

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and should not be used to distribute the papers in print or online or to submit the papers to another publication.

A prediction-based approach for online dynamic appointment scheduling: a case study in radiotherapy treatment

Tu-San Pham

Polytechnique Montréal and CIRRELT, Montréal, Canada, tu-san.pham@polymtl.ca,

Antoine Legrain

Polytechnique Montréal, GERAD, and CIRRELT, Montréal, Canada, antoine.legrain@polymtl.ca,

Patrick De Causmaecker

KU Leuven, Kortrijk, Belgium, patrick.decausmaecker@kuleuven.be,

Louis-Martin Rousseau

Polytechnique Montréal and CIRRELT, Montréal, Canada, louis-martin.rousseau@polymtl.ca,

Patient scheduling is a difficult task involving stochastic factors such as the unknown arrival times of patients. Similarly, the scheduling of radiotherapy for cancer treatments needs to handle patients with different urgency levels when allocating resources. High priority patients may arrive at any time, and there must be resources available to accommodate them. A common solution is to reserve a flat percentage of treatment capacity for emergency patients. However, this solution can result in overdue treatments for urgent patients, a failure to fully exploit treatment capacity, and delayed treatments for low-priority patients. This problem is especially severe in large and crowded hospitals. In this paper, we propose a prediction-based approach for online dynamic radiotherapy scheduling that dynamically adapts the present scheduling decision based on each incoming patient and the current allocation of resources. Our approach is based on a regression model trained to recognize the links between patients' arrival patterns, and their ideal waiting time in optimal offline solutions where all future arrivals are known in advance. When our prediction-based approach is compared to flat-reservation policies, it does a better job of preventing overdue treatments for emergency patients, while also maintaining comparable waiting times for the other patients. We also demonstrate how our proposed approach supports explainability and interpretability in scheduling decisions using SHAP values.

Key words: Operations Research; Radiotherapy scheduling; Integer Programming; Patient Scheduling; Explainability;

1. Introduction

According to the WHO ¹, cancer is the leading cause of death worldwide, responsible for nearly 10 million deaths in 2020. In Canada, the number of cancer deaths in 2020 was 83,300, accounting for about 30% of all deaths in the country. The number of cancer incidents is steadily rising in many countries, which leads to increased stress on treatment facilities and medical staff, while also making it more difficult to get timely access to necessary cancer treatments. Numerous studies (Coles et al. 2003, Mackillop 2007, Chen et al. 2008) have proven that a long waiting time preceding cancer treatment has a negative impact on the clinical outcome of that treatment (e.g. persistence of cancer symptoms, psychological distress, increase of death rate). In cancer treatment, waiting time is referred to as killing time (Mackillop 2007). Therefore, it is essential to improve the treatment scheduling process, as it will reduce the waiting time before patients start treatment.

In this paper, we focus on radiotherapy (RT), one of the most popular forms of cancer treatment. About 50% of cancer patients undergo RT treatment (Barton et al. 2014, Tyldesley et al. 2011). During RT treatment, a patient receives a high dose of radiation divided into small portions (called *fractions*) that are delivered to the cancer site over several consecutive days. This treatment is most commonly delivered by a linear accelerator (*linac*). As the number of linacs in a hospital is usually limited, the waiting time for cancer patients to start their RT treatments directly depends on treatment scheduling. We consider the RT scheduling problem arising at CHUM (Centre hospitalier de l'Université de Montréal), a large cancer center in Montréal, Canada. At CHUM, patients are categorized by the maximum recommended waiting time given by their doctor. Overdue treatments (treatments started after their recommended deadline) are highly discouraged. The objective is to minimize overdue treatments and patient waiting times. The main challenge presented by this problem is uncertainty in the demand for treatment by patients with different urgency levels. Palliative patients need urgent care to relieve intense pain, and so their treatment deadline is set to 1 to 3 days after their admission. In contrast, curative patients can wait for 2 to 4 weeks. Approximately 70% of patients treated at CHUM are curative. Despite accounting for only 30% of treatments, palliative treatments are the most challenging group to schedule due to their late arrivals and short treatment deadlines (Pham et al. 2021). Scheduling curative patients too early risks leading to overdue treatments in palliative patients by leaving too few reserved treatment slots available. However, delaying treatments in curative patients for too long could lead to inefficient capacity usage, while also increasing patient waiting times. Therefore, it is important to take

¹<https://www.who.int/news-room/fact-sheets/detail/cancer>

into account future patient arrivals while making scheduling decisions. Another challenge is the scheduling methodology employed at CHUM. Currently, scheduling tasks are done manually in an *online* fashion, i.e. patients are scheduled one by one as they are admitted. *Online scheduling* has many disadvantages compared to *batch scheduling*, where scheduling decisions are delayed to the end of the day or week, so that several patients can be scheduled at once. Batch scheduling can leverage accumulated information, and usually results in better schedules than online scheduling. However, it is difficult to adopt batch scheduling at CHUM. The scheduling staff need to call patients to confirm appointments and make adjustments when necessary. Online scheduling makes this communication easier. Therefore, it is required by CHUM to use online scheduling.

We propose a novel *online prediction-based approach* for *dynamic radiotherapy scheduling*, which uses the historical arrival patterns of patients to inform scheduling decisions. An Integer Programming (IP) model is proposed to derive optimal *offline schedules* for the variant of the problem where future arrivals are known in advance. Given a large number of instances with the same arrival rate and number of linacs, a regression model can be trained to detect the links between a patient's features, the (resource) allocation profile of the hospital, and the ideal schedule. The regression model can then predict a "good" waiting time for a patient. Patients are scheduled in an online fashion. Palliative patients are always scheduled on the earliest possible date, while curative patients are scheduled using a prediction model that suggests a waiting time based on their treatment plan and the present occupancy of the linacs. In other words, the prediction model is trained to dynamically delay treatments of curative patients, to make space for palliative ones based on the current allocation profile of the hospital. This results in a robust scheduling approach that is able to take future arrivals into account. We refer to our proposed machine learning-based approach as the *prediction-based approach*. We compare the prediction-based approach to myopic scheduling strategies, including the online greedy heuristic currently used at CHUM, a batch scheduling heuristic using a simple ordering rule, and two batch scheduling strategies using the IP model (daily and weekly). The algorithms are tested in a simulation with a rolling horizon to evaluate their performance over a fixed-time period. Test datasets are generated using real data from CHUM. The prediction-based approach outperforms the other approaches for most problem settings, which vary by the number of linacs and the arrival rate. It results in the lowest average overdue time and average waiting time for palliative patients, while being comparable to other approaches with respect to average overdue times and average waiting times for curative patients. This demonstrates its ability to efficiently integrate information about future arrivals into daily

scheduling decisions. Notably, the superiority of the prediction-based approach is more prominent in larger hospital settings with more linacs and more patients to treat (high arrival rate). The proposed approach also proves its robustness in a real problem instance, despite the high fluctuation in arrival rates observed in reality. Additionally, we demonstrate how our approach supports the explainability of the decision-making process using SHAP values. Explainability in decision-making enables better communication with clients. It is especially appreciated in the healthcare domain where most decisions are human-related. Explainability is a very active domain of research in Artificial Intelligence, yet receives little attention in the Operations Research community. As a hybrid method combining techniques from both domains, our approach supports explainable and interpretable scheduling decisions. Our contributions can be summarized as follows:

- We propose an online machine learning-based scheduling approach for radiation therapy. The machine learning model is trained on offline optimal solutions obtained via mathematical optimization to mimic the optimal policy without the online computational burden.
- The proposed approach successfully reduces overdue treatments and waiting times for patients compared to other optimization-based approaches. We also successfully solve instances with up to 8 linacs, which represents a larger hospital size than those considered in previous studies. Furthermore, it is robust to the high fluctuation in arrival rates observed in practice.
- The prediction-based approach is simple, easy to implement and maintain. It also offers solutions instantly, which makes it suitable for online scheduling problems.
- We demonstrate how our approach supports explainability and interpretability in scheduling decisions, to our knowledge for the first time. This is highly appreciated in the healthcare domain.

The rest of the paper is organized as follows. Section 2 provides a problem description and literature review. Section 3 presents the methodology. Section 4 presents numerical results. Section 5 discusses theoretical and practical implications of our approach, and Section 6 closes the paper with our conclusions.

2. Problem description and related work

In this section, we describe the radiotherapy scheduling problem arising at CHUM, and then present related work and situate our work in the literature.

2.1. Problem description

In RT treatment, a patient receives a series of radiation doses (fractions) delivered by a linac to the cancer site during the treatment course, which typically consists of several consecutive days

with breaks on the weekends. Multiple appointments are required. To simplify the problem, we assume that all linacs are identical. This does not affect the complexity of the problem, as special treatments that require special linacs are independent of the rest of the treatments, and can be scheduled separately in a sub-problem. In practice, each linac is associated with a technician team that will familiarize themselves with the technical settings of each patient at the beginning of their treatment. Therefore, hospital policy does not allow switching linacs during treatment. At CHUM, patients are divided into four categories with different waiting time targets based on their cancer type and condition. Palliative patients (categories P1 and P2) have their waiting time target set as 1 and 3 days, respectively. In curative patients (categories P3 and P4), the target is 14 and 28 days, respectively. To minimize the waiting time of patients, one only needs to assign the start date for treatment, as once a treatment starts it must be carried out daily. There are no treatments on weekends, but those gaps have been taken into account in the treatment protocols. The hospital does not allow gaps in treatment induced by planning. The waiting time of a patient is the time elapsed between their admission date, and the date of their first treatment. If a patient's treatment starts after the recommended deadline, the time delay is counted as overdue time.

Cancer treatment is a complicated process with many procedures and parties involved. The treatment workflow used at CHUM is illustrated in Figure 1. When a new consultation request is made, a patient first undergoes a consultation session where the doctor explains their condition and available options. If the patient agrees to proceed with the treatment, they will go through several preparation steps such as external consultation, exams, and a planification scan. Based on the patient's diagnosis, a treatment team then makes a personal treatment plan, which is verified and approved by a medical physicist before the treatments are prepared. The RT treatments are then carried out, with review and possible revision during the course of treatment. Finally, the treatment ends with post-treatments and follow-ups. Currently, the scheduling of treatments is performed manually at CHUM. Among various technical information, a treatment plan includes information about the number of fractions that a patient will receive, as well as the duration of each fraction. Since fraction durations at CHUM are always multiples of 5 minutes, we set the granularity of our schedules to 5-minute blocks. Each linac has a treatment capacity measured in time blocks.

2.2. Literature review

Patient scheduling has a long history and a rich literature. For a thorough review of the topic, we refer the readers to (Marynissen and Demeulemeester 2019, Cayirli and Veral 2003, Gupta and

² Figure provided by CHUM.

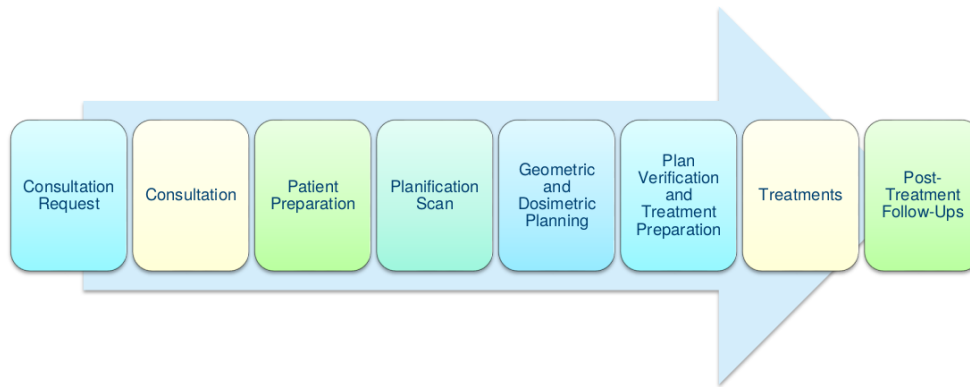


Figure 1 Radiotherapy treatment workflow at CHUM².

