Proceedings of the 2021 Winter Simulation Conference
S. Kim, B. Feng, K. Smith, S. Masoud, Z. Zheng, C. Szabo, and M. Loper, eds.

SUPPORTING EFFICIENT ASSIGNMENT OF MEDICAL RESOURCES IN CANCER TREATMENTS WITH SIMULATION-OPTIMIZATION

Leandro do C. Martins
Juliana Castaneda
Angel A. Juan
Abtin Tondar
Laura Calvet

IN3 – Computer Science Dept.
Universitat Oberta de Catalunya
Rambla Poblenou 156
08018 Barcelona, SPAIN

Barry B. Barrios

Stanford University
Stanford Center for Professional Development
408 Panama Mall
Stanford, CA 94305, USA

Jose Luis Sanchez-Garcia

Hospital Universitario Vall d'Hebron
Universitat Autonoma de Barcelona
Passeig de la Vall d'Hebron, 119
08035 Barcelona, SPAIN

ABSTRACT

When scheduling multi-period medical treatments for patients with cancer, medical committees have to consider a large amount of data, variables, sanitary and budget constraints, as well as probabilistic elements. In many hospitals worldwide, medical specialists regularly decide the optimal schedule of treatments to be assigned to patients by considering multiple periods and the number of available resources. Hence, decisions have to be made upon the priority of each patient, available treatments, their expected effects, the proper order and intensity in which they should be applied. Consequently, medical experts have to assess many possible combinations and, eventually, choose the one that maximizes the survival chances or expected life quality of patients. To support this complex decision-making process, this paper introduces a novel methodology that combines a biased-randomized heuristic with simulation, to return 'elite' alternatives to experts. A simplified yet illustrative case study shows the main concepts and potential of the proposed approach.

1 INTRODUCTION

Cancers are medical conditions in which some cells divide uncontrollably and may invade their surrounding tissues. In many cancers, malignant cells can spread to other organs by entering the bloodstream and lymph systems. As well-documented in many official sources (e.g., www.cancer.gov), there are many types of cancers. For example, leukemia is a type of cancer in which the tissues in the body that make white blood cells (WBCs), such as the lymphatic system and bone marrow, create too many WBCs. Multiple myeloma and lymphoma are the cancers of plasma cells and lymphocytes, respectively. Squamous cell carcinoma is the cancer of the skin. One of the most frequent is the adenocarcinoma that originates in the glandular epithelium, as in the gastrointestinal tract, colorectal, prostate, breast and some lung carcinomas. Sarcoma is a type of cancer that originates from the bone, cartilage, blood vessels, muscles, fat, and connective tissues. Tumors can also form in the central nervous system (CNS). The worldwide economic burden of cancer is unknown. However, the financial burden of cancer is estimated to be nearly 600 and 200 billion dollars annually in the United States and Europe, respectively (Hofmarcher et al. 2020).

There are many ways of treating cancers. For example, surgery, radiation therapy, chemotherapy, immunotherapy, targeted therapy, hormone therapy, and stem cell transplantation, to name the most popular ones. Personalized medicine is a vital component of treating cancers. Personalized cancer therapy aims to provide treatments specifically adapted to each patient's cancer biomarkers, such as cancer genes and proteins, and clinical symptoms throughout the course of treatments. Cancer patients can take genetic tests to show how their bodies would respond to different cancer treatment drugs. For instance, a patient's body may release a drug from itself faster than usual. In that case, physicians may use a higher dose of that drug to increase their half-life in the patient's body, resulting in higher treatment efficacy. This is called pharmacogenetics. Another advantage of personalized cancer therapy is minimizing the side effects of treatments by affecting less healthy cells and more cells involved in cancers. There is a substantial need to develop data-driven intelligent algorithms capable of supporting personalized cancer therapy. This includes choosing the best treatments for each patient, prioritizing patients, and using hospital resources in an optimal way that might depend on the cancer status of each patient.

Typically, each medical expert specializes in one part of the human body or in a pathology (e.g. oncology). Hence, whenever a patient shows multiple organs affected by a cancer, several decisions need to be made by a medical committee. In particular, this committee needs to design a multi-period treatment schedule or path, which has to consider dimensions such as current status and predicted evolution of the tumor in each part of the body, the patient's profile (genetic information, age, gender, and physical condition), recovery times, incompatibilities associated with each treatment, availability of resources at the hospital in each considered period, etc. Usually, there will be multiple patients to be scheduled, each of them with different characteristics and showing a variety of medical conditions. Thus, the number of possible combinations of treatments for them grows quite fast, making it difficult for a human mind to identify those treatment paths that could maximize the expected survival time of the set of patients or minimize the severity of the tumors affecting these patients at the end of the planned time horizon. In order to support such a complex decision making process, this paper proposes the use of intelligent hybrid algorithms capable of generating 'elite' multi-period scheduling paths for the set of patients. In particular, we propose a simheuristic algorithm (Juan et al. 2018) that combines a biased-randomized (BR) heuristic with simulation to generate a small collection of 'elite' scheduling paths for the set of patients. The committee of experts can then analyze these elite solutions and choose the one that better fits their medical criteria. Biased-randomization techniques allow us to introduce non-uniform randomness into a constructive heuristic. This is achieved by combining a skewed probability distribution with Monte-Carlo simulation. BR algorithms have been used to solve different combinatorial optimization problems, to provide alternative 'high-quality' solutions to complex optimization problems, and as a way to enhance some metaheuristic frameworks (Ferone et al. 2019; Quintero-Araujo et al. 2017; Juan et al. 2009).

