Gradient method for solving singular optimal control problems

Mariusz Bodzioch[0000−0002−1991−5062]

Faculty of Mathematics and Computer Science,
University of Warmia and Mazury in Olsztyn, Poland
mariusz.bodzioch@matman.uwm.edu.pl

Abstract. Solving an optimal control problem consists in finding a control structure and corresponding switching times. Unlike in a bang-bang case, switching to a singular control perturbs the control structure. The perturbation of one of the switching times affects any subsequent singular intervals in the control, as the trajectories move along different singular arcs with different values of singular controls. It makes the problem of finding optimal solutions extremely difficult. In this paper, we discuss a gradient method for solving optimal control problems, when singular intervals are present in the optimal structure. The method is based on applying the necessary conditions of optimality given by the Pontryagin Maximum Principle, where the control variable enters the Hamiltonian linearly. To demonstrate the method, we formulate a nonlinear optimal control problem and then, using the proposed algorithm, we solve the problem and find the optimal control structure and corresponding switching times. Lastly, we compare the results with results obtained using three popular optimisation modelling languages: Pyomo, AMPL and JuMP. These languages serve as interfaces for solving the optimal control problem with the non-linear optimisation algorithm Ipopt. Our case study shows that the presented method not only computes the switching times accurately, but also moves precisely along the singular arc.

Keywords: Gradient method · Optimal control · Singular control · Mathematical modelling

1 Introduction

In general, optimal control deals with the problem of finding a control that achieves a certain optimality criterion. Indeed, the need to control the dynamics of various objects arises in engineering, biological, ecological, and medical applications. In population control problems, for example, the aim is to minimise (or maximise) the size of one of the considered populations, however, other control objectives can also be considered. Solving an optimal control problem consists in finding a control structure and corresponding switching times. When singular controls do not appear in the optimal structure, computing bang-bang controls is quite simple. However, the appearance of singular controls makes the problem of finding optimal solutions extremely difficult.

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63775-9_10

Numerical methods for optimal control can be categorised into two main types: indirect and direct methods. Indirect methods are based on the Pontryagin Maximum Principle, where the optimal control problem is transcribed to a Hamiltonian boundary-value problem and then solved numerically using a differential-algebraic equation solver. On the other hand, direct methods involve approximating the problem and transforming it into a nonlinear programming problem, which is subsequently solved using well-developed software such as Ipopt [13]. For an overview and comparison of direct and indirect methods, we refer to [1,7,8,10] and references therein.

Note that, even if the method accurately computes the switching times, the obtained control may exhibit oscillations within the singular region, i.e. the control oscillates infinitely many times between the bounds [14]. This typically occurs when the optimal control consists of singular and bang-bang intervals. Thus, when applying a computational method without a priori information about the optimal control structure, both direct and indirect methods may produce incorrect results. In this paper, we present a gradient method for solving singular control problems in the presence of singular intervals in the optimal structure.

As an example, we formulate an optimal control problem for a heterogeneous population dynamics model with a non-standard objective functional introduced by us in [2]. This functional additionally penalises the size of the control-resistant population. From a biological standpoint, this is related to control-resistant subpopulations, such as chemotherapy-resistant cancer cells, virus-resistant, antibiotic-resistant strains of bacteria, or any pesticide-resistant insect populations. In our previous papers we showed that the resistance penalty gives rise to locally-optimal singular controls, see e.g. [3,4,5], as well as [9]. From a practical perspective, singular controls correspond to low-control time-varying control schedules. We analyse a simple model that allows us to use the gradient algorithm with the minimum required mathematical complexity. However, the algorithm can be applied to more detailed and complex models, where the control variable enters the Hamiltonian linearly.

This paper is organised as follows. Firstly, we formulate an optimal control problem and recall some formulas, concepts and notions, related to singular control. Next, we discuss gradient methods for finding optimal controls in the bang-bang form and derive a modified gradient method for solving optimal control problems when singular intervals are present in the optimal structure. We illustrate the applicability of the method using a simple mathematical model. Finally, we numerically solve a nonlinear optimal control problem, using both direct methods and the derived indirect method. Conclusions are drawn based on the performed numerical simulations.

2 Optimal control problem

A dynamic optimisation problem, in which the state $\mathbf{n} = \mathbf{n}(t)$ is linked in time to a control function $u = u(t)$, is called an optimal control problem. More precisely, the solution to ordinary differential equations $\dot{\mathbf{n}} = F(\mathbf{n}, u(t))$ is shaped by the

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63775-9_10

control with an objective $J = J(u)$ that is optimised over all possible responses subject to external controls. Let us now formulate the problem to be considered.

The dynamics is defined by a continuously differentiable vector field $\mathbf{f}$ and control vector field $\mathbf{g}$ as

$$\dot{\mathbf{n}} = \mathbf{f}(\mathbf{n}) + \mathbf{g}(\mathbf{n})u(t), \tag{1}$$

$\mathbf{n} \in \mathbb{R}^n$, with initial condition $\mathbf{n}(0) = \mathbf{n}_0$. The vector field $\mathbf{f}$ represents the uncontrolled dynamics, while the vector field $\mathbf{g}$ represents the influence of the control on the system. We define the objective (or cost) functional in so-called Bolza form as the integral of a Lagrangian L and a penalty term M,

$$J(u(\cdot)) = M(T, \mathbf{n}(T)) + \int_0^T L(\mathbf{n}(t), u(t))dt, \tag{2}$$

where T is a fixed terminal time. Then, the optimal control problem is as follows: minimise the objective (2) over all admissible controls $u : [0,T] \to [0, u_{\max}]$ subject to dynamics (1) over an interval $[0,T]$.

