Ecological Modelling 343 (2017) 54–67

Contents lists available at ScienceDirect

Ecological Modelling

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

Review

A guide to multi-objective optimization for ecological problems with an application to cackling goose management

Perry J. Williams[a,b,*], William L. Kendall[c]

[a] Colorado Cooperative Fish and Wildlife Research Unit, Department of Fish, Wildlife, and Conservation Biology, Colorado State University, Fort Collins, CO, United States
[b] Department of Statistics, Colorado State University, Fort Collins, CO, United States
[c] U.S. Geological Survey, Colorado Cooperative Fish and Wildlife Research Unit, Department of Fish, Wildlife, and Conservation Biology, Colorado State University, Fort Collins, CO, United States

A R T I C L E I N F O

Article history:
Received 15 May 2016
Received in revised form 3 September 2016
Accepted 12 October 2016
Available online 31 October 2016

Keywords:
Decision analysis
Decision theory
Dominated choice
Efficiency frontier
Multi-criteria decision making
Pareto frontier
Pareto optimality

A B S T R A C T

Choices in ecological research and management are the result of balancing multiple, often competing, objectives. *Multi-objective optimization* (MOO) is a formal decision-theoretic framework for solving multiple objective problems. MOO is used extensively in other fields including engineering, economics, and operations research. However, its application for solving ecological problems has been sparse, perhaps due to a lack of widespread understanding. Thus, our objective was to provide an accessible primer on MOO, including a review of methods common in other fields, a review of their application in ecology, and a demonstration to an applied resource management problem.

A large class of methods for solving MOO problems can be separated into two strategies: modelling preferences pre-optimization (the *a priori* strategy), or modelling preferences post-optimization (the *a posteriori* strategy). The *a priori* strategy requires describing preferences among objectives without knowledge of how preferences affect the resulting decision. In the *a posteriori* strategy, the decision maker simultaneously considers a set of solutions (the Pareto optimal set) and makes a choice based on the trade-offs observed in the set. We describe several methods for modelling preferences pre-optimization, including: the bounded objective function method, the lexicographic method, and the weighted-sum method. We discuss modelling preferences post-optimization through examination of the Pareto optimal set. We applied each MOO strategy to the natural resource management problem of selecting a population target for cackling goose (*Branta hutchinsii minima*) abundance. Cackling geese provide food security to Native Alaskan subsistence hunters in the goose's nesting area, but depredate crops on private agricultural fields in wintering areas. We developed objective functions to represent the competing objectives related to the cackling goose population target and identified an optimal solution first using the *a priori* strategy, and then by examining trade-offs in the Pareto set using the *a posteriori* strategy. We used four approaches for selecting a final solution within the *a posteriori* strategy; the most common optimal solution, the most robust optimal solution, and two solutions based on maximizing a restricted portion of the Pareto set. We discuss MOO with respect to natural resource management, but MOO is sufficiently general to cover any ecological problem that contains multiple competing objectives that can be quantified using objective functions.

* Corresponding author at: Colorado Cooperative Fish and Wildlife Research Unit, Department of Fish, Wildlife, and Conservation Biology, Colorado State University, Fort Collins, CO, United States.
E-mail address: perry.williams@colostate.edu (P.J. Williams).

http://dx.doi.org/10.1016/j.ecolmodel.2016.10.010
0304-3800/

Contents

1. Introduction

Ecological decisions that require balancing multiple objectives are pervasive. Examples include endangered species management (e.g., maximizing species persistence while minimizing cost; Maguire et al., 1987), managing harvested species (e.g., maximizing cumulative harvest while maintaining population objectives; Johnson et al., 1997), and choosing statistical models to infer ecological processes (i.e., maximizing model fit while minimizing model complexity; Williams, 2016). When a decision maker has multiple competing objectives, a solution that simultaneously optimizes each objective does not exist; improving one objective results in a trade-off from another. Solving multi-objective decision problems requires incorporating decision-maker preferences among objectives into the decision problem (either explicitly or implicitly) to reach a compromise solution.

