algorithmic and application-focused improvements, the SM should also be rigorously tested on radiotherapy optimization/inverse planning benchmark problems, like the TROTS dataset (Breedveld and Heijmen, 2017), as soon as it is able to handle the respective problem formulations. With this, transferability to other modalities like protons or volumetric modulated arc therapy (VMAT) is also within reach.

5. Conclusions

We introduced superiorization as a novel inverse planning technique, merging feasibility-seeking for linear inequality dose constraints with objective function reduction. Our initial comparison of superiorization with constrained minimization using linear dose-inequalities suggests possible dosimetric benefits and smoother convergence. Superiorization is thus a valuable addition to the algorithmic inverse treatment planning toolbox.

Acknowledgments

Appendix A. Objective Functions

The nonlinear objective functions used in this work correspond to those implemented in matRad (Wieser et al., 2017, Table 1):

$$f_{\text{sqdev}}(\boldsymbol{d}; \hat{d}) \quad = \frac{1}{n} \sum_i (d_i - \hat{d})^2$$

$$f_{\text{sqdev+}}(\boldsymbol{d}; \hat{d}) \quad = \frac{1}{n} \sum_i \Theta(d_i - \hat{d})(d_i - \hat{d})^2$$

$$f_{\text{sqdev-}}(\boldsymbol{d}; \hat{d}) \quad = \frac{1}{n} \sum_i \Theta(\hat{d} - d_i)(d_i - \hat{d})^2$$

$$f_{\text{minDVH}}(\boldsymbol{d}; \hat{d}, v) = \frac{1}{n} \sum_i \Theta(\hat{d} - d_i)\Theta(d_i - D_v(\boldsymbol{d}))(d_i - \hat{d})^2$$

$$f_{\text{maxDVH}}(\boldsymbol{d}; \hat{d}, v) = \frac{1}{n} \sum_i \Theta(d_i - \hat{d})\Theta(D_v(\boldsymbol{d}) - d_i)(d_i - \hat{d})^2$$

$$f_{\text{mean}}(\boldsymbol{d}) \quad = \frac{1}{n} \sum_i d_i$$

All objective functions are evaluated on a dose vector $\boldsymbol{d}$ of length n, which corresponds to the number of voxels in each VOI. $D_v(\boldsymbol{d})$ is the dose at least received by the volume fraction v and thus corresponds to the respective point in the DVH. The parameter $\hat{d}$ represents a prescribed/tolerance dose.

References

Agmon, S. (1954). "The Relaxation Method for Linear Inequalities." In: *Canadian Journal of Mathematics* 6, pp. 382–392. DOI: 10.4153/CJM-1954-037-2.

Alber, M. and R. Reemtsen (June 2007). "Intensity modulated radiotherapy treatment planning by use of a barrier-penalty multiplier method." In: *Optimization Methods and Software* 22.3, pp. 391–411. DOI: 10.1080/10556780600604940.

Bangert, M. et al. (June 5, 2020). *matRad.* Version 2.10.1. Heidelberg: Deutsches Krebsforschungszentrum. DOI: 10.5281/zenodo.3879616.

Bargetz, C., S. Reich, and R. Zalas (Jan. 2018). "Convergence properties of dynamic string-averaging projection methods in the presence of perturbations." In: *Numerical Algorithms* 77.1, pp. 185–209. DOI: 10.1007/s11075-017-0310-4.

Bauschke, H. H. and J. M. Borwein (Sept. 1, 1996). "On Projection Algorithms for Solving Convex Feasibility Problems." In: *SIAM Review* 38.3, pp. 367–426. DOI: 10.1137/S0036144593251710.

Bauschke, H. H. and V. R. Koch (2013). "Projection Methods: Swiss Army Knives for Solving Feasibility and Best Approximation Problems with Halfspaces." In: DOI: 10.1090/conm/636/12726.

Bazaraa, M. S., H. D. Sherali, and C. M. Shetty (June 13, 2006). *Nonlinear Programming: Theory and Algorithms.* 3rd ed. Hoboken, N.J: Wiley-Interscience. 872 pp.

Bortfeld, T. et al. (Oct. 1990). "Methods of image reconstruction from projections applied to conformation radiotherapy." In: *Physics in Medicine and Biology* 35.10, pp. 1423–1434. DOI: 10.1088/0031-9155/35/10/007.

