



HAL
open science

Explicit approximation of stochastic optimal feedback control for combined therapy of cancer

Mazen Alamir

► **To cite this version:**

Mazen Alamir. Explicit approximation of stochastic optimal feedback control for combined therapy of cancer. ACC 2024 - American Control Conference, IEEE, Jul 2024, Toronto, Canada. hal-04428200

HAL Id: hal-04428200

<https://hal.science/hal-04428200>

Submitted on 31 Jan 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Explicit approximation of stochastic optimal feedback control for combined therapy of cancer

Mazen Alamir

Abstract—In this paper, a tractable methodology is proposed to approximate stochastic optimal feedback treatment in the context of mixed immuno-chemotherapy therapy of cancer. The method uses a fixed-point value iteration that approximately solves a stochastic dynamic programming-like equation. It is in particular shown that the introduction of a variance-related penalty in the latter induces better results that cope with the consequences of softening the health safety constraints in the cost function. The convergence of the value function iteration is revisited in the presence of the variance related term. The implementation involves some Machine Learning tools in order to represent the optimal function and to perform complexity reduction by clustering. Quantitative illustration is given using a commonly used model of combined therapy involving twelve highly uncertain parameters.

Index Terms—Stochastic optimal control; Stochastic Dynamic Programming, Cancer therapy, Machine Learning, Clustering.

I. INTRODUCTION

Rationalizing drug delivery is an active research field that commonly involves population models with high number of uncertain parameters. Unfortunately, these parameters which are by nature highly variable between individuals, are inaccessible to identification because of the lack of excitation [9].

The great majority of applied mathematical-like works solve deterministic optimal control problems which are formulated using the nominal values of the parameters [15], [6], [11]. While this might be important to draw qualitative conclusions regarding the patterns of optimal strategies (intensive treatment, presence of singular arcs, etc); the resulting profiles do not accommodate for the high dispersion of the parameters.

Using repetitive solutions of such open-loop nominal scheduling strategies in a feedback mode through the receding-horizon principle (apply the first part of the optimal strategy and recompute a new optimal injection profile at the next decision instant and so on leading to the so called Model Predictive Control design (MPC).) obviously reduces the drawback of parameters mismatch [4], [5]. However, the performances can only be observed a posteriori as nothing is explicitly done to address the presence of uncertainties.

Another option is to use a parameterized state feedback law and to optimally tune its design parameters using explicitly the statistical description of the model's parameters dispersion by means of the randomized optimization framework [1]. Unfortunately, as far as available works are concerned, this is to be done for each initial state in the proposed frameworks.

Stochastic model predictive control (SMPC) offers an attractive alternative through the use of a set of scenarios in the on-line computation in order to minimize the approximate version of the expectation of the cost function [13], [12], [2]. This lead to a solution for each initial state but does not give a global overview on the performance on a whole region of the state space. On the other hand, this solution can be scalable in the dimension of the state which enables the use of almost arbitrarily complex models of the underlying dynamics.

Stochastic Dynamic Programming (SDP) [7], [3], [10] is a framework that, at least conceptually, outperforms the previous approaches in that it gives an 1) explicit state feedback law that 2) explicitly incorporates the statistical description of the parameters dispersion while being 3) defined on a whole region of interest within the state space.

Unfortunately, the price to pay to get these advantages is to represent the Bellman function over the extended space of state and control making this approach **non scalable** in the dimension of the state and control. Moreover, when minimizing the residual of the SDP equation, each call involves the computation of statistically defined quantities and this, for the computation to be relevant, should theoretically involve a high number of randomly sampled instances of the uncertain model's parameters. Note however that this last drawback is obviously shared by the SMPC approach.

The literature on solving the stochastic dynamic programming equation is huge and a complete survey of it is out of the scope of this paper. For a complete and recent survey, one can consult [3]. This paper does not claim a particularly novel algorithm to solve the SDP equation. Rather it focuses on the issue of introducing a variance-related term in the definition of the cost function, investigates its impact on handling the health-related constraint in the combined therapy of cancer and revisits the convergence proof of the underlying value function-related fixed-point iteration when such additional variance-related term is added.

More precisely, a simple framework is proposed for a moderate sized models¹ in which Machine Learning (ML) tools are used. More precisely, the model's structure uses a Support Vector Machine (SVM) regression model to represent the value function. On the other hand, a clustering approach is used to approximate the statistical quantities using a lower number of samples. This leads to a tractable approximate solution that can be obtained in less than 20 minutes when using a grid of 9604 points in the state/control space when no parallelization is used. Drastic reduction of the computation

This work was supported by the MIAI @ GrenobleAlpes under Grant ANR-19-P31A-003.