The process of solving multiple objective problems generally consists of identifying or soliciting the objectives of the decision maker, identifying potential actions, assessing the potential actions (or the predicted outcome of the actions) with respect to each objective, and making a choice. Scientific investigation can be used to predict the result of potential actions, but science alone is insufficient to address competing objectives because incorporating preferences among objectives requires value-based judgment (Holland-Bartels and Pierce, 2011). Some actions might result in obtaining one objective, while other actions might obtain another objective. Ultimately, a decision maker can implement one management action (or suite of actions). Thus, how can we formally combine a quantification of objectives, with objective preferences to select a final, optimal decision? *Multi-objective optimization* (MOO) is a sub-field of multi-criteria decision making that provides a formal mathematical framework for explicitly incorporating objectives and objective preferences to evaluate decisions. In contrast to other multi-criteria decision making methods (e.g., multi-criteria decision analysis), MOO is well suited for handling many objectives, and many (potentially infinite) alternative actions.

We outline the MOO framework and describe two strategies for solving MOO problems. Each strategy incorporates objective preferences into the decision problem. The strategies differ in the order in which preferences are incorporated; the first strategy (the *a priori* strategy) incorporates preferences pre-optimization and the second (the *a posteriori* strategy) incorporates them post-optimization. To demonstrate an application of MOO, we apply both strategies to a common natural resource management problem: selecting a *population target* for an animal population that affects multiple stakeholders differently. We define the population target as the abundance of animals that resource managers aspire to obtain and maintain through time. Our resource management problem was motivated by the management of cackling geese (*Branta hutchinsii minima*) across their range. Cackling geese nest on the coastal plain of the Yukon Kuskokwim (YK) Delta, Alaska. They constitute an important food source for Native Alaskan subsistence hunters. Ecosystem stability, the satisfaction of recreational hunters, and other non-consumptive users also depend on them. In their wintering area in Oregon and Washington (primarily in the Willamette Valley), cackling geese congregate on private agricultural fields and eat crops, resulting in loss of agricultural yield for landowners. The population target for the past 20+ years was 250,000 birds, however, the population target is currently being debated. Thus, we examine selecting a population target in a MOO framework. Multi-objective optimization is general, spanning many disciplines, and strategies used to solve MOO problems provide a framework for making defensible, transparent choices for natural resource management, and ecological decisions in general.

2. The multi-objective optimization problem

Multi-objective optimization assumes a decision maker can quantify the value of a decision with respect to the decision maker's objectives. Examples of objectives that have been explicitly quantified in natural resource management include: minimizing the probability of extinction (Maguire et al., 1987; Ewen et al., 2015; Larkin et al., 2016), maximizing the expected cumulative harvest of a hunted species (Johnson et al., 1997), maximizing the probability of successful population establishment of re-introduced species (Converse et al., 2013), maximizing biodiversity (Arponen et al., 2005; van Teeffelen and Moilanen, 2008; van Teeffelen et al., 2008; Cabeza et al., 2010; Tsai et al., 2015), maximizing habitat suitability (Williams, 1998; Holzkämper et al., 2006; Groot et al., 2007; Zsuffa et al., 2014), and maximizing habitat protection (Kennedy et al., 2008). A function that quantifies the value of the potential actions θ from a set of possible choices of actions Θ relative to an objective is termed an objective function (Keeney and Raiffa, 1976; Williams et al., 2002). Objective functions are inherently subjective because they are used to quantify the aim or interest of a decision maker (Hennig and Kutlukaya, 2007; Williams and Hooten, 2016). For consistency with MOO literature, we denote the objective function using $f(\theta)$ (notation definitions are also reported in Table 1 for reference). Objective functions are synonymous with loss functions, utility functions, or reward functions described in other fields (Williams et al., 2002; Berger, 2013; Williams, 2016; Williams and Hooten, 2016). The set of actions a decision maker can consider (Θ) can be either discrete or continuous. A specific action in the set of Θ is represented by θ. The value of the objective function (or utility) for a specific action is represented by $f(\theta)$. When a decision maker has one objective to maximize, and the objective function is unimodal, the decision maker can simply choose the value for θ that optimizes the objective function $f(\theta)$ (Fig. 1A). Decisions become