Optimizing immune cell therapies with artificial intelligence

Nicolas Houy* François Le Grand†

September 13, 2018

Abstract

Purpose. We determine an optimal injection pattern for anti-vascular endothelial growth factor (VEGF) and for the combination of anti-VEGF and unlicensed dendritic cells.

Methods. We rely on the mathematical model of Soto-Ortiz and Finley [45] for the interactions between the tumor growth, angiogenesis and immune system reactions. Our optimization algorithm belongs to the class of Monte-Carlo tree search algorithms. The objective consists in finding the minimal total drug doses for which an injection pattern yields tumor eradication.

Results. Our results are twofold. First, optimized injection protocols enable to significantly reduce the total drug dose for tumor elimination. For instance, for an early diagnosis date, a total dose equal to 58% of the standard anti-VEGF dose enables to eliminate the tumor. In the case of drug combination, associating 25% of the total standard anti-VEGF dose to 10% of the dendritic cell total standard dose eradicates tumor. Our second result is that administering a dose equal to the maximal standard dose allows for later diagnosis date compared to standard protocol. For instance, in the case of anti-VEGF injection, the optimal protocol postpones the maximal diagnosis date by more than one month.

Conclusions. Overall, our optimization based on artificial intelligence delivers significant gains in total drug administration or in the length of the therapeutic window. Our method is flexible and could be adapted to other drug combinations.

Keywords: Pharmacokinetics, Pharmacodynamics, Immunotherapy, Artificial Intelligence.

Conflicts of interest: None.

*University of Lyon, Lyon, F-69007, France; CNRS, GATE Lyon Saint-Etienne, F-69130, France. Email: houy@gate.cnrs.fr.

†emlyon business school, Écully, F-69130, France; ETH Zurich, Zurich, CH-8092, Switzerland. Email: legrand@em-lyon.com.

1 Introduction

Angiogenesis is the process of formation of new blood vessels from preexisting ones. It plays a crucial role in both physiological and pathological processes [24]. Indeed, angiogenesis is required on the one side for development of embryos and wound healing, and on the other side for tumor growth, in particular in carcinoma [46]. The role of angiogenesis for tumor growth has been initially hypothesized by Folkmann in the 70s (see [15] for the original paper and [16, 17] for later reviews) and demonstrated in a great number of studies ([25, 26, 35] among many others).

The vascular endothelial growth factor (VEGF, henceforth) is a signal protein playing a key role in the regulation of angiogenesis and matters for the growth of any type of blood vessels and in particular for the vascularization of cancer tumors [14, 37, 43]. The development of the tumor implies that some cells of the tumor may be in hypoxia due to a large distance to existing vascularization. This hypoxia fosters the production of VEGF, yielding high levels of circulating VEGF for a number of cancer patients [10]. This high level of VEGF activates the receptors of endothelial cell growth factors and ultimately leads to new blood vessels, tumor vascularization, and tumor growth [11].

VEGF plays a dual role in tumor proliferation. Besides its role in tumor angiogenesis, VEGF is also an immunosuppressive factor that dampens the body immune response, through two main channels [3, 13, 34]. First, T cells, which are suspected to play a major role in the immune response to cancer tumor [29, 33], are negatively affected by VEGF [48]. Second, high levels of circulating VEGF – as it is the case in tumor environments – also inhibit the immune action of dendritic cells [18–20, 32, 36].

Administering an anti-VEGF therapy will have a twofold impact on tumor growth [42]. On the one hand, an anti-angiogenic effect will perturb the new vasculature in tumor growth. On the other hand, dampening the immunosuppressive effect will favor the cytotoxic body response. For these two reasons, cancer therapeutic vaccines based on anti-VEGF are currently developed and investigated in clinical trials [22].

Anti-VEGF effects can be reinforced by immunotherapy and the injection of immune cells that can be natural or genetically modified. Typical examples of such cells are dendritic cells (DC). DC vaccine consists in re-injecting patient's dendritic cells after cell multiplication and activation with antigens of a specific tumor [40, 41]. The vaccines with DC injections, when used as a monotherapy, seem to face limitations [4], even though DC vaccines have benefited – and are benefiting – from significant advances [21]. One of the most promising roadmap seems to combine DC vaccines with other treatments, such as chemotherapy [12] among other options. In this respect, combining anti-VEGF with DC

vaccines, even though unexplored to a large extent, seems to be a promising route.

Soto-Ortiz and Finley [45] have developed a mathematical model to test in-silico the clinical benefits of the combination of anti-VEGF and DC injections. In-silico trials deliver two main results. First, injecting anti-VEGF only leads to tumor elimination under the condition that the protocol starts in a specific time window, while injecting DC only never yields tumor eradication. This result is in line with the clinical trial outcomes showing the failure of immunotherapy as monotherapy. Second, combining immunotherapy with anti-VEGF, and in particular proceeding to DC injections following anti-VEGF injections leads to tumor elimination and more importantly, the therapeutic window is significantly broader than with sole anti-VEGF injections. This last result indicates that a synergy between anti-VEGF and DC injections favors tumor elimination.

The whole analysis in [45] is based on a fixed protocol for anti-VEGF and DC injections. In the present paper, we use artificial intelligence methods to determine optimal protocols for anti-VEGF only and for the combination of anti-VEGF and DC. Our optimization embeds constraints that account for possible side effects of drug treatment (see [2, 31] for reviews of side effects of anti-angiogenic therapies). We base our constraints on the benchmark protocols for DC and anti-VEGF injections. In particular, if we allow for dose and schedule variations in the optimization, the daily and total doses should not exceed their benchmark counterpart, while the time interval between two injections cannot be shorter than in the benchmark protocols.

Regarding sole anti-VEGF injections, our results are twofold. First, we show that if we allow for the same total anti-VEGF dose, as the benchmark protocol, optimizing schedule and daily dosing enables to increase the therapeutic window by 39 days. In other words, an optimized protocol enables to postpone the tumor discovery by 39 days, which corresponds to a tumor diameter increase of 1.5 centimeters. Second, for any diagnosis date belonging to the therapeutic window of the standard protocol, our optimization enables to diminish the total anti-VEGF quantity required for tumor eradication. For instance, the optimized protocol leads to tumor elimination with as few as 58% of the total standard dose when the tumor is detected very early. Overall, optimizing anti-VEGF administration offers either a broader therapeutic window and possibly a later diagnosis date, or a sizable reduction in the total drug dose required for tumor elimination.

