Available online at www.sciencedirect.com
ScienceDirect
Linear Algebra and its Applications 428 (2008) 1207–1217
LINEAR ALGEBRA AND ITS APPLICATIONS
www.elsevier.com/locate/laa

A fast algorithm for solving a linear feasibility problem with application to intensity-modulated radiation therapy

Gabor T. Herman *, Wei Chen

Department of Computer Science, The Graduate Center, City University of New York, New York, NY 10016-4309, United States

Received 27 June 2006; accepted 12 November 2006
Available online 10 January 2007
Submitted by Y. Censor

Abstract

The goal of intensity-modulated radiation therapy (IMRT) is to deliver sufficient doses to tumors to kill them, but without causing irreparable damage to critical organs. This requirement can be formulated as a linear feasibility problem. The sequential (i.e., iteratively treating the constraints one after another in a cyclic fashion) algorithm ART3 is known to find a solution to such problems in a finite number of steps, provided that the feasible region is full dimensional. We present a faster algorithm called ART3+. The idea of ART3+ is to avoid unnecessary checks on constraints that are likely to be satisfied. The superior performance of the new algorithm is demonstrated by mathematical experiments inspired by the IMRT application.

Keywords: Intensity-modulated radiation therapy; Linear feasibility problem; Algebraic reconstruction techniques

1. Introduction

Intensity-modulated radiation therapy (IMRT) is a tool for treating cancer [1]. The goal of IMRT is to deliver a sufficiently high dose to the planning target volumes (PTVs), such as the tumor cells, but to avoid depositing a harmfully high dose to organs at risk (OARs). IMRT differs from older approaches to radiation therapy in as much that it uses a multi-leaf collimator (MLC) that allows the treatment planner to control the intensity of radiation within relatively thin beamlets.

* Corresponding author. Tel.: +1 212 817 8193; fax: +1 212 817 1510.
E-mail address: gabortherman@yahoo.com (G.T. Herman).

0024-3795/$ - see front matter
doi:10.1016/j.laa.2006.11.009

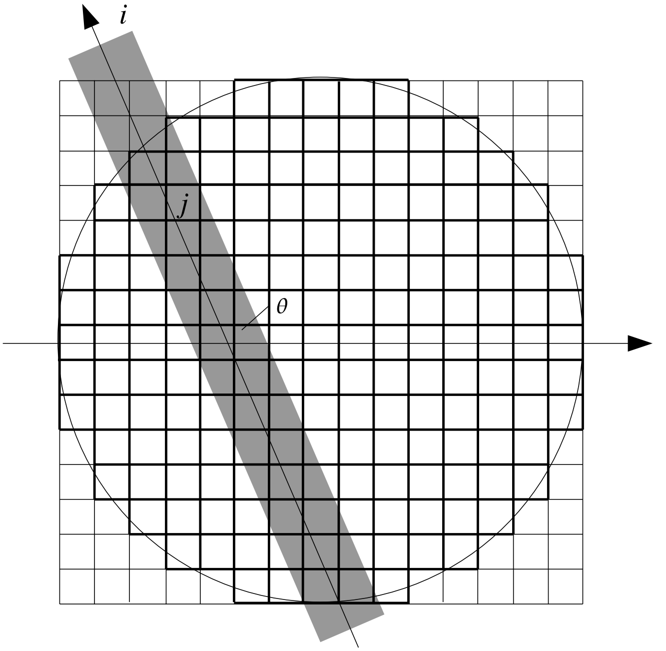

Fig. 1. The indicated circle encloses the 2D cross-section of the patient's body. The grid defines a subdivision into pixels j, for $1 \leqslant j \leqslant J$, these pixels are indicated by the heavy edges. The grey bar indicates the extent of beamlet i, for $1 \leqslant i \leqslant I$. The angle θ is between the direction of the beamlet and a fixed direction (in our case the positive horizontal axis).

We illustrate this idea by a schematic simplification of it. While IMRT planning is inherently three dimensional, due to unavoidable scattering of radiation within the body, the mathematical formulation that we discuss is formally identical in two dimensions; for this reason, we use a two dimensional version in this paper.

