Theoretical Computer Science 1069 (2026) 115797

Contents lists available at ScienceDirect

Theoretical Computer Science

journal homepage: www.elsevier.com/locate/tcs

SOBRA - Shielding Optimization for BRAchytherapy ⋆

Guillaume Blin [a], Adrian Miclăuș [b,*], Sebastian Ordyniak [c], Alexandru Popa [b,d]

[a] LaBRI, Université de Bordeaux CNRS UMR 5800, Talence, France
[b] Faculty of Mathematics and Computer Science, University of Bucharest, Bucharest, Romania
[c] School of Computer Science, University of Leeds, Leeds, UK
[d] National Institute for Research and Development in Informatics, Bucharest, Romania

ARTICLE INFO

Editor: Prof. Seth Pettie
Handling Editor: Kaite Harris

Keywords:
NP-hardness
APX-hardness
Approximation
FPT

ABSTRACT

In this paper, we study a combinatorial problem which arises in the development of innovative treatment strategies and equipment using tunable shields in internal radiotherapy. From an algorithmic point of view, the problem is related to circular integer word decomposition into circular binary words under constraints. We consider several variants of the problem, depending on constraints and parameters and present exact, approximation, fixed parameter tractable algorithms and NP-hardness and APX-hardness results.

1. Introduction

Radiotherapy - or *radiation therapy* - is generally used as part of cancer treatment. It uses ionizing (high-energy) radiation aiming at controlling or killing malignant cells, as a curative procedure or as part of adjuvant therapy. Radiation therapy can also damage normal cells, leading to side effects. For efficient treatment, radiation beams should precisely target the tumor site while sparing as much as possible the normal tissues. These include the vicinity of the tumor and the skin or organs that radiation must pass through to treat it (the so-called organs at risk). Although internal radiation therapy treatments are currently widespread and are considered routine, there is still room for related innovative developments. In particular, the precision of the irradiation might sometimes be improved.

Brachytherapy - sometimes named curietherapy - is a form of internal radiotherapy. It refers to a short distance (brachys in Greek) treatment of cancer with radiation from small, encapsulated radionuclide sources. Brachytherapy uses sealed radioactive sources (also called seeds) to deliver a high dose to tissues near the source. It is characterized by strong dose gradients, i.e., the dose becomes negligible in a very short distance from the source (about 10% per mm) [1] directly into or near the volume to be treated. The dose is then continuously delivered, either over a short period of time (temporary implants) or over the lifetime of the source to complete decay (permanent implants). There are many different techniques and sources available. Brachytherapy is commonly and effectively used for cervix, prostate, breast, and skin cancer. It can also be used to treat tumors in many other sites in the body.

We focus on High Dose Rate (HDR) implants. HDR brachytherapy is a form of internal radiation that temporarily exposes abnormal tissue to a large amount of radiation. The most common applications of HDR brachytherapy are in tumours of the cervix, esophagus,

⋆ A preliminary version of this article that was presented at International Workshop on Combinatorial Algorithms (IWOCA) 2016.

* Corresponding author
E-mail address: adrian.miclaus1@gmail.com (A. Miclăuș).

https://doi.org/10.1016/j.tcs.2026.115797
Received 29 June 2025; Received in revised form 7 January 2026; Accepted 29 January 2026
Available online 2 February 2026
0304-3975/© 2026 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

lungs, breasts and prostate. Under Computed Tomography (CT) and Fluoroscopy guidance, one end of a catheter is inserted, with a bronchoscope or a needle, into or next to the tumor site. Its other end is connected to a computerized machine. This machine passes a small radioactive metal seed through the catheter, which guides it to the tumor site. The seed is moved step by step through the catheter in order to cover the whole tumor site. The time spent at each position - also known as *dwell time* - is used to control the radiation dose received by each region of the tumor. At the end of a series of treatments, the catheter is removed, leaving no radioactive seeds in the body. The overall effect of HDR brachytherapy is to deliver short and precise amounts of high-dose radiation to a tumor while minimizing healthy tissue exposure.

One of the main drawbacks of the classical brachytherapy technique comes from the lack of precise modulation of the irradiation field and thus of conformation to the shape of the tumor site. Indeed, currently, the modulation of the radiation source is done by controlling the time spent at each position by the source along the catheter. The main problem is that, at any position, the irradiation is uniform and can be represented as a cylinder surrounding the catheter. This shape does not always conform to the relative placement of the tumor and the organs at risks (i.e., in the radiation field). We aim at studying the benefit of an innovative modulation technique in brachytherapy using tunable shields (as done in external radiotherapy). This approach will allow accumulating both the temporal modulation currently used and the shielding modulation. The aim is to provide treatment of better accuracy by adapting more precisely to the tumor shape. In the following, we consider a technique for modulating a unique radioactive source using a gear inspired by external radiotherapy. The use of shields will allow to preserve, for a given position along the catheter, some part of the surrounding area.

The Rotating Shield Brachytherapy (RSBT) was conceptually proposed by Ebert in [2] . In RSBT the dose is delivered through a partially shielded radiation source in an optimized step-shot fashion (as done in classical brachytherapy treatment) to improve tumor dose conformity. The intensity of radiation is modulated by the amount of time the shield is pointed in a given direction. Modern intensity-modulated brachytherapy (IMBT) or direction-modulated brachytherapy (DMBT) refers to HDR brachytherapy with dynamic shielding to shape the dose distribution. Techniques like Rotating Shield Brachytherapy (RSBT) for cervical or prostate cancer and other sector-based shielding methods have been actively studied over the past decade. Yang et al. [3] introduced the RSBT concept for cervical cancer, using a partially shielded source that rotates within a special applicator to spare organs-at-risk while treating the target. Liu et al. [4] developed a “rapid emission angle selection” method to discretize and select optimal shield angles for RSBT, reducing the problem size by focusing on a limited set of emission directions that still achieve a good dose distribution. Liu et al. [5] proposed an asymmetric dose-volume optimization with smoothness control, which added penalties to keep neighboring dwell angles at similar intensities (dwell times). This smoothness constraint improves deliverability by avoiding extreme modulation between adjacent angles. Researchers have generalized shield modulation to other designs. Liu et al. [6] presented paddle-based RSBT (P-RSBT), a novel system using a miniature multileaf collimator (“paddles”) around an electronic source. In P-RSBT, small tungsten leaves can retract to form variable emission windows (sectors) around the source, and the entire shield assembly can rotate (a conceptual design of a P-RSBT applicator can be seen in Fig. 1).The authors integrated a time constraint and “rotation stride” (discrete angle increment) into the optimization so that only a limited number of shield configurations are used, addressing the trade-off between dose quality and treatment time . Their planning was done with a custom inverse planning method (an asymmetric dose-volume constrained optimizer with smoothness regularization) and demonstrated that dynamic sector-based shielding could significantly escalate tumor dose before organ-at-risk limits are reached. Similarly, Webster et al. [7] explored a grooved shielding applicator for endorectal HDR brachytherapy, where channels in the applicator provided directional attenuation. They later demonstrated a dynamic modulation approach (termed DMBT) for rectal cancer, using a motorized rotation of the shielded applicator to achieve dose conformality improvements over conventional plans. These studies collectively showed that incorporating shields (whether continuously rotating or discrete sectors) can markedly improve target coverage and OAR sparing - e.g. dose to organs-at-risk reduced by 5-68% and target coverage improved by 19-72% in various IMBT implementations.

Introducing shield orientation as a decision variable transforms brachytherapy planning into a high-dimensional combinatorial optimization problem. Each dwell position may have multiple possible shield configurations (e.g. discrete angles or sector patterns), effectively creating a large mixed discrete-continuous search space (continuous dwell times with discrete angle choices). Blin et al. [8] formalized this problem in a paper titled “SOBRA - Shielding Optimization for BRAchytherapy”. They showed that even a simplified version of deciding which shield orientations to use at each dwell (with a goal similar to a linear dose penalty objective) is NP-complete. Intuitively, each dwell position can be seen as choosing from several binary mask settings (shield open or closed in each direction), and if one also tries to minimize the number of distinct shield configurations (to limit time or mechanical changes), it creates a complex combinatorial decision problem. Moreover, they provide approximation and quasi-polynomial algorithms for the problem mentioned above.

These formal results underscore that shield-modulated brachytherapy planning is computationally challenging (likely intractable in general), motivating the need for specialized optimization and heuristic approaches. To tackle the complexity, researchers have developed a range of optimization and heuristic methods for HDR brachytherapy with shielding. A common approach is to formulate the problem as a large-scale integer programming or mixed-integer programming (MIP) model - where binary variables might represent whether a particular shield orientation or segment is used, and continuous variables represent dwell times. Antaki et al. [9] introduced a Fast Mixed-Integer Optimization (FMIO) method for HDR brachytherapy planning. Their algorithm was designed to quickly solve the inverse planning problem with dose-volume constraints by exploiting problem structure. In a related study, Antaki et al. [10] applied a column-generation technique to the intensity-modulated brachytherapy planning problem. Another line of work has focused on faster continuous optimization once the discrete choices are fixed. Cho et al. [11] have developed convex optimization techniques for a given set of shield orientations. In particular, Cho et al. [11] reported a fast dose optimization method for RSBT based on the ADMM algorithm (using the Proximal Graph Solver, POGS) . By framing the inverse planning with dose-volume constraints

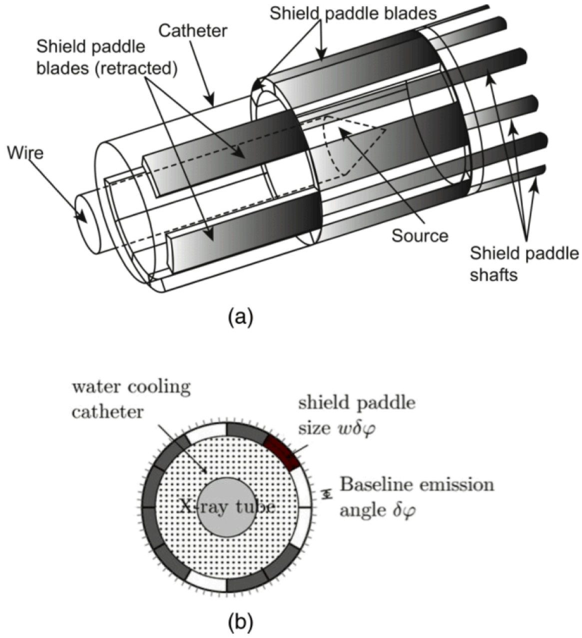

Fig. 1. A conceptual design of a P-RSBT applicator (a) 3D view. For the purpose of legibility, not all paddle shafts are drawn. (b) Cross-sectional view.

as a convex problem, they achieved a solution ≈ 18 × faster than a commercial MIP solver (CPLEX) with very similar dose metrics. Beyond optimizing dwell times, researchers have also looked at geometric optimization in shielded brachytherapy. For example, Yi et al. [12] addressed the problem of needle selection/placement in prostate RSBT. Their algorithm NEEPO ("NEEdle Position Optimization") starts with a fine "needle pool" and uses an iterative greedy removal strategy guided by a convex dose-optimization model. By adding a sparsity-promoting regularization term, NEEPO gradually drops needles that contribute least to dose coverage, resulting in plans that use fewer catheters while preserving the dosimetric gains of RSBT. This is an example of a tailored heuristic that addresses a combinatorial aspect (choosing a subset of needle tracks) on top of the shielded dose optimization. Similarly, Famulari et al. [13] and Robitaille et al. [14] developed an intermediate-energy (169-Yb) based RSBT system for prostate, which not only required new source designs but also new planning algorithms. In their workflow, precomputed dose maps for a discrete set of shield angles (e.g. 16 angles at 22.5° intervals) at each dwell position are used. Then a fluence map optimization is performed: essentially a MIP that picks the optimal dwell time at each dwell and each angle, subject to dose-volume constraints . They employed the above-mentioned FMIO method to solve this integer problem, and even investigated a combination with greedy heuristics - for example, iteratively removing unnecessary catheters to create less invasive plans while maintaining dosimetric quality. For more details see the survey by Moren et al. [15].

2. Problems definition and notations

Considering each dwell position of the irradiation source (denoted I), our main objective is to deliver to each part of the surrounding volume its proper irradiation dose. For this purpose, we use a paddle-based shielding equipment P of K paddles (also referred as sectors for ease) that can stop the radiation going through when they are not retracted. We consider the surrounding volume to be treated as a circular volume of interest divided in N subvolumes.

Definition 2.1. We say a shield sector is:

- open, if the paddle is retracted, allowing radiation going through.
- close, if the paddle is in place and radiation stopped.

Definition 2.2. A treatment plan for a given dwell position is a sequence of T shield configurations $((P^1,\tau^1),(P^2,\tau^2),\dots,(P^T,\tau^T))$, where $P^t, 1 \le t \le T$, is a paddle configuration and τ^t it its dwell time.