Let $\mathbf{p} : [0,T] \to (\mathbb{R}^n)^*$. We define the so-called Hamiltonian function

$$H(\mathbf{p}, \mathbf{n}, u) = \mathbf{p}^T \left(\mathbf{f}(\mathbf{n}) + \mathbf{g}(\mathbf{n})u\right) + L(\mathbf{n}, u). \tag{3}$$

The Pontryagin Maximum Principle provides the first-order necessary conditions for optimality for our optimal control problem.

Theorem 1. *If u^* is an optimal control with corresponding trajectory $\mathbf{n}^*$, then there exists a co-state vector $\mathbf{p} : [0,T] \to (\mathbb{R}^n)^*$, which satisfies the adjoint equation*

$$\dot{\mathbf{p}} = -\frac{\partial H}{\partial \mathbf{n}}, \tag{4}$$

with terminal condition $\mathbf{p}(T) = \frac{\partial M}{\partial \mathbf{n}}(T, \mathbf{n}^*)$, such that the Hamiltonian H is minimised a.e. on $[0,T]$ by u^* along the optimal trajectory with a constant minimum value c, i.e.

$$H(\mathbf{p}(t), \mathbf{n}^*(t), u^*(t)) = \min_{0 \le v(t) \le u_{\max}} H(\mathbf{p}(t), \mathbf{n}^*(t), v(t)) \equiv c.$$

The adjoint equation, together with the model dynamics, forms a two-point boundary value problem that is related to the optimal control through the minimising condition. It should be pointed out that the system may have multiple solutions. To determine the globally optimal control, we need to find all of them.

If the Hamiltonian is linear in control, its minimising property motivates the definition of the switching function

$$\Phi(t) = \frac{\partial H}{\partial u}. \tag{5}$$

If u^* is an optimal control, then $u^*(t) = \begin{cases} 0 & \text{if } \Phi(t) > 0, \\ u_{\max} & \text{if } \Phi(t) < 0. \end{cases}$ Thus, the condition $\Phi(t) = 0$ is the first-order necessary condition for the Hamiltonian to be minimal.

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63775-9_10

If $\Phi(\tau) = 0$ for some $\tau \in [0, T]$, but Φ is non-zero in some neighbourhood of τ, then the function Φ changes sign at time τ and the optimal control switches between 0 and $u_{\max}$: from 0 to $u_{\max}$ if $\dot{\Phi}(\tau) < 0$ and from $u_{\max}$ to 0 if $\dot{\Phi}(\tau) > 0$. This type of control is called bang-bang control. If Φ and all its derivatives vanish identically on some interval, then the control admits intermediate values between 0 and $u_{\max}$ and we say that the control is singular. In general, determining all points where $\Phi(t) = 0$, $t \in [0, T]$, is extremely complicated.

In many practical problems, optimal controls consist of finite concatenations of bang-bang and singular controls. Generally, the standard procedure for determining the structure of the optimal control is to analyse the switching function and its derivatives. A convenient tool in this analysis is the Lie bracket. We have

$$\begin{aligned}
\dot{\Phi} &= \mathbf{p}^T[\mathbf{f}, \mathbf{g}] - \frac{\partial L}{\partial \mathbf{n}}^T \mathbf{g}, \qquad (6)\\
\ddot{\Phi} &= \mathbf{p}^T\left[\mathbf{f}, [\mathbf{f}, \mathbf{g}]\right] - \frac{\partial L}{\partial \mathbf{n}}^T \left([\mathbf{f}, \mathbf{g}] + \frac{\partial \mathbf{g}}{\partial \mathbf{n}}\mathbf{f}\right) - \mathbf{f}^T \frac{\partial^2 L}{\partial \mathbf{n}^2}\mathbf{g} \\
&\quad + \left(\mathbf{p}^T\left[\mathbf{g}, [\mathbf{f}, \mathbf{g}]\right] - \frac{\partial L}{\partial \mathbf{n}}^T \frac{\partial \mathbf{g}}{\partial \mathbf{n}}\mathbf{g} - \mathbf{g}^T \frac{\partial^2 L}{\partial \mathbf{n}^2}\mathbf{g}\right) u.
\end{aligned}$$

The Lie bracket of two differentiable vector fields $\mathbf{f}$ and $\mathbf{g}$ can be defined as $[\mathbf{f}, \mathbf{g}](\mathbf{n}) = \frac{\partial \mathbf{g}}{\partial \mathbf{n}}\mathbf{f}(\mathbf{n}) - \frac{\partial \mathbf{f}}{\partial \mathbf{n}}\mathbf{g}(\mathbf{n})$. For the derivative of these formulas we refer to [9].

If a control is singular on some interval $I \subset [0, T]$, then $\Phi = \dot{\Phi} = \ddot{\Phi} = ... = 0$ identically on I. If we solve these equations for the co-state variable $\mathbf{p}$ and the term multiplying u is nonzero, the equation $\ddot{\Phi} = 0$ can be solved for u, determining the singular control

$$\begin{aligned}
&-\mathbf{p}^T\left[\mathbf{f}, [\mathbf{f}, \mathbf{g}]\right] + \frac{\partial L}{\partial \mathbf{n}}^T \left([\mathbf{f}, \mathbf{g}] + \frac{\partial \mathbf{g}}{\partial \mathbf{n}}\mathbf{f}\right) + \mathbf{f}^T \frac{\partial^2 L}{\partial \mathbf{n}^2}\mathbf{g} \\
&\qquad = \left(\mathbf{p}^T\left[\mathbf{g}, [\mathbf{f}, \mathbf{g}]\right] - \frac{\partial L}{\partial \mathbf{n}}^T \frac{\partial \mathbf{g}}{\partial \mathbf{n}}\mathbf{g} - \mathbf{g}^T \frac{\partial^2 L}{\partial \mathbf{n}^2}\mathbf{g}\right) u_{\text{sing}}, \qquad (7)
\end{aligned}$$

as a feedback formula that only depend on the state variables $\mathbf{n}$ of the system (1), and does not depend on the co-state variables $\mathbf{p}$ of the system (4).

