Journal of Applied Ecology 2008, **45**, 181–192 doi: 10.1111/j.1365-2664.2007.01367.x

Informed multi-objective decision-making in environmental management using Pareto optimality

Maureen C. Kennedy[1]*, E. David Ford[2], Peter Singleton[3], Mark Finney[4] and James K. Agee[2]

[1]University of Washington, Quantitative Ecology and Resource Management, Box 352100, Seattle, WA 98195–2100, USA; [2]University of Washington, College of Forest Resources, Box 352100, Seattle, WA 98195–2100, USA; [3]USFS Pacific North-west Research Station, 1133 N. Western Ave, Wenatchee WA 98801, USA; and [4]Rocky Mountain Research Station, Fire Sciences Laboratory, 5775 Hwy 10 West, Missoula MT 59801, USA

Summary

1. Effective decision-making in environmental management requires the consideration of multiple objectives that may conflict. Common optimization methods use weights on the multiple objectives to aggregate them into a single value, neglecting valuable insight into the relationships among the objectives in the management problem.

2. We present a multi-objective optimization procedure that approximates the non-dominated Pareto frontier without the use of weightings, allowing for visualization of the trade-offs among objectives. The non-dominated Pareto frontier is approximated by the simultaneous optimization of a vector objective function; two vector objective functions are defined as non-dominated if improvement with respect to one objective is at the detriment of another objective.

3. We demonstrate the method with a case study for the optimum distribution of forest fuels treatments that reduce the impact of fire on a forest. The multiple objectives are to protect habitat of an endangered species, protect late successional forest reserves and minimize the total area treated. In the comparison of three optimization searches, the number of non-dominated solutions increases with the dimensions of the objective space, but with only two objectives the search is ineffective in minimizing fire impact in the different landscape types. Key challenges include the extensive computation time required to approximate the non-dominated set, and reducing the number of solutions that are analysed in detail.

4. *Synthesis and applications*. The multi-objective optimization program presented can be adapted to other environmental management problems, and easily incorporates a wide range of quantifiable objectives. This tool provides decision-makers with a set of alternatives that estimates the full range of trade-offs among multiple objectives and provides a common ground from which dialogue can come to an informed compromise and decision in environmental management problems.

Key-words: fire management, fuels management, multiple objectives, non-dominance, optimization, Pareto frontier, spatial structure

Introduction

MULTI-OBJECTIVE OPTIMIZATION AND ENVIRONMENTAL MANAGEMENT

Multiple criteria analysis and multi-objective optimization have been utilized to design decision support systems for a variety of environmental management test cases (e.g. Chen & Chang 1998; Erickson *et al.* 2002; Seely *et al.* 2004; Xevi & Khan 2005; Chen *et al.* 2006; Higgs 2006; Linkov *et al.* 2006; Stirn 2006), including water management, contaminated sediments, location of waste facilities, air quality monitoring networks and forest management. Environmental management is a multi-objective problem because there are typically a number of objectives to be optimized and there are possible management actions that can be implemented; the potential effect of the management action is linked to the objectives through a model that quantifies the consequences of alternative actions. The challenge is in evaluating the performance of the action relative to the multiple objectives.

*Correspondence author. University of Washington, Quantitative Ecology and Resource Management. Box 352100 Seattle, WA 98195–2100, USA. Fax: 206 5433254. E-mail: mkenn@u.washington.edu

In many studies, the analysts solicit preferences from decision-makers that indicate the relative importance of each of the multiple objectives. Weightings for each of these are estimated from the preferences, and used to convert the objectives into a scalar-valued function that is optimized (Kim & Smith 2005). While this reduces the optimization problem and eases computation of the optimal solution, it also introduces a form of uncertainty in the optimal solution due to decisions required to quantify appropriate weights (Schoemaker & Waid 1982). The relative importance of the objectives may vary with which solicitation method is used, which decision-makers are asked, and even when preferences are solicited from a decision-maker at different times. Furthermore, the ranking of preferences can be sensitive to the weighting values (Hyde *et al.* 2005), producing variations in the preferred solution with small changes in the weights.