Definition 2.3. Each paddle configuration is a binary string of length K, $P^t = p^t_1 p^t_2 \dots p^t_K$, where p^t_k represents the state (open or closed) of the sector k of P^t.

$$p^t_k = \begin{cases} 0, & \text{if sector } k \text{ is closed} \\ 1, & \text{if sector } k \text{ is open} \end{cases} \quad \forall k,\ 1 \le k \le K.$$

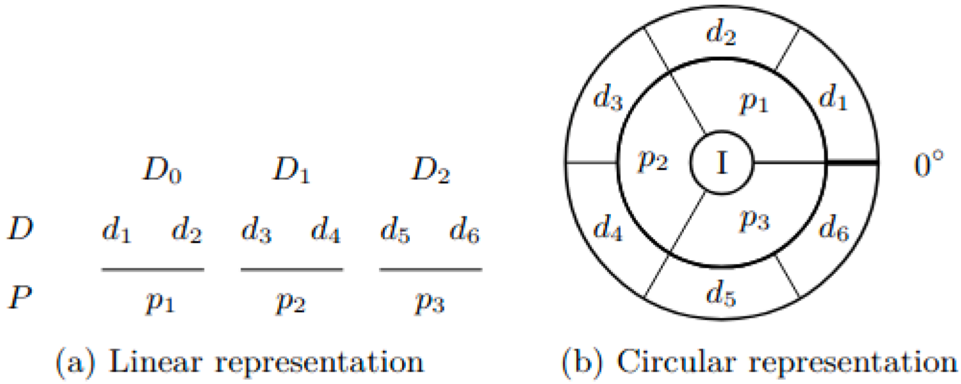

Fig. 2. Relation between P and D (K = 3, N = 6).

Definition 2.4. For each given step (P^t, τ^t) in the treatment plan, a corresponding received dose D^t by the surrounding volume is defined as a string of integers of length N, $D^t = d_1^t d_2^t \dots d_N^t$, where d_n^t corresponds to the total irradiation time the subvolume n was exposed during this step.

Definition 2.5. For a given treatment plan, the prescribed doses is a string of non-negative integers of length N, $\hat{D} = \hat{d}_1 \dots \hat{d}_N$, where $\hat{d}_n$ corresponds to the total irradiation time needed to achieve the right dose for the subvolume n.

Definition 2.6. For a given treatment plan, the total received doses is a string of integers of length N, $D = d_1 \dots d_N$, $d_n = \sum_{t=1}^{T} d_n^t$, $\forall n, 1 \le n \le N$.

Definition 2.7. The number of (consecutive) patient volumes covered by each shield sector is denoted $w = \frac{N}{K}$.

For simplicity, we assume that K divides N, so w is an integer.

Definition 2.8. Each shield sector p_k is associated with $D_k = D[(k-1) \cdot w + 1, k \cdot w]$ of length w.

The association between shield sectors and patient volumes is represented on an example in Fig. 2.

Informally, one may see P and D as circular strings, P placed inside D and representing a mask that can stop the radiation from going through.

Let us consider the practical case where one is applying a given shielded configuration (represented by P) on a patient (represented by D) for a given amount of time τ (expressed in a given unit of time).

Definition 2.9. Let us denote $D(P, \tau) = d_1, \dots, d_N$ the string of integers obtained by applying radiation for a time τ to D through the mask P.

Definition 2.10. For a mask $P = p_1 \dots p_K$,

$$p_i = \begin{cases} 0, & \text{denotes applying radiation} \\ 1, & \text{denotes no radiation} \end{cases} \qquad \forall i, 1 \le i \le K$$

to the area D_i.

Definition 2.11. For $D(P, \tau) = d_1 \dots d_N$,

$$d_i = \begin{cases} \tau, & \text{if radiation is applied for a time } \tau \\ 0, & \text{if no radiation is applied} \end{cases} \qquad \forall i, 1 \le i \le N.$$

In other words, each patient volume associated to an open sector (represented by a 1) is irradiated τ units of time, while volume associated to a closed sector (represented by a 0) is left in its previous state.

One may consider several variants of the problem, depending on constraints and parameters. First of all, the shield configuration can be considered as fixed or dynamic (one fixed mask or a minimal number of chosen masks) and provided with or without rotation capabilities (this last property is not considered here). These properties are related to manufacturing purposes and constraints.

Moreover, we can consider allowing or not irradiation overdoses ($d_n > \hat{d}_n$). Indeed, in practice, it is convenient to overdose a tumor region while one should try to not overdose organ at risks regions.

From a combinatorial point of view, there are two parameters that alter the overall treatment time; namely, the sum of the irradiation times and the number of configurations (as a transition between two configurations will require some time). In the following, we consider variants of the problem based on the previous observations. In the first two variants, the input consists of M shield configurations that are given and fixed. The goal is to decide what is the optimum amount of radiation that can be applied when allowing or disallowing overdoses.

Problem 2.1 (FixMasks). Given a prescribed dose represented as a string of non-negative integers $\hat{D} = \hat{d}_1 \ldots \hat{d}_N$ and M fixed shield configurations represented as binary strings $(P^1, P^2, \ldots, P^M)$, find a treatment plan $((P^1, \tau^1), \ldots, (P^T, \tau^T))$ minimizing $\sum_{n=1}^{N} |\hat{d}_n - d_n|$, where $d_n = \sum_{t=1}^{T} d_n^t$

While in FixMasks variant of the problem, $\hat{d}_n - d_n$ can be negative, that is overdoses are allowed, in FixMasks$^+$ variant, we moreover impose that $\forall n \le N$, $d_n \le \hat{d}_n$ - thus, forbidding overdoses. Moreover, we define the dual version of FixMasks$^+$ (in relation to the objective function).

Problem 2.2 (DFixMasks$^+$). Given a prescribed dose represented as a string of non-negative integers $\hat{D} = \hat{d}_1 \ldots \hat{d}_N$ and M fixed shield configurations represented as binary strings $(P^1, P^2, \ldots, P^M)$, find a treatment plan $((P^1, \tau^1), \ldots, (P^T, \tau^T))$ maximizing $\sum_{n=1}^{N} d_n$, where $d_n = \sum_{t=1}^{T} d_n^t$ and $d_n \le \hat{d}_n$, $\forall n$, $1 \le n \le N$.

As mentioned previously, two different criteria can be optimized in such a treatment plan one would like to either achieve the optimal difference between the prescribed dose and the actual total delivered dose using a minimal number of shield configurations or given an upper bound on the number of shield configurations, achieving the minimum reachable difference.

Problem 2.3 (MinFixMask$_{\text{OPT}}$). Given two non-negative integers K and *diff* and a string of integers $\hat{D} = \hat{d}_1 \ldots \hat{d}_N$ (with N being a multiple of K) find a treatment plan $((P^1, \tau^1), \ldots, (P^T, \tau^T))$ minimizing T such that $\sum_{n=1}^{N} |\hat{d}_n - d_n| \le diff$, where $d_n = \sum_{t=1}^{T} d_n^t$, $\forall n$, $1 \le n \le N$.

Problem 2.4 (MinFixMask$_{\text{BOUND}}$). Given two non-negative integers K and $T_{\max}$ and a string of integers $\hat{D} = \hat{d}_1 \ldots \hat{d}_N$ (with N being a multiple of K) find a treatment plan $((P^1, \tau^1), \ldots, (P^T, \tau^T))$, where $T \le T_{\max}$ minimizing $\sum_{n=1}^{N} |\hat{d}_n - d_n|$, where $d_n = \sum_{t=1}^{T} d_n^t$, $\forall n$ $1 \le n \le N$.

Similarly to FixMasks$^+$, in MinFixMask$_{\text{OPT}}$$^+$ and MinFixMask$_{\text{BOUND}}$$^+$ variants of the problem, we moreover impose that $d_n \le \hat{d}_n$, $\forall n$, $1 \le n \le N$ - thus, forbidding overdoses.

2.1. Our results

We show in Section 3, that the problems MinFixMask$_{\text{OPT}}$, MinFixMask$_{\text{BOUND}}$ are NP-complete. In Section 4, we show that FixMasks$^+$ is NP-complete and APX-hard. In Section 5, we prove that DFixMasks$^+$ is NP-complete, FPT with parameter the number of fixed shield configurations and present a polynomial $\frac{1}{M}$-approximation, an exponential $\frac{k}{M}$-approximation, where M is the number of fixed shield configurations and an $(1 - \frac{1}{e})$ by exploiting the monotonicity and submodularity of the objective function. We show in Section 6, that MinFixMask$_{\text{OPT}}$, MinFixMask$_{\text{OPT}}$$^+$, MinFixMask$_{\text{BOUND}}$ and MinFixMask$_{\text{BOUND}}$$^+$ can be solved in quasi-polynomial time if $\hat{d}_{\max}$ is bounded by a polynomial in the number of prescribed doses, where $\hat{d}_{\max}$ is the maximum prescribed dose to a subvolume of the patient. In the same section, we show that the problems MinFixMask$_{\text{OPT}}$ and $MinFixMasks^+_{OPT}$ can be approximated in polynomial time within a factor of $\log \hat{d}_{\max}$ of the optimum. We show, in Section 7, that the problems FixMask, FixMasks$^+$ with $M = 1$ and MinFixMask$_{\text{BOUND}}$, MinFixMask$_{\text{BOUND}}$$^+$ with $T_{\max} \in \{1, N\}$ can be solved in polynomial time. Finally, in Section 8, we present an experimental evaluation of MinFixMask$_{\text{BOUND}}$ for $w = 1$, comparing a MILP-based approach with a greedy heuristic on publicly available CT data.

3. Hardness of minfixmasks problems

First of all, we notice that the decision versions of both problems are the same.

Problem 3.1 (MinFixMasks decision version). Given three non-negative integers K, *diff*, $T_{\max}$ and a string of integers $\hat{D} = \hat{d}_1 \ldots \hat{d}_N$ (with N being a multiple of K) find a treatment plan $((P^1, \tau^1), \ldots, (P^T, \tau^T))$, such that $T \le T_{\max}$ and $\sum_{n=1}^{N} |\hat{d}_n - d_n| \le diff$, where, $d_n = \sum_{t=1}^{T} d_n^t$, $\forall n$, $1 \le n \le N$.

In the following, we prove that the previously mentioned problem is NP-complete. It implies that problems MinFixMask$_{\text{OPT}}$ and MinFixMask$_{\text{BOUND}}$ are also NP-complete. In the following, let us simplify a bit the formulation of the problem since there is no rotation allowed. Given a treatment plan $((P^1, \tau^1), \ldots, (P^T, \tau^T))$, since the paddle configuration is fixed, P^i which is a binary word can be seen as defining if a given paddle is open or not. In other words, P^i determines whether the corresponding dwell time τ^i is applied to a volume under each paddle or not. In the following, we will thus consider that a treatment plan corresponds to a set of dwell time $\mathcal{T}$ and a function that determine for each volume under a paddle if the corresponding dwell time has to be applied. Let us denote $f : \{1,\dots,\frac{N}{w}\} \times \{1,\dots,|\mathcal{T}|\} \rightarrow \{true, false\}$, where

$$f(i,j) = \begin{cases} true, \text{ if the dwell time } \tau^j \text{ is applied to a volume under the } i^{th} \text{ paddle} \\ false, \text{otherwise} \end{cases}$$

We use a reduction from the NP-complete problem MONOTONE 1–3 SAT [16]: Given a boolean formula $\phi = \{c_1, c_2, \dots, c_m\}$ in 3-CNF of m clauses built on a set $V = \{v_1, v_2, \dots, v_n\}$ of n variables such that every clause contains only unnegated literals, does there exist a truth assignment on V satisfying ϕ such that each clause is satisfied by exactly one of its three literals.

Given an instance (ϕ, V) of the MONOTONE 1–3 SAT problem, we build an instance of our problem as follows. Let q_i be an integer value computed using the following recurrence formula:

$$q_i = \begin{cases} w \cdot |V| + 1, i = 1 \\ w \cdot |V| + 1 + (w \cdot |V| + 1 + q_{i-1}) \cdot i, 2 \leq i \leq |V| \end{cases}$$

For ease of notation, in the following sequences, a^b will denote the concatenation of b occurrences of element a. For each variable $v_i \in V$, we build: $V_i^1 = q_i^w$ and $V_i^2 = (1+q_i)^w$. For ease, in the following, we will denote $\{V_i^1, V_i^2\}$ as $\mathcal{V}_i$. For each clause $c_i = (v_{a_i}, v_{b_i}, v_{c_i}) \in \phi$, $1 \leq i \leq m$, we build $C_i = (q_{a_i} + q_{b_i} + q_{c_i} + 2)^w$. The sequence $\hat{D}$ is obtained by concatenating in order $V_1^1 V_1^2 V_2^1 V_2^2 \dots V_{|V|}^1 V_{|V|}^2 C_1 C_2 \dots C_{|\phi|}$. Finally, we set $diff = w \cdot |V|$ and $T_{\max} = |V|$.