Singular control, however, is not necessarily minimising, but it can also be maximising. In that case, instead of being the optimal control, it would represent the worst possible option (see [9]). The strengthened Legendre-Clebsch condition provides a high-order necessary condition for the optimality of singular controls, allowing the distinction between these two classes. If a minimising control u is singular of order 1 on an open interval $I \subset [0, T]$, then $\frac{\partial}{\partial u}\frac{d^2}{dt^2}\frac{\partial H}{\partial u} < 0$ for all $t \in I$.

Three main types of solutions to single-input optimal control problems can be distinguished [6]: (i) singular controls do not exist, (ii) singular controls exist and are locally minimising (the strengthened Legendre-Clebsch condition is satisfied, $\frac{\partial}{\partial u}\frac{d^2}{dt^2}\frac{\partial H}{\partial u} < 0$), and (iii) singular controls exist, but are locally maximising (the strengthened Legendre-Clebsch condition is not satisfied, $\frac{\partial}{\partial u}\frac{d^2}{dt^2}\frac{\partial H}{\partial u} > 0$).

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63775-9_10

In case (i), the optimal control is bang-bang with a small number of switchings that can be easily established. In case (ii), the optimal control consists of both bang-bang and singular intervals, where a bang-bang control with a larger number of switchings is not optimal. In case (iii), the optimal control consists of a potentially very large number of switchings. The singular arc is the limit of bang-bang trajectories where the number of switchings increases and tends to infinity. In this limit, the singular arc corresponds to scenario (ii), which is not optimal. Thus, in this case, more switchings do not improve the result and typically the solution is bang-bang with a small number of switchings. This scenario is by far the most difficult of the three scenarios.

3 Gradient methods for finding optimal control

The computation of optimal controls using numerical methods is often inefficient and may fail when optimisation is performed without prior knowledge of their structure. When singular controls are not present in the optimal structure, computing bang-bang controls is quite simple. Following the ideas presented in [12] (cf. [9]), optimal switching times for bang-bang control problems with a specified number of switches can be found using iterative methods. In these methods, the gradient of the objective functional with respect to the switching times is computed in each iteration. The number of switches is taken arbitrarily at the beginning of the algorithm. Results obtained for different numbers of switches are then compared to identify the optimal solution. Indeed, defining an upper limit for the number of switches is always possible since too many switches are not applicable. In brief, we can describe this approach as follows. Arbitrarily select switching times $0 = t_0 < t_1 < ... < t_k < t_{k+1} = T$. Let u_i denote the value of the bang-bang control on the interval $[t_i, t_{i+1}]$ for $i = 0, ..., k$. Note that the control's value on the first interval determines the sequence $u = (u_1, ..., u_k)$. Solve the state equations (1) for a given control u. Using the transversality condition, $\mathbf{p}(T) = [\omega_1, \omega_2]^T$, the adjoint co-state variables $\mathbf{p}$ can be computed by integrating the adjoint equation (4) backward in time (since the equations depend only on the state variables $\mathbf{n}$). Subsequently, using formulas (5) and (6), the switching function Φ and its derivative $\dot{\Phi}$ can be easily evaluated at the switching times t_i. By introducing small changes, iteratively update the switching times until they agree with the zeros of the computed switching function. Formally, when a control u is perturbed by δu, the first variation δJ can be found as

$$\delta J = J(u + \delta u) - J(u) \approx \int_0^T \frac{\partial H}{\partial u} \delta u dt.$$

Note that the derivative with respect to δu is actually given by the switching function $\Phi = \frac{\partial H}{\partial u}$. Depending on the value of the switching function at the switching time t_i for the control used, increase or decrease the lengths of the intervals $[t_{i-1}, t_i]$, where the increment is simply taken as

$$\delta t_i = (-1)^i \alpha \frac{\partial H}{\partial u}\Big|_{t=t_i} = (-1)^i \alpha \Phi(t_i).$$

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63775-9_10

Here, the parameter α represents an adaptive step-size parameter (learning rate) ensuring the convergence of the procedure.

However, the method described above cannot be directly applied if locally optimal singular controls exist. The presence of any subsequent singular interval in the control will perturb the control structure and switching times. Furthermore, switching to singular control will cause the trajectory to follow a different singular arc with different values of the singular control.

Therefore, we propose a modified version of the aforementioned algorithm (cf. [2]). Suppose that the structure of a control is fixed, meaning that the optimal control consists of a finite sequence with elements from 0, 1, and S, corresponding to no control, full control and singular control, respectively. However, the structure of the optimal control does not provide any information about the switching times. Treating the objective functional as a function of switching times, with a slight abuse of notation, an approximate gradient of the objective functional with respect to the switching times can be computed using finite differences

$$\frac{\partial J}{\partial t_i} \approx \frac{J(t_1, ..., t_i + \Delta, ..., t_k) - J(t_1, ..., t_i, ..., t_k)}{\Delta} \tag{8}$$

for some positive constant $\Delta \ll 1$. Algorithm 1 describes the procedure.

Algorithm 1 Gradient method.

1. Assume k, structure of optimal control, initial switching times $\mathbf{t} = (t_1, ..., t_k)$ and initial value of learning rate α.
2. Solve the state equations (1) for a given u (defined by the assumed structure of optimal control), taking into account the singular control form, defined by (7).
3. Compute the co-state variables $\mathbf{p}$ by integrating the adjoint equation (4) backward in time using the terminal condition.
4. Evaluate the Hamiltonian H given by (3) along the controlled trajectory.
5. Compute the gradient $\nabla J(\mathbf{t})$, according to formula (8).
6. Adapt the learning rate α to ensure the gradient descent.
7. Compute new switching times $\mathbf{t} - \alpha \nabla J(\mathbf{t})$.
8. Repeat steps 2-7 until a prescribed tolerance has been reached.

