Stochastic Environmental Research and Risk Assessment
https://doi.org/10.1007/s00477-024-02769-1

ORIGINAL PAPER

Optimal decision rules for marked point process models

M. N. M. van Lieshout[1,2]

Accepted: 28 June 2024

Abstract

We study a Markov decision problem in which the state space is the set of finite marked point patterns in the plane, the actions represent thinnings, the reward is proportional to the mark sum which is discounted over time, and the transitions are governed by a birth-death-growth process. We show that thinning points with large marks maximises the discounted total expected reward when births follow a Poisson process and marks grow logistically. Explicit values for the thinning threshold and the discounted total expected reward over finite and infinite horizons are also provided. When the points are required to respect a hard core distance, upper and lower bounds on the discounted total expected reward are derived.

Keywords French thinning · Logistic growth · Marked point process · Markov decision process

Mathematics Subject Classification 60G55 · 90C40

1 Introduction

A realisation of a marked spatial point process is a finite list of locations in some compact subset of the plane together with real-valued or categorical marks attached to it (Chiu et al. 2013). Examples include earthquakes labelled by time of occurrence and magnitude, cells labelled by geometric marks (roundness, size etc) or atom locations in crystals labelled by their type.

In forestry, modelling and inference in terms of point processes marked by species label or diameter at breast height has a long history. The seminal monograph by Matérn (1986), a revision and extension of his earlier licentiate and PhD theses from 1947 and 1960, respectively, has been particularly influential. In this book, a class of isotropic spatial covariance functions and various models for marked point processes—including some based on mark-dependent thinning—that now bear Matérn's name were introduced. From an applied point of view, the sampling errors of line and area surveys were investigated. Others have looked at, for example, the characterisation of the amount of clustering/inhibition between trees (Loosmore and Ford 2006), at testing biodiversity axioms (Wiegand et al. 2020; De Jongh and Van Lieshout 2022), as well as at the modelling of hierarchical dependence and growth (Renshaw et al. 2009). Recent years have witnessed a growing interest in more complicated mark spaces such as function spaces (Ghorbani et al. 2021). An overview of the state of the art can be found in the discussion paper by Eckardt and Moradi (2024).

This paper is motivated by timber harvesting (Pretzch 2009) where the objective is to design a policy that maximises the profit or, equivalenty, the volume of timber over time. The classical strategy is to use discretised stand based growth tables and dynamic programming (Rönnqvist 2003). Optimisation of point process based policies has been rarer due to 'a lack of models and to difficulties in selecting trees to be removed' (Pukkala and Miina 1998) and tend to be simulation based (Fransson et al. 2020; Pukkala et al. 2015; Renshaw and Särkkä 2001; Renshaw et al. 2009).

A theory for optimal decision making using spatial point process models has been developed in other fields, for example in mobile network optimisation. However, the role of the point process tends to be auxiliary in that it is used to model the spatial distribution of users, base stations and so on, from which coverage probabilities and other performance characteristics of the network can be calculated (Baccelli and Blaszczyszyn 2009; Lee et al. 2020; Lu et al. 2021;

✉ M. N. M. van Lieshout
m.n.m.van.lieshout@cwi.nl

1 CWI, P.O. Box 94079, 1090 GB Amsterdam, The Netherlands

2 Department of Applied Mathematics, University of Twente, P.O. Box 217, 7500 AE Enschede, The Netherlands

Khloussy et al. 2015). Spatial point process models are also convenient in multi-target tracking (Van Lieshout 2008) and their void probabilities or divergence measures can form the basis for observer trajectory optimisation (Beard et al. 2017).

Our focus of interest is to assume that policies are defined directly on the marked point process in terms of a mark-dependent thinning (Matérn 1986; Myllimäki 2009). Such policies are well-known in forestry. German thinning, for instance, is supposed to enhance natural selection by felling a fraction of those trees whose diameter at breast height is smaller than some threshold; French thinning (also known as thinning from above) is similar, except that a fraction of trees with large rather than small sizes is removed to stimulate forest rejuvenation. In either case, picking a policy amounts to choosing the threshold. Simulations in Fransson et al. (2020) and Pukkala et al. (2015) suggest that French thinning might be the better strategy.

The plan is as follows. Section 2 reviews basic concepts from Markov decision theory and marked point processes. In Sect. 3, we give a formal definition of a decision process in which the actions consist of deleting a subset of the current points and the reward is proportional to the marks. The stochastic process that governs the dynamics is a birth-and-death process with independent deaths and a Poisson process of births; the marks grow logistically. We calculate the discounted total expected reward function over finite and infinite horizons and derive an optimal policy. In Sect. 4, we move on to allow interaction between the points and replace the Poisson birth process by one in which no point is allowed to come too close to another point. In this setting, we provide upper and lower bounds on the discounted total expected reward function over finite and infinite horizons. To not interrupt the flow of arguments, technical proofs are deferred to Sect. 6. Section 5.1 returns to the motivating example of timber harvesting and compares optimal French thinning to sub-optimal German thinning. In Sect. 5.2 the tightness of the bounds on the reward function is investigated by means of simulated examples. We conclude by mentioning some topics for further research.

2 Preliminaries and notation

2.1 Markov decision theory

A Markov decision process (Bertsekas 1995; Feinberg and Schwartz 2002; Puterman 1994) is defined as follows. Write $\mathcal{X}$ for the state space that contains all states that the process can be in. When the system is in state $x \in \mathcal{X}$, the decision maker can take an action in some set $A(x)$ that may depend on the state x. If action $a \in A(x)$ is chosen, a direct reward $r(x, a)$ is earned and a probability mass function $p(\cdot|x, a)$ on $\mathcal{X}$ governs the next state of the process. The fact that only the current state and action matter rather than the entire past history justifies the epithet 'Markov'.

A *policy* $\Phi = (\phi_i)_{i=0}^{\infty}$ is a procedure for the selection of an action at each decision time $i \in \mathbb{N}_0 = \{0, 1, 2, \dots\}$. Such a policy could be random or deterministic, and in principle take into account the entire history of the process. A policy is said to be *stationary* if its members $\phi_i \equiv \phi$ do not depend on the time i.

Let (X_i, Y_i) denote the stochastic process of states X_i and actions Y_i with $i = 0, 1, \dots$. Write $\mathbb{E}^{\Phi}$ for the expectation operator when the transitions are driven by policy Φ and let $0 \le \alpha < 1$ be a discount factor. Then an optimal policy maximises the α-discounted total expected reward

$$v_{\alpha}^{\Phi}(x) = \mathbb{E}^{\Phi}\left[\sum_{i=0}^{\infty} \alpha^i r(X_i, Y_i) | X_0 = x\right]. \tag{1}$$

Note that the value of a reward r decreases over time: r is only worth $\alpha^i r < r$ after i time units. If the reward function is bounded, then (1) is well-defined.

