Applying optimal control theory to complex epidemiological models to inform real-world disease management

E.H. Bussell, C.E. Dangerfield, C.A. Gilligan & N.J. Cunniffe

Affiliations: Department of Plant Sciences, University of Cambridge, Cambridge CB2 3EA, United Kingdom

Keywords: optimal control, feedback, model predictive control, disease management

Corresponding Author: Dr. Nik Cunniffe, njc1001@cam.ac.uk

Main Text

Summary

Mathematical models provide a rational basis to inform how, where and when to control disease. Assuming an accurate spatially-explicit simulation model can be fitted to spread data, it is straightforward to use it to test the performance of a range of management strategies. However, the typical complexity of simulation models and the vast set of possible controls mean that only a small subset of all possible strategies can ever be tested. An alternative approach – optimal control theory – allows the best control to be identified unambiguously.

However, the complexity of the underpinning mathematics means that disease models used to identify this optimum must be very simple. We highlight two frameworks for bridging the gap between detailed epidemic simulations and optimal control theory: open-loop and model predictive control. Both these frameworks approximate a simulation model with a simpler model more amenable to mathematical analysis. Using an illustrative example model we show the benefits of using feedback control, in which the approximation and control are updated as the epidemic progresses. Our work illustrates a new methodology to allow the insights of optimal control theory to inform practical disease management strategies, with the potential for application to diseases of humans, animals and plants.

1 Introduction

Mathematical modelling plays an increasingly important role in informing policy and management decisions concerning invading diseases [1, 2]. However, model-based identification of effective and cost-efficient controls can be difficult, particularly when models include highly detailed representations of disease transmission processes. There is a variety of mathematical tools for designing optimal strategies, but no standard for putting the results from mathematically motivated simplifications into practice. An open question is how to incorporate enough realism into a model to allow accurate predictions of the impact of control measures, whilst ensuring that the truly optimal strategy can still be identified [3]. In this paper we identify the difficulties – as well as potential solutions – in achieving a practically useful optimal strategy, highlighting the potential roles of open loop and model predictive control by way of a simple example.

Realistic simulation models

The optimisation of disease management involves determining the most appropriate control method(s), e.g. vaccination, quarantine or roguing, and the best deployment strategy for that method or combination of methods to minimise impacts of the disease. This minimisation can be difficult when resources are limited and there are economic costs associated with both control measures and disease. Methods that simulate the expected course of an epidemic and explicitly model effects of interventions can rapidly quantify the potential impact of a given strategy [4]. These simulation models accurately capture the dynamics of the real system and so have become important tools for assessing policy decisions relating to real-time management responses as well as to increased preparedness for future threats. Examples include vaccination policies for human papillomavirus in the UK [5, 6], livestock culling policies [7, 8] and vaccination optimisation [9, 10] for foot-and-mouth disease, and optimal host removal strategies for tree diseases of citrus [11–14] and sudden oak death [15].

Various complexities of disease dynamics, for example spatial heterogeneities and inherent individual differences in susceptibility and pathogen transmission (risk structure), have been shown to be important determinants of patterns and rates of epidemic spread [16–18]. To ensure accurate epidemic predictions, these factors must be included in simulation models designed to aid decision making. However, inclusion of these heterogeneities typically results in highly complex models with many possible control measures, making optimisation computationally infeasible when interventions can be combined, and particularly when control measures can also vary over time, in space or according to disease risk [19]. For most simulation models the only viable option is then to use the model to evaluate a small subset of plausible strategies that remain fixed

during the epidemic, potentially scanning over a single parameter such as a culling radius. We shall refer to this approach as 'Strategy Testing'. Using this approach makes it difficult to have high confidence in the best-performing strategy, since with no framework for choosing it, the set of strategies under test is likely to be biased. Further to this, as the set to test cannot span the entire space of control options, it is unlikely that the true optimum will be found.

