

Review



Cite this article: Sharp JA, Burrage K, Simpson MJ. 2021 Implementation and acceleration of optimal control for systems biology. *J. R. Soc. Interface* **18**: 20210241. <https://doi.org/10.1098/rsif.2021.0241>

Received: 22 March 2021

Accepted: 3 August 2021

Subject Category:

Reviews

Subject Areas:

systems biology, computational biology, biomathematics

Keywords:

convergence acceleration, forward–backward sweep method, optimal control, Wegstein, Aitken–Steffensen, Anderson

Author for correspondence:

Jesse A. Sharp

e-mail: jesse.sharp@hdr.qut.edu.au

Electronic supplementary material is available online at <https://doi.org/10.6084/m9.figshare.c.5557230>.

Implementation and acceleration of optimal control for systems biology

Jesse A. Sharp^{1,2}, Kevin Burrage^{1,2,3} and Matthew J. Simpson¹

¹School of Mathematical Sciences, and ²ARC Centre of Excellence for Mathematical and Statistical Frontiers, Queensland University of Technology, Brisbane, Australia

³Department of Computer Science, University of Oxford, Oxford OX2 6GG, UK

JAS, 0000-0002-2865-4853; KB, 0000-0002-8111-1137; MJS, 0000-0001-6254-313X

Optimal control theory provides insight into complex resource allocation decisions. The forward–backward sweep method (FBSM) is an iterative technique commonly implemented to solve two-point boundary value problems arising from the application of Pontryagin’s maximum principle (PMP) in optimal control. The FBSM is popular in systems biology as it scales well with system size and is straightforward to implement. In this review, we discuss the PMP approach to optimal control and the implementation of the FBSM. By conceptualizing the FBSM as a fixed point iteration process, we leverage and adapt existing acceleration techniques to improve its rate of convergence. We show that convergence improvement is attainable without prohibitively costly tuning of the acceleration techniques. Furthermore, we demonstrate that these methods can induce convergence where the underlying FBSM fails to converge. All code used in this work to implement the FBSM and acceleration techniques is available on GitHub at <https://github.com/Jesse-Sharp/Sharp2021>.

1. Introduction

Across the life sciences, we encounter systems over which we wish to exert control. Whether we consider outbreak control in epidemiology [1,2], chemotherapy in oncology [3–5], muscle contraction and gait regulation in biomechanics [6–8], engineering cellular processes in synthetic biology [9,10], cell population growth in tissue engineering [11,12], or biodiversity and invasive species management in ecology [13–15], we face decisions around how a particular intervention should be applied to best achieve desired outcomes. Using mathematical models of such systems, optimal control theory provides insight into these resource allocation decisions.

Optimal control is a science of trade-offs; between competing objectives, or in weighing up the benefits of control measures against their costs. We illustrate some key concepts of optimal control in figure 1. Suppose that without intervention, a crop yield will double, from x_0 to $2x_0$, between now and harvest time. We might consider applying a control, such as fertilizer, to increase the growth rate of the crop; thereby increasing the yield at harvest to $3x_0$. Of course, applying fertilizer comes at a cost, and this must be considered against the increase in crop yield. As such, it is not immediately apparent how much fertilizer should be applied, and over what time period. This depends entirely on our characterization of optimality: the *pay-off*. Depending on the pay-off, the optimal control may be continuous; whereby the strength can be readily and continuously adjusted throughout time, or bang-bang (discontinuous); whereby the control is applied at either a lower or upper bound with finitely many discrete switches between the two. The pay-off determines the objective(s) of control; which in our stylized example may be to maximize profits after cost of fertilizing is considered, or achieve a specific yield, for example $3x_0$, using the minimum amount of fertilizer.

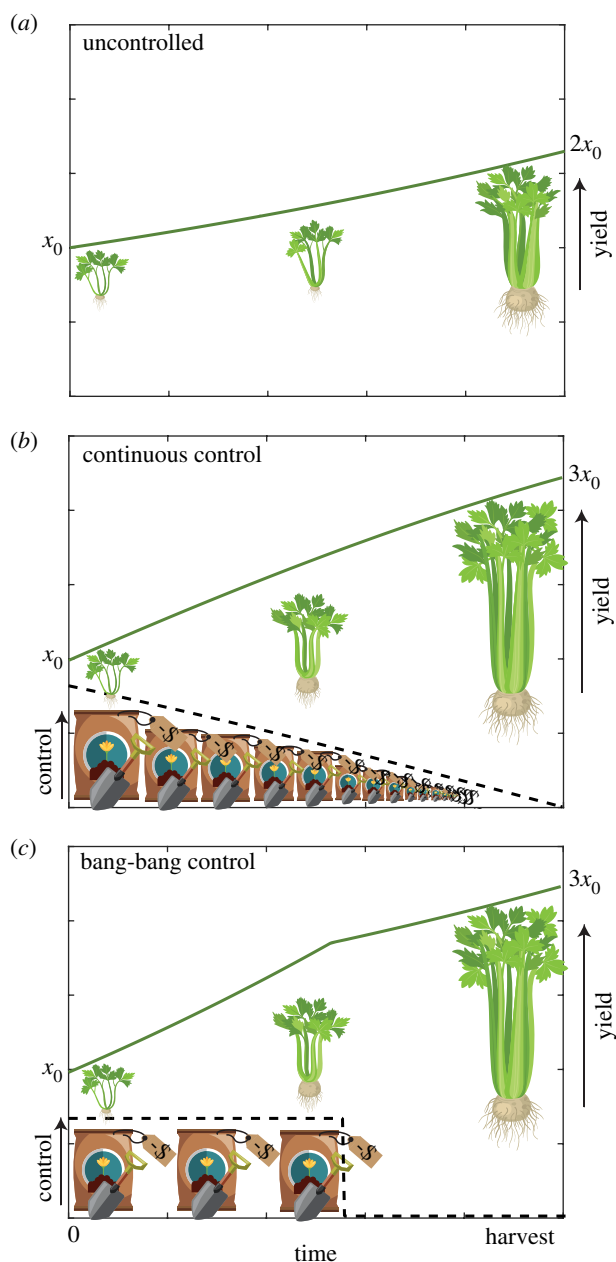


Figure 1. A pictorial example of optimal control for a growing crop. Suppose that initially, the crop yield is x_0 . We want to grow this crop to increase the yield, represented by the green line, come harvest time. Actions taken to increase the growth rate of the crop, such as applying fertilizer, are the controls, represented in black dash. Scenarios are presented for (a) no control, (b) continuous control and (c) bang-bang control. Optimal control theory helps us determine how best to apply these controls. Illustrations adapted from [ilyakalinin/iStock/Getty Images](#), [johavel/iStock/Getty Images](#).

Much of modern day optimal control theory stems from the seminal works of Pontryagin, through the Pontryagin maximum principle (PMP) [16], and Bellman, through the advent of dynamic programming and the Hamilton–Jacobi–Bellman equation [17], in the 1950s and 1960s. These foundations of optimal control are built upon centuries of development in the calculus of variations [18]. For brief but broad expositions of the theoretical roots of optimal control and developments following these seminal works, we direct the reader to articles such as [19,20].

Often we are unable to solve optimal control problems analytically, so we pursue computational approaches. Broadly, the numerical methods for optimal control can be classed as either

indirect or direct methods; for indirect methods optimality conditions are derived in the calculus of variations fashion via the PMP, leading to a two-point boundary value problem (TPBVP), while for direct methods the control problem is discretized and reformulated as a nonlinear programming problem [21]. For an early history of numerical methods in optimal control, including gradient and conjugate gradient methods, Newton–Raphson methods, quasi-linearization, feasible direction algorithms and feedback solutions we suggest [22]. Surveys [20,21] give an excellent overview of more recent developments in relevant numerical methods, including the forward–backward sweep method (FBSM), multiple-shooting methods, control parameterization, collocation and pseudospectral methods and complete discretization into finite-dimensional nonlinear programming problems.

Optimal control methodology and numerical solution techniques are continually being developed and improved. The growing popularity of artificial intelligence, machine learning and related disciplines has precipitated significant advances in computational techniques for handling large-scale systems with many variables, and related infinite-dimensional optimization problems. Nonlinear approximators, including neural networks, can be used to reduce infinite-dimensional optimization problems to finite-dimensional nonlinear programming problems. This approach is presented in [23], alongside other techniques that arise through unifying aspects of decision science, dynamic optimization, statistical and deterministic machine learning, nonlinear approximation theory and other fields. One example of control paired with machine learning arises in autonomous vehicles, where machine learning techniques can accelerate obtaining approximately optimal controls where computational power on-board is limited and controls satisfying strict safety constraints must be obtained rapidly [24]. Reinforcement learning, a technique from artificial intelligence resembling a model-free analogue of dynamic programming, has shown promising simulation results for the control of multi-species microbial communities in bioreactors [25].

Formulation and approximate solutions of fractional optimal control problems (FOCP)—optimal control of systems involving fractional derivatives—has also garnered wide interest recently within the control, numerical methods and applied mathematics communities; resulting in the development of new numerical approaches such as the non-standard two-step Lagrange interpolation method [26,27]; and amalgamations of new and existing techniques, such as pairing predictor-corrector methods for solving fractional differential equations with the FBSM for optimal control [28,29]. Applications involving FOCPs arise in areas of systems biology including epidemiology, where the incorporation of memory effects through fractional time derivatives may better describe disease transmission, by modelling the capacity for the population to learn from past outbreaks [26,30]; and in cancer therapy for determining optimal chemotherapeutic and immunotherapeutic treatment regimens [31,32].

The field of optimal control has historically focused on determining optimal interventions to apply to systems to meet specified objectives. More recently, however, optimal control techniques have been applied in a systems biology context to further our understanding of the underlying mechanisms or processes involved in a given system; for example via inverse optimal control, whereby exhibited behaviour observed in a system is used to elicit the underlying optimality

principles that may guide the system [33]. Optimality principles have been employed to investigate mechanisms in metabolism; for example, in [34], where optimal control techniques provide rationalization for experimentally and numerically observed sequential activation of metabolic pathways; in [35] where optimal control techniques predict enzyme activation times and metabolite concentrations; and in other work reviewed in [36], where further insights are gained regarding metabolic pathway activation and regulation. Optimal control has also provided insight into the emergence of persister cells in the presence of environmental volatility [37].

The FBSM is an iterative method for solving the TPBVPs that arise from the indirect PMP approach to optimal control. In systems biology, the FBSM for optimal control is very popular, owing particularly to its straightforward scalability to large systems, and to its moderate computational cost and mathematical complexity [38]. In this work, we review the implementation of the FBSM to solve optimal control problems, and investigate means of accelerating the convergence. To contextualize our discussion of the FBSM, we first consider the more familiar technique of successive over-relaxation (SOR). SOR is a generalization of the Gauss–Seidel method, and is widely applied in numerical linear algebra to accelerate convergence when solving linear systems iteratively [39]. Essentially, the process of SOR involves specifying an acceleration or relaxation parameter, $\beta \in (0, 2)$; a weighting factor that serves to reduce the spectral radius of the iterative matrix operator [40]. The error and rate of convergence of SOR is sensitive to this (problem dependent) choice of β , prompting investigation into theoretical convergence results and methods of determining β [40–42]. Despite challenges in identifying the optimal β , the SOR has historically been widely applied and studied in the literature due to the ease with which it can be implemented, and the rapid convergence it can deliver; even without identifying the optimal β [43,44].

This narrative closely parallels that of the FBSM in optimal control, where a weighting factor ω can be applied when updating the control between iterations to aid convergence. The optimal choice of ω is problem dependent, and significantly impacts the rate of convergence, or whether the FBSM converges at all. Nonetheless, the FBSM is frequently used in applied optimal control work as it is relatively straightforward to implement, and can still converge in absence of the optimal ω . Theoretical convergence results of the FBSM are available in the literature [45,46], although the focus is on the FBSM without weighted updating, with no consideration for choosing ω . Using regularization techniques, the FBSM is modified in [47] to improve convergence properties for large systems in a continuous setting, with a view to training deep neural networks in machine learning. These convergence results have recently been extended to the numerically discretized setting through symplectic Runge–Kutta discretization; taking advantage of the variational structure of optimal control problems [48]. The authors also demonstrate that the rate of convergence of the regularized FBSM with symplectic discretization can be improved with Anderson acceleration, an iterative acceleration technique. Although promising, this regularization introduces a regularization parameter, ρ . Similar to ω , the choice of ρ impacts convergence, and its choice is problem dependent. Understanding and implementing the regularization and symplectic techniques is not trivial, and

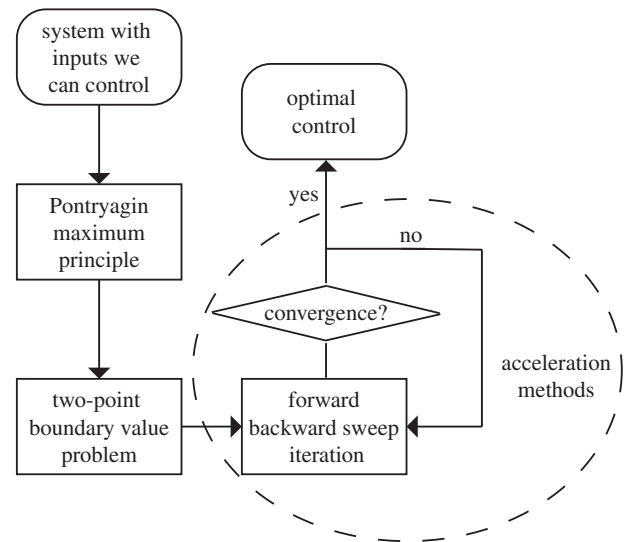


Figure 2. The process of optimal control via the Pontryagin maximum principle approach, with the incorporation of acceleration methods.

introduces conceptual complexity beyond what is necessary for many applied optimal control problems. As such, the standard FBSM remains an attractive choice for practitioners.

To this end, we aim to review acceleration techniques that can be paired with the standard FBSM. We implement such techniques alongside the FBSM with the goals of: (1) increasing the rate and frequency of convergence and (2) reducing the importance of, and challenges associated with, selecting ω . A graphical overview of the optimal control process we employ in this work, including the incorporation of acceleration methods, is presented in figure 2. We note that all code used to implement the algorithms presented in this review, the FBSM and the Wegstein, Aitken–Steffensen and Anderson acceleration methods, is available on GitHub (<https://github.com/Jesse-Sharp/Sharp2021>).

Throughout this work, we consider optimal control in the systems biology context. However, we note that optimal control is relevant to a wide variety of fields including chemical engineering [49], aeronautics and astronautics [21], management science and economics [50]. The FBSM, and by extension, the acceleration techniques we consider in this work, can be readily applied in any of these areas.

In §2, we review the PMP approach to optimal control, and the implementation of the FBSM. We provide a single-variable linear model, and a multi-variable nonlinear model in §3; and pose and solve example continuous, bang-bang (discontinuous), and fixed endpoint control problems. We review potential iterative acceleration methods in §4, and present the results of selected techniques in §5. We discuss the performance of these techniques in §6, and identify opportunities for application and further investigation.