When the state and action spaces are both finite, by Puterman (1994, Theorem 5.5.3b) it suffices to consider only *Markov policies* for which the actions chosen depend only on the current state and not on the past history. Furthermore, by Puterman (1994, Theorem 6.2.10), one may restrict oneself even further to the class of Markov policies that are deterministic and stationary. Note that such policies can be fully described by $\Phi = (\phi, \phi, \dots)$ for some mapping $\phi : \mathcal{X} \to A$ that assigns an action $\phi(x) \in A(x)$ to the current state x. The maximal α-discounted total expected reward can be found by policy iteration [e.g. Puterman (1994, Theorem 6.4.2)] or value iteration, also known as successive approximation or dynamic programming (Puterman 1994, Section 6.3). Policy iteration is based on successive improvement of the policy. If the current policy is $\Phi = (\phi, \phi, \dots)$, calculate its 'value' function v by solving the system of equations

$$v(x) - \alpha \sum_{y \in \mathcal{X}} v(y) p(y|x, \phi(x)) = r(x, \phi(x)), \quad x \in \mathcal{X}.$$

Any solution $\tilde{\Phi} = (\tilde{\phi}, \tilde{\phi}, \dots)$ to

$$\tilde{\phi}(x) = \operatorname*{argmax}_{a \in A(x)} \left\{ r(x, a) + \alpha \sum_{y \in \mathcal{X}} v(y) p(y|x, a) \right\}, \quad x \in \mathcal{X},$$

then yields an improved policy. This procedure is repeated until no further improvement is possible. Dynamic programming targets (1) directly. Specifically, if the current value function is v, find an improved $\tilde{v}$ as follows:

$$\tilde{v}(x) = \max_{a \in A(x)} \left\{ r(x, a) + \alpha \sum_{y \in \mathcal{X}} v(y) p(y|x, a) \right\}, \quad x \in \mathcal{X}.$$

The iteration stops when the improvement $||v - \tilde{v}||$ is smaller than a user-defined precision threshold.

When the cardinality of the state or action space is infinite, policy iteration is not guaranteed to converge in a finite number of steps [e.g. Puterman (1994, Section 6.4) or Bertsekas and Shreve (1978, p. 64)]. The dynamic programming approach on the other hand is amenable to generalisation to more general state and action spaces, and will form the basis for our exploration in the next sections.

2.2 Marked point processes

A point process on a compact set $W \subset \mathbb{R}^d$ is defined as a measurable mapping from some probability space into the set of boundedly finite integer-valued measures equipped with the Borel σ-algebra of the weak topology. A point process is *simple* if it almost surely does not contain multiple points in the sense that the measures assign mass 0 or 1 to sets consisting of a single point. In this case the measures may be identified with their support and realisations of the point process take the form $\mathbf{x} = \{w_1, \dots, w_n\} \subset W$ for $n \in \mathbb{N}_0$. For further details, we refer to (Daley and Vere-Jones 2008, Chapter 9.1).

A finite marked point process X (Daley and Vere-Jones 2003, Definition 6.4.1) with points in a compact set $W \subset \mathbb{R}^d$ and marks in a complete separable metric space L is a point process on the product space $W \times L$ such that the number of points in $W \times L$ is finite. When the ground process obtained from X upon disregarding the marks is simple, realisations are of the form $\{(w_1, m_1), \dots, (w_n, m_n)\}$ for $w_i \in W$, $m_i \in L$ $(i = 1, \dots, n)$ and $n \in \mathbb{N}_0$. In this case, we shall say that X itself is simple.

The *intensity measure* Λ of X is defined on product sets $A = B \times C$ in the Borel product σ-algebra on $W \times L$ by

$$\Lambda(A) = \mathbb{E}X(A) = \mathbb{E}X(B \times C),$$

the expected number of points in B with marks in C, and can be extended to general Borel sets through linear combinations and monotone limits. More generally, for any measurable function $f \geq 0$ on $W \times L$, and assuming that X is simple,

$$\mathbb{E}\left[\sum_{(w,m)\in X} f(w,m)\right] = \int_{W\times L} f(w,m)\, d\Lambda(w,m). \tag{2}$$

Here the left-hand side is taken to be infinite if and only if the right-hand side is. This identity is known as the *Campbell–Mecke formula* (Daley and Vere-Jones 2008, Section 9.5).

To close this section, we recall the definition of an independently marked Poisson process. Let $\beta : W \to [0, \infty)$ be a measurable, integrable function and ν a probability measure on L. Then, N, the total number of marked points is Poisson distributed with rate parameter $\int_W \beta(w)dw$ and given N, the locations of the marked points in W are independent and scattered according to a probability density proportional to β, the marks are independent and distributed according to ν. This point process is simple (Daley and Vere-Jones 2003, Section 5.4).

3 Marked Poisson process model with logistic growth

3.1 Definition of the model

Let the state space $\mathcal{X}$ consist of finite simple marked point patterns on a compact set $W \subset \mathbb{R}^2$ with marks in $L = [0, K]$ for some $K > 0$ (cf. Sect. 2.2). When at time $i \in \mathbb{N}_0$, the process is in state $\mathbf{x}$, a thinning action is carried out, resulting in a new state $\mathbf{a}$ that consists of all retained points $\mathbf{a} \subset \mathbf{x}$. Thus, in the notation and framework of Sect. 2.1, the action space $A(\mathbf{x})$ is finite and contains all subsets of $\mathbf{x}$. Define a reward function $r(\mathbf{x}, \mathbf{a})$ by

$$r(\mathbf{x}, \mathbf{a}) = R \sum_{(x,m)\in \mathbf{x}\setminus \mathbf{a}} m, \quad \mathbf{x} \in \mathcal{X}, \mathbf{a} \subset \mathbf{x}. \tag{3}$$

Thus, the reward is proportional to the sum of the marks of all removed points. When $R > 0$, the reward $r(\cdot,\cdot)$ takes non-negative values. Moreover, since the mark content in an $\mathbb{R}^+$-marked point process is a random variable by Daley and Vere-Jones (2003, Proposition 6.4.V), r is well-defined.

Upon taking action $\mathbf{a}$ in state $\mathbf{x}$, the dynamics that lead to the next state are modelled as a birth-death-growth process. Specifically, the marks of the retained points $(x, m) \in \mathbf{a}$ grow according to the well-known logistic model that was proposed around 1840 by Verhulst and Quetelet (Richards 1959). In this model, when the mark at time 0 is $m > 0$, the mark at time $n \in \mathbb{N}_0$ is

$$g^{(n)}(m) = \frac{K}{1 + \left(\frac{K}{m} - 1\right)e^{-\lambda n}}. \tag{4}$$

By convention, $g^{(n)}(0) = 0$. The parameter $\lambda > 0$ governs the rate of growth and $K \geq m \geq 0$ is an upper bound on the size. In combination with independent births and deaths, the next state is defined by the following dynamics:

- delete $\mathbf{x} \setminus \mathbf{a}$;
- independently of other points, let each $(x_i, m_i) \in \mathbf{a}$ die with probability $p_d \in (0, 1)$ (natural deaths) and otherwise grow to $(x_i, g^{(1)}(m_i))$ as in (4);
- add a Poisson process on W with intensity function $\beta : W \to \mathbb{R}^+$ and mark its points independently according to a probability measure ν on $[0, K]$.