The idea that there is a single solution to a multi-objective optimization problem is a fallacy of the weighted sum approach. This single solution does not exist because preferences can change and are not certain themselves, and there is no single answer that minimizes all of the objectives simultaneously. A method that yields optimal solutions regardless of weights is preferable. We present a method for multi-objective optimization based on approximating the non-dominated Pareto frontier (Cohon 1978; Fig. 1; Table 1) of decision variables. The Italian economist Vilfredo Pareto, studying economic efficiency and utility, originated the concept of Pareto optimality (Cirillo 1979). It is used extensively in economics, and has been adapted for engineering and design (e.g. Statnikov & Matusov 1995). Pareto optimality has been used more recently, but sparsely, in ecological model assessment (Reynolds & Ford 1999) and in a study of optimal foraging (Rothley *et al.* 1997).

The non-dominated Pareto frontier specifies the groups of decision variable values (the optimal set) that optimize the management objectives through simultaneous optimization of a vector-valued objective function. This method provides decision-makers with a range of multiple optimal alternatives *before* the relative importance of the objectives are specified, thus removing the consequences of uncertainty in the weighting of objectives. It preserves the role of the decision-maker in setting preferences and re-examining both them and the proposed management model. Two major challenges can prohibit the exploration of a full multi-objective optimization problem: (a) an efficient procedure to approximate the non-dominated Pareto frontier, and (b) effective exploration of the optimization results to allow for further analysis by decision-makers. We demonstrate how an evolutionary computation algorithm can approximate the non-dominated Pareto frontier and give a four-stage process for summarizing and presenting results.

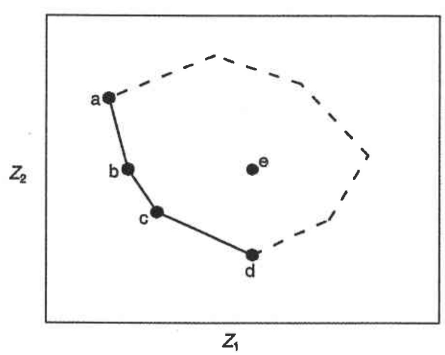

Fig. 1. Non-dominance example for the minimization of 2 objectives (Z_1, Z_2). The dashed line encloses the feasible space. Any solutions along the solid (a,b,c,d) line belong to the non-dominated Pareto frontier, where any improvement in Z_1 occurs at a cost to Z_2, and vice versa. The value (e) is an example of a dominated solution in the feasible space.

Table 1. 4-objective non-dominance example for FUELSOLVE case study. Each row lists objective values for unique fuels treatment distributions; the optimization problem is to minimize the objectives. In this set of 5, Fuels treatment 1–4 are not dominated by any of the other 5 results. Although in terms of dominance Fuels treatment 5 is indistinguishable from 3 and 4, it is dominated by Fuels treatment 1 and 2 (all four objectives are greater for 5 than for 1 and 2), and is not part of the non-dominated set

	LSR	Owl circle	Owl Core	Area treated
Fuels treatment 1	24·80	34·68	0·28	4518·18
Fuels treatment 2	1·08	86·56	1·76	5913·00
Fuels treatment 3	0·36	243·24	39·08	6894·72
Fuels treatment 4	53·60	22·64	0·10	4842·18
Fuels treatment 5	30·24	350·69	35·16	7864·00

MULTI-OBJECTIVE OPTIMIZATION USING PARETO OPTIMALITY

Multi-objective optimization using the concept of non-dominance (Cohon 1978) requires approximation of the Pareto frontier, i.e. the set of all non-dominated solutions. A solution is defined to be non-dominated if there exists no other feasible solution that will give an improvement in one objective without a subsequent degradation in at least one other objective (Cohon 1978; p. 70; Fig. 1; Table 1). The optimization can be formally presented as follows:

maximize (or minimize)

$$Z(x_1, x_2, \dots, x_n) = [Z_1(x_1, x_2, \dots, x_n), Z_2(x_1, x_2, \dots, x_n), \dots, Z_p(x_1, x_2, \dots, x_n)] \quad \text{eqn 1}$$

$$X = \{(x_1, \dots, x_n \mid x_j \text{ in } \mathbb{R}); j = 1, 2, \dots, n; \quad \text{eqn 2}$$

n decision variables (X), *p* objectives (Cohon 1978).

