When to monitor or control; informed invasive species management using a partially observable Markov decision process (POMDP) framework

Thomas K. Waring[1,2], **Vera L. J. Somers**[1,3], **Michael A. McCarthy**[4], **Christopher M. Baker**[1,2*]

[1]School of Mathematics and Statistics, University of Melbourne, VIC 3010, Australia

[2]Melbourne Centre for Data Science (MCDS), VIC 3010, Australia

[3]ARC Training Centre in Optimisation Technologies, Integrated Methodologies, and Applications (OPTIMA), VIC 3010, Australia

[4]School of Agriculture, Food and Ecosystem Sciences, University of Melbourne, VIC 3010, Australia

---

Abstract

1: For invasive species management, the trade-off between monitoring and control has high stakes, and existing decision-making methods are limited in this area. In particular, most current approaches are limited to a specific case study, or ignore the uncertainty and interaction between monitoring and control in the system. Hence, the field is missing an effective and general framework.

2: We propose a partially observable Markov decision process (POMDP) framework to help decision makers understand effective monitoring and control policy making. POMDPs can deal with uncertainty in both the model and state of the system, but are more challenging to solve due to the continuous and high-dimensional state space. We reduce the dimensionality of the state space by restricting the belief state to a parametrised family of probability distributions, aligning the mathematical representation of the problem with quantities of the physical process linked to human decision making.

3: The result of our model is a sequence of actions which minimises the expected cost incurred in managing the invasive species, where the recommendation depends on an estimate of the species' abundance, and the uncertainty in this estimate. We demonstrate the effectiveness of our proposed framework in two generic case studies of varying complexity. Furthermore, we investigate sensitivity of the results to the choices of control cost and efficacy, and monitoring cost and error.

4: The framework proposed by this paper makes the powerful machinery of POMDPs available to environmental managers. It computes the optimal course of action to manage a growing population of an invasive species, incorporating a varying time horizon and multiple control interventions. We sidestep the computational difficulties of general POMDPs to provide a clear, visual overview of decision-making recommendations, and how these decisions change in new situations. Initial results and scenario based analysis show promising results, and the framework could be extended to the related field of disease management.

Keywords: *invasive species, decision making, partially observable Markov decision processes, uncertainty*

1. Introduction

An uncontrolled invasive species can have a major impact on the economy and environment (Hoffmann and Broadhurst, 2016; Alvarez and Solís, 2018). However, perfect information of the severity of the scenario is rarely available, and so effective policy involves a challenging combination of both monitoring (information gathering) and control efforts. In particular, two main challenges arise with invasive species management. First, monitoring and control efforts have an influence on one another and can therefore, not easily be determined separately. And second, full knowledge of the system and the dynamics of its evolution is often unknown or unavailable (Büyüktahtakın and Haight, 2018).

In order to tackle this problem, different modeling methods and approaches have been taken: e.g., mathematical programming (Epanchin-Niell et al., 2012), analysing optimality conditions and optimal control (Rout et al., 2014; Mehta et al., 2007; Hauser and McCarthy, 2009; Mbah and Gilligan, 2010) and heuristics such as genetic algorithms (Carrasco et al., 2010).

All these modeling methods have their own strengths and weaknesses. Heuristics offer quick solutions, but which often fail to be optimal and/or robust. Optimal control methods, on the other hand, will find

*Corresponding author. *E-mail address: author4@email.com*

the optimal solution, but are more complex to solve. Furthermore, in all of these methods, models struggle to deal with the uncertainty in the system.

An alternative, promising method is to model the problem as a Markov Decision Process (MDP), which is a logical result of acknowledging the fact that the underlying spreading process is inherently stochastic. Stochastic Dynamic Programming (SDP) is often successfully applied to solve MDPs for stochastic systems, and its main advantage is that it can deal with and include uncertainty (Marescot et al., 2013; Hyytiäinen et al., 2013; Polasky, 2010; Rout et al., 2011; Moore et al., 2010). Its main shortcoming, however, lies in the fact that MDPs assume perfect knowledge of the state of the system.

Therefore, if we replace perfect knowledge of the state of the system by observations that are linked probabilistically to the underlying state, we obtain a partially observable Markov decision process (POMDP) (Monahan, 1982; Littman, 2009). POMPDPs are, however, far more difficult to solve than MDPs. First approaches are presented by Haight and Polasky (2010); Regan et al. (2011); Chadès et al. (2011); Rout et al. (2014), where Chadès et al. (2011) use a POMDP method to compute node priority for investment in small networks of pests, diseases and endangered species. On the other hand, Haight and Polasky (2010); Regan et al. (2011); Rout et al. (2014) tackle invasive species management using a discretized, short horizon POMPDP. To make the problem computationally tractable, these models have few management options and levels of infestation, and make major assumptions regarding eradication efficacy. All these papers offer specific recommendations based on limited scenario analysis and therefore, have different conclusions regarding the optimality of surveillance and intervention.