Bortfeld, T., W. Schlegel, and B. Rhein (1993). "Decomposition of pencil beam kernels for fast dose calculations in three-dimensional treatment planning." In: *Medical Physics* 20.2, pp. 311–318. DOI: 10.1118/1.597070.

Bortfeld, T. and C. Thieke (2006). "Optimization of Treatment Plans, Inverse Planning." In: *New Technologies in Radiation Oncology*. Ed. by W. Schlegel, T. Bortfeld, and A.-L. Grosu. Medical Radiology. Berlin/Heidelberg: Springer-Verlag, pp. 207–220. DOI: 10.1007/3-540-29999-8_17.

Breedveld, S. and B. Heijmen (June 2017). "Data for TROTS – The Radiotherapy Optimisation Test Set." In: *Data in Brief* 12, pp. 143–149. DOI: 10.1016/j.dib.2017.03.037.

Breedveld, S., B. van den Berg, and B. Heijmen (Nov. 2017). "An interior-point implementation developed and tuned for radiation therapy treatment planning." In: *Computational Optimization and Applications* 68.2, pp. 209–242. DOI: 10.1007/s10589-017-9919-4.

Brooke, M., Y. Censor, and A. Gibali (2020). "Dynamic string-averaging CQ-methods for the split feasibility problem with percentage violation constraints arising in radiation therapy treatment planning." In: *International Transactions in Operational Research*. DOI: 10.1111/itor.12929.

Carlsson, F. and A. Forsgren (2006). "Iterative regularization in intensity-modulated radiation therapy optimization." In: *Medical Physics* 33.1, pp. 225–234. DOI: 10.1118/1.2148918.

Cegielski, A. (Sept. 13, 2012). *Iterative Methods for Fixed Point Problems in Hilbert Spaces*. New York: Springer. 316 pp.

Censor, Y. (1999). "Mathematical Aspects of Radiation Therapy Treatment Planning: Continuous Inversion Versus Full Discretization and Optimization Versus Feasibility." In: *Computational Radiology and Imaging: Therapy and Diagnostics*. Ed. by C. Börgers and F. Natterer. The IMA Volumes in Mathematics and its Applications. New York, NY: Springer, pp. 101–112. DOI: 10.1007/978-1-4612-1550-9_6.

— (2003). "Mathematical Optimization For The Inverse Problem Of Intensity-Modulated Radiation Therapy." In: *Intensity-Modulated Radiation Therapy: The State of The Art*. Ed. by J. R. Palta and T. R. Mackie. American Association of Physicists in Medicine Medical Physics Monograph No. 29. Madison, Wisconsin, USA: Medical Physics Publishing, pp. 25–49.

— (Nov. 1, 2015). "Weak and Strong Superiorization: Between Feasibility-Seeking and Minimization." In: *Analele Universitatii "Ovidius" Constanta - Seria Matematica* 23.3, pp. 41–54. DOI: 10.1515/auom-2015-0046.

— (2017). "Can Linear Superiorization Be Useful for Linear Optimization Problems?" In: *Inverse Problems* 33.4, p. 044006. DOI: 10.1088/1361-6420/33/4/044006.

Censor, Y., M. D. Altschuler, and W. D. Powlis (Jan. 1, 1988). "A computational solution of the inverse problem in radiation-therapy treatment planning." In: *Applied*

Mathematics and Computation 25.1, pp. 57–87. DOI: `10.1016/0096-3003(88)90064-1`.

Censor, Y. and A. Cegielski (Nov. 2, 2015). "Projection methods: an annotated bibliography of books and reviews." In: *Optimization* 64.11, pp. 2343–2358. DOI: `10.1080/02331934.2014.957701`.

Censor, Y., K. E. Schubert, and R. W. Schulte (2021). "Developments in Mathematical Algorithms and Computational Tools for Proton CT and Particle Therapy Treatment Planning." In: *IEEE Transactions on Radiation and Plasma Medical Sciences*, pp. 1–1. DOI: `10.1109/TRPMS.2021.3107322`.

Censor, Y. and J. Unkelbach (Apr. 1, 2012). "From analytic inversion to contemporary IMRT optimization: Radiation therapy planning revisited from a mathematical perspective." In: *Physica Medica* 28.2, pp. 109–118. DOI: `10.1016/j.ejmp.2011.04.002`.