To ensure gradient descent, it is beneficial to reduce the learning rate as the training progresses. This can be done by implementing predefined learning rate schedules or employing adaptive learning rate algorithms. In our approach, we utilise the Adaptive Moment Estimation algorithm (Adam) and, if necessary, its modifications. However, other methods for adjusting the learning rate can also be considered (see e.g. [11]). It is important to note that, for the gradient to be well-defined, we typically require that $t_{i+1} - t_i > \Delta$ for all i. However, during iterations, the switching times change and potentially cross, i.e. $t_{t+1} - t_i$ may fall below Δ for some i. In this case the corresponding interval is removed from the control structure, resulting in trajectories with a reduced number of

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63775-9_10

switchings. In general, gradient methods are more suitable for finding local rather than global minima. Thus, to ensure optimality, the method should be run with different starting control structures and several random initial switching times for each control structure.

4 Numerical computations of optimal controls

4.1 Mathematical model

As an example, we consider a simple mathematical model describing the coexistence of control-sensitive and control-resistant subpopulations of the same species. The growth of the first one follows the exponential growth function, while the second population growths according to the logistic law, sharing resources with the first population. We assume that they differ in their reaction to control. In other words, we consider a scenario in which the growth of the control-sensitive subpopulation limits the growth of the control-resistant subpopulation. Intentionally, we choose as simple as possible mathematical model, because our aim is to investigate how the proposed objective functional affects the structure of the optimal control. The non-dimensional model considered in this study is as follows

$$\begin{aligned} \dot{n}_1 &= \lambda_1 n_1 - n_1 u, \\ \dot{n}_2 &= \lambda_2 n_2 (1 - n_2 - n_1), \end{aligned} \tag{9}$$

where n_1 and n_2 are the non-dimensional sizes of control-sensitive and control-resistant subpopulations, respectively; $u = u(t)$, $u : [0, T] \to [0, 1]$ is the non-dimensional control; T is the time horizon; λ_1 and λ_2 are growth rates.

4.2 Nonlinear objective functional

Mathematical models typically represent some underlying processes arising from medical, biomedical, physical, economical, or engineering problems. However, the form of objective functional is artificially imposed from the outside and usually there exist several options that can be used in a particular situation. The form of the objective depends on whether the system response and properties are satisfactory with respect to other criteria that were not included in the model dynamics. Let us now define the following objective functional

$$\begin{aligned} J(u(\cdot)) &= M\left(T, n_1(T), n_2(T)\right) + \int_0^T L\left(n_1(t), n_2(t), u(t)\right) dt \\ &= \omega_1 n_1(T) + \omega_2 n_2(T) \\ &\quad + \int_0^T \left(\eta_1 n_1(t) + \eta_2 n_2(t) + \xi G\left(\tfrac{n_2(t) - n_1(t)}{\epsilon}\right) + \theta u(t)\right) dt, \end{aligned} \tag{10}$$

where ω_1, ω_2, η_1, η_2, ξ, θ are non-negative parameters (weights), $\epsilon > 0$, and $G : \mathbb{R} \to (0, 1)$ is a twice continuously differentiable function. Terms $\omega_1 n_1(T)$,

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63775-9_10

$\omega_2 n_2(T)$ penalise the size of the entire population at the end of the assumed fixed control interval $[0, T]$, while $\eta_1 n_1(t)$, $\eta_2 n_2(t)$ penalise the size of the entire population during the control. The linear term $\theta u(t)$ represents the overall amount of control given and penalises side-effects (cumulative negative effects, toxicity or costs) related to their usage. It should be pointed out that the linear term containing the control u has a clear biological meaning, but makes the mathematical analysis difficult. The non-standard term $\xi G\left(\frac{n_2(t)-n_1(t)}{\epsilon}\right)$ is an activation function introduced to penalise time periods during which the population is control-resistant $(n_2 > n_1)$. We require the following properties

1. $G(x) \to 0$ as $x \to -\infty$; $G(x) \to 1$ as $x \to \infty$;
2. $G'(x) > 0$ for all x; $xG''(x) < 0$ for $x \neq 0$;
3. $G(0) = \frac{1}{2}$ and $G'(0) = \frac{1}{2}$.

Notice that $G(x) \approx 1$ whenever $x \gg 0$ and $G(x) \approx 0$ when $x \ll 0$. Thus, the term increases whenever $n_2 > n_1$ and stays roughly constant if $n_1 < n_2$. The function G can be thought of as a smoothed version of the Heaviside function, where the parameter ϵ controls the steepness of the slope, i.e. it determines how close the control-resistant population needs to become to the control-sensitive one for the penalty to be applied. The concept of this function was introduced by us in [2]. For numerical purposes, we can take $\frac{1}{2}\left(1 + \tanh \frac{x}{\epsilon}\right)$ as the function of G, as shown in Figure 1. However, other types of activation functions can also be applied.

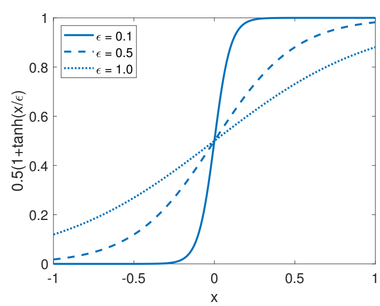

Fig. 1: Activation function $G(x) = \frac{1}{2}\left(1 + \tanh \frac{x}{\epsilon}\right)$.

4.3 Optimal control problem

As it was mentioned earlier, our goal is to control population dynamics (growth). The optimal control problem may be formulated as follows: for a fixed terminal time T find a measurable function $u : [0, T] \to [0, 1]$ minimising the objective functional (10) subject to dynamics (9).