Denton 2008). There are many ways to classify patient scheduling problems. To analyse related work, we adopt the classification of Saure et al. (2012), which divides patient scheduling into *allocation scheduling* (or *resource allocation*) and *appointment scheduling*. Appointment scheduling assigns specific appointment times to patients on a given service day. In most appointment scheduling problems, the main concern is maximizing patient comfort, satisfying precedence constraints, dealing with uncertainty in appointment duration and cancellations, etc. In contrast, allocation scheduling schedules patients in advance, taking medical capacity into account, without giving exact appointment times. Allocation scheduling usually involves patients of different types and urgency levels, therefore the main challenge is dealing with uncertainty, such as unknown future demands and cancellations. Algorithms for allocation scheduling are usually dynamic, meaning they focus on dynamically allocating medical resources on a rolling horizon. Our paper falls into the category of allocation scheduling. We review the literature on dynamic patient scheduling and radiotherapy scheduling. The literature is vast, so this not an extensive review, and only the most relevant papers are mentioned here.

2.2.1. Dynamic patient scheduling. There are different sources of uncertainty in healthcare: cancellations, uncertainty in service time, stochastic patient arrivals, deferrals in treatment plans, etc. While our paper focuses only on uncertainty resulting from stochastic patient arrivals, there are plenty of studies applying machine learning methods to tackle different sources of uncertainty. For example, Srinivas and Ravindran (2018) develop predictive models that identify the risk of a patient no-show and propose different scheduling rules that take into account the risk. A similar approach is used by Salah and Srinivas (2022) in an appointment scheduling system with the presence of stochastic factors from no-shows and uncertainty in service time.

Regarding uncertainty caused by stochastic patient arrivals, operating room (OR) planning has to deal with a high degree of uncertainty as ORs are usually shared between both elective (planned) procedures and emergency surgeries. Gerchak et al. (1996) propose a stochastic dynamic programming model for the OR planning problem, and then analyse the structure of the optimal policy. In the same problem, Lamiri et al. (2008) propose a Monte Carlo optimization method that combines Monte Carlo simulation with Mixed Integer Programming (MIP), and show that their approach yields about a 4% reduction of the overall cost compared to a deterministic model.

Vermeulen et al. (2009) present an adaptive approach to schedule CT scans at a radiology department. In their approach, the allocation of capacity to different patient groups is flexible and adaptive to the current and expected future situation. However, the algorithm is specific to their case study. Patrick et al. (2008) model a similar problem as a Markov decision process (MDP), and solve it using simulation-based approximate dynamic programming (ADP). Saure et al. (2012) and Sauré et al. (2020) extend this approach to more complicated problems. Schuetz and Kolisch (2012) investigate a capacity allocation problem in the service industry with requests of different priorities. The authors also adopt a similar approach, combined with a discrete event simulation.

In another problem family, patient admission scheduling (PAS), Ceschia and Schaerf (2012) propose a MIP model and Simulated Annealing heuristic to solve a PAS problem under uncertainty. They propose a *degree of dynamism*, adopted from dynamic vehicle routing problems, to denote a level of urgency. Ceschia and Schaerf (2016) then extend the problem with more realistic constraints and propose a local search-based approach. Zhu et al. (2019) study the compatibility of short and long-term objectives for the dynamic PAS. The authors propose a short-term MIP model which adjusts the objective functions based on patient delay or idle resources.

2.2.2. Radiotherapy scheduling. As opposed to patient scheduling, RT scheduling is a relatively new field of research. A review of RT scheduling can be found in (Kapamara et al. 2006, Vieira et al. 2016). The first two mathematical models for RT scheduling are proposed by Conforti et al. (2008) to maximize the number of treated patients within a given horizon with identical treatment times and a single linac. The model assigns patients to treatment slots of equal length, which is referred to as *block scheduling*. The authors then extend their models in (Conforti et al. 2010) to allow for more realistic settings where treatment duration can vary in length (*non-block scheduling*). Petrovic et al. (2006) and Petrovic and Leite-Rocha (2008) propose several constructive heuristics based on different prioritised rules, and improve the solutions using metaheuristics. They use block scheduling where treatment on a linac is of fixed duration,

and depends on the energy type of the linac. Burke et al. (2011) propose a non-block IP model to assign patients to days and linacs, and perform several simulations to investigate whether delaying the scheduling decision can lead to better schedules. Heuristic methods are also used by Kapamara and Petrovic (2009) to schedule patients in both pre-treatment and treatment stages. Pre-treatment for RT patients is also considered by Castro and Petrovic (2012). The authors solve a multi-objective problem by formulating it as a series of single-objective scheduling problems.

Another form of RT treatment is particle therapy, which uses a centralized machine to deliver an ion beam to different treatment rooms. Most studies in particle therapy focus on assigning specific appointment times for treatments to optimize the beam's utilization, i.e., minimizing the idle time of the machine. Maschler et al. (2016) propose an exact model, which is proven to be highly intractable. As a result, metaheuristics are widely used for particle therapy treatment scheduling (Maschler et al. 2016, Maschler and Raidl 2018, Vogl et al. 2019). In another context, Vieira et al. (2020) point out that most cancer centers in Europe are capable of treating their patients within the recommended waiting period. Therefore, they focus on meeting patients' preferences for treatment times. Frimodig and Schulte (2019) propose an IP and two constraint programming (CP) models for scheduling RT treatments. Their generated instances include an expected number of future arrivals, which they use to predict future resource utilization.

All of the aforementioned methods are tested on a static horizon without taking future arrivals into account. RT scheduling involves patients of different priorities, so using a dynamic resource allocation scheme that takes uncertainty of treatment demand into account is highly important. Pham et al. (2021) propose a two-phase approach for RT scheduling, and implement a simulation to evaluate the algorithm dynamically on a rolling horizon using different scheduling strategies. However, the approach is still myopic, i.e., it does not take future arrivals into account when making scheduling decisions. To handle high-priority patients, many cancer centers reserve a number of treatment slots per day for emergency treatments. This approach is used in several papers (Petrovic and Leite-Rocha 2008, Legrain et al. 2015, Pham et al. 2021, Frimodig and Schulte 2019). Nonetheless, it is difficult to set a proper threshold.

To the best of our knowledge, there are only two approaches for dynamic RT scheduling which take future events into account. The first approach is MDP, used by Saure et al. (2012) to provide a dynamic policy for RT scheduling. The problem is modeled as a discounted infinite-horizon MDP, and the equivalent Integer Programming (IP) model is solved using column generation. The same approach is also used by Gocgun (2018), but they also allow for the cancellation of treatments.

In contrast, Legrain et al. (2015) propose a hybrid method combining stochastic optimization and online optimization to solve the problem in an online fashion. Stochastic programming and MDPs are both prone to scaling problems which makes it difficult to apply either approach in large hospitals. The example studied by Saure et al. (2012) has an arrival rate of 8.25 requests per day with 120 appointment slots, which is equivalent to three linacs. Legrain et al. (2015) consider a 20-minute time slot for each patient, and test the algorithm on instances with up to two linacs and an arrival rate less than 3.5 requests per day. In addition, they are algorithmically heavy, and may be difficult to implement and maintain in real-world applications. In the literature, there are few attempts that tackle the scaling problem resulting from using stochastic programming with machine learning techniques. One example can be found in the context of two-stage stochastic programming (Larsen et al. 2021), where the authors use supervised machine learning to predict the computationally expensive second stage and apply the algorithm to a railway demand and capacity management problem. Our approach aims to fill this gap in the RT scheduling literature. We propose an algorithm that dynamically allocates resources based on expected future arrivals, and targets larger instances than those considered before in the literature, i.e., up to 8 linacs.

3. Methodology

Section 3.1 presents the mathematical model, and Section 3.2 presents several scheduling strategies. The remainder of Section 3 explains the prediction-based approach.

3.1. Mathematical model

The problem being solved consists of finding the best schedule of treatments for a set of patients \mathcal{P} , given a set of linacs \mathcal{L} . The sets of palliative patients and curative patients are denoted as \mathcal{P}^P and \mathcal{P}^C , respectively (so $\mathcal{P} = \mathcal{P}^P \cup \mathcal{P}^C$). Each instance consists of a set of fixed patients with appointments made from the previous scheduling decisions, denoted as $\bar{\mathcal{P}}$, while the set of new patients is denoted as $\hat{\mathcal{P}}$ (so $\mathcal{P} = \bar{\mathcal{P}} \cup \hat{\mathcal{P}}$). The set of working days in the planning horizon is denoted as \mathcal{T} . Each patient $i \in \mathcal{P}$ has an admission date a_i , a ready date r_i which is the earliest date the patient can begin treatment, and a due date d_i which is the recommended deadline to start the treatment. Each patient i has a personal treatment plan specifying the number of fractions I_i , and the duration of each fraction (p_i). Each linac l has a capacity (C_l^t) on day t , measured in blocks of 5 minutes. \hat{C}_l^t represents the available capacity of linac l on day t after deducting the fixed appointments from the previous scheduling decisions.

Our objective is to minimize overdue time and waiting time by determining the assignment of patients' first fractions to dates and linacs. Once treatment starts, it is carried out every day

Notation	Explanation	Notation	Explanation
\mathcal{L}	set of linacs	a_i	admission date of patient i
\mathcal{P}^P	set of palliative patients	α_i	admission time of patient i (consists of date and time)
\mathcal{P}^C	set of curative patients	r_i	ready date of patient i ($r_i \geq a_i$)
$\bar{\mathcal{P}}$	set of fixed patients	d_i	due date of patient i ($d_i \leq r_i$)
$\hat{\mathcal{P}}$	set of new patients	I_i	number of fractions of patient i
\mathcal{T}	set of working days in the planning horizon	p_i	duration of each fraction of patient i
C_l^t	total capacity of linac l on day t	\hat{C}_l^t	available capacity of linac l on day t

Table 1 List of notation.

until the required number of fractions is achieved. Subsequent fractions of each patient should be carried out on the same linac as the first treatment. A portion of linac capacity is reserved for palliative patients. Parameter γ ($\gamma < 1$) indicates the percentage of reserved capacity. We define a binary variable x_{tl}^i on the set of new patients $\hat{\mathcal{P}}$, that holds value 1 if patient i receives their first treatment on day t linac l . The model is as follows:

$$\begin{aligned} \text{minimize } & \sum_{i \in \hat{\mathcal{P}}} \sum_{t \in \mathcal{T}, t > r_i} \sum_{l \in \mathcal{L}} \omega_1(t - a_i) \log(t - a_i + 1) x_{tl}^i \\ & + \sum_{i \in \hat{\mathcal{P}}} \sum_{t \in \mathcal{T}, t > d_i} \sum_{l \in \mathcal{L}} \omega_2(t - d_i) \log(t - d_i + 1) x_{tl}^i \end{aligned} \quad (1)$$

Subject to

$$\sum_{t \in \mathcal{T}} \sum_{l \in \mathcal{L}} x_{tl}^i = 1 \quad \forall i \in \hat{\mathcal{P}} \quad (2)$$

$$x_{tl}^i = 0 \quad \forall i \in \hat{\mathcal{P}}, l \in \mathcal{L}, t \in \{0, \dots, r_i - 1\} \quad (3)$$

$$\sum_{i \in \hat{\mathcal{P}}} \sum_{t' = \max\{0, t - I_i + 1\}}^t p_i x_{t'l}^i \leq \hat{C}_l^t \quad \forall t \in \mathcal{T}, l \in \mathcal{L} \quad (4)$$

$$\sum_{i \in \mathcal{P}^C} \sum_{t' \in \{t - I_i + 1, \dots, t\}} p_i x_{t'l}^i \leq \max\{0, \hat{C}_l^t - \gamma C_l^t\} \quad \forall t \in \mathcal{T}, l \in \mathcal{L} \quad (5)$$

$$x_{tl}^i \in \{0, 1\} \quad \forall i \in \hat{\mathcal{P}}, t \in \mathcal{T}, l \in \mathcal{L} \quad (6)$$

The objective function (1) minimizes waiting time (the first term) and overdue time (the second term) with the respective weights ω_1 and ω_2 . By setting $\omega_2 \gg \omega_1$, we heavily penalize overdue time over waiting time. The log function is a non-linear function utilized to enforce an equal distribution of waiting time and overdue time between patients. Constraints (2) ensure all patients are scheduled within the planning horizon. Constraints (3) ensure no patients are scheduled before their ready date. Constraints (4) make sure linac capacity is respected for the entire course of treatment. Constraints (5) reserve a portion of linac capacity for palliative patients.