Censor, Y. and M. Zaknoon (Apr. 26, 2018). "Algorithms and Convergence Results of Projection Methods for Inconsistent Feasibility Problems: A Review."

Censor, Y. and A. J. Zaslavski (Jan. 1, 2013). "Convergence and perturbation resilience of dynamic string-averaging projection methods." In: *Computational Optimization and Applications* 54.1, pp. 65–76. DOI: `10.1007/s10589-012-9491-x`.

— (May 4, 2014). "String-averaging projected subgradient methods for constrained minimization." In: *Optimization Methods and Software* 29.3, pp. 658–670. DOI: `10.1080/10556788.2013.841693`.

Censor, Y. and S. A. Zenios (1998). *Parallel Optimization: Theory, Algorithms, and Applications*. Numerical Mathematics and Scientific Computation. New York, NY, USA: Oxford University Press.

Censor, Y. et al. (Apr. 2006). "A unified approach for inversion problems in intensity-modulated radiation therapy." In: *Physics in Medicine and Biology* 51.10, pp. 2353–2365. DOI: `10.1088/0031-9155/51/10/001`.

Censor, Y. et al. (Apr. 1, 2012). "On the effectiveness of projection methods for convex feasibility problems with linear inequality constraints." In: *Computational Optimization and Applications* 51.3, pp. 1065–1088. DOI: `10.1007/s10589-011-9401-7`.

Censor, Y. et al. (Mar. 1, 2014). "Projected Subgradient Minimization Versus Superiorization." In: *Journal of Optimization Theory and Applications* 160.3, pp. 730–747. DOI: `10.1007/s10957-013-0408-3`.

Cho, P. S. and R. J. Marks (Jan. 2000). "Hardware-sensitive optimization for intensity modulated radiotherapy." In: *Physics in Medicine and Biology* 45.2, pp. 429–440. DOI: `10.1088/0031-9155/45/2/312`.

Cisternas, E. et al. (2015). "matRad - a multi-modality open source 3D treatment planning toolkit." In: *IFMBE Proceedings*. Vol. 51, pp. 1608–1611. DOI: `10.1007/978-3-319-19387-8_391`.

Craft, D. et al. (Dec. 1, 2014). "Shared data for intensity modulated radiation therapy (IMRT) optimization research: the CORT dataset." In: *GigaScience* 3.1. DOI: `10.1186/2047-217X-3-37`.

Davidi, R., G. Herman, and Y. Censor (July 1, 2009). "Perturbation-resilient block-iterative projection methods with application to image reconstruction from projections." In: *International transactions in operational research : a journal of The International Federation of Operational Research Societies* 16.4, pp. 505–524. DOI: `10.1111/j.1475-3995.2009.00695.x`.

Davidi, R. et al. (2015). "Feasibility-Seeking and Superiorization Algorithms Applied to Inverse Treatment Planning in Radiation Therapy." In: *Contemporary Mathematics*. Ed. by S. Reich and A. Zaslavski. Vol. 636. Providence, Rhode Island: American Mathematical Society, pp. 83–92. DOI: `10.1090/conm/636/12729`.

Ezzell, G. A. et al. (2009). "IMRT commissioning: Multiple institution planning and dosimetry comparisons, a report from AAPM Task Group 119." In: *Medical Physics* 36.11, pp. 5359–5373. DOI: `10.1118/1.3238104`.

Fogliata, A. et al. (2007). "On the performances of different IMRT treatment planning systems for selected paediatric cases." In: *Radiation Oncology* 2.1, p. 7. DOI: `10.1186/1748-717X-2-7`.

Gadoue, S. M. et al. (2022). "A dose–volume constraint (DVC) projection-based algorithm for IMPT inverse planning optimization." In: *Medical Physics*. DOI: `10.1002/mp.15504`.

Gordon, D. and R. Gordon (Jan. 1, 2005). "Component-Averaged Row Projections: A Robust, Block-Parallel Scheme for Sparse Linear Systems." In: *SIAM Journal on Scientific Computing* 27.3, pp. 1092–1117. DOI: `10.1137/040609458`.

Guenter, M. et al. (2022). "Superiorization versus regularization: A comparison of algorithms for solving image reconstruction problems with applications in computed tomography." In: *Medical Physics* 49.2, pp. 1065–1082. DOI: `10.1002/mp.15373`.

