Environmental Modelling & Software 92 (2017) 125–141

Contents lists available at ScienceDirect

Environmental Modelling & Software

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

Direct policy search for robust multi-objective management of deeply uncertain socio-ecological tipping points

Julianne D. Quinn [a, *], Patrick M. Reed [a], Klaus Keller [b, c, d]

[a] School of Civil and Environmental Engineering, Cornell University, Ithaca, NY, USA
[b] Department of Geosciences, The Pennsylvania State University, University Park, PA, USA
[c] Department of Engineering and Public Policy, Carnegie Mellon University, Pittsburgh, PA, USA
[d] Earth and Environmental Systems Institute, The Pennsylvania State University, University Park, PA, USA

A R T I C L E   I N F O

Article history:
Received 15 June 2016
Received in revised form
17 February 2017
Accepted 20 February 2017

Keywords:
Socio-ecological management
Multi-objective decision making
Direct policy search
Tipping points
Robustness
Deep uncertainty

A B S T R A C T

Managing socio-ecological systems is a challenge wrought by competing societal objectives, deep uncertainties, and potentially irreversible tipping points. A classic, didactic example is the shallow lake problem in which a hypothetical town situated on a lake must develop pollution control strategies to maximize its economic benefits while minimizing the probability of the lake crossing a critical phosphorus (P) threshold, above which it irreversibly transitions into a eutrophic state. Here, we explore the use of direct policy search (DPS) to design robust pollution control rules for the town that account for deeply uncertain system characteristics and conflicting objectives. The closed loop control formulation of DPS improves the quality and robustness of key management tradeoffs, while dramatically reducing the computational complexity of solving the multi-objective pollution control problem relative to open loop control strategies. These insights suggest DPS is a promising tool for managing socio-ecological systems with deeply uncertain tipping points.

Software and data availability

Description and Availability: The Lake Problem optimization code, MORDM re-evaluation code, and best final reference sets are available on Github at https://github.com/julianneq/Lake_Problem_DPS. The optimization and hypervolume calculations can be replicated using the software code available for the Borg MOEA (http://borgmoea.org/), pareto.py (https://github.com/matthewjwoodruff/pareto.py) and the MOEA framework (http://moeaframework.org/)

Developer: The simulation code was adapted by Julianne Quinn from code developed by Riddhi Singh with adaptations by Tori Ward, Dave Hadka and Jon Herman

Funding Source: Development of the code was partially supported by the National Science Foundation, through the Network for Sustainable Climate Risk Management (SCRiM) under NSF cooperative agreement GEO-1240507 as well as the Penn State Center for Climate Risk Management

Source Language: The optimization code is written in C++ and the re-evaluation code in Python

License: GNU Lesser General Public License, Version 3

1. Introduction

As economic development continues globally, severe ecological consequences of human actions are manifesting themselves in many forms. Altered nutrient cycling, shifting biomes, and decreased biodiversity are just a few examples of the repercussions of anthropogenic activities (Parry, 2007). More responsible socio-

* Corresponding author. School of Civil and Environmental Engineering, 205 Hollister Hall, Ithaca, NY 14853, USA.
E-mail addresses: jdq8@cornell.edu (J.D. Quinn), patrick.reed@cornell.edu (P.M. Reed), klaus@psu.edu (K. Keller).

http://dx.doi.org/10.1016/j.envsoft.2017.02.017
1364-8152/

ecological management will require balancing conflicting objectives, some of which exhibit uncertain and precarious threshold behavior (e.g., Werners et al., 2013; Keller et al., 2008). For example, we are currently balancing a severe tradeoff between increasing energy production using fossil fuels and avoiding potentially irreversible ecological damages from crossing a threshold atmospheric CO_2 concentration (Solomon et al., 2009). In fact, Lenton et al. (2008) highlight eight components of the Earth System that could reach catastrophic tipping points as a result of global warming, with the areal extent of Artic summer sea-ice and the Greenland ice sheet facing the most imminent threat.

In environmental systems with thresholds, balancing conflicts in societal values or objectives is further complicated by severe uncertainties associated with identifying thresholds as well as the consequences of crossing them (Lenton, 2013; Keller & McInerney, 2008). These uncertainties are often considered “deep” or Knightian uncertainties, meaning planners cannot agree on prior probability density functions to describe the parameters of the system model, or even on the model itself (Lempert & Collins, 2007; Knight, 1921). In these cases, it is desirable to find robust management plans that perform well across a broad range of possible system conditions (Herman et al., 2015; Kwakkel et al., 2016).