The author is with the University of Grenoble Alpes, CNRS, Grenoble INP, GIPSA-lab, 38000 Grenoble, France (e-mail: mazen.alamir@grenoble-inp.fr).

¹the combined therapy model used involves four states and two controls which are commonly encountered sizes in all the related works.

time can be obtained since the proposed scheme is amenable to massive parallelization.

This paper is organized as follows: The problem of combined therapy is stated in Section II. Some recalls regarding SDP are proposed in Section III and some results are proposed for the convergence of a fixed-point iteration that is used to solve the associated equations. Section IV details the proposed approximate solution's framework while Section V shows the results on the combined therapy of cancer and discusses some issues. Finally, the paper ends with Section VI that summarizes the paper's contributions and gives some hints for further investigation.

II. PROBLEM STATEMENT

A. The dynamic model

Let us consider the population model used in [8] to describe the dynamics involved in the combined immuno-chemo therapy:

$$\dot{x}_1 = ax_1(1 - bx_1) - c_1x_4x_1 - k_3x_3x_1 \quad (1)$$

$$\dot{x}_2 = -\delta x_2 - k_2x_3x_2 + s_2 \quad (2)$$

$$\dot{x}_3 = -\gamma_0x_3 + u_2 \quad (3)$$

$$\dot{x}_4 = g\frac{x_1}{h+x_1}x_4 - rx_4 - p_0x_4x_1 - k_1x_4x_3 + s_1u_1 \quad (4)$$

where

x_1	tumor cell population
x_2	circulating lymphocytes population
x_3	chemotherapy drug concentration
x_4	effector immune cell population
u_1	rate of introduction of immune cells
u_2	rate of introduction of chemotherapy

The description of the role of each groups of term is given in Table I for the sake of clarity.

Eq.	Term	Description
(1)	$ax_1(1 - bx_1)$	Logistic tumor growth
(1)	$-c_1x_4x_1$	Death of tumor due to effector cells
(1)	$-k_3x_3x_1$	Death of tumor due to chemotherapy
(2)	$-\delta x_2$	Death of circulating lymphocytes
(2)	$-k_2x_3x_2$	Death of lymphocytes due to chemo
(2)	s_2	Constant source of lymphocytes
(3)	$-\gamma_0x_3$	Exponential decay of chemotherapy
(4)	$g\frac{x_1}{h+x_1}x_4$	Stimulation of tumor on effector cells
(4)	$-rx_4$	Death of effector cells
(4)	$-p_0x_4x_1$	Inactivation of effector cells by tumor
(4)	$-k_1x_4x_3$	Death of effector cells due to chemo

TABLE I

SIGNIFICATION OF THE TERMS INVOLVED IN THE DYNAMIC MODEL (1)-(4) [SOURCE [8]]

When referred to the nominal values of the parameters involved in the model (1)-(4), the following values are used:

$$\begin{aligned} a &= 0.25, b = 1.02 \times 10^{-14}, c_1 = 4.41 \times 10^{-10}, g = 1.5 \times 10^{-2} \\ h &= 20.2, k_2 = k_3 = 0.6, k_1 = 0.8, p_0 = 2 \times 10^{-11}, r = 0.04 \\ s_1 &= 1.2 \times 10^7, s_2 = 7.5 \times 10^6, \delta = 1.2 \times 10^{-2}, \gamma_0 = 0.9 \end{aligned}$$

Note that the dynamic model (1)-(4) involves 14 parameters. It is worth noting however that the order of magnitude of x_1 is about 10^9 while $h \approx 20$. This means that h acts only when the tumor is almost disappearing². That is the reason why in the remainder of this paper, this parameter is supposed to be known and hence it is not included in the set of uncertain parameters. As it is shown later, this enables a separable structure of the evolution map to be used [see (6)].

Gathering all the other parameters in a vector $p \in \mathbb{R}^{13}$ and using some discretization scheme with some *sufficiently small* sampling period τ , it is possible to put the dynamics above in the following condensed form for the easiness of notation:

$$x^+ = f(x, u, p) \quad (x, u, p) \in \mathbb{R}^4 \times \mathbb{R}^2 \times \mathbb{R}^{13} \quad (5)$$

$$= \Phi(x, u)\Psi(p) \quad (6)$$

for straightforward definition of Φ and Ψ .