Regarding the drug combination, the optimized protocol similarly offers a sizable reduction in total drug dosing. Indeed, tumor elimination can be reached with the combination of 25% of total anti-VEGF dose and of 10% of total DC dose. Optimizing the administration of drug combination significantly diminishes drug doses required for eliminating the tumor.

Finally, our contribution is also methodological since we illustrate the possibilities of importing artificial intelligence techniques in mathematical immunotherapy. More precisely, the algorithms we use belong to the family of Monte-Carlo Tree Search (MCTS, henceforth) algorithms, that have been popularized by their application in AlphaGo – the software that has defeated a number of Go champions in 2016-2017. See [6] for survey on MCTS methods, [44] for its application to Go software and [27] for applications in oncology. Other optimizations have been performed in immunotherapy. Optimal control theory has been used to determine optimal dosage [8, 9], but, contrary to our method, this approach requires the underlying mathematical model to be simplified. Moreover, some authors, as in [38] for instance, have relied on genetic algorithms to determine the optimal timing of injections, while the number and the dosage of the injections are fixed. Our approach enables to simultaneously optimize on the timing, the number and the dosage of injections. Finally, mathematical models have also been used for personalizing immunotherapy protocols (see [30] and [1] for a review).

2 Materials and methods

2.1 Model

The model we rely on stems from Soto-Ortiz and Finley [45]. This model simultaneously takes into account at the tumor level, angiogenesis, the body immune response and the immunosuppressive effects of the tumor. It builds on Robertson-Tessi and coauthors [39] for the interactions between tumor and the immune system, and on Cameron and Davis [7] for angiogenesis modeling.

The model mechanisms for the tumor growth and the immune system reactions in [39] can be summarized as follows. The tumor growth follows an hybrid dynamics between an exponential and a power law. Unlicensed DC collect antigen produced by tumor cells.[1] Once licensed, DC interact with three types of T cells: effector CD8+ T cells, which kill tumor cells; regulatory T cells (Tregs), which are the heart of the immunosuppressive effects; and helper T cells, which license DC and might be converted in Tregs. Tregs have a twofold immunosuppressive role. First, they directly harm the tumor cytotoxic activity of effector T cells. Second, they also produce suppressive cytokines TGF-β and IL-10 – which are also produced by tumor cells. These two cytokines suppress the activity of effector T cells.[2] Overall, the evolution of the tumor is determined by two competing

[1] We follow the terminology of [45] and immature DC are said to be unlicensed. Symmetrically, mature DC are said to be licensed.

[2] For further detail about immunosuppression, see [47] for the role of Tregs, [23] for the role of TGF-β

forces, which are on one side the cytotoxic activities of effector T cells and on the other side, the immunosuppressive effects of T-regs and cytokines. Three possible regimes are possible for tumor evolution: tumor suppression, unbounded tumor growth, and control where the tumor size remains constant, its growth being balanced by cytotoxic effects, net of immunosuppressive activity. A key parameter for the determination of the final tumor regime is the tumor's antigenicity, which characterizes the intensity in production of the antigen by the tumor and the magnitude of the response of the immune system to the antigen.

The angiogenesis modeling, that builds on Cameron and Davis [7], can be summed up as follows. The tumor growth positively depends on the length of tumor vasculature, and therefore on angiogenesis. Angiogenesis itself depends on three growth factors: VEGF, Angiopoietin-1 (Ang-1) and -2 (Ang-2). The primary growth factor is VEGF, whose role on endothelial cells has been explained in Introduction. Ang-1 contributes to mature vasculature by reinforcing connections between endothelial cells. Ang-2 is a natural antagonist to Ang-1 and normally contributes to vasculature regression. However, Ang-2 effects are reversed in presence of VEGF and Ang-2 then contributes to vessel sprouting. VEGF has therefore both a direct and an indirect effect in angiogenesis. Finally, in addition to the angiogenic effects of VEGF, Soto-Ortiz and Finley [45] also model its immunosuppressive effect – which is absent of Robertson-Tessi and coauthors [39] – similarly to the one of TGF-β.

2.2 Simulations

We use the same calibration as [45], for parameters and initial values. Two parameter values deserve particular attention: antigenicity and speed of tumor growth. We set the value of the tumor antigenicity parameter to 1×10^{-5}, which is the lowest value considered in [45]. This value corresponds to a low anti-tumor immune response. The tumor growth parameter is set to 0.69 (day^{-1}), which corresponds to a fast replication speed of tumor cells, and therefore to an aggressive cancer. This parameter can be set to any value in the range $[0.1 - 1]$ (in day^{-1}) and the larger the parameter value, the faster the replication of tumor cells. For the very low antigenicity parameter value that we set equal to 1×10^{-5}, values of tumor growth higher than 0.5 unambiguously leads to an unbounded tumor growth in absence of treatment. As a result, picking up a very low antigenicity and a fast tumor growth corresponds to a "realistic worst-case scenario" (as stated in [45]) for investigating the effectiveness of possible anti-tumor treatments and their combination.

and [28] for the role of IL-10.

Following [45], we consider two standard protocols, one for anti-VEGF – which typically consists of bevacizumab – and another one for unlicensed DC injections. The standard anti-VEGF protocol consists of 6 biweekly injections of 7×10^8 nanograms of anti-VEGF. This protocol is in line with phase-2 clinical trials involving bevacizumab [49]. The standard DC protocol consists of 15 biweekly injections of 5×10^6 unlicensed DC. This DC protocol is also consistent with typical DC regimen [5]. We will henceforth refer to these two protocols as *standard protocols*, while their schedule will be referred to as *standard schedules*. We will additionally consider variations around these standard protocols, in which we will prune the total dose and consider a total dose equal to $x\%$ of the standard dose. Hence, these pruned protocols follow a standard schedule and the only difference with standard protocols is that they are stopped after a total dose of $x\%$ of the standard total dose is reached. For instance, if the maximal dose is 80% of the standard dose, the protocol will consist of 12 biweekly injections of 5×10^6 unlicensed DC. We will refer to these protocols as *pruned standard protocols*.