The rest of the paper is structured as follows: Section 2 provides a review on scheduling in medical treatments. Section 3 presents more details on the specific decision-making problem being analyzed. Section 4 proposes a solving methodology based on the combination of a BR heuristic and simulation. The computational experiments carried out are described in Section 5, while Section 6 analyzes the obtained results. Finally, Section 7 draws the most relevant conclusions and identifies lines of future research.

2 RELATED WORK ON SCHEDULING CANCER THERAPIES

The number of cancer patients has been increasing during the last years, while the medical resources are still limited. This situation calls to improve the efficiency of patient scheduling. Most of the solving approaches to these scheduling optimization problems use exact methods to solve small instances, metaheuristics for the larger ones, and simulation methods to model these complex systems. Hence, for instance, Liang et al. (2015) describes the work with an oncology clinic in Burlington (USA) to develop a discrete event simulation model, which enables the evaluation of the operational performance in the clinic and the identification of initiatives for improvement in the process flow, scheduling, and staffing. A mathematical model is proposed, and computational experiments are carried out. Their approach balances doctors' workload, provides lower patient waiting times, and reduced total working times. Vogl et al. (2019) studies a real-world radiotherapy scheduling problem arising in a specialized ion beam center close to Vienna (Austria). In particular, the authors model this problem of scheduling recurring radiotherapy treatment appointments as a job shop scheduling problem with recurring activities, optional activities, and special time window constraints. The aim is to maximize facility usage and thereby minimize patients' waiting times before starting the treatment. Because of the complexity of the problem and the lack of reliable long-term stochastic data, they solve the deterministic version of the problem. A genetic algorithm (GA), an iterated local search (ILS), and a combination of both are designed, implemented, and compared using instances based on staff's knowledge about the underlying distribution functions. As expected, the two stand-alone metaheuristic approaches lead to reasonable solutions for small instances, while the combination of the GA and the ILS leads to significantly better results for larger instances.

Alvarado and Ntaimo (2018) points out the uncertainty in the appointment duration, the acuity levels of each appointment, and the availability of clinic nurses. The authors develop three mean-risk stochastic integer programming models for the problem of scheduling individual chemotherapy patient appointments and resources. These models were analyzed based on data from a real outpatient oncology clinic in Texas (USA). The experimental results show that the approach used can decrease patient waiting times and nurse over-time when compared to deterministic scheduling algorithms. Chang et al. (2020) presents a model aiming to make more efficient patient scheduling, which is defined as a stochastic online problem. The goal is to provide treatment for cancer at the earliest possibility to increase the patients' survival rate. An adaptive GA is proposed as the solving methodology. Small computational experiments are carried out based on data and characteristics of radiotherapy patient scheduling from a practical radiotherapy center. Similarly, Gocgun (2018) addresses the radiation therapy scheduling problem, where dynamically and stochastically arriving patients of different types are scheduled and treatment cancellations are considered. The authors formulate the problem as a Markov decision process (MDP). However, the MDP is intractable due to prominent state and action spaces, so they apply a simulation-based approximate dynamic programming approach. Unfortunately, the experiments are not entirely replicable, since the problem instances use randomly generated inputs.

In the context of particle therapy, Maschler and Raidl (2020) proposes an extension of the classical therapy patient scheduling problem in which the therapies should be provided on treatment days roughly at the same time. An iterated greedy metaheuristic is employed to deal with the problem. The authors compare the performance of different algorithms using a set of instances with up to 300 therapies (https://www.ac.tuwien.ac.at/research/problem-instances/). Similarly, Maschler et al. (2016) presents a heuristic approach based on a greedy randomized adaptive search procedure as well as an iterated greedy metaheuristic for the midterm planning part of the particle therapy patient scheduling problem (PTPSP). The objective is to schedule a set of therapies with all operational constraints satisfied, while minimizing treatment times and costs for used extended times of resources.

A different version of treatment scheduling for radiotherapy is presented in Vieira et al. (2020), which includes patient preferences regarding the time of their appointments. A mixed-integer linear programming model is presented. With up to 66 patients and 2 linear accelerators (linacs) per week, small-size instances are solved to optimality in reasonable computing times, while a heuristic is used to decompose the problem for more extensive instances. Computational experiments are described based on an instance generator that uses historical patient data from the Netherlands Cancer Institute radiotherapy department.

3 DETAILED PROBLEM DESCRIPTION

The efficient assignment of medical treatments to cancer patients is one of the critical problems that hospitals face today. In cases of patients with advanced, vital organ-compromised, or metastatic cancers, a committee of surgeons often meets to analyze each case and determine the optimal sequence of treatments that will be applied to maximize each patient's life expectancy and quality. In such cases, surgeons must consider how the combinations of different treatment regimes can impact each patient's treatment and overall health. Besides, recovery times between treatments may vary significantly. As illustrated in Figure 1, a specific treatment might be different from period to period (e.g., patient 2), or it can cover several periods in a row (e.g., patient 1 receives treatments that cover several consecutive periods). In other cases, the sequence of treatments for the patient starts in a period of time t, and several resting or monitoring periods need to be scheduled before a second intensive treatment can be applied.