Furthermore, while literature on each of surveillance and control is well advanced, research into the combination of them is limited (Büyüktahtakın and Haight, 2018). Studies that have looked into the relationship between monitoring and control reach varying and different conclusions regarding optimal management based on specific species studied and assumptions made, such as efficacy and cost of detection and control: see for example Rout et al. (2011, 2014); Carrasco et al. (2010); Mbah and Gilligan (2010); Moore et al. (2010); Polasky (2010); Hyytiäinen et al. (2013). That is, there is no general framework which addresses the wide variety of problems which face environmental managers making monitoring and control decisions for invasive species.

In this paper, we propose a POMDP framework which provides a general and interpretable method applicable to a wide range of circumstances and management options, and which is specifically built to deal with uncertainty. And although some parameter sensitivity is investigated in the literature, to the best of our knowledge there is no research looking into the boundary areas when the optimal decision switches from monitoring to control. Therefore, with an effective visual tool, we demonstrate where the cut-off lies to inform effective management policies.

Our approach is related to that of Moore et al. (2017), who obtained a general solution to the case where managers are choosing between two control actions, the benefits of which are uncertain. The current paper differs in that the abundance of the species in question is the uncertain variable, and the decision is between some set of control actions, whose efficacies are known, a monitoring action, or doing nothing.

2. Materials and Methods

2.1. Problem statement

We consider a control and surveillance, (i.e., monitoring) problem of invasive species management with some unknown abundance, and which is growing at a (possibly) uncertain rate. In response a decision maker or manager has some control actions available, each of which eliminates some fraction of the species, at some cost. Alternatively, the manager may opt to monitor the species, that is, to spend some resources to receive some estimate of the species' abundance.

We state this problem as a *Partially Observable Markov Decision Process* (POMDP) (Åström, 1965; Monahan, 1982). A POMDP is specified by sets of states, actions and observations, an objective function to be minimised (or maximised), and the probabilities which govern the state transitions and the observations.

The case at hand runs roughly as follows. The state at time t is the abundance N_t of the species, and we assume each unit of abundance incurs one unit of cost (that is, cost is measured in units of species abundance). The transitions are governed by a model for the growth of the species. For simplicity, we restrict ourselves to the case of exponential growth, so that

$$N_{t+1} = e^{r_t} \cdot N_t,$$

where the growth rate r_t may vary in time. To model an uncertain estimate of the value r, we draw each r_t from a normal distribution with mean r and variance Δr^2 (Keeling and Rohani, 2008, Section 6.2).

The actions available to the manager fall into two categories: *control* and surveillance, i.e., *monitoring*. Control actions, which we denote as a_i for $i = 1, 2, \ldots$, reduce the subsequent abundance by some fraction ρ_i, incurring some cost c_i. That is, if action a_i is taken at time t, the immediate cost is $N_t + c_i$, and the next-step abundance is

$$N_{t+1} = (1 - \rho_i)e^{r_t} \cdot N_t. \tag{1}$$

In this case, the manager receives no observation of the state. In the absence of any uncertainty, the decision is between doing nothing (if the abundance N_t is small) and controlling the species (if N_t is large). The threshold abundance where the decision changes is determined by the cost and efficacy of control actions, and how quickly the species' population is growing — this calculation is carried out at the beginning of section 2.3.2.

Monitoring actions, denoted A_i, allow the manager to observe the state at some cost C_i. We model this as a draw from a lognormal distribution around the true value, with some standard error e_i. In a fully general circumstance, monitoring actions might also have some control effects.

2.2. Belief Markov Decision Process

The conventional approach to solving a POMDP is to consider it as a conventional MDP, where each state is a probability distribution over the original states of the system Åström (1965). Therefore, the new 'belief state' space has one dimension for each original state, representing the probability that the system is in the given state, given the past history of actions and observations. Since higher dimensions are computationally more difficult, the belief state approch is forced to work with a small number of states (Chadès et al., 2021; Rout et al., 2014; Chadès et al., 2008), and certainly to avoid an unbounded state space.

Beyond the limits imposed by computational concerns, the results of the conventional belief state MDP can be difficult to interpret and communicate. For example, in a model which allows for low, medium and high abundance, the following unlikely belief is a point in the considered state space:

$$P(N = \text{low}) = 0.5$$
$$P(N = \text{medium}) = 0$$
$$P(N = \text{high}) = 0.5.$$

This belief is not one which corresponds to a real-world situation, and as such it complicates the interpretation of the model. In (Rout et al., 2014), for example, this problem is avoided by only giving scenarios for distributions where $P(N = \text{high}) = 0$, at the cost of a loss in generality. The interpretability problem is exacerbated by the difficulty of representing the value of a given action or state in more than three dimensions, as discussed in (Chadès et al., 2021, Section 5.2).