Consider Fig. 1. The MLC allows us to control independently the intensity of radiation x_i in each beamlet i, for $1 \leqslant i \leqslant I$. For each pixel j, $1 \leqslant j \leqslant J$, let a_j be the I-dimensional real-valued vector whose ith component is the dose contributed to pixel j if the intensity of radiation in beamlet i is unity. Then the total dose deposited into pixel j, is $\langle a_j, x\rangle = \sum_{i=1}^{I} a_{ji}x_i$. (Dose is measured in units denoted by Gy, we will not make this explicit in the rest of this paper.)

Roughly speaking, the treatment planning proceeds as follows. The grid shown in Fig. 1 is superimposed on an image (produced, for example, by computed tomography) of a cross-section of the patient and the oncologist identifies the PTVs and the OARs. For pixels in the PTVs the total deposited dose must be above a lower bound (to insure destruction of the tumor) and, without loss of generality, we may also state that it must be below an upper bound (an extreme amount of dose should not be delivered anywhere). For pixels in the OARs, the total deposited dose must be below an upper bound (so that the organ is not harmed beyond self-recovery) and must be above a lower bound (zero, since dose is by necessity nonnegative). In addition, we must have that, for $1 \leqslant i \leqslant I,\ x_i \geqslant 0$, since we cannot deliver a negative intensity by a beamlet. Again, without loss of generality, we may assume an upper bound on the x_i.

We have ended up with having to solve the following linear feasibility problem. Let I and K be positive integers (for the discussion above, $K = I + J$), let $\underline{p} \in \mathbb{R}^K,\ \bar{p} \in \mathbb{R}^K$ be such that

$\underline{p} < \bar{p}$ (meaning that $\underline{p}_k < \bar{p}_k$, for $1 \leqslant k \leqslant K$) and, for $1 \leqslant k \leqslant K$, let $a_k \in \mathbb{R}^I$ be such that $\|a_k\| = \sqrt{\langle a_k, a_k \rangle}$ is nonzero, and define

$$L_k = \{x \in \mathbb{R}^I | \underline{p}_k \leqslant \langle a_k, x \rangle \leqslant \bar{p}_k\}. \tag{1}$$

Task: Find an $x \in L \triangleq \bigcap_{k=1}^{K} L_k$.

ART3 is an iterative algorithm that solves such an interval linear feasibility problem in a finite number of steps, provided only that L is full dimensional (i.e., not a subset of any $(I-1)$-dimensional hyperplane) [2]. Below we present an alternative (we call it ART3+) that achieves the same thing, but faster. The idea of ART3+ is to avoid unnecessary checks on constraints that are likely to be satisfied.

In Section 2 we specify the algorithms ART3 and ART3+, and then we prove the finite convergence of the latter in Section 3. In Section 4 we specify some IMRT-based instances of the interval linear feasibility problem (1) and demonstrate the superiority of the performance of ART3+ over that of ART3. Our conclusions are in Section 5.

2. The algorithms

2.1. The iterative steps

Algebraic reconstruction techniques (ART) are sequential projection algorithms that, in a single step, refine the current estimate of the solution by projecting it towards the set of points satisfying a single constraint; see Chapter 11 of [3] or Section 10.4 of [4]. In our problem, every constraint is in the form of a linear interval inequality (1), which defines a hyperslab in the I-dimensional space. The ART3 and ART3+ algorithms both use projections and reflections with respect to hyperplanes according to the following definitions.

For any $a \in \mathbb{R}^I$ with $\|a\| \neq 0$ and any $b \in \mathbb{R}$, we define the hyperplane

$$H(a, b) = \{x \in \mathbb{R}^I | \langle a, x \rangle = b\}. \tag{2}$$

For $x \in \mathbb{R}^I$, the *projection* of x onto $H(a, b)$ is

$$P(x, a, b) = x + \frac{b - \langle a, x \rangle}{\|a\|^2} a \tag{3}$$

and the *reflection* of x in $H(a, b)$ is

$$R(x, a, b) = x + 2\frac{b - \langle a, x \rangle}{\|a\|^2} a. \tag{4}$$

For $1 \leqslant k \leqslant K$, let a_k, $\underline{p}_k$, $\bar{p}_k$ be specified as for (1). A single step of the ART3 and ART3+ makes use of the following operator to take care of the constraint in (1):