In the optimization method, vector objective values are calculated for feasible combinations of values for the decision variables (x_j), and then evaluated for the relative dominance status of each vector of objectives (Z). If the entire feasible space is evaluated and objective values are deterministic then the entire non-dominated Pareto frontier is calculated (e.g. Fig. 1). With the exception of rare cases and given the size of environmental management problems, calculation of the entire feasible space to determine the non-dominated Pareto frontier is computationally prohibitive. Search processes such as evolutionary algorithms are utilized to converge to an approximation of the non-dominated Pareto frontier; these techniques are used extensively in engineering and economics (Deb 2001).

The general optimization program can be conducted in the following steps:

1. Define and quantify the objectives.
2. Define the optimization decision variables.
3. Integrate search algorithm with the calculation of objectives.
4. Present results for post-processing.

We illustrate this method with a forest management example for the distribution of fuels reduction treatments in a watershed in the eastern Cascade Mountains, Washington State, USA, with the goal to reduce possible wildfire impact to multiple ecological values. This is a spatially complex environmental management optimization problem where management has multiple objectives.

BACKGROUND TO THE CASE STUDY

A century of fire exclusion, grazing, and selective removal of large, fire-tolerant trees in the dry forests of eastern Washington has resulted in forests of much higher density that are more prone to stand-replacing wildfires than those which occurred historically in the area (Arno & Brown 1991; Agee 1993; Wright & Agee 2004). Yet, these now dense, multi-canopied forests also serve biodiversity goals, including habitat for endangered species (Lehmkuhl *et al.* 2007). Federal legislation and regulation has defined both of these competing objectives (fuels reduction and habitat protection) without reference to the other.

The National Fire Plan, Healthy Forests Initiative & Healthy Forest Restoration Act of 2003 directs United States Federal agencies to treat these fire-prone forests by thinning and prescribed fire to reduce the risk of stand-destroying wildfire (O'Laughlin 2005). These documents all call for prioritizing hazardous fuels reduction, with an emphasis on rehabilitation and restoration of impacted forests. Yet the recovery plan for the northern spotted owl (and many other species defined as Federally threatened under the Endangered Species Act), known as the North west Forest Plan, provides protection for the owl through two mechanisms: an unmanaged buffer around each known nest site, and large protected areas known as Late Successional Reserves (LSRs) where management direction is to provide old growth, late successional conditions. This is in conflict with efforts to reduce fuels in the same landscape.

In addition, the structure of fuel treatments and their spatial distribution in a landscape are not well understood (Agee & Skinner 2005), and the effect of fuel treatments on landscape types of high ecological value (e.g. wildlife habitats and populations, old-growth forests; Huntzinger 2003; Lee & Irwin 2005) must also be considered and balanced against the risk of losing such landscape types to catastrophic fire. Policies recognize the imperative to consider ecological consequences of fuels reduction programs and subsequent fire risk (Franklin & Agee 2003). What choice should be made about which section of a forest is to be treated by thinning and prescribed fire in order to reduce the potential spread of a wildfire? The choice requires distributing the minimum area of treatment that will minimize possible fire damage to areas of special interest, but there is no single distribution that can achieve these minima simultaneously.

This problem is similar to many environmental management problems in that it has multiple objectives that are spatially complex and potentially conflicting. Furthermore, the context of the decision is framed by multiple policy documents and the objectives described by these documents are quantified in different currencies and are not amenable to simple addition or other combinations through weightings.

FUELSOLVE CASE STUDY

The case study is conducted through the US Forest Service project called FUELSOLVE, intended to integrate ecological values into the fuels and fire management decision-making process. We chose the 23162 ha Mission Creek watershed in the Okanogan and Wenatchee National Forests of Washington State to provide a real-world landscape for analysis. The first set of objectives this project considers are the ecological values (owl habitat, late successional reserves), whereas the second set of objectives measure the cost of treatment (area assigned treatment). In the next section, the method for the four stages are described, and the results given for the post-processing example.