Lemma 3.1. *Considering any subset of $\mathcal{V} \subseteq \{\mathcal{V}_j, 1 \leq j \leq |V|\}$, the optimal difference relative to $\mathcal{V}$ using at most $|\mathcal{V}|$ masks is $w \cdot |\mathcal{V}|$ by applying $\{\tau^i \in \{q_i, 1+q_i\} | \forall 1 \leq i \leq |V|,\ s.t.\ \mathcal{V}_i \in \mathcal{V}\}$.*

Proof. To clarify, if $\mathcal{V} = \{\mathcal{V}_1, \mathcal{V}_3, \mathcal{V}_4\}$ for example, one can achieve an optimal difference of $w \cdot 3$ using $\{\tau^1, \tau^3, \tau^4\}$.

By induction, using a single mask to minimize any $\mathcal{V}_j$, for any given $1 \leq j \leq |V|$ will induce a $diff \geq w$ since one cannot achieve a better dose difference than picking the median of q_j and $1 + q_j$ which will induce a difference of 0 for one of $\{V_j^1, V_j^2\}$ and w for the other.

Thus, the lemma is true for any set $\{\mathcal{V}_j\}$, $1 \leq j \leq |V|$. Let us assume that the lemma is correct for any subset $\mathcal{V}^\circ \subseteq \{\mathcal{V}_j, 1 \leq j \leq |V|\}$ where $|\mathcal{V}^\circ| = i - 1$ and let us prove its correctness when we considering any subset $\mathcal{V}^\bullet \subseteq \{\mathcal{V}_j, 1 \leq j \leq |V|\}$ where $|\mathcal{V}^\bullet| = i$.

Considering any such subset denoted as $\mathcal{V}^\bullet$, and pick any $\mathcal{V}_k \in \mathcal{V}^\bullet$. Consider now, the subset $\mathcal{V}^\circ = \mathcal{V}^\bullet - \mathcal{V}_k$.

Let us consider that a solution achieves the smallest difference relative to $\mathcal{V}^\circ$ which is 0. Then it has to use all the different τ values (i.e. i of them) to achieve this otherwise it contradicts our induction hypothesis since $|\mathcal{V}^\circ| = i - 1$.

Then, minimizing the difference for $\mathcal{V}_k$ can only been done using the τ values already defined for minimizing $\mathcal{V}^\circ$ - let us note these values $\mathcal{T}$. Let us furthermore note $\mathcal{T}_{<k}$ (resp. $\mathcal{T}_{>k}$), the set of values of $\mathcal{T}$ (resp. **only**) used to minimized a $\mathcal{V}_i$ with $i < k$ (resp. $i > k$).

Let us demonstrate that there is no benefit of using a value of $\mathcal{T}_{>k}$ for minimizing the difference for $\mathcal{V}_k$.

Consider the $\mathcal{V}_x \in \mathcal{V}^\circ$ with x being the largest value. Let us first assume that it not needs to use a $\tau \in \mathcal{T}$ which is not used by any other $\mathcal{V}_y \in \mathcal{V}^\circ - \mathcal{V}_x$. Then, we know that the overall contribution is $\sum \tau \in \mathcal{T} \leq (1 + q_{x-1} + w \cdot |V|) \cdot (x-1)$. Indeed, one should not achieve more than $1 + q_i + w \cdot |V|$ for any $\mathcal{V}_i$ or the overall difference will already be greater to $w \cdot |V|$. Moreover, $\sum \tau \in \mathcal{T} < q_x - (1 + q_{x-1})$ which means that $\mathcal{V}_x$ needs its own contribution that should be greater than $1 + q_{x-1}$. In other words, this contribution would not be used by other $\mathcal{V}_y \in \mathcal{V}^\circ - \mathcal{V}_x$.

This argument may be applied similarly for all the $\mathcal{V}_x \in \mathcal{V}^\circ$ considering them from the largest x value to the smallest.

Let us now prove that even using all the values of $\mathcal{T}_{<k}$ to minimize the difference for $\mathcal{V}_k$ is not enough. Any $\tau \in \mathcal{T}_{<k}$ are lower than $(1 + q_{k-1})$. By construction, the sum of all values in $\mathcal{T}_{<k}$ is lower than or equal to $(1 + q_{k-1}) \cdot (k-1)$. This means that applying all the values of $\mathcal{T}_{<k}$ to $\mathcal{V}_k$ will still induce a difference higher than $w \cdot |V|$.

Instead, if one uses only $i - 1$ masks for minimizing the difference for V° one can induce a difference of $w(i-1)$ for V°, and one can use the extra τ value for $\mathcal{V}_k$ to get a local difference of w and achieve a global difference of $w \cdot i$.

□

Using the reduction defined above, we are now ready to show the main theorem of this section.

Theorem 3.1. *The MINFIXMASKS decision version problem is NP-complete for any w.*

Proof. Clearly, the problem is contained in NP.

(⇒) Given an assignment satisfying ϕ such that each clause is satisfied by exactly one of its literals, we define the following set of doses $\{\tau^i | \tau^i = q_i if v_{istrueintheassignment} \tau^i = 1 + q_i otherwise\}$.

Let us consider the treatment plan where τ^i is applied to $\mathcal{V}_i$ and to any C_j such that $v_i \in c_j$ in ϕ, for all $1 \leq i \leq |V|$.

By definition of a truth assignment, the corresponding treatment applied either q_i or $1 + q_i$ to each $\mathcal{V}_i$. As stated in Lemma 3.1, this leads to an overall difference of $w \cdot |V|$ for $\mathcal{V} = \{\mathcal{V}_j, 1 \leq j \leq |V|\}$. Moreover, any C_j corresponding to a clause (v_a, v_b, v_c) receives a total dwell time of $q_a + q_b + q_c + 2$, since by hypothesis exactly one of $\{v_a, v_b, v_c\}$ is *true* in our assignment: that is either $q_a + (1 + q_b) + (1 + q_c)$ or $(1 + q_a) + q_b + (1 + q_c)$ or $(1 + q_a) + (1 + q_b) + q_c$. Thus, the difference for all C_j, $1 \leq j \leq |\phi|$, is null. Thus, the corresponding treatment is optimal.

(⇐) Given a treatment as a set of $|V|$ doses $\mathcal{T}$ inducing a total difference equal to $w \cdot |V|$. By Lemma 3.1, for $1 \leq i \leq |V|$, $\mathcal{T}$ contains either q_i or $1 + q_i$.

We define the following truth assignment : set $v_i = true$ if $q_i \in \mathcal{T}$; $v_i = false$ otherwise, for all $1 \leq i \leq |V|$.

Let us prove that the corresponding assignment is satisfying C.

First of all, by definition each variable is either *true* or *false*. Moreover, each clause c_j is satisfied since in order to achieve a total difference of $w \cdot |V|$, any C_j, $1 \leq j \leq |\phi|$, needs to induce a null difference. In order to get such null difference, any C_j corresponding to a clause (v_a, v_b, v_c) needs to receive a total dwell time of $q_a + q_b + q_c + 2$. This can only be satisfied using $q_a + (1 + q_b) + (1 + q_c)$ or $(1 + q_a) + q_b + (1 + q_c)$ or $(1 + q_a) + (1 + q_b) + q_c$, which induce by hypothesis that exactly one of $\{v_a, v_b, v_c\}$ is *true* in our assignment. □

In the sections 4 and 5 we focus on the case $w = 1$ because the cases with $w > 1$ are equivalent to it. Since we cannot have overdose, in a sector we can apply at most the minimum value among the values corresponding to that sector in the $\hat{D}$.

4. FixMasks$^+$

4.1. NP-Complete

In this subsection, we show that $FixMask^+$ is NP-complete by a reduction from MIS-3 (Maximum Independent Set in subcubic graphs) [17]. Let $G = (V, E)$ be an instance of MIS-3. We construct an instance of FixMasks$^+$ in the following way:

- $w = 1$
- $\hat{D} = \underbrace{1 \ldots 1}_{|V|+|E|}$
- for every vertex $v \in V$, we define a mask (vertex-voxel)

$$P_i^v = \begin{cases} 1, \text{ if } i = v, \\ 1, \text{ if } i = |V| + j \text{ and } e_j \text{ is incident with } v, \\ 0, \text{otherwise} \end{cases}$$

- for every edge e_j, we define a mask (edge-voxel)

$$P_i^{|V|+j} = \begin{cases} 1, \text{ if } i = |V| + j, \\ 0, \text{otherwise} \end{cases}$$

Let P be a treatment plan. We define $E(P) = \sum_{n=1}^{N} (\hat{d}_n - d_n)$, where $d_n = \sum_{t=1}^{|P|} d_n^t$. Given the construction mentioned above, we can show that G contains an independent set of dimension k if and only if there exists a treatment plan P such that $E(P) = |V| - k$. Before presenting the technical details, we show how the reduction works using an example.

Example 4.1. Let $G = (V, E)$, $V = \{1, 2, 3\}$, $E = \{(1, 2), (1, 3), (2, 3)\}$. Then, we obtain the following instance of FixMasks$^+$: $\hat{D} = 111111$, $(100110, 010101, 001011, 000100, 000010, 000001)$. The maximum independent set in G is 1 and the minimum error in the instance of FixMasks$^+$ is 2 by selecting the following treatment plan: $(100110, 1), (000001, 1)$.

Lemma 4.1. *G contains an independent set of size* $\geq k \iff \exists$ *a treatment plan* P *such that* $E(P) \leq |V| - k$.

Proof. ($\Rightarrow$). Let I be an independent set of G with $|I| \geq k$. Then we can construct the treatment plan $P = \underbrace{\{P^v \mid v \in I\}}_{A} \cup \underbrace{\{P^{|V|+j} \mid e_j = (x_j, y_j), x_j, y_j \notin I\}}_{B}$. The treatment plan P is feasible because any two masks in A cannot overlap the same edge-voxel (I is an independent set). Furthermore, by adding the masks in B we cover all the edges-voxel. Thus, $E(P) = |V| - |I| \leq |V| - k$.

($\Leftarrow$). Let S be a treatment plan such that $E(S) \leq |V| - k$. Construct a plan P by starting from S and, while there exists an edge-voxel j not covered by any mask in the current plan add to P the mask $P^{|V|+j}$. We apply this process until we cover all edges-voxel. A vertex-voxel can only be covered by its own mask P^v. Let $I = \{v \mid P^v \in P\}$. $E(P) = |V| - |I| \implies |I| = |V| - E(P) \geq |V| + k - |V| = k$. In addition, if there were two adjacent vertices $x, y \in I$ then the edge-voxel $|V| + j$ of the edge $e_j = (x, y)$ would receive dosage 2 which is impossible. Thus, I is an independent set of dimension at least k. □

Theorem 4.1. *FixMasks$^+$ problem is NP-complete.*

Proof. Clearly, the problem is contained in NP. The proof for the NP-hardness follows immediately from Lemma 4.1. □

4.2. APX-Hard

In this subsection we provide an L-reduction from MIS-3 [18] problem thus showing that the FixMasks$^+$ problem is APX-hard.

Theorem 4.2. *FixMasks$^+$ problem is APX-hard.*

Proof. We define $f : \{G = (V, E) \mid \Delta(G) \le 3\} \to \{(\hat{D}, P) \mid \hat{D} \in \{1\}^{|V|+|E|}, P \subseteq \{0,1\}^{|V|+|E|}\}$, $f(G) = ((\underbrace{1, \dots, 1}_{|V|+|E|}), P = (P^1, \dots, P^{|V|+|E|}))$,

where

$$P_i^v = \begin{cases} 1, & \text{if } i = v, \\ 1, & \text{if } i = |V| + j \text{ and } e_j \text{ is incident with } v, \\ 0, & \text{otherwise} \end{cases} \qquad P_i^{|V|+j} = \begin{cases} 1, & \text{if } i = |V| + j, \\ 0, & \text{otherwise} \end{cases}$$

and $g : \{\text{feasible plans } S \text{ for} f(G)\} \to \{I \subseteq V\}, g(S) = \{v \mid P^v \in S\}$. Both f and g can be computed in polynomial time. From Lemma 4.1 we have $OPT_{FixMasks^+} = |V| - OPT_{MIS-3}$. Then, $|\, OPT_{MIS-3}(G) - |g(S)|\, | = |\, OPT_{MIS-3} - (|V| - E(S))\, | = |\, |V| - OPT_{MIS-3}(G) - E(S)\, | \le |\, OPT_{FixMasks^+}(f(G)) - E(S)\, | \implies \alpha = 1$. Using the Caro-Wei inequality [19,20] and the fact that the graph is subcubic, we have $OPT_{MIS-3} \ge \frac{|V|}{4}$. Then, $OPT_{FixMasks^+} \le |V| - \frac{|V|}{4} \le 3 \cdot \frac{|V|}{4} \le 3 \cdot OPT_{MIS-3} \implies \beta = 3$. □

5. DFixMasks$^+$

Corollary 5.1. *DFixMasks$^+$ is NP-complete.*

Proof. Clearly, the problem is contained in NP. The proof for the NP-hardness immediately follows from $OPT_{FixMasks^+} = \sum_{i=1}^{N} \hat{d}_i$ - $OPT_{DFixMasks^+}$ and Theorem 4.1. □

5.1. ILP And FPT

In this subsection we provide an ILP formulation for the DFixMasks$^+$ problem.

$$\begin{aligned} \max \quad & \sum_{j=1}^{M} w_j x_j \\ \text{s.t.} \quad & \sum_{j=1}^{K} P_i^j x_j \le d_i, \qquad i = 1, \dots, N. \\ & x_j \in \mathbb{Z}_{\ge 0}, j = 1, \dots, M. \\ & w_j = \sum_{i=1}^{N} P_i^j, \qquad j = 1, \dots, M. \end{aligned}$$

Theorem 5.1. *DFixMasks$^+$ and FixMasks$^+$ can be solved in $2^{O(M^3)} \cdot (M \cdot \log(\hat{d}_{\max}))^{O(1)}$.*

Proof. This follows directly from Lenstra's theorem on Integer Linear Programming in fixed dimension [21]. Lenstra's result states that Integer Linear Programs with n variables can be solved in time $2^{O(n^3)} \cdot (m \cdot \log V)^{O(1)}$, where m is the number of constraints and V bounds the size of the input coefficients. In our case, both DFixMasks$^+$ and FixMasks$^+$ can be formulated as integer linear programs with the number of variables bounded by M, and the coefficients bounded by $\hat{d}_{\max}$. □

5.2. $\frac{1}{M}$-Approximation

In this subsection we present an $\frac{1}{M}$-approximation for the DFixMasks$^+$ problem. The algorithm works as follows: for each mask we compute the maximum dosage we can apply without overdose and choose the mask that offers us the highest dosage. We describe it formally in Algorithm 1.