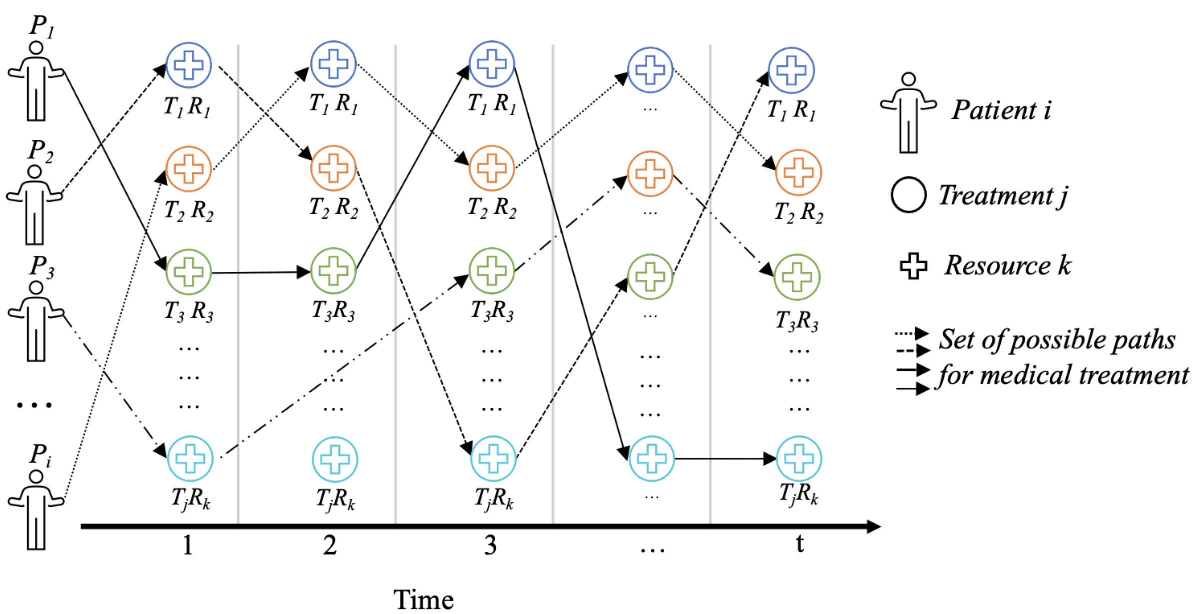

Figure 1: Representation of scheduling paths for medical treatment.

Under these circumstances, the main goal is to find a multi-period schedule (i.e., a patient-treatment assignment path) that respects the characteristics of each treatment and searches the maximization of the patients' welfare (or, equivalently, to minimize the aggregated level of illness caused by the tumors). As treatments are applied to different patients, their condition will need to be updated at each new period. Although not considered in this work, in a real-life scenario there might be changes in the list of patients over time, since new patients might arrive to the hospital and others can leave if they do not require additional treatment for a while. In any case, the availability of resources that can be assigned to the list of patients is limited, which means that once a specific resource has been assigned to a patient in a given period, it cannot be assigned to another patient in that same period (i.e., the resources assigned to a patient at any given time may not be available to another patient simultaneously). Therefore, it is imperative to determine the best treatment path for each patient, considering that some of them must be prioritized to use specific treatments due to the severity of their medical conditions. Notice that the problem described here could be easily extended to a more realistic scenario by considering dynamic lists of patients and resources, treatment costs, etc. Still, the initial version of the problem presented in this paper is already challenging due to the additional complexity introduced by the uncertainty regarding the response of each patient to the assigned treatment, which justifies the need for using a simulation-optimization approach.

Regarding additional details, we will assume that historical data is available for each patient (this might include overall physical health, age and gender, clinical symptoms, cancer-related genetic characteristics, histopathological data, cancer stage, metastasis, drug response, resistance, etc.). The patients' medical profiles are predefined, and each of them has associated a list of possible favorable or compatible procedures, such as treatment sessions provided at specific periods of time. The restrictions of these procedures are the following: *(i)* some of these actions might be interconnected (in fact, effective chemotherapy treatment may require multiple or a few sessions, each with a different level of intensity using various doses of one or multiple drugs); *(ii)* a tumor-removal surgery might not be applied prior to other actions in order to reduce the tumor size under a safe threshold level; *(iii)* some actions might be incompatible with others; *(iv)* some treatments can only be applied if the patient's resilience level is above a given threshold; and *(v)* some treatment procedures may pose a significant health risk to a patient.

4 OUR SIMULATION-OPTIMIZATION APPROACH

To solve this scheduling problem, we first propose a constructive heuristic to establish scheduling paths for a set of patients. In a second stage, this heuristic will be integrated into a BR multi-start framework by employing random sampling from a skewed probability distribution. The greedy heuristic performs as follows:

1. At the beginning of each period, the list of patients is sorted by urgency. The level of urgency of each patient is established according to the following hierarchical variables: *(i)* tumor type (benign, premalignant, or malignant); *(ii)* tumor status (in a scale of 0 to 10, where lower values represent a better condition); *(iii)* tumor location (some parts of the body might require a more urgent intervention when affected by a tumor); *(iv)* genetics predisposition (low, medium, or high); and *(v)* physical condition (low, medium, or high).
2. From the sorted list, patients are iteratively selected following a greedy criterion (from more urgent to less urgent). For each patient, a list of compatible treatments is created (notice that this list has to be updated every time a treatment resource is actually assigned to a patient, since it might not be available anymore in the current period for other patients). This treatments list will always contain the monitoring and the resting actions. It might also contain other more advanced treatments if they are compatible, e.g., surgery, chemotherapy, radiotherapy, etc. Typically, these advanced treatments also impose additional monitoring and / or resting treatments during subsequent periods. Eventually, the list of treatments is sorted according to the estimated reduction in the tumor status that each treatment will achieve. In this paper, we assume that sufficient historical data exist and that the actual reduction in the tumor status can be accurately fitted by a probability distribution, so the expected value can be considered as a reasonably good estimate. Finally, the first available treatment from the sorted list of compatible treatments is assigned to each patient. In the case of advanced treatments, the related monitoring or resting actions are also scheduled for future periods as prioritized treatments.
3. Once a treatment has been assigned to each patient, a more accurate estimate is assigned to the associated reduction in tumor status value. This enhanced estimate can be computed, for instance, by using regression models/machine learning methods or simulation tools such as CaTSiT (Bethge et al. 2015). CaTSiT is a computer simulation of metastatic progression and treatments published under the GNU General Public License. It is written in Java, the structure of the setup files is specified by an XML schema, and the results are saved in a spreadsheet file. CaTSiT allows a quantitative comparison of different metastasis formation models with clinical and experimental data and the effects of chemotherapy, external beam radiation, radioimmunotherapy, and radioembolization. This software, which is based on a discrete event simulation procedure, provides the number of