3.2. Scheduling strategies

In this section, we propose several scheduling strategies for radiotherapy treatments. All strategies utilize online scheduling for palliative patients, i.e., their treatments are scheduled for the earliest possible date on an available linac when the treatment request is submitted. A portion of linac capacity is reserved for palliative patients. The proposed strategies differ in the manner in which curative patients are handled. All scheduling strategies are summarized in Table 2.

Strategy	Scheduling palliative patients	Scheduling curative patients
offline	one time	one time
online-greedy	at admission	at admission
daily-greedy	at admission	every day
daily-IP	every day	every day
weekly-IP	every day	every Friday
prediction-based	at admission	at admission

Table 2 Scheduling strategies.

3.2.1. Online scheduling with a greedy heuristic This strategy is the closest one to the approach used by scheduling clerks at CHUM. Once a patient is admitted (usually when their treatment plan is approved), the algorithm looks up the following one or two weeks depending on the patient's category (P3 or P4). The first date that can accommodate all of the patient's fractions is then chosen. A date is considered eligible if on this day, and the following $(I_i - 1)$ days, there is a linac with enough remaining capacity to treat the patient, where I_i represents the number of fractions assigned to the patient. An illustration of the greedy heuristic can be found in Figure 2.

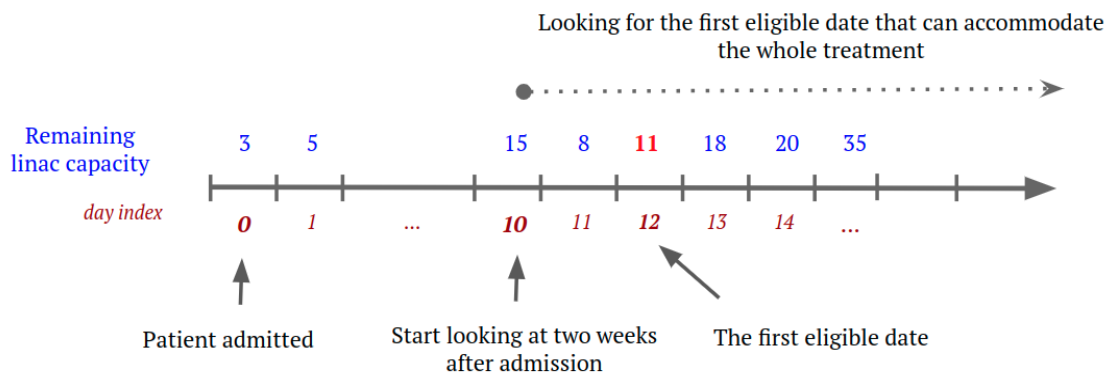


Figure 2 The greedy heuristic for an instance with one linac. A curative patient (category P4) with 3 fractions, each of which requires 10 time slots, is admitted on day 0. The algorithm looks forward two weeks, starting from day 10 (only business days are indexed). The patient is scheduled on day 12, the first day with enough capacity in subsequent days to accommodate the whole treatment.

3.2.2. Batch scheduling When batch scheduling, treatment requests are accumulated for a period of time, and then a scheduling decision is made by the end of the period, e.g., daily or weekly. Each batch of scheduling decisions includes a set of patients admitted the day after the previous scheduling decision was made. We propose two strategies to schedule patients in batches: greedy-based, and IP-based. The greedy-based heuristic sorts patients in descending order by their due day (which is determined by their priorities), numbers of fractions, and fraction lengths. Patients are then scheduled in a greedy manner following the given order. This prioritizes patients with more fractions and longer fraction lengths, as it is more difficult to find an eligible start day for them. In IP-based batch scheduling, patients are scheduled using the IP model in Section 3.1.

3.2.3. Offline scheduling In offline scheduling, we assume that all patient arrivals are known in advance. First, the algorithm schedules palliative patients for the earliest eligible day after their arrival. Curative patients admitted during the simulation period are then scheduled using the IP model in Section 3.1, without reserving a portion of linac capacity for palliative patients. The scheduling decision is made on the first day of the scheduling period. Having all future patient arrivals known in advance is an unrealistic assumption, but it provides a bound for other scheduling strategies. No strategy can bypass offline scheduling due to the lack of information on future arrivals. It also provides a picture of how a “perfect schedule” would look if we had a precise prediction of future arrivals.

3.2.4. Online scheduling using regression - a machine learning approach This scheduling strategy is the main contribution of the paper. Scheduling is carried out in an online fashion similar to the greedy heuristic. However, instead of looking forward one or two weeks, we apply a regression model to predict a start date for each patient. The algorithm then looks for an eligible date starting from the predicted day. An illustration of the algorithm is given in Figure 3. The details of the regression model will be presented in Section 3.3.

3.3. Prediction-based approach - a machine learning-based algorithm

In most cancer centers, the number of linacs is fixed. The arrival rate of patients is usually stable and can be modelled as a Poisson distribution. Given an instance with a number of linacs and an arrival rate, an optimal offline solution (referred to simply as an *offline solution*) given by offline scheduling (Section 3.2.3) provides an “ideal” schedule where all future arrivals are known in advance. All other scheduling strategies based on the gradual realization of patient arrivals share the same goal of achieving solutions as close as possible to offline solutions. The motivation for our approach is

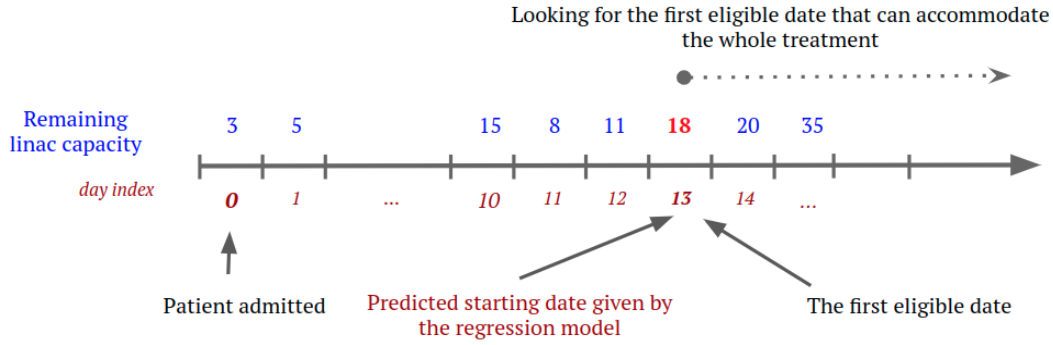


Figure 3 Prediction-based scheduling for the same instance as Figure 2. The algorithm starts the search from the predicted starting day (day 13). The patient is scheduled to the first eligible day (day 13).

the idea that given a large enough number of instances and their offline solutions, we might be able to learn the patterns in the ideal optimal waiting time of patients. Thus, given the present occupancy of linacs and the treatment plan of a patient, we can predict a “good” waiting time for the patient. In the following section, we describe our method of training the regression model and applying it in a machine learning-based scheduling algorithm, which we refer to as a *prediction-based approach*.

3.4. Training the regression model

A problem instance is characterized by a number of linacs $|\mathcal{L}|$ and an arrival rate λ . A pair $(|\mathcal{L}|, \lambda)$ is referred to as an *instance setting*, which represents a typical cancer center with a fixed number of linacs and a stable arrival rate. In all our instances, all linacs have a capacity of 120 blocks of 5 minutes a day. As each hospital has its own historical ideal scheduling patterns, an instance setting is associated with a separate regression model. For each instance in the training set, an offline solution is obtained by solving the offline scheduling as in Section 3.2.3. A set of training examples is then constructed from the instance and its offline solution.

To explain how training examples are constructed, we first define some terms. A problem instance has a set of simulation days (denoted as \mathcal{K}), each of which observes a set of patients admitted to the center (denoted as $\mathcal{P}_d, d \in \mathcal{K}$). A problem instance consists of the daily flows of patients admitted to the hospital during the simulation period, which is referred to as a patient flow ζ , where $\zeta = \{\mathcal{P}_d, \forall d \in \mathcal{K}\}$. A solution to an instance is the assignment of patients to a set of dates (\mathcal{D}) and linacs (\mathcal{L}): $s : \mathcal{P} \times \mathcal{D} \times \mathcal{L} \rightarrow \{0, 1\}$, where $s_{dl}^i = 1$ if patient i is assigned to day d linac l .

The *present time point* at any moment during the simulation period is denoted as ϕ , and the corresponding date is denoted as $d(\phi)$. At time point ϕ , the *present capacity* of linacs on date d is the available capacity after removing the treatment slots occupied by patients admitted (and scheduled) before the time point ϕ on the corresponding date. $\hat{c}_d^\phi = C - \sum_{i \in \mathcal{P}, l \in \mathcal{L} | s_{dl}^i = 1, \alpha_i < \phi} p_i$, where C is the

total linac capacity, and p_i and α_i are the fraction duration and admission time point of patient i , respectively. Note that α_i is the admission time point, which includes both admission date a_i and admission time. We distinguish between α_i and a_i , because the present capacity \hat{C}^ϕ varies by time points during the day. The *state* of an instance at time point ϕ , denoted as \hat{C}^ϕ , is the set of *present capacity* of all days in the *sample horizon*, $\hat{C}^\phi = \{\hat{c}_d^\phi, \forall d \in \mathcal{D}^\phi\}$, where \mathcal{D}^ϕ is the set of days in the sampling horizon at the time point ϕ . In our problem, the sampling horizon is set as 50 days after $d(\phi)$, $\mathcal{D}^\phi = \{d(\phi), d(\phi) + 1, \dots, d(\phi) + 49\}$, since in our realistic instance settings it is unlikely that a patient is not scheduled within 50 days of their admission. Therefore, when making a scheduling decision, the linac capacity search is restricted to a window of 50 days after their admission.

Given an instance and its offline solution s^* , a training example consisting of an input vector and a label is created for each curative patient in the patient flow. The label of patient i is the patient's waiting time w_i in the offline solution. The input vector consists of the *state* measured at the patient's admission time $\phi = \alpha_i$, and a vector of the patient's features. We select the features that are the most relevant to the scheduling task, i.e., the patient's ready date r_i , due date d_i , number of sessions I_i , and fraction length p_i . Thus, a training example for patient i is a tuple (X_i, w_i) , with input vector $X_i = \{\hat{C}^{\phi, \phi=\alpha_i}, r_i, I_i, d_i, p_i\}$. The collinearity of the features are tested in Appendix A.

The algorithm for creating training examples is presented in Algorithm 1. As the regression model is used to schedule curative patients only, one training example is created for each new curative patient in the instance. At the beginning of the algorithm, the present time point is set as the beginning of the first day in the simulation period \mathcal{K} , and the present capacity \hat{C}^ϕ is calculated accordingly. For each day $d \in \mathcal{K}$, the set of patients admitted on that day (\mathcal{P}_d) is iterated over in chronological order. If a patient is curative, a training example is created and added to the set of training examples. After each iteration, regardless of the patient's priority, the present capacity \hat{C}^ϕ is updated with the appointments of the corresponding patient from the offline solution s^* . The process continues until all patients in the patient flow are visited. Different regression models were tested (Appendix A), and XGBoost was chosen due to the low training time and high precision.

3.5. Using the regression model

The regression model is embedded in the prediction-based algorithm (Algorithm 2) to schedule RT treatments. When a new patient is admitted, there are two possibilities. If the patient is palliative, they will be scheduled on the first eligible date. Otherwise, an input vector is constructed as described in Section 3.4. The input vector is fed to the trained regression model to get a predicted waiting time, from which a predicted starting date is derived and utilized in the prediction-based algorithm to construct a schedule.