$$U(x, k) = \begin{cases} R(x, a_k, \underline{p}_k) & \text{if } \underline{p}_k - (\bar{p}_k - \underline{p}_k)/2 \leqslant \langle a_k, x \rangle < \underline{p}_k, \\ R(x, a_k, \bar{p}_k) & \text{if } \bar{p}_k < \langle a_k, x \rangle \leqslant \bar{p}_k + (\bar{p}_k - \underline{p}_k)/2, \\ P(x, a_k, (\bar{p}_k + \underline{p}_k)/2) & \text{if } |\langle a_k, x \rangle - (\bar{p}_k + \underline{p}_k)/2| > \bar{p}_k - \underline{p}_k, \\ x & \text{otherwise.} \end{cases} \tag{5}$$

Notice that the last line of this definition is applicable if, and only if, $x \in L_k$. Also, it is easy to see that $U(x, k) \in L_k$, for all $1 \leqslant k \leqslant K$.

The following result will be useful later on when proving the convergence of our proposed algorithm.

Lemma 1. *Let I and K be positive integers and, for $1 \leqslant k \leqslant K$, let a_k, $\underline{p}_k$ and $\bar{p}_k$ be specified as for* (1). *For any $x \in \mathbb{R}^I$, it is the case that if $z \in L = \bigcap_{k=1}^{K} L_k$, then*

$$\|U(x, k) - z\| \leqslant \|x - z\| \tag{6}$$

for all $1 \leqslant k \leqslant K$.

Proof. If we can show that (6) holds for $z \in L_k$, then the result follows since $L \subseteq L_k$. But the result for $z \in L_k$ can be trivially seen from the geometry of the situation (for a proof in a more general context, see [4, Lemma 5.5.1]). □

To control the order in which the constraints are used in the sequence of iterative steps, both algorithms make use of a data structure S, which is an ordered subset (possibly empty) of $\{1, 2, \ldots, K\}$ that may change during the execution of the algorithm. If S is not empty, we use $H(S)$ and $T(S)$ to denote its first and last element, respectively. If k occurs in S and $k \neq T(S)$, then $N(S, k)$ is the element in S that follows k. If $k = T(S)$, then $N(S, k)$ is undefined. The algorithms also make use of a special case of S that is denoted by S_0, in which each k, $1 \leqslant k \leqslant K$, occurs exactly once (i.e., S_0 is an ordered version of $\{1, 2, \ldots, K\}$). Similarly, we use x_0 to denote an arbitrary, but fixed, element of $\mathbb{R}^I$.

2.2. ART3

Algorithm 1 (*ART3*).

```
(1) x ← x0
(2) S ← S0
(3) repeat
(4)       feasible ← true
(5)       k ← H(S)
(6)       repeat
(7)          if x ∉ Lk then
(8)             x ← U(x, k)
(9)             feasible ← false
(10)         k ← N(S, k)
(11)      until k is undefined
(12) until feasible
```

The following is proved in [2].

Theorem 2. *Let I and K be positive integers and, for $1 \leqslant k \leqslant K$, let a_k, $\underline{p}_k$, $\bar{p}_k$ and L_k be specified as for* (1). *If $L = \bigcap_{k=1}^{K} L_k$ is full-dimensional, then the algorithm ART3 terminates in a finite number of steps and, at that time, $x \in L$.*

It is observed in practice that when using ART3 we get very rapidly to an x that satisfies most of the constraints. Nevertheless, the loop of Steps 6–11 has to be repeated K times and, in particular, all K constraints are checked in Step 7. This seems to be wasteful.

2.3. ART3+

This algorithm has been designed to avoid the wastefulness of ART3.

Algorithm 2 (*ART3+*).