metastases and the total tumor burden over time. Notice that we are only interested in enhancing the estimate of the selected treatment, since improving estimates of all possible treatments for each patient will significantly increase computational times. For our numerical experiments, and since we are building the first prototype of a new methodology, we have emulated the random behavior of the reduction in the tumor status by considering it as a random variable that follows a normal probability distribution. For the mean value, we have used pre-established values for each treatment. Regarding the variance, for our numerical experiments, we have considered it to be unitary (notice that the methodology will be the same if other values are used, and only the results will vary). With the enhanced estimates, the tumor status is updated for each patient accordingly. Also, whenever a monitoring or a resting action is assigned by a lack of other available options, a degradation in the tumor status is recorded.

4. The previous steps are repeated until all predefined periods have been considered. At each period, the list of available resources is reset, the list of patients is resorted according to their updated health status, and the imposed treatments (e.g., monitoring or resting ones) scheduled in previous periods are prioritized. Eventually, all last-period individual tumor status are aggregated to compute an estimate of the final tumor status of the patients as a collective.

As explained before, the heuristic tries to minimize the final (last-period) aggregated tumor status by prioritizing the patients by urgency and, for each of them, by selecting the most promising treatment available. In other words, the goal is to reduce the collective tumor status as much as possible by the end of the planned horizon. While each treatment contributes to the progressive improvement of the tumor status, applying a monitoring or resting action due to the lack of resources might worsen the tumor status.

The aforementioned constructive heuristic iterates over the sorted list of patients and the sorted list of compatible treatments in a greedy way, by always selecting the first feasible option in each list. Although this greedy strategy is able to provide a relatively good solution in extremely short computing times (usually below one second), even better solutions can be quickly obtained by introducing a biased-randomization process into the selection of each patient's treatment and then encapsulate the randomized heuristic into a multi-start framework, thus obtaining a probabilistic algorithm. This biased-randomization process makes use of Monte-Carlo simulation concepts to generate random samples from a skewed probability distribution (Grasas et al. 2017). In our case, we use a Geometric probability distribution with a parameter $\beta \in (0,1)$. As β approaches to 1, the performance of the probabilistic algorithm converges to the one of the greedy heuristic. On the contrary, as β approaches to 0, the probabilistic algorithm emulates a uniform-random behavior. As it will be shown in the next section, more interesting results can be achieved for intermediate values of the β parameter.

Figure 2 presents the overall flow of our methodology. As introduced, the greedy heuristic is extended into a BR algorithm, in which the selection of treatments to be assigned to patients is smoothed. Moreover, the benefit incurred by receiving any treatment is simulated. In this way, the solution method is able to generate a set of alternative treatment paths for decision-makers, each of them composed of different characteristics, such as solution variability, mean, and reliability.

5 COMPUTATIONAL EXPERIMENTS

The proposed algorithm was coded in Python 3.8 and the tests were performed on an Intel Core $i7$-8550U processor with 16 GB of RAM. For testing the concepts described in this paper, we generated a case study that considers the patient's medical parameters presented in Table 1 as input data. For each parameter, an order of importance is assigned, through which the patient's priority is determined in each period. The tumor type is the most important parameter employed during the hierarchical sorting of the list. It determines whether the tumor it is malignant, premalignant, or benign, with malignant being the highest priority and benign being the lowest priority. The tumor status is the second parameter through which the priority list is ordered in each period, and assumes values between 0 and 10. When a patient achieves a status of value

