Intelligent Search for Personalized Cancer Therapy Synthesis: an Experimental Comparison

Marco Esposito[1], Leonardo Picchiami[1]

[1]*Computer Science Dept., Sapienza University of Rome, via Salaria 113, 00198, Italy*

Abstract

CRC is one of the deadliest forms of tumor for adult humans, which often requires expensive and highly toxic treatments. In this work, we exploit a recent System Biology Markup Language (SBML) quantitative model of the tumor growth and its immune response to two immunotherapic drugs and define a non-linear, constrained optimization problem for the synthesis of personalized therapies. We compute a virtual population of physiologically admissible patients and define a parametric therapy model that enables the usage of standard parameter optimization algorithms. The whole approach is implemented with a novel methodology that requires a single tool, namely COPASI, for the therapy template modelling, the definition of the optimization problem and its solution. We present results of an experimental evaluation of our approach by comparing five search algorithms among those implemented in COPASI, in order to find out the most promising kind of algorithms for this problem. The results show that population-based search algorithms perform well and are able to compute personalized therapies within time bounds compatible with clinical practice.

Keywords

In Silico Clinical Trials, VPH models, AI search, Systems Biology, Simulation.

1. Introduction

In recent years, the scientific community has produced a large number of quantitative models of human patho-physiology (see, *e.g.*, the well-known BioModels [1] among the others).

As these models become more and more accurate and gradually obtain certification from the authorities (such as FDA, EMA, see, *e.g.*, [2, 3]), new approaches to the design of pharmacological treatments and the safety and efficacy assessment of therapies for high-impact diseases become possible.

Such computational methods, collectively referred to as In Silico Clinical Trials (ISCTs), exploit Artificial Intelligence (AI) and Simulation-based Optimization techniques to design personalized therapies based on the specific characteristics of each patient, thus enabling precision medicine. Such methods, however, are still far from adoption in clinical contexts, due to both a lack of specialized tools and low trust in modern computational techniques.

The Virtual Phisiological Human (VPH) models used in ISCTs are generally designed to consider inter-patient variability, *i.e.*, difference in their features that cause, for instance, different

IPS-RCRA 2021: 9th Italian Workshop on Planning and Scheduling and 28th International Workshop on Experimental Evaluation of Algorithms for Solving Problems with Combinatorial Explosion

esposito@di.uniroma1.it (M. Esposito); piccchiami@di.uniroma1.it (L. Picchiami)

0000-0003-4543-8818 (M. Esposito); 0000-0001-5477-6419 (L. Picchiami)

evolutions of the disease. The models also capture absorption of drugs and their interaction with the patient's organism, *i.e.*, the Pharmaco-Kinetics/Dynamics (PKPD), which is crucial to analyze the different responses to drugs.

Quantitative VPH models can be seen as *hybrid systems*, where systems of Ordinary Differential Equations (ODEs) describe the dynamics of biological quantities, state and time events capture discrete dynamics, and parameters encode the inter-patient variability.

VPH models are often described and distributed in open-standard modelling languages, among which System Biology Markup Language (SBML) [4] is the most widespread. There exist several tools for the design and analysis of SBML models. One of the most complete is the open-source software COPASI [5]. It supports most of the SBML standard for modelling biological processes, several analysis tasks, such as time-course simulation, parameter optimization and estimation, and sensitivity analysis. It also provides a graphical user interface as well as Application Program Interfaces (APIs) for several programming languages.

1.1. Motivation

Pharmacological therapies, especially for high-impact and chronic diseases, are generally considerably expensive and toxic. Thanks to the use of ISCTs, we can simulate the effects of a large number of candidate therapies without the risks and costs of *in vivo* experiments on human patients, with the goal of finding the best one for each individual.

Given a VPH model and a particular virtual patient, the task of computing a personalized pharmacological therapy that maximizes the efficacy of the treatment and its safety is a hard one. It is often the case, in fact, that high dosing regimens have high chances of yielding effective therapies, but provide low safety due to drug toxicity. Furthermore, the number of possible therapies is generally infinite and, even if dosing values and time are discretized, grows combinatorially in the number of drugs and the duration of the treatment.

Naïve approaches to the modelling of the problem and the search for the best personalized therapy fail to explore such space effectively and within time bounds compatible with clinical practice. Given the complexity of the problem, it is unclear what search strategy is the most promising, even if considering a particular VPH model or a single Virtual Patient (VP). Generally, the search space is highly constrained, and the objectives are conflicting and strongly non-linear.

1.2. Contribution

In this paper, we consider a recent VPH model of the immune response to Colorectal Cancer (CRC) and analyze the performance of different types of search algorithms to synthesize personalized therapies. We first generated a population of VPs, along the lines of [6] which makes it possible to set up ISCTs. The definition of the search space, *i.e.*, the set of candidate therapies, is a crucial step, in that it severely influences the effectiveness of the synthesis. We designed a parametric model of therapies and integrated it within the CRC SBML model. This enabled a new methodology that, within a single tool, namely COPASI, supports the definition of the therapy optimization problem and its solution with several different search algorithms, in order to find the optimal therapy for any given VP.

1.3. Paper Outline

The paper is organized as follows. In Section 2 we recall the state of the art on ISCTs, optimization and parameter estimation for VPH models, Section 3 introduces the CRC VPH model and in Section 4 we describe how we modelled the therapies to enable the parameter optimization approach. Then, in Section 5 we formalize the optimization problem and in Section 6 we show the experimental results of different kinds of search algorithms to solve it. Finally, in Section 7 we draw our conclusions.