Algorithm 1: Generating training examples

input : instance ins , its offline solution s^*
output: set of training examples \mathcal{E}

```

1  $\mathcal{E} \leftarrow \{\}$  ;
2  $\phi \leftarrow$  the beginning of the first day in the simulation period  $\mathcal{K}$  ;
3 calculate  $\hat{C}^\phi$  ;
4 foreach day  $d$  in the simulation period  $\mathcal{K}$  do
5   foreach patient  $i$  in  $\mathcal{P}_d$  do
6     if  $i$  is curative then
7        $X_i \leftarrow \{\hat{C}^\phi, r_i, I_i, d_i, p_i\}$  ;
8        $w_i \leftarrow$  waiting time of  $i$  in  $s^*$  ;
9        $\mathcal{E} \leftarrow \mathcal{E} \cup (X_i, w_i)$  ;
10    end
11     $\phi \leftarrow \alpha_i$  ;
12    update  $\hat{C}^\phi$  with appointments of patient  $i$  ;
13  end
14 end
```

Algorithm 2: Prediction-based scheduling

input : instance ins , a trained regression model \mathcal{M}
output: schedule s

```

1 Initialize an empty schedule  $s$  ;
2  $\phi \leftarrow$  the beginning of the first day in the simulation period  $\mathcal{K}$  ;
3 calculate  $\hat{C}^\phi$  ;
4 foreach day  $d \in \mathcal{K}$  do
5   foreach patient  $i$  in  $\mathcal{P}_d$  do
6     if  $i$  is palliative then
7        $d^* \leftarrow$  the first eligible date for  $i$  starting from  $\max(d, r_i)$  ;
8     end
9     if  $i$  is curative then
10       $X_i \leftarrow \{\hat{C}^\phi, d, r_i, I_i, d_i, p_i\}$  ;
11       $\bar{w} \leftarrow$  predicted waiting time given by  $\mathcal{M}(X_i)$  ;
12       $d^* \leftarrow$  the first eligible starting date for  $i$  from  $\max(d + \bar{w}, r_i)$  ;
13    end
14     $s \leftarrow$  update the schedule by scheduling all fractions of  $i$  with the starting date  $d^*$  ;
15     $\phi \leftarrow \alpha_i$  ;
16    update  $\hat{C}^\phi$  with new appointments of patient  $i$  ;
17  end
18 end
```

Category	Proportion (%)	Treatment deadline (days)	Percentage of overdue treatment (%)	Average waiting time (days)
P1	0.44	1	14.29	1.09
P2	27.14	3	79.89	6.91
P3	41.36	14	74.55	18.11
P4	31.06	28	29.89	22.59

Table 3 Overdue treatment and average waiting times for cancer patients at CHUM in 2017-2018.

4. Numerical results

Three experiments were carried out. The first experiment (Section 4.3) examines the behaviour of the proposed approaches at different reservation rates. The second experiment (Section 4.4) evaluates the algorithms on generated data in different scenarios. Finally, the third experiment (Section 4.5) shows the results on real patient flow at CHUM. For each instance setting, 500 instances were generated, 400 of which were used for training the regression model, while the remaining 100 were used for testing. We compare the six scheduling strategies listed in Section 3.2. All of the algorithms are implemented in Python. We used CPLEX 20.1 as the MIP solver. The experiments were run on a Linux-based PC cluster. To generate offline solutions, each instance was solved using a single thread with a budget of 15 Gb of RAM, and a time budget of 15 hours.

4.1. Data overview

We have access to historical data from CHUM for a 2-year period, from September 2017 to July 2019, which includes 4,538 patients. CHUM is equipped with 12 treatment rooms, 5 of which contain identical generic linacs. 4 rooms are dedicated to special treatments that must be bound to those specific rooms. The remaining 3 rooms can be used for generic treatments, as well as a small percentage of specialized treatments. Patients are classified into four categories with different treatment deadlines. The majority of patients are curative (more than 70%) with 14-day or 28-day treatment deadlines. The remaining patients are mostly palliative type P2 with 3-day treatment deadlines. A small portion of patients (type P1) requires urgent treatment with a 1-day deadline. In the given period, more than 70% of P2 and P3 patients at CHUM did not meet their recommended treatment deadline. Patient categories along with their waiting time targets, percentage of overdue treatments, and average waiting time, are listed in Table 3. The fraction length of patients at CHUM ranges from 10 to 165 minutes, with the majority of patients having 25-minute fractions (more than 50%). The number of fractions varies from 1 to up to 45 sessions. Most palliative patients have less than 5 fractions, while the number of fractions of curative patients ranges from 1 to 35 sessions, with peaks at 5, 15, 20, 25, 30, and 33 sessions.

4.2. Data generation

Before describing the process of generating our instances from CHUM's real data, we introduce some technical terms. A *personal treatment plan*, which provides information such as the number of fractions and fraction length, is tailored for an individual patient by a team of physicians and specialists based on the patient's diagnosis and condition. For convenience, we use the term *treatment plan* to refer to a *personal treatment plan*. A *patient pool* consists of all historical treatment plans taken from CHUM's dataset. A problem instance consists of an *instance scenario* and a *patient flow*. An instance scenario specifies a number of linacs and their capacity, along with a set of appointments made during the previous batch of scheduling decisions. A patient flow is the daily flow of new patients admitted to the hospital during a *simulation period*. Given an arrival rate λ and a set of simulation days \mathcal{K} , a set of $|\mathcal{K}|$ numbers is generated using a Poisson distribution with an event rate of λ . For each day $d \in \mathcal{K}$, a set of virtual patients is generated by randomly selecting a corresponding number of treatment plans from the patient pool. Then a ready date is generated for each patient based on their category. The ready date for P1 patients is always the same as their admission date. P2 patients have ready dates ranging from 0 days after their admission date to 2 days after their admission date, while curative patients (P3 and P4) have ready dates ranging from 5 to 7 days after their admission date. The due date is calculated from the admission date and the patient's category listed in Table 3.

An instance scenario is generated as follows. First, we generate an instance flow, which then serves as an input for "warm-up" simulation using the greedy heuristic (Section 3.2.1) to fill up the linacs, starting from an empty schedule. The warm-up simulation stops when the occupancy of the linacs reaches 90% of its capacity on any day. The 90% threshold is based on the real schedules at CHUM. The following day is usually almost full, with a couple of available slots that normally account for no more than 10% of the total capacity. All appointments on the dates after the end of the warm-up period are used as the new instance scenario. An illustration is shown in Figure 4.

4.3. Tuning the reservation rates for palliative patients

The reservation rate plays a different role in the prediction-based approach than it does in myopic approaches. In myopic approaches, the reserved capacity makes space for future arrivals of palliative patients. A low reservation rate might lead to a shortage of treatment slots for urgent cases. In contrast, a high reservation rate does not fully exploit the treatment facility, which leads to delayed treatments for curative patients. In the prediction-based approach, an implicit "dynamic reservation" is applied by delaying the treatments of curative patients dynamically to make space

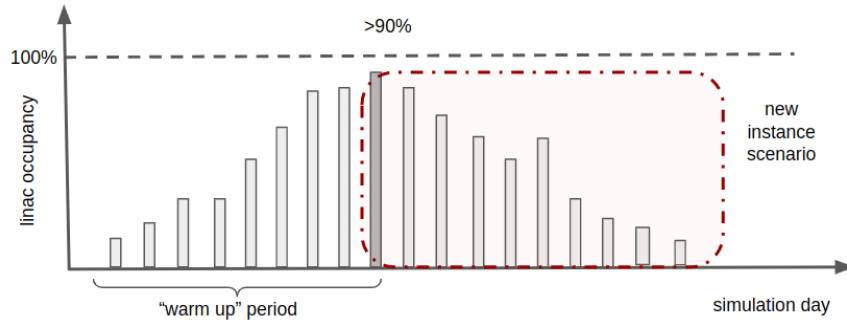


Figure 4 Illustration on how an instance scenario is generated.

for palliative ones based on the present allocation profile. Hence, the reserved capacity plays the role of a “safe corridor” to absorb the effect of unexpected high arrival. The goals of the experiment in this section are twofold: (1) to examine the behaviour of approaches with different reservation rates; and (2) to identify the reservation rate that yields the best performance for all approaches.

The experiment was carried out using a setting with 4 linacs and an arrival rate of 6.0. 100 instances were tested with five different reservation rates: 0%, 5%, 10%, 15%, and 20%. As there is no reservation in offline scheduling, five scheduling strategies were tested: online greedy heuristic, daily greedy heuristic, daily-IP, weekly-IP, and prediction-based. The results are shown in Figure 5. We can see from the figure that a higher reservation rate is indicative of a lower overdue time for palliative patients. However, the trade-off that comes with a low overdue rate, is a high number of overdue times for curative patients. Among all tested reservation rates, the prediction-based approach always gives lower overdue times in palliative patients by a large magnitude. Consider a reservation rate of 10%; other myopic approaches result in around 20 days of overdue time on average in palliative patients, while the prediction-based approach keeps that metric below 2 days. The prediction-based approach yields a slightly higher average overdue time for P3 compared to other approaches, while maintaining similar values for P4. A similar profile can be seen for other reservation rates. Nevertheless, one can observe that as the reservation rate increases, the overdue times for curative patients increase significantly when using myopic approaches, while the decrease for palliative patients is not proportional. From the aforementioned observations, we conclude that the prediction-based approach offers a dynamic reservation that outperforms the flat reservation policy of other myopic approaches, regardless of the reservation rate used. After discussing our results with a consultant from CHUM, we decided that a reservation rate of 10% results in the best trade-off between the delays in curative and palliative patients. This choice is further supported by the data showing that the total treatment duration of palliative patients

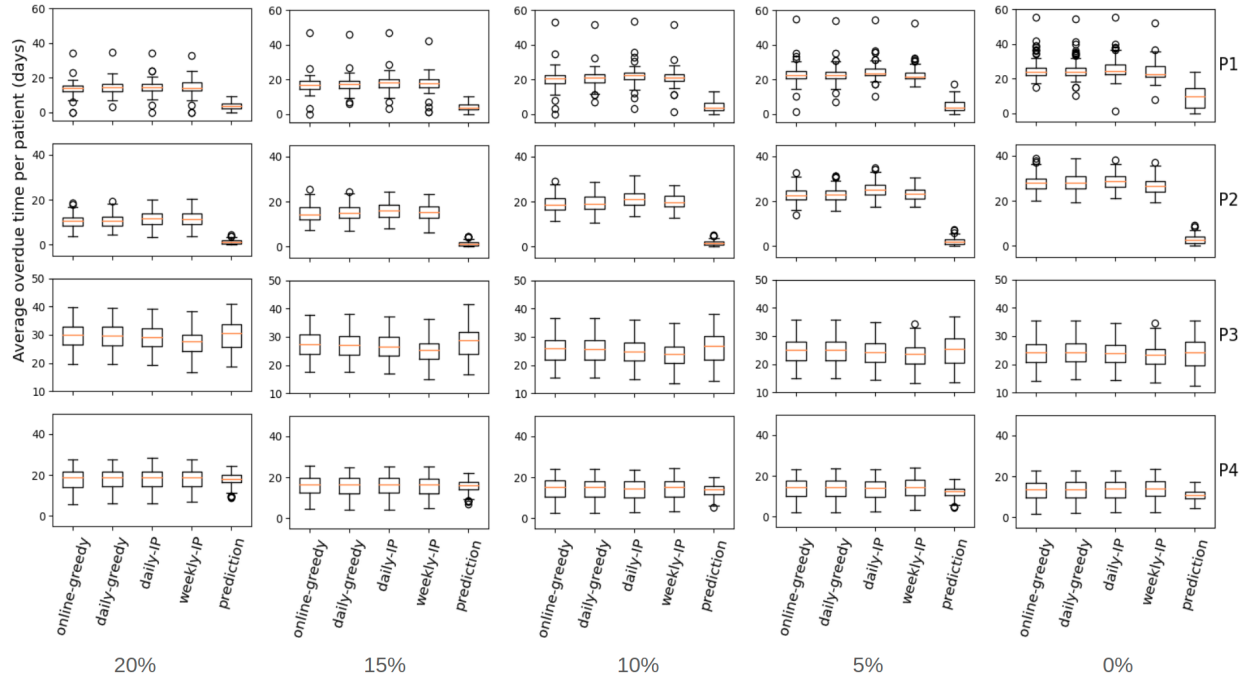


Figure 5 Average overdue time with different reservation rates - 4 linacs, arrival rate of 6.0.

accounts for around 8% of the total treatment duration for all patients. Therefore, we adopt a reservation rate of 10% for all scheduling approaches.