Optimal control of epidemiological models

Many mathematical techniques exist for characterising the true optimal control for a disease, such as equilibrium or final size analysis, depending on the system being analysed [16]. We here focus on optimising time-varying control of dynamical systems, for which optimal control theory (OCT) is widely used [20]. By analysing a set of equations describing the disease dynamics, OCT can mathematically characterise the optimal deployment strategy for a given control method and provide insight into the underlying dynamics, without the repeated simulation required to optimise simulation models. However, because of the underlying mathematical complexity, little progress can be made with OCT unless the underpinning models for disease spread are highly simplified. Early work in OCT focussed on optimal levels of vaccination and treatment [21], with extensions to consider further interventions including quarantine, screening, and health-promotion campaigns appearing later [22]. Disease models can also be coupled with economic effects [23–25], and within OCT this has been used to balance multiple costs, such as surveillance and control [26], or prophylactic versus reactive treatment [27].

The optimal strategies identified by OCT can be very complex, often specifying controls that switch strategies at specific times during the course of an epidemic. The added complexity of these switching controls can

significantly improve disease management when tested on a spatially explicit model, but can lead to poor performance if the exact time of the switch is not known [28], for example when parameter uncertainty gives a wide range of possible switch times. This demonstrates that uncertainties and additional complexities often prohibit OCT from being directly applicable to the real world. It is also unclear how insight from OCT alone could be translated into practical advice. To move towards robust strategies that could be used practically, more recent work has focussed on including additional features and heterogeneities into the models used in OCT, in particular spatial dynamics. Space is usually only included to a limited extent, for example by using metapopulation models (e.g. [29, 30]), or partial differential equations (e.g. [31]) to optimise spatial strategies, so whether the heterogeneities added are sufficient to identify robust and practical control strategies remains an open question.

Moving towards practical control

Despite finding the mathematically optimal control strategy, major simplifications to the system as modelled are required to allow progress to be made using OCT. It is therefore often unclear how these strategies would perform if adopted by policy makers. On the other hand, models with sufficient realism to inform policy directly are often impossible to optimise fully. Therefore, a framework is needed to combine the optimisation capabilities of OCT with the accurate predictions of simulation type models as required in policy making. The question is then how should we make practical use of OCT?

In §2 we describe two methods from control systems engineering for applying OCT results, and compare these versus Strategy Testing using a simple illustrative model in §3. We seek to answer how, under current computational constraints, results from OCT can be applied whilst maintaining the realism required for practical

application.

2 Applying optimal control to realistic systems

Outside of epidemiology, OCT has had wider use on approximate models of complex systems. A recent study reviews the use of OCT for agent-based models (ABMs) [32], a type of model that simulates the individual behaviour of autonomous agents. An *et al*. [32] suggest the use of a model that approximates the dynamics of the ABM, designed to be simple enough to allow mathematical analysis of the optimal control. A suitable approximate model is chosen and fitted either to real data, or to synthetic data from the ABM. The OCT results from the approximating model are then mapped onto the ABM to be tested: a process referred to as 'lifting', which could equally well apply to the detailed epidemic simulation models considered in this paper. We now describe two possible frameworks from control systems engineering for making use of this control lifting approach.

Open-loop control

The first method is the simplest application of control lifting, and the framework implicitly suggested by An *et al*. [32]. Control is optimised on the approximate model once using the initial conditions of the simulation model. The resulting optimal control strategy is lifted to the simulator and applied for the full simulation run time (figure 1). Repeated simulation of the OCT strategy on the simulation model allows assessment against other possible control strategies. The optimisation gives a single, time dependent strategy for all simulation realisations, and so does not incorporate any feedback. It is therefore referred to as 'open-loop' control, as it is fully specified by the simulation initial conditions and the trajectory predicted by the approximate model. Use

in epidemiology is uncommon, although Clarke *et al*. [33] use OCT in an approximate model to find optimal levels of Chlamydia screening and contact tracing which are then mapped onto a network simulation.

Model predictive control

Open-loop control requires the approximate model to remain accurate over the time scale of the entire epidemic. However, for tractability the approximate model must necessarily omit many heterogeneities present in the simulation model, such as spatial effects and risk structure. When strategies resulting from OCT are then applied to the simulation model or to the real system, the disease progress is likely to deviate systematically from the trajectory predicted by the approximate model. Model predictive control (MPC) is an optimisation technique incorporating system feedback that can take such perturbations into account [34, 35]. At regular update times the values of the state variables in the approximate model are reset to match those in the simulation at that time. The control is then re-optimised and the new control strategy is applied to the simulation until the next update time. The approximate and simulation models are therefore run concurrently, with multiple optimisations per realisation, to ensure that the approximate model and control strategy closely match each individual simulation realisation (figure 1). These multiple optimisations are computationally costly but tractable, unlike performing optimisation on the full simulation model.