Theorem 5.2. *Algorithm 1 computes a $\frac{1}{M}$-approximation for problem DFixMasks$^+$ in $O(M \cdot N)$.*

Proof. Let OPT be the an optimal solution for $DMinFixMasks^+$, $P = ((P^1, \tau^1), \dots, (P^T, \tau^t))$ a treatment plan corresponding to OPT and ALG the result of the Algorithm 1. Then, $\exists i, 1 \le i \le T$ such that $\tau^i \ge \frac{OPT}{T} \ge \frac{OPT}{M}$. Since $ALG \ge \tau^i, \forall i, 1 \le i \le T$ then $ALG \ge \frac{OPT}{M}$. The running time is $O(M \cdot N)$ because for every mask we compute in $O(N)$ the maximum dosage that can be applied. □

5.3. $\frac{k}{M}$-Approximation

In this subsection we present an exponential $\frac{k}{M}$-approximation algorithm. The algorithm is based on generating all subsets of masks of size k and solving the ILP formulation of the DFixMasks$^+$ problem for each subset and keeping the best solution. The Algorithm is described formally in Algorithm 2.

Theorem 5.3. *Algorithm 2 computes an $\frac{k}{M}$-approximation for DFixMasks$^+$ in $O(\binom{M}{k} \cdot 2^{O(k^3)} \cdot (k \cdot \log(\hat{d}_{\max}))^{O(1)})$.*

Algorithm 1 $\frac{1}{M}$-approximation for DFixMasks$^+$.

```
sol = −∞
for i ← 1 to M do
    val = −∞
    for j ← 1 to N do
        if P_j^i = 1 then
            if d_j > val then
                val = d_j
            end if
        end if
    end for
    if val > sol then
        sol = val
    end if
end for
return sol
```

Algorithm 2 $\frac{k}{M}$-approximation for DFixMasks$^+$.

```
for all subset S ⊆ {1, …, M} with |S| = k do
    Solve the ILP restricted to S:
        max Σ_{j∈S} w_j x_j   s.t.  Σ_{j∈S} P_i^j x_j ≤ d̂_i,  x_j ∈ ℤ_{≥0}
    Let val be its optimum value
    ALG ← max(ALG, val)
end for
return ALG
```

Proof. Let $(P^1,\tau^1),\dots,(P^T,\tau^T)$ be an optimal solution and $OPT = \tau^1 + \cdots + \tau^T$. We denote by $S_i = \sum_{j=1}^{i} \tau^i$, $\forall i$, $1 \le i \le T$. We show that $S_k \ge \frac{k}{M} \cdot OPT$, $\forall k$, $1 \le k \le T$. We assume without loss of generality that $\tau^1 \ge \tau^2 \ge \cdots \ge \tau^T$. Then $\tau^{k+1} + \cdots + \tau^T \le (T-k)\cdot\tau^k$. Moreover, because $\tau^k \le \frac{S_k}{k}$ we have $\tau^{k+1} + \cdots + \tau^T \le (T-k)\cdot\tau^k \le (T-k)\cdot\frac{S_k}{k}$. Thus, $OPT = S_k + \tau^{k+1} + \cdots + \tau^T \le S_k + (T-k)\cdot\frac{S_k}{k} \le \frac{T}{k}\cdot S_k$. Then, $S_k \ge \frac{k}{T}\cdot OPT$. Since $T \le M$, we have $S_k \ge \frac{k}{M}\cdot OPT$. The running time results directly from Theorem 5.1. □

5.4. $(1-\frac{1}{e})$-*Approximation*

In this subsection we present a randomized polynomial-time $(1-\frac{1}{e})$-approximation algorithm for the DFixMasks$^+$ problem (assuming the per-voxel dose caps $\hat{d}_i$ are polynomially bounded) by instantiating the Continuous-Greedy from [22] and Contention-Resolution framework of Chekuri, Vondrák and Zenklusen [23]. We begin by relaxing our discrete dose-packing objective to its multilinear extension $F(x)$ over the packing polytope defined by the per-voxel caps, and run the continuous-greedy process to obtain a fractional solution x^* with $F(x^*) \ge (1-\frac{1}{e}) \max_{Ax\le\hat{d}} F(x)$. We then round x^* to an integral plan by sampling each shot independently with probability x_e^* and applying the ordered contention-resolution scheme to prune any cap-violating selections, incurring at most another $(1-\frac{1}{e})$ factor loss. The result is a feasible treatment plan whose expected total dose is at least $(1-\frac{1}{e})$ times the optimum. The formal pseudocode is given in Algorithm 3.

Let $U = \{e_{j,l} = (P^j, 1) \mid i = 1,\dots M, l = 1,\dots,\hat{d}_{\max}\}$, where $\hat{d}_{\max} = \max_{1\le i\le N} \hat{d}_i$. For any subset $S \subseteq U$ we define the delivered dose at voxel i by $d_i(S) = \sum_{e_{j,l}\in S} P_i^j$. The total delivered dose (objective function) is $f(S) = \sum_{i=1}^{N} d_i(S)$. We allow only dose that not overdose any voxel: $d_i(S) \le \hat{d}_i$, $\forall i$, $1 \le i \le N$. We only consider S for which $d_i(S) \le \hat{d}_i$, $\forall i$, $1 \le i \le N$. Denote the feasible family by $\mathcal{F} = \{S \subseteq U \mid d_i(S) \le \hat{d}_i, 1 \le i \le N\}$. For every $e \in U$ we introduce a fractional variable $x_e \in [0,1]$ and the constraints voxel i receives at most $\hat{d}_i$ units become: $\sum_{e\in U} P_i^e \cdot x_e \le \hat{d}_i$. This defines a packing polytope $P = \{x \in [0,1]^U \mid A \cdot x \le \hat{d}\}$, where $A_{i,e} = P_i^e$.

Lemma 5.1. *The objective function f is nonnegative, monotone, and submodular.*

Proof. Let $S \subseteq U$. Since $d_i(S) \ge 0$, $\forall i$, $1 \le i \le N$ then $\sum_{i=1}^{N} f(S) \ge 0$, $\forall S \subseteq \mathcal{F}$ (f is nonegative). Let $S, T \subseteq U$ such that $S \subseteq T$. Then $d_i(S) \le d_i(S) \le d_i(T)$, $\forall i$, $1 \le i \le N$ (adding more elements can only increase its delivered dose). Summing over all voxels we have

Algorithm 3 Continuous-Greedy + Contention-Resolution for DFIXMASKS$^+$.

$x \leftarrow 0 \in \mathbb{R}^{|U|}$
for $t = 0$ to 1 **step** δ **do**
  Estimate gradient $\nabla F(x)$ by sampling
  Choose direction $v \in \arg\max_{v' \in P} \nabla F(x) \cdot v'$
  Update $x \leftarrow x + \delta\, v$
end for
$x^* \leftarrow x$
Sample a random set $R \subseteq U$ by including each e independently with probability x^*_e
Apply the ordered contention-resolution scheme to prune R into a feasible $S \subseteq R$
return S

$f(S) = \sum_{i=1}^{N} d_i(S) \le \sum_{i=1}^{N} d_i(T) = f(T)$ (f is monotone). Let $S, T \subseteq U$ such that $S \subseteq T$ and $e \notin T$. Let $a = d_i(S)$, $b = d_i(T)$, $p = P_i^e \in \{0, 1\}$. Adding e to S raises $d_i(S)$ from a to $a + p$, and to T from b to $b + p$. The marginal gain per-voxel is $\Delta_i(S) = (a + p) - a = p = (b + p) - b = \Delta_i(T)$, $\forall i\ 1 \le i \le N$. Summing over all voxels we obtain $f(S \cup \{e\}) - f(S) = \sum_{i=1}^{N} \Delta_i(S) = \sum_{i=1}^{N} \Delta_i(T) = f(T \cup \{e\}) - f(T)$ (f is modular which implies that it is also submodular). □

Theorem 5.4 (Continuous Greedy Approximation [22]). *Let $f : 2^U \to \mathbb{R}_{\ge 0}$ be a monotone submodular function and F its multilinear extension. Let $P \subseteq [0, 1]^U$ be a down-closed, solvable packing polytope. Then the continuous greedy algorithm returns a fractional solution $x^* \in P$ such that $F(x^*) \ge \left(1 - \frac{1}{e}\right) \cdot \max_{x \in P} F(x)$.*

Theorem 5.5 (Ordered Contention-Resolution Rounding [24]). *Let f be a monotone submodular function and $x^* \in P$ a fractional solution obtained by the continuous greedy algorithm. Then there exists an ordered contention-resolution scheme that rounds x^* to a feasible set $S \subseteq U$ such that $\mathbb{E}[f(S)] \ge \left(1 - \frac{1}{e}\right) \cdot F(x^*)$.*

Corollary 5.2. *Combining the above results, the continuous greedy algorithm followed by ordered contention-resolution yields a randomized polynomial-time algorithm for DFIXMASKS$^+$ that outputs a feasible set $S \in \mathcal{F}$ such that $\mathbb{E}[f(S)] \ge \left(1 - \frac{1}{e}\right) \cdot \max_{S' \in \mathcal{F}} f(S')$.*

6. Quasi-polynomial algorithms for *MinFixMasks* problems

In this section, we present exact algorithms for all variants of the *MinFixMasks* problem. The presented algorithms run in quasi-polynomial time if the values of the prescribed patient doses are bounded by a polynomial in the number of prescribed doses. As a by-product we show that the problems MINFIXMASKS$_{\text{OPT}}$ and MINFIXMASKS$_{\text{OPT}}{}^+$ can be approximated in polynomial-time within a factor of $\log \hat{d}_{\max}$ of the optimum where $\hat{d}_{\max}$ is the maximum prescribed dose to a subvolume of the patient, i.e., $\hat{d}_{\max} = \max_{\hat{d}_i \in \hat{D}} \hat{d}_i$.

Lemma 6.1. *For any instance of MINFIXMASK$_{\text{OPT}}$, MINFIXMASKS$_{\text{OPT}}{}^+$, MINFIXMASK$_{\text{BOUND}}$, MINFIXMASK$_{\text{BOUND}}{}^+$ there is an optimal solution $((P^1, \tau^1), \dots, (P^T, \tau^T))$ with $\tau_i \neq \tau_j$, $\forall i, j$, $1 \le i < j \le T$.*

Proof. Let $P = ((P^1, \tau^1), \dots, (P^T, \tau^T))$ be an optimal solution of an instance I of any of the mentioned variants of the *MinFixMasks* problem. We show that we can transform P into an equivalent treatment plan that does not use any dwell time more than once. Let $i, j \in [T]$ such that $i \neq j$ and $\tau^i = \tau^j$. We construct (P^i_*, τ^i_*) and (P^j_*, τ^j_*) such that $\tau^i_* = 2 \cdot \tau^j$, binary string P^i_* is obtained by the XOR of binary strings P^i, P^j and binary string P^j_* is obtained by the AND of binary strings P^i, P^j. Then, the treatment plan obtained from P by replacing (P^i, τ^i) with (P^i_*, τ^i_*) and (P^j, τ^j) with (P^j_*, τ^j_*) is also an optimal solution of I. Moreover, by applying this procedure iteratively we obtain an optimal solution of I such that all dwell times are pairwise distinct. □

Definition 6.1. Let S be a set of dwell times. We say that S is complete $\iff \exists S' \subseteq S$ such that $\forall i$, $1 \le i \le \hat{d}_{\max}$, $i = \sum_{x \in S''} x$ with $S'' \subseteq S'$.

