Biological Conservation 224 (2018) 9–15

Contents lists available at ScienceDirect

Biological Conservation

journal homepage: www.elsevier.com/locate/biocon

Adaptive management of ecological systems under partial observability

Milad Memarzadeh*, Carl Boettiger

Dept of Environmental Science, Policy, and Management, University of California Berkeley, Berkeley, CA 94720-3114, USA

ARTICLE INFO

Keywords:
Adaptive management
Decision making under uncertainty
Conservation
Fisheries

ABSTRACT

Adaptive management has a long history in ecology and conservation. Uncertainty in both the state of a system and the model defining its dynamics are fundamental challenges in adaptive management of complex ecological systems. Traditional approaches in conservation biology often ignore one or both sources of uncertainty due to the computational complexity involved. Here, we show that underestimating the role of uncertainty in both model estimation and decision-making results in aggressive decision rules which can potentially lead to the dramatic decline and possible collapse of a population, species, or ecosystem. We propose an approximate solution to adaptive management of ecological systems under both model and state uncertainties that is computationally feasible and applicable to complex management problems and provide a software for detailed implementation of our method, http://doi.org/10.5281/zenodo.1161521. We apply the proposed method in a marine ecosystem management context and show that by learning from historical data and arrival of new observations, decision makers can adapt their policies to avoid decline in the population and reach a sustainable population stability.

1. Introduction

Many studies in natural resource management emphasize the importance of incorporating uncertainty in the management process through adaptive management (Walters and Hilborn, 1978; Humphrey and Stith, 1990; Lancia et al., 1996; Williams, 2011a). In general, adaptive management is defined as a management that learns by doing (i.e. planning) and adopts management policies to reflect new observations (i.e. learning) (Walters and Holling, 1990; Williams and Brown, 2016). Adaptive management has a long history in ecological literature and have been widely used in behavioral ecology (Mangel and Clark, 1988), optimal harvesting in fisheries (Walters and Hilborn, 1976; Reed, 1979), as well as conservation biology and natural resource economics (Mangel, 1985; Moore et al., 2010; Britton et al., 2011). Walters and Hilborn (1976) are among the first implementations of adaptive management to control the population of Fraser River sockeye salmon under uncertainty in models describing the population dynamics.

Less computationally intensive approaches for natural resource management have since followed, such as management strategy evaluation (MSE) (Smith, 1994; Mapstone et al., 2008; Bunnefeld et al., 2011). Instead of attempting to solve for the optimal management policy in the space of all possible sequences of actions a manager could take, MSE evaluates consequences of a pre-determined set of strategies defined based on management objectives and provides the outcome of those simulations to the decision maker. Although this approach has proved to be promising in many applications in natural resource management, it is hard to quantify the quality of the selected management strategies compared to the optimal management.

Optimal management is the one strategy that results in the best outcome among all possible strategies and hence can be found by optimizing the management objectives over the life-span of the natural resource. Methods from the domains of decision theory and optimal control (Bertsekas, 1996; Sutton and Barto, 1998), such as Markov decision process (MDP), can be used to determine such an optimal policy, which has been a basis of adaptive management theory in ecology for years (Reed, 1979; Mangel, 1985; Clark and Kirkwood, 1986; Mangel and Clark, 1988; Sethi et al., 2005).

One of the main limitations of Markov decision process is the strong assumption about full observability of the system's state. For example, in fisheries management, the assumption is that the population biomass of the fishery of interest (i.e. state of the system) is measured without error, and fisher has certain knowledge of the exact population size before setting the harvest quota each year (i.e. management action). Relying on this assumption, Reed (1979) proves that the intuitive results of deterministic models (Beverton and Holt, 1957; Schaefer, 1957) remain optimal for maximizing the expected economic return of fisheries subject to stochastic growth. Clark and Kirkwood (1986) acknowledge the importance of the state uncertainty, due to error in the measurements, to the optimal ecological management, while observing

* Corresponding author.
E-mail address: miladm@berkeley.edu (M. Memarzadeh).

https://doi.org/10.1016/j.biocon.2018.05.009
Received 27 September 2017; Received in revised form 1 May 2018; Accepted 14 May 2018
0006-3207/ Published by Elsevier Ltd.

that if this assumption were relaxed, the “difficulty of the problem increases markedly”. Due to this difficulty and the conclusions of Reed (1979), most of the literature has focused on adaptive management with full observability of state space (e.g. certain knowledge of population biomass) and sidesteps this complexity in the decision making step of adaptive management (e.g. setting harvest rules) (Costello et al., 2016; Britten et al., 2017).

This has long begged the question: if the stochastic growth can be safely ignored in determining the optimal strategies for controlling the population (Reed, 1979), is the same true for measurement error? Since then, researchers have proposed approximate methods to incorporate the effect of measurement error on the optimal management policy and have come to a counter-intuitive conclusion that in presence of measurement error, the resulting policies are less conservative than the deterministic solutions (Ludwig and Walters, 1981; Clark and Kirkwood, 1986; Roughgarden and Smith, 1996; Engen et al., 1997; Sethi et al., 2005). Clark and Kirkwood (1986) acknowledge these counter-intuitive results: “results appear to contradict the conventional wisdom of renewable resource management, under which high uncertainty would call for increased caution in the setting of quotas”. We show that the reason for this counter-intuitive result arises from non-trivial assumptions made to simplify the decision optimization process of management and to get past the computational complexity. By fully incorporating the effect of measurement error, we illustrate that the resulting optimal strategies for controlling the system is significantly different (and not necessarily less conservative) with respect to the deterministic solutions, and ignoring this effect can result in dramatic declines in the species population and possibly total collapse of the ecosystem.

