Forest Policy and Economics 83 (2017) 107–120

Contents lists available at ScienceDirect

Forest Policy and Economics

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

Spatial interactions and optimal forest management on a fire-threatened landscape

Christopher J. Lauer[a,*], Claire A. Montgomery[b], Thomas G. Dietterich[c]

[a] Oregon State University, Applied Economics, 213 Ballard Extension Hall, Corvallis, OR 97331, United States
[b] Oregon State University, Forest Engineering, Resources & Management, 218 Snell Hall, Corvallis, OR 97331, United States
[c] Oregon State University, Electrical Engineering and Computer Science, 2067 Kelly Engineering Center, Corvallis, OR 97331, United States

ARTICLE INFO

Keywords:
Wildland fire
Spatial
Ecological disturbance
Risk
Approximate dynamic programming
Reinforcement learning
Forestry

ABSTRACT

Forest management in the face of fire risk is a challenging problem because fire spreads across a landscape and because its occurrence is unpredictable. Accounting for the existence of stochastic events that generate spatial interactions in the context of a dynamic decision process is crucial for determining optimal management. This paper demonstrates a method for incorporating spatial information and interactions into management decisions made over time. A machine learning technique called approximate dynamic programming is applied to determine the optimal timing and location of fuel treatments and timber harvests for a fire-threatened landscape. Larger net present values can be achieved using policies that explicitly consider evolving spatial interactions created by fire spread, compared to policies that ignore the spatial dimension of the inter-temporal optimization problem.

1. Introduction

Forest management is a dynamic problem; actions taken today have important consequences for the future value of forest land. Determining optimal management is further complicated by unpredictable ecological disturbances; one important disturbance on many forested landscapes is wildfire. While fire is a natural process that can be vital to the health of forest ecosystems, it can threaten values on the landscape such as timber, homes in the wildland urban interface, watershed health, air quality, and wildlife habitat. Significant amounts of money and resources have been devoted to the task of eliminating damaging wildfire. According to the 2015 Fire Budget Report, produced by the USDA Forest Service, > 50% of the Forest Service's 2015 operating budget was devoted to fire-related activities, compared to just 16% in 1995. The USDA Forest Service spent more than $1.7 billion on fire suppression costs alone in 2015 and has spent more than $1 billion in eight of the last ten years on fire suppression activity (USDA, 2015).

Value on a forest landscape can be diminished, or destroyed, by unpredictable events like fire, and the extent of the damage is at least partially outside of manager control. Routledge (1980) and Reed (1984) demonstrated that the optimal rotation age for timber harvest determined by Faustmann (1968) can be adjusted to account for the possibility of an unpredictable natural disaster damaging or destroying a stand. They showed that the optimal rotation age decreases as the probability of stand destruction increases. Reed built upon his model in subsequent papers to demonstrate how optimal management will be affected if fire arrival rate is a function of stand characteristics. He also explored how to determine the optimal schedule of investment in fire protection, such as fuel treatment or fire-fighting infrastructure (Reed, 1989, 1993). Many authors build upon Reed's insights; examples include Amacher et al. (2005) and Garcia-Gonzalo et al. (2014) who look at how fuel treatment and silivicultural interventions affect optimal rotation age for fire threatened forest stands. Daigneault et al. (2010) examine how carbon sequestration is affected by manager response to fire risk.

An important shortcoming of these models is that they focus on management at the stand level and do not account for fire's ability to spread between stands. Because fire can travel large distances across a landscape, fire risk on an individual stand is a function of the condition of the entire landscape. Several authors have created models to account for fire behavior that includes spatial interactions. Wei et al. (2008) separated fire arrival probability into ignition and spread probabilities. These probabilities account for varying ecological factors such as slope and wind direction. They used the conditional spread probabilities to find locations of fuel treatments in the current time period to minimize expected loss to fire on the modeled landscape. Ager et al. (2010) used repeated simulation to compare the damage probabilities for structures in the wildland-urban interface under different fuel management

* Corresponding author.
E-mail addresses: lauerc@oregonstate.edu (C.J. Lauer), claire.montgomery@oregonstate.edu (C.A. Montgomery), tgd@cs.orst.edu (T.G. Dietterich).

http://dx.doi.org/10.1016/j.forpol.2017.07.006
Received 22 December 2016; Received in revised form 5 May 2017; Accepted 11 July 2017
Available online 17 July 2017
1389-9341/ © 2017 Elsevier B.V. All rights reserved.

strategies when fire spreads. These models are static because they consider only treatment in the current period and the expected value of the treated landscape. Chung et al. (2013) developed a model that places fuel treatments over two time periods to minimize expected loss. The vegetation evolves over time, but no fire actually occurs. While this model is intertemporal, it is not dynamic because it does not account for how optimal management will adjust on a post-fire landscape, should fire actually occur.

Stochastic dynamic programing is a natural method for determining the optimal management of a fire-threatened forest landscape. Actions are selected to maximize the current period value plus the expected future value of the treated landscape given that optimal choices are made in future periods as new knowledge becomes available (e.g., stochastic events are realized). This optimization strategy depends on the principle of optimality (Bellman, 1957) which is the assumption that no matter what the current action is, actions in all future states will be optimally chosen. Solving a stochastic dynamic programming problem requires complete enumeration of all possible outcomes in each time period; this can make even very simple problems intractably large – the so-called "curse of dimensionality." Konoshima et al. (2008, 2010) was able to formulate and solve the optimal fuel treatment and harvest problem as a stochastic dynamic problem by considering only 2-periods, 7 management units, 4 stand age classes, 4 treatment options, and 2 weather conditions. They demonstrated that landowners will try to protect on-site timber values by shortening rotations, as suggested by Reed (1984), but will try to protect adjacent timber values by postponing harvest to avoid high spread rates associated with young stands. This study demonstrated the importance of accounting for stochasticity and spatial interactions in a dynamic decision framework, but its practical usefulness for decision-making and policy analysis is limited.