2. State of the art

A preliminary study for this work has been performed in [7]. In that work, the feasibility of the approach has been studied using a simpler therapy template model and a smaller number of algorithms. The problem of synthesizing personalized therapies using VPH models has been approached with a variety of techniques, *e.g.*, reinforcement learning [8] and genetic and evolutionary algorithms [9, 10, 11, 12]. A simulation-based backtracking algorithm has been proposed in [13, 14] for the problem of computing an optimal personalized therapy for assisted reproduction (an area where therapeutic success depends on many factors hard to be manually taken under control [15, 16, 17]) on a patient *digital twin*, defined as the subset of the VPs of a complete population best matching (up to a used-defined error threshold) the human patient clinical data. One of the most studied problems in systems biology is the estimation of unknown model parameters to fit experimental data. [18, 19, 20, 21] present several specialized optimization algorithms for this task. Real-world case studies include [22, 23, 24, 25]. Several tools are explicitly designed to solve parameter optimisation problems in systems biology; the most notable are AMICI [26], AMIGO2 [27], MEIGO [28], PESTO [29], Data2Dinamics [30], dMod [31], pyABC [32] and SBML2Julia [33]. Most of these use SBML and allow the definition of the optimisation problem via the PEtab [34] open standard. There exists a large number of SBML simulators. Among the most commonly used, besides COPASI, we mention BioUML [35] and LibRoadRunner [36]. SBML2Modelica [37] is a tool that focuses on the interoperability between systems biology and (hybrid) Cyber-Physical Systems (CPSs) domains by automatically translating SBML models to Modelica [38] and the FMI/FMU open standard. This enables the exploitation of well-established techniques for CPS optimization and verification. For instance, it is possible to use backtracking optimization algorithms [39] that exploit the efficient storing and retrieval of intermediate simulator states [40], as well as verification algorithms for systems also in presence of uncontrollable events (*e.g.*, [41, 42, 43, 44, 45, 46]). The solution of large-scale optimization problems via logic- [47], automata- [48] or constraint-based [49] formalisms (*e.g.*, [50, 51, 52, 53, 54] among others) has been largely studied in the literature, on several industry-relevant application domains (*e.g.*, [55, 56, 57, 58, 59, 60] among others). However, when the problem model is a complex VPH model, such approaches cannot be applied, since the model cannot be accurately defined an it is only available as a simulable black-box. The problem of finding an optimal therapy can be stated as an optimal planning model in hybrid domains. In fact, the VPH models we considered are hybrid systems whose inputs are discrete event sequences [61, 62]. Existing approaches to solve this problem [63, 64, 65, 66] are not applicable in our setting, as symbolic methods fail to capture the complexity of large ODE-based VPH

real-world models.

3. VPH model of the immune response to colorectal cancer

The VPH model proposed in [67] describes the dynamics of the growth of CRC and the immune response of the patient's immune system. The model also includes the PKPD of two drugs, namely Atezolizumab and Cibisatamab, and the patient's response to their administration. The change of the tumor size over time, as well as the total quantity of drugs administered, can be observed via simulation thanks to dedicated model variables.

As it is common for models of this kind, the dynamics of the tumor growth and the immune system's response to the drugs vary significantly from patient to patient. The CRC model presents 23 real-valued parameters, which jointly define the space of possible human patients. An assignment to those parameters represents a possible patient, and the evolution of their biological quantities can be computed by simulation.

However, as discussed in [68, 69], not all assignments yield patients that are of interest and plausible. In this model, it may be the case that a particular assignment leads to a patient that does not develop a tumor or that shows dynamics that are incompatible with the laws of biology. We used the approach from [6] to compute a population of physiologically admissible VPs. We adopted the admissibility criteria from [67], *i.e.*, admissible patients must develop a tumor of a given size within 8000 days and have plausible dynamics. Figure 1 shows the relative growth of the tumor in the 25 VPs we considered in our experiments, without any drug administration. It can be observed that for all of them the condition gets progressively worse over time.

We filtered out all VPs that showed no response to the two drugs among the whole virtual population. In this model, in fact, 85% of all admissible patients lack the biological characteristic to take advantage of immunotherapy. We excluded such patients because our ultimate goal is to assess the ability of optimization algorithms to synthesize effective treatments and not to study the efficacy of the drugs over the whole population (as done in [67]). Including non-responding VPs in the population would significantly alter the results since no optimization technique could ever find an effective therapy.

The CRC model makes it possible to simulate any therapy over a given VP. However, safety constraints exist over therapies (defined thanks to clinical trials on human patients), which deal with the toxicity of the drugs. In particular, it is not safe to administer more than 24 mg/kg of Atezolizumab and 160 mg of Cibisatamab in a single dose or cumulatively within a certain period (four weeks for Atezolizumab and three weeks for Cibisatamab).

In general, any VP is simulated for 400 days from the day in which the tumor has reached the given threshold (one of the 23 parameters). Such a period is deemed long enough to verify whether patients respond positively to therapies or not.

Figure 1 shows the relative growth of the tumor diameter in the 25 VPs we used in our experiments when no therapy is administered. For all of them, the condition get worse and worse over time.

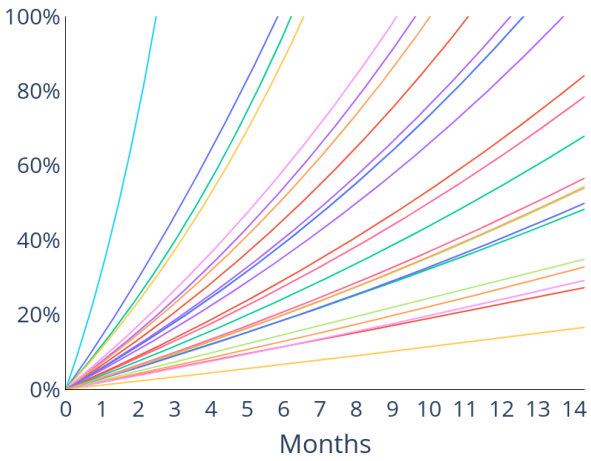

Figure 1: Relative growth of tumor size in the 25 patients without therapy.

4. Modelling the therapies

The set of all possible therapies involving any given number of drugs is infinite. It corresponds, in fact, to the set of all functions mapping each time instant (inside a finite interval) to the quantity of each drug administered at that time.

In order to find a personalized therapy for a given VP, we chose to parametrize the space of therapies. Defining a parametric template model for therapies enables, in fact, the use of standard parameter optimization techniques to compute an optimized therapy. The choice of one template model (among the many possible ones) is sensible since it has several repercussions over the success of our approach and the usefulness of the result. For our purposes, a desirable template model should meet a set of requirements.

Number of parameters. The number of parameters of the template must be as low as possible. The optimization problem generally gets harder and harder as the parameter space gets larger since there may be a combinatorial explosion in the number of possible therapies.

Expressiveness. The template model should be able to represent the largest possible portion of the therapies that are of interest. Since the goal is to find the optimal therapy, the set of expressible therapies should include it.

Constraints and symmetries. The number of constraints that define the space of admissible therapies should be as low as possible. Ideally, any assignment of the parameters should encode a therapy that is safe and of interest. Moreover, it would be desirable not to have symmetries (see, *e.g.*, [70]), *i.e.*, different parameter assignments that correspond to the same actual therapy.