4.4. Numerical results on generated instances

To examine the algorithms' performance on different hospital sizes and crowding levels, we generated instances with a wide range of linac numbers and arrival rates. First, we identified a realistic range of arrival rates for a hospital with a given number of linacs using two types of *capacity simulation*. For the first capacity simulation, we ran the greedy algorithm in Section 3.2.1 without imposing the capacity constraints. We then measured the average weekly occupancy rate. Figure 6a shows the results of the first capacity simulation for 4 linacs with arrival rates of 5.0, 5.5, 6.0, and 6.5, respectively. The red dotted line represents the linac capacity. Arrival rates 6.0 and 6.5 result in overload for multiple weeks during the simulation period. An arrival rate of 5.0 is in the safe zone, with the occupancy rate remaining mostly below the linac capacity. However, an arrival rate of 5.5 can cause overload in some weeks, but generally stays within the safe zone.

For the second capacity simulation, we ran the greedy algorithm with capacity constraints and then observed the progression of the weekly average waiting time. If the average waiting time increased over time, the arrival rate was too high for a hospital of the given size to handle. A stable average waiting time indicates that the arrival rate is reasonable. Figure 6b shows the results of the

second capacity simulation for 4 linacs. Observe from the figure that an arrival rate of 6.5 causes the average waiting time to increase rapidly over time. This matches the observation from the first capacity simulation. The remaining arrival rates result in stable lines, with some peaks in the arrival rates of 5.5 and 6.0. The two capacity simulations complement each other, and we can conclude that arrival rates from 5.0 to 6.0 are reasonable for a hospital with 4 linacs.

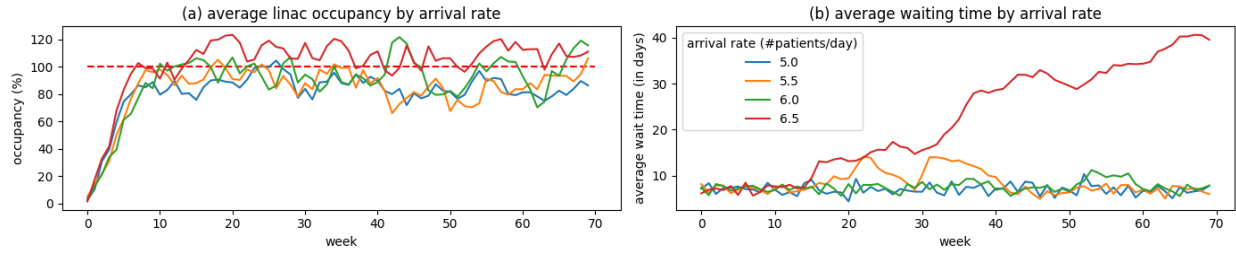


Figure 6 Simulation on 4 linacs with arrival rates of 5.0, 5.5, 6.0 and 6.5.

After identifying a reasonable range of arrival rates, we performed experiments using our realistic instance settings. Figure 7 shows the results for 4 linacs with an arrival rate of 5.0. Observe from the figure that in an ideal situation where all future arrivals are known in advance, the overdue time is insignificant in all patient categories. The greedy heuristic and the three batch scheduling strategies are similar in both waiting time and overdue time. Overdue treatments mostly occur in P2 and P3 patients. The prediction-based approach shows a slightly better average waiting time and overdue time in all patient categories. It also results in less deviation in waiting time and overdue time between patients. For each priority group, we performed an one-way ANOVA test. The results confirm that the difference between the tested algorithms are statistically significant, with p-values in all four priority groups less than $1E-5$. The post hoc paired t-tests confirm that the results for the prediction-based approach for P1, P2, and P4 are substantially different from the results for the other four approaches. However, in priority group P3, there is no significant difference between the prediction-based approach and the weekly-IP-based approach (p-value = 0.33).

Figure 8 shows results for the same hospital size in a more crowded scenario with an arrival rate of 6.0. We can observe from the figure that the more crowded a hospital is, the better the prediction-based approach performs compared to batch scheduling strategies and the greedy heuristic. The most significant difference is for P2 patients, the most challenging category for other algorithms. As P2 patients account for almost 30% of patients, and their treatment deadlines are urgent (only 3 days after their admission date), it is difficult to decide how much space should be reserved for their

arrivals when scheduling curative patients. Despite having less information than batch scheduling strategies, the prediction-based approach demonstrates an ability to learn good decision-making from offline schedules. Specifically, it learns to dynamically delay treatment for curative patients in order to make space for palliative treatments based on the present state of the schedule. This explains the lower overdue rates associated with prediction-based scheduling in palliative patients.

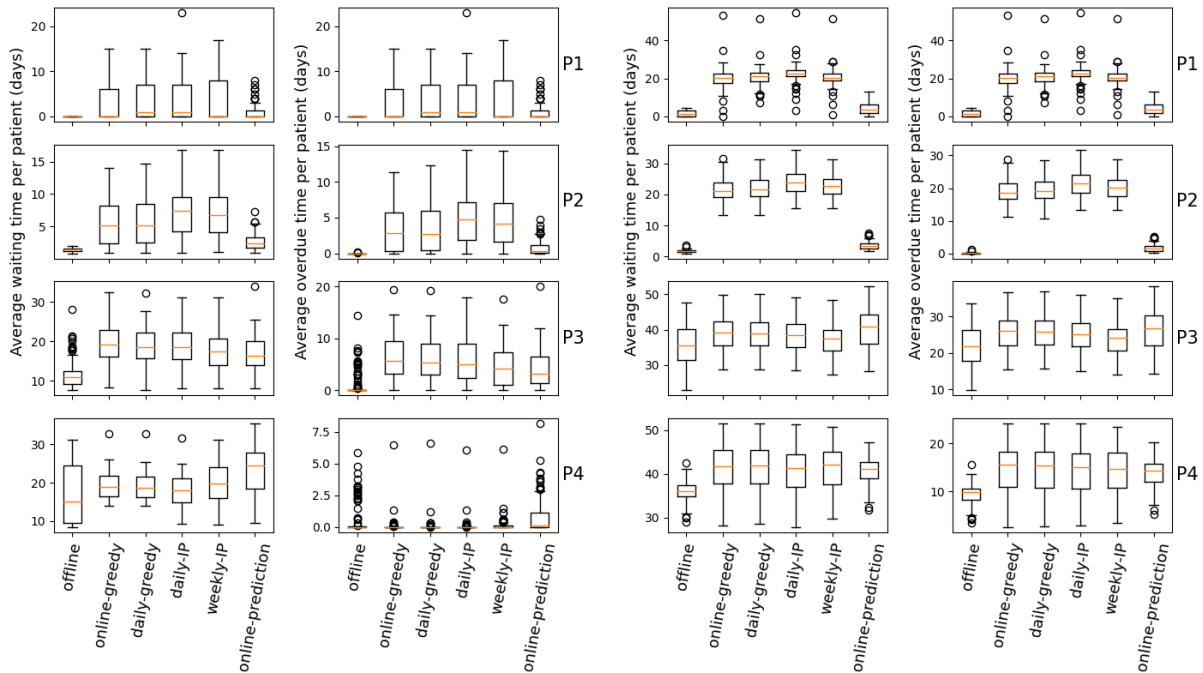


Figure 7 Simulation on 4 linacs - arrival rate 5.0. **Figure 8** Simulation on 4 linacs - arrival rate 6.0.

For the next experiment, we tested the algorithms on a much larger hospital with 8 linacs. The reasonable arrival rates suggested by the capacity simulations are less than 12 patients per day, as an arrival rate of 13.0 causes the weekly average waiting time to increase over time (Figure 9). We tested two settings with arrival rates of 10.0, and 12.0. The results are shown in Figures 10, and 11, respectively. We derive two observations from these results. First, the experiment matches our earlier observation from experiments using 4 linacs that the larger and more crowded a hospital is, the better the prediction-based approach performs compared to other scheduling strategies. Second, our approach scales well on large instances. To our knowledge, most experiments in the literature tested their ideas on much smaller hospital sizes. Furthermore, most real hospitals do not have more than 8 linacs operating full-time. In large hospitals, as the number of linacs and patients grows, the scheduling task becomes extremely challenging for most algorithms. Our prediction-based approach offers a fast online approach with superior results.

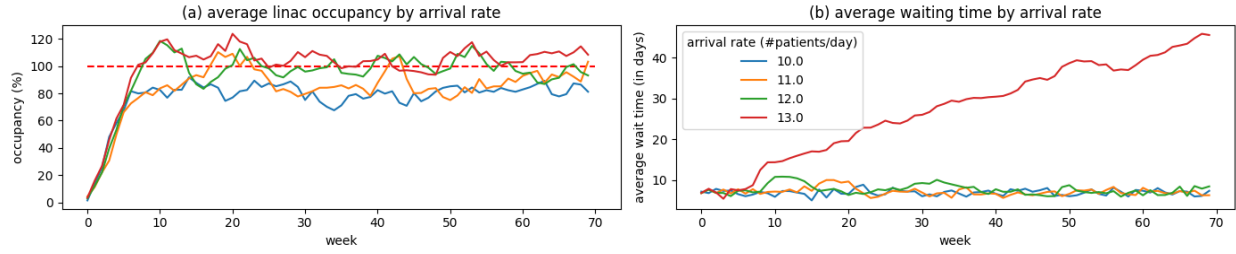


Figure 9 Capacity simulation on 8 linacs with arrival rates of 10, 11, 12 and 13. Starting from an arrival rate of 13, the average waiting time of patients increases over time.

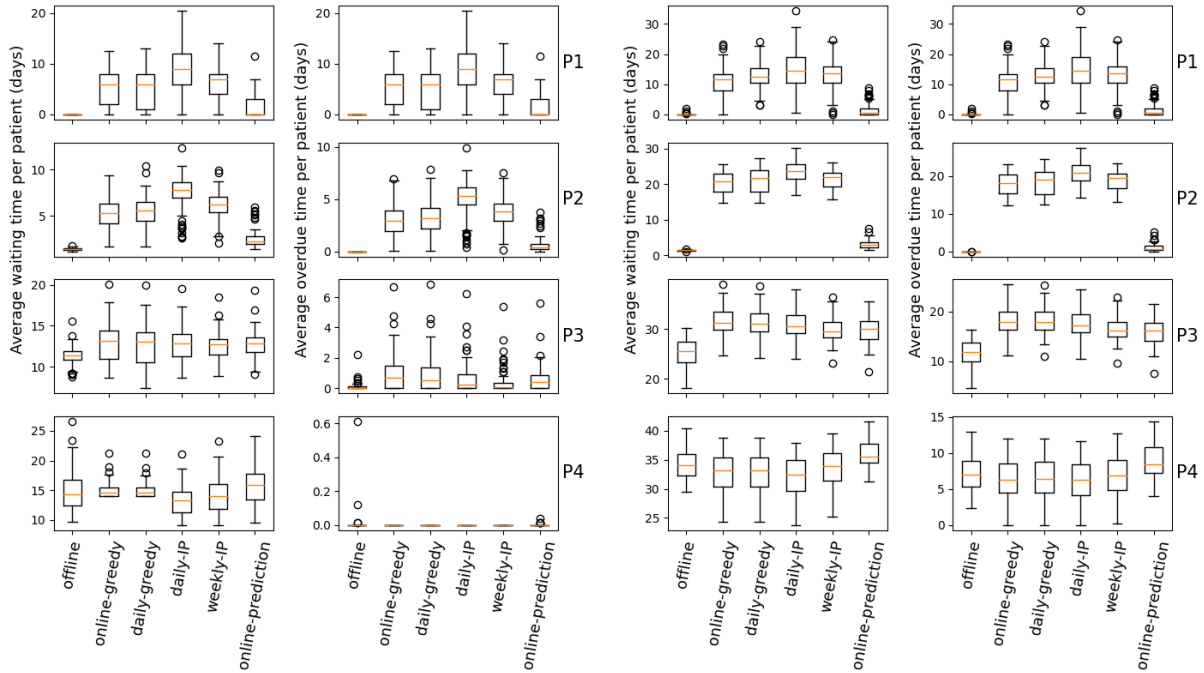


Figure 10 Simulation on 8 linacs - arrival rate 10.0. **Figure 11** Simulation on 8 linacs - arrival rate 12.0.