(1) $x \leftarrow x_0$
(2) **repeat**
(3) $S \leftarrow S_0$
(4) *feasible* $\leftarrow$ *true*
(5) $k \leftarrow H(S)$
(6) **repeat**
(7) **if** $x \notin L_k$
(8) **then**
(9) $x \leftarrow U(x, k)$
(10) *feasible* $\leftarrow$ *false*
(11) **else** remove k from S
(12) $k \leftarrow N(S, k)$
(13) **until** k is undefined
(14) **if** not *feasible* **then**
(15) **repeat**
(16) $k \leftarrow H(S)$
(17) **repeat**
(18) **if** $x \notin L_k$
(19) **then** $x \leftarrow U(x, k)$
(20) **else** remove k from S
(21) $k \leftarrow N(S, k)$
(22) **until** k is undefined
(23) **until** S is empty
(24) **until** *feasible*

The essential difference of ART3+ from ART3 is the removal of satisfied constraints in Steps 11 and 20. However, to insure finite convergence we need to restart the process in Step 2, to make sure that all the constraints are simultaneously satisfied.

3. The finite convergence of ART3+

Theorem 3. *Let I and K be positive integers and, for $1 \leqslant k \leqslant K$, let a_k, $\underline{p}_k$, $\bar{p}_k$ and L_k be specified as for (1). If $L = \bigcap_{k=1}^{K} L_k$ is full-dimensional, then the algorithm ART3+ terminates in a finite number of steps and, at that time, $x \in L$.*

Prior to proving this theorem, we put our approach of moving from ART3 to ART3+ into a more general context. Consider the following.

Let X be an arbitrary nonempty set, K be a positive integer and, for $1 \leqslant k \leqslant K$, let L_k be a closed nonempty subset of X. Let U be an operator such that, for any $x \in X$ and for $1 \leqslant k \leqslant K$, $U(x, k) \in X$. Let S, $H(S)$, $T(S)$ and $N(S, k)$ be defined for $\{1, 2, \ldots, K\}$ exactly as before. In the special case of ART3 and ART3+, $X = \mathbb{R}^I$, the L_k are defined by (1), and $U(x, k)$ is defined by (5). However, this clearly does not have to be: whatever choices we make for X, the L_k and U, we get legitimate algorithms, let us call them ALG and ALG+. The following is a useful general result.

Theorem 4. *Assume the conditions of the previous paragraph. Suppose that, for any $x_0 \in X$ and for any nonempty ordered subset S_0 of $\{1, 2, \ldots, K\}$, ALG terminates in a finite number of steps. Then ALG+ has the properties that, for any $x_0 \in X$ and for any nonempty ordered subset S_0 of $\{1, 2, \ldots, K\}$,*

(i) *if Step* 6 *of ALG+ is entered, then the condition in Step* 13 *will be satisfied within a finite number of steps, and*
(ii) *if Step* 15 *of ALG+ is entered, then the condition in Step* 23 *will be satisfied within a finite number of steps.*

Proof. The truth of (i) is trivial. When Step 6 is entered, either from Step 5 or from Step 13, we have that $1 \leqslant k \leqslant K$. In either case, Step 12 replaces k by $N(S, k)$. Since S is finite, this can take place only a limited number of times without k becoming undefined (and so satisfying the condition in Step 13).

To show (ii), suppose that when Step 15 is entered S is nonempty (otherwise the condition in Step 23 is satisfied). By the argument given for (i) above, if Step 17 is entered then the condition in Step 22 will be satisfied within a finite number of steps. During those steps, it is possible that Step 20 was executed, but it is also possible that it was not. Since Step 20 cannot be executed infinitely often without eventually satisfying the condition in Step 23, the only way it is possible for (ii) to be violated is that there comes a moment in the execution of the algorithm when Step 15 is entered and after that moment Step 20 is never executed. Let x_0' and S_0' denote the values of x and S at that moment. Now consider running ALG with $x_0 = x_0'$ and $S_0 = S_0'$. Comparing line by line ALG and Steps 15–23 of ALG+, we see that never executing Step 20 in ALG+ is equivalent to ALG not terminating in a finite number of steps, contrary to the hypothesis of the theorem. □

Corollary 5. *Let I and K be positive integers and, for $1 \leqslant k \leqslant K$, let a_k, $\underline{p}_k$, $\bar{p}_k$ and L_k be specified as for* (1). *If $L = \bigcap_{k=1}^{K} L_k$ is full dimensional, then the algorithm ART*3+ *has the property that for any $x_0 \in \mathbb{R}^I$ and for any nonempty ordered subset S_0 of $\{1, 2, \ldots, K\}$,*

(i) *if Step* 6 *is entered, then the condition in Step* 13 *will be satisfied within a finite number of steps, and*
(ii) *if Step* 15 *is entered, then the condition in Step* 23 *will be satisfied within a finite number of steps.*

Proof. This is an immediate consequence of Theorems 2 and 4, and the observation that if $L = \bigcap_{k=1}^{K} L_k$ is full dimensional then the intersection of those L_k for which k is an element of a nonempty subset S_0 of $\{1, 2, \ldots, K\}$ is also full dimensional. □

In order to provide a proof of Theorem 3, we need one more preliminary result.