It is clear, then, that it is valuable from a computational point of view to reduce the dimensionality of our state space. Furthermore, these dimensions should reflect quantities that have real-world meaning, in order to better reflect the actual 'belief state' of a manager.

2.3. Proposed solution

We propose that the 'belief state' of a manager should be restricted to a parametrised family of probability distributions. Reflecting the goals stated above, we choose as parameters the abundance of the controlled species, and the uncertainty in this value. More specifically, we choose a lognormal distribution, parametrised by $n, \Delta n$ such that $\log N$ follows a normal distribution with mean n and standard deviation Δn. This choice is made for convenience: exponential growth and multiplicative control correspond to translations of the distribution. Such a simplification fulfills the desiderata of the previous section, as demonstrated in Section 3. In short, given an estimate with error of $\log N$ our model computes the optimum policy amongst options of monitoring or control.

2.3.1. Problem setup

We now re-state our decision problem with our suggested simplification. If the state is $n, \Delta n$, then the cost is, as above, the expected abundance. For a lognormal distribution, this is

$$\mathbb{E}[N] = \exp(n + (\Delta n)^2/2).$$

Following the original model of (1), growth of the invasive population translates the estimate n_t of $\log N$, so that under a control action a_i, we have $n_{t+1} = n_t + r + \log(1 - \rho_i)$. Accounting for uncertainty in r, we have $\Delta n_{t+1} = \Delta n_t + \Delta r$.

Again following our above model, under a monitoring action, we draw n_{t+1} from $\mathrm{Norm}(n + r, \Delta n)$. If the monitoring error is e, then the new error estimate is given by

$$\Delta n_{t+1} = \frac{\Delta n \cdot e}{\Delta n + e} + \Delta r. \tag{2}$$

2.3.2. Solution: action value functions

Following a standard method for solving Markov Decision Problems (Bellman, 2003), we aim to calculate the *action-value function* $Q_T(n, \Delta n \mid a)$, defined as the minimal (expected) cost of taking action a in state $(n, \Delta n)$, over some finite time horizon T. To do so, we follow Bellman's principle of optimality, which reads:

$$Q_T(n, \Delta n \mid a) = \mathrm{cost}(n, \Delta n, a) + \gamma \cdot \mathbb{E}\big[\min_{a'} Q_{T-1}(n', \Delta n' \mid a') \big]. \tag{3}$$

Here, the expectation is over the subsequent state $n', \Delta n'$, which (depending on the action a) may be uncertain, and γ is a *discount factor*, which de-emphasises costs incurred further in the future. (Note that the subsequent state has subscript $T-1$, being one step *closer* to the chosen horizon.) Having calculated this value for each actions, the best action in any given state is simply the one with the smallest cost function.

For $T = 0$ the second term in (3) is zero, so that $Q_0(n, \Delta n \mid a) = \exp(n + (\Delta n)^2/2) + \mathrm{cost}(a)$. For $T > 0$, we begin by evaluating the expected cost incurred if only control actions are taken, so that all randomness disappears. Then, to avoid the "min" in (3), we can compute the value of a *sequence* $\mathbf{a}$ of actions, with length T, by the same process (noting that the optimal action with $T = 0$ will always be to do nothing). We claim that, in this case, the expected cost is always a linear function of $N = \exp(n)$, with coefficients (i.e., slope and intercept) depending on Δn. This proceeds by induction on T, the length of the action sequence $\mathbf{a}$. If $T = 0$, then as above

$$Q_0(n, \Delta n \mid -) = \exp(n + (\Delta n)^2/2) = e^{(\Delta n)^2/2} N.$$

Then, suppose that we have computed the slope A and intercept B associated to δn and some sequence of actions $\mathbf{a}$, viz:

$$Q_T(n, \Delta n \mid \mathbf{a}) = A(\Delta n, \mathbf{a}) \cdot N + B(\mathbf{a}).$$

To extend this expression to length $T+1$, we add an extra action a_i and substitute into (3) to obtain

$$\begin{aligned} Q_{T+1}(n, \Delta n \mid a_i \mathbf{a}) &= \exp(n + (\Delta n)^2/2) + c_i + \gamma Q_T(n + r + \log(1 - \rho_i), \Delta n + \Delta r \mid \mathbf{a}) \\ &= \Big[e^{(\Delta n)^2/2} + \gamma(1 - \rho_i) e^r A_T(\Delta n + \Delta r, \mathbf{a}) \Big] \cdot N + [c_i + \gamma B_T(\mathbf{a})]. \end{aligned}$$

For a fixed $\Delta n, T$ and list of M possible actions, the full library of these coefficients can be calculated recursively in $O(MT)$ time. For fixed Δn and T, the *state-value function* $V_T(n, \Delta n) = \min_{\mathbf{a}} Q_T(n, \Delta n \mid \mathbf{a})$ is piecewise-linear and convex (in N), making it again straightforward to compute recursively.