Lemma 6.2. *Let S be a set of dwell times. Then an S-restricted treatment plan minimizing $\sum_{n=1}^{N} |\hat{d}_n - d_n|$ can be found in time $O((\hat{d}_{\max})^2 \cdot |S| + K \cdot W + K \cdot \hat{d}_{\max})$. Moreover, the same applies to an S-restricted treatment satisfying the additional constraint that $\hat{d}_n - d_n \ge 0$, $\forall n$, $1 \le n \le N$. Finally, if S is complete then the S-restricted treatment plans returned by the above algorithms are optimal among all (not necessarily S-restricted) treatment plans.*

Proof. Because every treatment plan has to apply the same dose to every patient volume under the same paddle, we obtain that $d_{(i-1)\cdot w+1} = d_{i \cdot w+1} = \dots = d_{i \cdot w}$ $\forall i$, $1 \le i \le K$. Hence, finding a treatment plan minimizing $\sum_{n=1}^{N} |\hat{d}_n - d_n|$ is equivalent to finding K natural numbers $\overline{d}_1, \ldots, \overline{d}_K$ minimizing $|\hat{d}_n - \overline{d}_{\lceil \frac{n}{w} \rceil}|$ such that each number $\overline{d}_i$ can be realized by a treatment plan. Since we are only considering S-restricted treatment plans, it holds that the set of realizable numbers (for any such treatment plan) is equal to the set of numbers that can be written as a sum of a subset of S. Moreover, because we can choose the masks arbitrarily it holds that for every K numbers $\overline{d}_1, \ldots, \overline{d}_K$, which can be written as the sum of subsets of S, there is an S-restricted treatment plan realizing these numbers. To see this let $\overline{d}_1, \ldots, \overline{d}_K$ be any numbers such that for each $\overline{d}_i$ there is a subset $S(\overline{d}_i)$ of S with $\overline{d}_i = \sum_{s \in S(\overline{d}_i)} s$, then $(P^1, s_1), \ldots P^T, s_{|S|})$, where $s_1, \ldots, s_{|S|}$ is an arbitrary ordering of the numbers in S and the i-th bit of the binary string P^j is 1 if and only if $s_j \in S(\overline{d})_i$ is an S-restricted treatment plan realizing the numbers $\overline{d}_1, \ldots, \overline{d}_K$. It follows that given the set of allowed dwell times S, we can minimize each of the sums $\sum_{n=(i-1)\cdot w+1}^{i\cdot w} |\hat{d}_n - \overline{d}_{\lceil \frac{n}{w} \rceil}|$ separately for every i with $1 \leq i \leq K$. Because of Lemma 7.1, the smaller $\sum_{n=(i-1)\cdot w+1}^{i\cdot w} |\hat{d}_n - \overline{d}_{\lceil \frac{n}{w} \rceil}|$ is, the closer $\overline{d}_{\lceil \frac{n}{w} \rceil}$ is to a median of the sequence $\hat{d}_{(i-1)\cdot w+1}, \ldots \hat{d}_{i\cdot w}$. Moreover, if we consider the case where we have the additional constraint that $\hat{d}_n - d_n \geq 0\ \forall n,\ 1 \leq n \leq N$, then the optimal value for $\overline{d}_{\lceil \frac{n}{w} \rceil}$ is the one that is closest to the minimum of the sequence $\hat{d}_{(i-1)\cdot w+1}, \ldots \hat{d}_{i\cdot w}$. These considerations naturally lead to the following algorithm to find an S-restricted treatment plan minimizing $\sum_{n=1}^{N} |\hat{d}_n - d_n|$. The algorithm first computes a table that contains for every number i between 1 and $\hat{d}_{\max}$ either a subset S' of S such that $i = \sum_{s \in S'} s$ or *Nil* if no such subset S' exists for i. Using the standard text-book algorithm for the Subset Sum problem [25] (running in $O(s \cdot N)$ time, where s is the sum that needs to be obtained and N is the size of the set of integers), this can be achieved in $O(\hat{d}^2_{\max}|S|)$ time. In the case that we have no additional constraint the algorithm then computes the (at most two medians) m_i^1 and m_i^2 of the sequence $\hat{d}_{(i-1)\cdot w+1}, \ldots \hat{d}_{i\cdot w}\ \forall i\ 1 \leq i \leq K$. This can be achieved in time at most $O(K \cdot w)$ [26] time. In the case that we have the additional constraint $\hat{d}_n - d_n \geq 0\ \forall n,\ 1 \leq n \leq K$, the algorithm computes the minimum $\min_i$ for each of these sequences in $O(K \cdot w)$ time. Finally, the algorithm computes for every $1 \leq i \leq K$, the number $\overline{d}_i$ as the number between 1 and $\hat{d}_{\max}$ that has a non-nil entry in the table and either: (1) is closest to one of the medians m_i^1 and m_i^2 (in the case with no additional constraints) or (2) is closest to the minimum $\min_i$ (in the case that overdoses are not allowed). Using the table and the medians m_i^1 and m_i^2 respectively the minima $\min_i$ this can be achieved in time $O(K\hat{d}_{\max})$. Considering an arbitrary ordering $s_1, \ldots, s_{|S|}$ of the numbers in S and denoting by $S(\overline{d}_i)$ the subset of S contained in the table for the number $\overline{d}_i$, the algorithm then outputs the treatment plan $\left((P^1, s_1), \ldots, (P^T, s_{|S|})\right)$ where $s_1, \ldots, s_{|S|}$ is an arbitrary ordering of the numbers in S and the i-th bit of the binary string P^j is 1 if and only if $s_j \in S(\overline{d}_i)$. This completes the description of the algorithm. The running time of the algorithm is at $O((\hat{d}_{\max})^2 \cdot |S| + K \cdot W + K \cdot \hat{d}_{\max})$. Finally, if S is complete then the S-restricted treatment plan returned by the algorithm minimizes $\sum_{n=1}^{N} |\hat{d}_n - d_n|$ under all possible treatment plans. The same applies when considering the additional constraint $\hat{d}_n - d_n \geq 0,\ \forall\ n,\ 1 \leq n \leq N$. □

Lemma 6.3. *There is a treatment plan minimizing* $\sum_{n=1}^{N} |\hat{d}_n - d_n|$ *using at most* $\lfloor \log \hat{d}_{\max} \rfloor + 1$ *steps. Moreover, such a treatment plan can be found in polynomial time. The same holds for a treatment plan minimizing* $\sum_{n=1}^{N} |\hat{d}_n - d_n|$ *under the additional constraint that* $\hat{d}_n - d_n$, $\forall n$ *with* $1 \leq n \leq N$.

Proof. Because the set $S = \{2^i \mid 0 \leq i \leq \lfloor \log \hat{d}_{\max} \rfloor\}$ is complete and has size $\lfloor \hat{d}_{\max} \rfloor + 1$, this follows immediately from Lemma 6.2. □

Because any non-trivial instance of MinFixMask$_{\text{OPT}}$ and MinFixMasks$_{\text{OPT}}{}^+$ require at least one step, we obtain the following corollary from the above lemma.

Corollary 6.1. *MinFixMask$_{OPT}$ and MinFixMasks$_{OPT}{}^+$ can be approximated in polynomial time within a factor of* $\log \hat{d}_{\max}$ *of the optimum.*

Theorem 6.1. *MinFixMask$_{OPT}$, MinFixMasks$_{OPT}{}^+$, MinFixMask$_{BOUND}$, and MinFixMask$_{BOUND}{}^+$ can be solved in time* $O(\hat{d}_{\max}^{\lfloor \log \hat{d}_{\max} \rfloor + 1}(\hat{d}^2_{\max}(\lfloor \log \hat{d}_{\max} \rfloor + 1) + K \cdot w + K \cdot \hat{d}_{\max}))$.

Proof. The algorithm goes over all sets S containing at most $\lfloor \log \hat{d}_{\max} \rfloor + 1$ (respectively at most $\min\{\lfloor \log \hat{d}_{\max} \rfloor + 1, T_{\max}\}$ in the case of MinFixMask$_{\text{BOUND}}$ and MinFixMask$_{\text{BOUND}}{}^+$) dwell times between 1 and $\hat{d}_{\max}$. For every such set S, the algorithm then uses Lemma 6.2 to compute the optimal (the meaning of optimal here depends on the considered problem) S-restricted treatment plan. Finally, in the case of MinFixMask$_{\text{BOUND}}$ and MinFixMask$_{\text{BOUND}}{}^+$ the algorithm returns a shortest treatment plan satisfying $\sum_{n=1}^{N} |\hat{d}_n - d_n| \leq \text{diff}$ and in the case of MinFixMask$_{\text{BOUND}}$ and MinFixMask$_{\text{BOUND}}{}^+$ returns a treatment plan minimizing $\sum_{n=1}^{N} |\hat{d}_n - d_n|$ found for any of the considered sets S. The stated running of the algorithm follows because there are at most $\hat{d}_{\max}^{\lfloor \log \hat{d}_{\max} \rfloor + 1}$ such sets S and because of Lemma 6.2 for each set S, we require time at most $O(\hat{d}^2_{\max}|S| + K \cdot w + K \cdot \hat{d}_{\max})$. The correctness of the algorithm follows from Lemmas 6.1, 6.2 and 6.3. □

Corollary 6.2. *MinFixMask$_{OPT}$, MinFixMasks$_{OPT}{}^+$, MinFixMask$_{BOUND}$, and MinFixMask$_{BOUND}{}^+$ can be solved in quasi-polynomial time if* $\hat{d}_{\max}$ *is bounded by a polynomial in the number of prescribed doses.*

7. Polynomial algorithms

In this section, we show that FIXMASKS, FIXMASKS$^+$ with $M = 1$ and MINFIXMASK$_{\text{BOUND}}$, MINFIXMASK$_{\text{BOUND}}$$^+$ with $T_{\max} \in \{1, N\}$ are solvable in polynomial time.

Theorem 7.1. *FIXMASKS$^+$ with $M = 1$ and MINFIXMASK$_{BOUND}$$^+$ with $T_{\max} = 1$ can be solved in $O(N)$ time.*

Proof. Because we are not allowed to apply overdoses, we obtain that the maximum and also the optimum irradiation time is equal to the minimum of all prescribed doses $\hat{d}_j$ of $\hat{D}$ for which the corresponding paddle is open (in the case of the problem MINFIXMASK$_{\text{BOUND}}$$^+$, the minimum of all values of $\hat{d}$). The mask corresponding to the optimal solution for problem MINFIXMASK$_{\text{BOUND}}$$^+$ contains 1 at the position where minimum is located and 0 at the other positions. Since the minimum of these doses can be obtained in linear time, the result follows. □

The main observation required to show that FIXMASKS with $M = 1$ and MINFIXMASK$_{\text{BOUND}}$ with $T_{\max} = 1$ can also be solved in polynomial time is given in the following lemma, which can be considered folklore and is stated here only for the convenience of the reader.

Lemma 7.1. *For a sequence S of integer numbers and a integer number x, consider the function $f(x) = \sum_{i=1}^{|S|} |x - S[i]|$. Then $f(x)$ has a unique minimum, which is only reached by any number in between the at most two medians of S.*

The above lemma implies that an optimum dwell time for an instance of $FixMask$ is a median of the subsequence of $\hat{D}$ containing all prescribed doses for which the paddles are open (in the case of the problem MINFIXMASK$_{\text{BOUND}}$, the median of all values of $\hat{d}$). .