Han, W., Q. Wang, and W. Cai (May 1, 2021). "Computed tomography imaging spectrometry based on superiorization and guided image filtering." In: *Optics Letters* 46.9, pp. 2208–2211. DOI: `10.1364/OL.418355`.

Herman, G. T. (May 28, 2014). "Superiorization for Image Analysis." In: *Proceedings of the 16th International Workshop on Combinatorial Image Analysis*. Vol. 8466. IWCIA 2014. Berlin, Heidelberg: Springer-Verlag, pp. 1–7. DOI: `10.1007/978-3-319-07148-0_1`.

Herman, G. T. et al. (Aug. 22, 2012). "Superiorization: An optimization heuristic for medical physics." In: *Medical Physics* 39.9, pp. 5532–5546. DOI: `10.1118/1.4745566`.

Küfer, K.-H. et al. (2005). *Multicriteria optimization in intensity modulated radiotherapy planning*. ITWM Report 77. Fraunhofer (ITWM).

Luenberger, D. G. and Y. Ye (2008). *Linear and nonlinear programming*. 3rd ed. International Series in Operations Research and Management Science. New York, NY: Springer. 546 pp.

Motzkin, T. S. and I. J. Schoenberg (1954). "The Relaxation Method for Linear Inequalities." In: *Canadian Journal of Mathematics* 6, pp. 393–404. DOI: `10.4153/CJM-1954-038-x`.

Nocedal, J. and S. J. Wright (1999). *Numerical Optimization*. Springer series in operations research. New York: Springer. 636 pp.

Pakkaranang, N. et al. (Dec. 1, 2020). "Superiorization methodology and perturbation resilience of inertial proximal gradient algorithm with application to signal recovery." In: *The Journal of Supercomputing* 76.12, pp. 9456–9477. DOI: `10.1007/s11227-020-03215-z`.

Penfold, S. and Y. Censor (Oct. 30, 2015). "Techniques in Iterative Proton CT Image Reconstruction." In: *Sensing and Imaging* 16.1, p. 19. DOI: `10.1007/s11220-015-0122-3`.

Penfold, S. et al. (May 7, 2017). "Sparsity constrained split feasibility for dose-volume constraints in inverse planning of intensity-modulated photon or proton therapy." In: *Physics in Medicine and Biology* 62.9, pp. 3599–3618. DOI: `10.1088/1361-6560/aa602b`.

Powlis, W. D. et al. (Jan. 1, 1989). "Semi-automated radiotherapy treatment planning with a mathematical model to satisfy treatment goals." In: *International Journal of Radiation Oncology*Biology*Physics* 16.1, pp. 271–276. DOI: `10.1016/0360-3016(89)90042-4`.

Schultze, B. et al. (Feb. 2020). "An Improved Method of Total Variation Superiorization Applied to Reconstruction in Proton Computed Tomography." In: *IEEE Transactions on Medical Imaging* 39.2, pp. 294–307. DOI: `10.1109/TMI.2019.2911482`.

Thieke, C. et al. (Nov. 2007). "A new concept for interactive radiotherapy planning with multicriteria optimization: first clinical evaluation." In: *Radiotherapy and Oncology: Journal of the European Society for Therapeutic Radiology and Oncology* 85.2, pp. 292–298. DOI: `10.1016/j.radonc.2007.06.020`.

Unkelbach, J. et al. (Nov. 12, 2018). "Robust radiotherapy planning." In: *Physics in Medicine and Biology* 63.22, 22TR02. DOI: `10.1088/1361-6560/aae659`.

Wächter, A. and L. T. Biegler (Mar. 28, 2006). "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming." In: *Mathematical Programming* 106.1, pp. 25–57. DOI: `10.1007/s10107-004-0559-y`.

Wieser, H.-P. et al. (2017). "Development of the open-source dose calculation and optimization toolkit matRad." In: *Medical Physics* 44.6, pp. 2556–2568. DOI: `10.1002/mp.12251`.

Wu, Q. and R. Mohan (2000). "Algorithms and functionality of an intensity modulated radiotherapy optimization system." In: *Medical physics* 27.4, pp. 701–711. DOI: `10.1118/1.598932`.

Yang, Q., W. Cong, and G. Wang (Mar. 2017). "Superiorization-based multi-energy CT image reconstruction." In: *Inverse Problems* 33.4, p. 044014. DOI: `10.1088/1361-6420/aa5e0a`.