4.5. Numerical results on real instances

In a real-world scenario, we might observe a high fluctuation in arrival rate. In this experiment, we aim to test the model's sensitivity to real patient flow with an unstable arrival rate. We extract the real patient flow for the period from May 2018 to June 2019. As mentioned in Section 4.1, CHUM's facility consists of 12 linacs, 3 of which are for special treatments. Therefore, we remove all patients bound to special linacs, which accounts for 9% of patients. Among the remaining 9 linacs, only 5 generic linacs operate full-time (10 hours a day). The remaining 4 linacs operate only half of each working day on average due to a lack of technicians. We construct real instances using the extracted patient flow and 7 linacs operating full-time. Then we use the first 100 days to initialize a partially-filled schedule using the method described in Section 4.2. Finally, the

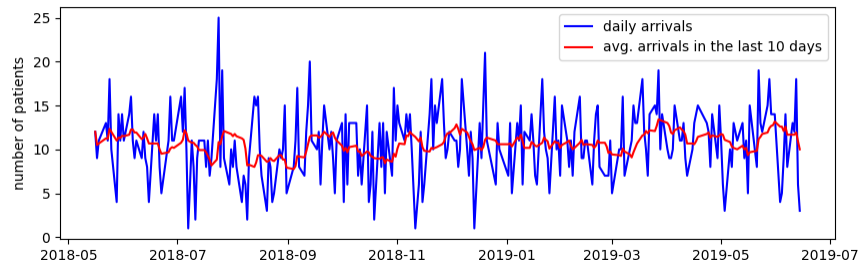


Figure 12 Real patient flow at CHUM (from May 2018 to June 2019).

remaining days of data are used as the patient flow for the instance. The patient flow of the real instance consists of patient arrivals for 180 days, from October 2018 to June 2019. In Figure 12, we plot the real arrival rates at CHUM for the period from May 2018 to June 2019. The blue line represents the real daily arrivals of patients, and the red line represents the average arrival rate of a 10-day interval prior to a given date. Observe from the figure that the average arrival rate fluctuates significantly, reaching 7.75 patients/day at its lowest, and reaching 13.42 patients/day at its highest.

The regression model for the real instance is trained on the setting with 7 linacs and an arrival rate of 10.1, which is the average arrival rate in the first 3 months of the data. Similar to Section 4.4, we ran all proposed scheduling strategies on 100 test instances. The results are presented in Figure 13. Similar to the results from generated instances in Section 4.4, the results here show that the prediction-based approach outperforms other approaches in palliative patients while maintaining a comparable waiting time and overdue time for curative patients. Observed from Figure 12 that the arrival rate fluctuates by time in real life, we also train a regression model on a variable arrival rate. The training instances are generated with a dynamic arrival rate that changes every 10 days within the range of $[\lambda - 1.5, \lambda + 1.5]$, with λ being the average arrival rate (10.1 patients per day in our case). The parameters of the data generator are chosen based on the observation of the real arrival rate's fluctuation (the red line in Figure 12). The 10 days-interval and the deviation of ± 1.5 in arrival rate mimic the fluctuation of the real data while preventing sudden changes in arrival rate. The online prediction-based approach using the variable arrival rate is labeled as *prediction-based**. As the fluctuated model is trained on a different dataset, we do not include its results in Figure 13.

Then we test the trained regression model on the real patient flow. Due to the long simulation period, the result of the offline scheduling is not available. We first analyse the results given by the prediction-based approach with a stable arrival rate. The results presented in Table 4 and Figure 14 show that the prediction-based approach results in the lowest average overdue time in both

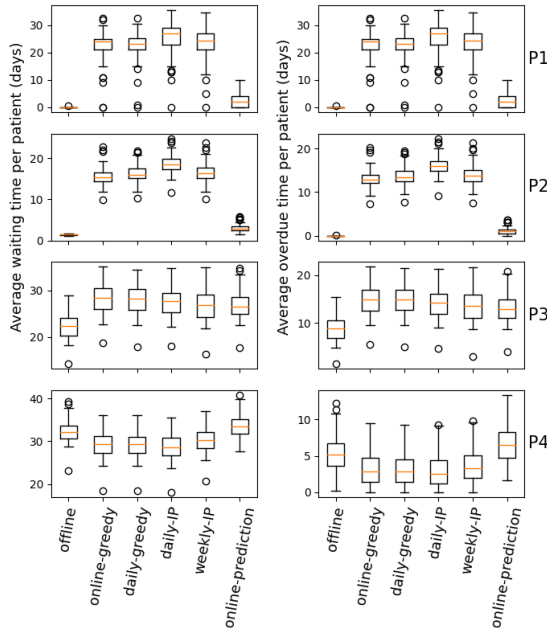


Figure 13 Simulation on 7 linacs, arrival rate 10.1.

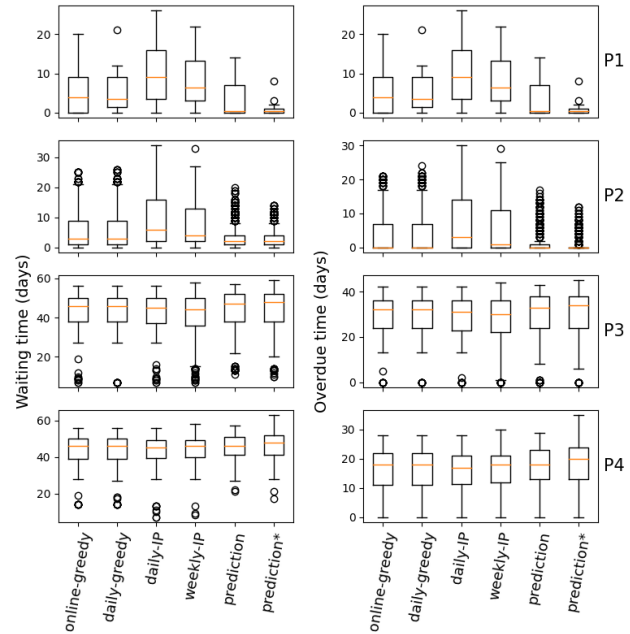


Figure 14 Simulation on the real instance.

P1 (3.29 days) and P2 (1.99 days). The greedy algorithms (both online and daily) perform relatively well on palliative patients, but result in 2 extra days of average overdue time compared to the prediction-based approach. Statistical tests (one-way ANOVA and post-hoc paired t-test) confirm the differences between the results of the tested algorithms on P2 (all p-values are less than $7E-31$) are statistically significant. On P1, even though the one-way ANOVA finds no significant differences in the results (p-value= 0.073), the statistical differences between the prediction-based approach and the other four approaches are significant. The paired t-tests between the results on P1 given by the prediction-based approach and other approaches all have p-values of less than 0.05. In curative patients, we observe little difference in the results. The average overdue time in P3 given by the prediction-based approach is about one day higher than other approaches. This small delay is compensated for by the large reduction in overdue treatments of palliative patients.

Looking at the results given by the prediction-based approach trained on a variable arrival rate (prediction-based*), we observe that the approach reduces even further the overdue time in palliative patients (1.21 days in P1 and 1.0 in P2, compared to 3.29 days in P1 and 1.99 days in P2 given by the model trained on a stable arrival rate). The box-plots in Figure 14 shows that there are also less outliers in overdue time of palliative patients. However, there is a trade-off of a higher overdue time in curative patients (0.27 days higher in P3 and 0.86 days higher in P4) compared to the prediction-based approach with a stable arrival rate. We conclude that the prediction approach

Scheduling strategy	avg. occupancy	Waiting time (days)					Overdue time (days)				
		overall	P1	P2	P3	P4	overall	P1	P2	P3	P4
online-greedy	97.45	33.02	5.14	6.13	3.67	44.02	17.80	5.14	3.91	29.74	16.18
daily-greedy	97.51	32.91	6.00	6.23	43.48	43.80	17.71	6.00	3.99	29.58	16.00
daily-IP	97.72	33.53	9.79	9.63	42.87	43.44	18.25	9.79	7.15	28.93	15.65
weekly-IP	97.61	33.04	7.86	7.72	42.42	44.10	17.76	7.86	5.37	28.51	16.19
prediction-based	97.14	32.93	3.29	4.05	44.21	44.94	17.69	3.29	1.99	30.22	16.96
prediction-based*	96.89	33.02	1.21	3.03	44.46	45.80	17.79	1.21	1.00	30.49	17.82

Table 4 Average waiting time of patients for different scheduling strategies on the real instance.

shows its robustness in real-life scenarios where there is a high fluctuation in patient arrivals. This allows us to avoid the cost of retraining the model, as that would be necessary only when hospitals upgrade their infrastructure or experience a significant change in arrival rates.

4.6. Explainable decision-making with SHAP value

In this section, we use SHAP (SHapley Additive exPlanations) to analyse the regression model used in our prediction-based approach, and show how it can assist explainability in scheduling decisions. SHAP values represent the relative impact of a variable on the outcome. To demonstrate the methodology for our analysis, we use the regression model for the instance setting with 6 linacs and an arrival rate of 9.0. We first look at a global interpretation of the model, which gives us an overview of the impact each feature has on the decision. A beeswarm plot (Figure 15) is an information-dense summary of how top features impact the model's output. Each point corresponds to an observation in the training set. The y-axis indicates the feature name, and the higher a feature is, the more impact it has on the output. The x-axis indicates the SHAP values of the data points. All the values to the left of the centre line represent the observations that shift the output in the negative direction (reducing the prediction of waiting time), while the points on the right shift the prediction in a positive direction (increasing the waiting time). The colors represent the values of the features from low to high. Observe from Figure 15 that for the given model, the number of sections has the highest impact on the decision. A high number of sections has a positive impact on the output, i.e. a high number of sections is likely to result in a large waiting time and vice versa. The contribution of the remaining features can be interpreted similarly, e.g, a large due day has a high and positive impact on the waiting time (patients with high due days have more waiting times), while a low remaining capacity on day 49 results in a higher waiting time.

To explain an individual scheduling decision we need to look at a plot that offers local interpretability. A waterfall plot visualizes how the feature values contribute to the final prediction of an individual observation. Two examples of waterfall graphs for two different patients can be found in Figures 16 and 17. The x-axis shows the value of the output (number of waiting days

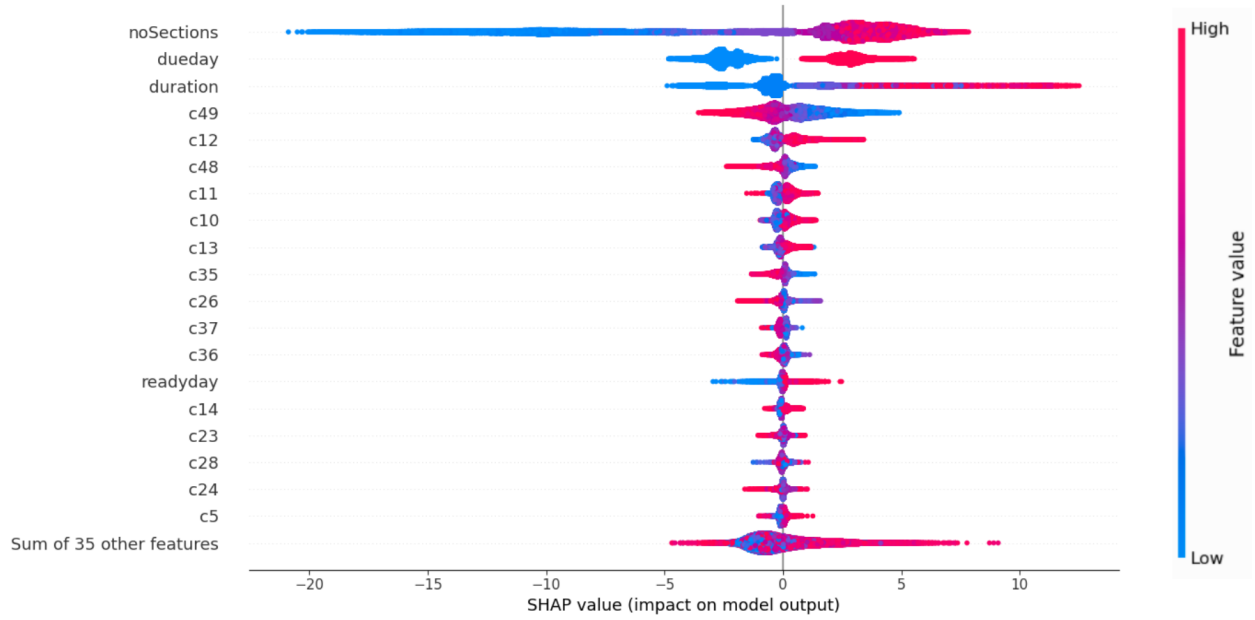


Figure 15 Beeswarm plot for the instance setting with 6 linacs and an arrival rate of 9.0.

in our problem). $E[f(x)]$ is the base value, which is generally the average output of the training set. To explain a decision, we start from the base value at the bottom of the plot, then add up the contribution of each feature (positive for red and negative for blue) one row at a time, until we reach the actual model output $f(x)$ at the top of the plot. The gray text before the feature names shows the value of each feature for this sample. Thus, the waterfall plot explains how much each feature contributes to the final decision.