2. Forward–backward sweep method

In an optimal control problem with one state variable, $x(t)$, one control, $u(t)$, over a fixed time interval, $t \in [t_0, t_N]$, such as the crop growth example presented in figure 1, we seek the optimal control $u^*(t)$ that minimizes or maximizes a specified pay-off function, J , subject to the dynamics of the state. In this section, we briefly review the PMP approach to such an optimal control problem, and the standard implementation of

the FBSM for solving the resulting two-point boundary value problem. The FBSM is readily extended to problems with multiple state variables, multiple controls, state constraints and free end-times [5,38,46,51]; however for this overview, we restrict ourselves to the single variable, single control, fixed end-time case for clarity.

The pay-off typically comprises a cost function $\mathcal{L}(t, x(t), u(t))$ integrated over the time interval, and/or a function, ϕ , of the state at final time: $\phi(x(t_N))$. As such, we seek to minimize or maximize J , subject to

$$J = \phi(x(t_N)) + \int_{t_0}^{t_N} \mathcal{L}(t, x(t), u(t)) dt \quad (2.1)$$

and

$$\frac{dx}{dt} = f(x(t), u(t), t), \quad x(t_0) = x_0. \quad (2.2)$$

Applying the PMP, we construct the Hamiltonian $H(t, x(t), u(t), \lambda(t)) = \mathcal{L}(t, x(t), u(t)) + \lambda(t)f(x(t), u(t), t)$, where $\lambda(t)$ is the co-state variable linking our state to our pay-off. The necessary conditions for optimal control are obtained from the Hamiltonian:

- (1) The optimal control, $u^*(t)$, is obtained by minimizing the Hamiltonian

$$\frac{\partial H}{\partial u} = 0. \quad (2.3)$$

- (2) The co-state is found by setting

$$\frac{d\lambda}{dt} = -\frac{\partial H}{\partial x},$$

- (3) satisfying the transversality condition

$$\lambda(t_N) = \lambda_N = \left. \frac{\partial \phi}{\partial x} \right|_{t=t_N}. \quad (2.4)$$

Following these steps yields a TPBVP to solve for $x(t), \lambda(t)$, subject to $x(t_0) = x_0$, and $\lambda(t_N) = \lambda_N$. To solve this numerically, we discretize t into $N + 1$ time points separated by a step-size $dt = (t_N - t_0)/N$; $\mathbf{t} = [t_0, t_0 + dt, \dots, t_0 + Ndt] = [t_0, t_1, \dots, t_N]$. Here, we consider a uniform discretization in time; although this is not strictly necessary, as discussed in §3. Using superscripts to denote the iteration number, provide an initial guess of the control at each t ; $\mathbf{u}^{(0)} = [u_0^{(0)}, u_1^{(0)}, \dots, u_N^{(0)}]$. From $\mathbf{u}^{(0)}$, solve equation (2.2) numerically from t_0 to t_N to obtain $\mathbf{x}^{(0)} = [x_0^{(0)}, x_1^{(0)}, \dots, x_N^{(0)}]$. Now, using $\mathbf{x}^{(0)}$, solve for $\boldsymbol{\lambda}^{(0)} = [\lambda_0^{(0)}, \lambda_1^{(0)}, \dots, \lambda_N^{(0)}]$ backwards in time from t_N to t_0 , starting from λ_N . With the optimality condition from equation (2.3), generate a temporary update for the control, $\hat{\mathbf{u}}^{(1)}$. The next iteration begins with an updated guess for the control, $\mathbf{u}^{(1)}$. These steps are repeated until a convergence condition is satisfied. The algorithm for the FBSM is summarized in §1 of the electronic supplementary material.

In some instances, directly updating the control, such that

$$\mathbf{u}^{(k)} = \hat{\mathbf{u}}^{(k)}, \quad k = 1, 2, \dots \quad (2.5)$$

is sufficient; however more commonly a weighted update is performed [5,38], such that in the $(k + 1)$ th iteration,

$$\mathbf{u}^{(k+1)} = \omega \mathbf{u}^{(k)} + (1 - \omega) \hat{\mathbf{u}}^{(k+1)}, \quad k = 1, 2, \dots, \omega \in [0, 1]. \quad (2.6)$$

This weighted updating is also referred to as applying a relaxation factor, similar to SOR as discussed in §1. An appropriate choice of ω in equation (2.6) can accelerate convergence relative to equation (2.5), or in some cases induce convergence where equation (2.5) leads to divergence. The weighting parameter, ω , can be held constant between iterations, although faster convergence may be achieved by updating ω . For example, by reducing ω as the system approaches convergence, a greater portion of the updated control is maintained relative to the control from the previous iteration [38], possibly accelerating convergence. A challenge commonly faced in implementing this control updating scheme is that the best choice for ω is problem dependent, and often is determined heuristically in practice. We address the extent to which the proposed acceleration algorithms address this issue in §4.

To facilitate the following discussion regarding acceleration, we note that the FBSM can be thought of as a generalized fixed point iteration [46], where each iteration comprises a forward and backward sweep and a control update. As such, for a control problem with one control, discretized into $N + 1$ time points, each iteration of the FBSM can be thought of as the application of a nonlinear operator, \mathcal{F} , of dimension $N + 1$, such that $\mathbf{u}^{(k+1)} = \mathcal{F}(\mathbf{u}^{(k)})$, or:

$$\begin{bmatrix} u_0^{(k+1)} \\ u_1^{(k+1)} \\ \vdots \\ u_N^{(k+1)} \end{bmatrix} = \begin{bmatrix} f_0(u_0^{(k)}, u_1^{(k)}, \dots, u_N^{(k)}) \\ f_1(u_0^{(k)}, u_1^{(k)}, \dots, u_N^{(k)}) \\ \vdots \\ f_N(u_0^{(k)}, u_1^{(k)}, \dots, u_N^{(k)}) \end{bmatrix},$$

where $\mathcal{F} = (f_0, f_1, \dots, f_N)^T$. However, in general, we are not able to write down an explicit expression for \mathcal{F} . Viewing the FBSM as a fixed point iteration process informs the choice of acceleration methods discussed in §4.

Importantly, we use the term *function evaluation* in this work to refer to the process of solving the system of ODEs for the state forward in time and the system of ODEs for the co-state backwards in time, once. This aligns with a single iteration of the standard FBSM. The function evaluation nomenclature becomes convenient when discussing the FBSM in the context of acceleration algorithms that typically focus on reducing the number of times expensive functions are evaluated. Producing numerical solutions to the ODE systems is by far the most computationally expensive component of the FBSM. This computational expense increases with the size and complexity of the systems; reducing the number of times these systems must be solved becomes more advantageous as the size and complexity of the systems increases. The function evaluation description also facilitates comparison between acceleration methods that require solving the ODE systems a different number of times per iteration. Throughout this work, we use \mathcal{N} to denote the total number of function evaluations a given method takes to achieve convergence.

2.1. Adapted forward–backward sweep method

The FBSM can be extended to handle problems where we aim not only to minimize or maximize a given quantity over time but also ensure that a specific state is reached at final time. This aligns with the crop growth example from figure 1 if the objective is to achieve a specific yield of $3x_0$ at harvest, rather than to maximize yield. In this case, we may have an integral term in the pay-off as described in equation (2.1);

however, the function of the final state, $\phi(x(t_N))$, is redundant in a control problem with a prescribed final state. Equation (2.2) is also modified to incorporate the additional constraint

$$J = \int_{t_0}^{t_N} \mathcal{L}(t, x(t), u(t)) dt, \quad \text{subject to}$$

$$\frac{dx}{dt} = f(x(t), u(t), t), \quad x(t_0) = x_0, \quad x(t_N) = x_N. \quad (2.7)$$

Here, x_N is the specified state that must be reached at final time. Since we have introduced an additional boundary value to the system, we no longer obtain the transversality condition from equation (2.4). Instead, we seek the final time condition on the co-state, λ_N , and associated optimal control that satisfies equation (2.7). We proceed by considering an adapted FBSM that takes as an input a guess for this final time condition, $\hat{\lambda}_N$, and solves the corresponding control problem. If we denote this application of the FBSM as the function $V(\hat{\lambda}_N)$, and the corresponding final value of the state, \hat{x}_N , then the adapted FBSM is an iterative process that solves for the root of $V(\hat{\lambda}_N)$; the value of $\hat{\lambda}_N$ for which $x_N - \hat{x}_N = 0$. This outer iterative process can be solved using standard techniques such as the bisection method or secant method; the former converging more reliably provided that the initial guesses for $\hat{\lambda}_N$ form an interval that brackets the root, the latter converging in fewer iterations [38]. Each of these outer iterations necessitates solving a boundary value problem to convergence, often involving numerous iterations of the FBSM. In this work, we apply the secant algorithm as presented in [38] without modification, for the adapted FBSM. The acceleration techniques described in §4 are applied only to the inner FBSM processes, reducing \mathcal{N} for each internal FBSM problem, leaving the outer secant iterations unchanged. Using $\mathcal{N}^{(k)}$ to denote the number of function evaluations in the k th internal FBSM problem, we can express the cumulative function evaluations required for convergence of the adapted FBSM as Σ , such that $\Sigma = \mathcal{N}^{(1)} + \mathcal{N}^{(2)} + \dots$.

The adapted FBSM can also be used to solve control problems with isoperimetric constraints; integral constraints of the form

$$\int_{t_0}^{t_N} h(t, x(t), u(t)) dt = K,$$

where K is a prescribed constant. For example, if $h(t, x(t), u(t)) = u(t)$, then K represents a specific and known amount of control that must be applied. The approach to solve problems with isoperimetric constraints, as outlined in [38], is to introduce an additional state variable, z , with

$$\frac{dz}{dt} = h(t, x(t), u(t)), \quad z(t_0) = 0, \quad z(t_N) = K.$$

This transforms the problem with an isoperimetric constraint into a problem with a fixed endpoint, that can be solved using the adapted FBSM as described.

3. Control problems

To investigate the robustness and effectiveness of the iterative acceleration techniques that we will discuss in §4, we consider two distinct systems, and for each system we study three example control problems. The first system is a single species linear differential equation subject to a control. We

later demonstrate that under certain conditions we are able to obtain exact solutions for control problems applied to this model. The second system is a three species model for acute myeloid leukaemia (AML) governed by a coupled non-linear system of differential equations, subject to a control. We construct the linear model to examine the behaviour of the acceleration techniques as applied to a simple idealized set of control problems. We include the AML model, variations upon which have been considered in the literature [5,51,52], to examine how the acceleration techniques perform when applied to problems more reflective of those considered in applied optimal control. For each model, we consider three distinct control problems: continuous control, bang-bang control and continuous control with fixed endpoint.

For all control results presented in this work, convergence is deemed to be achieved when the error, ϵ , measured as the Euclidean norm of the difference between subsequent controls, falls below a tolerance of 1×10^{-10} . Numerical solutions to ODEs are obtained using a fourth-order Runge–Kutta method [53] with constant time-stepping. A uniform time discretization is sufficient for all control problems considered in this work. However, the FBSM and acceleration methods readily generalize to a non-uniform discretization. If the desired discretization for the state equations differs from that of the co-state equations, it is necessary to perform interpolation within each iteration of the FBSM to obtain values at corresponding time points. This can be computationally expensive and introduce an additional source of error. Where the desired discretizations for the state and co-state differ, numerical schemes with internal interpolation such as continuous Runge–Kutta methods may be appropriate [54,55].

3.1. Single-variable linear model

The linear model is a single species model for the growth of $x(t)$, subject to control $u(t)$ that increases the growth rate. This model could represent our stylized crop growth example presented in §1. We suppress the explicit time dependence of the state and co-state variables and the control in the following equations for notational convenience. For numerical results, we solve the linear problems on the domain $0 \leq t \leq 1$, with time-step $dt = 3.91 \times 10^{-3}$, giving $N = 257$ time points.

$$\frac{dx(t)}{dt} = \gamma x(t) + u(t), \quad x(0) = x_0, \quad \gamma > 0, \quad 0 \leq t \leq 1. \quad (3.1)$$

In the absence of control, $u(t) \equiv 0$, this model admits the solution $x(t) = x_0 e^{\gamma t}$, describing exponential growth.

3.1.1. Continuous control

We seek to maximize a quadratic cost function J , subject to

$$J = \int_0^1 (ax^2 - bu^2) dt, \quad a > 0, \quad b > 0. \quad (3.2)$$

Following the standard PMP approach for solving optimal control problems, we form the Hamiltonian and derive the co-state equation, transversality condition and optimality condition. The Hamiltonian is given by

$$H = ax^2 - bu^2 + \lambda(\gamma x + u). \quad (3.3)$$

The co-state equation is

$$\frac{d\lambda}{dt} = -\frac{\partial H}{\partial x} = -2ax - \lambda\gamma, \quad (3.4)$$




Figure 3. Solution to the linear continuous control problem. The optimal control, $u^*(t)$, is shown in black dash and the corresponding state, $x(t)$, in blue. This solution is produced with model parameter $\gamma = 0.5$, time-step $dt = 3.91 \times 10^{-3}$, over the interval $0 \leq t \leq 1$. The contributions of the state and the control to the pay-off are equally weighted, with $a = b = 1$.

with transversality condition $\lambda(1) = 0$. In this case, the optimality condition is

$$\frac{\partial H}{\partial u} = \lambda - 2bu = 0,$$

such that the optimal control is given by

$$u^*(t) = \frac{\lambda(t)}{2b}. \quad (3.5)$$

For model parameter $\gamma = 0.5$ and pay-off weightings $a = b = 1$, with initial condition $x_0 = 1$, we are able to solve the control problem analytically using standard techniques for linear systems with complex eigenvalues [56]. The process is laborious so we present the approach and analytical solution in §2 of the electronic supplementary material. In the electronic supplementary material, we also plot the analytical results against the numerical results to demonstrate the excellent agreement. The numerical solution to the linear continuous control problem is presented in figure 3. Convergence via the FBSM requires $\mathcal{N} = 57$ iterations.

3.1.2. Bang-bang control

For the bang-bang control, we consider the same state equation as in equation (3.1), and incorporate bounds on the control.

$$\begin{aligned} \frac{dx(t)}{dt} &= \gamma x(t) + u(t), & x(0) &= x_0, \\ \gamma > 0, & 0 \leq t \leq 1, & 0 \leq u(t) &\leq 2. \end{aligned}$$

We seek to maximize a cost function J that is linear in u ,

$$J = \int_0^1 (ax^2 - bu) dt, \quad a > 0, \quad b > 0.$$

We form the Hamiltonian and derive the co-state equation and transversality condition

$$H = ax^2 - bu + \lambda(\gamma x + u).$$

The co-state equation is

$$\frac{d\lambda}{dt} = -\frac{\partial H}{\partial x} = -2ax - \lambda\gamma,$$

with transversality condition $\lambda(1) = 0$.




Figure 4. Solution to the linear bang-bang control problem. The optimal control, $u^*(t)$, is shown in black dash and the corresponding state, $x(t)$, in blue. This solution is produced with model parameter $\gamma = 0.5$, time-step $dt = 3.91 \times 10^{-3}$, over the interval $0 \leq t \leq 1$, with pay-off weightings of $a = 1$ for the state, and $b = 3$ for the control. The bang-bang control has prescribed bounds of $0 \leq u^*(t) \leq 2$.