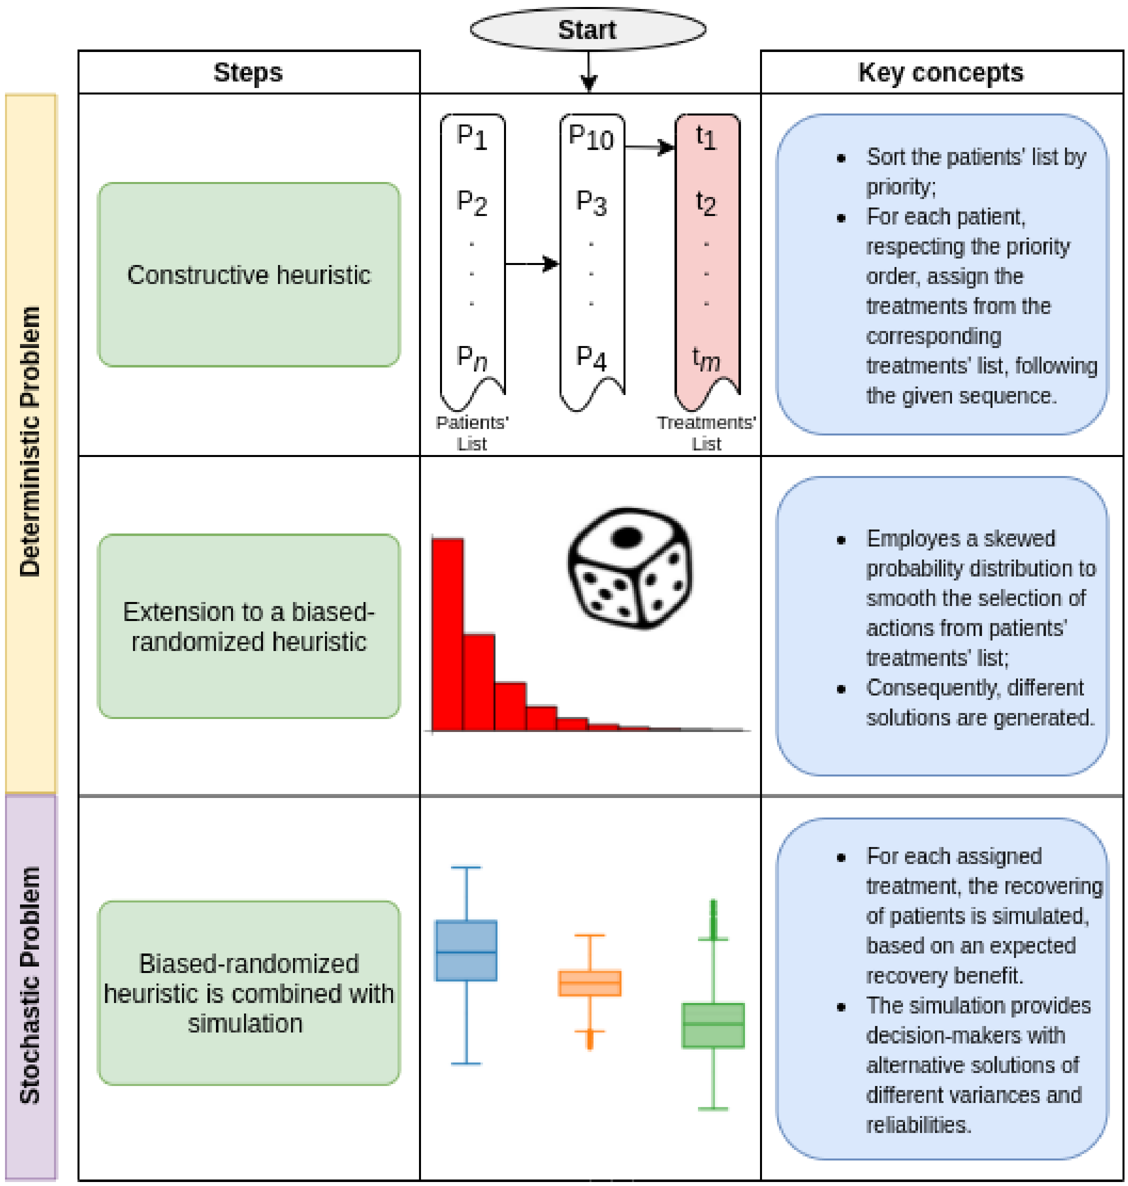

Figure 2: Optimization-simulation approach flow.

0, she is assumed to be in the best possible condition inside her type of tumor. On the contrary, a status of 10 indicates that the patient is in the worst possible condition inside her type of tumor. The location of the tumor is the third parameter that is used to sort the priority list, with neurological cancer being the highest priority, followed by digestive organ cancer, reproductive organ cancer (gynecologic / urologic), and finally breast cancer. Genetic predisposition, physical condition, and age are also used, in this order, to undo possible matches in the previous sorting criteria.

Table 1: Parameters used from the patient's clinical history.

Sort priority	Parameter	Description	Posible values
1	Tumor type	Type of tumor cell	{benign, premalignant, malignant}
2	Tumor status	Stage of cancer which can be from cured (0) to terminal (10)	{1:10}
3	Tumor location	Location of cancer cells	{breast, gynecologic / urologic, digestive, neurologic}
4	Genetics	Genetic predisposition of the patient to tumors	{low, medium, high}
5	Physical	Physical condition of the patient	{low, medium, high}
6	Age	Patient Age	{0:120}

Five possible treatments are considered, which are chemotherapy, surgery, radiotherapy, monitoring and resting. The first three are associated with an initially defined set of subsequent actions:

- A complete chemotherapy treatment might require three sessions (at different time periods) with a resting action between each chemotherapy sessions (i.e., 6 actions in total, covering 6 periods).
- A tumor-removal action or surgery has associated a sequence of treatments that are resting and monitoring sessions. In case there are no resources for the monitoring, the patient will receive a resting action as post-surgery treatment, and the monitoring is considered for being assigned in the following period.
- Similarly to chemotherapy, a complete radiotherapy treatment might require three sessions (at different time periods) with a rest between each chemotherapy sessions (i.e., 6 actions in total, covering 6 periods).
- When all the treatments on the initial list have been allocated, a monitoring, followed by a resting action are scheduled for each patient until the last period.

The case study also includes a total of 15 patients and 6 periods. Regarding the resources' availability, this is considered to be 'intermediate', i.e., there might be situations in which some patients must wait for the next periods to receive the treatment. In this way, it is possible to investigate the dynamism of the environment.

Table 2 presents the parameters' setting for our experiments. Regarding the β for selecting the treatments to be assigned to patients through biased-randomization, its values are selected randomly between 0.6 and 0.9. For the simulation of the patients' evolution, the parameter μ assumes the values -3, -2, 0 or 1, in case the treatments are chemotherapy/radiotherapy, surgery, rest/monitoring, or rest due to the lack of resources, respectively. As mentioned, the latter 'rest' treatment incurs the patients to be waiting for resources, then worsening their health status. Regarding the parameter σ, we have considered the value 1.0, whose which the μ parameter depends to be modified during the simulation process.