Lemma 6. *Let I be a positive integer. If L is a full dimensional subset of $\mathbb{R}^I$ and $x^0, x^1, x^2, \ldots$ is an infinite sequence of elements of $\mathbb{R}^I$, such that*

$$\|x^{m+1} - x\| \leqslant \|x^m - x\| \quad \textit{for all } x \in L, \quad m = 0, 1, 2, \ldots, \tag{7}$$

then the sequence $x^0, x^1, x^2 \ldots$ *is convergent.*

Proof. Proof can be provided easily by following the proof in [5] of the lemma on p. 397 in that paper. □

Proof of Theorem 3. If ART3+ stops, then the x at that time is clearly an element of L.

Now define a sequence $x^0, x^1, x^2, \ldots$ as follows: $x^0 = x_0$ and if x^m is the value of x when Step 9 or 19 is entered, then $x^{m+1} = U(x^m, k)$; i.e., it is the value of x when that step is exited. Using this definition we are going to show that the assumption that ART3+ does not stop leads to a contradiction.

From Corollary 5 it follows that if ART3+ does not stop, then Step 3 is entered infinitely often and the x at that time is such that $x \notin L_k$ for some k that occurs in S_0. But that means that Step 9 will be executed after that time and so the sequence $x^0, x^1, x^2, \ldots$ will be infinite. From Lemmas 1 and 6 it follows that the sequence converges to a point x^*.

If $x^* \notin L$, then $x^* \notin L_k$ for at least one $k \in \{1, 2, \ldots, K\}$. Since L_k is a closed set, there must exist an integer t such that $x^m \notin L_k$, for all $m \geqslant t$. Applying the same consequence of Corollary 5 as we used in the previous paragraph, we see that we are bound to come across an $m \geqslant t$ such that $x^{m+1} = U(x^m, k)$, and so $x^{m+1} \in L_k$. This contradiction shows that $x^* \in L$.

Now define the positive real numbers r_k, for $1 \leqslant k \leqslant K$, by

$$r_k = \begin{cases} \min\{\bar{p}_k - \langle a_k, x^*\rangle, \langle a_k, x^*\rangle - \underline{p}_k\} & \text{if } x^* \text{ is in the interior of } L_k, \\ (\bar{p}_k - \underline{p}_k)/2 & \text{otherwise,} \end{cases} \tag{8}$$

and the ball B by

$$B = \left\{ x \in \mathbb{R}^I \,|\, \|x - x^*\| \leqslant \min_{1 \leqslant k \leqslant K} \frac{r_k}{\|a_k\|} \right\}. \tag{9}$$

By convergence, there is an integer t such that $x^m \in B$, for all $m \geqslant t$. Consider the case when we enter Step 7 (respectively, Step 18) with $x = x^m$, with $m \geqslant t$.

$$\langle a_k, x^m\rangle - \underline{p}_k = \langle a_k, x^m - x^*\rangle + (\langle a_k, x^*\rangle - \underline{p}_k). \tag{10}$$

If x^* is in the interior of L_k, then since in (10) the second term at the right hand side is positive and the absolute value of the first term is not greater than that of the second term by (8) and (9), we have that $\langle a_k, x^m\rangle - \underline{p}_k \geqslant 0$. Similarly, $\bar{p}_k - \langle a_k, x^m\rangle \geqslant 0$, and so $x^m \in L_k$. So the condition of Step 7 (respectively, Step 18) is not satisfied and we move onto Step 11 (respectively, Step 20) without defining x^{m+1} at this time.