In seeking the optimality condition, we find

$$\frac{\partial H}{\partial u} = \lambda - b. \quad (3.6)$$

As equation (3.6) does not depend on u , we define a switching function

$$\psi(t) = \lambda - b,$$

and produce an expression for the control, based on the bounds on u and the sign of the switching function:

$$u^*(t) = \begin{cases} 0, & \psi(t) < 0, \\ 2, & \psi(t) > 0. \end{cases} \quad (3.7)$$

If $\psi(t)$ is zero over any finite interval excluding isolated points, the optimal control is singular rather than bang-bang. Over such intervals, minimization of the Hamiltonian does not provide sufficient information to determine the optimal control, and further conditions must be considered [38,57]. We restrict our focus in this work to non-singular bang-bang optimal control problems. The numerical solution to the linear bang-bang control problem is presented in figure 4. Convergence to this solution via the FBSM required $\mathcal{N} = 8$ iterations.

3.1.3. Continuous control with fixed endpoint

For the fixed endpoint problem, we proceed with the same state equation; however we now impose a terminal condition on x .

$$\begin{aligned} \frac{dx(t)}{dt} &= \gamma x(t) + u(t), \\ x(0) &= x_0, \quad x(1) = 10, \quad \gamma > 0, \quad 0 \leq t \leq 1. \end{aligned}$$

We seek to maximize the same quadratic cost function J , as considered in equation (3.2). As such, we form the same Hamiltonian given in equation (3.3) and derive the same co-state, equation (3.4), and expression for the control, equation (3.5). Note however that we do not prescribe a final time condition on the co-state equation via the transversality condition; as the system already has two boundary conditions, doing so would cause it to be overdetermined.

Instead, we make two guesses for $\lambda(1)$; for example, $\lambda^{(0)}(1) = -10$ and $\lambda^{(1)}(1) = 10$. We proceed by applying the adapted FBSM outlined in §2, using these guesses to initialize the secant method. Numerical results for the linear fixed endpoint control problem are presented in figure 5. Convergence of the adapted FBSM is achieved after $\Sigma = 177$ iterations.

3.2. Multiple-variable nonlinear model

The AML model is a nonlinear coupled multi-species model describing the interactions between progenitor blood cells, $A(t)$, and leukaemic stem cells, $L(t)$, that occupy the same

niche in the bone marrow, thereby competing for space and resources. Haematopoietic stem cells, $S(t)$, act as upstream production of $A(t)$. These dynamics have been explored in the literature both experimentally [58,59], and through mathematical modelling [52,60]. We subject the model to a chemotherapy-like control, $u(t)$, that acts as an additional death term for $L(t)$. The state can be expressed as $x(t) = [S(t), A(t), L(t)]^T$. As there are now three state equations, we require three co-state equations: $\lambda(t) = [\lambda_1(t), \lambda_2(t), \lambda_3(t)]^T$. We suppress the explicit time dependence of the state and co-state variables and the control in the following equations for notational convenience:

$$\left. \begin{aligned} \frac{dS}{dt} &= \underbrace{\rho_S S(1-S)}_{\text{logistic growth}} - \underbrace{\delta_S S}_{\text{differentiation}}, \\ \frac{dA}{dt} &= \underbrace{\delta_S S}_{\text{upstream production}} + \underbrace{\rho_A A(1-A-L)}_{\text{logistic growth with competition}} - \underbrace{\delta_A A}_{\text{differentiation}}, \\ \text{and } \frac{dL}{dt} &= \underbrace{\rho_L L(1-A-L)}_{\text{logistic growth with competition}} - \underbrace{\delta_L L}_{\text{differentiation}} - \underbrace{\frac{\alpha L}{\gamma+L}}_{\text{immune response}} - \underbrace{uL}_{\text{chemotherapy control}}. \end{aligned} \right\} \quad (3.8)$$

For each control problem associated with the AML model, we use initial conditions that yield a coexisting steady state in the absence of control (all three species non-zero): $S(0) = 1 - \delta_S/\rho_S$, $A(0) = 0.3255$ and $L(0) = 0.3715$. We solve the AML problems numerically on the domain $0 \leq t \leq 10$, with time-step $dt = 4.88 \times 10^{-4}$, giving $N = 20481$ time points. Model parameters are specified in table 1.

3.2.1. Continuous control

For the AML continuous control problem, we seek to minimize a quadratic cost function J that accounts for both the cost of applying the control and the cost of the leukaemic burden, subject to

$$J = \int_0^{10} (a_1 u^2 + a_2 L^2) dt, \quad a_1 > 0, a_2 > 0. \quad (3.9)$$

We form the Hamiltonian and derive the co-state equation, transversality condition and optimality condition. The Hamiltonian is given by

$$\begin{aligned} H &= a_1 u^2 + a_2 L^2 + (\rho_S S(1-S) - \delta_S S)\lambda_1 \\ &\quad + (\delta_S S + \rho_A A(1-A-L) - \delta_A A)\lambda_2 \\ &\quad + (\rho_L L(1-A-L) - \delta_L L - \frac{\alpha L}{\gamma+L} - uL)\lambda_3. \end{aligned} \quad (3.10)$$

The co-state equations are

$$\left. \begin{aligned} \frac{d\lambda_1}{dt} &= -\frac{\partial H}{\partial S} = -\rho_S \lambda_1 + 2\rho_S \lambda_1 S + \delta_S \lambda_1 - \delta_S \lambda_2, \\ \frac{d\lambda_2}{dt} &= -\frac{\partial H}{\partial A} = -\rho_A \lambda_2 + 2\rho_A \lambda_2 A + \rho_A \lambda_2 L + \delta_A \lambda_2 + \rho_L \lambda_3 L \\ \text{and } \frac{d\lambda_3}{dt} &= -\frac{\partial H}{\partial L} = -2a_2 L + \rho_A \lambda_2 A - \rho_L \lambda_3 + \rho_L \lambda_3 A + 2\rho_L \lambda_3 L \\ &\quad + \delta_L \lambda_3 + \frac{\alpha \gamma \lambda_3}{(\gamma+L)^2} + \lambda_3 u, \end{aligned} \right\} \quad (3.11)$$

with transversality conditions $\lambda_1(10) = \lambda_2(10) = \lambda_3(10) = 0$, obtained in the usual way. In this case, the optimality condition is

$$\frac{\partial H}{\partial u} = 2a_1 u - \lambda_3 L = 0, \quad (3.12)$$

such that the optimal control is given by

$$u^*(t) = \frac{\lambda_3 L}{2a_1}. \quad (3.13)$$

Numerical solutions for the AML continuous control problem are presented in figure 6. These solutions are obtained via the FBSM, requiring $\mathcal{N} = 38$ iterations with $\omega = 0.55$. This choice of ω minimizes \mathcal{N} for the AML continuous control problem solved with the FBSM without acceleration techniques. We discuss the choice of ω further in §5.

3.2.2. Bang-bang control

For the bang-bang AML problem, we consider the same states as in equation (3.8), and incorporate bounds, $0 \leq u \leq 0.3$, on the control. We seek to minimize a cost function J that is linear in the control and the state variable L :

$$J = \int_0^{10} (a_1 u + a_2 L) dt, \quad a_1 > 0, a_2 > 0. \quad (3.14)$$

We form the Hamiltonian and derive the co-state equations, transversality conditions and optimality condition. The Hamiltonian is given by

$$\begin{aligned} H &= a_1 u + a_2 L + (\rho_S S(1-S) - \delta_S S)\lambda_1 \\ &\quad + (\delta_S S + \rho_A A(1-A-L) - \delta_A A)\lambda_2 \\ &\quad + (\rho_L L(1-A-L) - \delta_L L - \frac{\alpha L}{\gamma+L} - uL)\lambda_3. \end{aligned} \quad (3.15)$$

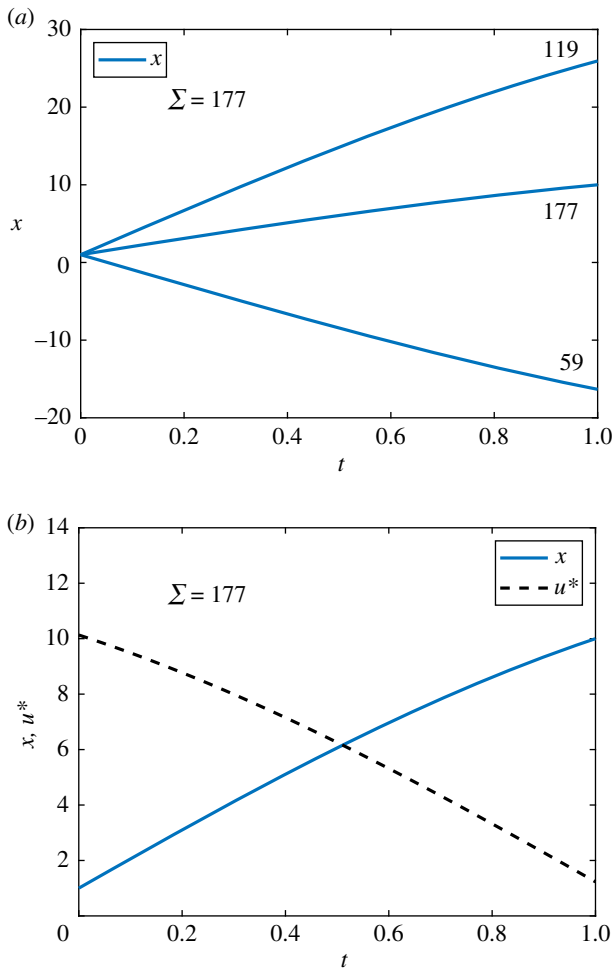


Figure 5. Results are presented for the linear problem with specified terminal state value, $x(t_N) = 10$, solved using the adapted FBSM. Underlying FBSM problems are solved with time-step $dt = 3.91 \times 10^{-3}$, over the interval $0 \leq t \leq 1$, with pay-off weightings of $a = b = 1$. In (a), the $x(t)$ iterates of the adapted FBSM are presented. We annotate the cumulative function evaluations after the first ($\mathcal{N}^{(1)} = 59$) and second ($\mathcal{N}^{(1)} + \mathcal{N}^{(2)} = 119$) iterations of the adapted FBSM, based on initial guesses for $\lambda(t_N)$ of $\lambda(t_N) = -10$ and $\lambda(t_N) = 10$. The total cumulative function evaluations required for convergence of the adapted FBSM, $\Sigma = \mathcal{N}^{(1)} + \mathcal{N}^{(2)} + \mathcal{N}^{(3)} = 177$, is indicated. The converged result for $x(t)$, satisfying $|x(t_N) - 10| \leq 1 \times 10^{-10}$ is presented in (b); this figure also includes the optimal control, $u^*(t)$.

The co-state equations are

$$\left. \begin{aligned} \frac{d\lambda_1}{dt} &= -\rho_S \lambda_1 + 2\rho_S \lambda_1 S + \delta_S \lambda_1 - \delta_S \lambda_2, \\ \frac{d\lambda_2}{dt} &= -\rho_A \lambda_2 + 2\rho_A \lambda_2 A + \rho_A \lambda_2 L + \delta_A \lambda_2 + \rho_L \lambda_3 L \\ \text{and } \frac{d\lambda_3}{dt} &= -a_2 + \rho_A \lambda_2 A - \rho_L \lambda_3 + \rho_L \lambda_3 A + 2\rho_L \lambda_3 L \\ &\quad + \delta_L \lambda_3 + \frac{\alpha \gamma \lambda_3}{(\gamma + L)^2} + \lambda_3 u, \end{aligned} \right\}$$

with transversality conditions $\lambda_1(10) = \lambda_2(10) = \lambda_3(10) = 0$. In this case, the switching function is

$$\psi(t) = \frac{\partial H}{\partial u} = a_1 - \lambda_3 L, \quad (3.16)$$

such that the optimal control is given by

$$u^*(t) = \begin{cases} 0, & \psi(t) > 0, \\ 0.3, & \psi(t) < 0. \end{cases} \quad (3.17)$$




Figure 6. Solution to the AML continuous control problem. The optimal control, $u^*(t)$, is shown in black dash and the corresponding state equations for $S(t)$, $A(t)$ and $L(t)$ are shown in blue, red and yellow, respectively. This solution is produced with model parameters given in table 1, time-step $dt = 4.88 \times 10^{-4}$, over the interval $0 \leq t \leq 10$, with pay-off weightings of $a_1 = 1$ for the control, and $a_2 = 2$ for state variable $L(t)$.

Table 1. AML model parameters. Parameters correspond to those presented with the original model [52], with immune response parameters introduced in subsequent work [5].

description	variable	value	dimension
proliferation of S	ρ_S	0.5	$[T^{-1}]$
proliferation of A	ρ_A	0.43	$[T^{-1}]$
proliferation of L	ρ_L	0.27	$[T^{-1}]$
differentiation of S into A	δ_S	0.14	$[T^{-1}]$
differentiation of A	δ_A	0.44	$[T^{-1}]$
differentiation of L	δ_L	0.05	$[T^{-1}]$
characteristic rate of the immune response	α	0.015	$[T^{-1}]$
half-saturation constant of the immune response	γ	0.1	$[-]$

Note that the correspondence between the sign of $\psi(t)$ and the chosen bound is reversed in equation (3.17) relative to equation (3.7) as we are now performing minimization rather than maximization. Numerical solutions for the AML bang-bang control problem are presented in figure 7. These solutions are obtained via the FBSM, requiring $\mathcal{N} = 34$ iterations with $\omega = 0.4$. This choice of ω minimizes \mathcal{N} for the AML bang-bang control problem solved with the FBSM without acceleration techniques. We discuss the choice of ω further in §5.

3.2.3. Continuous control with fixed endpoint

For the fixed endpoint problem, we proceed with the same state equations as for the AML continuous control problem given in equation (3.8); however we now impose a terminal condition on the leukaemic population: $L(10) = 0.05$. We seek to minimize the same quadratic cost function J , as considered in equation (3.9). We form the same Hamiltonian given in equation (3.10) and derive the same co-state, equation (3.11), and expression for the control, equation (3.13).




Figure 7. Solution to the AML bang-bang control problem. The optimal control, $u^*(t)$, is shown in black dash and the corresponding state equations for $S(t)$, $A(t)$ and $L(t)$ are shown in blue, red and yellow, respectively. This solution is produced with model parameters given in table 1, time-step $dt = 4.88 \times 10^{-4}$, over the interval $0 \leq t \leq 10$, with pay-off weightings of $a_1 = 1$ for control, and $a_2 = 2$ for the state variable $L(t)$.

We obtain final time conditions, $\lambda_1(10) = \lambda_2(10) = 0$, via the transversality conditions as usual; however we do not prescribe $\lambda_3(10)$. Instead, we make two guesses for $\lambda_3(10)$; for instance, $\lambda_3^{(0)}(10) = 0$ and $\lambda_3^{(1)}(10) = 10$. We then proceed by applying the adapted FBSM outlined in §2, using these guesses to initialize the secant method. Numerical results for the AML fixed endpoint control problem are presented in figure 8. These results are produced using the adapted FBSM with $\omega = 0.55$ in $\Sigma = 434$ iterations. This choice of ω minimizes Σ for the AML fixed endpoint control problem solved with the FBSM without acceleration techniques. We discuss this further in §5.