Theorem 7.2. *FIXMASKS with $M = 1$ and MINFIXMASK$_{BOUND}$ with $T_{\max} = 1$ can be solved in $O(N)$ time.*

Proof. Because of Lemma 7.1 the best possible value that we can achieve for $\sum_{n=1}^{N} |\hat{d}_n - d_n|$ is obtained by setting the dwell time τ to any median of the subsequence of $\hat{D}$ containing only the prescribed doses for which the paddles are open (in the case of the problem MINFIXMASK$_{\text{BOUND}}$, the median of all values of $\hat{d}$). The mask corresponding to the optimal solution for problem MINFIXMASK$_{\text{BOUND}}$ contains 1 at the position where median is located and 0 at the other positions. It is known [26] that a median of N numbers can be found in linear time. □

Theorem 7.3. *MINFIXMASK$_{BOUND}$ and MINFIXMASK$_{BOUND}$$^+$ with $T_{\max} = N$ can be solved in $O(N^2)$ time.*

Proof. We construct a treatment plan of dimension N with the following structure: each value τ^i corresponds to a mask P^i containing value 1 at position i and 0 at the other positions. This plan achieves absolute difference 0 which is optimal. □

8. Experiments

In this section we present our experimental results. The experiments focus on the $MinFixMask_{BOUND}$ problem with $w = 1$ and we compare a MILP algorithm with a greedy one. We start by defining the problem as a mixed integer linear program. Let N be the number of voxels, $\hat{d}_n$ the target dose at voxel n, $T_{\max}$ the upper bound on the number of masks and $M = \max_{1 \le i \le N} \hat{d}_i$. The program definition is as follows:

$$
\begin{aligned}
\min_{P,\tau,d,\delta} \quad & \sum_{n=1}^{N} \delta_n \\
\text{s.t.} \quad & P_{n,t} \in \{0,1\}, && n = 1,\dots,N,\ t = 1,\dots,T_{max}. \\
& \tau_t \in \{0,1,\dots,M\}, && t = 1,\dots,T_{max}. \\
& d_{n,t} \ge 0, \quad \delta_n \ge 0, && n = 1,\dots,N,\ t = 1,\dots,T_{max}. \\
& P_{n,t} = 1 \implies d_{n,t} = \tau_t, && n = 1,\dots,N,\ t = 1,\dots,T_{max}. \\
& P_{n,t} = 0 \implies d_{n,t} = 0, && n = 1,\dots,N,\ t = 1,\dots,T_{max}. \\
& \delta_n \ge \hat{d}_n - \sum_{t=1}^{T_{max}} d_{n,t}, && n = 1,\dots,N. \\
& \delta_n \ge \sum_{t=1}^{T_{max}} d_{n,t} - \hat{d}_n, && n = 1,\dots,N.
\end{aligned}
$$

Next, we briefly explain the MILP formulation. The goal is to choose up to $T_{\max}$ masks together with their (integer) dwell times so as to minimize the total absolute deviation from the prescribed dose vector $\hat{D}$. For each time step $t \in \{1,\dots,T_{\max}\}$ and voxel $n \in \{1,\dots,N\}$, the binary variable $P_{n,t} \in \{0,1\}$ indicates whether voxel n is exposed (i.e., unshielded) by the mask used at step t. The integer variable $\tau_t \in \{0,1,\dots,M\}$ represents the dwell time (dose intensity) assigned to that mask. We use auxiliary variables $d_{n,t} \ge 0$ to model the dose delivered to voxel n at step t: if $P_{n,t} = 1$ then $d_{n,t} = \tau_t$, and if $P_{n,t} = 0$ then $d_{n,t} = 0$. Hence the total delivered dose at voxel n is $\sum_{t=1}^{T_{\max}} d_{n,t}$. Finally, the nonnegative slack variables δ_n linearize the absolute error at each voxel by enforcing $\delta_n \ge$

$\hat{d}_n - \sum_{t=1}^{T_{\max}} d_{n,t}$ and $\delta_n \geq \sum_{t=1}^{T_{\max}} d_{n,t} - \hat{d}_n$, so that, at optimality, $\delta_n = \left| \hat{d}_n - \sum_{t=1}^{T_{\max}} d_{n,t} \right|$. The objective $\min \sum_{n=1}^{N} \delta_n$ therefore minimizes the total ℓ_1 deviation from the target dose across all voxels. To implement the MILP algorithm we used Gurobi [27].

Algorithm 4 Greedy for MINFIXMASK$_{\text{OPT}}$ with $w = 1$.

τ_list $\leftarrow$ [ ], *mask*_list $\leftarrow$ [ ]
for $t \leftarrow 1$ **to** $T_{\max}$ **do**
  for $i \leftarrow 1$ **to** N **do**
    $r_i \leftarrow \hat{d}_i - d_i$
  end for
  if $\max_{1 \leq i \leq N} r_i = 0$ **then**
    break
  end if
  Let $\left(r_{(1)}, \ldots, r_{(N)}\right)$ be $\{r_n\}$ sorted descending
  $max_{val} \leftarrow 0$
  $max_k \leftarrow 0$
  for $k \leftarrow 1$ **to** N **do**
    $val \leftarrow k \cdot r_{(k)}$
    if $val > max_{val}$ **then**
      $max_{val} \leftarrow val$
      $max_k \leftarrow k$
    end if
  end for
  $\tau \leftarrow r_{(max_k)}$
  Build mask P by $P_i = \begin{cases} 1, & r_i \geq \tau, \\ 0, & \text{otherwise.} \end{cases}$
  for $i \leftarrow 1$ **to** N **do**
    $d_i \leftarrow d_i + P_i \cdot \tau$
  end for
  Append τ to τ_list, append P to *mask*_list
end for
return τ_list, *mask*_list

Next, we briefly explain the greedy algorithm. The main idea of this algorithm is to choose at each step the (mask, dwell time) pair that maximizes the immediate reduction in total absolute error. At each iteration, we compute $r_n = \hat{d}_n - d_n$, where d_n is the delivered dose on voxel n until this step. We sort these N residuals in descending order. We observe that choosing dwell-time r_i then the total absolute error decreases by $i \cdot r_i$. Thus, we iterate from 1 to N and choose as dwell time the value r_k that maximizes the expression $k \cdot r_k$. We set $\tau = r_k$ and binary mask P, where $P_i = 1$ if $r_i \geq \tau$ and 0 otherwise. Therefore, the time complexity is $O(T_{\max} \cdot N \cdot \log N)$. The algorithm is described formally in Algorithm 4.

We used the publicly available DICOM test volumes from the SlicerRtData repository [28]. For each CT slice, the 2D pixel matrix was then flattened into a one-dimensional vector. To focus our optimization on regions of non-zero dose, we discarded any non-positive intensity values, retaining only the positive elements for subsequent MILP and greedy algorithm input. To ensure our slice-by-slice MILP optimization remains clinically practical, we impose a 45 seconds per-slice time limit on the solver. This choice is informed by typical brachytherapy planning workflows, which perform full-case contouring, reconstruction, and optimization in ≈ 175 min over roughly 75 CT slices (≈ 2.3 min/slice) [29]. In experiments, $T_{\max}$ can take the following five values: $log_2 N$, $\sqrt{N}$, $\frac{N}{4}$, $\frac{N}{2}$, $\frac{3 \cdot N}{4}$. The experiments were performed on two data sets from SlicerRtData repository: plastimatch_tiny-rt-study and oncentra-4.1.6-prostate-ct-tilted-plane. In order for the MILP algorithm to obtain results within 45 seconds, we limited the value of N for the oncentra-4.1.6-prostate-ct-tilted-plane dataset to 250. The results obtained on the dataset plastimatch_tiny-rt-study can be seen in Tables 3, 4 and Fig. 3 and for dataset oncentra-4.1.6-prostate-ct-tilted-plane in Tables 5,6 and Fig. 4. Also, the greedy algorithm was tested on all values in the dataset oncentra-4.1.6-prostate-ct-tilted-plane and the results can be seen in Tables 1, 2.

Across both datasets, we observe a consistent quality-time trade-off. For small $T_{\max}$, the MILP baseline often achieves lower error than the greedy heuristic within the available time budget, reflecting the benefit of global optimization when the search space is small enough. As $T_{\max}$ increases, the greedy method remains fast (seconds per slice) and frequently reaches very small (often zero) error once sufficient masks are allowed (empirically, $T_{\max} > \sqrt{N}$ on these instances), while the MILP may degrade due to the fixed 45 s time limit and the rapidly expanding search space. These results suggest that greedy methods can be attractive building blocks for clinically constrained workflows where runtime is a dominant constraint, while MILP remains useful as a benchmark (and potentially for small $T_{\max}$ regimes). All data and source code can be found in [30].

Table 1
Greedy on oncentra-4.1.6-prostate-ct-tilted-plane.

File	T_{max}	Error	Time (s)
178.txt	15	147 163	1.4
	239	0	3.4
	14 322	0	3.0
	28 644	0	3.5
	42 966	0	3.8
223.txt	15	94 473	1.3
	236	0	2.9
	14 018	0	2.9
	28 037	0	2.9
	42 056	0	2.9
248.txt	15	129 092	1.1
	236	0	2.9
	13 946	0	3.1
	27 893	0	3.2
	41 839	0	4.2
285.txt	15	132 360	1.7
	236	0	4.6
	13 966	0	4.0
	27 933	0	3.8
	41 899	0	4.7
305.txt	15	109 222	1.7
	236	0	4.2
	13 958	0	4.6
	27 917	0	4.5
	41 875	0	4.3
359.txt	15	144 527	1.5
	236	0	4.0
	14 002	0	3.9
	28 005	0	4.2
	42 007	0	3.9

Table 2
Greedy on oncentra-4.1.6-prostate-ct-tilted-plane.

File	T_{max}	Error	Time (s)
378.txt	15	138 718	1.3
	237	0	4.0
	14 046	0	3.7
	28 092	0	3.7
	42 138	0	4.0
416.txt	15	119 637	1.4
	236	0	4.1
	13 956	0	4.3
	27 913	0	3.5
	41 869	0	3.5
502.txt	15	111 034	1.4
	236	0	4.1
	13 938	0	3.9
	27 876	0	3.6
	41 814	0	3.8
508.txt	15	99 858	1.5
	237	0	3.4
	14 067	0	3.3
	28 134	0	3.7
	42 201	0	3.5
653.txt	15	149 667	1.2
	238	0	3.8
	14 179	0	3.9
	28 358	0	2.8
	42 537	0	3.7
734.txt	15	114 193	1.4
	236	0	3.5
	13 988	0	3.3
	27 976	0	3.7
	41 964	0	3.8

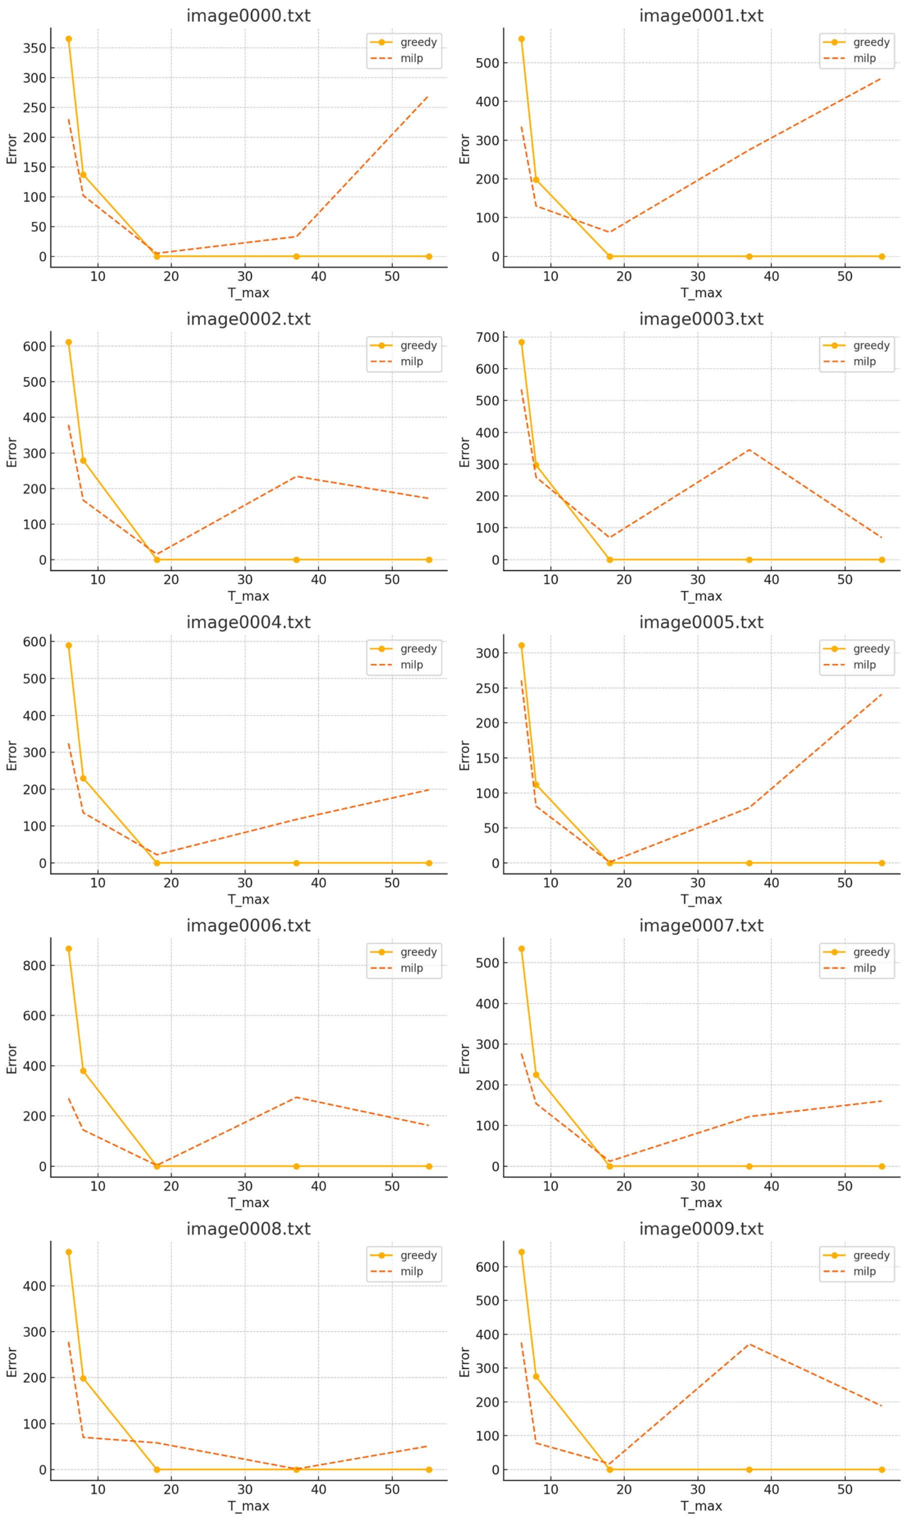

Fig. 3. MILP vs Greedy on plastimatch_tiny-rt-study.

