Environmental Modelling & Software 74 (2015) 114–129

Contents lists available at ScienceDirect

Environmental Modelling & Software

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

An open source framework for many-objective robust decision making

David Hadka [a,*], Jonathan Herman [f], Patrick Reed [b], Klaus Keller [c,d,e]

[a] Applied Research Laboratory, The Pennsylvania State University, University Park, PA, USA
[b] School of Civil and Environmental Engineering, Cornell University, Ithaca, NY, USA
[c] Department of Geosciences, The Pennsylvania State University, University Park, PA, USA
[d] Department of Engineering and Public Policy, Carnegie Mellon University, Pittsburgh, PA, USA
[e] Earth and Environmental Systems Institute, The Pennsylvania State University, University Park, PA, USA
[f] Department of Civil & Environmental Engineering, University of California, Davis, CA, USA

ARTICLE INFO

Article history:
Received 9 December 2014
Received in revised form
22 July 2015
Accepted 23 July 2015
Available online 1 October 2015

Keywords:
Robust decision making
Deep uncertainty
Threshold
Scenario discovery
Risk assessment
Robustness

ABSTRACT

This study introduces a new open source software framework to support bottom-up environmental systems planning under deep uncertainty with a focus on many-objective robust decision making (MORDM), called OpenMORDM. OpenMORDM contains two complementary components: (1) a software application programming interface (API) for connecting planning models to computational exploration tools for many-objective optimization and sensitivity-based discovery of critical deeply uncertain factors; and (2) a web-based visualization toolkit for exploring high-dimensional datasets to better understand system trade-offs, vulnerabilities, and dependencies. We demonstrate the OpenMORDM framework on a challenging environmental management test case termed the “lake problem”. The lake problem has been used extensively in the prior environmental decision science literature and, in this study, captures the challenges posed by conflicting economic and environmental objectives, a water quality “tipping point” beyond which the lake may become irreversibly polluted, and multiple deeply uncertain factors that may undermine the robustness of pollution management policies. The OpenMORDM software framework enables decision makers to identify policy-relevant scenarios, quantify the trade-offs between alternative strategies in different scenarios, flexibly explore alternative definitions of robustness, and identify key system factors that should be monitored as triggers for future actions or additional planning. The web-based OpenMORDM visualization toolkit allows decision makers to easily share and visualize their datasets, with the option for analysts to extend the framework with customized scripts in the R programming language. OpenMORDM provides a platform for constructive decision support, allowing analysts and decision makers to interactively discover promising alternatives and potential vulnerabilities while balancing conflicting objectives.

Software availability

- **Name of Software**: OpenMORDM
- **Description**: OpenMORDM is an open-source R library for multiobjective robust decision making (MORDM). It includes support for loading datasets from a number of sources including CSV, XLS, XLSX, databases, and R matrices and data frames; visualizing the data sets using various 2D and 3D plots; performing scenario discovery and trade-off analysis; and computing uncertainty/robustness metrics. OpenMORDM also includes a web-based data exploration and visualization toolkit.
- **Developer: D**. Hadka (dmh309@psu.edu) with contributions by P. Reed and K. Keller.
- **Funding Source**: Development 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**: R
- **Supported Systems**: Unix, Linux, Windows, Mac
- **License**: GNU General Public License, Version 3

* Corresponding author.
E-mail addresses: dmh309@psu.edu (D. Hadka), jdherman@ucdavis.edu (J. Herman), patrick.reed@cornell.edu (P. Reed), klaus@psu.edu (K. Keller).

http://dx.doi.org/10.1016/j.envsoft.2015.07.014

- **Availability**: http://github.com/dhadka/OpenMORDM

1. Introduction

A critical component of environmental planning and management is the search for robust solutions capable of withstanding deviations from our best projections of the future. This challenge is amplified by the presence of deep uncertainty, where the suite of all possible future events as well as their associated probability distributions are themselves uncertain (e.g., future climatological and hydroeconomic factors [Knight (1921), Lempert (2002), Olson et al. (2012)]). These challenges have led to several “bottom-up” decision support frameworks [Nazemi and Wheater (2014), Weaver et al. (2013)], which move beyond trying to predict the most probable future(s) to discover which states of the world (SOWs) may lead to high consequence system vulnerabilities. This step helps with the task of evaluating the likelihoods of the discovered system vulnerabilities as it can help to focus the analysis on a subset of plausible future scenarios (e.g., Lempert et al. (2012)). Bottom-up or robustness-based approaches include Decision Scaling [Brown (2010), Brown et al. (2012)], Information-Gap (Info-Gap) [Ben-Haim (2004)], Robust Decision Making (RDM) [Lempert (2002), Lempert et al. (2013, 2006), Groves and Lempert (2007), Lempert and Collins (2007)], and Many-Objective Robust Decision Making (MORDM) [Kasprzyk et al. (2013)]. As highlighted by Herman et al. (2015), these bottom-up frameworks can be generalized into four steps: identifying decision alternatives, sampling states of the world, specifying robustness measures, and performing scenario discovery to identify the most important uncertainties. The final step, scenario discovery, is commonly used to find policy-relevant controls by determining the ranges of each uncertainty leading to system failure [Lempert et al. (2006)]. Herman et al. (2015) note that while these methods are often defined at a conceptual level, specific implementations share a number of potentially interchangeable concepts which should be compared to understand consequences for decision support. This work addresses the need for software and visualization tools to flexibly support the quantitative components of these “bottom-up” environmental systems planning frameworks, which share the goal of identifying robust solutions. The following paragraphs introduce the conceptual frameworks for decision support under deep uncertainty, while the quantitative methods implemented in this work are described in Section II.