4. Iterative accelerators

In this section, we outline several techniques for acceleration of iterative schemes. Where appropriate, we first present the univariate/scalar version of the method for familiarity, then provide the multivariate/vector analogue of the method for use with accelerating the FBSM. We attempt to use notation that aligns most closely with commonly used notation in the literature, while maintaining internal consistency in this work. In the scalar case, we consider the iterative process $x^{(k+1)} = f(x^{(k)})$, where $x^{(k)}$ is the k th iterate and f is the iterating function. In the vector case, we consider $X^{(k+1)} = F(X^{(k)})$, where $X^{(k)} = [x_0^{(k)}, x_1^{(k)}, \dots, x_N^{(k)}]^T$ is the k th iterate, consisting of $N + 1$ values, and $F = [f_0, f_1, \dots, f_N]^T$ is the $N + 1$ dimensional operator of the iterative process. For clarity, we stress that in the context of the acceleration algorithms applied to the FBSM, $X^{(k)}$ is the discretized control in the k th iteration.

The acceleration methods considered in this work apply either to problems stated as fixed point iterations (as above), or as root-finding problems. For acceleration via root-finding algorithms, we can consider the complementary problems in the scalar and vector setting, respectively: $g(x) := x - f(x) = 0$ and $G(X) := X - F(X) = \mathbf{0}$, where $\mathbf{0}$ is the zero column vector of length $N + 1$.

We note that many of the methods presented here can be written in several different forms. While some forms better facilitate analysis of aspects such as convergence speed and

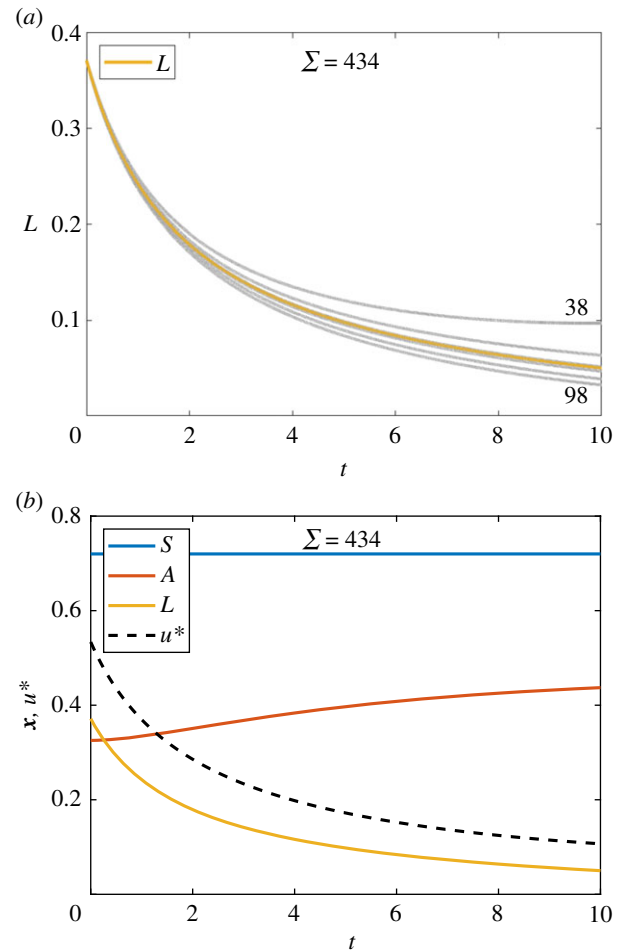


Figure 8. Results are presented for the AML problem with specified terminal state, $L(t_N) = 0.05$, solved using the adapted FBSM. Each underlying FBSM problem is solved with model parameters given in table 1, time-step $dt = 4.88 \times 10^{-4}$, over the interval $0 \leq t \leq 10$, with pay-off weightings of $a_1 = 1$ for the control, and $a_2 = 2$ for the state variable $L(t)$. In (a), the $L(t)$ iterates of the adapted FBSM are presented in grey; the converged solution satisfying $L(t_N) = 0.05$ is plotted in yellow. We annotate \mathcal{N} for the first ($\mathcal{N}^{(1)} = 38$) and second ($\mathcal{N}^{(1)} + \mathcal{N}^{(2)} = 98$) iterations of the adapted FBSM, based on initial guesses for $\lambda_3(t_N)$ of $\lambda_3(t_N) = 0$ and $\lambda_3(t_N) = 10$. Due to the close proximity, subsequent iterations are not annotated. The cumulative function evaluations required for convergence of the adapted FBSM ($\Sigma = 434$) is indicated. The converged result for $L(t)$, satisfying $|L(t_N) - 0.05| \leq 1 \times 10^{-10}$, is presented in (b); this figure also includes the optimal control, $u^*(t)$, and trajectories for $S(t)$ and $A(t)$.

numerical stability, others emphasize ease of understanding and implementation. In this work, we prioritize usability and present methods and algorithms in forms reflective of their implementation where possible. For the purpose of this work, we feel it is sufficient to present the methods and discuss their implementation without delving into their derivation or rigorous theoretical convergence results. For readers interested in these aspects, we suggest these articles [61,62], and numerical analysis texts [63,64].

4.1. Newton and quasi-Newton methods

Newton's method is one of the most prevalent root-finding algorithms, due to its relatively straightforward implementation and potential for quadratic convergence [64]. For a univariate

function, Newton's method is given by

$$x^{(k+1)} = x^{(k)} - \frac{f(x^{(k)})}{f'(x^{(k)})}. \quad (4.1)$$

We arrive at the scalar secant method by replacing the derivative term, $f'(x^{(k)})$, in equation (4.1) with a finite difference approximation

$$x^{(k+1)} = x^{(k)} - f(x^{(k)}) \frac{x^{(k)} - x^{(k-1)}}{f(x^{(k)}) - f(x^{(k-1)})}.$$

Newton's method for multivariate systems is

$$X^{(k+1)} = X^{(k)} + \Delta X^{(k)},$$

where $\Delta X^{(k)}$ is obtained by solving

$$J_k \Delta X^{(k)} = -F(X^{(k)}).$$

Here, J_k is the Jacobian matrix of F evaluated at $X^{(k)}$ [64]. Setting aside the interpretation of the Jacobian in the context of the FBSM, numerically approximating an $N \times N$ Jacobian matrix using finite differences requires $\mathcal{O}(N^2)$ FBSM iterations at each Newton step. A range of quasi-Newton methods have been developed to minimize the computational expense associated with computing the Jacobian at each Newton step. It is not immediately apparent how the secant method should be extended to multivariate systems, but one such interpretation is the quasi-Newton Broyden's method. Broyden's method reduces the number of function evaluations required at each Newton step by forming the full Jacobian only initially, then updating the Jacobian matrix via a rank-one update based on the secant method [63,65]. We later discuss the Wegstein method [66], which is another interpretation of the secant method in multivariate settings.

In the context of accelerating the FBSM, techniques that require forming or approximating a full Jacobian, even once, are not appropriate. We have an $N+1$ dimensional system, where $N+1$ is the number of time points in the discretization of the ODEs, so we expect N to be large, relative to the number of iterations required for the FBSM to converge without acceleration techniques, via equation (2.6). As such, we restrict our focus to Jacobian-free methods in the remainder of this section; in particular, we discuss and implement the Wegstein and Aitken–Steffensen methods and Anderson acceleration. We provide a broad overview alongside the key equations here, and provide complete algorithms alongside notes for implementation in §4 of the electronic supplementary material.

4.2. Wegstein method

Wegstein's method can be thought of as an element-wise extension of the secant method to multivariate systems [67]. Although Wegstein's method appears less popular than other methods considered in this work, it has found practical utility, particularly in chemical and process engineering software [68,69]. We include it here due to the striking similarity it bears to the control update with relaxation presented in equation (2.6). It is also one of the more straightforward techniques, both in conception and implementation:

$$\left. \begin{aligned} \hat{x}^{(k+1)} &= f(x^{(k)}), \\ x^{(k+1)} &= q^{(k)} x^{(k)} + (1 - q^{(k)}) \hat{x}^{(k+1)}, \end{aligned} \right\} \quad (4.2)$$

$$\text{where } q^{(k)} = \frac{a^{(k)}}{a^{(k)} - 1}, \text{ and } a^{(k)} = \frac{f(x^{(k)}) - f(x^{(k-1)})}{x^{(k)} - x^{(k-1)}}. \quad (4.3)$$

In implementation, from an initial value x_0 , it is necessary to perform two function evaluations, i.e. $x_1 = f(x_0)$, and $f(x_1)$, before it is possible to compute equation (4.3) for the Wegstein method [66]. In subsequent iterations only one new function evaluation is required.

The extension of Wegstein's method to multivariate systems follows exactly the process outlined in equations (4.2) and (4.3), as it is extended element-wise. While convergence is guaranteed when using Wegstein's method for a single non-linear equation, the uncoupling implied by the element-wise extension can lead to divergence [70].

In equations (4.2) and (4.3), $q^{(k)}$ denotes q in the k th iteration; however, we note that it is not necessarily most effective to update q every iteration. As such, in this work, we explore various updating regimes. There is also the option of applying bounds on q . Bounds of $-5 < q_i < 0$, $\forall i$, where i denotes the i th element of the system, are frequently applied when implementing Wegstein's method [71,72]. This bounding appears to work reasonably well for the small non-linear test systems we consider in §5 of the electronic supplementary material, although we were not able to identify a theoretical result supporting this specific choice. For the control problems we consider, this bounding is not effective, so we apply different bounds, discussed further in §5. The univariate Wegstein method can be thought of as a modification of the Aitken method, which at the time the Wegstein method was developed, was only understood for the univariate case [73].

4.3. Aitken–Steffensen method

Aitken's Δ^2 method, also referred to as Aitken's delta-squared process or Aitken extrapolation, was originally posed by Aitken in 1927 as a means of extending Bernoulli's method of approximating the largest root of an algebraic equation. This extension facilitates numerically approximating not only the largest root, but all roots of the equation [74]. Aitken's method generates a new sequence, \hat{x} , in parallel to the fixed point iteration.

$$\begin{aligned} \hat{x}^{(k)} &= x^{(k)} - \frac{(x^{(k+1)} - x^{(k)})^2}{x^{(k+2)} - 2x^{(k+1)} + x^{(k)}}, \quad \text{or} \\ \hat{x}^{(k)} &= x^{(k)} - \frac{(\Delta x^{(k)})^2}{\Delta^2 x^{(k)}}, \end{aligned} \quad (4.4)$$

where Δ is the difference operator; $\Delta x^{(k)} = x^{(k+1)} - x^{(k)}$, and the higher order operator is applied recursively; $\Delta^2 x^{(k)} = \Delta(\Delta x^{(k)}) = \Delta x^{(k+1)} - \Delta x^{(k)}$ [64]. From an initial value, $x^{(0)}$, two function evaluations, iterations of the underlying fixed point process, must be performed to obtain $x^{(1)}$ and $x^{(2)}$, before equation (4.4) can be computed.

The derivation of Aitken's method assumes an underlying linearly converging series of iterates. The order of convergence of the resulting Aitken accelerated series is still linear; however, this series converges faster than the original series [63]. We discuss Aitken's Δ^2 method and Steffensen iteration together, as Steffensen iteration is a straightforward extension of Aitken's method, whereby the Aitken value, $\hat{x}^{(k)}$, is used to continue the fixed point iteration, i.e. $x^{(k+1)} = \hat{x}^{(k)}$. Despite the striking similarity, Steffensen's method was seemingly developed shortly after (1933) and without knowledge of Aitken's method [75]. Steffensen iteration can achieve quadratic convergence [64,76,77]. Further theoretical convergence results for the Steffensen

method are established by Nievergelt [78] and in a series of papers by Noda [79–81].

Aitken and Steffensen iteration can be extended to multivariate systems [64]. In the following statements, we outline the method for an $N+1$ dimensional system, $X^{(k)} = [x_0^{(k)}, x_1^{(k)}, \dots, x_N^{(k)}]^T \in \mathbb{R}^{N+1}$, as appropriate for use with the FBSM

$$\hat{X}^{(k)} = X^{(k)} - \Delta \mathcal{X}^{(k)} (\Delta^2 \mathcal{X}^{(k)})^{-1} \Delta X^{(k)}, \quad (4.5)$$

where $\Delta X^{(k)} = X^{(k+1)} - X^{(k)}$, $\mathcal{X}^{(k)}$ is a matrix constructed with columns $(X^{(k)}, X^{(k+1)}, \dots, X^{(k+N)})$, such that $\mathcal{X}^{(k)}$ is a square matrix of dimension $N+1$, with $\Delta \mathcal{X}^{(k)} = \mathcal{X}^{(k+1)} - \mathcal{X}^{(k)}$, and $\Delta^2 \mathcal{X}^{(k)} = \Delta \mathcal{X}^{(k+1)} - \Delta \mathcal{X}^{(k)}$.

In the form given by equation (4.5), there are glaring issues with using the Steffensen method to accelerate convergence of the FBSM. Setting aside the question of whether $\Delta^2 \mathcal{X}^{(k)}$ is invertible, forming $\mathcal{X}^{(k)}$ would require $\mathcal{O}(N)$ iterations of the FBSM to be performed, and since N relates to the number of time points in the discretization of the ODEs in the FBSM, we expect N to be large, relative to the number of iterations required for the FBSM to converge without acceleration.

We instead consider a modification of the Steffensen method, requiring fewer function evaluations per iteration. Introduce $m < N$, and define $\Delta X^{(k)} = X^{(k+1)} - X^{(k)}$ as before, $\mathcal{X}^{(k)}$ is now a rectangular matrix constructed with columns $(X^{(k)}, X^{(k+1)}, \dots, X^{(k+m+1)})$, such that $\mathcal{X}^{(k)} \in \mathbb{R}^{(N+1) \times (m+2)}$, with $\Delta \mathcal{X}^{(k)} = \mathcal{X}^{(k+1)} - \mathcal{X}^{(k)}$, and $\Delta^2 \mathcal{X}^{(k)} = \Delta \mathcal{X}^{(k+1)} - \Delta \mathcal{X}^{(k)}$, both of dimension $(N+1) \times m$. We now interpret the matrix inverse in equation (4.5) as the Moore–Penrose pseudoinverse [82], a generalization of the matrix inverse for singular and rectangular matrices; we discuss this further in §3 of the electronic supplementary material. This partial implementation requires only $m+1$ function evaluations per iteration. For the remainder of this document when referring to the Steffensen method we are specifically referring to this partial Steffensen implementation. We present the derivation of the multivariate Aitken–Steffensen method and outline where the partial implementation differs in §3 of the electronic supplementary material.