Let us introduce the following notations

$$\mathbf{n} = \begin{bmatrix} n_1 \\ n_2 \end{bmatrix}, \quad \mathbf{f}(\mathbf{n}) = \begin{bmatrix} \lambda_1 n_1 \\ \lambda_2 n_2(1 - n_2 - n_1) \end{bmatrix}, \quad \mathbf{g}(\mathbf{n}) = \begin{bmatrix} -n_1 \\ 0 \end{bmatrix}, \quad \mathbf{p} = \begin{bmatrix} p_1 \\ p_2 \end{bmatrix}.$$

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63775-9_10

Following the theory described in Section 2, by elementary but tedious calculations, we get

$$H = p_1(\lambda_1 - u)n_1 + p_2\lambda_2 n_2(1 - n_2 + n_1) + \eta_1 n_1 + \eta_2 n_2 + \tfrac{\xi}{\epsilon}G'\left(\tfrac{n_2-n_1}{\epsilon}\right) + \theta u,$$

where p_1 and p_2 satisfy

$$\begin{aligned}\dot{p}_1 &= -p_1(\lambda_1 - u) + p_2\lambda_2 n_2 - \eta_1 + \tfrac{\xi}{\epsilon}G'\left(\tfrac{n_2-n_1}{\epsilon}\right),\\ \dot{p}_2 &= -p_2\lambda_2(1 - 2n_2 - n_1) - \eta_2 - \tfrac{\xi}{\epsilon}G'\left(\tfrac{n_2-n_1}{\epsilon}\right),\end{aligned} \tag{11}$$

with terminal conditions $p_1(T) = \omega_1$, $p_2(T) = \omega_2$. From the definition of the switching function Φ and relations for its derivatives, we can state that the singular control has the following form

$$u_{\text{sing}} = \lambda_1 - \frac{\lambda_2 n_2(1 - n_2 - n_1)}{n_1} + \frac{\epsilon^2\lambda_2(\eta_2 - \eta_1)n_2 + 2\epsilon\xi\lambda_2 n_2 G'\left(\frac{n_2-n_1}{\epsilon}\right)}{\xi n_1 G''\left(\frac{n_2-n_1}{\epsilon}\right)} \tag{12}$$

and the corresponding singular trajectory lies in the following singular arc

$$\begin{aligned}0 = F_{\text{arc}}(n_1, n_2; c) := \epsilon\left(\eta_1 + \lambda_1\theta - c + (\eta_2 - \eta_1)n_2\right)&\\ + \epsilon\xi G\left(\tfrac{n_2-n_1}{\epsilon}\right) - \xi(1 - n_1 - n_2)G'\left(\tfrac{n_2-n_1}{\epsilon}\right).&\end{aligned} \tag{13}$$

This is a consequence of the Pontryagin Maximum Principle and the condition of constancy of the Hamiltonian, $H \equiv c = const$. Note that the above formulas are obtained as feedback formulas that only depend on the state variables $\mathbf{n}$ of the system (9), but do not depend on the co-state variables $\mathbf{p}$.

As it was mentioned earlier, $\frac{\partial H}{\partial u} = \Phi$, thus to verify if the Legendre-Clebsch condition is fulfilled along the singular arc, we need to determine the coefficient next to the control u in the expression for $\ddot{\Phi}$, i.e. $\frac{\partial}{\partial u}\frac{d^2\Phi}{dt^2} = -\frac{n_1^2}{\epsilon^2}\xi G''\left(\frac{n_2-n_1}{\epsilon}\right)$. Using the assumptions about G, we see that the Legendre-Clebsch condition is satisfied only if $n_1 > n_2$.

Regarding the problem under consideration, we can have any of these three cases: (i) if the term G that penalises the resistance is omitted from the objective functional ($\xi = 0$), the singular control does not exist and the optimal control is bang-bang, (ii) if $n_1 > n_2$, then the Legendre-Clebsch condition is satisfied and the optimal control consists of bang-bang and singular intervals, and (iii) if $n_1 < n_2$, then the Legendre-Clebsch condition is not satisfied and the singular control is not optimal.

4.4 Numerical optimisation

In this section we solve the optimal control problem using numerical computations. The goal is to minimise the objective functional under the model dynamics with the terminal non-dimensional time T chosen to be 14. To solve the optimal control problem and compute the optimal solution, we use the gradient method described in the previous section. Possible optimal control structures for different numbers of switches k are listed in Table 1.

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63775-9_10

Table 1: Possible structures of optimal controls for different numbers of switches k. Characters 0, 1, S denote no control, full control and singular control.

k	possible structures											
0	0				1				S			
1	01		$0S$		10		$1S$		$S0$		$S1$	
2	010	$01S$	$0S0$	$0S1$	101	$10S$	$1S0$	$1S1$	$S01$	$S0S$	$S10$	$S1S$
3	0101	$01S0$	$0S01$	$0S10$	1010	$10S0$	$1S01$	$1S10$	$S010$	$S0S0$	$S101$	$S1S0$
	$010S$	$01S1$	$0S0S$	$0S1S$	$101S$	$10S1$	$1S0S$	$1S1S$	$S01S$	$S0S1$	$S10S$	$S1S1$

The gradient method was run for each of the control structures with several random initial switching times. An adaptive learning rate algorithm, specifically the Adam Optimisation Algorithm, was employed to ensure effective gradient descent. The initial value of α was empirically set in the range of $0.0001 - 0.001$, while the exponential decay rates for the first and second moments were set within the range of $0.9 - 0.9999$. However, alternative methods for adjusting the learning rate could also be considered. To solve equations and find switching times numerically, the standard MATLAB solver ode45 and the MATLAB event model were used with an error tolerance set to 10^{-9}. Each simulation used the following parameter values: $\lambda_1 = 0.2$, $\lambda_2 = 0.1$, $\omega_1 = 5$, $\omega_2 = 10$, $\eta_1 = 2$, $\eta_2 = 3$, $\theta = 0.5$, $\epsilon = 0.01$, $\xi = 1$.