Methods

MULTI-OBJECTIVE OPTIMIZATION OF A RESOURCE MANAGEMENT PROBLEM

1. Define the objectives

(a) Landscape types. Owl activity centres have been recorded for the Mission Creek Watershed, reflecting observed activity in the years 2000–04. These centres are surrounded by two buffers of increasing size: cores (radius = 1127 m around activity centre) and circles (radius = 2931 m around activity centre); cores are given special protection in the region. LSRs are areas designated as old-growth or set aside as potential old-growth, and the study area includes both LSRs and owl activity centres (Fig. 2), with some overlap between landscape types. The fire impact objectives have effect on owl circles, owl cores and LSRs. Cost of treatment is included as an objective because an optimal solution for reducing fire spread would be to treat the entire study area, but this solution would have negative ecological and economic consequences.

(b) Quantify objectives. The fire impact objectives are estimated by the results of fire spread simulated on each treated landscape.

Requirements for fire spread simulations

To run the fire spread models, multiple GIS layers are required (Table 2) and the landscape data were gathered from multiple

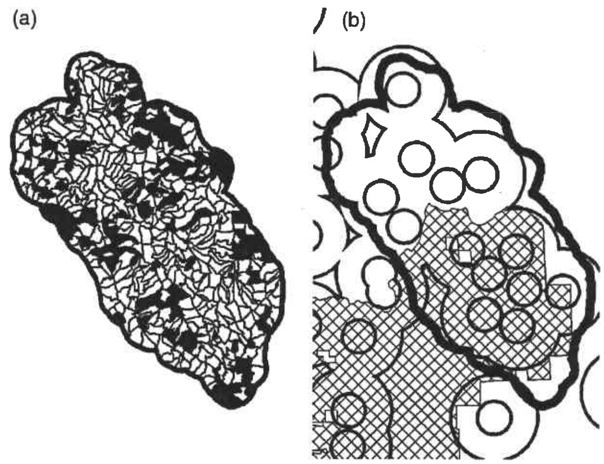

Fig. 2. (a) Study area and window for the Mission Creek Watershed. White polygons are the discrete landscape areas available for treatment. Black polygons are not available for treatment. (b) Valuable landscape types used as objectives in the optimization. The smaller circles are owl cores (1127 m buffer), the larger circles are owl circles (2931 m buffer). The shaded grid near the bottom is Late Succession Reserve (LSR).

Table 2. Landscape attributes modified for fuel treatment

Attribute	Fuels treatment value
Surface fuel model	34 (TL1)
Canopy cover	50%
Canopy height	30 m
Canopy base height	5·4 m
Canopy bulk density	0·03 kg m^{-3}

sources using ArcMap (ESRI 2004). Using the 40-model classification system (Scott & Burgan 2005), surface fuel model types were assigned to the Mission Creek watershed with shrub cover, elevation, aspect, and canopy cover. Weather files were collected from the Swauk weather station during a period from 1994 to 2003, and represented by dates ranging from 15 July to 31 August (the fire season). Only extreme weather conditions were used, under which wildfires escape initial attack and grow rapidly in size, in conjunction with a constant wind speed (10 mph) and wind direction (south-west). A study area was defined for the watershed (23162·76 ha) and a rectangular window was drawn around the study area to be included in the fire spread simulations (Fig. 2). The landscape is reduced to an ASCII grid file with 277 rows and 215 columns. Each cell in the grid represents a 90 m × 90 m pixel that contains the landscape attribute and each attribute has a separate file. The grid is orientated such that the cell (1,1) is the north-east corner of the landscape. In addition, the fire spread simulations require a particular landscape file format (*.lcp) and this was generated for the base Mission Creek Landscape (MissionCreek.lcp).

For the calculation of fire spread we utilize the command-line version of the gui program *flammap* (Finney 2002, 2003), which has two subprograms: *flammap* calculates the fire spread variables, whereas *randig* calculates realized spread based on random ignition points. The fire spread simulations generate k ignitions placed randomly on the landscape. In the algorithm the minimum travel time of the fire to a pixel is calculated, and if that time is less than the specified fire duration the fire is determined to have spread to that pixel. Each of the k ignitions is spread independently on the unburned treated landscape, so that the fires do not interact. Given that the actual location of ignitions is difficult to predict, the goal is to characterize possible fire behaviour through multiple ignitions assigned randomly across the landscape. The more ignitions simulated, the greater the coverage of the landscape; this, however, is at great computational cost. This requires that the objective function incorporates the multiple ignitions in a manner that reduces the variability in the objective function, without undue computational burden. This is evaluated below.