Table 2: Problem Parameters

β	μ	σ	**# periods**
$[0.6, 0.9]$	$\{-3, -2, 0, 1\}$	1.0	6

6 ANALYSIS OF RESULTS

The best solutions to the problem are those with the lowest total value of the patients' tumor status in the last period. Figure 3 presents the box-plots of the best 5 solutions obtained for the set of 15 patients over 6 periods, for both the greedy and BR assignment of treatments to patients. As one can notice, the first set of solutions provided by the greedy strategy (Figure 3a) presents greater and similar variability of the final tumor status. Moreover, for the cases in which the variability of the collective tumor status is lower, i.e., solutions 1 and 2, the resulting outliers represent patients without a good progressive recovering, reaching a tumor status close to 6. While the accumulated tumor status generated by greedy assignment varies from 15.3 to 16.7, for the case in which the selection is biased by a probability distribution (Figure 3b), this measure is in between 12.6 and 14.8. It allows us to certify which solving methodology is able to provide solutions with a better recovering of patients. Besides, it can be seen in Figure 3b, that solution 1, with the lowest patient tumor status, presents an outlier even though the interquartile range is small. This indicates that, despite the fact that the collective tumor status is much reduced for the last period, it is a solution path that does not improve the whole set of patients. Similarly to the greedy results, solution 1 has an outlier with tumor status close to 6. On the other hand, the variability of solutions 2 and 3 are very similar, although the median of solution 2 is lower than the median of solution 3. This means that the treatment path in solution 2 generates a significant improvement, since half of the patients have a tumor status less than or equal to 2, with the other half above this value. Solution 4 has the lowest variability of the best solutions presented, but its median is the highest. It also has two outliers that are far away from the other values in that solution, and thus from the average of the tumor status in general. This suggests that it is a solution that does not offer a distribution of resources and allocation of treatment that contributes to the recovery of the whole patient population. On the contrary, solution 5 does not present any outliers, and its variability is higher than solution 1 and 4. The distribution of the results is asymmetric, with a low median of around a tumor status of 1.5, indicating that many of the patients have a tumor status at that value, and only a few are above it. Solutions such as 5 and 2 are considered to be good solutions, since they offer a lower tumor status for most of the patients in the last period, and their maximum variability does not exceed the median possible tumor status value. Both solutions have a median of less than 2, indicating that 50% of the patients have a tumor status below this value. These two solutions, despite being in an acceptable range for tumor status, do not reach the value of zero for any patient. On the contrary, solution 3 has its lower limit at 0, indicating that there are patients who recover completely.

The evolution of the tumor status of the patients over the 6 evaluated periods is monitored individually. In Figure 4, for each solution strategy –greedy and BR–, the tumor status evolution of patients 4 and 5 is depicted for the considered periods. For the first patient, both strategies present similar convergence of the tumor status over time, reaching the total recovering at period 5. Specifically, the initial treatment list of patient 4 has chemotherapy and surgery. The patient has recovered only with the chemotherapy treatment, which requires 3 sessions and a rest between each one, so it is not necessary to schedule the surgery in the following periods –in the hypothetical case that there would be more periods or another cycle of 6 periods. On the other hand, for patient 5 (Figure 4b), the evolution of the tumor status has constant recovery behavior over the periods that treatment is applied. In this case, it can be observed how fast the tumor status decreases over time when the BR strategy is applied, contrary to the greedy strategy, which presents slower convergence. Like patient 4, patient 5 has chemotherapy as his/her initial treatment, but he/she does not recover completely in the 6 treatment periods. For a complete recovery, this patient needs more periods and treatments. However, the tumor status is reduced from 10 to less than 2 when the BR assignment is applied, while the greedy strategy allows the achievement of a final tumor status of 4.

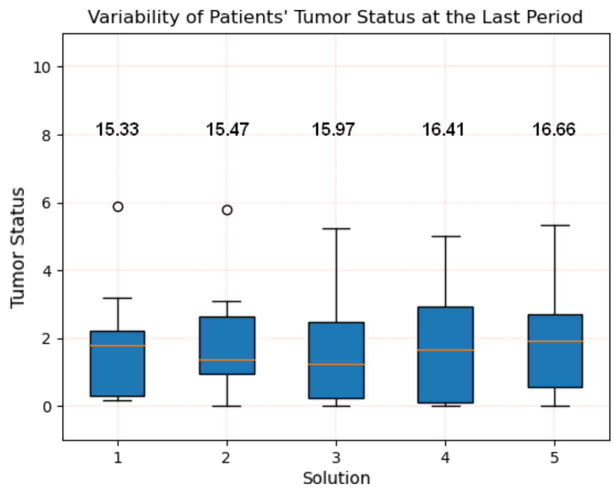

(a) Greedy assignment.

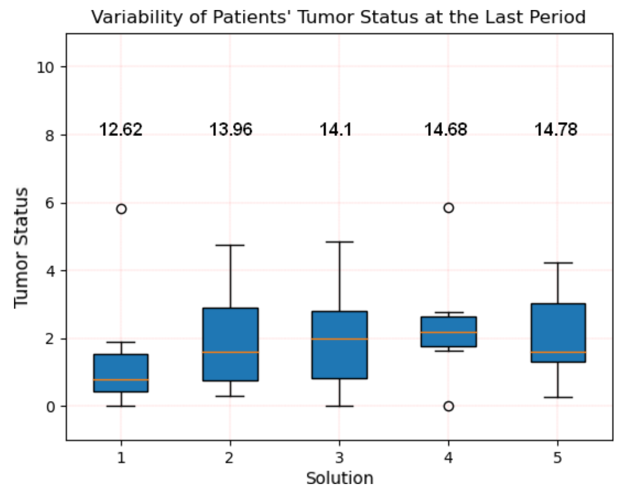

(b) BR assignment.

Figure 3: Box-plots of the patients' tumor status in the last period for each of the best 5 solutions.