MPC has had some use within the epidemiological literature, the majority being for control of drug applications for single individuals rather than control of epidemics at the population level. Examples include finding management strategies for HIV that are robust to measurement noise and modelling errors [36, 37], and control of insulin delivery in patients with diabetes [38]. These studies highlight the benefits of MPC for robust control, i.e. control that remains effective despite system perturbations. However, only one study concentrates

on epidemic management [39], and that does not explicitly test the feedback control on simulations.

3 Optimising strategies on an illustrative network model

Methods

To demonstrate open-loop and MPC for epidemic management we use a stochastic SIR network model including host demography and risk structure. The model is deliberately kept simple to show how the underpinning idea is broadly applicable across human, animal and plant diseases. Whilst the model and its parameters are arbitrary and do not represent a specific disease, we use it to represent a scenario in which a simulation model has already been fitted to a real disease system; the network model is therefore used here as a proxy for a potentially very detailed simulation model.

Simulation Model

In our model, infection spreads stochastically across a network of nodes that are clustered into three distinct regions (figure 2a). Each node contains a host population stratified into high and low risk groups. The infection can spread between individuals within nodes and between connected nodes. The net rate of infection of risk group r in node i is given by:

$$S_i^r \sum_j \sigma_{ij} \left(\rho^{rH} I_j^H + \rho^{rL} I_j^L \right) , \tag{1}$$

where S and I are numbers of susceptible and infected hosts respectively, subscripts identify the node, and superscripts specify high (H) or low (L) risk group. The sum is over all connected nodes including the focal node itself, with the relative transmission strength into node i from node j given by σ_{ij}, and risk structure given

by the 2×2 matrix ρ. Full details of the model are given in the supplementary material. Although not limited to these applications, the model in Equation 1 could represent crop or livestock diseases spreading through farms, or sexually transmitted infections spreading through towns, cities or countries.

Mass vaccination is the only intervention we consider, with the potential to target based on both risk group and region but randomised across host infection status (i.e. the vaccine is given to all hosts but is only effective on susceptibles). Logistical and economic constraints are included through a maximum total vaccination rate ($\eta_{\max}$) that can be divided between risk groups and regions. Within each group susceptibles are vaccinated at rate: $f\eta_{\max}S/N$, where f is the proportion of control allocated to that group, and N is the total group population.

Optimal allocation of the vaccination resources minimises an epidemic cost J representing the disease burden of the epidemic across all infected hosts over the simulation time (T): $J = \int_{t=0}^{T} I(t)\mathrm{d}t$. In common with the particular control we consider and the risk and spatial structures, this simple choice of objective function was made merely to illustrate our methods, but the framework generalises immediately to more complex settings.

Approximate Models

Exhaustive optimisation of control using the simulation model, across space, risk group and time, is clearly very computationally expensive. To assess the best level of approximation, we consider two different deterministic approximate models of the simulator. The first model is purely risk structured, factoring out all spatial information and leaving one high risk and one low risk population group. This model is deterministic and based on the assumption that all nodes are spatially well-mixed with each other. The second approximate

model is more complex, in as much as it is also deterministic and risk structured, but additionally includes a first approximation to the host spatial structure by including the regional host information. Spatial dynamics are included between but not within the three regions to maintain enough simplicity to obtain optimal control results, thereby assuming that nodes are spatially well-mixed within each region. This could represent, for example, optimising control at the country level, but not at the regional level. We refer to this model as the spatial approximate model. A single set of parameters is fitted for each model to data from an ensemble of simulation model runs. We then test which of the two approximate models is the more useful for control optimisation. Full details of the approximate models, fitting and optimisation procedures are given in the supplementary material.