Write $(X_i, Y_i)_{i=0}^{\infty}$ for the sequence of successive states X_i and actions Y_i. A randomised policy $\Phi = (\phi_i)_{i=0}^{\infty}$ is a sequence of conditional probability kernels $\phi_i(\cdot|X_0, Y_0, \ldots, X_{i-1}, Y_{i-1}, X_i)$ on A to generate Y_i based on the history of the process such that $\phi_i(A(\mathbf{x}_i)|\mathbf{x}_0, \mathbf{a}_0, \ldots, \mathbf{x}_i) = 1$. If the policy is Markov and deterministic, Y_i is simply a function of X_i, and one may write $Y_i = \phi_i(X_i)$. Recalling the definitions of the action spaces and of the reward function, for $0 \le \alpha < 1$, the infinite horizon α-discounted total expected reward function (1) under policy $\Phi = (\phi_i)_{i=0}^{\infty}$ with initial state $X_0 = \mathbf{x}$ reads

$$v_\alpha^\Phi(\mathbf{x}) = \mathbb{E}^\Phi\left[\sum_{i=0}^{\infty} \alpha^i \left(R \sum_{(x,m)\in X_i\setminus Y_i} m\right) \mid X_0 = \mathbf{x}\right]. \tag{5}$$

The following lemma shows that the model just described is well-defined.

Lemma 1 *The infinite horizon α-discounted total expected reward function $v_\alpha^\Phi(\mathbf{x})$, $\mathbf{x} \in \mathcal{X}$, defined in (5) is finite for all $0 \le \alpha < 1$, all $R > 0$ and all policies Φ.*

Proof Pick $\mathbf{x} \in \mathcal{X}$ and write $n(\mathbf{x}) < \infty$ for its cardinality. Since the growth function (4) is bounded by K,

$$\mathbb{E}^\Phi\left[\sum_{(x,m)\in X_0\setminus Y_0} m \mid X_0 = \mathbf{x}\right] \le Kn(\mathbf{x}).$$

For $i > 0$, X_i is the union of survivors from $\mathbf{x}$, survivors from subsequent generations starting with $X_1 \setminus X_0$ up to $X_{i-1} \setminus X_{i-2}$ and points born in the time since the last decision. Therefore, recalling the birth and death dynamics,

$$\mathbb{E}^\Phi\left[\sum_{(x,m)\in X_i\setminus Y_i} m|X_0 = \mathbf{x}\right] \le Kn(\mathbf{x})(1-p_d)^i + K\int_W \beta(w)dw \sum_{k=0}^{i-1}(1-p_d)^k.$$

Hence,

$$v_\alpha^\Phi(\mathbf{x}) \le RKn(\mathbf{x})\sum_{i=0}^{\infty}\alpha^i(1-p_d)^i + RK\int_W \beta(w)dw \sum_{i=1}^{\infty}\alpha^i\sum_{k=0}^{i-1}(1-p_d)^k.$$

For all $p_d \in (0, 1)$, the first series in the right-hand side converges to $1/(1-\alpha(1-p_d))$. Since

$$\sum_{i=1}^{\infty}\alpha^i\sum_{k=0}^{i-1}(1-p_d)^k = \sum_{i=1}^{\infty}\alpha^i\frac{1-(1-p_d)^i}{p_d} \le \frac{1}{p_d}\sum_{i=1}^{\infty}\alpha^i < \infty$$

for all $p_d \in (0, 1)$ and β is integrable, $v_\alpha^\Phi(\mathbf{x})$ is finite. □

The reward function r itself is not bounded, so the (N) regime of Bertsekas and Shreve (1978, Chapter 9) applies.

3.2 Optimal policy and reward

Our objective in this section is to find the best policy in the sense that it maximises the total expected reward, discounted for elapsed time.

Formally, the optimal α-discounted total expected reward $v_\alpha^*(\mathbf{x})$ is defined as the supremum of $v_\alpha^\Phi(\mathbf{x})$ over all policies, including randomised ones. By Bertsekas and Shreve (1978, Proposition 9.1), the supremum in the definition of $v_\alpha^*(\mathbf{x})$ may be taken over the class of Markov policies, and, by Bertsekas and Shreve (1978, Proposition 9.8), satisfies the equation

$$v_\alpha^*(\mathbf{x}) = \max_{\mathbf{a}\subset\mathbf{x}}\left\{R\sum_{(x,m)\in\mathbf{x}\setminus\mathbf{a}} m + \alpha\mathbb{E}\left[v_\alpha^*(X) \mid \mathbf{x}, \mathbf{a}\right]\right\}, \quad \mathbf{x} \in \mathcal{X}, \tag{6}$$

where X is distributed according to the one step birth-death-growth dynamics from state $\mathbf{x}$ under action $\mathbf{a}$. Observe that the optimality equations (6) are not sufficient conditions for v_α^*. Nevertheless, $v_\alpha^*(\mathbf{x})$ can be calculated as the limit of a dynamic programming algorithm. Set $v_0(\mathbf{x}) = 0$ for all $\mathbf{x} \in \mathcal{X}$ and set $n = 1$. Define, for every $\mathbf{x} \in \mathcal{X}$,

$$v_n(\mathbf{x}) = \max_{\mathbf{a}\subset\mathbf{x}}\left\{R\sum_{(x,m)\in\mathbf{x}\setminus\mathbf{a}} m + \alpha\mathbb{E}\left[v_{n-1}(X) \mid \mathbf{x}, \mathbf{a}\right]\right\}.$$

Then set $n = n + 1$ and repeat. This algorithm converges to $v_\alpha^*(\mathbf{x})$ as $n \to \infty$ by Bertsekas and Shreve (1978, Proposition 9.14) but, in general, is of little help in constructing an optimal policy, let alone a stationary one. Given a stationary policy $\Phi = (\phi, \phi, \ldots)$, a necessary and sufficient condition for it to be optimal according to Bertsekas and Shreve (1978, Proposition 9.13) is that

$$v_\alpha^\Phi(\mathbf{x}) = \max_{\mathbf{a}\subset\mathbf{x}}\left\{R\sum_{(x,m)\in\mathbf{x}\setminus\mathbf{a}} m + \alpha\mathbb{E}\left[v_\alpha^\Phi(X)|\mathbf{x}, \mathbf{a}\right]\right\}. \tag{7}$$

For our model, the dynamic programming algorithm does suggest an optimal deterministic and stationary Markov policy.