4.4. Anderson acceleration

Anderson acceleration or Anderson mixing, originally denoted as the extrapolation algorithm by Anderson in the 1960s [83], is a technique developed for accelerating convergence of fixed point iteration problems with slowly converging Picard iterations [84]. Anderson acceleration is of particular interest in this work, as it has recently been implemented to accelerate the convergence of a regularized version of the FBSM [48]. In contrast to a standard fixed point iteration, whereby the next iterate depends only on the immediately preceding iterate, Anderson acceleration has ‘memory’ through the inclusion of previous iterates [85]. Unlike other methods considered in this work, Anderson acceleration explicitly uses the differences between residuals of subsequent iterates alongside iterates and their differences in computing future iterates.

Anderson acceleration involves solving a least-squares problem at each iteration. The problem can be expressed in both constrained and unconstrained forms, with the updating step dependent on the form [86,87]. We solve the following unconstrained least-squares problem in each iteration of

Anderson acceleration:

$$\gamma = \arg \min_{\gamma} (\|G - \gamma dG\|), \quad (4.6)$$

where $\arg \min(\cdot)$ returns the argument, γ , that minimizes the expression in equation (4.6). The corresponding updating step is

$$X^{(k+1)} = X^{(k)} + G^{(k)} - (dX^{(k-1)} + dG^{(k-1)})\gamma, \quad (4.7)$$

where $G^{(k)} = F(X^{(k)}) - X^{(k)}$ is the residual, $dX^{(k)}$ is a matrix with columns $(\Delta X^{(k-m)}, \Delta X^{(k-m+1)}, \dots, \Delta X^{(k)})$, and $dG^{(k)}$ is a matrix with columns $(\Delta G^{(k-m)}, \Delta G^{(k-m+1)}, \dots, \Delta G^{(k)})$, and m indicates the number of previous iterates that are incorporated.

4.5. Acceleration methods applied to typical fixed point problems

As a precursor to implementing these acceleration methods for control problems, we apply them to solve example non-linear systems of dimension 2×2 , 3×3 and 4×4 . We provide these systems and the results of the acceleration methods compared to standard fixed point iteration in §5 of the electronic supplementary material. We do not discuss these results in detail, although broad comparisons regarding the application of the acceleration methods to these systems and to control problems are made in §6. We provide code on GitHub (<https://github.com/Jesse-Sharp/Sharp2021>) for implementing the acceleration algorithms to solve systems of arbitrary size.

5. Acceleration results

In this section, we discuss the results of applying the acceleration algorithms. When discussing results we are solely focused on reducing \mathcal{N} , the number of function evaluations required for the control problems to reach convergence; as in all convergent cases we arrive at the same optimal control results. We first discuss the aspects of each method that can be tuned, then outline the results of the standard FBSM with the best choice of ω but without acceleration methods applied, to establish a baseline against which to compare the acceleration methods. A detailed suite of results for each control problem and each acceleration method, for various combinations of tuning parameters, is provided in §6 of the electronic supplementary material.

5.1. Tuning

Each method we consider has parameters that can be tuned to improve performance for a given problem. For the FBSM without acceleration, we can select $\omega \in [0, 1)$; the parameter that weights the contribution of the control from the previous iteration, and the newly calculated control, to the control used in the next iteration, as stated in equation (2.6). Control problems based on the linear model are able to converge via direct updating, as given in equation (2.5), equivalent to $\omega = 0$. Increasing ω in this case only serves to increase \mathcal{N} , so we do not attempt to tune ω when considering the linear model. Using the standard FBSM without acceleration the continuous linear problem requires $\mathcal{N} = 57$ while the bang-bang linear problem requires only $\mathcal{N} = 8$.

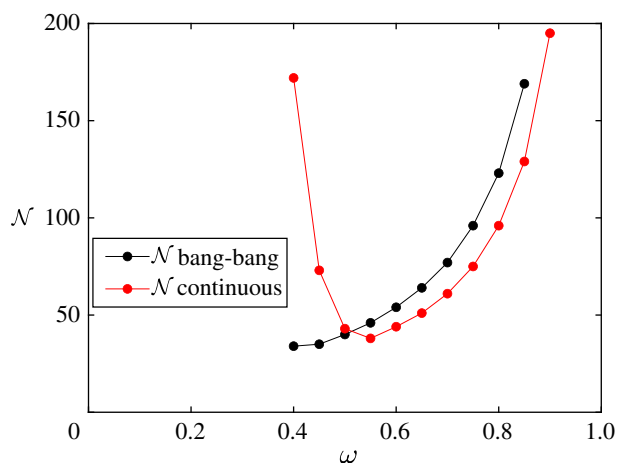


Figure 9. Here, we plot \mathcal{N} against $\omega \in [0, 1)$ in increments of 0.05, for the AML continuous (in red) and bang-bang (in black) control problems, using FBSM with no acceleration. Results correspond to model parameters given in table 1, time-step $dt = 4.88 \times 10^{-4}$, over the interval $0 \leq t \leq 10$. The continuous problem is solved with pay-off weightings of $a_1 = 1$ for control, and $a_2 = 2$ for the state variable $L(t)$, while the bang-bang problem is solved with $a_1 = 1$ and $a_2 = 3$. Where an ω value does not have a corresponding marker, this indicates that the procedure fails to converge within $\mathcal{N} = 200$.

In figure 9, we plot \mathcal{N} against $\omega \in [0, 1)$, for the continuous and bang-bang AML problems. As expected, for small ω we find that the problem does not converge, and for large ω , \mathcal{N} increases. For the continuous AML problem we identify $\omega = 0.55$ as the best choice, with $\mathcal{N} = 38$. For the bang-bang AML problem, we find that $\omega = 0.4$ is best, with $\mathcal{N} = 34$.

Recall that fixed endpoint problems are solved using the adapted FBSM; this entails solving several control problems to convergence with the FBSM. Each of these control problems can have a different optimal ω . In this instance, $\omega = 0.55$ also happens to be best for the AML fixed endpoint problem if holding ω constant, when considering $\omega \in [0, 1)$ at increments of 0.05. These ω values will not coincide in general. When applying the acceleration methods to the fixed endpoint problems, we employ the tuning parameters that perform best for the continuous problem. This does not imply that we are using the best tuning parameters for the acceleration methods in the context of the fixed endpoint problem. Importantly, this demonstrates whether or not the techniques can effectively reduce Σ , the cumulative function evaluations required for convergence of the adapted FBSM for fixed endpoint problems, without requiring prohibitive tuning.

With the Wegstein method, we only select ω for the two FBSM iterations required for initialization, and specify n , such that we update q every n th iteration. We generate results for $n \in \{1, 2, \dots, 10\}$. We also bound q ; however identifying suitable bounds is challenging. In this work, we select bounds that perform reasonably, but acknowledge that we do not search for optimal bounds, nor do we think that attempting to do so is realistic. This drawback of Wegstein's method contributes to its inconsistent performance relative to other methods. For the partial Aitken–Steffensen methods, we choose ω , and the parameter m that specifies the dimension of the $(N+1) \times m$ matrices in the updating step, requiring $m+1$ function evaluations per iteration. We generate results for $m \in \{1, 2, \dots, 10\}$. Similarly, for Anderson acceleration we select ω and M , where M determines the

maximum number of previous iterations to retain when solving the least-squares problem and performing the updating step. We produce results for $M \in \{1, 2, \dots, 10\}$.

5.2. Wegstein method

For the continuous linear problem, we apply bounds of $-2 \leq q \leq 0$. For the bang-bang linear problem, we leave q unbounded. For both AML problems, we apply bounds $-1 \leq q \leq 1$. We explore the effect of updating q every n th iteration, $n \in \{1, 2, \dots, 10\}$. For the continuous linear problem, $n = 4$ minimizes \mathcal{N} , although $n \in \{1, 2, \dots, 5\}$ all perform well. For the linear bang-bang problem, the Wegstein method converges without bounding on q , and varying n does not affect convergence. The Wegstein method outperforms other acceleration methods for the linear bang-bang problem with $\mathcal{N} = 9$, but does not improve upon $\mathcal{N} = 8$ for the FBSM without acceleration.

For the continuous AML problem, the performance of Wegstein's method is inconsistent. With $n = 6$ and $\omega = 0.55$, the Wegstein method achieves convergence with $\mathcal{N} = 26$, outperforming the FBSM; however, almost every other combination of tuning parameters considered with $\omega \in [0, 1)$ and $n \in \{1, 2, \dots, 10\}$ require larger \mathcal{N} than the FBSM without acceleration. Generally, increasing n produces worse outcomes. We do, however, observe that the Wegstein method can induce convergence for $\omega < 0.4$, where the standard FBSM does not converge. For the bang-bang AML problem, the Wegstein method appears more robust; consistently outperforming the standard FBSM across most of the tuning parameter space. The best result requires only $\mathcal{N} = 9$, with $\omega = 0$ and $n = 7$, although several other combinations of tuning parameters are similarly successful. For $\omega \geq 0.4$, corresponding to values that the underlying FBSM converges, we find that moderate $n \in \{3, 4, \dots, 7\}$ produces the best results; while for smaller ω , larger $n \in \{6, 7, \dots, 10\}$ consistently performs best. Once again we observe that convergence is achieved for ω values where the underlying FBSM would not converge.

For the linear fixed endpoint problem, the adapted FBSM with the Wegstein method consistently generates a moderate reduction in Σ , compared to the adapted FBSM without acceleration, for all $n \in \{1, 2, \dots, 10\}$, $-2 \leq q \leq 0$. For the AML fixed endpoint problem we do not observe improvement. Using the tuning parameters that perform best for the continuous AML problem, we find that Σ for the adapted FBSM with Wegstein's method is more than double that of the adapted FBSM without acceleration. This results from the inconsistency of Wegstein's method with poor tuning. In §6 of the electronic supplementary material, it can be seen that some control problems within the adapted FBSM that require $\mathcal{N} \approx 50$ without Wegstein's method, require $\mathcal{N} \approx 200$ with the specified Wegstein tuning parameters.

5.3. Partial Aitken–Steffensen method

Both Aitken and Steffensen methods significantly and consistently outperform the FBSM without acceleration for the continuous linear problem. The Aitken method performs best for $m \in \{1, 2, 3\}$, requiring $\mathcal{N} = 12$. Steffensen's method performs best when $m = 6$, requiring only $\mathcal{N} = 8$. In the linear bang-bang case, both Aitken and Steffensen methods perform marginally worse than the FBSM without acceleration, which requires only $\mathcal{N} = 8$. In the best cases, with $m = 1$ the Aitken

method requires $\mathcal{N} = 10$, and with $m = 7$ the Steffensen method requires $\mathcal{N} = 9$.

For the continuous AML problem, we observe a stark difference between the Aitken and Steffensen methods; while the Steffensen method is able to achieve convergence for values of ω where the underlying FBSM fails to converge, $\omega \leq 0.35$, particularly for $m \in \{1, 2, 3, 4\}$, the Aitken method only converges to the optimal control for ω values where the underlying FBSM converges. For $\omega \leq 0.35$, the Aitken method achieves apparent convergence; the iterative procedure terminates as the convergence criteria are met. However, explicitly calculating the pay-off associated with these controls via equation (3.9), and comparing this result to the pay-off associated with the control obtained via the standard FBSM, indicates that the controls obtained via the Aitken method for $\omega \leq 0.35$ are not optimal, as they fail to minimize J . The best result for the Aitken method, with $\omega = 0.5$ and $m = 5$, requires $\mathcal{N} = 30$, marginally improving on the FBSM without acceleration, requiring $\mathcal{N} = 38$. Steffensen's method produces more significant improvements, requiring only $\mathcal{N} = 19$ with $\omega = 0.5$ and $m = 5$. In each case, neighbouring combinations of tuning parameters also yield equivalent or comparable improvement over the standard FBSM. In the bang-bang AML problem, we observe similar behaviour; for ω values that the underlying FBSM fails to converge, the Steffensen method consistently converges. The Aitken method achieves apparent convergence for these values of ω ; the iterative procedure terminates as the convergence criteria are met; however, the resulting controls contain intermediate values between the lower and upper bounds. As such the resulting controls are not bang-bang, so we treat these results as failing to converge. At best, Aitken's method requires $\mathcal{N} = 8$, with $\omega = 0.5$ and $m = 1$, while Steffensen's requires only $\mathcal{N} = 7$, with $\omega = 0.5$ and $m = 5$. The vast majority of tuning parameter combinations yield improvements over the $\mathcal{N} = 34$ of the standard FBSM.

Aitken and Steffensen methods consistently offer significant improvement over the standard adapted FBSM for the linear fixed endpoint problem for $m \in \{1, 2, \dots, 10\}$, with the exception of $m = 1$ for the Steffensen method, which yields only marginal improvement. Using the best performing tuning parameters for the continuous AML problem, we find that both Aitken and Steffensen methods improve upon the standard adapted FBSM for the AML fixed endpoint problem. Relative to $\Sigma = 434$ required without acceleration, the $\Sigma = 360$ required with Aitken's method reflects a modest improvement, while the $\Sigma = 238$ required with the Steffensen method is a significant improvement.

5.4. Anderson acceleration

Anderson acceleration performs exceptionally well on the continuous linear problem, requiring only $\mathcal{N} = 7$ for $M \in \{4, 5, \dots, 10\}$ compared to $\mathcal{N} = 57$ for the standard FBSM. For the linear bang-bang problem, however, it is the worst performing acceleration method; achieving at best $\mathcal{N} = 11$, with $M = 1$.

Similarly to the Wegstein and Steffensen methods, Anderson acceleration achieves convergence in both the continuous and bang-bang AML problems for ω values where the underlying FBSM fails to converge. Anderson acceleration achieves the best individual result for the continuous AML problem, requiring only $\mathcal{N} = 17$, with $\omega = 0.85$ and $M = 6$. Again, we observe comparable improvement over a wide range of

tuning parameters. For the bang-bang AML problem Anderson acceleration consistently outperforms FBSM without acceleration, particularly for $\omega < 0.7$, at best requiring $\mathcal{N} = 17$, with $\omega = 0.35$ and $M \in \{7, 8, 9, 10\}$, with other non-neighbouring tuning parameter combinations also yielding $\mathcal{N} = 17$.

For both the linear and AML fixed endpoint problems Anderson acceleration produces the most significant reduction in Σ , and improves upon the adapted FBSM over a wide range of tuning parameters. In the linear case, Anderson acceleration requires only $\Sigma = 24$ for $M \in \{4, 5, \dots, 10\}$. In the AML fixed endpoint problem, using the tuning parameters that perform best for the continuous AML problem, Anderson acceleration converges in only $\Sigma = 204$; less than half as many as the standard adapted FBSM.

5.5. Method comparison with best tuning

Results presented in figures 10 and 11 provide comparison of the error, ε , as each method approaches convergence, for the linear and AML problems, respectively. Error is measured as the Euclidean norm of the difference between subsequent controls; $\varepsilon = \|F(X^{(k)}) - X^{(k)}\|$, with the exception of Aitken's method, where error is measured as the difference between subsequent values in the Aitken series; $\varepsilon = \|\hat{X}^{(k)} - \hat{X}^{(k-1)}\|$. Convergence is achieved when $\varepsilon \leq 1 \times 10^{-10}$, marked in black dash. In each case, we are plotting the result that minimizes \mathcal{N} for each method, over the space of tuning parameters considered, including the best tuning of ω for the FBSM without acceleration. Error is plotted on a logarithmic scale. For the linear bang-bang problem with the Wegstein and Anderson methods, and the AML bang-bang problem with the Wegstein method, the error after the final iteration is $\varepsilon = 0$, as two subsequent iterates for the control are identical. This is represented on the logarithmic scale as a line that intersects the horizontal axis.