Both the spatial and intertemporal aspects of solving for the optimal management of a fire-threatened landscape contribute to the curse of dimensionality. Heuristic solution methods, such as simulated annealing or tabu search, are widely used to identify management that is approximately optimal in a given time period for large landscapes with spatial interactions. Chung et al. (2013) is one example. See also Bettinger et al. (2003). But these methods alone are not sufficient for solving dynamic problems in which managers respond to new information as it becomes available. For that, we need an approximate approach to solving dynamic programming problems.

In this study, we develop a method for incorporating spatial interactions into a dynamic decision process that accounts for the stochastic nature of fire by using a machine learning technique known as approximate dynamic programming (Sutton and Barto, 1998; Powell, 2007, 2009). We applied our method to model optimal timing and placement of timber harvest and fuel treatment on a 64-unit landscape that we parameterized to represent the ecological conditions of southwest Oregon. The value functions we estimate provide a way to model the expected benefits, costs, and externalities associated with different management actions which have uncertain consequences in multiple locations on the landscape. The value function implies an optimal policy for management actions. We evaluate the effectiveness of our method by comparing simulated outcomes over 150-year time horizons to outcomes generated using Reed and Faustmann timber harvest rotations and to a rule-of-thumb approach to placing fuel treatments. The policies generated by our model lead to higher average net present values (NPV) on the landscape compared to the benchmark policies. In the concluding section, we highlight the relevance of our key findings, describe model limitations, and discuss plans to extend the basic model to address relevant questions in forest management.

2. Model

A bio-economic model can provide insight into factors that drive landowner behavior and can be used to determine an optimal set of actions for a decision-making agent when landscape conditions give rise to fire risk that affects stand value. Our model accounts for the financial incentives faced by the agent. It could also be specified to include non-market incentives. The ecological processes that determine the evolution of the landscape can be thought of as a constraint on the agent's actions. Vegetation evolves and stochastic disturbances (in this case, fire) occur over time so that future costs and revenues depend on actions taken today and, likewise, the present value depends on those future rewards. Finally, because fire spreads across the landscape creating interactions that affect stand value, our model is explicitly spatial.

2.1. Markov decision process and Bellman's equation

We represent the economic and ecological components of this problem as a Markov Decision Process (MDP) (Puterman, 1994) that has five different components.

1. A set of possible *States*, which depends on the attributes of the individual stands contained by the landscape. The state S_t describes the conditions of the landscape at time *t*.
2. A set of *Actions* that describes what a land manager can do. In our setting, the overall action at time *t* is a vector x_t of the management activities applied to each stand in the landscape. For each stand, there are four possible management activities: harvest timber (clearcut), treat fuel to reduce fire risk, implement both activities, or do nothing; in each time-step one of these four options must be chosen for each stand.
3. A *Reward Function*, $C(S_t, x_t)$ that describes the immediate (i.e. current period) costs or revenues associated with a particular action for a particular state.
4. A *State Transition Model*, $S_{t+1} = S'(S_t, x_t, W_t)$, that describes how the state evolves over time. The transition from S_t to S_{t+1}, is a function of the current state, S_t, the current action, x_t, and a vector of stochastic events, W_t, which includes fire arrival and weather.
5. A *Discount Factor*, δ, that determines how current rewards are valued relative to future rewards.

To analyze this MDP, we employ Bellman's equation (Bellman, 1957), a recursive equation that assigns a value to a particular state Eqs. (1a), (1b) by computing an expectation over the values of future states. This equation represents a landowner's decision-making process. It can be broken into two parts – immediate cost or benefit of a management action, which is captured by the reward function, and the expectation of the maximized value of the next period's state, $V(S'(S_t, x_t, W_t))$. We employ the so-called action-value representation, also known as the "*Q*-value" (Watkins and Dayan, 1992). The quantity $Q(S_t, x_t)$ is the expected return of taking action x_t in state S_t and then behaving optimally thereafter. In approximate dynamic programming, it is often useful to first compute $Q(S_t, x_t)$ for each possible action x_t and then choose the action that has the highest *Q* value. This is the optimal action in state S_t, and it defines the value $V(S_t)$.

$$V(S_t) = \max_{x_t} Q(S_t, x_t) \tag{1a}$$

where $Q(S_t, x_t)$ is the value of performing action x_t on the landscape in state S_t and behaving optimally thereafter:

$$Q(S_t, x_t) = C(S_t, x_t) + \mathbb{E}_{S_{t+1}}[\delta V(S'(S_t, x_t, W_t))] \tag{1b}$$

2.2. Spatial interactions

To incorporate spatial interactions, we decompose $Q(S_t, x_t)$ into separate action-value functions for each stand in the landscape. Let $Q_i(S_t, x_t)$ denote the contribution of stand $i = 1, 2, \ldots, I$ to the value of the overall landscape at time *t*. Note that the state S_t, action x_t, and stochastic event W_t are not indexed by *i* because they represent the

Download English Version:

https://daneshyari.com/en/article/6459676

Download Persian Version:

https://daneshyari.com/article/6459676

Daneshyari.com