We simulate the mathematical model over a time window of 4,000 days. All computations are implemented in C++. The evaluation of the protocol efficacy is solely based on the number of tumor cells at day 4,000. A high efficacy corresponds to a small number of tumor cells at day 4,000. In line with [45], we will additionally say that a protocol is *successful* if the number of tumor cells at day 4,000 is strictly smaller than one. In that case, the tumor will be considered to be eradicated.[3]

2.3 Optimization

The optimization aims at determining protocols which offer the best efficacy and therefore yield the smallest number of tumor cells at day 4,000. Optimal protocols will be parametrized by the diagnosis date and the maximal total drug quantity. The diagnosis date can be interpreted at the date at which the tumor is discovered. In other words, the diagnosis date is a lower bound on the first injection date. The main reason motivating the choice of the diagnosis date rather the starting date is that it simplifies the interpretation of the optimal protocol. Consider an optimal protocol characterized by a diagnosis date t and a given total drug dose D. By construction, all protocols with the same diagnosis date t but a total dose higher than D will lead to a number of tumor cells smaller than with the initial protocol. Dually, all protocols with the same total dose D but a diagnosis date

[3]The tumor dynamics is continuous, such that the number of tumor cells is not restricted to be a natural number. In particular, it might be the case that in some simulations the number of cells drops below 1 at some date, while the number of cells at day 4,000 is above one, meaning a protocol failure. We address this question in Section 2 of the Appendix.

earlier than t will also lead a number of tumor cells at day 4,000 smaller than with the initial protocol. As shown in [45], the latter interpretation is not valid when we parametrize protocol by the standard definition of starting date.

The optimization will consist in choosing for a given diagnosis date and a given maximal total drug dose, the injection schedule (injection dates and doses) that minimizes the number of tumor cells at day 4,000. Additionally, all optimizations are subject to a set of constraints that aim at limiting drug side effects. Loosely speaking, these constraints prevent injections from being more frequent or more concentrated than in the standard schedules and can be formulated as follows:

C.1 the delay between two injections cannot be smaller than 14 days, which is the time interval between two injections in the standard protocol;

C.2 the doses per injection cannot exceed the daily doses in the standard protocol (7×10^8 nanograms for anti-VEGF, 5×10^6 cells for unlicensed DC).

Finally, our optimization relies on a MCTS algorithm, adapted from the one used in [27] for chemotherapy. Since the optimization bears upon schedules and quantities, the dimensionality of the optimization exercise is too high for the optimization to be conducted with standard techniques, such as dynamic programming. Using artificial intelligence techniques allows us to overcome the curse of dimensionality and to compute optimal protocols. We conduct two separate optimization exercises, one for sole anti-VEGF injections and another one for the combination of anti-VEGF and unlicensed DC.

A detailed presentation of our algorithm, though in a different framework, can be found in [27]. The mechanics of the algorithm can be summarized as follows. We start with a patient with a tumor of a given mass. The evolution of this patient is simulated using the PK/PD model until a decision can be made – at each day – when the algorithm is required to determine the optimal treatment dose. The dose is typically a continuous variable lying in a compact interval of the form $[0, d_{\max}]$, where 0 corresponds to no treatment and $d_{\max}$ to the maximum tolerated dose. The algorithm considers a discretization of the dosing space, with typically 10 to 20 possible treatment doses (including obviously no-treatment and the maximal dose). The algorithm then computes the best dose among these discrete alternatives. To do so, it associates to each possible dose a fictive patient, which is an exact copy of the current patient to which the treatment dose under consideration has been administered. These fictive patients are then simulated applying a default continuation treatment policy (see below for further details about this simulation). Finally, the recommended dose is the one corresponding to the patient with the most favorable outcome after simulation – typically, the smallest tumor mass. The recommended dose is then administered to the

actual patient, which is simulated to the next day. The algorithm ends when it reaches the final horizon date. The tricky part of the algorithm is the choice of default continuation treatment policy. In this paper, the default policy is a standard protocol with starting dates (for anti-VEGF and DC when relevant) determined by a genetic algorithm.

We now discuss the two separate optimization exercises.

Optimal anti-VEGF protocols. As discussed previously, any optimal anti-VEGF protocol is parametrized by a diagnosis date and a total anti-VEGF quantity. More precisely, every optimal anti-VEGF protocol will be denoted $AV(t, av)$ where:

- $t \in (0, 4000)$ in days is the diagnosis date;
- $av \in [0\%, 100\%]$ characterizes the maximal total dose of anti-VEGF, relatively to the standard total dose equal to 4.2×10^9 nanograms.

In other words – and in addition to constraints C.1 and C.2 given above –, in a protocol $AV(t, av)$, as the tumor is diagnosed at day t, the treatment cannot start before day t, and the total anti-VEGF dose cannot exceed $av\%$ of the total standard dose, equal to 4.2×10^9 nanograms.

We report in Figure 1 examples of standard and optimized protocols. The two bottom panels represent the injection pattern as a function of time for the standard and optimal protocols respectively. The top panel plots the corresponding evolution of the number of tumor cells in the case of the two protocols. The diagnosis date is day 580 and the total injection is 100% of standard total dose. With our notation, the optimized protocol therefore corresponds to $AV(580, 100\%)$.

In Figure 1, we can observe the 6 identical biweekly injections of the standard protocol (middle panel), that implies the tumor size evolution of the top panel (gray line). The final tumor size at day 4,000 amounts approximately to 340 cells. This number of tumor cells exceeding the one-cell threshold, the standard protocol is considered non successful. This is in line with the finding of [45], who show that day 580 is outside of the therapeutic window. The optimized protocol $AV(580, 100\%)$ corresponds to a very different injection pattern that is displayed in the bottom panel. If the two first injections are the same (for both the timing and the quantity) as in the standard protocol, the following ones are different. The optimal protocol tends to diminish per-injection quantities but to increase the number of injections. Unsurprisingly, the constraint bearing on the total anti-VEGF quantity is binding, such that the total drug quantity is identical in both protocols. Besides having a proper injection schedule, the optimized protocol outcome greatly differs from the standard one. Indeed, the tumor size at day 4,000 is approximately equal to 0.1 cell. This quantity