The trajectory of the system starting at initial state x under a control profile u and when the parameter value p holds is denoted by $x^u(k|x, p)$, namely:

$$x^u(k+1|x, p) = f(x^u(k|x, p), u(k), p)$$

$$x^u(0|x, p) = x$$

B. The control objective and constraints

The aim of the combined therapy is to reduce the size of the tumor population x_1 while keeping the level of lymphocyte cells x_2 above a lower bound x_2^{min} . This has to be done using quantized drug delivery:

$$u \in \mathbb{U} := \{0, u_1^{max}\} \times \{0, u_2^{max}\} \quad (7)$$

so that an infinite sequence of control lying inside \mathbb{U} is denoted hereafter by \mathbb{U}^∞ .

In order to address the control objective, the following stage cost function can be used that incorporates the constraint on x_2 as a soft constraint:

$$L(x, u) := x_1^2 + \rho_c \max\{0, x_2^{min} - x_2\} + \rho_1 u_1 + \rho_2 u_2 \quad (8)$$

so that **if the model's parameters were perfectly known**, an optimization problem can be defined for all given initial state x as follows:

$$\mathcal{P}(x|p) :$$

$$\min_{u \in \mathbb{U}^\infty} J(u|x, p) := \sum_{k=0}^{\infty} \gamma^k L(x^u(k|x, p), u(k)) \quad (9)$$

for some $\gamma \in (0, 1]$. Now, since the parameters are supposed to be unknown, the stochastic control approach amounts at replacing the deterministic cost function (9) by a statistically defined one such as:

$$J_\alpha(u|x, \Pi) := \mu(J(u|x, \cdot)) + \alpha\sigma(J(u|x, \cdot)) \quad (10)$$

$$=: \mathcal{S}_\alpha^{(J)}(u, x) \quad (11)$$

where $\mu(\cdot)$ and $\sigma(\cdot)$ stand respectively for the expectation and the variance of their argument when p is sampled according to

²On the other hand, without h , the effector immune population will keep growing while the tumor is absent which contradicts its very definition as a tumor stimulated immune population.

some supposedly known Probability Density Function (pdf) Π . Note that (11) simply introduces a short notation of the r.h.s of (10).

This leads to the following stochastic optimization problem:

$$\mathcal{P}_\alpha(x|\Pi) : \min_{\mathbf{u} \in \mathbb{U}^\infty} J_\alpha(\mathbf{u}|x, \Pi) \quad (12)$$

Note that in many stochastic control-related works, attention is focused on the expectation of the cost function so that $\alpha = 0$ is widely used in (10). This is not totally appropriate in the case of cancer therapy. Indeed, expectation is a relevant indicator only when high number of realizations are expected to take place³ in which only the global average is of interest. In the case of cancer therapy, a scenario is a patient being treated. This is why one should obviously be interested in the risk that **each patient afford** during the treatment. In such cases, it is precisely those *bad* scenarios with non negligible likeliness that really matter. This has to do not only with the expectation but definitively with the associated variance as well. This explains the use of $\alpha > 0$ in (12).

In the next section, some recalls are proposed on the Stochastic Dynamic Programming equations that might be invoked to approximately solve the stochastic optimization problem (12).

III. RECALLS AND PRELIMINARY RESULTS

A. Stochastic Dynamic Programming

SDP attempts to solve the following functional equation in which the unknown function $V(\cdot)$ is the optimal solution of (12). More precisely:

$$V(x) := \min_{u \in \mathbb{U}} Q(x, u) \quad \text{where} \quad (13)$$

$$Q(x, u) := L(x, u) + \gamma \min_{v \in \mathbb{U}} \left[\mathcal{S}_\alpha^{(Q)}(x, u, v) \right] \quad (14)$$

$$\text{where } x^+ = f(x, u, p) \quad (15)$$

in which $\mathcal{S}^{(Q)}(x, u, v)$ is defined in a similar way as $\mathcal{S}^{(J)}(\mathbf{u}, x)$ [see (11)]:

$$\mathcal{S}_\alpha^{(Q)}(x, u, v) := \mu(Q(f(x, u, \cdot), v)) + \alpha \sigma(Q(f(x, u, \cdot), v)) \quad (16)$$

where here again, $\mu(\cdot)$ and $\sigma(\cdot)$ are respectively the expectation and the variance of their argument when p (involved in the definition of $f(x, u, p)$) is sampled according to the pdf Π .