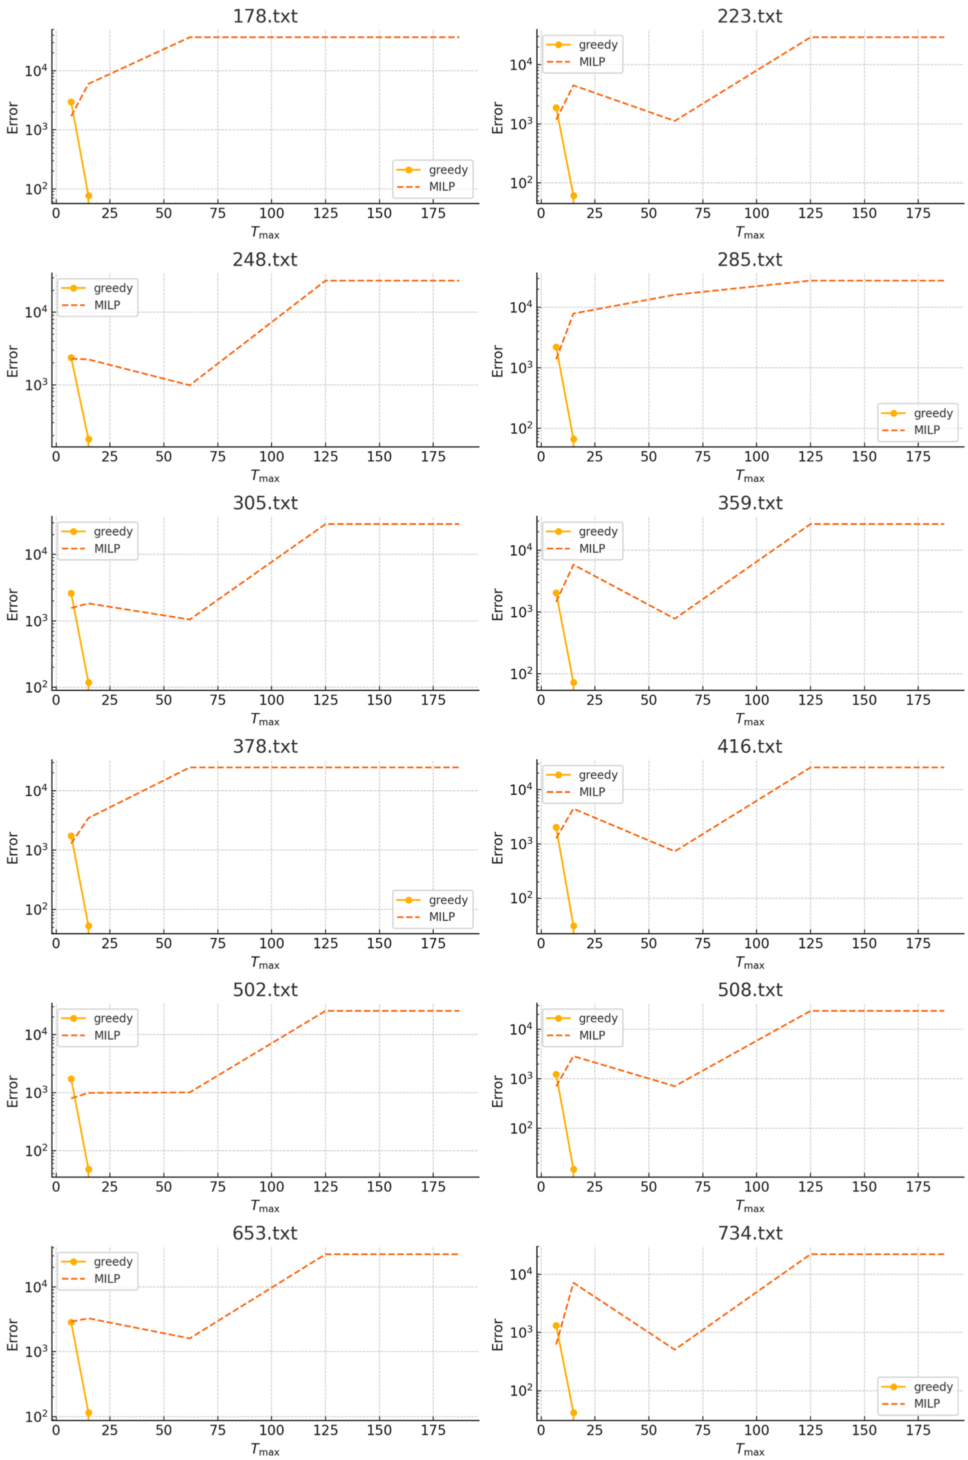

Fig. 4. MILP vs Greedy on oncentra-4.1.6-prostate-ct-tilted-plane.

Table 3
MILP vs Greedy on plastimatch_tiny-rt-study.

file	N	T_{max}	method	error	time (s)
image0000.txt	74	6	greedy	365	0.0015
image0000.txt	74	6	milp	230	45
image0000.txt	74	8	greedy	137	0.0015
image0000.txt	74	8	milp	102	45
image0000.txt	74	18	greedy	0	0.0049
image0000.txt	74	18	milp	5	45
image0000.txt	74	37	greedy	0	0.0050
image0000.txt	74	37	milp	33	45
image0000.txt	74	55	greedy	0	0.0050
image0000.txt	74	55	milp	270	45
image0001.txt	74	6	greedy	562	0.0024
image0001.txt	74	6	milp	335	45
image0001.txt	74	8	greedy	198	0.0029
image0001.txt	74	8	milp	130	45
image0001.txt	74	18	greedy	0	0.0045
image0001.txt	74	18	milp	62	45
image0001.txt	74	37	greedy	0	0.0024
image0001.txt	74	37	milp	275	45
image0001.txt	74	55	greedy	0	0.0025
image0001.txt	74	55	milp	460	45
image0002.txt	74	6	greedy	612	0.0012
image0002.txt	74	6	milp	379	45
image0002.txt	74	8	greedy	279	0.0015
image0002.txt	74	8	milp	167	45
image0002.txt	74	18	greedy	0	0.0031
image0002.txt	74	18	milp	15	45
image0002.txt	74	37	greedy	0	0.0029
image0002.txt	74	37	milp	234	45
image0002.txt	74	55	greedy	0	0.0029
image0002.txt	74	55	milp	172	45
image0003.txt	74	6	greedy	684	0.0016
image0003.txt	74	6	milp	535	45
image0003.txt	74	8	greedy	297	0.0018
image0003.txt	74	8	milp	259	45
image0003.txt	74	18	greedy	0	0.0031
image0003.txt	74	18	milp	69	45
image0003.txt	74	37	greedy	0	0.0047
image0003.txt	74	37	milp	345	45
image0003.txt	74	55	greedy	0	0.0026
image0003.txt	74	55	milp	69	45
image0004.txt	74	6	greedy	590	0.0015
image0004.txt	74	6	milp	324	45
image0004.txt	74	8	greedy	230	0.0015
image0004.txt	74	8	milp	136	45
image0004.txt	74	18	greedy	0	0.0053
image0004.txt	74	18	milp	22	45
image0004.txt	74	37	greedy	0	0.0053
image0004.txt	74	37	milp	118	45
image0004.txt	74	55	greedy	0	0.0054
image0004.txt	74	55	milp	198	45

Table 4
MILP vs Greedy on plastimatch_tiny-rt-study.

file	N	T_{max}	method	error	time (s)
image0005.txt	74	6	greedy	311	0.0023
image0005.txt	74	6	milp	261	45
image0005.txt	74	8	greedy	112	0.0029
image0005.txt	74	8	milp	81	45
image0005.txt	74	18	greedy	0	0.0025
image0005.txt	74	18	milp	1	45
image0005.txt	74	37	greedy	0	0.0027
image0005.txt	74	37	milp	79	45
image0005.txt	74	55	greedy	0	0.0026
image0005.txt	74	55	milp	241	45
image0006.txt	74	6	greedy	866	0.0012
image0006.txt	74	6	milp	270	45
image0006.txt	74	8	greedy	379	0.0015
image0006.txt	74	8	milp	144	45
image0006.txt	74	18	greedy	0	0.0030
image0006.txt	74	18	milp	3	45
image0006.txt	74	37	greedy	0	0.0032
image0006.txt	74	37	milp	274	45
image0006.txt	74	55	greedy	0	0.0031
image0006.txt	74	55	milp	162	45
image0007.txt	74	6	greedy	535	0.0012
image0007.txt	74	6	milp	277	45
image0007.txt	74	8	greedy	225	0.0016
image0007.txt	74	8	milp	154	45
image0007.txt	74	18	greedy	0	0.0027
image0007.txt	74	18	milp	12	0
image0007.txt	74	37	greedy	0	0.0027
image0007.txt	74	37	milp	122	45
image0007.txt	74	55	greedy	0	0.0028
image0007.txt	74	55	milp	160	45
image0008.txt	74	6	greedy	474	0.0013
image0008.txt	74	6	milp	278	45
image0008.txt	74	8	greedy	199	0.0015
image0008.txt	74	8	milp	70	45
image0008.txt	74	18	greedy	0	0.0029
image0008.txt	74	18	milp	58	45
image0008.txt	74	37	greedy	0	0.0028
image0008.txt	74	37	milp	1	45
image0008.txt	74	55	greedy	0	0.0028
image0008.txt	74	55	milp	51	45
image0009.txt	74	6	greedy	644	0.0013
image0009.txt	74	6	milp	376	45
image0009.txt	74	8	greedy	275	0.0029
image0009.txt	74	8	milp	78	45
image0009.txt	74	18	greedy	0	0.0053
image0009.txt	74	18	milp	17	45
image0009.txt	74	37	greedy	0	0.0028
image0009.txt	74	37	milp	371	45
image0009.txt	74	55	greedy	0	0.0027
image0009.txt	74	55	milp	188	45

Table 5
MILP vs Greedy on oncentra-4.1.6-prostate-ct-tilted-plane.

file	N	T_{max}	method	error	time (s)
178.txt	250	7	greedy	2949	0.0038
178.txt	250	7	milp	1699	
178.txt	250	15	greedy	77	0.0092
178.txt	250	15	milp	5989	45
178.txt	250	62	greedy	0	0.0114
178.txt	250	62	milp	36,859	45
178.txt	250	125	greedy	0	0.0111
178.txt	250	125	milp	36,859	45
178.txt	250	187	greedy	0	0.0122
178.txt	250	187	milp	36,859	45
223.txt	250	7	greedy	1888	0.0057
223.txt	250	7	milp	1179	45
223.txt	250	15	greedy	61	0.0078
223.txt	250	15	milp	4482	45
223.txt	250	62	greedy	0	0.0188
223.txt	250	62	milp	1120	45
223.txt	250	125	greedy	0	0.0165
223.txt	250	125	milp	29,247	45
223.txt	250	187	greedy	0	0.0155
223.txt	250	187	milp	29,247	45
248.txt	250	7	greedy	2371	0.0058
248.txt	250	7	milp	2283	45
248.txt	250	15	greedy	179	0.0108
248.txt	250	15	milp	2238	45
248.txt	250	62	greedy	0	0.0128
248.txt	250	62	milp	986	45
248.txt	250	125	greedy	0	0.0176
248.txt	250	125	milp	27,295	45
248.txt	250	187	greedy	0	0.0164
248.txt	250	187	milp	27,295	45
285.txt	250	7	greedy	2237	0.0045
285.txt	250	7	milp	1386	45
285.txt	250	15	greedy	67	0.0139
285.txt	250	15	milp	7929	45
285.txt	250	62	greedy	0	0.0125
285.txt	250	62	milp	16,187	45
285.txt	250	125	greedy	0	0.0164
285.txt	250	125	milp	27,706	45
285.txt	250	187	greedy	0	0.0106
285.txt	250	187	milp	27,706	45
305.txt	250	7	greedy	2626	0.0059
305.txt	250	7	milp	1566	45
305.txt	250	15	greedy	119	0.0092
305.txt	250	15	milp	1836	45
305.txt	250	62	greedy	0	0.0108
305.txt	250	62	milp	1051	45
305.txt	250	125	greedy	0	0.0184
305.txt	250	125	milp	28,729	45
305.txt	250	187	greedy	0	0.0125
305.txt	250	187	milp	28,729	45
359.txt	250	7	greedy	2058	0.0053
359.txt	250	7	milp	1458	45
359.txt	250	15	greedy	72	0.0106
359.txt	250	15	milp	5861	45
359.txt	250	62	greedy	0	0.0121
359.txt	250	62	milp	777	45
359.txt	250	125	greedy	0	0.0130
359.txt	250	125	milp	26,748	45
359.txt	250	187	greedy	0	0.0163
359.txt	250	187	milp	26,748	45