Control Scenarios

We test six different control scenarios, which compare Strategy Testing of controls based purely on the simulation model (scenarios 1 and 2) with open-loop and MPC applied using both of our approximate models (scenarios 3 to 6):

1. 'High': exclusively vaccinate high risk individuals
2. 'Split': partition control resources between high and low risk groups based on an optimisation performed in advance
3. 'Risk OL': open-loop control using the risk based approximate model
4. 'Risk MPC': MPC using the risk based approximate model
5. 'Space OL': open-loop control using the spatial approximate model

6. 'Space MPC': MPC using the spatial approximate model

The optimal constant allocation for the 'Split' strategy was found by running many simulation model realisations for each of a range of partition values, as in [11], and selecting the value that gave the lowest average epidemic cost (supplementary figure S8). The six strategies are assessed by repeatedly running the simulation model under each control scenario.

Results

The OCT results for optimising the vaccination strategy in the risk based approximate model lead to initial vaccination of high risk individuals only, before switching priorities and treating the more populous low risk group almost exclusively. The OCT results from the spatial approximate model show this same switch (figure 2b), but a number of spatial switches are also seen, allowing control to track the epidemic as it progresses through the three regions (supplementary figure S9). The spatial strategies are therefore much more complex than the risk based controls.

Applying the control scenarios to the simulation model and comparing epidemic costs shows that incorporating greater realism, through a more complex approximate model as well as by using MPC, allows for improved disease management (figure 3 and supplementary figure S10). Of the constant and purely simulation based 'user-defined' strategies, splitting control between risk groups is slightly more effective than just vaccinating the high risk group. The optimal allocation to the high risk group used in the 'Split' strategy is 63% of vaccination resources, with the rest used to vaccinate low risk individuals, although this does occur in a broad minimum of epidemic cost (supplementary figure S8). Applying the optimisations from the risk based approximate model to the simulation model gives an improvement over either of the 'user-defined' strategies,

although there is little difference in epidemic cost between the open-loop and MPC frameworks (see below). Adding space into the approximate model improves control further, leading to the smallest epidemic costs when the spatial MPC framework is used.

The illustrative model demonstrates the management improvements that can be achieved by combining OCT with both open-loop and MPC. The key results of the OCT analyses are the control switching times. Using the switching controls from either approximate model with open-loop control gives lower epidemic costs than the naively chosen 'user-defined' strategies. The feedback present in the MPC controllers allows further reductions to the epidemic cost. By re-evaluating the timing of the switches during the epidemic, and potentially including additional switches, the control can respond more closely to the exact trajectory of the current simulation realisation (figures 2b–d). This gives control that is more robust to uncertainty and systematic errors in the approximate model, and hence performs better on the complex simulation model.

In the risk based strategies there is little difference between open-loop and MPC. This is because the precise timing of the switch from high to low risk group vaccination does not significantly affect the epidemic cost (supplementary figure S11). The timings of disease introduction into regions B and C are highly variable between simulation runs (supplementary figure S2). The potential for additional switches in the spatial approximate model gives more flexibility for the MPC controller to respond to this variability, and so spatial MPC shows a significant improvement over open-loop which cannot adapt to perturbations. The performance of the control is closely linked to the accuracy of the approximate model. In our example, spatial dynamics are clearly important because of the timing of spread between regions, and so the more informed controls of the spatial model outperform the risk based strategies.

4 Discussion

Our results show that the choice of approximate model affects the performance of both open-loop and MPC strategies. Here we have found a suitable approximate model in an ad hoc manner, but a key challenge for the future is to develop a more formal method for choosing the most appropriate approximate model. A more accurate model may give better predictions, and hence control that is closer to the true optimum, but simpler models are often sufficient [40] and accuracy must be balanced with added complexity and optimisation constraints. One difficulty in doing this is that it is not always clear where the boundary of mathematical or computational feasibility is, and so how complex the model can be made in practice. It is also difficult to determine mathematically, in a systematic way, which aspects of the dynamics are important to capture accurately. This key issue must be considered though, since the implications relate directly to applications in the real world.

Practical disease control requires surveys of the real system to assess the state of the epidemic. Both open-loop and MPC optimise control using predictions of the future dynamics, making them both feed-forward controllers. The approximate model underlying these frameworks allows more informed decisions between surveys, resulting in control that is closer to the true optimum. Accurate predictions can avoid continuous or very frequent surveys which may be expensive or logistically challenging. As discussed previously, the repeated updates in the feedback loop of MPC improve these predictions and hence the performance of the control. However, each update will require surveillance of the real system, so the frequency of updates must be chosen so as to balance improved knowledge of the system with any surveillance constraints.