To add in monitoring actions, we make one further simplifying assumption: that in any sequence of actions, monitoring only occurs once. That is, the error e in monitoring and rate Δr at which uncertainty grows are small enough that monitoring once ensures it is not necessary again. This ensures that the term inside the expectation of (3) is a piecewise linear function of e^n, which makes it analytically calculable. Since $n = \log N$ is normally distributed, if $\varphi_{n,\Delta n}(x)$ is the probability distribution function for $x \sim \mathrm{Norm}(n, \Delta n)$, then the relevant integrals are (Jawitz, 2004, Section 2.2)

$$\int_l^u \varphi_{n,\Delta n}(x)dx = \frac{1}{2}\left(\mathrm{erf}\left(\frac{u-n}{\Delta n\sqrt{2}}\right) - \mathrm{erf}\left(\frac{l-n}{\Delta n\sqrt{2}}\right)\right)$$

and

$$\int_l^u e^x\varphi_{n,\Delta n}(x)dx = \frac{\exp(n+\Delta n^2/2)}{2}\left(\mathrm{erf}\left(\frac{u-n-\Delta n^2}{\Delta n\sqrt{2}}\right) - \mathrm{erf}\left(\frac{l-n-\Delta n^2}{\Delta n\sqrt{2}}\right)\right).$$

These permit the calculation of the value of monitoring (noting that we can no longer predict what future actions will be, since the outcome of monitoring is stochastic):

$$Q_T(n, \Delta n \mid \mathrm{monitor}) = \exp(n + (\Delta n)^2/2) + \mathrm{cost}(\mathrm{monitor}) + \gamma\int_{-\infty}^{\infty} V_{T-1}(x+r, \Delta n')\varphi_{n,\Delta n}(x)dx,$$

where $\Delta n'$ is given by (2). For (control) actions *before* a monitoring step, the process of backwards induction works in the same way.

2.4. Model Parameters

The parameters specifying the model are collected in Table 1. Fixing these based on the scenario in question, the optimum action depends on the abundance of the species, and the uncertainty in that value. If we denote the abundance by N, then given an estimate n of $\log N$ with approximate error Δn, our model computes the optimum action.

3. Results

To illustrate our results, we plot the optimum sequence of actions for pairs $(\Delta n, \overline{N})$ of uncertainty and expected abundance. For the following examples, we take $r = 1.2 \pm 0.05$ and $\gamma = 0.9$. Control is assumed to cost $c_e = 4$ units with effectiveness $\rho = 0.6$, and monitoring costs $c_m = 1$ unit, with error $e = 0.1$.

The simplest interesting case is over $T = 2$ time steps: at the first step a decision must be made

Notation	Variable name in code	Definition
r	`r`	Disease growth rate
Δr	`dr`	Uncertainty in growth rate
ρ_i	`rho`	Control efficacy
c_i	`ce`	Control cost
—	`interventions`	A dictionary containing pairs `(ce,rho)` of cost & efficacy
c_m	`cm`	Cost of monitoring
e	`err_mon`	Error in monitoring
—	`mon_dict`	A dictionary of triples `(cm,err_mon,rho_mon)` of cost, error & control efficacy
T	`num_steps`	Decision horizon
γ	`gamma`	Discount factor

Table 1: Model parameters.

between supressing the outbreak (a control intervention), collecting information to inform the second decision (a monitoring intervention), or doing nothing. The outcome of this analysis is given in fig. 1.

Qualitatively, the displayed results match intuition. For small values of the uncertainty (in this case $\Delta n < 0.5$), monitoring is never optimal, as enough is known about the problem already. In this case, the decision between control and doing nothing is made based on the expected future abundance of the species. As uncertainty increases (rightwards on the graph), monitoring actions become more valuable, and as such optimal over a larger range of parameter values, matching the orange region on the right of the graph. Furthermore, the marginal benefit of monitoring is greater when the decision is less "obvious" — that is, when the estimated abundance is not very high or very low. In these extreme cases, even with a large "factual uncertainty" in the prevalence of the species, the "decision uncertainty" is small, so monitoring interventions are not worthwhile.

The simple case of fig. 1 illustrates the qualitative success of our proposed method. For practical applications, however, it could be useful to implement more complexity. In fig. 2 we illustrate some of the ways this is possible. To clarify the entries in the legend of fig. 2a, that is, the actions available to the manager, we include a tree representation of the same, in fig. 2b.

Figure 2 has distinct regions corresponding to '`monitor`' and '`control, monitor`', which arise as we have added one extra time step ($T = 3$). This is informed by the fact that uncertainty increases over time[1], but also by the benefits of suppressing an exponentially growing outbreak earlier rather than later.