Regularity and Explainability. The therapies that the template model can express must be as regular as possible. Very irregular therapies (*e.g.*, doses changing drastically from

week to week) may fail to obtain the physician's trust, which is crucial as the ultimate responsibility for the treatment lies on them.

Integration. The therapy template model must be described using the SBML language and integrated within the CRC model.

We designed a therapy template model that offers a good trade-off between all the requirements. We define 32 parameters: 30 of them, $\left(d_a^1, \ldots, d_a^{15}, d_c^1, \ldots, d_c^{15}\right)$, are real-valued and define the doses of each drug that can be administered during each of the 15 months, while the remaining two, f_a and f_c, are integer numbers that define the frequency of administration of, respectively, Atezolizumab and Cibisatamab. For each drug x, the administered dose at week $k \in [1, 60]$, denoted by $\text{adm}_x(\tau, k)$, is equal to $d_x^m \times \text{a}(x, k)$, where $m = \lceil \frac{k}{4} \rceil$ and $\text{a}(x, k)$ is 1 if $k \bmod f_x = 1$ and 0 otherwise.

Within this therapy template, we model the safety requirements about the maximum dose that can be administered in a single day as simple bound constraints (aka box constraints) over the 30 dose parameters.

The other two safety requirements, which limit the total quantity of drug administered within a certain period, must be modelled as non-linear constraints. Although COPASI supports the definition of non-linear constraints in the optimization problem, they are only evaluated after the simulation and not symbolically. Hence, the most straightforward way to model them is as functional (*i.e.*, black-box) constraints as follows. Let τ be a therapy. We introduce two new variables in the model, denoted by $\text{safety}_{ate}(\tau, t)$ and $\text{safety}_{cibi}(\tau, t)$, which, at time t, will hold 0 if the safety constraint over, respectively Atezolizumab and Cibisatamab is satisfied and a positive number otherwise.

We introduce an additional linear constraint that makes sure that the feasible therapies administer at least one of the two drugs within the first week as they would otherwise not be realistic.

The proposed model has a relatively low number of parameters and expresses very regular and explainable therapies since the doses stay constant over each month. The fact that it cannot represent some classes of therapies (*e.g.*, with doses changing within one month) is not a compelling limitation as the set of expressible therapies is still very large and diverse.

We described this therapy template model using the SBML language and integrated it into the model. Therefore, the resulting VPH model has 55 parameters that can be governed: 23 of them define the VPs and the remaining 32 define the space of therapies.

Finally, we note that our template model is able to represent all therapies for CRC (involving Atezolizumab and Cibisatamab) that have been tested in past clinical trials and that are summarized in [67].

5. Optimization

Given a VP λ in the set of all possible VPs $\Lambda \subseteq \mathbb{R}^{23}$, let t_0 be the starting time of therapies for λ, $H = t_0 + 400$ be the time horizon of the therapies for λ, and $\mathrm{T} = \mathbb{R}^{30} \times \mathbb{N}^2$ be the space of the therapy template model parameters. We want to compute a therapy τ which is optimal under two conflicting objectives: maximizing the efficacy and minimizing the total quantity of

administered drugs. We first define a measure of the *inefficacy* of a therapy τ as

$$\text{ineff}(\lambda, \tau, H) = \frac{\text{tumorsize}(\lambda, \tau, H)^2}{\text{tumorsize}(\lambda, \tau, t_0)},$$

where $\text{tumorsize}(\lambda, \tau, t)$ is a model variable that holds the size of the tumor at any time t during the simulation. We chose such a measure so that we can compare the inefficacy of therapies over different patients. In fact, only considering the final tumor size as the inefficacy measure would ignore the fact that the initial tumor size varies from patient to patient. Instead, our measure also considers the tumor reduction, and thus, given two patients with the same final tumor size, it gives a lower (better) inefficacy value to the one with the larger initial tumor.

We now define a measure of the quantity of drugs that is administered with therapy τ. Thanks to the safety constraints on each drug, we know the maximum total amounts of each drug that can be administered while still keeping the therapy safe. Let Max_{ate} and Max_{cibi} be such values. We define

$$\text{totdrugs}(\tau, H) = 100 \times \left(0.5 \times \frac{\text{total}_{ate}(\tau, H)}{\text{Max}_{ate}} + 0.5 \times \frac{\text{total}_{cibi}(\tau, H)}{\text{Max}_{cibi}}\right)$$

where $\text{total}_{ate}(\tau, t)$ and $\text{total}_{cibi}(\tau, t)$ are model variables that hold the total quantity of, respectively Atezolizumab and Cibisatamab administered up to time t with therapy τ.

We are now able to define our objective function: $J : \Lambda \times \mathrm{T} \to \mathbb{R}$ such that

$$J(\lambda, \tau \in \mathrm{T}) = \alpha \times \text{ineff}(\lambda, \tau, H) + \beta \times \text{totdrugs}(\tau, H)$$

with $\alpha, \beta \in \mathbb{R}_+$.

Given the maximum single-day doses μ_{ate} and μ_{cibi} for, respectively Atezolizumab and Cibisatamab, solving the personalized therapy synthesis problem means finding a therapy $\tau = \left(d_a^1, \ldots, d_a^{15}, d_c^1, \ldots, d_c^{15}, f_a, f_c\right) \in \mathrm{T}$ that minimizes $J(\lambda, \tau)$ while satisfying all constraints. In particular, τ must be such that

- $0 \leq d_x^i \leq \mu_x \quad \forall\, i \in [1, 15] \;\forall\, x \in \{ate, cibi\}$
- $\text{safety}_x(\tau, H) \leq 0 \quad \forall\, x \in \{ate, cibi\}$
- $d_a^1 + d_c^1 > 0.$

The latter constraint ensures that at least one of the two drugs is administered in the first week. In fact, regardless of the frequency of administration of drug x, if d_x^1 is > 0, the therapy administers a dose equal to d_x^1 at day 1.

5.1. Optimization algorithms

Among the different parameter optimization algorithms included in COPASI, we chose five in the attempt to cover different search techniques as much as possible.

We now describe each algorithm and, in particular, the salient characteristic of their implementation in COPASI.

Random Search. We used this algorithm as a baseline for comparing all studied optimization methods. The algorithm samples a random assignment of the therapy parameters at each step and evaluates the objective function if all constraints are met.