In this paper we have focussed on a top-down approach, finding robust, practically-applicable strategies by

making use of OCT to optimise simulation models. Equally, many studies use OCT without simulation models, rarely considering practical application of the resulting optimal controls. With this bottom-up approach, a system for testing the results on realistic systems is vital to ensure that these results are robust to additional realism. Using an MPC framework as considered here could be one way in which OCT researchers could demonstrate the potential impact of their work to a wider audience.

Exhaustive testing of alternative simulation model parameterisations is beyond the scope of this study, but we generally find that spatial MPC also performs best across other reasonable parameter sets (supplementary material §3). We have assumed throughout that an accurate simulation model of the real system in question can be built, and that a single set of parameters can be fitted for the chosen deterministic approximate model. In reality there may be considerable uncertainty in parameters for the simulator so fitting a single deterministic model may be challenging. A question for future study would be how to handle these uncertainties, perhaps also incorporating improved knowledge of parameters as the simulation proceeds [41].

The strategies found by OCT are highly dependent on the exact form of the objective function, which we have here chosen to be very simple. Extending the objective to include costs associated with control as well as with each switch in strategy would allow a more detailed assessment of the practicality of implementing these complex strategies. More research is needed into how to quantify the balancing of very different costs though, for example treatment costs and disease burden [29]. In human disease, cost-effectiveness analyses are usually based on quality adjusted life years [42]. A similar concept could perhaps be used for plant and animal diseases, including calculations of yield losses [43] as well as effects on welfare, biodiversity and tourism for example [44]. The methods we have described however, are not dependent on the form of the control or objective function. For an appropriate approximate model, the feedback in MPC ensures accurate predictions and so

should always improve performance over open-loop. The frameworks we describe can be used to provide an additional, unbiased control scenario to the Strategy Testing process that is already in common use.

In this paper we have shown that coupling feedback control with simulation models and OCT can help to design effective and robust intervention strategies for managing pathogens of human, animal and plant populations. Whilst these techniques may be able to transfer optimal control results to more realistic simulations and so to practical application, it does raise the issue of communicability of results. With complex feedback strategies between two models, one complex in structure and the other mathematically complex, the overall result is no longer simple to explain. Future research must therefore focus on improving the accuracy of simulation models, and analysing their reliability, so that simulations can be used to establish conclusively the benefit of these complex OCT based strategies.

Additional Information

Acknowledgements

We thank Andrew Craig, Eleftherios Avramidis and Hola Adrakey for helpful discussions. We also thank two anonymous reviewers for their helpful and constructive comments.

Data Accessibility

All code and animations are available at `https://github.com/ehbussell/Bussell2018Model`.

Authors' Contributions

E.H.B., C.E.D. and N.J.C. designed the study, E.H.B. conducted the analysis and wrote the initial draft of the manuscript. All authors contributed to data interpretation, manuscript editing and discussion.

Competing Interests

We have no competing interests.

Funding

E.H.B. acknowledges the Biotechnology and Biological Sciences Research Council of the United Kingdom (BBSRC) for support via a University of Cambridge DTP Ph.D. studentship.

References

[1] Heesterbeek H, et al. 2015. Modeling infectious disease dynamics in the complex landscape of global health. *Science* **347**, aaa4339 doi:10.1126/science.aaa4339

[2] Metcalf CJE, Edmunds WJ, Lessler J. 2015. Six challenges in modelling for public health policy. *Epidemics* **10**, 93–96 doi:10.1016/j.epidem.2014.08.008

[3] Cunniffe NJ, Koskella B, E Metcalf CJ, Parnell S, Gottwald TR, Gilligan CA. 2015. Thirteen challenges in modelling plant diseases. *Epidemics* **10**, 6–10 doi:10.1016/j.epidem.2014.06.002

[4] Lofgren ET, et al. 2014. Opinion: mathematical models: a key tool for outbreak response. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 18095–18096 doi:10.1073/pnas.1421551111

[5] Jit M, Choi YH, Edmunds WJ. 2008. Economic evaluation of human papillomavirus vaccination in the United Kingdom. *BMJ* **337**, a769 doi:10.1136/bmj.a769