The fact that V satisfying (13)-(14) is a solution to (12) is a direct consequence of the Bellman principle and the fact that $L = \mu(L) + \sigma(L) = L + 0$ since $L(x, u)$ is a deterministic p -unrelated term.

From the above, it comes out that the truly unknown function to find is $Q(\cdot, \cdot)$ since the optimal value function $V(\cdot)$ as well as the optimal control are then recovered from the static optimization problem (13).

Now it is not hard to see from (14) that Q is a solution of a fixed-point iteration:

$$Q^{(i+1)} = F(Q^{(i)}) \quad (17)$$

where the operator F (commonly called the Bellman operator) is defined by:

$$F(Q)(x, u) := L(x, u) + \gamma \min_{v \in \mathbb{U}} \left[\mathcal{S}^{(Q)}(f(x, u, \cdot), v) \right] \quad (18)$$

In the next section, some convergence results regarding the above fixed-point iterations are derived in a general conceptual setting before a specific implementation is proposed in Section IV for the combined therapy of cancer.

B. Fixed-Point iteration convergence analysis

For the sake of brevity, all the proofs of the forthcoming stated results can be examined in the arXiv version of the present paper (<https://arxiv.org/abs/1905.04937>). Nevertheless, the working assumptions and the results are stated in this section. First of all, the following continuity assumption of the dynamics is used:

Assumption 1 (Continuous dynamics): The map f describing the discrete-time state evolution (5) is continuous in its arguments.

Regarding the uncertain vector p , the following assumption is needed

Assumption 2 (Finitely supported uncertainties): there exists a compact set \mathbb{P} to which belong all possible realizations of the uncertain vector p .

Definition 3.1 (Bounded excursion): A map G defined on $\mathbb{X} \times \mathbb{U}$ is said to have *B-bounded excursions w.r.t the dynamics* f if the following inequality holds:

$$|G(f(x, u, p), v) - \mu(G(f(x, u, \cdot), v))| \leq B \quad (19)$$

for all $(x, u, v, p) \in \mathbb{X} \times \mathbb{U}^2 \times \mathbb{P}$.

This definition simply states that the excursion of the realizations $G(f(x, u, p), v)$ from the mean value $\mu(G(f(x, u, \cdot), v))$ is bounded. It goes without saying that, under Assumption 1, such a bound exists as soon as the map G is continuous. since all the involved bounding sets are compact.

Lemma 3.1 (Fixed-Point convergence): Under Assumptions 1 and 2, if the fixed point map F invoked in (17) produces only **continuous** maps $Q^{(i)}$ with *B-bounded excursions* for some $B > 0$, then the fixed point iteration converges on $\mathbb{Z} := \mathbb{X} \times \mathbb{U}$ provided that the penalty α on the variance term in (16) satisfies:

$$\alpha < \frac{1 - \gamma}{2\gamma B} \quad (20)$$

and for sufficiently small sampling period τ that is used to derive the discrete-time dynamics (5). \diamond

From the above lemma, one can clearly derive the standard convergence result that is known to hold when no penalty on the variance is used (see for instance [7], [10], [3] and the references therein).

It is worth underlying that the convergence results cited above depend on the assumption according to which the successive iterates $Q^{(i)}$ produced by the Bellman map F admit a *B-bounded excursions* with common upper bound B . The way this condition is enforced is discussed in the next section.

³such as saving energy in building management for instance despite of bad knowledge of exogenous parameters such as power demands and whether conditions.

IV. THE PROPOSED COMPUTATION FRAMEWORK

A. Parametrization of $Q(x, u)$

In the previous section, the SDP equation and the related fixed-point iteration have been defined in a conceptual framework. In this framework, the unknown function $Q(x, u)$ one is looking for has no specific finite-dimensional representation that would be compatible with concrete computation schemes. In this section, this parameterization aspect is addressed.

For the sake of simplicity, the notation $z = (x, u)$ and $\mathbb{Z} := \mathbb{X} \times \mathbb{U}$ is used so that $Q(z)$, $f(z, p)$ stands for $Q(x, u)$ and $f(x, u, p)$ respectively. The finite dimensional parameterization of Q is obtained through the following steps:

✓ Consider a fixed grid of points $\mathcal{Z} := \{z^{(i)}\}_{i=1}^{n_g}$ is defined on the compact subset \mathbb{Z} of interest.