Levenberg-Marquardt. This algorithm is an adaptive-step gradient descent method, a hybrid between Steepest Descent and Newton methods. At each step, it uses finite differences to estimate the first and second derivatives in the current point. Bound constraints are handled by forcing parameters to stay within their bounds at each step. On the other hand, functional constraints are converted to non-negative penalties on the objective function.

Nelder-Mead. The algorithm constructs a simplex, *i.e.*, a polytope of $N + 1$ vertices in N dimensions. At each step, the objective function is evaluated at each vertex, and the polytope is modified according to the results. Constraints are handled analogously to the Levenberg-Marquardt algorithm.

Genetic Algorithm. A population-based metaheuristic search inspired by evolution. The initial population is constructed by including randomly sampled individuals and one individual given by the user. A new generation is obtained by randomly pairing individuals and producing two new individuals from each pair. The number of crossover points is chosen randomly between zero and 50% of the number of parameters, while a random parameter is mutated. In the selection phase, each individual's fitness (at this point, the number of individuals is double the initial number) is compared with that of a number of others equal to 20% of the population size. Then, all individuals are sorted by the number of individuals they outperformed, and finally, the population size is reduced to the initial number of individuals by removing those who performed worse. The algorithm always forces parameters inside their bounds, both when sampling random individuals and when applying mutations. If an individual does not meet all functional constraints, its fitness value is set to infinity, and thus it is effectively excluded. We set the population size to 64 and the mutational variance to 0.1.

Particle Swarm Optimization. A population-based metaheuristic search method in which a set of particles (points in the parameter space) collectively explore the search space in order to find the global minimum. The algorithm terminates if the standard deviation of the objective values of the particles is smaller than a given value. This implementation handles both bound and functional constraints like the Genetic Algorithm implementation. We used a swarm size of 30 and a standard deviation for termination of 0.1.

6. Results

In this section we show the outcome of the execution of all search algorithms over the 25 VPs. In each experiment, the initial point was set as the reference therapy described in [67], which administers the maximum allowed doses of drugs. The coefficients α and β of the objective function were set, respectively to 0.9 and 0.1. In order to mimic a real clinical setting, we

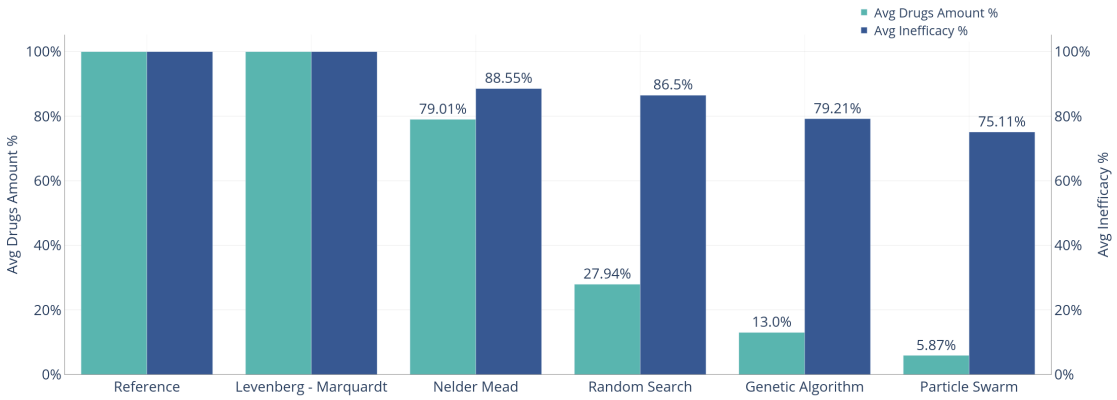

Figure 2: Average inefficacy and average total quantity of drug for each algorithm.

assigned to each experiment a time budget of 1 hour. All experiments were run on two identical machines equipped with Intel (R) Xeon (R) 8-core CPU E5430 and 32 GB RAM.

Figure 2 shows the average of the total quantity of drugs and the average inefficacy over all patients of the optimal therapies synthesized by each algorithm. In the chart, 100% of drugs quantity represents the maximum safe total quantity that can be administered, *i.e.*, the quantity of drugs administered by the reference therapy. For each patient, 100% of inefficacy indicates the inefficacy value reached on that particular patient with the reference therapy (which administers the maximum allowed quantity of drugs).

We first note that the Levenberg-Marquardt algorithm never improves over the reference therapy for all patients. Nelder-Mead, on the other hand, performs worse than Random Search, with an average reduction of the total drug quantity of just over 20%, while Random Search is able to reduce the total dose of around 72%. Genetic Algorithm and Particle Swarm Optimization both manage to drastically reduce the average total drug quantity while significantly improve on the inefficacy. The poor performance of the Levenberg-Marquardt and Nelder-Mead algorithms may be explained by the naïve handling of the bound and functional constraints. In fact, both algorithms were not designed to work in constrained space and the COPASI implementation makes ineffective choices whenever constraints are found to be violated. We also note that pure gradient-descent methods, like Levenberg-Marquardt, are naturally unfit for our problem, since the objective function is highly non-linear and the space is constrained. Not being able to escape from local minima and unfeasible regions is likely to make the problem unapproachable for this whole family of algorithms.

Figure 3 shows the Response Evaluation Criteria in Solid Tumors (RECIST) classification of the VPs with the reference therapy and with the optimal therapies. RECIST [71] is a set of classification criteria for the patients response to cancer therapies, based on the final diameter of the tumor after the treatment. RECIST classifies patients in four different classes: Complete Response (CR) if the final tumor diameter is < 10 mm, Partial Response (PR) if the tumor shrinks by at least 30%, Stable Disease (SD) if the diameter decreases by less than 30% or increases at most by 20% and Progressive Disease (PD) if the diameter grows by more than 20%. The