For the initial condition $[0.44, 0.05]$, representing an initially control-sensitive population, it was found that the minimal value of the objective functional, 11.8659, is achieved for a control of type $1S1$ with two switching times: 2.0822 and 13.8674. By evaluating the Hamiltonian H along the obtained controlled trajectory and solving the adjoint equations (11) backwards in time using the terminal condition, we verified that the Hamiltonian is indeed constant ($H \equiv c = 1.0112$, constant up to the method order). Additionally, we confirmed that the switching function has appropriate signs. The optimal solution with corresponding control, trajectory and switching function are shown in Figure 2. The optimal control starts with a full control interval as it penalises the control-sensitive subpopulation. Then, the control switches to singular part, defined by (12). As depicted in Figures 2a and 2c, a singular control maintains the size of the control-sensitive subpopulation just above the size of the control-resistant one. It follows directly from the penalty-activation function G included in the objective functional. The singular interval in the middle is crucial for preserving the sensitive subpopulation. The control ends with a full dose interval to penalise the population size at the terminal time. The red dotted curve in Figure 2c represents the singular arc, given by (13), where the right part minimises, while the left part – maximises the objective. By computing the derivative of the switching function on co-state and state variables, we determine that in the region between the two parts of the singular arc, the control can only switch from 0 to 1, while in the other regions (outside), it may switch only from 1 to 0.

The initial condition $[0.24, 0.25]$ was chosen to represent an initially control-resistant population. In this case, the Legendre-Clebsch condition is not satisfied

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63775-9_10

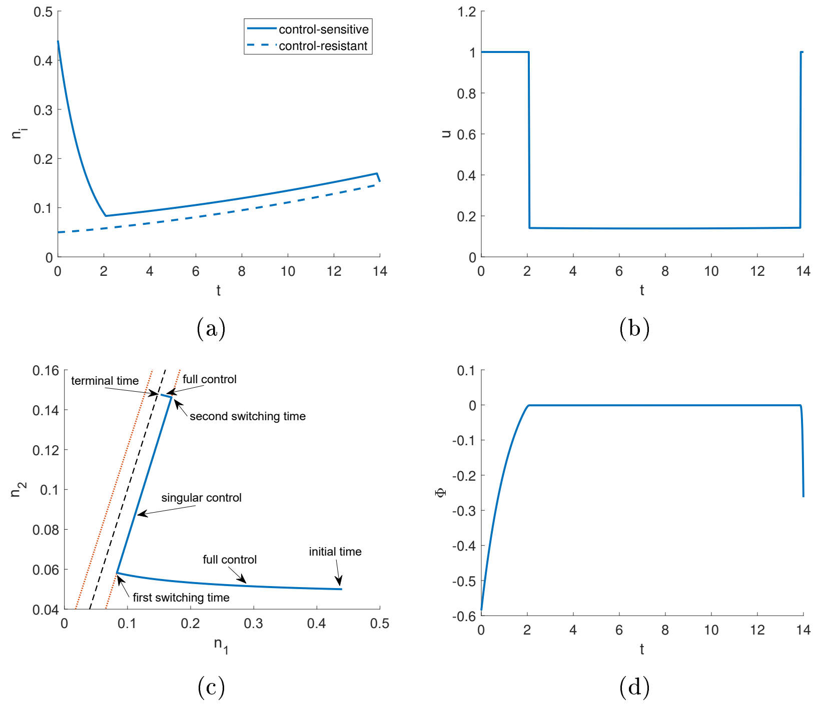

Fig. 2: Optimal solution (a), together with the corresponding optimal control (b), trajectory (c) and the switching function (d). Red dotted curve in (c) depicts the singular arc. Initial condition was chosen to be $n_1(0) = 0.44$, $n_2(0) = 0.05$.

(for $n_2 > n_1$), indicating that the optimal control starting from the singular interval cannot be optimal. This reduces the number of possible control structures. In this scenario, the optimal trajectory was numerically found to be $0S1$. The minimal value of the objective cost is 30.9052, while switching times are 1.0180 and 13.5811. The optimal solution, corresponding control, trajectory and switching function are shown in Figure 3 (the constant value of $H \equiv c = 2.2095$). The control begins with a no-control interval, and when the population becomes sufficiently control-sensitive, a singular control is used.

4.5 Comparative analysis

Now, we utilise three popular languages of optimisation modelling:

- Pyomo: A Python-based open-source optimisation modelling language with its differential algebraic equation extension Pyomo.DAE.
- AMPL: A Mathematical Programming Language, specifically designed for expressing and solving mathematical programming problems.