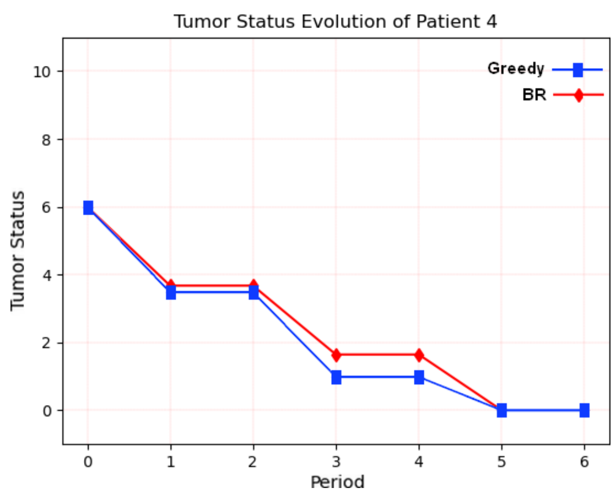

(a) Patient 4.

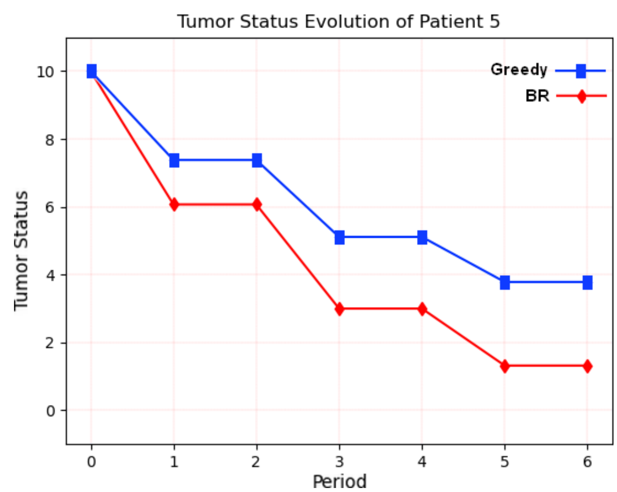

(b) Patient 5.

Figure 4: Tumor status evolution of patients resulting from the greedy and BR treatments assignment for six periods.

In general, the results present different treatment paths that, with an optimal allocation of resources, offer patients partial or total recovery in a given period of time. The information obtained allows for planning not only the allocation of hospital resources –which is usually limited for cancer patients–, but also to foresee the patient's recovery based on a sequence of treatments that can be scheduled.

7 CONCLUSIONS

Nowadays, there is an increasing need for data-driven intelligent algorithms capable of supporting decision-making in cancer therapy. This support includes choosing the best treatments for each patient, prioritizing them according to their medical condition, and making optimal use of the hospital's resources. This work has proposed an approach, which combines a BR heuristic with simulation, that has significant potential to evaluate numerous possible outcomes of changes in many variables that may affect both patients and medical centers. The aim is to suggest the most efficient and effective hypothetical types and sequences of operations and treatments, resulting in maximizing patients' recovery and optimizing medical centers' performance. An illustrative case study has been carried out to show the main concepts and potential of our approach.

According to our results, the use of the proposed strategy for supporting the assignment of cancer treatments to numerous patients allows finding an efficient treatment pathway that guarantees the patient's improvement or total recovery in a given time frame while ensuring the efficient use of the hospitals' resources. Regarding the different approaches, the BR strategy showed to be more efficient than the greedy selection in terms of both variability and accumulated tumor status of patients at the end of the planning horizon. Moreover, individually, patients are more likely to recover more efficiently when BR is incorporated into the decision-making process.

Several lines of research stem from this work. Our first goal is to integrate the CaTSiT simulator (Bethge et al. 2015), a computer simulation of metastatic progression and treatments, into our methodology, so that the estimated reduction in tumor status can be more accurate. Also, we plan to introduce other enhancements in our model in order to increase its realism, among others: *(i)* allow new patients to enter the system, as well as other customers to leave the it; *(ii)* increase the number of treatments as well as the number of options associated with each of them; *(iii)* increase the number of tumor locations to be considered; and *(iv)* consider side effects and the drug resistance of patients to the different treatments. In addition, budget restrictions or cost minimization could be introduced, since they constitute relevant factors. From a methodological point of view, a powerful metaheuristic framework that includes local search strategies could help to improve the results. Finally, more computational experiments are required to assess the performance of our methodology under a wide range of scenarios.

ACKNOWLEDGMENTS

This work has been partially funded by the Spanish Ministry of Science, Innovation, and Universities (PID2019-111100RB-C21/AEI/ 10.13039/501100011033 and RED2018-102642-T), and the Erasmus+ Program (2019-I-ES01-KA103-062602).

REFERENCES

Alvarado, M., and L. Ntaimo. 2018. "Chemotherapy Appointment Scheduling under Uncertainty using Mean-risk Stochastic Integer Programming". *Health Care Management Science* 21(1):87–104.

Bethge, A., U. Schumacher, and G. Wedemann. 2015. "Simulation of Metastatic Progression using a Computer Model including Chemotherapy and Radiation Therapy". *Journal of Biomedical Informatics* 57:74–87.

Chang, W., S. M. Lo, T.-L. Chen, J. C. Chen, and H.-N. Wu. 2020. "Utilizing Online Stochastic Optimization on Scheduling of Intensity-modulate Radiotherapy Therapy (IMRT)". *Journal of Biomedical Informatics* 108:103499.

Ferone, D., A. Gruler, P. Festa, and A. A. Juan. 2019. "Enhancing and Extending the Classical GRASP Framework with Biased Randomisation and Simulation". *Journal of the Operational Research Society* 70(8):1362–1375.

Gocgun, Y. 2018. "Simulation-based Approximate Policy Iteration for Dynamic Patient Scheduling for Radiation Therapy". *Health Care Management Science* 21(3):317–325.