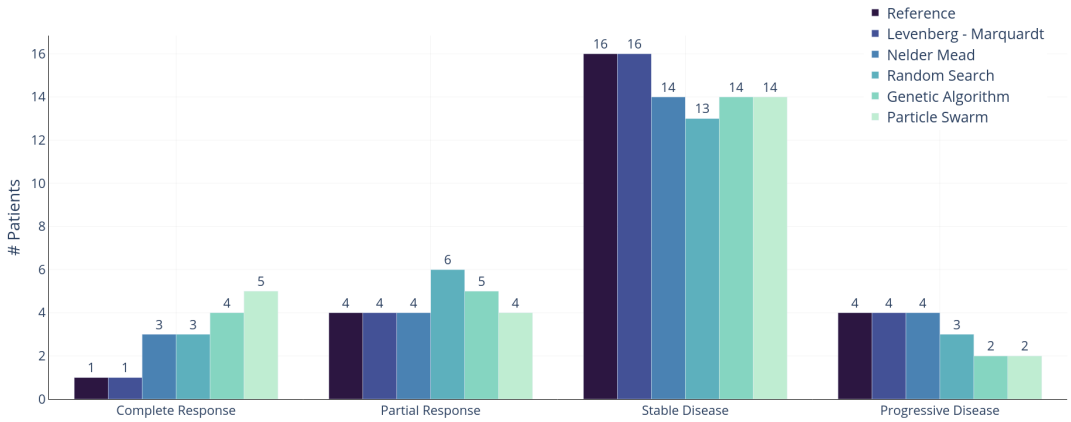

Figure 3: RECIST results with each algorithm.

RECIST classification of the 25 VPs with the optimal therapies are consistent with the previously shown results. In particular, we note that, with Genetic Algorithm and Particle Swarm, 3 and 4 VPs respectively change from PD to CR when the optimal therapy is administered.

Figure 4 shows the growth of the tumor of each VP with the reference therapy and with the optimal therapies synthesized with each algorithm, excluding Levenberg-Marquardt. As expected from the results shown in Figure 2, we observe that the tumor size tends to decrease most with Random Search, Genetic Algorithm and Particle Swarm. However, the latter also manages, for some patients, to reduce the tumor size more quickly, thus computing therapies that are safer and more effective.

Finally, Figure 5 shows the optimal therapies found for patient #3 with each algorithm. In each plot, 100% of a drug indicates the maximum quantity of that drug that can be administered in a single dose.

7. Conclusions

In this work, we exploited a SBML model of Colorectal Cancer to design an ISCT and compared five different search algorithms in the computation of personalized pharmacological therapies. The whole approach was implemented by using the tool COPASI. The approach required the modelling of a non-trivial parametric model of therapies, which allowed to reduce the search space to therapies that are actually of interest and enable the use of standard parameter optimization methods. We conclude, in our experimental analysis, that gradient-based algorithms, as well as algorithm that do not employ a smart handling of constraints, fail to solve the problem effectively. On the other side, population-based algorithms such as Genetic Algorithms and Particle Swarm Optimization show good and promising performance. Possible future research directions include the comparison of other families of algorithms (also outside of COPASI) and the modelling of more sophisticated therapy template models.

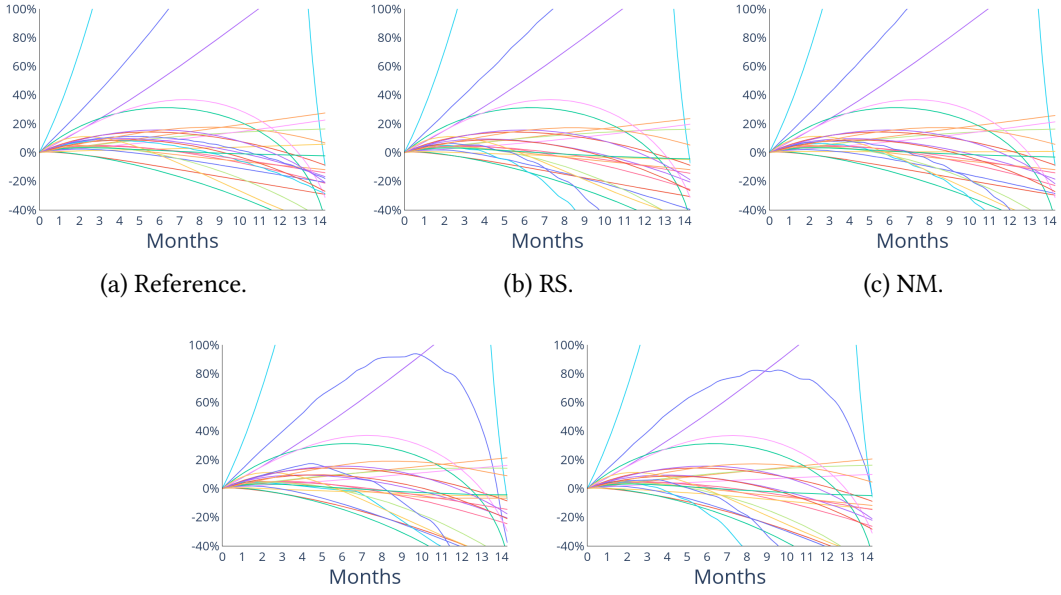

(a) Reference. (b) RS. (c) NM.

(d) G. (e) PSO.

Figure 4: Relative tumor size over time with reference therapy and the optimal therapies synthesized with each algorithm.

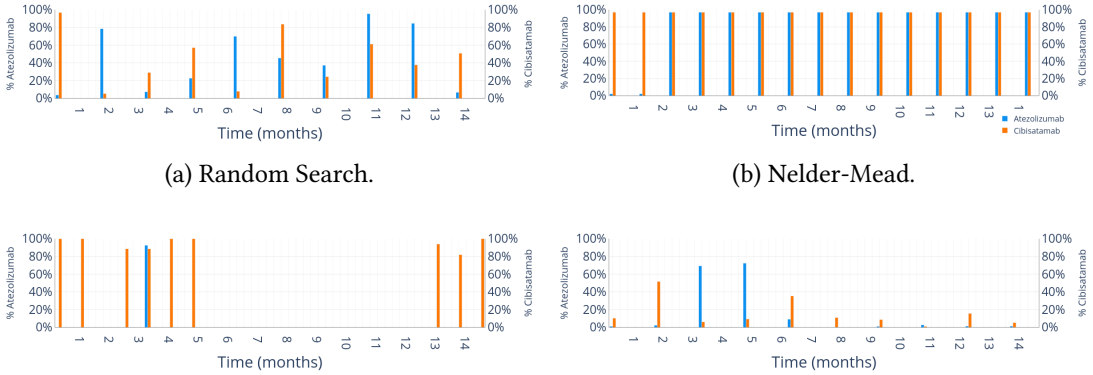

(a) Random Search. (b) Nelder-Mead.