✓ Denote by $\mathbf{q} := \{q_i\}_{i=1}^{n_g}$ the vector of values of the function Q at the grid points, namely $q_i = Q(z^{(i)})$.

✓ Choose some regression model (Polynomial, Gaussian, Support Vector Machine Regressor (SVMR), etc.) and denoted it by \hat{Q} . As far as this work is concerned, a SVMR with gaussian basis is used with the default parameters used in the scikit-learn Machine Learning (ML) library [14].

✓ This choice enables to associate to each vector of values $\mathbf{q} \in \mathbb{R}^{n_g}$ a function $\hat{Q}(\cdot|\mathbf{q})$ defined on \mathbb{Z} which is precisely the one identified (or *learned* to invoke a ML vocabulary) using the learning data $(\mathcal{Z}, \mathbf{q})$. Based on the above items, it is now possible to state the following definition:

Definition 4.1 (Admissible regression model): A regression model is said to be admissible if and only if for all set \mathcal{Z} of finite grid points in \mathbb{Z} and any compact set of function values $\mathbb{Q} \subset \mathbb{R}^{n_g}$, there exists $B > 0$ such that all identified maps $\hat{Q}(\cdot|\mathbf{q})$ with $\mathbf{q} \in \mathbb{Q}$ shows B -bounded excursions in the sense of Definition 3.1.

B. Main convergence result

Using the above preliminary results and definitions, the following main result regarding the convergence of the proposed Fixed-Point iteration scheme can be stated as follows:

Proposition 4.1 (Main result): Under Assumptions 1 and 2, if an admissible regression model (in the sense of Definition 4.1) is used over some grid points $\mathcal{Z} := \{z^{(i)}\}_{i=1}^{n_g}$ with the initialization $q_i = L(z^{(i)})$ to solve the stochastic optimal feedback of the combined therapy of cancer as stated in Section II, then the resulting fixed-point iteration (18) with $\gamma \in (0, 1)$ is convergent provided that the variance related penalty α and the sampling period τ are taken sufficiently small. \diamond

The only remaining task is to find a regression model that is admissible in the sense of Definition 4.1. Fortunately, there are many options. The one that is used in the remainder of

this paper is the Support Vector Machine regressor (SVMR) which shows the following structure:

$$\hat{Q}(z|\mathbf{q}) := \beta_0(\mathbf{q}) + \sum_{i=1}^{n_g} \beta_i(\mathbf{q}) G(z^{(i)}, z) \quad (21)$$

where the coefficients $\beta_i(\mathbf{q})$ admits uniform bound if \mathbf{q} belongs to some compact set \mathbb{Q} .

C. Computation of the expectation and variance

The fixed-point iteration (18) requires the computation of the map $\mathcal{S}^{(Q)}(f(x, u, \cdot), v)$ which involves the computation of the expectations and variances of quantities over the probability space of parameter values. In this section, the way this is done concretely is explained.

Recall that the r.h.s of the dynamics in (6) can be decomposed into two multiplicative terms $x^+ = \Phi(x, u)\Psi(p)$ where only the second term is parameter-dependent. Consequently, the computation of the expectation of x^+ needs the statistics of $\Psi(p) \in \mathbb{R}^{14}$ to be approximated. This is done by identifying n_{cl} clusters based on a set \mathcal{S}_ψ of n_s samples of the vector $\Psi(\cdot)$ that are drawn using the pdf Π . The centers of the computed clusters, denoted by $\Psi^{(j)}$, $j = 1, \dots, n_{cl}$ can be used as representatives of their clusters. Moreover, the ratio π_j between the sizes of these clusters and the total population n_s can be used as a measure of their probabilities, namely

$$\pi_j := \frac{1}{n_s} \text{card} \left\{ \psi \in \mathcal{S}_\psi \mid \text{cl}(\psi) = \text{cl}(\Psi^{(j)}) \right\} \quad (22)$$

Based on the above definitions, the following expressions are used in the approximation of terms in (16):

$$\hat{\mu}(x, u, v|\mathbf{q}) := \sum_{j=1}^{n_{cl}} \pi_j \hat{Q}(\Phi(x, u)\Psi^{(j)}, v|\mathbf{q}) \quad (23)$$

$$\hat{\sigma}(x, u, v, |\mathbf{q}) := \sum_{j=1}^{n_{cl}} \pi_j \left[\hat{Q}(\Phi(x, u)\Psi^{(j)}, v|\mathbf{q}) - \hat{\mu}(x, u, v|\mathbf{q}) \right]^2 \quad (24)$$