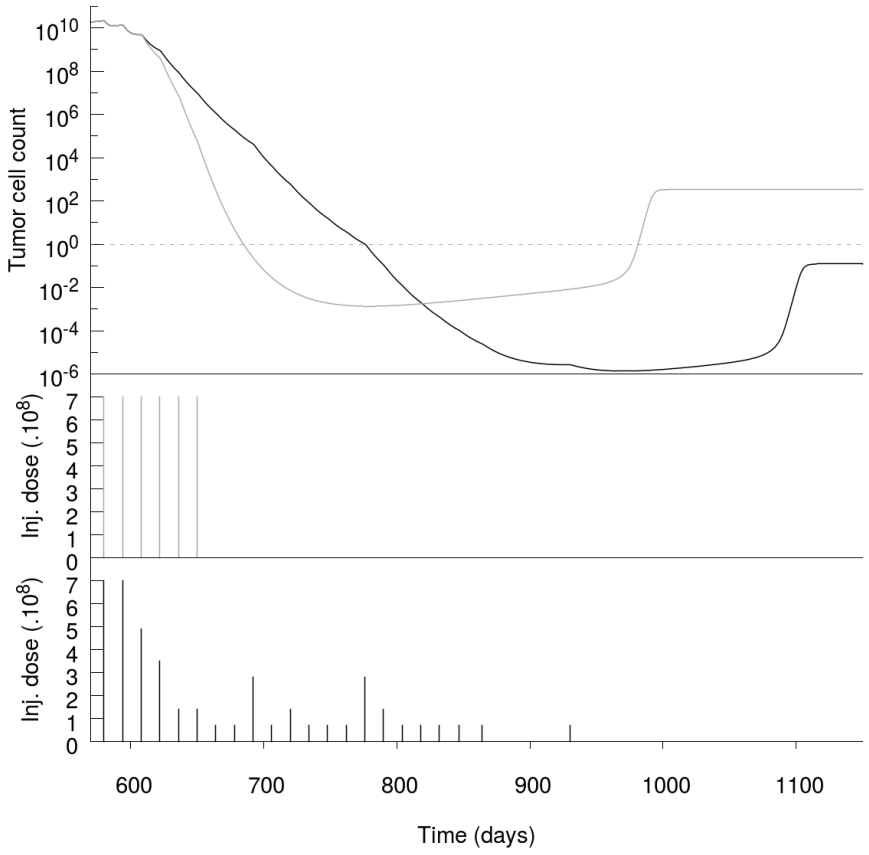

Figure 1: Example of standard and optimal anti-VEGF protocols for a starting date at day 580 and 100% of the total standard dose.
The top panel plots the evolution of tumor cells over the protocol period for the standard protocol (in gray) and the optimized protocol $AV(580, 100\%)$ (in black).
The middle panel plots the injection pattern for the standard protocol, while the bottom panel plots the injection pattern for the optimized one.

being below the one-cell threshold, the protocol can be considered successful. With the optimal protocol, the tumor is eradicated, while this is not the case with the standard protocol.

Optimal combined protocols. We now discuss optimal protocols in the case of a combination of anti-VEGF and unlicensed DC. Every optimal combined protocol is parametrized by three variables: the diagnosis date, the total anti-VEGF quantity, and the total DC quantity. Of note, the two first variables are exactly the same as in the case of sole anti-VEGF injections. This reflects the fact that the protocol now embeds injections of both anti-VEGF and unlicensed DC. More precisely, every optimal combination protocol will be denoted $C(t, av, dc)$ where:

- $t \in (0, 4000)$ and $av \in [0\%, 100\%]$ have the exact same meaning as for the sole anti-VEGF injection;
- $dc \in [0\%, 100\%]$ characterizes the total dose of unlicensed DC, relatively to the standard total dose, amounting to 7.5×10^7 cells;

In other words, in a protocol $C(t, av, dc)$, the treatment is diagnosed at – and cannot start before – day t, the total anti-VEGF dose cannot exceed $av\%$ of the total standard dose of anti-VEGF and the total DC dose cannot exceed $dc\%$ of the total DC standard dose. Of course, constraints C.1 and C.2 given above still apply. Finally, let us note that anti-VEGF and DC injections are independent of each other and we do not impose any constraint in the sequence of injections.

We report in Figure 2 an example of optimized protocol, which is $C(0, 30\%, 15\%)$, starting at date 0, with 30% of the total standard dose of anti-VEGF and 15% for the dose of DC. The injection patterns for DC and anti-VEGF as a function of time are plotted in the two bottom panels. The top panel plots the corresponding evolution of the number of tumor cells. We observe that the protocol involves a number of tumor cells at day 4,000 approximately equal to 2.5×10^{-5}. Consequently, this number being smaller than one, the protocol can be considered to be successful and to yield tumor eradication.

3 Results

As seen in the example of Figure 1 for injection of anti-VEGF only, the optimized protocol can possibly – for the same total drug quantity and the same diagnosis date – yield tumor eradication, while it is not the case for the standard protocol. We now provide

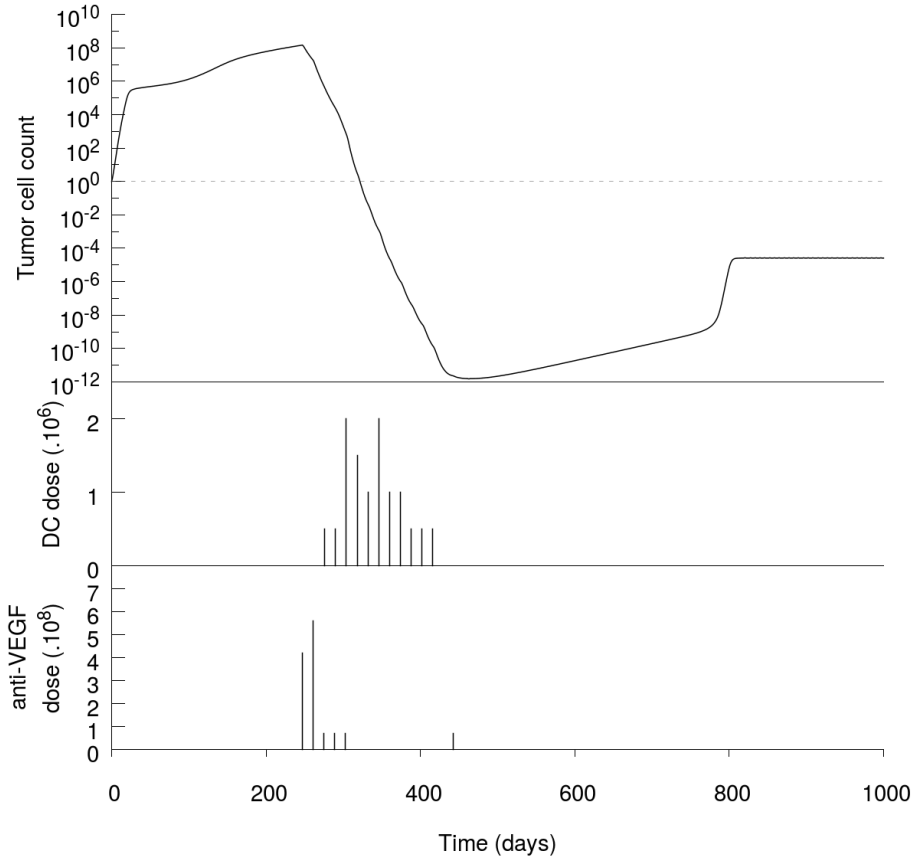

Figure 2: Example of an optimal combined protocol $C(0, 30\%, 15\%)$, for a diagnosis date at day 0, 30% of the standard anti-VEGF dose and 15% of the standard DC dose.

general results in the case both of anti-VEGF solely and of the combination of anti-VEGF and DC.