Grasas, A., A. A. Juan, J. Faulin, J. de Armas, and H. Ramalhinho. 2017. "Biased Randomization of Heuristics Using Skewed Probability Distributions: A Survey and Some Applications". *Computers & Industrial Engineering* 110:216–228.

Hofmarcher, T., P. Lindgren, N. Wilking, and B. Jönsson. 2020. "The Cost of Cancer in Europe 2018". *European Journal of Cancer* 129:41–49.

Juan, A. A., J. Faulin, R. Ruiz, B. Barrios, M. Gilibert, and X. Vilajosana. 2009. "Using Oriented Random Search to Provide a Set of Alternative Solutions to the Capacitated Vehicle Routing Problem". In *Operations Research and Cyber-Infrastructure*, 331–345. Springer.

Juan, A. A., W. D. Kelton, C. S. Currie, and J. Faulin. 2018. "Simheuristics Applications: Dealing with Uncertainty in Logistics, Transportation, and other Supply Chain Areas". In *Proceedings of the 2018 Winter Simulation Conference*, edited by M. Rabe, A. A. Juan, N. Mustafee, A. Skoogh, S. Jain, and B. Johansson, 3048–3059. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Liang, B., A. Turkcan, M. E. Ceyhan, and K. Stuart. 2015. "Improvement of Chemotherapy Patient Flow and Scheduling in an Outpatient Oncology Clinic". *International Journal of Production Research* 53(24):7177–7190.

Maschler, J., and G. R. Raidl. 2020. "Particle Therapy Patient Scheduling with Limited Starting Time Variations of Daily Treatments". *International Transactions in Operational Research* 27(1):458–479.

Maschler, J., M. Riedler, M. Stock, and G. R. Raidl. 2016. "Particle Therapy Patient Scheduling: First Heuristic Approaches". In *Proceedings of the 11th Int. Conference on the Practice and Theory of Automated Timetabling. Udine, Italy*, 223–244. EURO Working Group on Automated Timetabling.

Quintero-Araujo, C. L., J. P. Caballero-Villalobos, A. A. Juan, and J. R. Montoya-Torres. 2017. "A Biased-Randomized Metaheuristic for the Capacitated Location Routing Problem". *International Transactions in Operational Research* 24(5):1079–1098.

Vieira, B., D. Demirtas, J. B. van de Kamer, E. W. Hans, L.-M. Rousseau, N. Lahrichi, and W. H. van Harten. 2020. "Radiotherapy Treatment Scheduling considering Time Window Preferences". *Health Care Management Science* 23(4):520–534.

Vogl, P., R. Braune, and K. F. Doerner. 2019. "Scheduling Recurring Radiotherapy Appointments in an Ion Beam Facility". *Journal of Scheduling* 22(2):137–154.

AUTHOR BIOGRAPHIES

LEANDRO DO C. MARTINS is a Postdoctoral Researcher at the Universitat Oberta de Catalunya (UOC), in Barcelona, Spain. He holds a PhD in Network and Information Technologies from UOC and a MSc in Computer Science from the Federal University of Ouro Preto, Brazil. His main research interests include metaheuristics, biased-randomized algorithms, vehicle routing, and scheduling. His email address is leandrocm@uoc.edu.

JULIANA CASTANEDA is a first-stage researcher at the ICSO research group at the Internet Interdisciplinary Institute in the Computer Science Dept. at the Universitat Oberta de Catalunya. She holds a bachelor's degree in industrial engineering and a master's degree in logistics and supply chain management. Her research interests are mainly focused on the application of intelligent algorithms and data science in the logistics and supply chains decision-making processes with efficiency and sustainability criteria. Her email address is jcastanedaji@uoc.edu.

ANGEL A. JUAN is a Full Professor of Operations Research & Industrial Engineering in the Computer Science Dept. at the UOC (Barcelona, Spain). He is also the Director of the ICSO research group at the Internet Interdisciplinary Institute and Lecturer at the Euncet Business School. Dr. Juan holds a Ph.D. in Industrial Engineering and an M.Sc. in Mathematics. He completed a predoctoral internship at Harvard University and postdoctoral internships at Massachusetts Institute of Technology and Georgia Institute of Technology. His main research interests include applications of simheuristics and learnheuristics in computational logistics and finance. He has published more than 100 articles in JCR-indexed journals and over 250 papers indexed in Scopus. His website address is http://ajuanp.wordpress.com and his email address is ajuanp@uoc.edu.

BARRY B. BARRIOS holds a B.Sc. in Physics from the Massachusetts Institute of Technology and a M.Sc. in Industrial and Systems Engineering from Virginia Tech. His main research lines are related to Discrete-Event Simulation for Enterprise Scale Systems. He is currently pursuing a graduate certificate from Stanford University in Artificial Intelligence. His e-mail address is barriosb@stanford.edu.

ABTIN TONDAR is a biomedical researcher who has worked in academia and a few pharmaceutical companies in the past few years. He currently pursues a Ph.D. in bioinformatics at the Universitat Oberta de Catalunya. His email address is mtondar@uoc.edu.

LAURA CALVET is an Assistant Professor of Statistics in the Computer Science Dept. at the UOC. Her main lines of research are: *i*) design of optimization algorithms relying on metaheuristics, machine learning and/or simulation applied to sustainable logistics & computing; *ii*) applied statistics & economics. Her email address is lcalvetl@uoc.edu.

JOSE LUIS SANCHEZ-GARCIA (MD, PhD) is a digestive-colorectal surgeon who works in the Department of General Surgery in the Hospital Universitario Vall d'Hebron (Barcelona, Spain). He is an Associate Professor in the Department of Surgery of the Medical School in the Universitat Autònoma de Barcelona. His email address is 30208jsg@comb.cat.