Using these expression, the weighted sum (16) can be approximated by:

$$\hat{\mathcal{S}}_\alpha^{(Q)}(x, u, v|\mathbf{q}) := \hat{\mu}(x, u, v|\mathbf{q}) + \alpha \hat{\sigma}(x, u, v|\mathbf{q}) \quad (25)$$

The updated values of q_i at the underlying grid points $z^{(i)}$ can therefore be obtained through:

$$q_i^+ := L(z^{(i)}) + \gamma \min_{v \in \mathbb{U}} \left[\hat{\mathcal{S}}_\alpha^{(Q)}(z^{(i)}, v|\mathbf{q}) \right] \quad (26)$$

for each element $z^{(i)}$ of the underlying grid \mathcal{Z} . Note that the optimization problem in v can be solved by simple enumeration over the four elements of the set \mathbb{U} .

This implements a fixed-point iteration on the vector of values \mathbf{q} than can be shortly written as follows:

$$\mathbf{q}^+ = F(\mathbf{q}) \quad (27)$$

where the notation F used in (17) is overloaded using the finite-dimensional parametrization \mathbf{q} of Q . This is the fixed-point iteration that is proved to be convergent under the assumptions of Proposition 4.1.

In Section V, the instantiation of the framework on the combined therapy of cancer is detailed and the corresponding results are commented.

V. NUMERICAL INVESTIGATION ON THE COMBINED THERAPY OF CANCER

A. Parameters used in the numerical investigation

As far as stochastic controller design is concerned (nominal controller is also investigated), the above framework has been investigated using the following parameters. A number $n_s = 10^4$ samples have been drawn to perform the clustering of set of values of Ψ (see Section IV-C). The number of identified clusters was taken equal to $n_{cl} = 20$.

Regarding the pdf Π , the following definition has been used $p_i = (1 + \nu)p_i^{nom}$ where ν is a normal distribution with 0 mean and variance $\sigma = 0.4$. The grid \mathcal{Z} has been taken as the cartesian product of a uniform grid on x (having 7 uniformly distributed values over the normalized interval $[0, 1]$ per component) and the four values control set \mathbb{U} . This induces a set \mathcal{Z} composed of $7^4 \times 4 = 9604$ elements.

Clustering has been performed using the scikit-learn library's KNN (k-nearest neighbors) utility. The scikit learn library's SVM regression model has been used with radial basis function's kernel and the default values of the remaining parameter. The sampling period $\tau = 0.25$ Day is used and the lower bound $x_2^{min} = 0.05$ is used for the normalized⁴ lymphocyte population size.

Three different controllers are compared. Namely:

- 1) The nominal controller for which Π is a dirac distribution centered at the nominal vector of parameters. The is referred to as Controller #1.
- 2) The expectation-based controller for which the distribution of the parameter is the one described above V-A where no penalty on the variance term ($\alpha = 0$). The is referred to as Controller #2.
- 3) The mixed expectation/variance-based controller for which the penalty $\alpha = 0.1$ is used on the variance-related term. The is referred to as Controller #3.

B. Results

Despite the fact that the convergence of the fixed-point iteration has been proved only for $\gamma \in (0, 1)$ and for sufficiently small α and τ , the value $\gamma = 1$ has been used to show that convergence does occur for this ideal value (no discount on the cost function). On the other hand $\alpha = 0.1$ for Controller #3.

First of all, Figure 1 shows a typical convergence curve for the fixed-point iteration. The difference between two successive solutions during the iterations, namely $\|\mathbf{q}^{(i+1)} - \mathbf{q}^{(i)}\|_\infty$ is shown. Note that because of the use of $\gamma = 1$, no strict contraction is obtained but eventually, contraction is settled and the solution is obtained.

⁴The state components have been normalized using respectively the reference values $(10^9, 10^9, 1, 10^9)$.

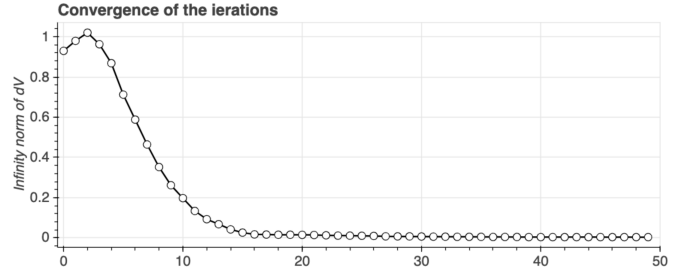


Fig. 1. Typical convergence curve for fixed-point iterations: Evolution of the difference $\|\mathbf{q}^{(i+1)} - \mathbf{q}^{(i)}\|_\infty$ between two successive iterations. $\gamma = 1$, $\alpha = 0.1$.