To determine an appropriate fire duration for the optimization problem, preliminary fires were spread on the untreated landscape; at 4·5 days (6480 min) the mean area burned was 7500 ha. This is a reasonable size relative to other major fires in the area and all simulations were conducted with duration of 6480 min.

Fire impact

The fire spread simulation produces new landscape files whose attributes are the arrival time of the fire and the proportion (q) of the k fires that reach each pixel in the landscape (q_{rl}; $r = 1, 2, \ldots, 277$; $l = 1, 2, \ldots, 215$; $q = \{0, 1/k, 2/k, \ldots, 1\}$). For the fire impact objective function, the q-value is summed across all pixels that are also of a target landscape type (*fire_sum*, e.g. sum the q-values for all pixels that are in LSR). This is divided by k for a per ignition value.

$$\textit{fire\_sum}:\ Z_p(x_1, x_2, \ldots, x_n) = \frac{1}{k}\sum_{r_p l_p} q_{r_p l_p} \qquad \text{eqn 3}$$

for $p = \{1, 2, 3\}$, r_p indexes the rows relevant to objective p, and l_p indexes the columns relevant to objective p in the landscape grid file, and n is the number of units searched for treatment. The r and l values give the coordinates for each of the objective areas in the study landscape.

To determine the value for k, the variation in the per ignition *fire_sum* value for fires simulated on the untreated landscape declines steeply between 3 and 5 ignitions, then levels off at ignitions greater than 5 (Fig. 3). Given the computational cost of simulating multiple ignitions, $k = 5$ ignitions were used for all optimizations.

The cost objective (area treated) is the value of area treated under the current treatment alternative:

$$Z_p(x_1, x_2, \ldots, x_n) = 0.81^* \sum_{m=1}^{n} \sum_{r_m l_m} I_{r_m l_m}$$
$$I_{r_m l_m} \begin{cases} 1 \text{ if } x_m = 1 \text{ and } pixel_{r_m l_m} \text{ is treated} \\ 0 \text{ otherwise} \end{cases} \qquad \text{eqn 4}$$

where r_m, l_m index the row and columns for the pixels of treatable unit x_m, and 0·81 is the hectare per pixel. This can be calculated for the total area, or for the area treated of a specific landscape type. In order to demonstrate the method, we chose the simplest representation of cost of treatment as area treated. Actual cost has more complex components that depend, for example, on contiguous size, location and accessibility. When quantified these components could be included in the search.

2. Define the optimization decision variables

In the case study, there is one type of decision variable, whether or not an area of forest should be treated to reduce the potential