Table 6
MILP vs Greedy on oncentra-4.1.6-prostate-ct-tilted-plane.

file	N	$T_{\max}$	method	error	time (s)
378.txt	250	7	greedy	1754	0.0038
378.txt	250	7	milp	1284	45
378.txt	250	15	greedy	53	0.0128
378.txt	250	15	milp	3460	45
378.txt	250	62	greedy	0	0.0157
378.txt	250	62	milp	24,654	
378.txt	250	125	greedy	0	0.0142
378.txt	250	125	milp	24,654	45
378.txt	250	187	greedy	0	0.0111
378.txt	250	187	milp	24,654	45
416.txt	250	7	greedy	2011	0.0051
416.txt	250	7	milp	1269	45
416.txt	250	15	greedy	31	0.0075
416.txt	250	15	milp	4426	45
416.txt	250	62	greedy	0	0.0104
416.txt	250	62	milp	726	45
416.txt	250	125	greedy	0	0.0095
416.txt	250	125	milp	25,444	45
416.txt	250	187	greedy	0	0.0090
416.txt	250	187	milp	25,444	45
502.txt	250	7	greedy	1738	0.0060
502.txt	250	7	milp	797	45
502.txt	250	15	greedy	48	0.0076
502.txt	250	15	milp	990	45
502.txt	250	62	greedy	0	0.0127
502.txt	250	62	milp	1009	45
502.txt	250	125	greedy	0	0.0143
502.txt	250	125	milp	25,624	45
502.txt	250	187	greedy	0	0.0168
502.txt	250	187	milp	25,624	45
508.txt	250	7	greedy	1233	0.0039
508.txt	250	7	milp	706	45
508.txt	250	15	greedy	15	0.0093
508.txt	250	15	milp	2870	45
508.txt	250	62	greedy	0	0.0118
508.txt	250	62	milp	710	45
508.txt	250	125	greedy	0	0.0150
508.txt	250	125	milp	23,597	45
508.txt	250	187	greedy	0	0.0099
508.txt	250	187	milp	23,597	45
653.txt	250	7	greedy	2873	0.0042
653.txt	250	7	milp	2911	45
653.txt	250	15	greedy	115	0.0136
653.txt	250	15	milp	3270	45
653.txt	250	62	greedy	0	0.0175
653.txt	250	62	milp	1604	45
653.txt	250	125	greedy	0	0.0130
653.txt	250	125	milp	31,853	45
653.txt	250	187	greedy	0	0.0164
653.txt	250	187	milp	31,853	45
734.txt	250	7	greedy	1320	0.0059
734.txt	250	7	milp	620	45
734.txt	250	15	greedy	42	0.0105
734.txt	250	15	milp	7136	45
734.txt	250	62	greedy	0	0.0111
734.txt	250	62	milp	506	45
734.txt	250	125	greedy	0	0.0142
734.txt	250	125	milp	22,055	45
734.txt	250	187	greedy	0	0.0162
734.txt	250	187	milp	22,055	45

9. Conclusions and future work

In this paper, we studied the computational complexity of several variants of radiotherapy shielding mask optimization problems. Extending the work from Blin et al. [8], we generalized their NP-hardness proof for arbitrary w and introduced various problem formulations, each yielding distinct complexity and approximation results.

From the clinical perspective motivating this work help delineate which planning subproblems are computationally intractable in general and thus require approximation algorithms or heuristics. In particular, our experimental comparison illustrates the practical trade-off between optimization strength and runtime under clinically realistic per-slice budgets: MILP can be competitive when the number of allowed masks is small, whereas a simple greedy strategy scales to larger instances and can produce very low-error plans in seconds. These findings support the broader goal outlined in the introduction: enabling fast, deliverable planning methods that better exploit directional shielding to improve dose conformity while maintaining practical computation times.

An intriguing open problem arising from this work is the determination of the computational complexity of FixMasks. Although related variants are known to be computationally challenging, a formal NP-hardness proof for this particular case remains to be established.

CRediT authorship contribution statement

Guillaume Blin: Writing – review & editing, Writing – original draft; **Adrian Miclăuș:** Writing – review & editing, Writing – original draft; **Sebastian Ordyniak:** Writing – review & editing, Writing – original draft; **Alexandru Popa:** Writing – review & editing, Writing – original draft.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

[1] R. Potter, C. Haie-Meder, E. Van Limbergen, I. Barillot, M. De Brabandere, J. Dimopoulos, I. Dumas, B. Erickson, S. Lang, A. Nulens, P. Petrow, J. Rownd, C. Kirisits, Recommendations from gynaecological (GYN) GEC-ESTRO working group (II): concepts and terms in 3D image-based treatment planning in cervix cancer brachytherapy-3D dose-volume parameters and aspects of 3D image-based anatomy, radiation physics, radiobiology, Radiother. Oncol. 78 (1) (2006) 67–77. https://doi.org/10.1016/j.radonc.2005.11.014

[2] M.A. Ebert, Possibilities for intensity-modulated brachytherapy: technical limitations on the use of non-isotropic sources, Phys. Med. Biol. 47 (14) (2002) 2495. http://stacks.iop.org/0031-9155/47/i=14/a=309.

[3] W. Yang, Y. Kim, X. Wu, Q. Song, Y. Liu, S.K. Bhatia, W. Sun, R.T. Flynn, Rotating-shield brachytherapy for cervical cancer, Phys. Med. Biol. 58 (11) (2013) 3931. http://stacks.iop.org/0031-9155/58/i=11/a=3931.

[4] Y. Liu, R.T. Flynn, W. Yang, Y. Kim, S.K. Bhatia, W. Sun, X. Wu, Rapid emission angle selection for rotating-shield brachytherapy, Med. Phys. 40 (5) (2013). https://doi.org/10.1118/1.4802750

[5] Y. Liu, R. Flynn, Y. Kim, X. Wu, Asymmetric dose-volume optimization with smoothness control for rotating-shield brachytherapy, Med. Phys. 41 (11) (2014). Publisher Copyright: © 2014 American Association of Physicists in Medicine. © 2014, Blackwell Publishing. All rights reserved. https://doi.org/10.1118/1.4897617

[6] Y. Liu, R.T. Flynn, Y. Kim, H. Dadkhah, S.K. Bhatia, J.M. Buatti, W. Xu, X. Wu, Paddle-based rotating-shield brachytherapy, Med. Phys. 42 (10) (2015).

[7] M.J. Webster, S. Devic, T. Vuong, D. Yup Han, J.C. Park, D. Scanderbeg, J. Lawson, B. Song, W. Tyler Watkins, T. Pawlicki, W.Y. Song, Dynamic modulated brachytherapy (DMBT) for rectal cancer, Med. Phys. 40 (1) (2013). https://doi.org/10.1118/1.4769416

[8] G. Blin, M. Gasparoux, S. Ordyniak, A. Popa, SOBRA - Shielding Optimization for BRAchytherapy, in: Combinatorial Algorithms 27Th International Workshop, IWOCA 2016, Helsinki, Finland, August 17–19, 2016, Proceedings, Combinatorial Algorithms 27th International Workshop, IWOCA 2016, Helsinki, Finland, August 17–19, 2016, Proceedings, Helsinki, Finland, 2016. https://hal.science/hal-01331027.

[9] M. Antaki, L.D. Christopher, S.A. Enger, Fast mixed integer optimization (FMIO) for high dose rate brachytherapy, Phys. Med. Biol. 65 (21) (2020) 215005. https://doi.org/10.1088/1361-6560/aba317

[10] M. Antaki, M.-A. Renaud, M. Morcos, J. Seuntjens, S.A. Enger, Applying the column generation method to the intensity modulated high dose rate brachytherapy inverse planning problem, Phys. Med. Biol. 68 (6) (2023) 065007. https://doi.org/10.1088/1361-6560/acbc63

[11] M. Cho, X. Wu, H. Dadkhah, J. Yi, R.T. Flynn, Y. Kim, W. Xu, Fast dose optimization for rotating shield brachytherapy, Med. Phys. 44 (10) (2017) 5384–5392. https://doi.org/10.1002/mp.12486

[12] J. Yi, Q.E. Adams, K.M. Hopfensperger, R.T. Flynn, Y. Kim, J.M. Buatti, W. Xu, X. Wu, Optimal Needle Placement for Prostate Rotating-Shield Brachytherapy (RSBT), 2021. https://arxiv.org/abs/2110.07841.

[13] G. Famulari, M. Duclos, S.A. Enger, A novel 169yb-based dynamic-shield intensity modulated brachytherapy delivery system for prostate cancer, Med. Phys. 47 (3) (2020) 859–868. https://doi.org/10.1002/mp.13959

[14] M. Robitaille, C. Ménard, G. Famulari, D. Béliveau-Nadeau, S.A. Enger, 169Yb-based high dose rate intensity modulated brachytherapy for focal treatment of prostate cancer, Brachytherapy 23 (5) (2024) 523–534. https://doi.org/10.1016/j.brachy.2024.05.005

[15] B. Morén, T. Larsson, A.C. Tedgren, Optimization in treatment planning of high dose-rate brachytherapy - review and analysis of mathematical models, Med. Phys. 48 (5) (2021) 2057–2082. https://doi.org/10.1002/mp.14762

[16] T.J. Schaefer, The complexity of satisfiability problems, in: Proceedings of the Tenth Annual ACM Symposium on Theory of Computing, STOC '78, Association for Computing Machinery, New York, NY, USA, 1978, p. 216-226. https://doi.org/10.1145/800133.804350

[17] R.M. Karp, Reducibility among Combinatorial Problems, Springer US, Boston, MA, 1972, pp. 85–103. https://doi.org/10.1007/978-1-4684-2001-2_9

[18] P. Alimonti, V. Kann, Some APX-completeness results for cubic graphs, Theor. Comput. Sci. 237 (1) (2000) 123–134. https://doi.org/10.1016/S0304-3975(98)00158-3

[19] Y. Caro, New results on the independence number, Technical Report, Technical Report, Tel-Aviv University, 1979.

[20] V.K. Wei, A lower bound on the stability number of a simple graph, 1981.

[21] H.W. Lenstra, Integer programming with a fixed number of variables, Math. Oper. Res. 8 (4) (1983) 538–548. http://www.jstor.org/stable/3689168.

[22] J. Vondrák, Optimal approximation for the submodular welfare problem in the value oracle model, in: Proceedings of the 40Th Annual ACM Symposium on Theory of Computing (STOC), 2008, pp. 67–74.

[23] C. Chekuri, J. Vondrák, R. Zenklusen, Submodular function maximization via the multilinear relaxation and contention resolution schemes, SIAM J. Comput. 43 (6) (2014) 1831–1879. https://doi.org/10.1137/110839655

[24] C. Chekuri, J. Vondrák, R. Zenklusen, Submodular function maximization via the multilinear relaxation and contention resolution schemes, SIAM J. Comput. 43 (6) (2014) 1831–1879.

[25] T. Cormen, C. Leiserson, R. Rivest, C. Stein, Introduction to Algorithms, MIT Press, 2014.

[26] M. Blum, R.W. Floyd, V. Pratt, R.L. Rivest, R.E. Tarjan, Time bounds for selection, J. Comput. Syst. Sci. 7 (4) (1973) 448–461. https://doi.org/10.1016/S0022-0000(73)80033-9

[27] L. Gurobi Optimization, Gurobi Optimizer Reference Manual, 2023, (https://www.gurobi.com).

[28] C. Pinter, SlicerRtData, 2017, (https://github.com/SlicerRt/SlicerRtData).

[29] S.M. van Vliet-Pérez, R. van Paassen, L.S.G.L. Wauben, R. Straathof, N.J. van de Berg, J. Dankelman, B.J.M. Heijmen, I.K. Kolkman-Deurloo, R.A. Nout, Time-action and patient experience analyses of locally advanced cervical cancer brachytherapy, Brachytherapy 23 (3) (2024) 274–281. https://doi.org/10.1016/j.brachy.2024.01.007

[30] G. Blin, A. Miclăuş, S. Ordyniak, A. Popa, SOBRA, 2025, https://github.com/adimiclaus15/SOBRA.