6. Discussion

Modelling processes in systems biology is complex; frequently involving large state systems consisting of several ODEs [88–90], including canonical examples such as the mitogen-activated protein kinase cascade [91], Wnt/ β -catenin signalling pathway [92], early incarnations of whole-cell models [93,94], and other cellular signalling, metabolic and regulatory processes and mechanisms [95,96]. The acceleration methods we implement act only on the control term; the number and form of state equations has no bearing on the mathematical and computational complexity of the acceleration methods. As such, the methods scale excellently with system complexity. In this section, we discuss the results presented in §5, and draw insights into the convergence behaviour of the FBSM when augmented with acceleration techniques. We highlight opportunities for application of these methods, and outline several avenues for further investigation.

6.1. Acceleration outcomes

In evaluating the performance of each acceleration method, we are interested in: (1) how significantly they are able to reduce \mathcal{N} , (2) method robustness and (3) method accessibility. In this context, we use robustness to refer to how consistently the method outperforms the best tuned FBSM over the range of tuning

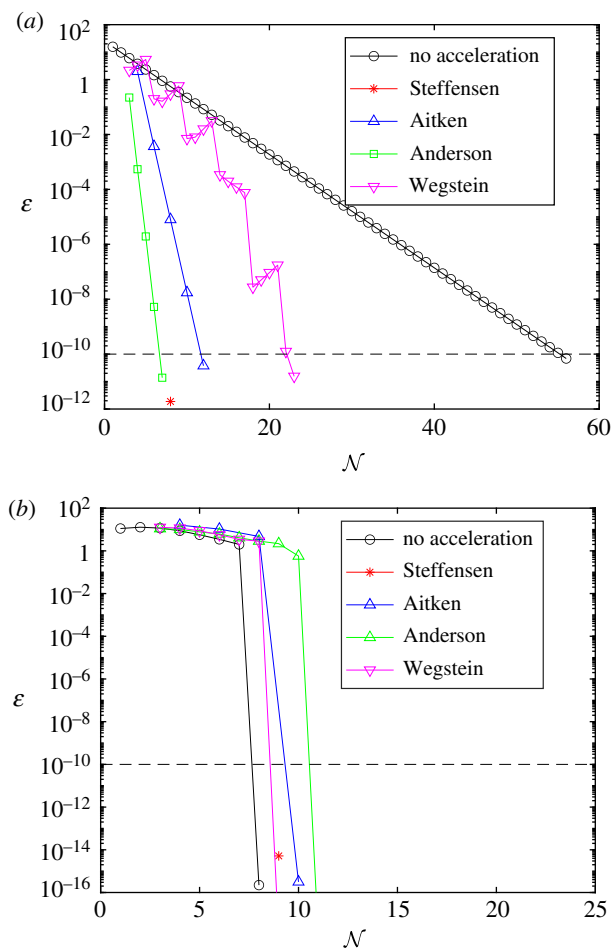


Figure 10. Convergence rates for the result that minimizes \mathcal{N} , for each acceleration method, when applied to the linear control problems. Results in this plot are produced with model parameter $\gamma = 0.5$, time-step $dt = 3.91 \times 10^{-3}$, over the interval $0 \leq t \leq 1$, with pay-off weighting $a = b = 1$ for the continuous control (a), and $a = 1, b = 3$ for the bang-bang control (b). The tolerance of 1×10^{-10} required for convergence is marked in black dash. As the methods do not necessarily use the same number of function evaluations per iteration, markers indicate each time ε is computed. Continuous control results correspond to the FBSM with no acceleration, the partial Steffensen method with $m = 6$, partial Aitken method with $m = 1$, Anderson acceleration with $M = 4$ and Wegstein with bounds $-2 \leq q \leq 0$, updating q every 4th iteration. Bang-bang results correspond to the FBSM with no acceleration, the partial Steffensen method with $m = 7$, partial Aitken method with $m = 1$, Anderson acceleration with $M = 1$ and Wegstein without bounds on q , updating q every iteration. The standard FBSM outperformed all acceleration methods in solving the linear bang-bang control problem. We attribute this to how few iterations were required ($\mathcal{N} = 8$) for convergence without acceleration.

parameters considered. We judge the accessibility of each method based on implementation and conceptual complexity. Overall, we find that the acceleration methods, particularly Anderson and Steffensen, significantly and robustly reduce \mathcal{N} . Anderson acceleration appears most effective for continuous control, while the Steffensen method appears best for bang-bang control. The Aitken method occasionally outperforms Steffensen, but overwhelmingly the Steffensen method appears to be the better option of the two for the range of parameters we consider. Implementing the Anderson and Steffensen methods introduces challenge beyond that of the underlying FBSM, although it is not prohibitively difficult; particularly with reference to the code where we implement these methods, that we

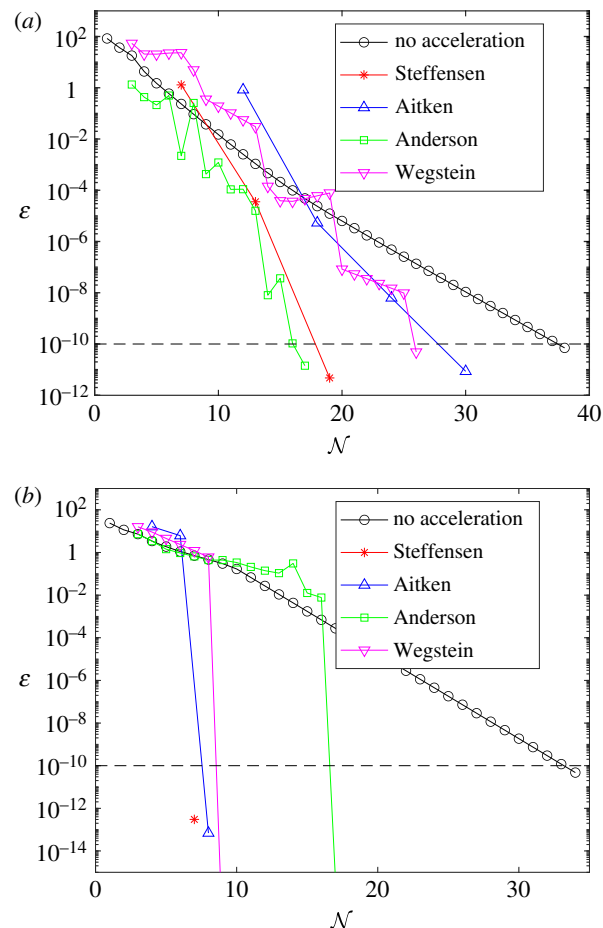


Figure 11. Convergence rates for the converged result that minimizes \mathcal{N} , for each acceleration method, when applied to the AML control problems. Results in this plot are produced with model parameter $\gamma = 0.5$, time-step $dt = 4.88 \times 10^{-4}$, over the interval $0 \leq t \leq 1$, with pay-off weighting $a = b = 1$ for the continuous control (a), and $a = 1, b = 3$ for the bang-bang control (b). The tolerance of 1×10^{-10} required for convergence is marked in black dash. As the methods do not necessarily use the same number of function evaluations per iteration, markers indicate each time ε is computed. Continuous control results correspond to the FBSM with no acceleration, $\omega = 0.55$, the partial Steffensen and partial Aitken methods with $m = 5$ and $\omega = 0.5$, Anderson acceleration with $M = 6$ and $\omega = 0.85$, and Wegstein method with $\omega = 0.55$, bounds $-1 < q < 1$, updating q every 6th iteration. Bang-bang results correspond to the FBSM with no acceleration, $\omega = 0.4$, the partial Steffensen method with $m = 5$ and $\omega = 0.5$, partial Aitken method with $m = 1$ and $\omega = 0.5$, Anderson acceleration with $M = 7$ and $\omega = 0.35$, and Wegstein method with $\omega = 0$, bounds $-1 < q < 1$, updating q on the 7th iteration.

make available on GitHub (<https://github.com/Jesse-Sharp/Sharp2021>). Both methods introduce conceptual complexity, perhaps marginally less so for the Steffensen method due to the similarities it shares with the familiar Newton's method.

We produce heatmaps to visualize the convergence behaviour of the acceleration methods across the range of tuning parameters considered. Figure 12 corresponds to the AML continuous control problem, while figure 13 corresponds to the AML bang-bang control problem. Recall that with the tuning of ω that minimizes \mathcal{N} , the FBSM with no acceleration requires $\mathcal{N} = 38$ for the AML continuous control problem, and $\mathcal{N} = 34$ for the AML bang-bang control problem. Tuning parameter combinations that reflect a reduction in \mathcal{N} relative to these FBSM results are shaded in the green spectrum, while




Figure 12. This heatmap provides insight into the convergence behaviour of the acceleration methods for the AML continuous control problem. Here, we visualize \mathcal{N} against ω and the method specific tuning parameter: n for Wegstein (a), m for partial Aitken (b) and partial Steffensen (c), and M for Anderson acceleration (d). Tuning parameter combinations requiring $\mathcal{N} = 38$, equivalent to the best tuned FBSM without acceleration, are shaded yellow. Colours in the green-yellow spectrum represent a reduction in \mathcal{N} relative to FBSM without acceleration, while colours in the yellow-red spectrum represent an increase in \mathcal{N} .

worse performing combinations are shaded in the red spectrum. The midpoint of the colour spectra, yellow, corresponds to the FBSM result with the best tuning, without acceleration. Simulations are terminated when \mathcal{N} exceeds 100; reflecting a combination of tuning parameters that do not yield convergence within this specified maximum. Data supporting these heatmaps, and similar results for the linear control problems are provided in §6 of the electronic supplementary material.

In identifying tuning parameter combinations that yield significant reductions in \mathcal{N} , we are looking for bright green areas in the heatmaps. We assess the robustness of each method by considering whether we observe large contiguous areas in the green spectrum, such as in figure 12c, indicating robustness, or patchy areas with both green spectrum and red spectrum, such as figure 13d, suggesting a lack of robustness.

In table 2, we provide our subjective but informed rating of the methods against the criteria of reduction in \mathcal{N} , robustness and accessibility. We consider the continuous and bang-bang control cases separately in terms of reduction in \mathcal{N} and robustness.

Despite its conceptual simplicity and straightforward implementation, Wegstein's method is significantly hampered by the difficulty in choosing bounds. If there were a more

informed approach for identifying suitable bounds, Wegstein's method could be particularly useful for bang-bang control problems. Due to the effect of ω , intermediate control iterates of the FBSM do not appear bang-bang; as such the bulk of \mathcal{N} are incurred in refining the control about the switching points. Wegstein's method can accelerate this refinement by adaptively setting $q_i = 0$ where appropriate.

6.2. Convergence insights

As outlined in §5, the linear model control problems converge with $\omega = 0$. It may at first seem counterintuitive that Wegstein's method can improve upon this, given that the computed q in Wegstein's method acts as a stand-in for ω . There are two aspects of distinction that enable Wegstein's method to generate improvement in this case: first, while ω is held constant both within the time discretization and between iterations, the element-wise nature of Wegstein's method enables each element of the discretization to have a different value, q_i , $i \in \{0, 1, \dots, N\}$, and q can be updated between iterations; second, observing the values of q_i in Wegstein's method indicates that $q < 0$ can be appropriate. This suggests that $\omega < 0$ could also be used to accelerate the standard implementation of the FBSM. Preliminary investigation suggests that this is true for the linear model;




Figure 13. This heatmap provides insight into the convergence behaviour of the acceleration methods for the AML bang-bang control problem. Here, we visualize \mathcal{N} against ω and the method specific tuning parameter: n for Wegstein (a), m for partial Aitken (b) and partial Steffensen (c), M for Anderson acceleration (d). Tuning parameter combinations requiring $\mathcal{N} = 34$, equivalent to the best tuned FBSM without acceleration, are shaded yellow. Colours in the green-yellow spectrum represent a reduction in \mathcal{N} relative to FBSM without acceleration, while colours in the yellow-red spectrum represent an increase in \mathcal{N} .

Table 2. Method comparison. We rate the methods considered in this work against key factors such as the reduction in \mathcal{N} that they deliver and how robustly they perform over the range of tuning parameters considered, for both continuous (Cts) and bang-bang (BB) control problems. We also consider how accessible the methods are from the standpoints of ease of implementation (Imp) and conceptual complexity. Methods are rated as being either strongly positive ($\checkmark\checkmark$), positive (\checkmark), neutral (\sim), negative (\times) against each aspect.

method	reduction in \mathcal{N}		robustness		accessibility	
	Cts	BB	Cts	BB	Imp	complexity
FBSM	\sim	\sim	\checkmark	\checkmark	\checkmark	\checkmark
Wegstein	\sim	$\checkmark\checkmark$	\times	\checkmark	\checkmark	\sim
Aitken	\checkmark	\checkmark	\sim	\checkmark	\sim	\sim
Steffensen	$\checkmark\checkmark$	$\checkmark\checkmark$	\checkmark	$\checkmark\checkmark$	\sim	\sim
Anderson	$\checkmark\checkmark$	\checkmark	\checkmark	\checkmark	\sim	\times

however, we do not pursue this further as we expect it to be of limited applicability beyond contrived problems.