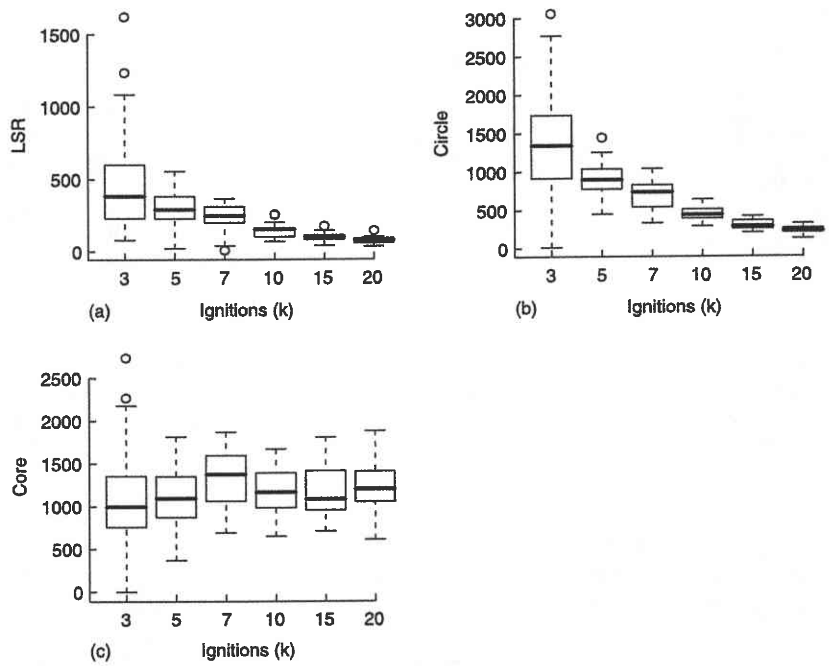

Fig. 3. Variation in the per ignition value of *fire_sum* with increasing number of ignitions for (a) LSR, (b) owl circle, and (c) owl core. Twenty simulations were conducted for each number of ignitions.

spread of fire. The study area was divided into 509 polygons of similar vegetation attributes to serve as discrete treatable areas (Fig. 2a; polygon size ranged from 19 to 95 ha). Of these, if > 30% of the pixels in a polygon contained > 40% canopy cover, that polygon was considered available to treatment. This left 413 polygons to be searched for whether treatment is assigned; each decision variable x_j ($j = 1, 2, \ldots, 413$) in the optimization problem is a single potentially treated polygon.

The single fuel treatment is defined by five landscape attributes of a thin and burn fuel treatment (Agee & Skinner 2005; Table 2). In the model, if a treatable polygon is assigned fuel treatment, then all pixels in the polygon are modified to match the treatment attributes (Table 2). The possible values for the decision variables are:

$x_j = \{0, 1\}$ for $j = 1, 2, \ldots, 413$, where 0 = no treatment, 1 = assigned treatment. eqn 5

The optimization algorithm searches combinations of fuel treatment assignments that yield the spatial distribution of treatments that minimize the fire impact and treatment cost (area).

3. Integrate search algorithm with the calculation of objectives

The program used to solve the multi-objective problem is a multi-objective evolutionary algorithm, *Pareto_evolve* (see Appendix S1 in Supplementary material), first developed for ecological process model assessment by Reynolds & Ford (1999). The optimization algorithm initializes by randomly generating a large number of decision variable combinations (i.e. spatial allocation of treatments). There are typically 100 individual decision variable combinations, and this set is called a **population**. Once initialized, the algorithm has two major stages: evaluation and breeding.

In the evaluation stage, the effectiveness of each individual in the current population is evaluated for how well it achieves the optimization objectives. The individuals are ranked, where the non-dominated individuals of the current population are assigned Rank 1. They are assigned a fitness based on their non-dominated ranking and a measure that reduces the fitness of individuals that are similar and increases the fitness of unique individuals, and individuals are then chosen randomly by their fitness to enter the breeding stage of the algorithm.

The individuals chosen to enter the breeding stage are called **parents** of the next population that will be evaluated. The next population is produced either through **mutation** of the parent vector (small changes in randomly chosen decision variable values, e.g. change from a zero to a 1 or vice versa), or through **cross-over** between two vectors (decision variable values are exchanged between the two parents, e.g. either the variable value doesn't change or it switches to or from 0 and 1, but it retains neighbouring values in the vector). This new population of individuals then enters the evaluation stage. Each cycle of evaluation and breeding is termed a **generation**. The algorithm is terminated either when it reaches a specified maximum generation or it converges to a unique optimum.

Pareto_evolve is a generic optimization program that the user modifies to coordinate with the user-supplied evaluation code (i.e. the fire spread simulations). It has been utilized to assess various process-based models, including stand development (Reynolds & Ford 1999), competition (Turley 2001) and shoot extension (Komuro *et al.* 2006). To configure the algorithm for a specific optimization problem the user must define the decision variable search ranges (e.g. 0,1; step 2 above), and objective target values (e.g. 0 for minimization; step 1 above). The final output of the search is the approximated non-dominated Pareto frontier, which is presented in step 4.

4. Presentation of results for post-processing

In order to evaluate the effectiveness of the algorithm and to compare the results of different objective function combinations, three searches were conducted (Table 3). For each search the population size is 100 individuals for each generation. Given the constraints on computing power, the most generations conducted for a search is 400.

In the first search, 6 objective functions were optimized simultaneously (Search 1): fire impact on LSR, owl core and owl circle and area treated in the total study area, in LSR and in owl core. This