Aside from the uncertainty in the estimating population biomass due to the measurement error, another challenge in adaptive management is uncertainty in the models defining dynamics of the population (i.e. model uncertainty). Walters and Hilborn (1978) provide an early non-mathematical overview that explores the effect of model uncertainty on management outcomes in natural resource management. They notice that once model uncertainty is introduced, finding the optimal management becomes very complex and hence solving the decision optimization under model uncertainty becomes intractable. This problem is known as the curse of dimensionality (Bertsekas, 1996) in decision theory. As a result, most of attempts to include model uncertainty in adaptive management of natural resources have focused on simplified examples (McDonal-Madden et al., 2010; Runge, 2011).

Here, we propose an approximate solution to adaptive management of complex ecological systems under both model and state uncertainties that is computationally scalable and applicable to the complex real-world management problems, which we call PLUS (Planning and Learning for Uncertain Systems). PLUS is comprised of two steps: model estimation (i.e. learning) and decision optimization (i.e. planning). We also provide a publicly available software and detailed implementation of our method, as an R package at http://doi.org/10.5281/zenodo.1161521. The proposed method builds upon recent advancements in adaptive management in domains of engineering and computer sciences (Doshi-Velez et al., 2012; Memarzadeh et al., 2014) and addresses several gaps in the literature regarding adaptive management and its application to natural resource management such as: (1) proposed adaptive management approaches being prone to curse of dimensionality and as a result are only applied to simplified examples, (2) proposed approaches make non-trivial assumptions and/or approximations to get past the computational barrier in the cost of losing flexibility and generality, which might also result in finding bad decision rules that can potentially lead to dramatic decline in population of biological species and possibly total collapse of the ecosystem.

We build the decision optimization step based on the framework of partially observable Markov decision process (POMDP) (Smallwood and Sondik, 1973; Sondik, 1978). POMDPs overcome one of the main limitations of Markov decision process (MDP, that are very popular in adaptive management of natural resources) and incorporate measurement error in the decision making by allowing partial observability of the state space through the introduction of belief state: manager's knowledge about system's state in a shape of probability distribution. This means that POMDP incorporates uncertainty in estimating the population biomass in the probabilistic manner. The reader should note that the belief state is well-known in the family of state space models for estimating the dynamics model under state uncertainty (Costello et al., 2016; Britten et al., 2017), although has been ignored and marginalized out in the decision making step. POMDPs are able to incorporate state uncertainty in decision making step with introduction of much higher computational complexity. Detailed formulation of POMDPs is reported in Appendix A. POMDPs have recently been used widely in ecological literature for control of animal population (Runge, 2013), natural resource management (Williams, 2009, 2011b), fishery management (Kling et al., 2017), and conservation in the face of climate change (Conroy et al., 2011). Although these studies emphasize partial observability and environmental variability as the challenging sources of uncertainty in the management of natural resources, they do not specifically propose an efficient method and software to get past the computational barrier. Moreover, POMDPs still require perfect knowledge of the models defining the system's dynamics, and as a result cannot incorporate model uncertainty. The approach taken by Walters and Hilborn (1976) of including the model parameters into the state space can be also implemented in POMDPs, however finding optimal strategies under such formulation is intractable and suffers from the curse of dimensionality. Here, we develop a heuristic based on minimization of the Bayes risk to overcome the computational complexity of decision making under model uncertainty and allow incorporation of both state and model uncertainties feasibly. The detailed formulation of the heuristic is reported in Appendix B. Overall goal of the decision optimization step is to find a decision (i.e. catch quota in the beginning of each year) that maximizes the long-term expected economic return of managing the system (i.e. fishery).

As mentioned before, the second step corresponds to estimating the dynamics model from historical observations and reducing the model uncertainty. In this step, the goal is to represent the historical observations in a shape of probability distribution over population dynamics model, and adaptively update the model and reduce the uncertainty once a new set of observations arrive in the following years. One of the main challenges in this step is that observations are only noisy measurements of the state of the system and hence need careful consideration. In general, one might not be able to represent the posterior distribution of the model parameters given the historical observations in closed-form, and as a result, need to adapt an approximate learning scheme based on Markov chain Monte Carlo (MacKay, 2003; Memarzadeh et al., 2014, Memarzadeh et al., 2016; Costello et al., 2016; Britten et al., 2017). Although in this article, we assume that the uncertainty is among a set of candidate models (as in Walters and Hilborn, 1976) and as a result, learning can be done in closed-form using Bayes rule (refer to Appendix C for detailed formulations).

2. Model formulation

We model the population dynamics of the fishery as follow,

$$x_{t+1} = f(x_t) - c_t + \sigma_t^X \tag{1}$$

where $x_t \in X$ denotes the population biomass, $c_t \in C$ is the harvested biomass (i.e. catch quota), with subscript t denoting years, σ_t^X is the annual biomass deviation due to stochasticity in the population dynamics (which we call growth noise), and f is the function governing the dynamics. We consider two candidate functions for f as the Ricker model (Ricker, 1954),

$$f(x_t) = x_t \exp\left(r\left(1 - \frac{x_t}{K}\right)\right) \tag{2}$$