Theorem 1 *Consider the Markov decision process with state space $\mathcal{X}$, action spaces $A(\mathbf{x}) = \{\mathbf{y} \in \mathcal{X} : \mathbf{y} \subset \mathbf{x}\}$, $\mathbf{x} \in \mathcal{X}$, reward function (3) with $R > 0$, and birth-death-growth dynamics based on independent deaths with probability $p_d \in (0, 1)$, a Poisson birth process with measurable, integrable intensity function $\beta : W \to \mathbb{R}^+$ marked independently according to probability measure ν on $[0, K]$ for $K > 0$ and logistic growth function (4). Then, for $0 \le \alpha < 1$,*

$$v_\alpha^*(\mathbf{x}) = \frac{R\alpha}{1-\alpha}\int_W\int_0^K s(m)\,\beta(w)dwd\nu(m) + R\sum_{(x,m)\in\mathbf{x}} s(m), \tag{8}$$

where

$$s(m) = \sup_{n\in\mathbb{N}_0}\left\{\frac{K\alpha^n(1-p_d)^n}{1+\left(\frac{K}{m}-1\right)e^{-\lambda n}}\right\}, \quad m\in[0,K].$$

Furthermore, the optimal α-discounted total expected reward corresponds to a stationary policy that removes all points with a mark that is at least

$$d_\alpha^* = \sup_{n\in\mathbb{N}_0}\left\{\frac{K}{1-e^{-n\lambda}}\left(\alpha^n(1-p_d)^n - e^{-n\lambda}\right)\right\}$$

under the convention that $0/0=0$.

We close this section with three remarks. Firstly, for $\alpha = 1$, the total expected reward $v_1^*(\mathbf{x})$ is infinite. Secondly, in the proof of Theorem 1 in Sect. 6, we derive an optimal policy and corresponding α-discounted total expected reward for finite time horizons too. Finally, note that the suprema in $s(m)$ and d_α^* are attained.

4 Hard core models with logistic growth

In this section, we refine the Poisson model of the previous section to the case where births are governed by a hard core process whose realisations are not allowed to contain a pair of points $\{x_1, x_2\}\subset W$ such that $||x_1-x_2||\le K$. For the motivating example from forestry in which the marks correspond to the diameter at breast height, the condition ensures that all trees can grow to their maximal size.

As in Sect. 3.1, when at time $i\in\mathbb{N}_0$ the process is in state $\mathbf{x}$, a thinning action is carried out, resulting in a new state $\mathbf{a}$ that consists of all retained points. The reward function (3) also remains unchanged.

The dynamics are modified in such a way that the hard core is respected. Thus, let the process be in some state $\mathbf{x}$ satisfying the hard core constraint and suppose that action $\mathbf{a}$ is taken. The next state is then governed by the following birth-death-growth process:

- delete $\mathbf{x}\setminus\mathbf{a}$;
- independently of other points, let each $(x_i, m_i)\in\mathbf{a}$ die with probability $p_d\in(0,1)$ and otherwise grow to $(x_i, g(m_i))$ for some bounded, continuous function $g:[0,K]\to[0,K]$ satisfying $m\le g(m)$ for $m\in[0,K]$;
- add a hard core process on W with hard core distance K and measurable, integrable intensity function $\beta: W\to\mathbb{R}^+$; mark its points independently according to a probability measure ν on $[0,K]$ and remove all points that fall within distance K to a point in $\mathbf{a}$.

In this framework, the reward function is bounded since the hard core condition implies an upper bound on the number of points that can be alive at any time. We are therefore in the (D) regime of Bertsekas and Shreve (1978, Chapter 9).

As before, define $v_\alpha^*(\mathbf{x})$ as the supremum of (5) over all policies Φ. By Bertsekas and Shreve (1978, Proposition 9.1) it suffices to consider Markov policies only, and $v_\alpha^*(\mathbf{x})$ is the limit of the dynamic programming algorithm (Bertsekas and Shreve 1978, Proposition 9.14). The optimality condition (7) applies. Moreover, since the action sets are finite, Corollary 9.17.1 in Bertsekas and Shreve (1978) guarantees the existence of an optimal deterministic stationary policy. An explicit expression seems hard to obtain. However, the following bounds on the finite horizon discounted total expected reward function are available.

Theorem 2 *Consider the Markov decision process with state space $\mathcal{X}$, action spaces $A(\mathbf{x}) = \{\mathbf{y}\in\mathcal{X}:\mathbf{y}\subset\mathbf{x}\}$, $\mathbf{x}\in\mathcal{X}$, reward function (3) with $R>0$, and birth-death-growth dynamics described above. Write $g^{(n)}$ for the n-fold composition of the growth function g.*

For $\alpha\in[0,1)$ and $\mathbf{x}\in\mathcal{X}$ containing no pair of K-close points, initialise $v_0(\mathbf{x})=0$. Define, for $n\in\mathbb{N}$,

$$v_n(\mathbf{x}) = \max_{\mathbf{a}\subset\mathbf{x}}\left\{R\sum_{(x,m)\in\mathbf{x}\setminus\mathbf{a}} m + \alpha\mathbb{E}\left[v_{n-1}(X)\mid\mathbf{x},\mathbf{a}\right]\right\},$$

where X is distributed according to the one step birth-death-growth dynamics from $\mathbf{x}$ under action $\mathbf{a}$. Then $\tilde v_n(\mathbf{x})\le v_n(\mathbf{x})\le\hat v_n(\mathbf{x})$ where

$$\begin{aligned}\tilde v_n(\mathbf{x}) =& R\sum_{(x,m)\in\mathbf{x}}\tilde s_n(x,m)\\ &+R\sum_{k=1}^{n-1}\alpha^k\int_W\int_0^K\tilde s_{n-k}(w,m)\,\beta(w)dwd\nu(m)\\ \hat v_n(\mathbf{x}) =& R\sum_{(x,m)\in\mathbf{x}}\hat s_n(m)\\ &+R\sum_{k=1}^{n-1}\alpha^k\int_W\int_0^K\hat s_{n-k}(m)\,\beta(w)dwd\nu(m)\end{aligned}$$

with $\tilde s_0=\hat s_0=0$ and, for $n\in\mathbb{N}$,

$$\hat s_n(m) = \max\left\{m, \alpha(1-p_d)g^{(1)}(m),\ \dots,\ \alpha^{n-1}(1-p_d)^{n-1}g^{(n-1)}(m)\right\}$$