The alternative is that x^* is not in the interior of L_k, and so we must have that either $\langle a_k, x^*\rangle = \underline{p}_k$ or $\langle a_k, x^*\rangle = \bar{p}_k$. Consider the first case. By (8) and (9) we have either that $x^m \in L_k$ and so x^{m+1} is not defined at this time or that $\underline{p}_k - (\bar{p}_k - \underline{p}_k)/2 \leqslant \langle a_k, x^m\rangle < \underline{p}_k$ and so $x^{m+1} = R(x^m, a_k, \underline{p}_k)$. Then we see from (4) and the fact that $\langle a_k, x^*\rangle = \underline{p}_k$ that $\|x^{m+1} - x^*\| = \|x^m - x^*\|$. The same result can be similarly derived if $\langle a_k, x^*\rangle = \bar{p}_k$.

What we have shown is that, for all $m \geqslant t$, $\|x^{m+1} - x^*\| = \|x^m - x^*\|$. Since this is only possible for a sequence $x^0, x^1, x^2, \ldots$ that converges to x^* if, for $m \geqslant t$, $x^m = x^* \in L$, we get the contradiction that the algorithm terminates in a finite number of steps. □

4. Experiments

We applied both ART3 and ART3+ to some mathematical problems that are motivated by IMRT. For both algorithms we always use x_0 to be the vector of all zeros and S_0 is defined by the "natural ordering" of the pixels in Fig. 1 (row-by-row and column-by-column within each row) for the first J items, followed by the nonnegativity constraints on the beamlets.

4.1. Experimental setup

For our experiments, the square region in Fig. 1, is subdivided into 405×405 pixels of size 1 mm^2. The circle identifies 128,153 pixels (i.e., $J = 128{,}153$).

We always use the same five beam directions: $0°, 72°, 144°, 216°, 288°$. (This is a simplification of radiation therapy practice, but that is not relevant for our aim of comparing the relative speed of the two algorithms ART3 and ART3+.) For each direction, there are 103 beamlets of width 4 mm, with the center of the center beamlet going through the center of the square region. Therefore $I = 5 \times 103 = 515$ is the total number of beamlets and $K = 128{,}668$.

For $1 \leqslant i \leqslant I$ and $1 \leqslant k \leqslant J$, we set the ith component of a_k to 1 if the center of the kth pixel is within the ith beamlet, and to 0 otherwise. (A physical restatement of this is that we assume that a unit intensity of radiation injected into the ith beamlet will deposit a dose of 1 unit into exactly those pixels whose centers are contained in the ith beamlet. This is also a much simplified version of reality.) For $J + 1 \leqslant k \leqslant K$, the ith component of a_k is 1 if $k = J + i$ and is 0 otherwise. Also, for these values of k, $\underline{p}_k = 0$ and $\bar{p}_k = 10$.

The way our experiments differ from each other is in the choice of $\underline{p}_k$ and $\bar{p}_k$ for $1 \leqslant k \leqslant J$. The vectors $\underline{p}_k$ and $\bar{p}_k$ can be represented as grey-value images of the same size as illustrated in Fig. 1. This is what we do to report on our experiments: with each we associate four images, the top-left represents $\underline{p}$, the top-right $\bar{p}$, the bottom-left the dose distribution produced by ART3, and the bottom-right the dose distribution produced by ART3+, see Figs. 2–4. To be precise, for $1 \leqslant j \leqslant J$, the greyness in the jth pixel of one of these bottom images is determined by $\langle a_j, x\rangle$, where x is the output of the algorithm. In all cases, the treatment plan was defined by the $\underline{p}$ and $\bar{p}$ such that L was full-dimensional, and so both ART3 and ART3+ terminate in a finite number of steps.

4.2. The treatment plans

4.2.1. Experiment 1

The treatment plan for this experiment (see Fig. 2) was suggested to us by James M. Galvin, D.Sc. from Department of Radiation Oncology, Thomas Jefferson University, Philadelphia. It has a small PTV on the left and a big OAR on the right. For the pixels j whose centers are in the PTV, we set $\underline{p}_j = 9$ (this in practice would be 9 Gy), and for the pixels j whose centers are in the OAR, we set $\bar{p}_j = 2.5$.