[6] Choi YH, Jit M, Gay N, Cox A, Garnett GP, Edmunds WJ. 2010. Transmission dynamic modelling of the impact of human papillomavirus vaccination in the United Kingdom. *Vaccine* **28**, 4091–4102 doi:10.1016/j.vaccine.2009.09.125

[7] Keeling MJ, Woolhouse ME, Shaw DJ, Matthews L, Chase-Topping M, Haydon DT, Cornell SJ, Kappey J, Wilesmith J, Grenfell BT. 2001. Dynamics of the 2001 UK foot and mouth epidemic: stochastic dispersal in a heterogeneous landscape. *Science* **294**, 813–817 doi:10.1126/science.1065973

[8] Keeling MJ. 2005. Models of foot-and-mouth disease. *Proc. R. Soc. Lond. B Biol. Sci.* **272**, 1195–1202 doi:10.1098/rspb.2004.3046

[9] Keeling MJ, Woolhouse MEJ, May RM, Davies G, Grenfell BT. 2003. Modelling vaccination strategies against foot-and-mouth disease. *Nature* **421**, 136–42 doi:10.1038/nature01343

[10] Tildesley MJ, Savill NJ, Shaw DJ, Deardon R, Brooks SP, Woolhouse ME, Grenfell BT, Keeling MJ. 2006. Optimal reactive vaccination strategies for a foot-and-mouth outbreak in the UK. *Nature* **440**, 83–86 doi:10.1038/nature04324

[11] Cunniffe NJ, Stutt RO, DeSimone RE, Gottwald TR, Gilligan CA. 2015. Optimising and communicating options for the control of invasive plant disease when there is epidemiological uncertainty. *PLoS Comput. Biol* **11**, e1004211 doi:10.1371/journal.pcbi.1004211

[12] Hyatt-Twynam SR, Parnell S, Stutt RO, Gottwald TR, Gilligan CA, Cunniffe NJ. 2017. Risk-based management of invading plant disease. *New Phytol.* **214**, 1317–1329 doi:10.1111/nph.14488

[13] Adrakey HK, Streftaris G, Cunniffe NJ, Gottwald TR, Gilligan CA, Gibson GJ. 2017. Evidence-based controls for epidemics using spatio-temporal stochastic models in a Bayesian framework. *J. R. Soc. Interface* **14**, 20170386 doi:10.1098/rsif.2017.0386

[14] Craig AP, Cunniffe NJ, Parry M, Laranjeira FF, Gilligan CA. 2018. Grower and regulator conflict in management of the citrus disease Huanglongbing in Brazil: a modelling study. *J. Appl. Ecol.* **55**, 1956–1965 doi:10.1111/1365-2664.13122

[15] Cunniffe NJ, Cobb RC, Meentemeyer RK, Rizzo DM, Gilligan CA. 2016. Modeling when, where, and how to manage a forest epidemic, motivated by sudden oak death in California. *Proc. Natl. Acad. Sci. U.S.A.* **113**, 5640–5645 doi:10.1073/pnas.1602153113

[16] Keeling MJ, Rohani P, 2008. *Modeling Infectious Diseases in Humans and Animals*. Princeton University Press

[17] Anderson RM, Medley GF, May RM, Johnson AM. 1986. A preliminary study of the transmission dynamics of the human immunodeficiency virus (HIV), the causative agent of AIDS. *Math. Med. Biol.* **3**, 229–263 doi:10.1093/imammb/3.4.229

[18] Smith DL, Lucey B, Waller LA, Childs JE, Real LA. 2002. Predicting the spatial dynamics of rabies epidemics on heterogeneous landscapes. *Proc. Natl. Acad. Sci. U.S.A.* **99**, 3668–3672 doi:10.1073/pnas.042400799

[19] Bellman R, 2013. *Dynamic Programming*. Courier Corporation

[20] Lenhart S, Workman JT, 2007. *Optimal Control Applied to Biological Models*. CRC Press

[21] Sethi SP, Staats PW. 1978. Optimal control of some simple deterministic epidemic models. *J. Oper. Res. Soc.* **29**, 129–136 doi:10.1057/jors.1978.27

[22] Behncke H. 2000. Optimal control of deterministic epidemics. *Optim. Control Appl. Methods* **21**, 269–285 doi:10.1002/oca.678