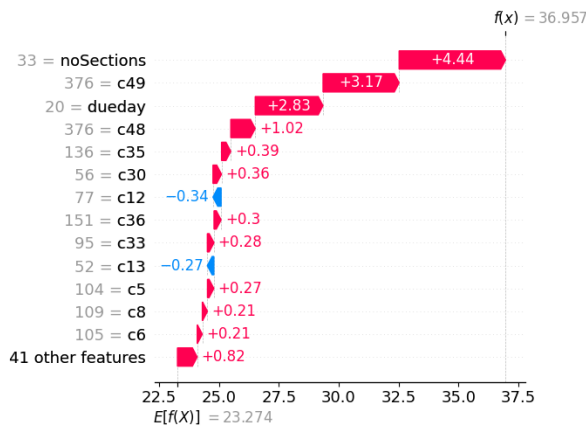


Figure 16 The waterfall plot for a P4-patient.

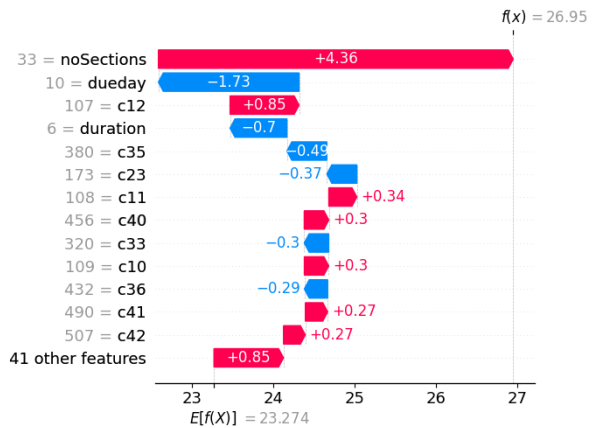


Figure 17 The waterfall plot for a P3-patient.

5. Discussion and future work

In this section, we discuss the potential impact and challenges of our implementation in the real world, as well as possible future research directions. Currently, our industrial partner, the oncology department of CHUM, is collaborating with Gray³ (a patient scheduling start-up), to develop an application for their scheduling tasks. The goal of the collaboration is to automate the scheduling tasks that are currently done manually. We have been actively collaborating with Gray from the beginning. The scheduling system is based on a modified IP model. It was deployed in early 2022, and Gray is constantly using feedback from CHUM to improve the system. The system is still myopic, i.e. it does not take stochastic information into account when making scheduling decisions. Overdue treatments, especially in P2 and P3 patients, remain a big challenge. We hope to be able to conduct tests using the predictive model in the second half of 2022. Supporting explainable decisions is another advantage of our approach. From our experience deploying the system, building confidence is the most challenging task, as it is difficult to convince the scheduling staff to “trust” the algorithm, especially when decisions are human-sensitive. Therefore, transparent scheduling systems are highly appreciated by the industry.

When implementing a machine learning-based approach for a real world application, it is important to consider computing costs. In our approach, the most expensive cost is associated with computing offline solutions. For our largest instance size (8 linacs, arrival rates of 12.0), the average run time is 47,522 seconds (~ 13 hours). Generating offline solutions for 500 instances requires about 7000 hours of CPU time. It is our opinion that this is both reasonable and feasible for large hospitals, and can be done easily with a computer cluster or server. Additionally, the regression model adapts well to fluctuation in arrival rates and does not need to be retrained often, as shown in Section 4.5. However, for hospitals of larger size than the one in this study, generating offline solutions might become a problem. Moving forward, we plan to apply unsupervised learning to train the regression model and overcome this obstacle.

Patient scheduling consists of allocation scheduling and appointment scheduling. Solving both problems at once is difficult due to the high complexity. Studies that consider both problems either (1) decompose the problem into 2 phases (Pham et al. 2021, Vieira et al. 2020) where appointment scheduling is solved either in the second phase or as a post-processing step, or (2) use (meta)heuristics (Maschler et al. 2016, Maschler and Raidl 2018, Vogl et al. 2019). However,

³<https://www.graysuite.com/>

those approaches do not guarantee the best performance. It is an open question if compromising some objectives in allocation scheduling would lead to a better solution in appointment scheduling in terms of the number of fixed appointments that were rescheduled, patient satisfaction, etc. We plan to investigate this in future research.

6. Conclusions

Scheduling radiotherapy treatment is difficult due to the stochastic arrival rate of patients with different priorities who require multiple appointments. In this paper, we proposed an online machine learning-based approach that makes scheduling decisions dynamically based on the present allocation profile of the hospital. The algorithm was evaluated on a rolling horizon and compared with several myopic approaches, both heuristic-based and IP-based. The test instances were generated based on real data from CHUM to simulate different hospital sizes and crowding levels. The results show that our prediction-based approach outperforms other approaches. The prediction-based approach results in a lower average overdue time and lower average waiting time for palliative patients when compared with myopic strategies, while maintaining the same performance for curative patients. Our approach also scales decently with instance size, and we successfully solve instances with up to 8 linacs. We proved the robustness of our prediction-based approach in a real problem instance from CHUM, despite the high fluctuation of arrival rates observed in reality. Finally, we demonstrated how our approach supports explainability and interpretability.

References

- Barton MB, Jacob S, Shafiq J, Wong K, Thompson SR, Hanna TP, Delaney GP (2014) Estimating the demand for radiotherapy from the evidence: a review of changes from 2003 to 2012. *Radiotherapy and oncology* 112(1):140–144.
- Burke EK, Leite-Rocha P, Petrovic S (2011) An integer linear programming model for the radiotherapy treatment scheduling problem. *arXiv preprint arXiv:1103.3391*.
- Castro E, Petrovic S (2012) Combined mathematical programming and heuristics for a radiotherapy pre-treatment scheduling problem. *Journal of Scheduling* 15(3):333–346.
- Cayirli T, Veral E (2003) Outpatient scheduling in health care: a review of literature. *Production and operations management* 12(4):519–549.
- Ceschia S, Schaerf A (2012) Modeling and solving the dynamic patient admission scheduling problem under uncertainty. *Artificial intelligence in medicine* 56(3):199–205.
- Ceschia S, Schaerf A (2016) Dynamic patient admission scheduling with operating room constraints, flexible horizons, and patient delays. *Journal of Scheduling* 19(4):377–389.

- Chen Z, King W, Pearcey R, Kerba M, Mackillop WJ (2008) The relationship between waiting time for radiotherapy and clinical outcomes: a systematic review of the literature. *Radiotherapy and Oncology* 87(1):3–16.
- Coles C, Burgess L, Tan L (2003) An audit of delays before and during radical radiotherapy for cervical cancer—effect on tumour cure probability. *Clinical oncology* 15(2):47–54.
- Conforti D, Guerriero F, Guido R (2008) Optimization models for radiotherapy patient scheduling. *4OR* 6(3):263–278.
- Conforti D, Guerriero F, Guido R (2010) Non-block scheduling with priority for radiotherapy treatments. *European Journal of Operational Research* 201(1):289–296.
- Frimodig S, Schulte C (2019) Models for radiation therapy patient scheduling. *International Conference on Principles and Practice of Constraint Programming*, 421–437 (Springer).
- Gerchak Y, Gupta D, Henig M (1996) Reservation planning for elective surgery under uncertain demand for emergency surgery. *Management Science* 42(3):321–334.
- Gocgun Y (2018) Simulation-based approximate policy iteration for dynamic patient scheduling for radiation therapy. *Health care management science* 21(3):317–325.
- Gupta D, Denton B (2008) Appointment scheduling in health care: Challenges and opportunities. *IIE transactions* 40(9):800–819.
- Kapamara T, Petrovic D (2009) A heuristics and steepest hill climbing method to scheduling radiotherapy patients. *Proceedings of the 35th International Conference on Operational Research Applied to Health Services, Catholic University of Leuven, Belgium* (Citeseer).
- Kapamara T, Sheibani K, Haas OC, Reeves CR, Petrovic D (2006) A review of scheduling problems in radiotherapy. *Proceedings of the Eighteenth International Conference on Systems Engineering (ICSE2006), Coventry University, UK*, 201–207.
- Lamiri M, Xie X, Dolgui A, Grimaud F (2008) A stochastic model for operating room planning with elective and emergency demand for surgery. *European Journal of Operational Research* 185(3):1026–1037.
- Larsen E, Lachapelle S, Bengio Y, Frejinger E, Lacoste-Julien S, Lodi A (2021) Predicting tactical solutions to operational planning problems under imperfect information. *INFORMS Journal on Computing*.
- Legrain A, Fortin MA, Lahrichi N, Rousseau LM (2015) Online stochastic optimization of radiotherapy patient scheduling. *Health care management science* 18(2):110–123.
- Mackillop WJ (2007) Killing time: the consequences of delays in radiotherapy. *Radiotherapy and Oncology* 84.1 1(4).
- Marynissen J, Demeulemeester E (2019) Literature review on multi-appointment scheduling problems in hospitals. *European Journal of Operational Research* 272(2):407–419.
- Maschler J, Raidl GR (2018) Particle therapy patient scheduling with limited starting time variations of daily treatments. *International Transactions in Operational Research*.
- Maschler J, Riedler M, Stock M, Raidl GR (2016) Particle therapy patient scheduling: First heuristic approaches. *Proceedings of the 11th Int. Conference on the Practice and Theory of Automated Timetabling. Udine, Italy*.

- Patrick J, Puterman ML, Queyranne M (2008) Dynamic multipriority patient scheduling for a diagnostic resource. *Operations research* 56(6):1507–1525.
- Petrovic S, Leite-Rocha P (2008) Constructive approaches to radiotherapy scheduling. *Proceedings of the World Congress on Engineering and Computer Science*, 722–727.
- Petrovic S, Leung W, Song X, Sundar S (2006) Algorithms for radiotherapy treatment booking. *25th Workshop of the UK planning and scheduling special interest group*, 105–112.
- Pham TS, Rousseau LM, De Causmaecker P (2021) A two-phase approach for the radiotherapy scheduling problem. *Health Care Management Science* 1–17.
- Salah H, Srinivas S (2022) Predict, then schedule: Prescriptive analytics approach for machine learning-enabled sequential clinical scheduling. *Computers & Industrial Engineering* 108270.
- Sauré A, Begen MA, Patrick J (2020) Dynamic multi-priority, multi-class patient scheduling with stochastic service times. *European Journal of Operational Research* 280(1):254–265.
- Saure A, Patrick J, Tyldesley S, Puterman ML (2012) Dynamic multi-appointment patient scheduling for radiation therapy. *European Journal of Operational Research* 223(2):573–584.
- Schuetz HJ, Kolisch R (2012) Approximate dynamic programming for capacity allocation in the service industry. *European Journal of Operational Research* 218(1):239–250.
- Srinivas S, Ravindran AR (2018) Optimizing outpatient appointment system using machine learning algorithms and scheduling rules: a prescriptive analytics framework. *Expert Systems with Applications* 102:245–261.
- Tyldesley S, Delaney G, Foroudi F, Barbera L, Kerba M, Mackillop W (2011) Estimating the need for radiotherapy for patients with prostate, breast, and lung cancers: verification of model estimates of need with radiotherapy utilization data from british columbia. *International Journal of Radiation Oncology* Biology* Physics* 79(5):1507–1515.
- Vermeulen IB, Bohte SM, Elkhuisen SG, Lameris H, Bakker PJ, La Poutre H (2009) Adaptive resource allocation for efficient patient scheduling. *Artificial intelligence in medicine* 46(1):67–80.
- Vieira B, Demirtas D, van de Kamer JB, Hans EW, Rousseau LM, Lahrichi N, van Harten WH (2020) Radiotherapy treatment scheduling considering time window preferences. *Health care management science* 1–15.
- Vieira B, Hans EW, van Vliet-Vroegindeweij C, Van De Kamer J, Van Harten W (2016) Operations research for resource planning and-use in radiotherapy: a literature review. *BMC medical informatics and decision making* 16(1):1–11.
- Vogl P, Braune R, Doerner KF (2019) Scheduling recurring radiotherapy appointments in an ion beam facility. *Journal of Scheduling* 22(2):137–154.
- Zhu YH, Toffolo TA, Vancroonenburg W, Berghe GV (2019) Compatibility of short and long term objectives for dynamic patient admission scheduling. *Computers & Operations Research* 104:98–112.