4.2.2. Experiment 2

The treatment plan for this experiment (see Fig. 3) was adopted from the 2002 AAPM (American Association of Physicists in Medicine) poster of "Experience with an IMRT Head and Neck QA Phantom",[1] Andrea Nelson et al., Department of Radiation Physics, The University of Texas

[1] See http://rpc.mdanderson.org/RPC/Publications/2002 AAPM Posters/Nelson AAPM phantom presentation.pdf.

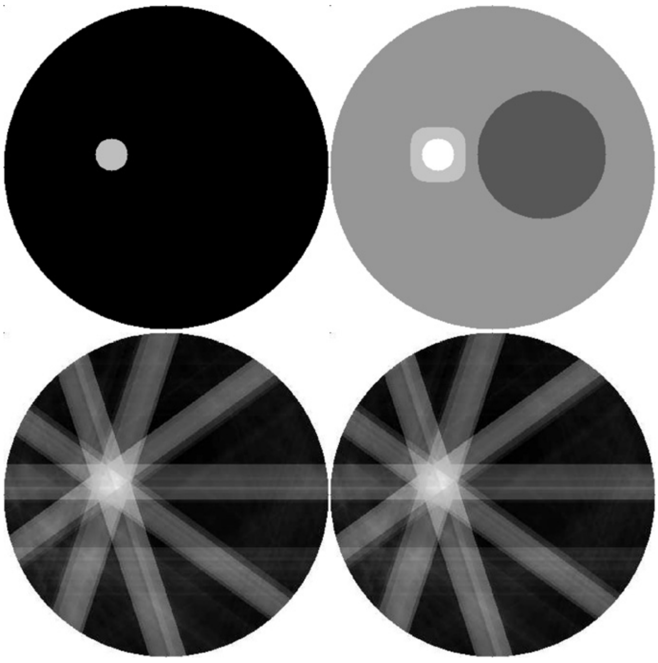

Fig. 2. The first experiment.

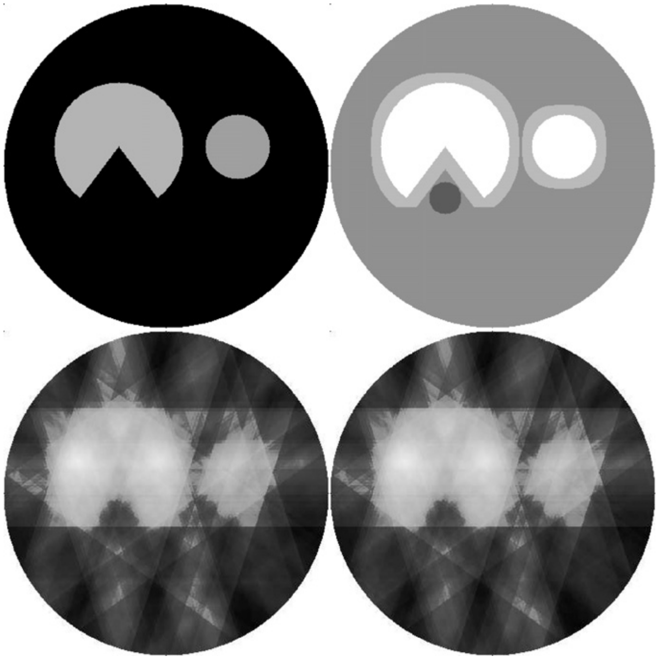

Fig. 3. The second experiment.

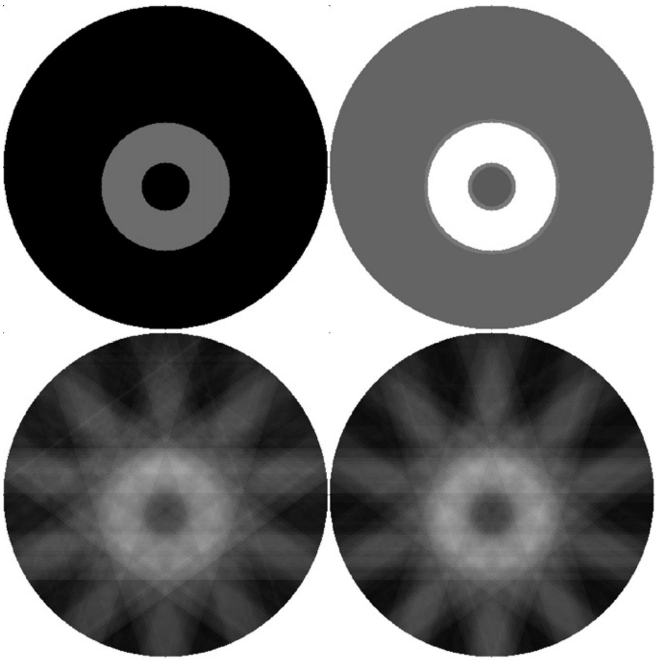

Fig. 4. The third experiment.