(c) Genetic algorithm. (d) Particle Swarm Optimization.

Figure 5: Optimal therapies found for patient #3.

Acknowledgments

This work was partially supported by: Italian Ministry of University and Research under grant "Dipartimenti di eccellenza 2018–2022" of the Department of Computer Science of Sapienza University of Rome. INdAM "GNCS Project 2020"; Sapienza University projects RG11816436BD4F21, RG11916B892E54DB, RP11916B8665242F; Lazio POR FESR projects E84G20000150006, F83G17000830007.

References

[1] N. Le Novère, B. Bornstein, A. Broicher, M. Courtot, M. Donizelli, H. Dharuri, L. Li, H. Sauro, M. Schilstra, B. Shapiro, J. Snoep, M. Hucka, BioModels Database: A free, centralized database of curated, published, quantitative kinetic models of biochemical and cellular systems, Nucl. Ac. Res. 34 (2006). doi:10.1093/nar/gkj092.

[2] U.S.A. Food and Drug Administration, Physiologically based pharmacokinetic analyses – format and content guidance for industry, 2018. FDA-2016-D-3969.

[3] European Medicines Agency, Reporting of physiologically based pharmacokinetic (PBPK) modelling and simulation, 2019. URL: https://www.ema.europa.eu/en/reporting-physiologically-based-pharmacokinetic-pbpk-modelling-simulation, eMA/CHMP/458101/2016.

[4] M. Hucka, F. Bergmann, A. Dräger, S. Hoops, S. Keating, N. Le Novère, C. Myers, B. Olivier, S. Sahle, J. Schaff, L. Smith, D. Waltemath, D. Wilkinson, The Systems Biology Markup Language (SBML): Language specification for Level 3 Version 2 Core, JIB 15 (2018). doi:10.1515/jib-2017-0081.

[5] C. Lee, L. Xu, M. Singhal, P. Mendes, S. Hoops, J. ù Pahle, N. Simus, R. Gauges, S. Sahle, U. Kummer, COPASI —- A COmplex PAthway SImulator, Bioinformatics 22 (2006). doi:10.1093/bioinformatics/btl485.

[6] S. Sinisi, V. Alimguzhin, T. Mancini, E. Tronci, B. Leeners, Complete populations of virtual patients for in silico clinical trials, Bioinformatics 36 (2020). doi:10.1093/bioinformatics/btaa1026.

[7] M. Esposito, L. Picchiami, Simulation-based synthesis of personalised therapies for colorectal cancer, in: OVERLAY, 2021. To appear.

[8] A. Jalalimanesh, H. S. Haghighi, A. Ahmadi, M. Soltani, Simulation-based optimization of radiotherapy: Agent-based modeling and reinforcement learning, Mathematics and Computers in Simulation 133 (2017) 235–248.

[9] R. Schmucker, G. Farina, J. Faeder, F. Fröhlich, A. S. Saglam, T. Sandholm, Combination treatment optimization using a pan-cancer pathway model, bioRxiv (2020).

[10] N. Noman, P. Moscato, Designing optimal combination therapy for personalised glioma treatment, Memetic Computing 12 (2020) 317–329.

[11] T. Cassidy, M. Craig, Determinants of combination gm-csf immunotherapy and oncolytic virotherapy success identified through in silico treatment personalization, PLoS computational biology 15 (2019) e1007495.

[12] A. L. Jenner, F. Frascoli, C.-O. Yun, P. S. Kim, Optimising hydrogel release profiles for viro-immunotherapy using oncolytic adenovirus expressing il-12 and gm-csf with immature dendritic cells, Applied Sciences 10 (2020) 2872.

[13] T. Mancini, F. Mari, A. Massini, I. Melatti, I. Salvo, S. Sinisi, E. Tronci, R. Ehrig, S. Röblitz, B. Leeners, Computing personalised treatments through in silico clinical trials. A case study on downregulation in assisted reproduction, in: RCRA 2018, volume 2271 of *CEUR W.P.*, CEUR, 2018.

[14] S. Sinisi, V. Alimguzhin, T. Mancini, E. Tronci, F. Mari, B. Leeners, Optimal personalised treatment computation through in silico clinical trials on patient digital twins, Fundam. Inform. 174 (2020). doi:10.3233/FI-2020-1943.

[15] B. Leeners, T. Kruger, K. Geraedts, E. Tronci, T. Mancini, F. Ille, M. Egli, S. Roeblitz, L. Saleh, K. Spanaus, C. Schippert, Y. Zhang, M. Hengartner, Lack of associations between female hormone levels and visuospatial working memory, divided attention and cognitive bias across two consecutive menstrual cycles, Front. Behav. Neuro. 11 (2017). doi:10.3389/fnbeh.2017.00120.

[16] M. Hengartner, T. Kruger, K. Geraedts, E. Tronci, T. Mancini, F. Ille, M. Egli, S. Roeblitz, R. Ehrig, L. Saleh, K. Spanaus, C. Schippert, Y. Zhang, B. Leeners, Negative affect is unrelated to fluctuations in hormone levels across the menstrual cycle: Evidence from a multisite observational study across two successive cycles, J. Psycho. Res. 99 (2017). doi:10.1016/j.jpsychores.2017.05.018.

[17] B. Leeners, T. Krüger, K. Geraedts, E. Tronci, T. Mancini, M. Egli, S. Röblitz, L. Saleh, K. Spanaus, C. Schippert, Y. Zhang, F. Ille, Associations between natural physiological and supraphysiological estradiol levels and stress perception, Front. Psycol. 10 (2019). doi:10.3389/fpsyg.2019.01296.

[18] L. Schmiester, D. Weindl, J. Hasenauer, Efficient gradient-based parameter estimation for dynamic models using qualitative data, Bioinformatics (2021). URL: https://doi.org/10.1093/bioinformatics/btab512. doi:10.1093/bioinformatics/btab512.

[19] A. F. Villaverde, F. Fröhlich, D. Weindl, J. Hasenauer, J. R. Banga, Benchmarking optimization methods for parameter estimation in large kinetic models, Bioinformatics 35 (2019) 830–838.

[20] Y. Schälte, P. Stapor, J. Hasenauer, Evaluation of derivative-free optimizers for parameter estimation in systems biology, IFAC-PapersOnLine 51 (2018) 98–101.