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63775-9_10

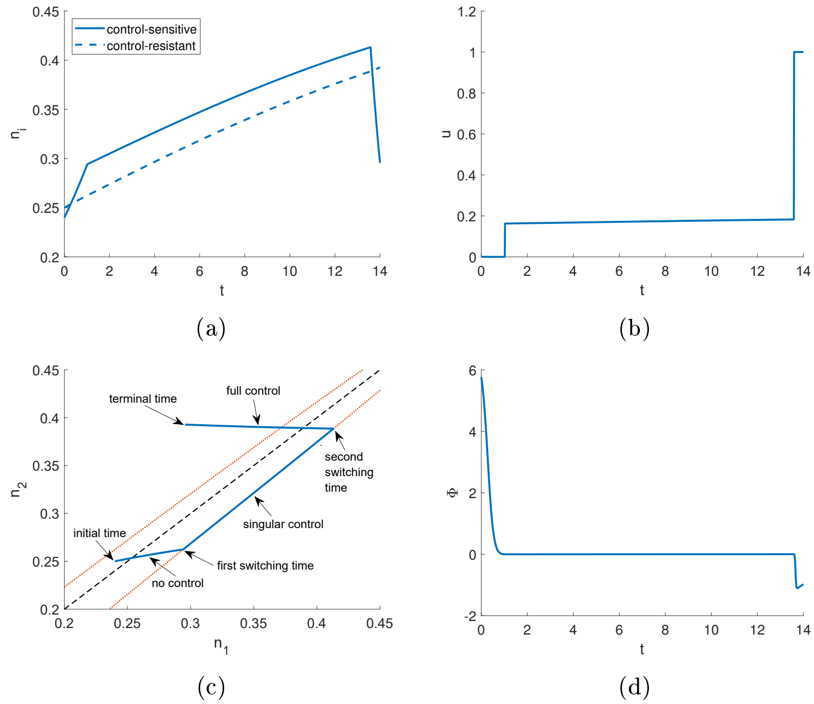

Fig. 3: Optimal solution (a), together with the corresponding optimal control (b), trajectory (c) and the switching function (d). Red dotted curve in (c) depicts the singular arc. Initial condition was chosen to be $n_1(0) = 0.24$, $n_2(0) = 0.25$.

– JuMP: An algebraic modelling language that is a collection of supporting packages for mathematical optimisation embedded in the Julia language.

These modelling languages serve as interfaces for solving the optimal control problem using the non-linear optimisation algorithm Ipopt, with a regularly spaced grid of 1000 points and the forward Euler method. Ipopt is an implementation of an interior point method and is capable of solving constrained non-linear programming problems (see [13] for details).

Figures 4a-4b depict the results obtained using the Pyomo modelling language with the Ipopt solver, for initial conditions $[0.44, 0.05]$ and $[0.24, 0.25]$, respectively. When penalising the side effects (control costs, $\theta = 0.5 > 0$), a control with a large number of "fast" switchings is obtained. It was verified that increasing the number of discretisation points (grid points) increases the number of switching points and has virtually no effect on the solution. The optimal singular trajectory is the limit of bang-bang trajectories with an increasing number of switchings tending to infinity. It is important to note that using the approximation method to determine a singular arc based on oscillations is generally not feasible. This phenomenon supports the concept of "chattering" (see [14]),

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63775-9_10

where the control oscillates many times between the bounds. Figures 4c-4d show the results obtained using the AMPL language, while Figures 4e-4f showcase the result obtained using the JuMP language. These results closely align with those obtained using the gradient method, but numerical artifacts on the singular interval can still be observed. This should be treated as artifacts caused by the appearance of a singular arc, where the method is unable to switch on to the optimal singular trajectory correctly and then switch to the bang-bang control.

In each optimisation result, the first and last switches follow those obtained from the gradient method, up to the grid-step size. Additionally, the mean optimal control values differ by less than 0.5% across methods, with the smallest value obtained using Pyomo and the largest using AMPL. The number of grid points typically has a negligible impact on the qualitative results, primarily contributing to a more accurate numerical calculation of the objective cost (through numerical integration) and switching times. Increasing the number of grid points may have an impact on the results, but a noticeable improvement in the result occurs with a very significant increase in the number of grid points and is sensitive to changes in model parameters, in particular weights in the objective functional. For example, when solving the problem with $\theta = 0$, each modelling method gives almost identical results, with minor differences attributable to the grid step size. In this scenario, the optimal trajectories closely approximate the singular arc defined by (13).

5 Conclusions

In this paper, we discussed a modified gradient method for solving the optimal control problems when singular intervals are present in the optimal structure. The appearance of singular controls complicates the problem of finding optimal solutions, because trajectories move along different singular arcs with different values of singular controls. However, the existence of singular controls leads to intermediate-value controls, which can significantly impact the optimisation results. The structure of optimal control depends primarily on the form of the objective functional, but also changes based on initial conditions and model dynamics. Even if the optimal control is initially bang-bang, it may change to bang-singular, when the dynamics change in response to the control. Singular controls are often more natural candidates for optimality than bang-bang controls, which arise only when singular controls do not exist or are inadmissible.

The numerical computation of optimal strategies poses a problem of high complexity. Direct methods can successfully run if all optimal control variables are bang-bang. Many different approaches can be applied after discretising the optimal control problem, even when the exact number of switching times is not known a priori. However, numerical procedures for solving singular control problems are usually problematic. Direct forward-backward sweep methods are not advisable. Indeed, some methods can detect the structure of the optimal solution and accurately compute the switching times without prior information about optimal control. As we showed, it is possible that the discretisation of