and, writing $b(x,K)$ for the closed ball centred at x with radius K,

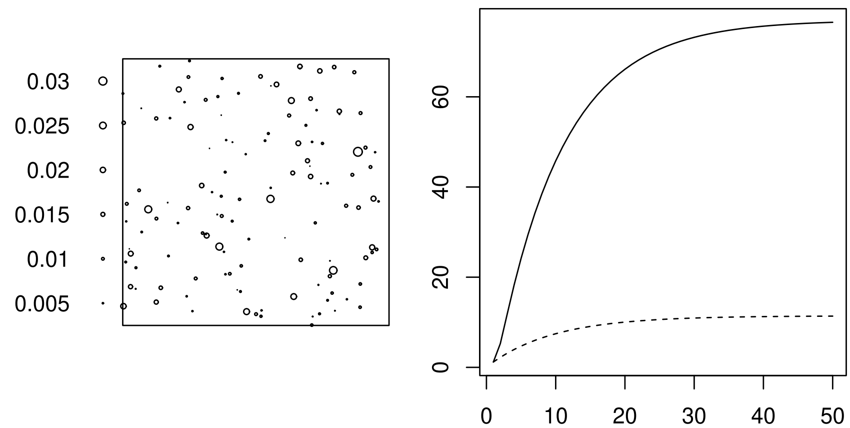

Fig. 1 Left panel: sample $\mathbf{x}$ from a Poisson process with intensity $\beta = 5.0$ on $[0,5]^2$ marked independently according to a Beta distribution on $[0, 0.1]$ with shape parameters $\lambda_1 = 2.0$ and $\lambda_2 = 20.0$. Right panel: graphs of the finite horizon α-discounted total expected reward $v_n(\mathbf{x})$ against n for the birth-death-growth dynamics of Sect. 5 under the optimal policy (solid line) and under German thinning at threshold level K (dotted line)

$$\tilde{s}_n(x,m) = \max\{m, \alpha(1-p_d)g^{(1)}(m) - \alpha K \int_{b(x,K)\cap W} \beta(w)dw, \ldots, \alpha^{n-1}(1-p_d)^{n-1} g^{(n-1)}(m) - \alpha K \int_{b(x,K)\cap W} \beta(w)dw \sum_{i=0}^{n-2} \alpha^i (1-p_d)^i\}.$$

The proof of Theorem 2 carries over to the case that g and $\hat{s}_n$ are space-dependent, in other words, $g(x,m)$, $\hat{s}_n(x,m)$ are functions of locations and marks. When the growth function is logistic,

$$\tilde{s}_n(x,m) = \max_{i=0,\ldots,n-1} \left\{ \frac{K\alpha^i(1-p_d)^i}{1+\left(\frac{K}{m}-1\right)e^{-\lambda i}} - \alpha K \left(\frac{1-\alpha^i(1-p_d)^i}{1-\alpha(1-p_d)}\right) \int_{b(x,K)\cap W} \beta(w)dw \right\};$$

$$\hat{s}_n(m) = \max_{i=0,\ldots,n-1} \left\{ \frac{K\alpha^i(1-p_d)^i}{1+\left(\frac{K}{m}-1\right)e^{-\lambda i}} \right\}.$$

Over an infinite time horizon, the optimal α-discounted total expected reward is bounded by $\tilde{v}$ and $\hat{v}$ which have the same functional forms as $\tilde{v}_n$ and $\hat{v}_n$ in Theorem 2 for $\tilde{s}$ and $\hat{s}$ given in the following result. The bounds coincide if $\alpha = 0$.

Corollary 1 *The functions $\hat{s}_n$ and $\tilde{s}_n$ defined in Theorem 2 take values in $[0,K]$ and increase monotonically to*

$$\hat{s}(m) = \sup_{n\in\mathbb{N}_0} \left\{\alpha^n(1-p_d)^n g^{(n)}(m)\right\}, \quad m \in [0,K],$$

and, for $x \in W$ and $m \in [0,K]$,

$$\tilde{s}(x,m) = \sup_{n\in\mathbb{N}_0} \left\{\alpha^n(1-p_d)^n g^{(n)}(m) - \alpha K \int_{b(x,K)\cap W} \beta(w)dw \sum_{i=0}^{n-1} \alpha^i (1-p_d)^i \right\}$$

as $n \to \infty$.

5 Examples

5.1 A comparison between French and German thinning

Let us return to the motivating example in which the marks represent the diameter at breast height of a tree and the reward is the value of the timber harvest. Suppose that the dynamics are those outlined in Sect. 3. Then, according to Theorem 1, the optimal policy is a French thinning, that is, to fell all trees whose diameter is at least as large as some threshold value.

An alternative approach is that of German thinning or thinning from below in which small rather than large trees are being harvested. Using the fact that trees can only grow bigger, so that any tree is felled upon first appearance if at all, it is not hard to see that the α-discounted total expected reward function is of the form (8) with

$$s(m) = m\, 1\{m \le d\},$$

where d is the threshold for harvesting. Thus, it is best to cut all trees immediately, that is, to set $d = K$.

To evaluate what is lost compared to the the optimal policy, the right panel of Fig. 1 shows the finite horizon α-discounted total expected reward as a function of the time horizon for French and for German thinning. The initial pattern $\mathbf{x}$, shown in the left panel of Fig. 1, was a sample from a Poisson process with intensity $\beta = 5$ on $W = [0,5]^2$. For the mark dynamics, we used a logistic growth function with $\lambda = 2$ and maximal size $K = 0.1$; the initial marks were sampled from a Beta distribution on [0, 0.1] with shape parameters $\lambda_1 = 2$ and $\lambda_2 = 20$. The death rate was $p_d = 0.05$, the discount factor $\alpha = 0.9$ and the reward parameter $R = 1$. Taking a time horizon up to $n = 50$, the solid line in the right panel represents the α-discounted total expected reward (10) under the optimal policy that at each decision epoch fells all trees whose diameter at breast height is at least

$$d_n = \max\left\{0, K\frac{\alpha(1-p_d)-e^{-\lambda}}{1-e^{-\lambda}}, \dots, K\frac{\alpha^{n-1}(1-p_d)^{n-1}-e^{-(n-1)\lambda}}{1-e^{-(n-1)\lambda}}\right\}$$
$$=0\times\mathbf{1}\{n=1\}+0.083\times\mathbf{1}\{n\geq 2\}, \quad n=1,2,\dots,50,$$

the dotted line is the graph of

$$v_n(\mathbf{x}) = R\sum_{(x,m)\in\mathbf{x}} m + \frac{\lambda_1 RK\beta|W|}{\lambda_1+\lambda_2}\frac{\alpha-\alpha^n}{1-\alpha}, \quad n=1,2,\dots,50,$$

corresponding to the greedy German thinning. Note that the latter is by far the inferior policy. Its discounted total expected reward tends to 11.4 compared to 76.5 for the optimal policy.