[21] A. Yazdani, L. Lu, M. Raissi, G. E. Karniadakis, Systems biology informed deep learning for inferring parameters and hidden dynamics, PLoS computational biology 16 (2020) e1007575.

[22] O. D. Sánchez, E. Ruiz-Velázquez, A. Y. Alanís, G. Quiroz, L. Torres-Treviño, Parameter estimation of a meal glucose–insulin model for tidm patients from therapy historical data, IET systems biology 13 (2019) 8–15.

[23] C. Audigier, T. Mansi, H. Delingette, S. Rapaka, V. Mihalef, D. Carnegie, E. Boctor, M. Choti, A. Kamen, D. Comaniciu, et al., Parameter estimation for personalization of liver tumor radiofrequency ablation, in: International MICCAI Workshop on Computational and Clinical Challenges in Abdominal Imaging, Springer, 2014, pp. 3–12.

[24] T. Chen, N. F. Kirkby, R. Jena, Optimal dosing of cancer chemotherapy using model predictive control and moving horizon state/parameter estimation, Computer methods and programs in biomedicine 108 (2012) 973–983.

[25] M. Raissi, N. Ramezani, P. Seshaiyer, On parameter estimation approaches for predicting disease transmission through optimization, deep learning and statistical inference methods, Letters in Biomathematics 6 (2019) 1–26.

[26] F. Fröhlich, D. Weindl, Y. Schälte, D. Pathirana, Ł. Paszkowski, G. T. Lines, P. Stapor, J. Hasenauer, Amici: High-performance sensitivity analysis for large ordinary differential equation models, arXiv preprint arXiv:2012.09122 (2020).

[27] E. Balsa-Canto, D. Henriques, A. Gábor, J. R. Banga, Amigo2, a toolbox for dynamic modeling, optimization and control in systems biology, Bioinformatics 32 (2016) 3357–3359.

[28] J. A. Egea, D. Henriques, T. Cokelaer, A. F. Villaverde, A. MacNamara, D.-P. Danciu, J. R. Banga, J. Saez-Rodriguez, Meigo: an open-source software suite based on metaheuristics for global optimization in systems biology and bioinformatics, BMC bioinformatics 15 (2014) 1–9.

[29] P. Stapor, D. Weindl, B. Ballnus, S. Hug, C. Loos, A. Fiedler, S. Krause, S. Hroß, F. Fröhlich, J. Hasenauer, PESTO: Parameter EStimation TOolbox, Bioinf. 34 (2018) 705–707.

[30] A. Raue, B. Steiert, M. Schelker, C. Kreutz, T. Maiwald, H. Hass, J. Vanlier, C. Tönsing, L. Adlung, R. Engesser, et al., Data2dynamics: a modeling environment tailored to parameter estimation in dynamical systems, Bioinformatics 31 (2015) 3558–3560.

[31] D. Kaschek, W. Mader, M. Fehling-Kaschek, M. Rosenblatt, J. Timmer, Dynamic modeling, parameter estimation, and uncertainty analysis in r, Journal of Statistical Software 88 (2019) 1–32.

[32] E. Klinger, D. Rickert, J. Hasenauer, pyabc: distributed, likelihood-free inference, Bioinformatics 34 (2018) 3591–3593.

[33] P. F. Lang, S. Shin, V. M. Zavala, Sbml2julia: interfacing sbml with efficient nonlinear julia modelling and solution tools for parameter optimization, arXiv preprint arXiv:2011.02597 (2020).

[34] L. Schmiester, Y. Schälte, F. T. Bergmann, T. Camba, E. Dudkin, J. Egert, F. Fröhlich, L. Fuhrmann, A. L. Hauber, S. Kemmer, et al., Petab—interoperable specification of parameter estimation problems in systems biology, PLoS computational biology 17 (2021) e1008646.

[35] F. Kolpakov, I. Akberdin, T. Kashapov, l. Kiselev, S. Kolmykov, Y. Kondrakhin, E. Kutumova, N. Mandrik, S. Pintus, A. Ryabova, et al., Biouml: an integrated environment for systems biology and collaborative analysis of biomedical data, Nucleic acids research 47 (2019) W225–W233.

[36] E. Somogyi, J.-M. Bouteiller, J. Glazier, M. König, J. Medley, M. Swat, H. Sauro, libRoadRunner: A high performance SBML simulation and analysis library, Bioinformatics 31 (2015). doi:10.1093/bioinformatics/btv363.

[37] F. Maggioli, T. Mancini, E. Tronci, SBML2Modelica: Integrating biochemical models within open-standard simulation ecosystems, Bioinformatics 36 (2020). doi:10.1093/bioinformatics/btz860.

[38] P. Fritzson, V. Engelson, Modelica—a unified object-oriented language for system modeling and simulation, in: European Conference on Object-Oriented Programming, Springer, 1998, pp. 67–90.

[39] T. Mancini, F. Mari, A. Massini, I. Melatti, F. Merli, E. Tronci, System level formal verification via model checking driven simulation, in: CAV 2013, volume 8044 of *LNCS*, Springer, 2013. doi:10.1007/978-3-642-39799-8_21.

[40] S. Sinisi, V. Alimguzhin, T. Mancini, E. Tronci, Reconciling interoperability with efficient verification and validation within open source simulation environments, Simul. Model. Pract. Theory 109 (2021). doi:10.1016/j.simpat.2021.102277.

[41] T. Mancini, F. Mari, A. Massini, I. Melatti, E. Tronci, SyLVaaS: System level formal