Robust Decision Making (RDM), like other “bottom-up” approaches, seeks to distinguish robust solutions which provide satisfactory performance across many plausible SOWs [Lempert (2002), Lempert et al. (2013), Groves and Lempert (2007), Lempert and Collins (2007)]. Given a pre-specified set of alternatives to analyze, RDM subjects each to an ensemble of SOWs that are treated as exploratory samples over plausible ranges of uncertain factors [Bryant and Lempert (2010), Groves and Lempert (2007), Lempert et al. (2006, 2012)]. The goal — as is generally the case in Decision Scaling and Info-Gap — is to identify future scenarios that may cause the system to fail. RDM studies often adopt a “satisficing” approach [Simon (1959)], in which solutions must satisfy performance requirements across many plausible futures rather than provide optimal performance in a single future. Using a satisficing approach, robustness can be quantified with the domain criterion [Schneller and Sphicas (1983), Starr (1962)] which aims to maximize the volume of the uncertain factor space in which performance requirements are satisfied [Lempert and Collins (2007)]. Additionally, RDM analyses typically employ the Patient Rule Induction Method (PRIM) [Friedman and Fisher (1999)] to perform a high-dimensional sensitivity analysis for scenario discovery in order to identify the ranges of uncertain factors most likely to cause system failure [Bryant and Lempert (2010), Groves and Lempert (2007), Lempert et al. (2006, 2008)]. RDM builds upon exploratory modeling [Bankes (1993), Kwakkel and Pruyt (2013)] by providing a systematic approach to identifying vulnerabilities.

Info-Gap analysis aims to quantify the maximum allowable deviation of deeply uncertain system factors that can be tolerated while still satisfying performance requirements [Ben-Haim (2004), Hipel and Ben-Haim (1999), Hall et al. (2012)]. Uncertain factors are sampled radially outward from a baseline (expected) future state of the world until a failure condition is reached; the distance from the baseline at which this occurs is termed α, or the “uncertainty horizon” [Hall et al. (2012), Korteling et al. (2013)]. Note that in this definition, there is no mention of probability distributions for the uncertainties. Rather, α defines the distance in the space of deeply uncertain factors between the baseline (expected) state of the world and the nearest state of the world in which the model predicts system failure. It assumes that a larger value of α implies the system is more resilient to perturbations in the deeply uncertain parameters. However, α fails to identify which specific uncertain factors, or combinations of factors, predict system failure. Recent examples of Info-Gap applications in water resources planning problems include Hine and Hall (2010) and Matrosov et al. (2013).

Decision Scaling, like other “bottom-up” approaches, inverts the decision making process. Rather than focusing on predictive distributions (derived, for example, by downscaling of Atmospheric-Ocean General Circulation Model (AOGCM) projections), Decision Scaling first aims to identify thresholds likely to trigger consequential system risks. The approach is a three-step process of (1) identifying key concerns and decision thresholds, (2) modeling the response to changing environmental conditions, and (3) estimating the relative probability of the critical environmental thresholds being crossed [Brown et al. (2012)]. Decision Scaling studies typically focus on uncertain climate factors, though recent work extends the approach to include hydroeconomic factors [Ghile et al. (2014), Lownsbery (2014)]. Decision Scaling's most significant difference from the other decision support frameworks is its assumption that the likelihoods associated with changes in temperature and precipitation can be inferred as subjective probabilities. The subjective probabilities are developed via expert evaluations of how SOWs attained from statistical weather generators relate to AOGCM projections [Brown et al. (2012)]. Decision Scaling has been most widely used as a discrete choice framework for choosing between pre-specified design alternatives [e.g., Moody and Brown (2013)] or as a vulnerability analysis to characterize the risks of existing systems [Ghile et al. (2014), Turner et al. (2014)].

In the Decision Scaling and Info-Gap frameworks, it is common to analyze a relatively small set of discrete decision alternatives that are pre-specified by stakeholders. This reflects a high degree of knowledge about system behavior under uncertainty, and may cause an analysis to be vulnerable to a significant status quo bias [Brill et al. (1990)]. Furthermore, pre-specified alternatives may overlook important trade-offs between conflicting objectives that reflect decision relevant performance requirements or tensions between stakeholders [Herman et al. (2014)]. RDM analyses can also suffer from these issues if the practitioner explores only a fixed set of alternatives. To overcome these challenges, Kasprzyk et al. (2013) propose Many-Objective Robust Decision Making (MORDM), in which alternatives are discovered via many-objective optimization in the projected future state of the world. MORDM supports constructive learning to improve decisions for complex, ill-defined environmental planning and management problems. This follows the framework of Many-Objective Visual Analytics (MOVA) [Woodruff et al. (2013)], a foundation for constructive decision aiding [Tsoukias (2008), Roy (1999)] in which problem framing is performed interactively with stakeholder feedback.