5.2 Tightness of bounds for hard core models

The key idea in the proof of Theorem 2, given in full in Sect. 6, is to bound integrals of the form

$$\int_{W\times[0,K]} s_n(w,l)1\{w\notin\cup_{(x,m)\in\phi(\mathbf{x})}b(x,K)\}\beta(w)dwd\nu(l)$$

from below by

$$\int_{W\times[0,K]} s_n(w,l)\beta(w)dwd\nu(l) - K\sum_{(x,m)\in\phi(\mathbf{x})}\Lambda(b(x,K))$$

and from above by

$$\int_{W\times[0,K]} s_n(w,l)\beta(w)dwd\nu(l)$$

to obtain, respectively, the lower and upper bounds $\tilde{v}_n(\mathbf{x})$ and $\hat{v}_n(\mathbf{x})$ on the finite horizon discounted total expected reward function. One would expect the tightness to depend on the amount of overlap in balls around the points of the process. When the intensity function is small, there tend to be few points that are spread out well. On the other hand, when β gets larger, more points will be located closer together which may affect the accuracy of the approximation.

To back up this conjecture, we calculated $\hat{v}_n(\mathbf{x})$ and $\tilde{v}_n(\mathbf{x})$ in two regimes, a dense one and a sparse one. For specificity, we considered the stationary case with $\beta > 0$ constant. For the inital pattern $\mathbf{x}$, a sample from a Strauss process (Kelly and Ripley 1976) on $W = [0,5]^2$ with interaction parameter set to zero was chosen. The activity parameter was set to give the required intensity: $\beta = 1.0$ in the sparse regime and $\beta = 4.3$ in the dense regime. For the mark dynamics, we used a logistic growth function with $\lambda = 2$ and maximal size $K = 0.1$; the initial marks were sampled from a Beta distribution on $[0, K]$ with shape parameters $\lambda_1 = 2$ and $\lambda_2 = 20$. The death rate was set to $p_d = 0.05$. Finally, we used discount factor $\alpha = 0.9$ and reward parameter $R = 1$.

The results are plotted in Fig. 2. The left panels show the pattern $\mathbf{x}$. In the right panels, the solid lines are the graphs of $\hat{v}_n(\mathbf{x})$ as a function of n, the dotted lines show $\tilde{v}_n(\mathbf{x})$ plotted against n. Integrals were estimated by the Monte Carlo method with 1, 000 samples. In the sparse regime, the approximation is quite good, for the denser regime, the gap between the two graphs is quite wide except for very small n. In both cases, the dynamic programming algorithm converges rapidly.

6 Proofs

Proof of Theorem 1 After initialising $v_0(\mathbf{x}) = 0$ for all $\mathbf{x}\in\mathcal{X}$, clearly the optimal expected reward after one action is $v_1(\mathbf{x}) = R\sum_{(x,m)\in\mathbf{x}} m$, which is attained for action $\mathbf{a} = \emptyset$, or, in other words, by removing all points with mark greater than or equal to $d_1 = 0$. The proof proceeds by induction. Set, for $n\in\mathbb{N}$,

$$d_n = \max\left\{0, K\frac{\alpha(1-p_d)-e^{-\lambda}}{1-e^{-\lambda}}, \dots, K\frac{\alpha^{n-1}(1-p_d)^{n-1}-e^{-(n-1)\lambda}}{1-e^{-(n-1)\lambda}}\right\} \tag{9}$$

and suppose that the optimal α-discounted total expected reward $v_n(\mathbf{x})$ over n actions is attained by removing all points whose mark is greater than or equal to d_n and is given by

$$v_n(\mathbf{x}) = R\sum_{k=1}^{n-1}\alpha^k\int_W\int_0^K s_{n-k}(m)\,\beta(w)dwd\nu(m) + R\sum_{(x,m)\in\mathbf{x}} s_n(m), \tag{10}$$

where, for $1\leq k\leq n$,

$$s_k(m) = \max\left\{m, \alpha\left(1-p_d\right)g^{(1)}(m), \dots, \alpha^{k-1}\left(1-p_d\right)^{k-1}g^{(k-1)}(m)\right\}.$$

Now, for $n+1$, the optimal finite horizon α-discounted total expected reward is

$$v_{n+1}(\mathbf{x}) = \max_{\mathbf{a}\subset\mathbf{x}}\left\{R\sum_{(x,m)\in\mathbf{x}\setminus\mathbf{a}} m + \alpha\mathbb{E}\left[v_n(X)\mid\mathbf{x},\mathbf{a}\right]\right\},$$

where X is distributed according to the one step birth-death-growth dynamics from $\mathbf{x}$ under action $\mathbf{a}$. By the induction

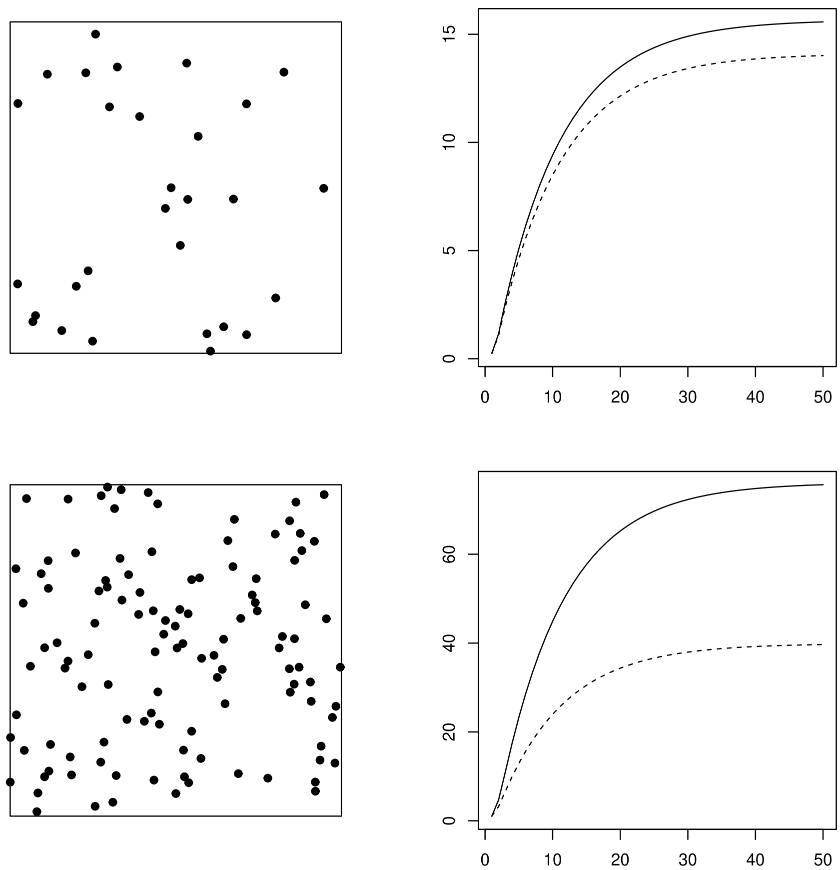

Fig. 2 Left panels: samples $\mathbf{x}$ from a Strauss hard core process with intensity $\beta = 1.0$ (top) and $\beta = 4.3$ (bottom) on $[0, 5]^2$. Right panels: graphs of $\hat{v}_n(\mathbf{x})$ (solid lines) and $\tilde{v}_n(\mathbf{x})$ (dotted lines) against n for the birth-death-growth dynamics of Sect. 5