Anti-VEGF protocols. Our main results for sole anti-VEGF injections are plotted in Figure 3, where we represent the areas where optimal and standard protocols are successful. More precisely, on the left hand-side panel 3a, we plot for any diagnosis date t, the smallest total anti-VEGF dose av, for which an optimal protocol $AV(t, av)$ (black line) is successful and yields tumor eradication. Of note, we do not represent the actual final number of tumor cells, and as in [45], we solely focus on whether protocols are successful or not. For instance, the right-most point of Figure 3 for optimized protocols corresponds to $AV(580, 100\%)$ presented in Figure 1. Consistently with tumor size evolutions presented

on the top panel of Figure 1, $AV(580, 100\%)$ is successful, while it is not the case for its standard counterpart. We also draw the same plot for "pruned" standard protocols. The pruned standard protocols we consider are constructed such that the total dose is the minimal dose for which the tumor is eradicated with a standard injection pattern. Furthermore, in order to make both standard and optimal protocols more comparable, we also use the convention of diagnosis date, instead of protocol starting date.[4] On the right hand side panel 3b, we plot for each tumor diameter, the smallest total anti-VEGF dose for which an optimal protocol with that dose yields tumor eradication. We do the same plot for a pruned standard protocol constructed as before. We deduce panel 3b from panel 3a as follows. For any day, we compute the number of tumor cells at that day using the mathematical model of the tumor evolution (without treatment) and we then deduce the tumor diameter using a simple geometric formula.[5]

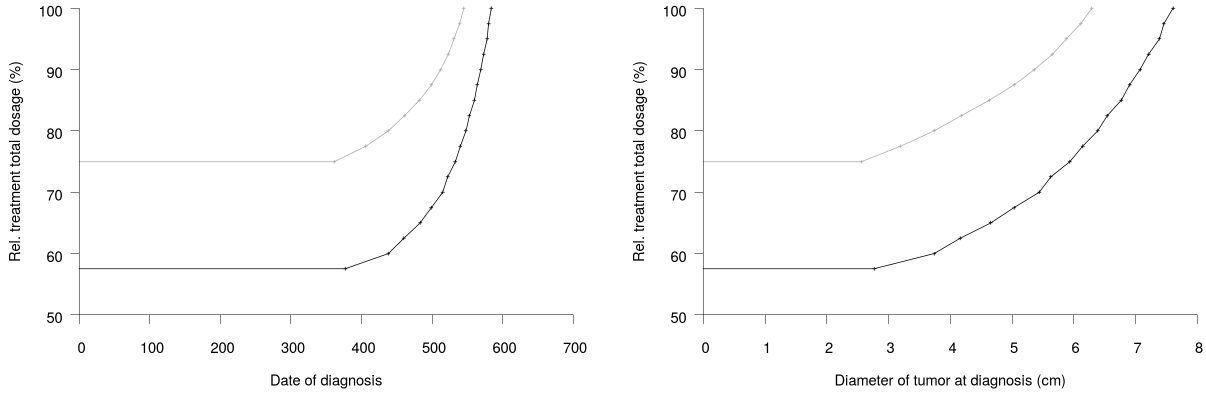

(a) Smallest total anti-VEGF dose required for tumor eradication as a function of the diagnosis date.

(b) Smallest total anti-VEGF dose required for tumor eradication as a function of the tumor diameter.

Figure 3: Comparison of optimized and pruned standard protocols.
For every diagnosis date or every tumor size, we compute the smallest total drug dose that yields tumor eradication with an optimized and a standard protocol. In both graphs, the black curve refers to the optimal protocols and the gray curve to standard protocols.

We can draw two main lessons from Figure 3. First, we observe that if tumor eradication

[4]There is therefore no contradiction with the results of [45], who report that starting the standard protocol early does not yield tumor elimination. For early diagnosis dates in our graph, the injections in the standard protocol will therefore start at a later date.

[5]The diameter of a tumor d is assumed to be related to the number of tumor cells, n, by the following formula $d \approx 2 \times \left(\frac{3n}{4\pi B}\right)^{1/3}$, where $B = 10^8$ (in cells/cubic centimeter) characterizes the density of the tumor in tumor cells. The formula can be found in [45].

can be achieved by a pruned standard protocol for a given diagnosis date, tumor eradication can also be achieved by an optimal protocol with a smaller total anti-VEGF dose. The gain in the total dose is approximately constant for all diagnosis dates, and amounts approximately to 17% of the standard dose for early diagnosis and 20% for the latest ones. Of note, for the sake of fairness of the comparison between both protocols, we have not used the standard dose for the standard schedule, but the smallest share of the standard dose that yields tumor eradication for a pruned standard protocol. For instance, for early diagnosis dates, pruned standard protocols require at least 75% of the standard dose for tumor eradication, while the optimized protocol only requires 58% of the total dose. For the latest diagnosis dates (in which both protocols are still successful), the pruned standard protocol requires 100% of the standard dose, while the optimized one only needs 80%.

The second lesson from Figure 3 is that for a given total drug dose, the optimized protocol offers a later diagnosis date, or equivalently a larger tumor mass at the diagnosis. However, the higher the total dose, the smaller the postponement of diagnosis dates. For instance, for a 75% dose, the diagnosis date can be postponed by 172 days – or almost 6 months – from day 361 to day 533, while for the 100% dose, the diagnosis date can be postponed by 39 days, from day 545 to day 584. The gain is also very sizable in terms of tumor diameter at diagnosis. For a 75% dose, the tumor diameter at diagnosis must be below 2.56 centimeters for the standard protocol, while it can be almost as large as 5.94 centimeters with the optimal protocol, which is more than twice larger. For a 100% dose, the tumor diameter at diagnosis must be smaller than 6.29 centimeters with the standard protocol, while it can reach 7.61 centimeters with an optimal protocol. Optimization therefore offers a sizable increase in the tumor size at diagnosis.

We summarize all above results in Table 1.

Dose	Gain in diagnosis date	Gain in tumor size (days)
75%	172 days	3.38 cm
100%	39 days	1.32 cm

Table 1: Gains in diagnosis length and tumor size when switching from standard to optimal protocol, for two specific doses.