We apply the acceleration methods to small nonlinear test systems in §5 of the electronic supplementary material. We know these systems have multiple fixed points; all methods we consider aside from Aitken's method, in some of our

examples, reach different fixed points to fixed point iteration. By contrast, when applied to accelerate control problems, we observe only the Aitken method converging to a result other than the optimal control obtained via the FBSM, as discussed in §5. This apparent convergence of the Aitken method to controls that are not optimal is a significant deterrent to

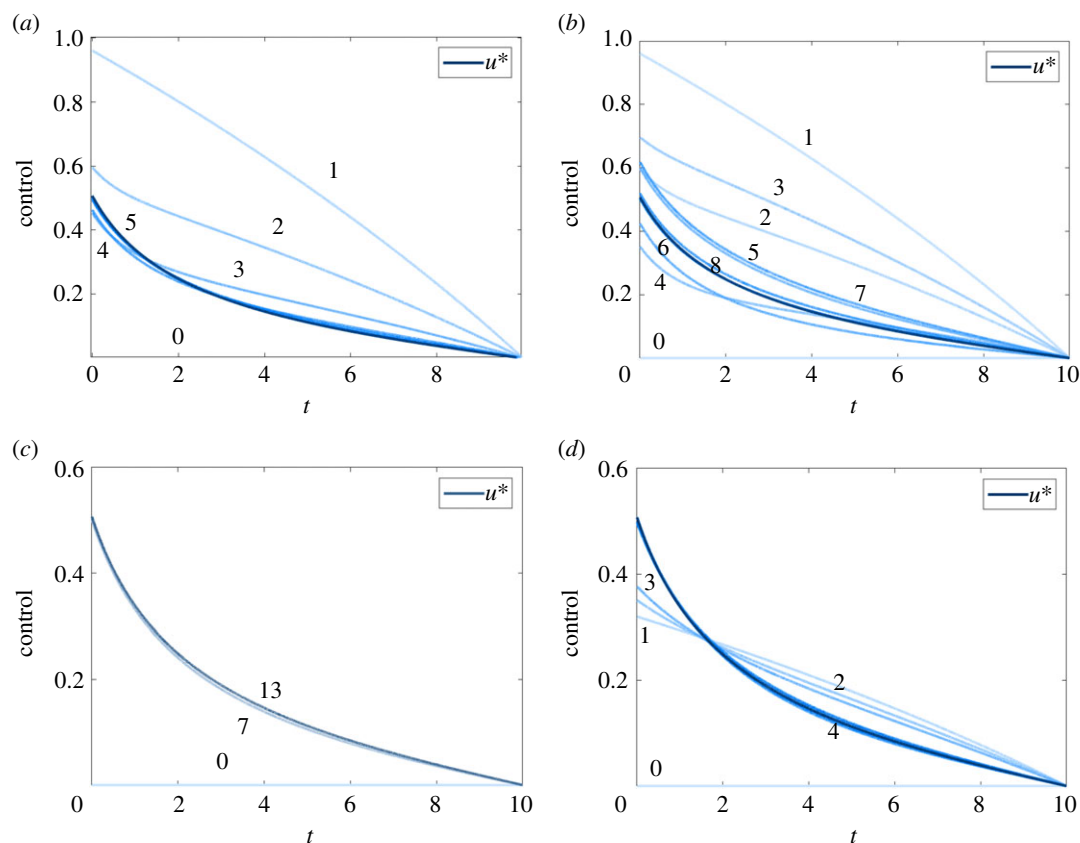


Figure 14. Here we observe the iterates of the control in the AML continuous control problem as it converges, for (a) the FBSM with no acceleration, (b) the Wegstein method, (c) the partial Steffensen method and (d) Anderson acceleration. Initial iterates are shown in light blue, while darker blue denotes later iterates. Results for the Aitken method are not shown as they are visually similar to the Steffensen result. We present the results corresponding to the tuning parameters that minimize \mathcal{N} , outlined in §5. Where it is visually distinguishable, we indicate the number of function evaluations corresponding to a particular iterate. While all methods produce the same eventual result for u^* , they follow considerably different series of iterates. Note that the vertical scale in (a) and (b) differs from that of (c) and (d).

using the Aitken method in situations where the optimal control is not known *a priori*. Outside of this issue with the Aitken method, each acceleration method produces the same optimal control for a given problem. However, they each approach the converged control differently. In figure 14, we plot the control as it converges for the FBSM and acceleration methods. In the code, we provide on GitHub (<https://github.com/Jesse-Sharp/Sharp2021>), users can view the control iterates of each method as they approach convergence. Visualizing these methods as they converge gives insight into how they may be able to arrive at different fixed points; under certain circumstances the accelerated series of iterates may leave the basin of attraction for the fixed point found via fixed point iteration.

6.3. Summary and outlook

In this work, we review the theory and implementation of the FBSM for solving TPBVPs that arise from application of PMP in solving optimal control problems. We study a single-variable linear model and a multiple-variable nonlinear model and consider continuous, bang-bang and fixed endpoint control problems. Conceptualizing the FBSM as a fixed point iteration, we leverage and adapt existing acceleration methods to significantly and robustly increase the convergence rate of the FBSM for a range of optimal control problems. The Anderson and partial Steffensen methods appear to perform best, without requiring prohibitive tuning.

Accelerating the convergence of the FBSM, and reducing the importance of appropriately selecting ω for a single control problem, is promising. That said, the real utility of the robust acceleration methods in this work is in application to families of control problems. We provide a glimpse of this benefit through considering fixed endpoint control problems, though other excellent opportunities for application arise due to the uncertainty prevalent in the life sciences. First, it is common for there to be uncertainty around model parameters and structure [97,98]. In this case, solving optimal control problems over several model structures and sets of model parameters provides insight into the sensitivity of the control strategy [99–102]. Secondly, when performing multi-objective optimization, a trade-off is made between objectives. For example, in equation (3.14), we seek to minimize the cumulative negative impact of leukaemia and of the control; parameters a_1 and a_2 weight the relative importance of each objective. In practical applications, it is not always clear how to determine these weightings. It can therefore be useful to generate a family of optimal controls that are each optimal for their specific combination of pay-off weighting parameters, akin to a Pareto frontier [49,103,104]. Producing these sets of control results benefits significantly from acceleration techniques such as the Anderson and Steffensen methods, where a consistent reduction in \mathcal{N} is obtained without optimal tuning.

In this work, multi-objective optimization is considered in the form of a control problem with a single cost function comprising a scalar combination of state and control terms. More

generally, multi-objective optimization can be formulated as a control problem with a vector-valued cost function, with the goal of minimizing each component simultaneously. There are a range of strategies for handling multi-objective optimal control problems formulated in this way, and we direct readers to [105] for a recent and extensive survey.

Here, we have only considered systems subject to a single control. While this is reflective of the vast majority of applications featured in the control literature, there are instances where we are interested in applying multiple controls simultaneously [13,14,51]. The FBSM can be readily applied to solve problems with multiple controls [51]; a logical extension of this work is to adapt the acceleration methods or identify suitable alternative methods for accelerating convergence of the FBSM for problems with multiple controls.

Over a range of tuning parameters the Wegstein, Steffensen and Anderson methods are able to induce convergence where the underlying FBSM fails to converge; such as in the AML control problems with $\omega = 0$. This behaviour has been documented for Anderson acceleration [87] and Wegstein's method [67] when applied to standard fixed point iteration problems. This presents an opportunity for future exploration, in identifying control problems that cannot be solved via the FBSM for any ω , and attempting to produce solutions using these acceleration techniques.

The examples we consider in this work include a variety of control problem formulations that arise in systems biology. However, it is worth noting that the examples are not exhaustive. Further challenges can be introduced; either through the formulation of the control problem, or as a result of the

behaviour of the underlying system. Examples of such challenges include control problems with singular arcs, path constraints, multiple local solutions, discontinuous dynamics and sensitivity to the initial guess of the control [35]. These challenges can introduce numerical difficulties, and complications in terms of the optimal control theory; for example, control problems with singular arcs typically require additional necessary conditions for optimality beyond those obtained from the PMP [38]. A thorough assessment of the appropriateness of the FBSM as a method for solving control problems with such complications is an avenue for further investigation. We stress that the acceleration techniques that we develop and survey in this work are able to accelerate convergence when compared to a naive FBSM implementation, and in some cases induce convergence where the naive FBSM fails to converge. We anticipate that these trends will persist if these acceleration techniques are applied to appropriately conceived implementations of the FBSM for the various complications outlined here.

Data accessibility. This article has no additional data.

Competing interests. We declare we have no competing interests.

Funding. J.A.S. acknowledges support from the Australian Government Research Training Program and the AF Pillow Applied Mathematics Trust. J.A.S. and K.B. acknowledge support from the Australian Centre of Excellence for Mathematical and Statistical Frontiers (CE140100049). M.J.S. is supported by the Australian Research Council (grant no. DP200100177).

Acknowledgements. We thank two anonymous referees for their helpful comments.

References

- Abbasi Z, Zamani I, Amiri Mehra AH, Shafieirad M, Ibeas A. 2020 Optimal control design of impulsive SQEIR epidemic models with application to COVID-19. *Chaos, Solitons Fractals* **139**, 110054. (doi:10.1016/j.chaos.2020.110054)
- Neilan RLM, Schaefer E, Gaff H, Fister KR, Lenhart S. 2010 Modeling optimal intervention strategies for cholera. *Bull. Math. Biol.* **72**, 2004–2018. (doi:10.1007/s11538-010-9521-8)
- Aspirin AP, Yangjin K. 2021 Polytherapeutic strategies with oncolytic virus–bortezomib and adjuvant NK cells in cancer treatment. *J. R. Soc. Interface* **18**, 20200669. (doi:10.1098/rsif.2020.0669)
- Burnett AK, Goldstone AH, Stevens RMF, Hann IM, Rees JHK, Gray RG, Wheatley K. 1998 Randomised comparison of addition of autologous bone-marrow transplantation to intensive chemotherapy for acute myeloid leukaemia in first remission: results of MRC AML 10 trial. *Lancet* **351**, 700–708. (doi:10.1016/S0140-6736(97)09214-3)
- Sharp JA, Browning AP, Mapder T, Burrage K, Simpson MJ. 2019 Optimal control of acute myeloid leukaemia. *J. Theor. Biol.* **470**, 30–42. (doi:10.1016/j.jtbi.2019.03.006)
- Falisse A, Serranoli G, Dembia CL, Gillis J, Jonkers I, De Groot F. 2019 Rapid predictive simulations with complex musculoskeletal models suggest that diverse healthy and pathological human gaits can emerge from similar control strategies. *J. R. Soc. Interface* **16**, 20190402. (doi:10.1098/rsif.2019.0402)
- Lai AKM, Dick TJM, Biewener AA, Wakeling JM. 2021 Task-dependent recruitment across ankle extensor muscles and between mechanical demands is driven by the metabolic cost of muscle contraction. *J. R. Soc. Interface* **18**, 20200765. (doi:10.1098/rsif.2020.0765)
- Patil NS, Dingwell JB, Cusumano JP. 2020 Task-level regulation enhances global stability of the simplest dynamic walker. *J. R. Soc. Interface* **17**, 20200278. (doi:10.1098/rsif.2020.0278)
- Del Vecchio D, Dy AJ, Qian Y. 2016 Control theory meets synthetic biology. *J. R. Soc. Interface* **13**, 20160380. (doi:10.1098/rsif.2016.0380)
- He F, Murabito E, Westerhoff HV. 2016 Synthetic biology and regulatory networks: where metabolic systems biology meets control engineering. *J. R. Soc. Interface* **13**, 20151046. (doi:10.1098/rsif.2015.1046)
- Chapman LAC, Shipley RJ, Whiteley JP, Ellis MJ, Byrne HM, Waters SL. 2014 Optimising cell aggregate expansion in a perfused hollow fibre bioreactor via mathematical modelling. *PLoS ONE* **9**, e105813. (doi:10.1371/journal.pone.0105813)
- Shipley RJ, Davidson AJ, Chan K, Chaudhuri JB, Waters SL, Ellis MJ. 2011 A strategy to determine operating parameters in tissue engineering hollow fibre bioreactors. *Biotechnol. Bioeng.* **108**, 1450–1461. (doi:10.1002/bit.23062)
- Baker CM, Armsworth PR, Lenhart SM. 2017 Handling overheads: optimal multi-method invasive species control. *Theor. Ecol.* **10**, 493–501. (doi:10.1007/s12080-017-0344-1)
- Bussell EH, Cunniffe NJ. 2020 Applying optimal control theory to a spatial simulation model of sudden oak death: ongoing surveillance protects tanoak while conserving biodiversity. *J. R. Soc. Interface* **17**, 20190671. (doi:10.1098/rsif.2019.0671)
- Carrasco LR, Baker R, MacLeod A, Knight JD, Mumford JD. 2010 Optimal and robust control of invasive alien species spreading in homogeneous landscapes. *J. R. Soc. Interface* **7**, 529–540. (doi:10.1098/rsif.2009.0266)
- Pontryagin LS, Boltyanskii VG, Gamkrelidze RV, Mischenko EF. 1962 *The mathematical theory of optimal processes [English translation]*. New York, NY: Interscience.
- Bellman RE. 1957 *Dynamic programming*. Princeton, NJ: Princeton University Press.
- Goldstine HH. 1980 *A history of the calculus of variations from the 17th through to the 19th century*. New York, NY: Springer.
- Bryson AE. 1996 Optimal control—1950 to 1985. *IEEE Control Syst. Mag.* **16**, 26–33. (doi:10.1109/37.506395)