verification as a service, in: PDP 2015, IEEE, 2015. doi:10.1109/PDP.2015.119.
[42] T. Mancini, F. Mari, A. Massini, I. Melatti, E. Tronci, SyLVaaS: System level formal verification as a service, Fundam. Inform. 149 (2016). doi:10.3233/FI-2016-1444.
[43] T. Mancini, F. Mari, A. Massini, I. Melatti, E. Tronci, Anytime system level verification via parallel random exhaustive hardware in the loop simulation, Microprocessors and Microsystems 41 (2016). doi:10.1016/j.micpro.2015.10.010.
[44] T. Mancini, F. Mari, A. Massini, I. Melatti, I. Salvo, E. Tronci, On minimising the maximum expected verification time, Inf. Proc. Lett. 122 (2017). doi:10.1016/j.ipl.2017.02.001.
[45] T. Mancini, F. Mari, I. Melatti, I. Salvo, E. Tronci, J. Gruber, B. Hayes, M. Prodanovic, L. Elmegaard, User flexibility aware price policy synthesis for smart grids, in: DSD 2015, IEEE, 2015. doi:10.1109/DSD.2015.35.
[46] T. Mancini, I. Melatti, E. Tronci, Any-horizon uniform random sampling and enumeration of constrained scenarios for simulation-based formal verification, IEEE TSE (2021). doi:10.1109/TSE.2021.3109842.
[47] C. Barrett, C. Tinelli, Satisfiability modulo theories, in: Handbook of Model Checking, Springer, 2018.
[48] E. Clarke, T. Henzinger, H. Veith, Handbook of Model Checking, Springer, 2016.
[49] F. Rossi, P. Van Beek, T. Walsh (Eds.), Handbook of Constraint Programming, Elsevier, 2006.
[50] M. Cadoli, T. Mancini, Combining relational algebra, SQL, constraint modelling, and local search, TPLP 7 (2007). doi:10.1017/S1471068406002857.
[51] M. Cadoli, T. Mancini, F. Patrizi, SAT as an effective solving technology for constraint problems, in: ISMIS 2006, volume 4203 of *LNCS*, Springer, 2006.
[52] T. Mancini, M. Cadoli, D. Micaletto, F. Patrizi, Evaluating ASP and commercial solvers on the CSPLib, Constraints 13 (2008).
[53] T. Mancini, P. Flener, J. Pearson, Combinatorial problem solving over relational databases: View synthesis through constraint-based local search, in: SAC 2012, ACM, 2012. doi:10.1145/2245276.2245295.
[54] T. Mancini, F. Mari, I. Melatti, I. Salvo, E. Tronci, An efficient algorithm for network vulnerability analysis under malicious attacks, in: ISMIS 2018, Springer, 2018. doi:10.1007/978-3-030-01851-1_29.
[55] T. Mancini, E. Tronci, A. Scialanca, F. Lanciotti, A. Finzi, R. Guarneri, S. Di Pompeo, Optimal fault-tolerant placement of relay nodes in a mission critical wireless network, in: RCRA 2018, volume 2271 of *CEUR W.P.*, CEUR, 2018.
[56] Q. Chen, A. Finzi, T. Mancini, I. Melatti, E. Tronci, MILP, pseudo-boolean, and OMT solvers for optimal fault-tolerant placements of relay nodes in mission critical wireless networks, Fundam. Inform. 174 (2020). doi:10.3233/FI-2020-1941.
[57] T. Mancini, F. Mari, I. Melatti, I. Salvo, E. Tronci, J. Gruber, B. Hayes, M. Prodanovic, L. Elmegaard, Demand-aware price policy synthesis and verification services for smart grids, in: SmartGridComm 2014, IEEE, 2014. doi:10.1109/SmartGridComm.2014.7007745.
[58] B. Hayes, I. Melatti, T. Mancini, M. Prodanovic, E. Tronci, Residential demand management using individualised demand aware price policies, IEEE Trans. Smart Grid 8 (2017). doi:10.1109/TSG.2016.2596790.

[59] T. Mancini, F. Mari, I. Melatti, I. Salvo, E. Tronci, J. Gruber, B. Hayes, L. Elmegaard, Parallel statistical model checking for safety verification in smart grids, in: SmartGridComm 2018, IEEE, 2018. doi:10.1109/SmartGridComm.2018.8587416.

[60] I. Melatti, F. Mari, T. Mancini, M. Prodanovic, E. Tronci, A two-layer near-optimal strategy for substation constraint management via home batteries, IEEE Trans. Ind. Elect. (2021). doi:10.1109/TIE.2021.3102431.

[61] T. Mancini, F. Mari, A. Massini, I. Melatti, E. Tronci, System level formal verification via distributed multi-core hardware in the loop simulation, in: PDP 2014, IEEE, 2014. doi:10.1109/PDP.2014.32.

[62] T. Mancini, F. Mari, A. Massini, I. Melatti, E. Tronci, Anytime system level verification via random exhaustive hardware in the loop simulation, in: DSD 2014, IEEE, 2014. doi:10.1109/DSD.2014.91.

[63] M. Fox, D. Long, Modelling mixed discrete-continuous domains for planning., JAIR 27 (2006).

[64] M. Vallati, D. Magazzeni, B. De Schutter, L. Chrpa, T. McCluskey, Efficient macroscopic urban traffic models for reducing congestion: A PDDL+ planning approach, in: AAAI 2016, AAAI, 2016.

[65] S. Bogomolov, D. Magazzeni, A. Podelski, M. Wehrle, Planning as model checking in hybrid domains., in: AAAI 2014, AAAI, 2014.

[66] S. Bogomolov, D. Magazzeni, S. Minopoli, M. Wehrle, PDDL+ planning with hybrid automata: Foundations of translating must behavior., in: ICAPS 2015, AAAI, 2015.

[67] H. Ma, H. Wang, R. J. Sové, J. Wang, C. Giragossian, A. S. Popel, Combination therapy with t cell engager and pd-l1 blockade enhances the antitumor potency of t cells as predicted by a qsp model, Journal for immunotherapy of cancer 8 (2020).

[68] E. Tronci, T. Mancini, I. Salvo, S. Sinisi, F. Mari, I. Melatti, A. Massini, F. Davi', T. Dierkes, R. Ehrig, S. Röblitz, B. Leeners, T. Krüger, M. Egli, F. Ille, Patient-specific models from inter-patient biological models and clinical records, in: FMCAD 2014, IEEE, 2014. doi:10.1109/FMCAD.2014.6987615.

[69] T. Mancini, E. Tronci, I. Salvo, F. Mari, A. Massini, I. Melatti, Computing biological model parameters by parallel statistical model checking, in: IWBBIO 2015, volume 9044 of *LNCS*, Springer, 2015. doi:10.1007/978-3-319-16480-9_52.

[70] T. Mancini, M. Cadoli, Detecting and breaking symmetries by reasoning on problem specifications, in: SARA 2005, volume 3607 of *LNCS*, Springer, 2005.

[71] L. H. Schwartz, S. Litière, E. De Vries, R. Ford, S. Gwyther, S. Mandrekar, L. Shankar, J. Bogaerts, A. Chen, J. Dancey, et al., Recist 1.1—update and clarification: From the recist committee, European journal of cancer 62 (2016) 132–137.