In addition, in fig. 2 we have added a second control action. Specifically, this is modeled as the same intervention performed twice, with independent results, with a cost of $2c_e$ and an efficacy of $1 - (1 - \rho)^2$. As the abundance increases (upwards on the plot), these more intense control actions become optimal.

[1]Note that, with different parameter values, the action `ignore,monitor` appears on the plot.

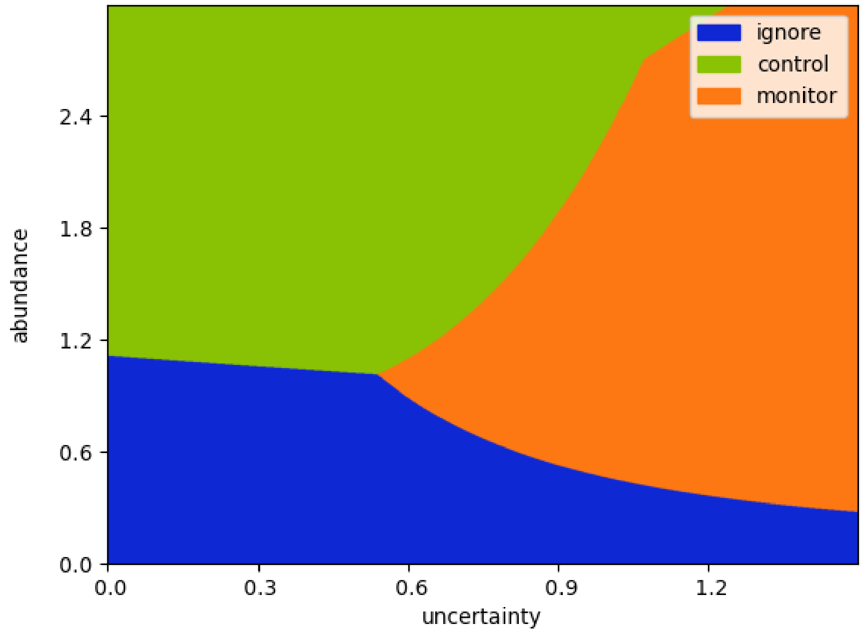

Figure 1: Optimal management actions for a sequence of $T = 2$ time steps for given uncertainty and abundance. Management options consist of 'ignore', 'monitor' or 'control'.

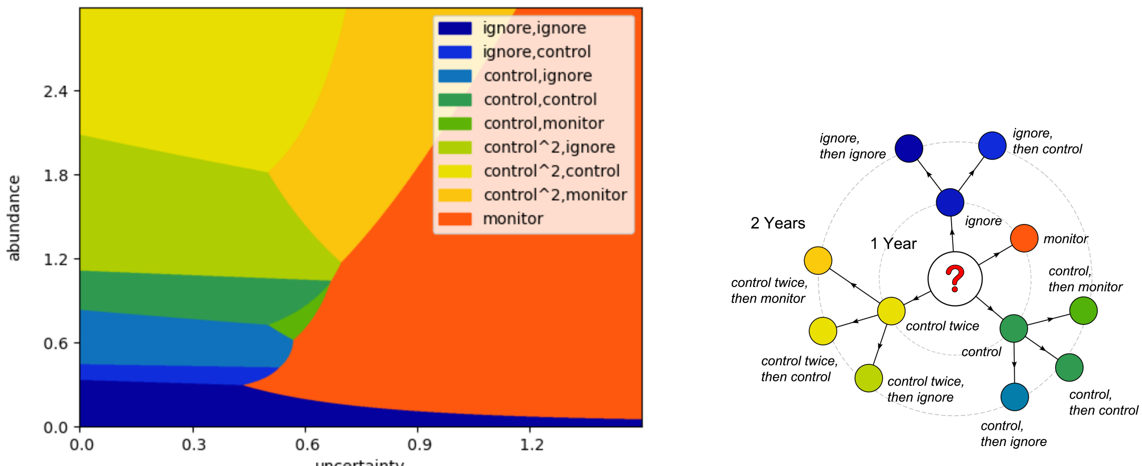

(a) Optimal management actions over three time steps

(b) A tree, representing the choices available to the manager.

Figure 2: Optimal management actions for further complexity, here $T = 3$, illustrating the quantitative power of our model.

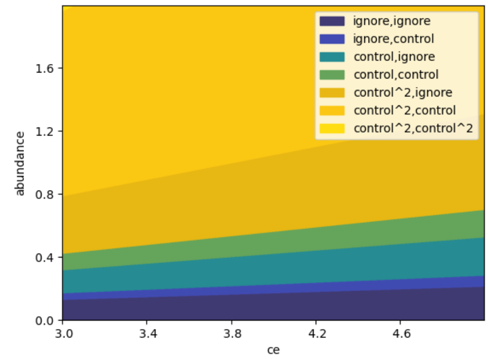

(a) Optimal actions, for varying c_e.

