A Tool for Decision Support in Dynamic Conservation Management

Ilija Bogunovic[1], Andreas Krause[1], and Sarah J. Converse[2]

[1] ETH Zürich, Zürich, Switzerland
[2] USGS Patuxent WRC,Laurel, MD, USA

Abstract. We present an algorithm for generating diverse recommendations in the reserve design problem. We also present a tool for decision support in dynamic conservation management that incorporates the recommendation algorithm.

1 Introduction

Conservation management is a resource allocation problem under uncertainty prototypical for many problems in computational sustainability. One of the main tasks of conservation management is to determine how to recommend patches of land for conservation in such a way that the long-term persistence of endangered species on this land is ensured. Making such a long-term decision is a challenging problem since it depends on various aspects such as uncertainty in the species distribution, their dynamics, as well as the availability of patches and financial resources over time. Further, recommending an optimal reserve, a set of patches that maximizes survival probability of endangered species over time, is difficult due to combinatorial number of possibilities, and NP-hardness of the optimization. As shown by [1], however, natural formulations of the problem are submodular, and a simple greedy algorithm can be used to obtain a near optimal solution. However, finding only one solution leaves the conservation management community without possibility to explore alternatives.

In this study, we tackle this problem using the randomized *Best-K* algorithm. This algorithm is able to produce diverse reserve recommendations which are near-optimal.

More, we develop a tool (see Fig. 1) for decision support in dynamic conservation management which incorporates the Best-K algorithm. This tool can significantly ease the process of decision making in conservation management by allowing managers to choose between different conservation strategies and to better utilize their available resources.

2 Related Work

The problem of protecting rare species by recommending patches of land for conservation has been studied by Krause et al. [1] They define the objective function

$f(R) = \sum_{i \in I} f^{(i)}(R)$ where $f^{(i)} : 2^P \to \mathbb{R}$ quantifies the probability that species i is still present at a patch in the reserve $R \in P$ after some prediction horizon T (e.g. after 50 years). Hence the conservation planning problem the authors address is to select a set R of patches from the set of feasible patches P as a reserve:

$$R^* = \operatorname*{arg\,max}_{R:c(R) \leq b} f(R), \tag{1}$$

that maximizes the persistence probability while respecting a budget constraint b on the total cost $c(R) = \Sigma_{p \in R} c(p)$. More, by proving that the objective function is a monotonic, submodular function, they suggest a simple greedy algorithm that produces a near-optimal reserve. However, as a deterministic approach, it produces only a single solution. Because of the real-world nature of the conservation planning problem some reserves could be more appealing than others. Here, we propose a randomized variant of the greedy algorithm that allows users to explore a diverse set of alternatives.

There are several powerful tools available for conservation planning, including Marxan [2] and Zonation [3]. However, none of those tools currently implements complex patch dynamics models of species persistence. Here we propose a tool that supports such models, while also providing (near-)real-time optimization performance. We demonstrate it for a case study of protecting three Federal Candidate taxa inhabitating a remnant prairie ecosystem in the South Puget Sound: the Taylor's checkerspot (Euphydryas editha taylori), Mazama pocket gopher (Thomomys mazama) and streaked horned lark (Eremophila alpestris strigata).

3 Approach

The Best-K algorithm approximately solves problem (1) while recommending different reserves in every run. It replaces deterministic greedy selection by non-uniform random selection. In every step, the algorithm efficiently finds K patches which have the highest marginal benefit, i.e. the highest ratio between the gain in the objective function maximization and patch's cost. One patch out of these K patches is picked non-uniformly at random and added to the reserve. The probability of adding patch p to reserve R is given by

$$\Pr[p] \propto \exp\left(L \cdot \frac{f(R \bigcup \{p\}) - f(R)}{c(p)}\right).$$

The desired reserve similarity over multiple runs can be controlled via an input parameter L: $L = 0$ corresponds to uniformly random selection, thus significantly reducing similarity between recommended reserves. In the limit $L \to \infty$, the algorithm reduces to deterministic (greedy) selection. Due to submodularity of f, guarantees about the performance of *Best-K* can be stated, which degrade for small values of L.