Optimal combined protocols. We now turn to the optimal combined protocols. In this case, the dimensionality of the optimization problem is much larger than the one for sole anti-VEGF protocols. Consequently, we determine for every pair of total drug dose for anti-VEGF and DC, if we can find a successful protocol yielding tumor elimination. In other words, for standard schedules, we seek two starting dates for the standard schedules

of both anti-VEGF and DC injections, such that the combined treatment yields tumor elimination. For optimized protocols, for any given drug quantities av and dc, we look for the optimal protocol $C(0, av, dc)$ and check whether it yields to tumor elimination. We plot our results on Figure 4. For instance, the gray square with 30% of standard anti-VEGF dose and 15% of standard DC dose corresponds to the optimized protocol $C(0, 30\%, 15\%)$ presented in Figure 2. Consistently with the tumor size evolution plotted on Figure 2, this protocol is successful.

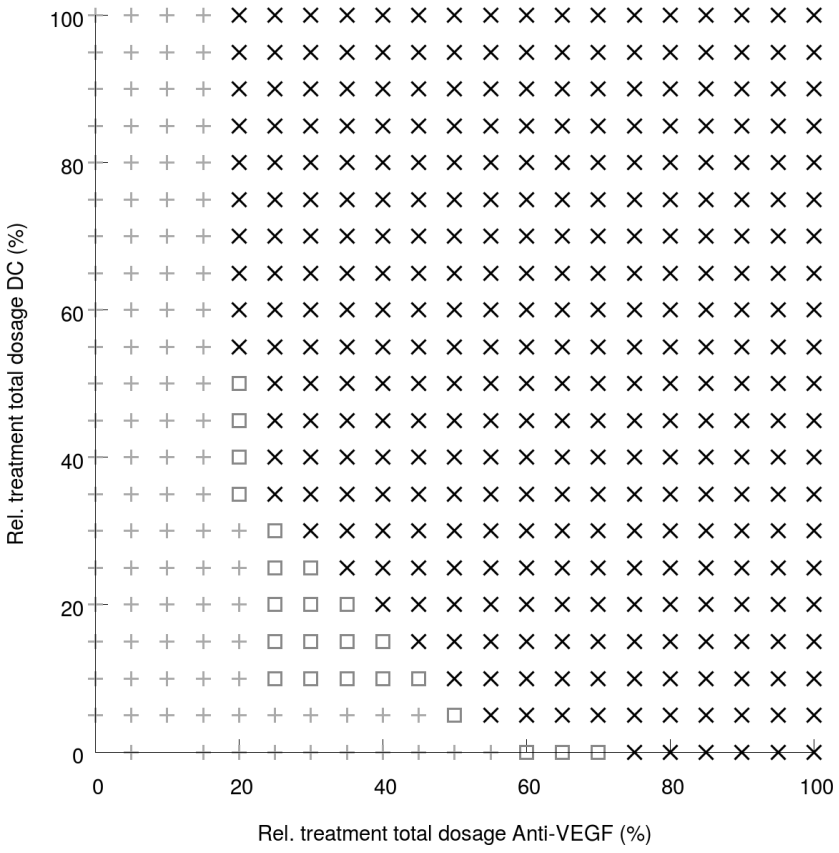

Figure 4: Existence of protocols yielding tumor elimination for all pairs of quantities of anti-VEGF and DC. Drug quantities vary between 0% and 100% of the total standard quantity with a step of 5% for each drug.
A gray plus (+) indicates that neither optimized nor standard schedule yields tumor elimination. A gray square (□) indicates that only optimized protocol yields tumor elimination. A black times (×) indicates that both optimized and standard protocols yield tumor elimination.

We can draw several conclusions from Figure 4. First, we recover part of our results for sole anti-VEGF injections. Indeed, in line with our results for sole anti-VEGF injections

(corresponding to a combined protocol with a zero DC dose), we observe that for the standard schedule to yield tumor eradication, the total anti-VEGF dose must be greater than 75% of the standard anti-VEGF quantity, while the optimized schedule is successful as soon as the anti-VEGF dose is above 60% of the standard total dose. Second, we observe that the optimization also enables to take advantage of the complementarity between DC and anti-VEGF injections. An optimized protocol can eliminate tumor with only 25% of standard anti-VEGF dose and 10% of total DC dose. In consequence, adding only a small amount of DC with the proper optimized schedule enables to significantly reduce the dose of anti-VEGF that is required. The 10% dose in DC enables to diminish the anti-VEGF dose from 60% to 25% of the standard dose.

These results are summed up in Table **??**.

	Anti-VEGF dose	
DC dose	Optimal protocol	Standard protocol
0%	60%	75%
10%	25%	60%

Table 2: Anti-VEGF dose needed for tumor elimination.

4 Discussion

Importing artificial intelligence techniques for optimization purposes in immunotherapy has delivered two main outcomes. First, an optimized schedule allows for decreasing the total drug use, both for sole anti-VEGF and for the combination of anti-VEGF and DC, without impacting the overall outcome and tumor elimination. As a result, optimizing the schedule may help contribute to reduce drug adverse effects by diminishing doses, as well as making these immunotherapy treatments more affordable by decreasing the total treatment cost. Second, the optimized schedule enables to yield tumor elimination, while it would be too late or the tumor would be too large to do so for a standard schedule, even with a full dose. As a consequence, the optimized schedule offer a longer therapeutic window and makes tumor elimination possible for later diagnosis dates.

In this article, we have developed an algorithm based on the family of MCTS algorithms. Three main reasons have motivated our choice to work with such an algorithm. First, this algorithm takes advantage of the sequential and finite time features of the problem at hand. In particular, it allows us to handle a full-fledged PK/PD model without any simplification. Second, the algorithm does not require an evaluation function that would be tricky to design in the highly non-linear framework we work with. Finally, our algorithm offers a very high

degree of flexibility. This feature is particularly interesting for future implementation in an applied contexts. A first possible application would be to consider informational issues. For instance, the PK/PD parameters' values of patients may not be known with certainty and have to be inferred from information arrival during treatment. As in [27], our algorithm may easily be coupled with Bayesian learning. Another possible application would be to consider probabilistic constraints. Such constraints may for instance reflect the possible absence of patients on a treatment day. The protocol will then account for these contingencies.

References

[1] Z. Agur, K. Halevi-Tobias, Y. Kogan, and O. Shlagman. Employing dynamical computational models for personalizing cancer immunotherapy. *Expert Opin Biol Ther*, 16(11):1373–1385, Nov 2016.