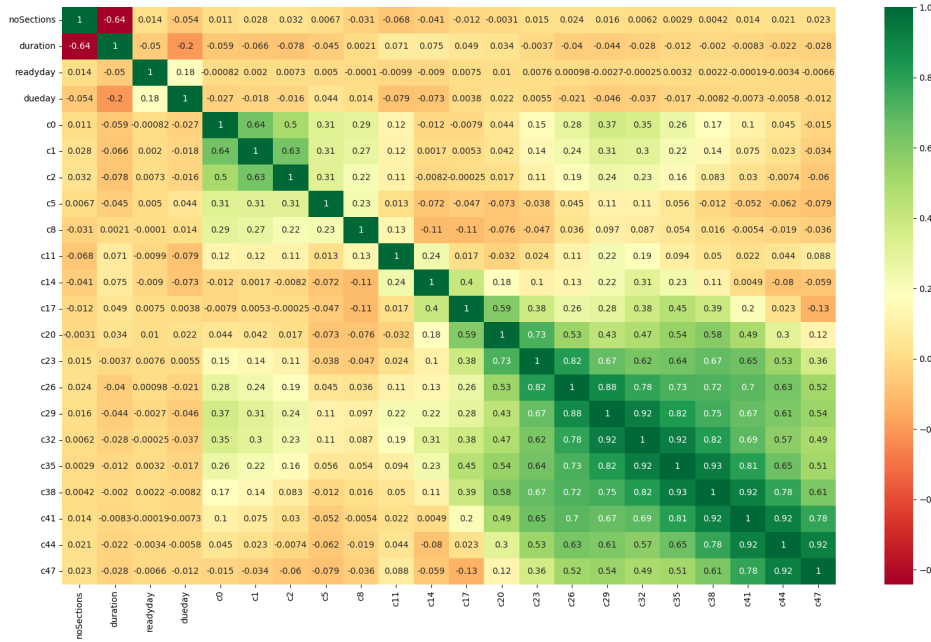


Figure 18 The correlation heatmap of input features.

Appendix A: Training regression models

Before training the regression model, we test the collinearity of the features to verify if there are highly correlated features that should be removed. A correlation heatmap can be found in Figure 18. We can see from the figure that there is a certain correlation between the available capacity of consecutive days (the capacity reduces gradually over time). We plotted only one-third of the considered days. However, we consider that information is relevant to the decision.

After confirming the independence of input features, different regression models are tested. We present the results on different instance settings, from easiest to hardest. Eight regression models were tested. For each model and instance setting, we report the training time, mean squared error (MSE), mean absolute error (MAE), and R-squared value on both the training set and the testing set.

Three instance settings were tested, from the easiest to the most difficult:

- 4 linacs, arrival rate of 5.0
- 6 linacs, arrival rate of 9.0
- 8 linacs, arrival rate of 12.0

The training results for the three settings are presented in Table 5, 6, and 7, respectively. Random Forest and XGBoost consistently gave the best results on both the training set and the testing set. However, Random Forest requires more time for training, and it performs worse in difficult settings, i.e., when the size of the training test increases (69.12, 212.88, and 294.18 seconds in instance settings of 4, 6, and 8 linacs, respectively). Meanwhile, XGBoost takes only 6.56 seconds in 4 linacs and 2.86 seconds in 8 linacs. Random Forest gives the lowest error rates, but the MAE in the testing set given by Random Forest and XGBoost are very close. A lower differences between the error rates in the

	Training				Testing		
	Training time	MSE	MAE	R-squared	MSE	MAE	R-squared
MLP	29.7	6.28	1.8	0.81	5.94	1.76	0.82
SGD	0.15	12.82	2.78	0.61	12.38	2.75	0.62
Lasso	0.96	12.79	2.78	0.61	12.36	2.75	0.62
ElasticNet	0.72	13.06	2.81	0.6	12.62	2.77	0.61
Decision Tree	1.22	3.6	0.59	0.89	9.85	1.8	0.7
Random forest	69.12	0.6	0.49	0.98	4.43	1.35	0.86
XGBoost	6.56	1.5	0.85	0.95	4.3	1.34	0.87

Table 5 Training results for instance setting with 4 linacs, and an arrival rate of 5.0.

	Training				Testing		
	Training time	MSE	MAE	R-squared	MSE	MAE	R-squared
MLP	139.93	10.55	2.34	0.85	10.31	2.29	0.85
SGD	0.65	24.6	3.68	0.66	23.71	3.61	0.65
Lasso	1.65	24.31	3.66	0.66	23.4	3.59	0.66
ElasticNet	1.13	24.51	3.66	0.66	23.56	3.58	0.66
Decision Tree	3.54	8.11	1.16	0.89	14.42	2.13	0.79
Random forest	212.88	0.92	0.56	0.99	6.7	1.55	0.9
XGBoost	2.37	3.45	1.26	0.95	6.78	1.64	0.9

Table 6 Training results for instance setting with 6 linacs, and an arrival rate of 9.0.

	Training				Testing		
	Training time	MSE	MAE	R-squared	MSE	MAE	R-squared
MLP	159.35	7.27	1.91	0.87	7.46	1.94	0.87
SGD	0.89	18.93	3.25	0.67	19.34	3.27	0.67
Lasso	1.3	18.88	3.23	0.67	19.26	3.25	0.67
ElasticNet	0.57	19.03	3.23	0.67	19.45	3.26	0.67
Decision Tree	4.82	5.49	0.9	0.9	10.11	1.76	0.83
Random forest	294.18	0.57	0.44	0.99	4.36	1.27	0.93
XGBoost	2.86	2.57	1.08	0.95	4.54	1.37	0.92

Table 7 Training results for instance setting with 8 linacs, and an arrival rate of 12.0.

training test and the setting test from XGBoost also indicate that it is less prone to overfitting compared to Random Forest. The R-squared values given by XGBoost in the three test cases are approximate 90%, which shows that the data fits the regression model well. Based on the above analysis, we choose XGBoost as the regression model for our algorithm.

Appendix B: Additional results on generated instances

In this section, we provide the results for additional instance settings with a hospital size of 6 linacs. The realistic arrival rates identified through the capacity simulations are between 7.0 and 9.0 (see Figure 19).

The waiting times and overdue times given by the proposed algorithms using those settings are shown in Figures 20, and 21, respectively. Observe that the differences between the algorithms are more profound than in the previous smaller settings with 4 linacs, especially when the arrival rate increases. The prediction-based approach offers better waiting times and overdue times for palliative patients (both P1 and P2) than scheduling methods, while maintaining a comparable waiting time for curative patients. Similar to the case with 4 linacs, the more crowded a hospital is, the better the prediction-based approach performs over other approaches. Another observation from the boxplots is that the average waiting time (and overdue time) of P4 patients in greedy heuristic and batch scheduling is lower than the others, which also leads to higher overdue time in palliative patients. This shows that the greedy heuristic and batch scheduling

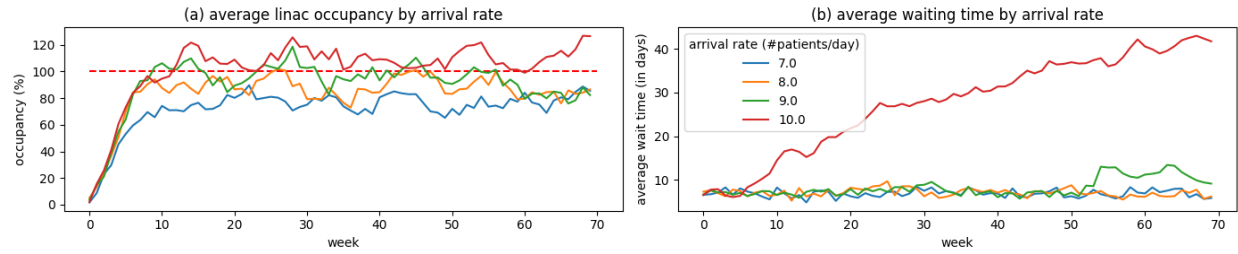


Figure 19 Capacity simulation on 6 linacs with arrival rates of 7.0, 8.0, 9.0 and 10.0. Starting from an arrival rate of 10.0, the average waiting time of patients increases over time.

are not effective in delaying curative treatments to accommodate palliative ones. The results are in line with the results in Section 4.4, and confirm that the prediction-based approach performs better in large and crowded hospitals.

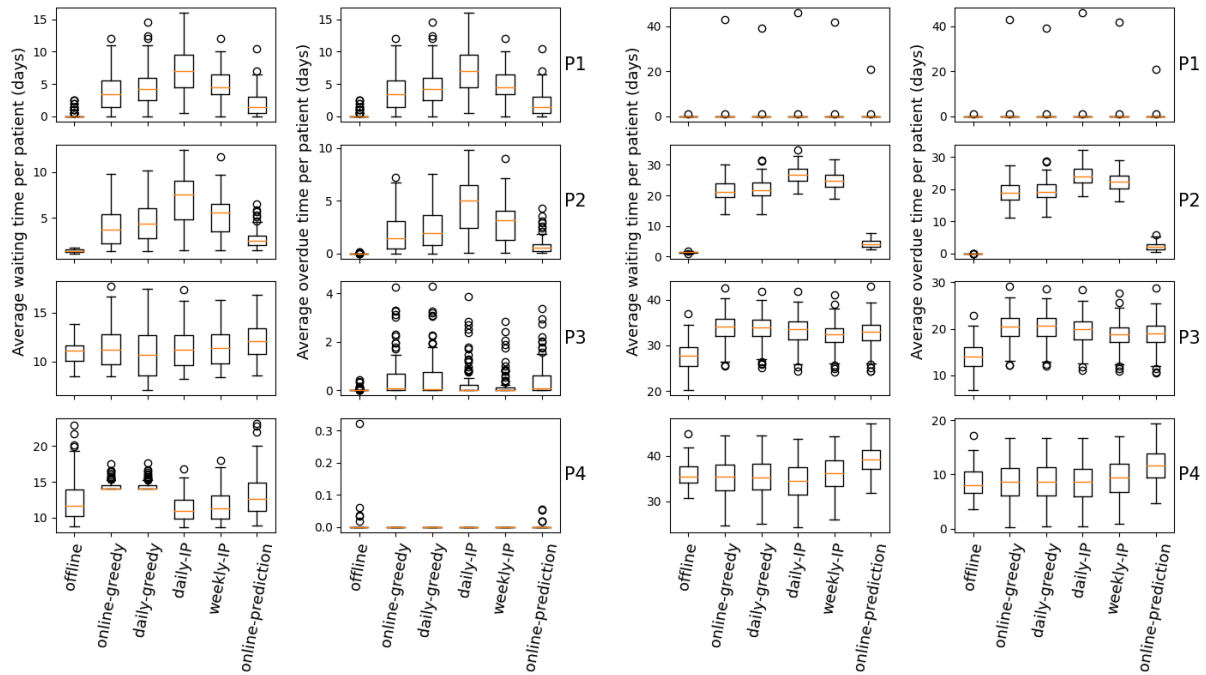


Figure 20 Simulation on 6 linacs - arrival rate 7.0. **Figure 21** Simulation on 6 linacs - arrival rate 9.0.