assumption, the discounted expectation $\alpha\mathbb{E}\left[v_n(X) \mid \mathbf{x}, \mathbf{a}\right]$ is the sum of

$$\alpha R \sum_{k=1}^{n-1} \alpha^k \int_W \int_0^K s_{n-k}(m)\, \beta(w) dw d\nu(m)$$
$$= R \sum_{k=2}^{n} \alpha^k \int_W \int_0^K s_{n+1-k}(m)\, \beta(w) dw d\nu(m)$$

and contributions from the points in $\mathbf{a}$ that survive as well as from points born in the interval between decisions n and $n+1$. These contributions are, respectively,

$$\alpha R \sum_{(x,m)\in\mathbf{a}} \left(1 - p_d\right) s_n(g^{(1)}(m))$$

and, using the Campbell–Mecke formula (2),

$$\alpha R \int_W \int_0^K s_n(m)\, \beta(w) dw d\nu(m).$$

The optimal action is to assign a point $(x, m) \in \mathbf{x}$ to $\mathbf{x} \setminus \mathbf{a}$ if and only if $m \geq \alpha(1 - p_d) s_n(g^{(1)}(m))$. By the induction assumption and (4), this is the case if and only if

$$m \geq \alpha^k (1 - p_d)^k g^{(k)}(m) \Leftrightarrow m \geq K \frac{\alpha^k (1 - p_d)^k - e^{-k\lambda}}{1 - e^{-k\lambda}} \qquad (11)$$

for all integers $1 \leq k \leq n$. Consequently, d_{n+1} has the required form. For this allocation rule, the reward is $\max\left\{m, \alpha\left(1 - p_d\right) s_n(g^{(1)}(m))\right\} = s_{n+1}(m)$ and the induction step is complete.

Next, let n go to infinity and fix $m \in [0, K]$. Note that $s(m)$ is finite for all $p_d \in (0, 1)$ and $0 \leq \alpha < 1$. Additionally, $\lim_{n\to\infty} s_n(m) = s(m)$. Thus, for any $\mathbf{x} \in \mathcal{X}$,

$$R \sum_{(x,m)\in\mathbf{x}} s_n(m) \to R \sum_{(x,m)\in\mathbf{x}} s(m)$$

as $n \to \infty$. Furthermore,

$$\sum_{k=1}^{n-1} \alpha^k \int_0^K s_{n-k}(m)\, d\nu(m) \to \sum_{k=1}^{\infty} \alpha^k \int_0^K s(m)\, d\nu(m), \quad n \to \infty,$$

because of dominated convergence applied to the doubly indexed sequence $a_{k,n}$ defined by $\mathbf{1}\{k \leq n-1\}\alpha^k \int s_{n-k}\, d\nu$. In conclusion, by Bertsekas and Shreve (1978), Proposition 9.14), for each $\mathbf{x} \in \mathcal{X}$, $\lim_{n\to\infty} v_n(\mathbf{x}) = v_\alpha^*(\mathbf{x})$, the optimal α-discounted total expected reward, and $v_\alpha^*(\mathbf{x})$ has the claimed form.

To complete the proof, we need to show that $v_\alpha^*(\mathbf{x})$ is attained by the stationary deterministic policy Φ that retains all points with mark smaller than d_α^*. Denote its infinite horizon α-discounted total expected reward by

$$v_\alpha^{d^*}(\mathbf{x}) = \mathbb{E}^\Phi\left[R\sum_{i=0}^\infty \alpha^i \sum_{(x,m)\in X_i\setminus Y_i} m \mid X_0 = \mathbf{x}\right]$$

and focus on the contributions of each generation of points. A point $(x,m)\in\mathbf{x}$, the initial generation, yields a reward $R\alpha^n(1-p_d)^n g^{(n)}(m)$ precisely when $g^{(n-1)}(m) \geq g^{(n-2)})(m) \geq \cdots \geq m$ are less than d_α^* but $g^{(n)}(m) \geq d_\alpha^*$ with obvious modification for $n=0$. Since, as in (11), $g^{(n)}(m) \geq d_\alpha^*$ if and only if

$$g^{(n)}(m) \geq \alpha^k(1-p_d)^k g^{(n+k)}(m)$$

for all $k\in\mathbb{N}_0$, we conclude that every point of $\mathbf{x}$ contributes $R\,s(m)$. The points that are born in the time before the next decision (generation 1) yield the same total expected reward, but this is discounted by α due to the later birth date. Similarly, the total expected reward of points belonging to the second generation is discounted by α^2, and so on. Tallying up, the α-discounted total expected reward of generations $k=1,2,\ldots$ is

$$R\sum_{k=1}^\infty \alpha^k \int_W\int_0^K s(m)\,\beta(w)dwd\nu(m)$$

on application of the Campbell–Mecke formula. Finally add the contribution from the initial generation to conclude that the threshold d_α^* defines an optimal policy. Condition (7) is readily verified.

Proof of Theorem 2 The proof proceeds by induction. For $n=0$, evidently $\tilde v_0 \leq v_0 \leq \hat v_0$. Assume that $\tilde v_k(\mathbf{x}) \leq v_k(\mathbf{x}) \leq \hat v_k(\mathbf{x})$ for all $k\leq n$ and all $\mathbf{x}\in\mathcal{X}$ satisfying the hard core condition and that $\tilde v_k$, $\hat v_k$ have the required form. Since

$$v_{n+1}(\mathbf{x}) = \max_{\mathbf{a}\subset\mathbf{x}}\left\{R\sum_{(x,m)\in\mathbf{x}\setminus\mathbf{a}} m + \alpha\mathbb{E}\left[v_n(X)\mid \mathbf{x},\mathbf{a}\right]\right\} \qquad (12)$$

and $v_n(X)\geq \tilde v_n(X)$, let us consider the expectation of $\tilde v_n(X)$ under the hard core birth-death-growth dynamics when action $\mathbf{a}$ is taken in state $\mathbf{x}$. By the definition of $\tilde v_n$ and distinguishing between surviving and new-born points, $\mathbb{E}\left[\tilde v_n(X)\mid\mathbf{x},\mathbf{a}\right]$ reads