Since its seminal inception (Bankes, 1993; Lempert et al., 2002, 2010; Walker et al., 2003), the field of decision making under deep uncertainty has emphasized a transition from classical “predict then act” risk management frameworks to exploratory modeling frameworks (e.g., see Dessai et al., 2009). These methods move beyond planning for a single expected future and instead emphasize investigating the response of system management plans to a wide range of deeply uncertain states-of-the-world (SOWs) in order to discover robust actions for avoiding unacceptable outcomes (Bryant & Lempert, 2010; Lempert et al., 2006; Hall et al., 2012). In their recent review, Herman et al. (2015) highlight the rapid growth in new methodologies and applications of decision analysis frameworks focused on robustness or deep uncertainty, such as robust decision making (RDM) (Lempert et al., 2006), dynamic adaptive policy pathways (Haasnoot et al., 2013), many-objective robust decision making (MORDM) (Kasprzyk et al., 2013), and decision scaling (Brown et al., 2012). Despite the growing diversity of robustness-focused frameworks, the taxonomy of methods presented by Herman et al. (2015) emphasizes the commonalities between them and the importance of bridging their capabilities to advance the field. These approaches share four core methodological components: (1) eliciting or searching for alternative management actions, (2) using exploratory modeling to broadly sample possible SOWs that could impact the performance of alternative policies or actions, (3) eliciting robustness measures that distinguish SOWs of concern, and (4) potentially using sensitivity analysis to clarify the key factors that most strongly influence robustness for subsequent monitoring (Herman et al., 2015).

This study advances the MORDM framework (Kasprzyk et al., 2013) with a specific focus on two technical contributions: (1) demonstrating the value and use of direct policy search (DPS) (Rosenstein & Barto, 2001) for identifying adaptive robust operational control strategies for socio-ecological systems and (2) demonstrating how nonlinear environmental thresholds, or tipping points, pose fundamental challenges for balancing economic benefits and their consequent risks to socio-ecological systems. As initially developed by Kasprzyk et al. (2013), the MORDM framework focuses on aiding decision makers and stakeholders in learning how to frame complex, ill-defined environmental planning problems and in discovering robust decisions that perform well across a broad array of possible SOWs. A distinguishing feature of MORDM relative to other frameworks is its use of many-objective evolutionary optimization to identify approximately Pareto optimal management decisions. Pareto optimal, or non-dominated, decisions represent those management actions for which improvement in one objective is only possible with degrading performance in one or more other objectives (Pareto, 1896). These solutions are first discovered through multi-objective optimization to one's best estimate of the true SOW. The solutions are then re-evaluated in alternative SOWs to determine how robust they are to uncertainties in system parameters. At its core, the MORDM framework provides *a posteriori* decision support, meaning it first presents explicit representations of key system tradeoffs and robustness challenges and then elicits stakeholder preferences in selecting management actions (i.e., generate first, choose later, as classified by Cohon & Marks (1975)).

In this study, we demonstrate the value of exploiting DPS in the MORDM framework using the classical shallow lake problem (Carpenter et al., 1999). In this didactic example, a hypothetical town situated on a lake attempts to balance the economic benefits it receives from discharging phosphorus (P) into the lake with the environmental costs of irreversibly tipping the lake into a eutrophic state. The behavior of this stylized model of lake eutrophication is representative of many socio-ecological systems with tipping points, such as harvested fish populations, grasslands consumed by cattle on rangelands, and global carbon cycle dynamics (Carpenter et al., 2015; Anderies et al., 2013). Early work on the lake problem (e.g., Carpenter et al., 1999; Lempert and Collins, 2007) has focused on optimizing the town's pollution control policy to maximize the expected net present value of a utility function which rewards economic benefits and penalizes pollution using a monetary valuation of displaced ecological benefits. Collapsing these objectives into a single expected utility function poses several problems. First, it assumes *a priori* knowledge of stakeholders' values, and agreement among stakeholders on those values. Monetizing environmental benefits and costs to find a single “optimal” solution can fail to capture the full range of achievable objective values, which would better embody the range of preferences among different stakeholders (see, e.g., Admiraal et al., 2013). Second, maximizing the expected value of a utility function requires agreement on the probability distribution of stochastic inputs, which poses severe challenges for systems with deeply uncertain characteristics (Lempert & Collins, 2007; Knight, 1921).

Recent many-objective extensions of the lake problem have sought to explicitly capture the tradeoffs between economic and environmental objectives (e.g., Singh et al., 2015; Hadka et al., 2015; Ward et al., 2015), as well as deep uncertainty in the lake model parameters (e.g., Singh et al., 2015; Hadka et al., 2015). Ward et al. (2015) find that when optimizing pollution control strategies for the town as a time series of P release decisions, several state-of-the-art multi-objective evolutionary algorithms (MOEAs) fail to find effective policies due to the high dimensional decision space for candidate pollution control action, weak system responses to late period decisions (i.e., temporal salience structure as discussed by Thierens et al. (1998)), and the non-linear pollution threshold. One way to potentially overcome these challenges is to employ a closed loop control method in which knowledge of the system state is used as a feedback control to inform the decision at each time step (Bertsekas, 1995). Not only can the additional information provided at each time step improve the signal of the late-period decisions, but it can also allow for a different set of P release decisions under different realizations of stochastic P inflows. The open loop intertemporal pollution control strategy employed by Ward et al. (2015), however, only finds one vector of pollution control decisions that perform best in expectation, and is reflective of the methodologies used in many environmental policy studies (e.g., Nordhaus, 2013).

Here, we employ a closed loop control strategy called direct policy search (DPS), that has proven to be a simple and