M.D. Anderson Cancer Center. It has a big PTV on the left and a small PTV on the right, and an even smaller OAR lies in the concave mouth of the left PTV, which makes the planning hard. For the pixels j whose centers are in the left PTV, we set $\underline{p}_j = 6.6$, for the pixels j whose centers are in the right PTV, we set $\underline{p}_j = 5.4$, and for the pixels j whose centers are in the OAR, we set $\bar{p}_j = 4.5$.

4.2.3. Experiment 3

The treatment plan for this experiment (see Fig. 4) was adopted from the talk of "IMRT Optimization Based on Physical Criteria",[2] Thomas Bortfeld et al., in the 2003 AAPM summer school on IMRT. It has a big ring shaped PTV and a small OAR right inside the PTV. For the pixels j whose centers are in the PTV, we set $\underline{p}_j = 5.4$, and for the pixels j whose centers are in the OAR, we set $\bar{p}_j = 4.5$.

4.3. Comparison of performance of ART3 and ART3+

The experiments were conducted using an Intel Xeon 1.7 MHz processor, 1 G RAM workstation. The total times needed by the two algorithms for each of the experiments are shown in Table 1. It can be seen that in these cases ART3 needed more than 1.45 times the time needed by ART3+ to get to a solution.

[2] See http://www.aapm.org/meetings/03SS/Presentations/Bortfield.pdf.

Table 1
Timing for finding solutions with the various algorithms

Algorithm	ART3 time (s)	ART3 + time (s)	Ratio
Experiment 1	3.124	2.160	1.45
Experiment 2	64.996	33.270	1.95
Experiment 3	3.048	1.800	1.69

Table 2
Timing with decreasing upper bound with the various algorithms

$\bar{p}_j$ for OAR	ART3 time (s)	ART3 + time (s)	Ratio
4.5	3.048	1.800	1.69
4.4	4.188	1.988	2.11
4.3	7.436	2.796	2.66
4.2	49.515	15.605	3.17

As the difficulty of the task increases (in the sense that L is chosen to be smaller), so does the relative advantage of ART3+ over ART3 increase. This is illustrated in Table 2 for Experiment 3. As we decrease the upper bound for the pixels in the OAR from 4.5 to 4.2, the L gets smaller and smaller and the ratio of the execution time of ART3 to that of ART3+ increases from 1.69 to 3.17.

5. Conclusions

We have treated the IMRT problem as a feasibility problem of finding a solution of a linear interval inequality system. The experiments have shown that ART3+ performs considerably faster than ART3. The reason is that ART3+ saves time by not over-frequently checking whether the current estimate satisfies all the constraints.

We point out that Theorem 4 indicates that the same idea may be also applicable to other finitely convergent algorithms.

Acknowledgments

The authors gratefully acknowledge useful discussions with Yair Censor, James M. Galvin, Ivan Kazantsev, Patrick L. Combettes. The work of the authors is supported by NIH Grant HL070472 and NSF Grant DMS0306215.

References

[1] J.R. Palta, T.R. Mackie, Intensity-Modulated Radiation Therapy: The State of the Art, Medical Physics Publishing, Madison, WI, 2003.
[2] G.T. Herman, A relaxation method for reconstructing objects from noisy X-rays, Math. Program. 8 (1975) 1–19.
[3] G.T. Herman, Image Reconstruction from Projections: The Fundamentals of Computerized Tomography, Academic Press, New York, 1980.
[4] Y. Censor, S.A. Zenios, Parallel Optimization: Theory, Algorithms, and Applications, Oxford University Press, New York, 1997.
[5] T.S. Motzkin, I.J. Schoenberg, The relaxation method for linear inequalities, Canad. J. Math. 6 (1954) 393–404.