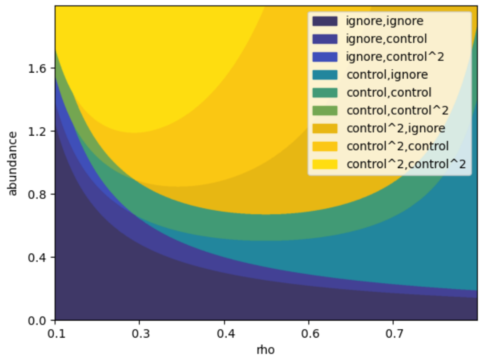

(b) Optimal actions, for varying ρ.

Figure 3: Parameter sensitivity analysis for control parameters (the cost of controlling c_e and control efficacy ρ) with no initial uncertainty ($\Delta n = 0$), for varying abundance $\overline{N}$ for $T = 3$. Apart from c_e (left) and ρ (right), parameters are kept the same to the example in fig. 2.

In fig. 3, we test the effect that variation of the parameters c_e and ρ has on the optimal control actions — ignoring for the moment any uncertainty, and therefore any monitoring actions. For the cost parameter, fig. 3a demonstrates that as the cost increases (moving rightwards on the graph), the species must be more abundant to justify a given control action. Figure 3b analyses the efficacy parameter ρ in an analogous way. For a small ρ, the abundance must be very high to justify any control action. As ρ increases, it becomes worthwhile to control the species at smaller abundances. For ρ close to 1, controlling once tends to be sufficient — as such, the regions corresponding to more expensive control actions (yellow and gold) become smaller.

To examine the effect of varying parameters related to monitoring, we can plot the border of the region where monitoring is optimal, as abundance and uncertainty vary. This corresponds, for example, to the edge of the orange region in fig. 2. Figure 4 plots these frontiers for varying monitoring cost c_m and error e. As either of these parameters increase, monitoring becomes less worthwhile, and as such needs more uncertainty to justify it. This is demonstrated by the fact that the frontiers corresponding to larger cost and error (colored yellow) are further to the right.

4. Discussion

In this paper, we have presented a partially observable Markov decision process (POMDP) framework for informed invasive species management. In particular, we developed an intuitive, visual tool that can help to decide between when to monitor and when to control given an estimate (with uncertainty) of the

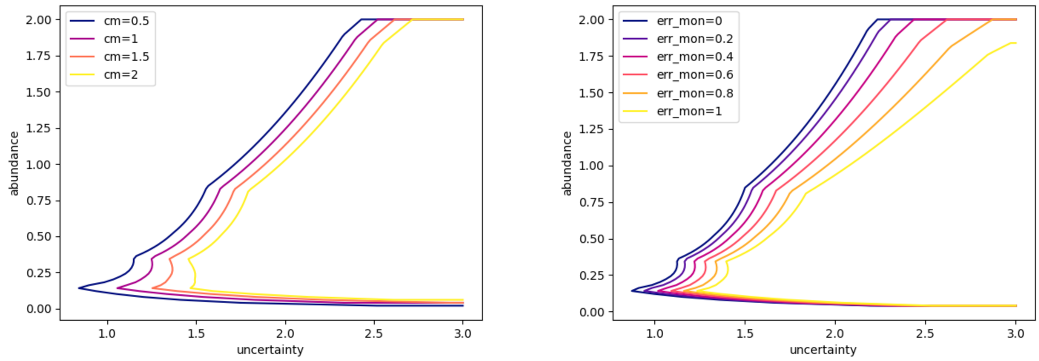

(a) Monitoring frontiers, for varying c_m.

(b) Monitoring frontiers, for varying e.

Figure 4: Parameter sensitivity analysis for monitoring parameters (the cost of monitoring c_m and the monitoring error e) showing the frontier where monitoring becomes optimal, plotted on the same axes as figs. 1 and 2. Other parameters are kept at the values used in the example in fig. 2.

species abundance. It is natural to parametrise an ecological problem, such as an invasive species, by its growth rate and the relative costs of control or monitoring compared to doing nothing at all. As such, our framework provides a very general and broadly applicable starting point for decision making in the presence of uncertainty about the state of the system being managed.

Importantly, we model abundance and uncertainty in a more general way than previous work. Previous POMDP approaches have simplified the ecological system dynamics to make the problem computational tractable Haight and Polasky (2010); Regan et al. (2011); Chadès et al. (2011); Rout et al. (2014). Instead, we simplify the computational problem by assuming a parametric distribution for the population size. Therefore, a major advantage of our POMDP approach is that it can deal with uncertainty in the state of the system, allowing analysis of the complex tradeoff between investment in monitoring and investment in control. Our model provides easily interpretable results without the restrictions of a specific scenario.