[23] Perrings C, et al. 2014. Merging economics and epidemiology to improve the prediction and management of infectious disease. *EcoHealth* **11**, 464–475 doi:10.1007/s10393-014-0963-6

[24] Epanchin-Niell RS. 2017. Economics of invasive species policy and management. *Biol. Invasions*, 1–22 doi:10.1007/s10530-017-1406-4

[25] Cunniffe NJ, Laranjeira FF, Neri FM, DeSimone RE, Gilligan CA. 2014. Cost-effective control of plant disease when epidemiological knowledge is incomplete: modelling Bahia bark scaling of citrus. *PLOS Comput. Biol.* **10**, 1–14 doi:10.1371/journal.pcbi.1003753

[26] Ndeffo Mbah ML, Gilligan CA. 2010. Balancing detection and eradication for control of epidemics: sudden oak death in mixed-species stands. *PLoS ONE* **5**, 1–8 doi:10.1371/journal.pone.0012317

[27] Brown VL, White KAJ. 2011. The role of optimal control in assessing the most cost-effective implementation of a vaccination programme: HPV as a case study. *Math. Biosci.* **231**, 126–134 doi:10.1016/j.mbs.2011.02.009

[28] Forster G, Gilligan CA. 2007. Optimizing the control of disease infestations at the landscape scale. *Proc. Natl. Acad. Sci. U.S.A.* **104**, 4984–4989 doi:10.1073/pnas.0607900104

[29] Rowthorn RE, Laxminarayan R, Gilligan CA. 2009. Optimal control of epidemics in metapopulations. *J. R. Soc. Interface* **6**, 1135–1144 doi:10.1098/rsif.2008.0402

[30] Ndeffo Mbah ML, Gilligan CA. 2011. Resource allocation for epidemic control in metapopulations. *PLoS ONE* **6**, doi:10.1371/journal.pone.0024577

[31] Neilan RM, Lenhart S. 2011. Optimal vaccine distribution in a spatiotemporal epidemic model with an application to rabies and raccoons. *J. Math. Anal. Appl.* **378**, 603–619 doi:10.1016/j.jmaa.2010.12.035

[32] An G, Fitzpatrick BG, Christley S, Federico P, Kanarek A, Neilan RM, Oremland M, Salinas R, Laubenbacher R, Lenhart S. 2017. Optimization and control of agent-based models in biology: a perspective. *Bull Math Biol* **79**, 63–87 doi:10.1007/s11538-016-0225-6

[33] Clarke J, White KAJ, Turner K. 2013. Approximating optimal controls for networks when there are combinations of population-level and targeted measures available: chlamydia infection as a case-study. *Bull. Math. Biol.* **75**, 1747–1777 doi:10.1007/s11538-013-9867-9

[34] Camacho EF, Bordons C, 2012. *Model Predictive Control in the Process Industry*. Springer Science & Business Media

[35] Lee JH. 2011. Model predictive control: review of the three decades of development. *Int. J. Control Autom. Syst.* **9**, 415–424 doi:10.1007/s12555-011-0300-6

[36] Zurakowski R, Teel AR. 2006. A model predictive control based scheduling method for HIV therapy. *J. Theor. Biol.* **238**, 368–382 doi:10.1016/j.jtbi.2005.05.004

[37] David J, Tran H, Banks HT. 2011. Receding horizon control of HIV. *Optim. Control Appl. Meth.* **32**, 681–699 doi:10.1002/oca.969

[38] Hovorka R, et al. 2004. Nonlinear model predictive control of glucose concentration in subjects with type 1 diabetes. *Physiol. Meas.* **25**, 905–920 doi:10.1088/0967-3334/25/4/010

[39] Sélley F, Besenyei Á, Kiss IZ, Simon PL. 2015. Dynamic control of modern, network-based epidemic models. *SIAM J. Appl. Dyn. Syst.* **14**, 168–187 doi:10.1137/130947039

[40] Thompson RN, Hart WS. 2018. Effect of confusing symptoms and infectiousness on forecasting and control of ebola outbreaks. *Clin. Infect. Dis.*, ciy248 doi:10.1093/cid/ciy248

[41] Thompson RN, Gilligan CA, Cunniffe NJ. 2018. Control fast or control smart: when should invading pathogens be controlled? *PLoS Comput. Biol.* **14**, e1006014 doi:10.1371/journal.pcbi.1006014