20. Sargent RWH. 2000 Optimal control. *J. Comput. Appl. Math.* **124**, 361–371. (doi:10.1016/S0377-0427(00)00418-0)
21. Rao AV. 2009 A survey of numerical methods for optimal control. *Adv. Astronaut. Sci.* **135**, 497–528.
22. Polak E. 1973 An historical survey of computational methods in optimal control. *SIAM Rev.* **15**, 553–584. (doi:10.1137/1015071)
23. Zoppoli R, Sanguineti M, Gnecco G, Parisini T. 2019 *Neural approximations for optimal control and decision*. Cham, Switzerland: Springer.
24. Ibrahim M, Kallies C, Findeisen R. 2020 Learning-supported approximated control for autonomous vehicles in the presence of state dependent uncertainties. *2020 European Control Conf., St Petersburg, Russia, 12–15 May 2020*, pp. 338–343. (doi:10.23919/ECC51009.2020.9143737).
25. Treloar NJ, Fedorec AJH, Ingalls B, Barnes CP. 2020 Deep reinforcement learning for the control of microbial co-cultures in bioreactors. *PLoS Comput. Biol.* **16**, e1007783. (doi:10.1371/journal.pcbi.1007783)
26. Bonyah E, Sagoe AK, Kumar D, Deniz S. 2021 Fractional optimal control dynamics of coronavirus model with Mittag-Leffler law. *Ecol. Complex.* **45**, 100880. (doi:10.1016/j.ecocom.2020.100880)
27. Sweilam NH, Al-Mekhlafi SM, Baleanu D. 2019 Optimal control for a fractional tuberculosis infection model including the impact of diabetes and resistant strains. *J. Adv. Res.* **17**, 125–137. (doi:10.1016/j.jare.2019.01.007)
28. Jajarmi A, Baleanu D. 2019 On the fractional optimal control problems with a general derivative operator. *Asian J. Control* **23**, 1062–1071. (doi:10.1002/asjc.2282)
29. Jajarmi A, Ghanbari B, Baleanu D. 2019 A new and efficient numerical method for the fractional modeling and optimal control of diabetes and tuberculosis co-existence. *Chaos* **29**, 093111. (doi:10.1063/1.5112177)
30. Sweilam NH, Al-Mekhlafi SM, Baleanu D. In press. A hybrid fractional optimal control for a novel Coronavirus (2019-nCov) mathematical model. *J. Adv. Res.* (doi:10.1016/j.jare.2020.08.006)
31. Hassani H, Tenriero Machado JA, Mehrabi S. 2021 An optimization technique for solving a class of nonlinear fractional optimal control problems: application in cancer treatment. *Appl. Math. Model.* **93**, 868–884. (doi:10.1016/j.apm.2021.01.004)
32. Sweilam NH, Al-Mekhlafi SM, Assiri T, Atangana A. 2020 Optimal control for cancer treatment mathematical model using Atangana-Baleanu-Caputo fractional derivative. *Adv. Differ. Equ.* **2020**, 334. (doi:10.1186/s13662-020-02793-9)
33. Tsiantis N, Balsa-Canto E, Banga JR. 2018 Optimality and identification of dynamic models in systems biology: an inverse optimal control framework. *Bioinformatics* **34**, 2433–2440. (doi:10.1093/bioinformatics/bty438)
34. Oyarzún DA, Ingalls BP, Middleton RH, Kalamatianos D. 2009 Sequential activation of metabolic pathways: a dynamic optimization approach. *Bull. Math. Biol.* **71**, 1851–1872. (doi:10.1007/s11538-009-9427-5)
35. Tsiantis N, Banga JR. 2020 Using optimal control to understand complex metabolic pathways. *BMC Bioinf.* **21**, 472. (doi:10.1186/s12859-020-03808-8)
36. Ewald J, Bartl M, Kaleta C. 2017 Deciphering the regulation of metabolism with dynamic optimization: an overview of recent advances. *Biochem. Soc. Trans.* **45**, 1035–1043. (doi:10.1042/BST20170137)
37. Browning AP, Sharp JA, Mapder T, Baker CM, Burrage K, Simpson MJ. 2021 Persistence as an optimal hedging strategy. *Biophys. J.* **120**, 133–142. (doi:10.1016/j.bpj.2020.11.2260)
38. Lenhart S, Workman JT. 2007 *Optimal control applied to biological models*. London, UK: Chapman & Hall/CRC, Taylor & Francis.
39. Varga RS. 2000 *Matrix iterative analysis*, 2nd edn. Heidelberg, Germany: Springer.
40. Reid JK. 1966 A method for finding the optimum successive over-relaxation parameter. *Comput. J.* **9**, 200–204. (doi:10.1093/comjnl/9.2.200)
41. Carré BA. 1961 The determination of the optimum accelerating factor for successive over-relaxation. *Comput. J.* **4**, 73–78. (doi:10.1093/comjnl/4.1.73)
42. Kulsrud HE. 1961 A practical technique for the determination of the optimum relaxation factor of the successive over-relaxation method. *Commun. ACM* **4**, 184–187. (doi:10.1145/355578.366504)
43. Hadjidimos A. 2000 Successive overrelaxation (SOR) and related methods. *J. Comput. Appl. Math.* **123**, 177–199. (doi:10.1016/S0377-0427(00)00403-9)
44. Young DM. 1971 *Iterative solutions of large linear systems*. Orlando, FL: Academic Press, Inc.
45. Hackbush W. 1978 A numerical method for solving parabolic equations with opposite orientations. *Computing* **20**, 229–240. (doi:10.1007/BF02251947)
46. McAsey M, Mou L, Han W. 2012 Convergence of the forward-backward sweep method in optimal control. *Comput. Optim. Appl.* **43**, 207–226. (doi:10.1007/s10589-011-9454-7)
47. Li Q, Chen L, Tai CEW. 2018 Maximum principle based algorithms for Deep Learning. *J. Mach. Learn. Res.* **18**, 1–29. (doi:jmlr.org/papers/v18/17-653)
48. Liu X, Frank J. 2021 Symplectic Runge–Kutta discretization of a regularized forward–backward sweep iteration for optimal control problems. *J. Comput. Appl. Math.* **383**, 113113. (doi:10.1016/j.cam.2020.113133)
49. Logist F, Houska B, Diehl M, Van Impe JF. 2011 Robust multi-objective optimal control of uncertain (bio)chemical processes. *Chem. Eng. Sci.* **66**, 4670–4682. (doi:10.1016/j.ces.2011.06.018)
50. Sethi PS. 2019 *Optimal control theory—applications to management science and economics*, 3rd edn. Cham, Switzerland: Springer.
51. Sharp JA, Browning AP, Mapder T, Baker CM, Burrage K, Simpson MJ. 2020 Designing combination therapies using multiple optimal controls. *J. Theor. Biol.* **497**, 110277. (doi:10.1016/j.jtbi.2020.110277)
52. Crowell HL, MacLean AL, Stumpf MPH. 2016 Feedback mechanisms control coexistence in a stem cell model of acute myeloid leukaemia. *J. Theor. Biol.* **401**, 43–53. (doi:10.1016/j.jtbi.2016.04.002)
53. Press WH. 2007 *Numerical recipes: the art of scientific computing*. New York, NY: Cambridge University Press.
54. Bellen A, Vermiglio R. 1996 Some applications of continuous Runge–Kutta methods. *App. Numer. Math.* **22**, 63–80. (doi:10.1016/S0168-9274(96)00026-8)
55. Zennaro M. 1986 Natural continuous extensions of Runge–Kutta methods. *Math. Comput.* **46**, 119–133. (doi:10.2307/2008218)
56. Kapitula T. 2015 *Ordinary differential equations and linear algebra: a systems approach*. Philadelphia, PA: Society for Industrial and Applied Mathematics.
57. Bryson AE, Ho YC. 1975 *Applied optimal control: optimization, estimation, and control*. Abingdon, UK: Taylor & Francis.
58. Ishikawa F *et al.* 2007 Chemotherapy-resistant human AML stem cells home to and engraft within the bone-marrow endosteal region. *Nat. Biotechnol.* **25**, 1315–1321. (doi:10.1038/nbt1350)
59. Sipkins DA, Wei X, Wu JW, Runnels JM, Côté D, Means TK, Luster AD, Scadden DT, Lin CP. 2005 In vivo imaging of specialized bone marrow endothelial microdomains for tumour engraftment. *Nature* **435**, 969–973. (doi:10.1038/nature03703)
60. MacLean AL, Celso CL, Stumpf MPH. 2013 Population dynamics of normal and leukaemia stem cells in the haematopoietic stem cell niche show distinct regimes where leukaemia will be controlled. *J. R. Soc. Interface* **10**, 20120968. (doi:10.1098/rsif.2012.0968)
61. Brezinski C, Redivo-Zaglia M. 2019 The genesis and early developments of Aitken’s process, Shanks’ transformation, the ϵ -algorithm, and related fixed point methods. *Numer. Algorithms* **80**, 11–133. (doi:10.1007/s11075-018-0567-2)
62. Ramière I, Helfer T. 2015 Iterative residual-based vector methods to accelerate fixed point iterations. *Comput. Math. Appl.* **70**, 2210–2226. (doi:10.1016/j.camwa.2015.08.025)
63. Burden RL, Faires DJ. 1985 *Numerical analysis*. Boston, MA: PWS Publishing Company.
64. Henrici P. 1964 *Elements of numerical analysis*. New York, NY: John Wiley & Sons.
65. Broyden CG. 1965 A class of methods for solving nonlinear simultaneous equations. *Math. Comput.* **19**, 577–593. (doi:10.1090/S0025-5718-1965-0198670-6)
66. Wegstein JH. 1958 Accelerating convergence of iterative processes. *Commun. ACM* **1**, 9–13. (doi:10.1145/368861.368871)
67. Gutzler CH. 1959 *An iterative method of Wegstein for solving simultaneous nonlinear equations*. Masters thesis, Oregon State College.
68. Manzoor H, Selam MA, Rahman FBA, Adham S, Castier M, Abdel-Wahab A. 2020 A tool for assessing the scalability of pressure-retarded osmosis (PRO) membranes. *Renew. Energy* **149**, 987–999. (doi:10.1016/j.renene.2019.10.098)
69. Skorych V, Dosta M, Heinrich S. 2020 Dyssol—an open-source flowsheet simulation framework for

- particulate materials. *SoftwareX* **12**, 100572. (doi:10.1016/j.softx.2020.100572)
70. Orbach O, Crowe CM. 1971 Convergence promotion in the simulation of chemical processes with recycle—the dominant eigenvalue method. *Can. J. Chem. Eng.* **49**, 509–513. (doi:10.1002/cjce.5450490414)
 71. Arman A. 1986 *Acceleration algorithms for process design simulations*. Masters thesis, Oklahoma State University.
 72. Rosen EM, Pauls AC. 1977 Computer aided chemical process design: the FLOWTRAN system. *Comput. Chem. Eng.* **1**, 11–21. (doi:10.1016/0098-1354(77)80003-3)
 73. Crowe CM, Nishio M. 1975 Convergence promotion in the simulation of chemical processes—the general dominant eigenvalue method. *Am. Inst. Chem. Eng. J.* **21**, 528–533. (doi:10.1002/cjce.5450490414)
 74. Aitken AC. 1927 On Bernoulli's numerical solution of algebraic equations. *Proc. R. Soc. Edinb.* **46**, 289–305. (doi:10.1017/S0370164600022070)
 75. Steffensen JF. 1933 Remarks on iteration. *Scand. Actuar. J.* **1**, 64–72. (doi:10.1080/03461238.1933.10419209)
 76. Acklele AS, Kearfott RB, Allen EJ, Padmanabhan S. 2010 *Classical and modern numerical analysis: theory, methods and practice*. London, UK: Chapman & Hall/CRC, Taylor & Francis.
 77. Nievergelt Y. 1991 Aitken's and Steffensen's accelerations in several variables. *Numerische Mathematik* **59**, 295–310. (doi:10.1007/BF01385782)
 78. Nievergelt Y. 1995 The condition of Steffensen's acceleration in several variables. *J. Comput. Appl. Math.* **58**, 291–305. (doi:10.1016/0377-0427(94)00004-K)
 79. Noda T. 1984 The Steffensen iteration method for systems of nonlinear equations. *Proc. Jpn Acad.* **60**, 18–21. (doi:10.3792/pjaa.60.18)
 80. Noda T. 1986 The Aitken-Steffensen formula for systems of nonlinear equations. III. *Proc. Jpn Acad.* **62**, 174–177. (doi:10.3792/PJAA.62.174)
 81. Noda T. 1987 The Steffensen iteration method for systems of nonlinear equations. II. *Proc. Jpn Acad.* **63**, 186–189. (doi:10.3792/pjaa.63.186)
 82. Penrose R. 1955 A generalized inverse for matrices. *Math. Proc. Camb. Phil. Soc.* **51**, 406–413. (doi:10.1017/S0305004100030401)
 83. Anderson DG. 1965 Iterative procedures for nonlinear integral equations. *J. ACM* **12**, 547–560. (doi:10.1145/321296.321305)
 84. Anderson DG. 2019 Comments on 'Anderson Acceleration, Mixing and Extrapolation'. *Numer. Algorithms* **80**, 135–235. (doi:10.1007/s11075-018-0549-4)
 85. Eyert V. 1996 A comparative study on methods for convergence acceleration of iterative vector sequences. *J. Comput. Phys.* **124**, 271–285. (doi:10.1006/jcph.1996.0059)
 86. Fang H, Saad Y. 2008 Two classes of multisecond methods for nonlinear acceleration. *Numer. Linear Algebra Appl.* **16**, 197–221. (doi:10.1002/nla.617)
 87. Walker HF, Ni P. 2011 Anderson acceleration for fixed-point iterations. *SIAM J. Numer. Anal.* **49**, 1715–1735. (doi:10.1137/10078356X)
 88. Benner P, Findeisen R, Flockerzi D, Reichl U, Sundmacher K. 2014 *Large-scale networks in engineering and life sciences*. Basel, Switzerland: Birkhäuser Verlag.
 89. Domijan M, Rand DA. 2015 Using constraints and their value for optimization of large ODE systems. *J. R. Soc. Interface* **12**, 20141303. (doi:10.1098/rsif.2014.1303)
 90. Snowden TJ, van der Graaf PH, Tindall MJ. 2017 Methods of model reduction for large-scale biological systems: a survey of current methods and trends. *Bull. Math. Biol.* **79**, 1449–1486. (doi:10.1007/s11538-017-0277-2)
 91. Huang CYF, Ferrell JE. 1996 Ultrasensitivity in the mitogen-activated protein kinase cascade. *Proc. Natl Acad. Sci. USA* **93**, 10 078–10 083. (doi:10.1073/pnas.93.19.10078)
 92. Mirams GR, Byrne HM, King JR. 2010 A multiple timescale analysis of a mathematical model of the Wnt/ β -catenin signalling pathway. *J. Math. Biol.* **60**, 131–160. (doi:10.1007/s00285-009-0262-y)
 93. Atlas JC, Nikolaev EV, Browning ST, Shuler ML. 2008 Incorporating genome-wide DNA sequence information into a dynamic whole-cell model of *Escherichia coli*: application to DNA replication. *IET Syst. Biol.* **2**, 369–382. (doi:10.1049/iet-syb:20070079)
 94. Karr JR, Sanghvi JC, Macklin DN, Gutschow MV, Jacobs JM, Bolival B, Assad-Garcia N, Glass JL, Covert MW. 2012 A whole-cell computational model predicts phenotype from genotype. *Cell* **150**, 389–401. (doi:10.1016/j.cell.2012.05.044)
 95. Lee JM, Gianchandani EP, Eddy JA, Papin JA. 2008 Dynamic analysis of integrated signaling, metabolic and regulatory networks. *PLoS Comput. Biol.* **4**, e1000086. (doi:10.1371/journal.pcbi.1000086)
 96. Singh A, Jayaraman A, Hahn J. 2006 Modeling regulatory mechanisms in IL-6 signal transduction in hepatocytes. *Biotechnol. Bioeng.* **95**, 850–862. (doi:10.1002/bit.21026)
 97. Doyle FJ, Stelling J. 2006 Systems interface biology. *J. R. Soc. Interface* **3**, 603–616. (doi:10.1098/rsif.2006.0143)
 98. Gutenkunst RN, Waterfall JJ, Casey FP, Brown KS, Myers CR, Sethna JP. 2007 Universally sloppy parameter sensitivities in systems biology models. *PLoS Comput. Biol.* **3**, e189. (doi:10.1371/journal.pcbi.0030189)
 99. Bokil VA, Allen LJS, Jeger MJ, Lenhart S. 2019 Optimal control of a vectored plant disease model for a crop with continuous replanting. *J. Biol. Dyn.* **13**, 325–353. (doi:10.1080/17513758.2019.1622808)
 100. Hamelin FM, Bowen B, Bernhard P, Bokil VA. 2021 Optimal control of plant disease epidemics with clean seed usage. *Bull. Math. Biol.* **83**, 46. (doi:10.1007/s11538-021-00872-w)
 101. Mwanga GG, Haario H, Capasso V. 2015 Optimal control problems of epidemic systems with parameter uncertainties: application to a malaria two-age-classes transmission model with asymptomatic carriers. *Math. Biosci.* **261**, 1–12. (doi:10.1016/j.mbs.2014.11.005)
 102. Rawson T, Wilkins KE, Bonsall MB. 2020 Optimal control approaches for combining medicines and mosquito control in tackling dengue. *R. Soc. Open Sci.* **7**, 181843. (doi:10.1098/rsos.181843)
 103. Alvarez-Vázquez LJ, García-Chan N, Martínez A, Vázquez-Méndez ME. 2010 Multi-objective Pareto-optimal control: an application to wastewater management. *Comput. Optim. Appl.* **46**, 135–157. (doi:10.1007/s10589-008-9190-9)
 104. Keane AJ, Scanlan JP. 2007 Design search and optimization in aerospace engineering. *Phil. Trans. R. Soc. A* **365**, 2501–2529. (doi:10.1098/rsta.2007.2019)
 105. Peitz S, Dellnitz M. 2018 A survey of recent trends in multiobjective optimal control—surrogate models, feedback control and objective reduction. *Math. Comput. Appl.* **23**, 30. (doi:10.3390/mca23020030)