[2] B. Al-Husein, M. Abdalla, M. Trepte, D. L. Deremer, and P. R. Somanath. Antiangiogenic therapy for cancer: an update. *Pharmacotherapy*, 32(12):1095–1111, Dec 2012.

[3] G. Bergers and L. E. Benjamin. Tumorigenesis and the angiogenic switch. *Nat. Rev. Cancer*, 3(6):401–410, Jun 2003.

[4] K. F. Bol, G. Schreibelt, W. R. Gerritsen, I. J. de Vries, and C. G. Figdor. Dendritic cell-based immunotherapy: State of the art and beyond. *Clin. Cancer Res.*, 22(8):1897–1906, Apr 2016.

[5] P. Brossart, S. Wirths, G. Stuhler, V. L. Reichardt, L. Kanz, and W. Brugger. Induction of cytotoxic T-lymphocyte responses in vivo after vaccinations with peptide-pulsed dendritic cells. *Blood*, 96(9):3102–3108, Nov 2000.

[6] C. Browne, E. Powley, D. Whitehouse, S. Lucas, P. I. Cowling, S. Tavener, D. Perez, S. Samothrakis, and S. Colton. A survey of monte carlo tree search methods. *IEEE Trans. Comput. Intellig. and AI in Games*, 4(1):1–43, 2012.

[7] M. A. Cameron and A. L. Davis. A Mathematical Model of Angiogenesis in Glioblastoma Multiforme. Technical report, Arizona State University.

[8] A. Cappuccio, M. Elishmereni, and Z. Agur. Optimization of interleukin-21 immunotherapeutic strategies. *J. Theor. Biol.*, 248(2):259–266, Sep 2007.

[9] F. Castiglione and B. Piccoli. Cancer immunotherapy, mathematical modeling and optimal control. *J. Theor. Biol.*, 247(4):723–732, Aug 2007.

[10] H. F. Dvorak. Vascular permeability factor/vascular endothelial growth factor: A critical cytokine in tumor angiogenesis and a potential target for diagnosis and therapy. *Journal of Clinical Oncology*, 20(21):4368–4380, 2002.

[11] H. F. Dvorak, L. F. Brown, M. Detmar, and A. M. Dvorak. Vascular permeability factor/vascular endothelial growth factor, microvascular hyperpermeability, and angiogenesis. *Am. J. Pathol.*, 146(5):1029–1039, May 1995.

[12] L. A. Emens and G. Middleton. The interplay of immunotherapy and chemotherapy: harnessing potential synergies. *Cancer Immunol Res*, 3(5):436–443, May 2015.

[13] N. Ferrara. The role of VEGF in the regulation of physiological and pathological angiogenesis. *EXS*, (94):209–231, 2005.

[14] N. Ferrara and W. J. Henzel. Pituitary follicular cells secrete a novel heparin-binding growth factor specific for vascular endothelial cells. *Biochem. Biophys. Res. Commun.*, 161(2):851–858, Jun 1989.

[15] J. Folkman. Tumor angiogenesis: therapeutic implications. *N. Engl. J. Med.*, 285(21):1182–1186, Nov 1971.

[16] J. Folkman. The role of angiogenesis in tumor growth. *Semin. Cancer Biol.*, 3(2):65–71, Apr 1992.

[17] J. Folkman. Role of angiogenesis in tumor growth and metastasis. *Semin. Oncol.*, 29(6 Suppl 16):15–18, Dec 2002.

[18] D. Gabrilovich, T. Ishida, T. Oyama, S. Ran, V. Kravtsov, S. Nadaf, and D. P. Carbone. Vascular endothelial growth factor inhibits the development of dendritic cells and dramatically affects the differentiation of multiple hematopoietic lineages in vivo. *Blood*, 92(11):4150–4166, 1998.

[19] D. I. Gabrilovich, H. L. Chen, K. R. Girgis, H. T. Cunningham, G. M. Meny, S. Nadaf, D. Kavanaugh, and D. P. Carbone. Production of vascular endothelial growth factor by human tumors inhibits the functional maturation of dendritic cells. *Nat. Med.*, 2(10):1096–1103, Oct 1996.

[20] D. I. Gabrilovich, T. Ishida, S. Nadaf, J. E. Ohm, and D. P. Carbone. Antibodies to vascular endothelial growth factor enhance the efficacy of cancer immunotherapy by improving endogenous dendritic cell function. *Clinical Cancer Research*, 5(10):2963–2970, 1999.

[21] A. D. Garg, P. G. Coulie, B. J. Van den Eynde, and P. Agostinis. Integrating next-generation dendritic cell vaccines into the current cancer immunotherapy landscape. *Trends Immunol.*, 38(8):577–593, 08 2017.

[22] J. V. Gavilondo, F. Hernandez-Bernal, M. Ayala-Avila, A. V. de la Torre, J. de la Torre, Y. Morera-Diaz, M. Bequet-Romero, J. Sanchez, C. M. Valenzuela, Y. Martin,

K. H. Selman-Housein, A. Garabito, O. C. Lazo, J. V. Gavilondo, F. Hernandez-Bernal, M. Ayala Avila, A. de la Torre, J. Soriano, A. V. de la Torre, Y. Martin, A. Garabito, E. L. de Armas, E. Ibanez, L. J. Lopez, A. I. Zamora, R. Ortiz, E. Cardenas, U. Vazquez, B. Hernandez, T. Bartomeu, R. Moya, M. Hurtado, J. de la Torre, K. H. Selman-Housein, L. Borrero, A. M. Gonzalez, O. Diaz, K. Aroche, M. Catala, M. Izquierdo, J. Pinero, O. Ramirez, Y. Cruz, J. Jordan, M. Mesa, R. Gonzalez, Y. Gonzalez, J. Aguero, R. Llanes, D. Quincoses, F. Hernandez-Bernal, C. Valenzuela, O. Lazo, C. Bermudez, R. Hernandez, I. Campa, E. Garcia, A. Herrera, G. Melo, I. Raices, K. Cruz, P. Lopez-Saura, J. V. Gavilondo, M. Ayala Avila, Y. Morera, M. Bequet-Romero, J. Sanchez, L. Perez, A. Musacchio, M. Perez, G. Hernandez, S. Padron, L. Trimino, A. Sanchez, and D. Perez. Specific active immunotherapy with a VEGF vaccine in patients with advanced solid tumors. results of the CENTAURO antigen dose escalation phase I clinical trial. *Vaccine*, 32(19):2241–2250, Apr 2014.