[42] Whitehead SJ, Ali S. 2010. Health outcomes in economic evaluation: the QALY and utilities. *Br. Med. Bull.* **96**, 5–21 doi:10.1093/bmb/ldq033

[43] Savary S, Ficke A, Aubertot JN, Hollier C. 2012. Crop losses due to diseases and their implications for global food production losses and food security. *Food Secur.* **4**, 519–537 doi:10.1007/s12571-012-0200-5

[44] Boyd IL, Freer-Smith PH, Gilligan CA, Godfray HCJ. 2013. The consequence of tree pests and diseases for ecosystem services. *Science* **342**, 1235773 doi:10.1126/science.1235773

Figure and table captions

Figure 1

Open-loop and model predictive control (MPC). The model hierarchy is shown, with optimised controls from the approximate model directly lifted to the simulation model. The real system is in green, the models and fitting processes are in blue, and the control framework is in orange. Without the orange dashed feedback loop, this is open-loop control. MPC resets the state of the approximate model at regular update steps, before re-optimising and lifting controls to the simulation model until the next update time.

Figure 2

(a) shows the network used for the illustrative simulation model, including region labels. The epidemic is seeded in the red node in region A, and can spread between connected nodes (grey lines). In **(b)** the control

allocation is shown for a single space based MPC run, with the corresponding open-loop allocation indicated by the black dotted line. **(c)** shows the total number of infected individuals under a single run of space based open-loop control. Control is based on the prediction of the approximate model starting from the initial conditions. **(d)** shows the number of infected individuals in the simulation and space based approximate model corresponding to the MPC control carried out in **(b)**. Here the prediction is reset to match the simulation at every update step (0.5 time units) and the control is re-optimised. By repeatedly correcting for differences between short-term model predictions and realised numbers of infected individuals – rather than relying on a potentially increasingly inaccurate prediction made at the initial time – MPC gives better predictions of the simulation state as well as improved control when compared to open-loop (note different y axis scales).

Figure 3

Results of different control optimisation schemes on the illustrative simulation model. Spatial MPC performs best, showing an improvement over both open-loop and user-defined strategies.

Figures

Figure 1

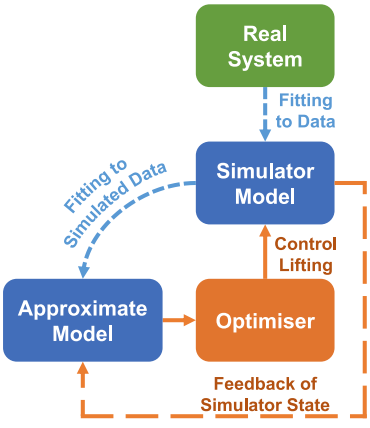

Figure 2

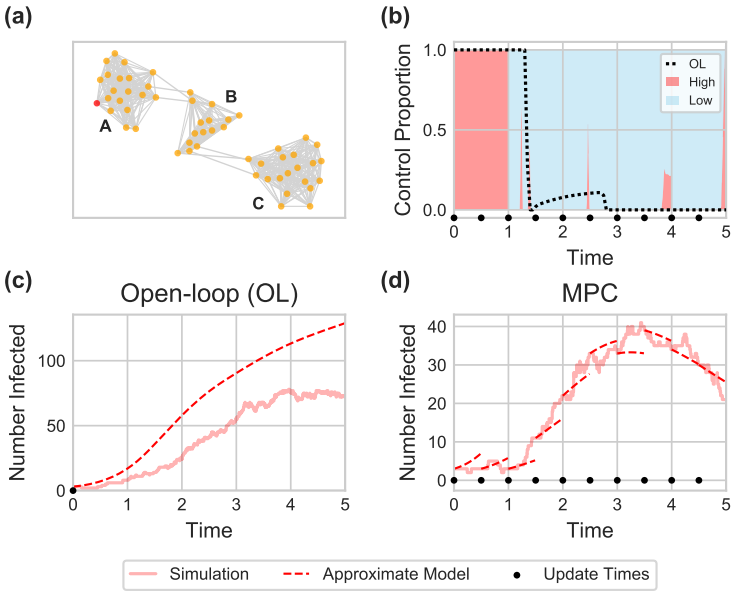

Figure 3