Another constraint which is not included in the optimization problem (1) is a constraint on the target persistence probability per species. The value of

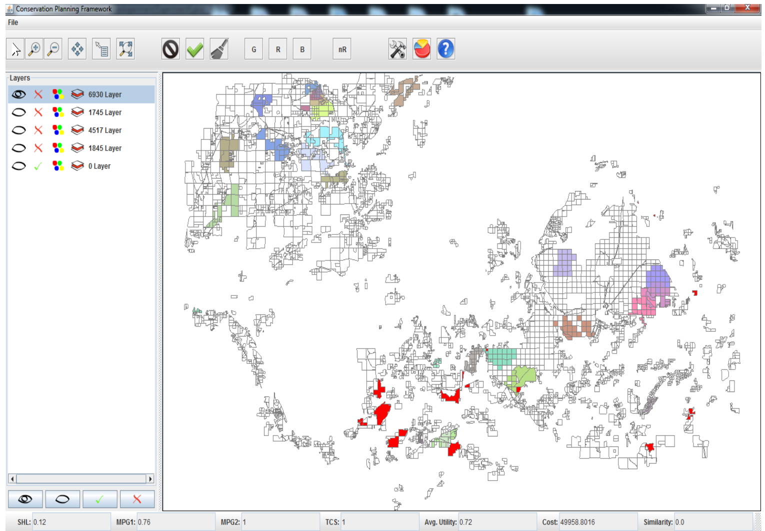

Fig. 1. Tool GUI: A region map is partitioned into parcels (white cells) which are grouped into contiguous patches (colored regions). Currently, one possible reserve consisting of recommended patches is showed.

this constraint can be specified as the input parameter of the algorithm. Once for all species the target persistance probability is achieved or a given budget is exceeded, the algorithm returns the recommended patches.

More, the algorithm can adapt its reserve recommendation depending on the input sets such as the set of patches that must be included in the reserve, the set of patches that cannot be included and the set of patches that are already managed by the conservation planning organization. These sets are explained in more details in section 4.

4 Decision Support Tool

The main goal of our decision support tool is to allow interactive optimization in (near-) real-time, so that conservation managers can explore their decisions in the space of possible recommendations. In order to enable this, we separate survival simulations from the decision support tool. In the simulation part, we generate 5500 candidate patches by picking a random parcel as a seed and then growing the patch (i.e. a set of contiguous parcels) up to a certain size. To model population dynamic among the parcels we use a patch dynamic model (see [1]) which predicts whether species i is present or absent on certain parcel at time t. We say that species i survives on a patch p if it is present on at least one of its parcels at time t. Next, we run a series of simulations based on this model. Simulation results are pairs (p, S) where S is a set of simulations in which species i survives on a patch p after a period of T years.

Once we were able to model the population dynamics of the species, the decision support tool allows us to solve optimization problems in real-time, e.g. to recommend a reserve when the budget has changed, patch availability has changed etc. The tool generates recommendations in just a few seconds per optimization problem instance.

When specialized to our case study, the decision support tool presents a map of the South Puget Sound Region. The land on this map is divided into parcels. The tool integrates two algorithms for a reserve recommendation: the greedy [1] and the Best-K algorithm. It also has an option of recommending various different reserves and presenting them in different application layers. The tool settings window allows conservation managers to set constraints on the maximal available budget and a maximal species survival probability. Moreover they can also set various algorithm input parameters, e.g. the numbers K and L in the Best-K algorithm. Some of the parcels may be already owned by a conservation organization. We build an additional patch that we refer to as *managed* and which contains all the owned parcels. This patch is always a part of recommended reserves. Further on, the tool allows additional management constraints to be imposed:

- *Parcel availability constraint*: when some parcel is disallowed, the algorithms will only recommended reserves so that no selected patch contains this parcel.
- *Patch selection constraint*: when a patch is selected, the recommended reserve must contain this patch.

Once these constraints are set, one of the algorithms can be used to recommend the rest of patches that should be a part of the reserve. This tool feature introduces additional management functions that have an influence on how patches are recommended. The tool also provides: the parcel description option, a real-time insight into measures such as the single species survival probabilities, the average species survival probability, the cost of the currently presented reserve and, if multiple alternative solutions have been generated, their similarity. Lastly, the tool can export a report containing statistics and other details of all currently recommended reserves.

We believe this tool can facilitate decision support in conservation management by allowing conservation management community to make reserve recommendations by solving real-time optimization problems.

References

[1] D. Golovin, A. Krause, B. Gardner, S. J. Converse, and S. Morey, "Dynamic resource allocation in conservation planning," in *AAAI*, 2011.

[2] I. Ball, H. Possingham, and W. M., "Spatial conservation prioritisation: Quantitative methods and computational tools." *Oxford University Press. chapter Marxan and relatives: Software for spatial conservation prioritisation*, 2009.

[3] A. Moilanen and H. Kujala, "Zonation: Spatial conservation planning framework and software," *www.helsinki.fi/bioscience/ConsPlan*, 2008.