Similar problems and discussions can be found in the related field of disease management, where an analogous trade-off exists between monitoring, i.e. testing members of the population to gauge the spread of disease, and control, say by deploying medication, vaccination or implementing quarantine strategies. Although a variety of strategies have been proposed, POMDP solutions are still limited (Nowzari et al., 2016; Zino and Cao, 2021). Diseases whose management involves both surveillance (i.e. testing) and control (vaccination, medications and other control strategies) elements, such as HIV and COVID, could be particularly suitable for our method. One initial interesting proposal in this field is due to Hauskrecht and Fraser (2000), who apply a POMDP framework for diagnosis and treatment of ischemic heart disease.

Our framework is currently mainly limited by the assumptions of our model — for example, we assume that the population of the species grows exponentially, and ignore interactions with the rest of the environment. Our model could be extended to more general models of species infestation. However, we only solve our model over 3 time-steps, at the most, and, over such time-frames, exponential growth may be generally applicable.

Our model also contains the assumption that certain parameters, such as the control efficacy ρ, are known exactly, which could be relaxed with further work. One way include parameter uncertainty would be to link our work to an adaptive management framework, such as described by Moore et al. (2017). In such an 'adaptive management' framework, r or ρ would be added as dimensions to the model, so that the belief state of the manager includes a distribution over possible values of the parameters (Chadès et al., 2021, Section 4.3). In this situation, monitoring actions have the additional benefit of reducing parameter uncertainty, such as about the efficacy of management.

In this paper, we formulate a POMDP model which is applicable to a wide range of invasive species scenarios. In contrast with traditional POMDP approaches, we limit our model to dimensions of abundance and its uncertainty, retaining sufficient complexity to model the trade-off between monitoring and control, while remaining analytically solvable. This idea is not limited to the field of invasive species management, and has the potential to be applied to the 'reverse' problem of endangered species conservation, and to management of disease. POMDP approaches to the conservation of Sumatran tigers are proposed in Chadès et al. (2008); McDonald-Madden et al. (2011), with the similar limitations to the model of the underlying system. Hence, the framework proposed in this paper could be altered to investigate these scenarios in a more general way.

Author Contributions

T.K.W. developed the model based on preliminary work by M.A.M. in consultation with all authors. V.L.J.S and T.K.W. drafted the manuscript. T.K.W. ran the simulations and generated the results with input from V.L.J.S. All authors edited and improved the manuscript.

Data Availability Statement

The model and results presented in this paper are available at `https://github.com/thomaskwaring/to-monitor-or-control`.

References

S. Alvarez and D. Solís. Rapid response lowers eradication costs of invasive species. *Choices*, 33(4):1–9, 2018.

K. J. Åström. Optimal Control of Markov Processes with Incomplete State Information I. *Journal of mathematical analysis and applications*, 10:174–205, 1965. ISSN 0022-247X. doi: {10.1016/0022-247X(65) 90154-X}.

R. Bellman. *Dynamic Programming*. Dover Books on Computer Science Series. Dover Publications, 2003. ISBN 9780486428093.

I. E. Büyüktahtakın and R. G. Haight. A review of operations research models in invasive species management: state of the art, challenges, and future directions. *Annals of Operations Research*, 271: 357–403, 12 2018. ISSN 15729338. doi: 10.1007/S10479-017-2670-5/TABLES/2.

L. R. Carrasco, R. Baker, A. Macleod, J. D. Knight, and J. D. Mumford. Optimal and robust control of invasive alien species spreading in homogeneous landscapes. *Journal of the Royal Society Interface*, 7: 529–540, 2010. doi: 10.1098/rsif.2009.0266.

I. Chadès, E. McDonald-Madden, M. A. McCarthy, B. Wintle, M. Linkie, and H. P. Possingham. When to stop managing or surveying cryptic threatened species. *Proceedings of the National Academy of Sciences*, 105(37):13936–13940, 2008.

I. Chadès, T. G. Martin, S. Nicol, M. A. Burgman, H. P. Possingham, and Y. M. Buckley. General rules for managing and surveying networks of pests, diseases, and endangered species. *Proceedings of the National Academy of Sciences*, 108(20):8323–8328, 2011.

I. Chadès, L. V. Pascal, S. Nicol, C. S. Fletcher, and J. Ferrer-Mestres. A primer on partially observable Markov decision processes (POMDPs). *Methods in Ecology and Evolution*, 12(11):2058–2072, 2021. ISSN 2041-210X. doi: 10.1111/2041-210X.13692.

R. S. Epanchin-Niell, R. G. Haight, L. Berec, J. M. Kean, and A. M. Liebhold. Optimal surveillance and eradication of invasive species in heterogeneous landscapes. *Ecology Letters*, 15:803–812, 8 2012. ISSN 1461023X. doi: 10.1111/J.1461-0248.2012.01800.X.