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63775-9_10

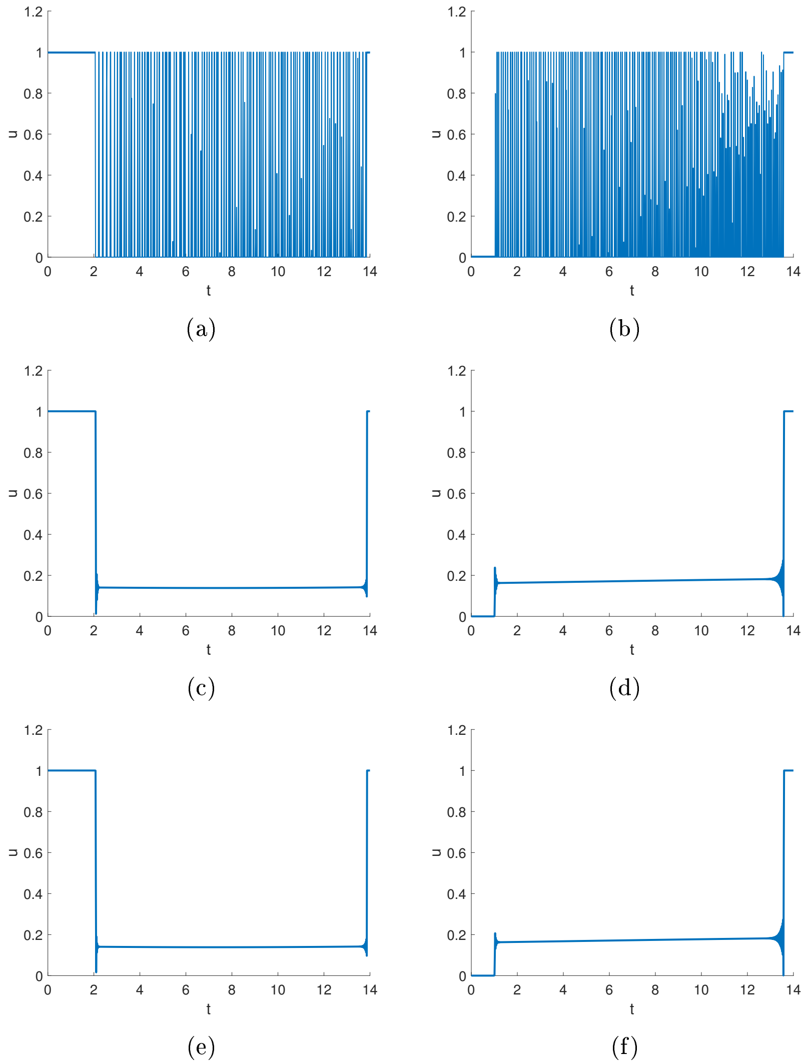

Fig. 4: Optimal controls of the model (9)-(10) obtained by Pyomo (a)-(b), AMPL (c)-(d) and JuMP (e)-(f). The initial condition in (a), (c) and (e) was set to be $[0.44, 0.05]$, while in (b), (d) and (f) was set to be $[0.24, 0.25]$.

an optimal control problem generates numerical artifacts and determining the exact optimal trajectory becomes impossible. In contrast to other numerical algorithms and modelling interfaces, the presented algorithm, after switching to the singular interval, moves along the particular singular arc that can be derived analytically. This approach provides the best possible approximation of optimal controls. A priori assumption that the optimal control structure is

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63775-9_10

fixed effectively eliminates non-optimal controls, which could only be numerical artifacts.

References

1. Atkins, S., Aghaee, M., Martcheva, M., Hager, W.: Computational and Mathematical Population Dynamics. World Scientific (2023). https://doi.org/doi.org/10.1142/13045
2. Bajger, P., Bodzioch, M., Foryś, U.: Singularity of controls in a simple model of acquired chemotherapy resistance. Discrete and Continuous Dynamical Systems Series B **24**(5), 2039–2052 (2019). https://doi.org/10.3934/dcdsb.2019083
3. Bajger, P., Bodzioch, M., Foryś, U.: Numerical optimisation of chemotherapy dosage under antiangiogenic treatment in the presence of drug resistance. Mathematical Methods in the Applied Sciences **43**, 10671–10689 (2020). https://doi.org/10.1002/mma.6958
4. Bodzioch, M., Bajger, P., Foryś, U.: Angiogenesis and chemotherapy resistance: optimizing chemotherapy scheduling using mathematical modeling. Journal of Cancer Research and Clinical Oncology **147**(8), 2281–2299 (2021). https://doi.org/10.1007/s00432-021-03657-9
5. Bodzioch, M., Bajger, P., Foryś, U.: Competition between populations: preventing domination of resistant population using optimal control. Applied Mathematical Modelling **114**, 697–693 (2023). https://doi.org/10.1016/j.apm.2022.10.016
6. Bonnard, B., Chyba, M.: Singular Trajectories and their Role in Control Theory, Mathematics and Applications, vol. 40. Springer (2003)
7. Falcone, M., Ferretti, R., Grüne, L., McEneaney, W.M. (eds.): Numerical Methods for Optimal Control Problems. Springer (2019). https://doi.org/10.1007/978-3-030-01959-4
8. Foroozandeh, Z., do Rosário de Pinho, M., Shamsi, M.: On numerical methods for singular optimal control problems: An application to an AUV problem. Discrete and Continuous Dynamical Systems Series B **24**(5), 2219–2235 (2019). https://doi.org/10.3934/dcdsb.2019092
9. Ledzewicz, U., Schättler, H.: Optimal Control for Mathematical Models of Cancer Therapies. Springer (2015). https://doi.org/10.1007/978-1-4939-2972-6
10. Pager, E.R., Rao, A.V.: Method for solving bang-bang and singular optimal control problems using adaptive Radau collocation. Computational Optimization and Applications **81**, 857–887 (2022). https://doi.org/10.1007/s10589-022-00350-6
11. Ruder, S.: An overview of gradient descent optimization algorithms. CoRR **abs/1609.04747** (2016), http://arxiv.org/abs/1609.04747
12. Śmieja, J., Świerniak, A., Duda, Z.: Gradient method for finding optimal scheduling in infinite dimensional models of chemotherapy. Journal of Theoretical Medicine **3**, 25–36 (2000). https://doi.org/10.1080/10273660008833062
13. Wächter, A., Biegler, L.: On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. Mathematical Programming **106**, 25–57 (2006). https://doi.org/10.1007/s10107-004-0559-y
14. Zelikin, M., Borisov, V.: Theory of chattering control: with applications to astronautics, robotics, economics, and engineering. Birkhäuser (1994). https://doi.org/10.1007/978-1-4612-2702-1

ICCS Camera Ready Version 2024
To cite this paper please use the final published version:
DOI: 10.1007/978-3-031-63775-9_10