[23] L. Gorelik and R. A. Flavell. Immune-mediated eradication of tumors through the blockade of transforming growth factor-beta signaling in T cells. *Nat. Med.*, 7(10):1118–1122, Oct 2001.

[24] A. Hoeben, B. Landuyt, M. S. Highley, H. Wildiers, A. T. Van Oosterom, and E. A. De Bruijn. Vascular endothelial growth factor and angiogenesis. *Pharmacol. Rev.*, 56(4):549–580, Dec 2004.

[25] L. Holmgren, G. Jackson, and J. Arbiser. p53 induces angiogenesis-restricted dormancy in a mouse fibrosarcoma. *Oncogene*, 17(7):819–824, Aug 1998.

[26] L. Holmgren, M. S. O'Reilly, and J. Folkman. Dormancy of micrometastases: balanced proliferation and apoptosis in the presence of angiogenesis suppression. *Nat. Med.*, 1(2):149–153, Feb 1995.

[27] N. Houy and F. Le Grand. Optimal dynamic regimens with artificial intelligence: The case of temozolomide. *PLoS ONE*, 13(6):e0199076, 2018.

[28] C. L. Hsieh, D. S. Chen, and L. H. Hwang. Tumor-induced immunosuppression: a barrier to immunotherapy of large tumors by cytokine-secreting tumor vaccine. *Hum. Gene Ther.*, 11(5):681–692, Mar 2000.

[29] K. L. Knutson and M. L. Disis. Tumor antigen-specific T helper cells in cancer immunity and immunotherapy. *Cancer Immunol. Immunother.*, 54(8):721–728, Aug 2005.

[30] Y. Kogan, K. Halevi-Tobias, M. Elishmereni, S. Vuk-Pavlovi?, and Z. Agur. Reconsidering the paradigm of cancer immunotherapy by computationally aided real-time personalization. *Cancer Res.*, 72(9):2218–2227, May 2012.

[31] Y. Kubota. Tumor angiogenesis and anti-angiogenic therapy. *Keio J Med*, 61(2):47–56, 2012.

[32] S. Kusmartsev and D. I. Gabrilovich. Immature myeloid cells and cancer-associated immune suppression. *Cancer Immunol. Immunother.*, 51(6):293–298, Aug 2002.

[33] D. Mougiakakos, A. Choudhury, A. Lladser, R. Kiessling, and C. Johansson. Chapter 3 - regulatory t cells in cancer. volume 107 of *Advances in Cancer Research*, pages 57 – 117. Academic Press, 2010.

[34] J. E. Ohm and D. P. Carbone. VEGF as a mediator of tumor-associated immunodeficiency. *Immunol. Res.*, 23(2-3):263–272, 2001.

[35] M. S. O'Reilly, L. Holmgren, C. Chen, and J. Folkman. Angiostatin induces and sustains dormancy of human primary tumors in mice. *Nat. Med.*, 2(6):689–692, Jun 1996.

[36] T. Oyama, S. Ran, T. Ishida, S. Nadaf, L. Kerr, D. P. Carbone, and D. I. Gabrilovich. Vascular endothelial growth factor affects dendritic cell maturation through the inhibition of nuclear factor-κB activation in hemopoietic progenitor cells. *The Journal of Immunology*, 160(3):1224–1232, 1998.

[37] J. Plouet, J. Schilling, and D. Gospodarowicz. Isolation and characterization of a newly identified endothelial cell mitogen produced by AtT-20 cells. *EMBO J.*, 8(12):3801–3806, Dec 1989.

[38] M. Qomlaqi, F. Bahrami, M. Ajami, and J. Hajati. An extended mathematical model of tumor growth and its interaction with the immune system, to be used for developing an optimized immunotherapy treatment protocol. *Math Biosci*, 292:1–9, Oct 2017.

[39] M. Robertson-Tessi, A. El-Kareh, and A. Goriely. A mathematical model of tumor-immune interactions. *J. Theor. Biol.*, 294:56–73, Feb 2012.

[40] G. Schuler. Dendritic cells in cancer immunotherapy. *Eur. J. Immunol.*, 40(8):2123–2130, Aug 2010.

[41] G. Schuler, B. Schuler-Thurner, and R. M. Steinman. The use of dendritic cells in cancer immunotherapy. *Curr. Opin. Immunol.*, 15(2):138–147, Apr 2003.

[42] S. Selvaraj, M. Raundhal, A. Patidar, and B. Saha. Anti-VEGF antibody enhances the antitumor effect of CD40. *Int. J. Cancer*, 135(8):1983–1988, Oct 2014.

[43] D. R. Senger, S. J. Galli, A. M. Dvorak, C. A. Perruzzi, V. S. Harvey, and H. F. Dvorak. Tumor cells secrete a vascular permeability factor that promotes accumulation of ascites fluid. *Science*, 219(4587):983–985, Feb 1983.

[44] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587):484–489, January 2016.

[45] L. Soto-Ortiz and S. D. Finley. A cancer treatment based on synergy between anti-angiogenic and immune cell therapies. *J. Theor. Biol.*, 394:197–211, Apr 2016.

[46] H. M. Verheul and H. M. Pinedo. The role of vascular endothelial growth factor (VEGF) in tumor angiogenesis and early clinical development of VEGF-receptor kinase inhibitors. *Clin. Breast Cancer*, 1 Suppl 1:S80–84, Sep 2000.

[47] C. T. Viehl, T. T. Moore, U. K. Liyanage, D. M. Frey, J. P. Ehlers, T. J. Eberlein, P. S. Goedegebuure, and D. C. Linehan. Depletion of CD4+CD25+ regulatory T cells promotes a tumor-specific immune response in pancreas cancer-bearing mice. *Ann. Surg. Oncol.*, 13(9):1252–1258, Sep 2006.

[48] T. Voron, E. Marcheteau, S. Pernot, O. Colussi, E. Tartour, J. Taieb, and M. Terme. Control of the immune response by pro-angiogenic factors. *Front Oncol*, 4:70, 2014.

[49] J. J. Vredenburgh, A. Desjardins, J. E. Herndon, J. M. Dowell, D. A. Reardon, J. A. Quinn, J. N. Rich, S. Sathornsumetee, S. Gururangan, M. Wagner, D. D. Bigner, A. H. Friedman, and H. S. Friedman. Phase II trial of bevacizumab and irinotecan in recurrent malignant glioma. *Clin. Cancer Res.*, 13(4):1253–1259, Feb 2007.