R. G. Haight and S. Polasky. Optimal control of an invasive species with imperfect information about the level of infestation. *Resource and Energy Economics*, 32(4):519–533, 2010.

C. E. Hauser and M. A. McCarthy. Streamlining 'search and destroy': cost-effective surveillance for invasive species management. *Ecology Letters*, 12:683–692, 7 2009. ISSN 1461-0248. doi: 10.1111/J.1461-0248.2009.01323.X.

M. Hauskrecht and H. Fraser. Planning treatment of ischemic heart disease with partially observable markov decision processes. *Artificial intelligence in medicine*, 18(3):221–244, 2000.

B. D. Hoffmann and L. M. Broadhurst. The economic cost of managing invasive species in australia. *NeoBiota*, 31:1–18, 2016.

K. Hyytiäinen, M. Lehtiniemi, J. K. Niemi, and K. Tikka. An optimization framework for addressing aquatic invasive species. *Ecological Economics*, 91:69–79, 2013.

J. Jawitz. Moments of truncated continuous univariate distributions. *Advances in Water Resources - ADV WATER RESOUR*, 27:269–281, 03 2004. doi: 10.1016/j.advwatres.2003.12.002.

M. J. Keeling and P. Rohani. *Modeling Infectious Diseases in Humans and Animals*. Princeton University Press, 2008. ISBN 978-0-691-11617-4. doi: 10.2307/j.ctvcm4gk0.

M. L. Littman. A tutorial on partially observable markov decision processes. *Journal of Mathematical Psychology*, 53(3):119–125, 2009.

L. Marescot, G. Chapron, I. Chadès, P. L. Fackler, C. Duchamp, E. Marboutin, and O. Gimenez. Complex decisions made simple: a primer on stochastic dynamic programming. *Methods in Ecology and Evolution*, 4(9):872–884, 2013.

M. L. N. Mbah and C. A. Gilligan. Balancing detection and eradication for control of epidemics: Sudden oak death in mixed-species stands. *PLOS ONE*, 5:e12317, 2010. ISSN 1932-6203. doi: 10.1371/JOURNAL.PONE.0012317.

E. McDonald-Madden, I. Chadès, M. A. McCarthy, M. Linkie, and H. P. Possingham. Allocating conservation resources between areas where persistence of a species is uncertain. *Ecological Applications*, 21(3):844–858, 2011.

S. V. Mehta, R. G. Haight, F. R. Homans, S. Polasky, and R. C. Venette. Optimal detection and control strategies for invasive species management. *Ecological Economics*, 61(2-3):237–245, 2007.

G. E. Monahan. State of the art—a survey of partially observable markov decision processes: theory, models, and algorithms. *Management science*, 28(1):1–16, 1982.

A. L. Moore, L. Walker, M. C. Runge, E. McDonald-Madden, and M. A. McCarthy. Two-step adaptive management for choosing between two management actions. *Ecological Applications*, 27(4):1210–1222, 2017. doi: https://doi.org/10.1002/eap.1515.

J. L. Moore, T. M. Rout, C. E. Hauser, D. Moro, M. Jones, C. Wilcox, and H. P. Possingham. Protecting islands from pest invasion: optimal allocation of biosecurity resources between quarantine and surveillance. *Biological Conservation*, 143:1068–1078, 2010. doi: 10.1016/j.biocon.2010.01.019.

C. Nowzari, V. M. Preciado, and G. J. Pappas. Analysis and Control of Epidemics: A Survey of Spreading Processes on Complex Networks. *IEEE Control Systems Magazine*, 36(1):26–46, 2016. ISSN 1066033X. doi: 10.1109/MCS.2015.2495000.

S. Polasky. A model of prevention, detection, and control for invasive species. In *Bioinvasions and Globalization: Ecology, Economics, Management, and Policy*. Oxford University Press, 2010.

T. J. Regan, I. Chadès, and H. P. Possingham. Optimally managing under imperfect detection: a method for plant invasions. *Journal of Applied Ecology*, 48(1):76–85, 2011.

T. M. Rout, J. L. Moore, H. P. Possingham, and M. A. Mccarthy. Allocating biosecurity resources between preventing, detecting, and eradicating island invasions. *Ecological Economics*, 71:54–62, 2011. doi: 10.1016/j.ecolecon.2011.09.009.

T. M. Rout, J. L. Moore, and M. A. Mccarthy. Prevent, search or destroy? a partially observable model for invasive species management. *Journal of Applied Ecology*, 51:804–813, 6 2014. ISSN 1365-2664. doi: 10.1111/1365-2664.12234.

L. Zino and M. Cao. Analysis, prediction, and control of epidemics: A survey from scalar to dynamic network models. *IEEE Circuits and Systems Magazine*, 21(4):4–23, 2021.