Figures 2 and 3 show the benefit from using the stochastic formulation compared to the nominal one. In order to show this, $N_{sam} := 20000$ different simulations are performed by drawing random samples of pairs of initial states and parameter vectors and the resulting closed-loop cost function is simulated before statistics are drawn. More precisely, 100 initial states are sampled and for each of these initial state, 200 random parameter vectors samples are drawn. Namely, for each given pair $(x^{(i)}, p^{(i)})$, $i = 1, \dots, N_{sam}$

$$J_{cl}(x^{(i)}, p^{(i)}) := \sum_{k=0}^{N_{sim}} L(x(k), u(k)) \quad (28)$$

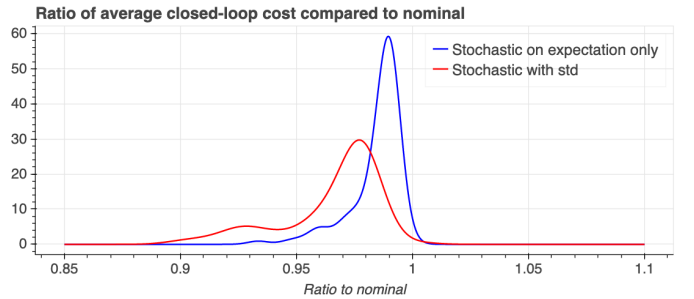


Fig. 2. Comparison between the closed-loop performance ratios relatively to the nominal controller for the two stochastic formulations for the 20000 randomly sampled initial states and parameter vectors.

where $N_{sim} = 50$ is the simulation horizon length (12.5 Days) while $x(k)$ and $u(k)$ stand for the closed-loop trajectories at instant k starting at initial instant $k = 0$ at $x^{(i)}$ and when the parameter vector $p^{(i)}$ is used. Then different statistical comparisons are done when the above procedure is executed using the three controllers mentioned above, namely, the nominal, the stochastic controller based on the expectation term only ($\alpha = 0$) and finally the stochastic controller using a penalty $\alpha = 0.1$ on the variance.

More precisely, Figure 2 shows the histogram of the ratios between the closed-loop performance as defined in (28) when the stochastic controllers are used compared to the nominal. One can clearly observe the following features:

- 1) Both stochastic controllers lead to statistically better results than the nominal.

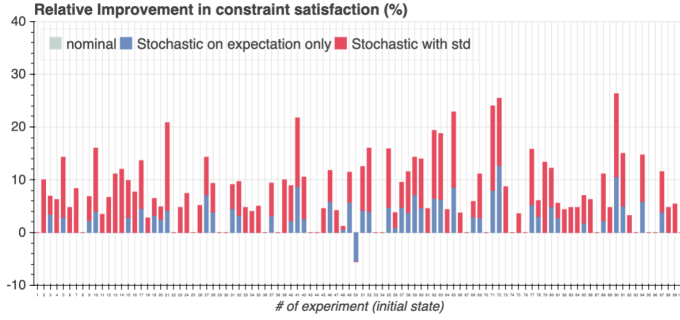


Fig. 4. Safety constraints satisfaction improvement relatively to the nominal controller for the two stochastic formulations. Abscissa represents the initial state number (100 are randomly selected) and the percentage refers to the 200 parameter vectors that are simulated for the same initial state.

- 2) The stochastic Controller #3 that incorporates penalty on the variance shows better results than the one using only the expectation-related penalty term.
- 3) It seems clear that even when focusing on the mean value of the cost function, the controller # 3 shows a better average than the controller # 2. This might seem contradictory with the fact that the latter is supposed to minimize this expectation. It can be conjectured that this comes from the fact that the expectation cannot be sufficiently well represented by the average over the 20 clusters (see Section IV-C) that are used to induce tractable computation of the expectation and the variance and considering the penalty on the variance correct this bias in a favorable way.

Figure 3 shows the comparison between the average levels of the lymphocytes population cells sizes (compared to the nominal) when the stochastic controllers are used. Here again, it comes out that the performances of the controller # 3 are slightly better than that of the second expectation-only-based controller.

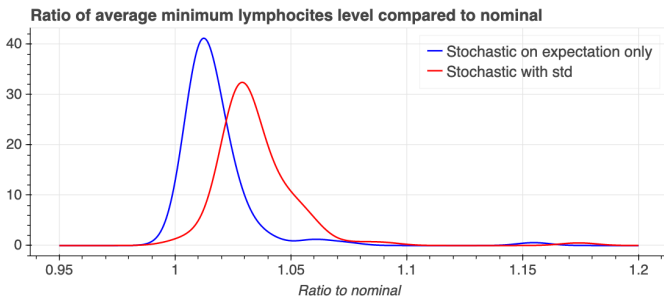


Fig. 3. Comparison between the average lymphocytes population level ratios relatively to the nominal controller for the two stochastic formulations for the 20000 randomly sampled initial states and parameter vectors.

Another way to look at the constraint satisfaction is to count the number of scenarios where the health constraint $x_2 \geq x_2^{min}$ is satisfied. Figure 4 shows the percentage of improvement in the constraints satisfaction compared to the nominal when the stochastic controllers are used.

VI. CONCLUSION AND FUTURE WORKS

In this paper a tractable stochastic control design for the combined therapy of cancer is proposed. The method is based on an approximate solution of the SDP equations through a value function fixed-point iteration. The convergence of the latter in the presence of variance related penalty in the cost function is proved under mild technical assumption. The result suggests that the inclusion of a variance-related term in the cost function might give better result by partially compensating for the errors induced by the approximation of the statistical quantities through computations involving a reduced number of clusters. Undergoing work involves the use of GPU-based computation framework to handle a higher order of magnitude of the samples size in order to examine how these qualitative results scale with the dimension of the populations over which they are evaluated.

REFERENCES

- [1] M. Almir. On probabilistic certification of combined cancer therapies using strongly uncertain models. *Journal of Theoretical Biology*, 384:59 – 69, 2015.
- [2] M. Almir. On the use of supervised clustering in stochastic NMPC design. *CoRR*, abs/1811.09069, 2018.
- [3] D. P. Bertsekas. *Dynamic programming and optimal control*. 2017.
- [4] S. Chareyron and M. Almir. Mixed immunotherapy and chemotherapy of tumors: Feedback design and model updating schemes. *Journal of Theoretical Biology*, 45:444–454, 2009.
- [5] Tao Chen, Norman F. Kirkby, and Raj Jena. Optimal dosing of cancer chemotherapy using model predictive control and moving horizon state/parameter estimation. *Computer Methods and Programs in Biomedicine*, 108(3):973 – 983, 2012.
- [6] Lisette G. de Pillis, Weiqing Gu, and Ami E. Radunskaya. Mixed immunotherapy and chemotherapy of tumors: modeling, applications and biological interpretations. *Journal of theoretical biology*, 238(4):841–862, 2006.
- [7] T. Jaakoola, M. I. Jordan, and S. P. Singh. On the convergence of stochastic iterative dynamic programming algorithms. *Journal of Machine Learning Research*, 6(6):1185–1201, 1994.
- [8] K. Kassara and A. Moustafid. Angiogenesis inhibition and tumor-immune interactions with chemotherapy by a control set-valued method. *Mathematical Biosciences*, 231(2):135 – 143, 2011.
- [9] Kanchi Lakshmi Kiran and S. Lakshminarayanan. Global sensitivity analysis and model-based reactive scheduling of targeted cancer immunotherapy. *Biosystems*, 101(2):117 – 126, 2010.
- [10] Robert Kirkby. Convergence of discretized value function iteration. *Computational Economics*, 49(1):117–153, Jan 2017.
- [11] U. Ledzewicz, H. Schattler, and A. d’Onofrio. Optimal control for combination therapy in cancer. In *47th IEEE Conference on Decision and Control*, 2008.
- [12] A. Mesbah. Stochastic model predictive control with active uncertainty learning: A survey on dual control. *Annual Reviews in Control*, 45:107 – 117, 2018.
- [13] A. Mesbah, S. Streif, R. Findeisen, and R. Braatz. Stochastic nonlinear model predictive control with probabilistic constraints. In *Proceedings of the American Control Conference ACC2014*, June 2014.
- [14] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [15] G. W. Swan. General applications of optimal control in cancer chemotherapy. *IMA Journal of mathematics Applied in Medicine & Biology*, 5:303–316, 1988.