$$\begin{aligned}
&R\,\mathbb{E}\left[\sum_{(x,m)\in X}\tilde s_n(x,m)\mid \mathbf{x},\mathbf{a}\right]\\
&\quad + R\sum_{k=1}^{n-1}\alpha^k\int_W\int_0^K \tilde s_{n-k}(w,m)\,\beta(w)dwd\nu(m)\\
&= R\sum_{(x,m)\in\mathbf{a}}\left(1-p_d\right)\tilde s_n(x,g^{(1)}(m))\\
&\quad + R\sum_{k=1}^{n-1}\alpha^k\int_W\int_0^K \tilde s_{n-k}(w,m)\,\beta(w)dwd\nu(m)\\
&\quad + R\int_W\int_0^K \tilde s_n(w,m)\,\mathbf{1}\{w\notin U_K(\mathbf{a})\}\,\beta(w)dwd\nu(m),
\end{aligned}$$

where the symbol $U_K(\mathbf{a})$ signifies the union of closed balls with radius K around the points in $\mathbf{a}$. The calculation of the last term above relies on the Campbell–Mecke formula (2). Now, the integral in the last line above can be written as

$$\begin{aligned}
&R\int_W\int_0^K \tilde s_n(w,m)\,\beta(w)dwd\nu(m)\\
&\quad - R\int_W\int_0^K \tilde s_n(w,m)\,\mathbf{1}\{w\in U_K(\mathbf{a}\}\,\beta(w)dwd\nu(m)
\end{aligned}$$

and is bounded from below by

$$\begin{aligned}
&R\int_W\int_0^K \tilde s_n(w,m)\,\beta(w)dwd\nu(m)\\
&\quad - RK\sum_{(x,m)\in\mathbf{a}}\int_W\int_0^K \mathbf{1}\{w\in b(x,K)\}\,\beta(w)dwd\nu(l)
\end{aligned} \qquad (13)$$

where the induction assumption is invoked for the inequality $\tilde s_n \leq K$. Next, return to (12). The induction assumption and the bound on $\mathbb{E}\left[\tilde v_n(X)\mid\mathbf{x},\mathbf{a}\right]$ imply that

$$v_{n+1}(\mathbf{x}) \geq \max_{\mathbf{a}\subset\mathbf{x}}\left\{R\sum_{(x,m)\in\mathbf{x}\setminus\mathbf{a}} m + \alpha\mathbb{E}\left[\tilde v_n(X)\mid\mathbf{x},\mathbf{a}\right]\right\},$$

which in turn is greater than or equal to

$$\begin{aligned}
&\max_{\mathbf{a}\subset\mathbf{x}}\Bigg\{R\sum_{(x,m)\in\mathbf{x}\setminus\mathbf{a}} m + \alpha R\sum_{(x,m)\in\mathbf{a}}\Big[(1-p_d)\tilde s_n(x,g^{(1)}(m))\\
&\quad -K\int_{b(x,K)\cap W}\beta(w)dw\Big]\Bigg\}\\
&\quad + R\sum_{k=1}^{n}\alpha^k\int_W\int_0^K \tilde s_{n+1-k}(w,m)\beta(w)dwd\nu(m).
\end{aligned}$$

The action that assigns (x,m) to $\mathbf{x}\setminus\mathbf{a}$ if and only if

$$m \geq \alpha\left[\left(1-p_d\right)\tilde s_n(x,g^{(1)}(m)) - K\int_{b(x,K)\cap W}\beta(w)dw\right]$$

optimises the right hand side and, with

$$\tilde{s}_{n+1}(x,m) = \max\left\{m, \alpha\left(1-p_d\right)\tilde{s}_n(x, g^{(1)}(m)) - \alpha K \int_{b(x,K)\cap W} \beta(w)dw\right\},$$

one sees that

$$v_{n+1}(\mathbf{x}) \geq \tilde{v}_{n+1}(\mathbf{x}) = R \sum_{(x,m)\in\mathbf{x}} \tilde{s}_{n+1}(x,m) + R\sum_{k=1}^{n} \alpha^k \int_W \int_0^K \tilde{s}_{n+1-k}(w,m)\,\beta(w)dwd\nu(m),$$

an observation that completes the induction argument and therefore the proof of the lower bound.

For the upper bound $v_n \leq \hat{v}_n$, use a similar induction argument based on $\hat{s}_n$ but with (13) replaced by the upper bound

$$R\int_W \int_0^K \hat{s}_n(m)\,\beta(w)dwd\nu(m).$$

□

7 Conclusion

In this paper we considered optimal policies for Markov decision problems inspired by forest harvesting. We proved that French thinning is optimal when births follow a Poisson process and marks grow logistically. When the points are required to respect a hard core distance, we derived upper and lower bounds on the discounted total expected reward function for general birth-death-growth dynamics.

In future, it would be of interest to study configuration-dependent asymmetric birth and growth models (Van Lieshout 2008, 2009; Renshaw et al. 2009). Indeed, in a forestry setting, the growth of well-established, large trees may hardly be hampered by the emergence of saplings close by, while it would be harder for young and small trees to flourish near large ones. Moreover, the natural environment, such as the availability of nutrients, might play a role. Finally, refinements of the action space that allow for different thresholds in different mark strata could be investigated.

Acknowledgements Thanks are due to the anonymous referees for their thoughtful remarks.

Funding This work was not funded by any grant.

Code availability The R-scripts used in Sect. 5 are available on request.

Declarations

Conflict of interest The author has no conflict of interest to declare.

References

Baccelli F, Blaszczyszyn B (2009) Stochastic geometry and wireless networks, in two volumes. NOW

Beard M, Vo BT, Vo BN, Arulampalam S (2017) Void probabilities and Cauchy–Schwarz divergence for generalized labeled multi-Bernoulli models. IEEE Trans Signal Process 65:5047–5061

Bertsekas DP (1995) Dynamic programming and optimal control. Prentice and Hall, Hoboken

Bertsekas DP, Shreve SE (1978) Stochastic optimal control: the discrete time case. Academic Press, New York

Chiu SN, Stoyan D, Kendall WS, Mecke J (2013) Stochastic geometry and its applications, 3rd edn. Wiley, New York

Daley DJ, Vere-Jones D (2003) An introduction to the theory of point processes: elementary theory and methods. Springer, Berlin

Daley DJ, Vere-Jones D (2008) An introduction to the theory of point processes: general theory and structure. Springer, Berlin

de Jongh MC, van Lieshout MNM (2022) Testing biodiversity using inhomogeneous summary statistics and global envelope tests. Spat Stat 50:100607

Eckardt M, Moradi M (2024) Marked spatial point processes: current state and extensions to point processes on linear networks. J Agric Biol Environ 29:346–378

Feinberg EA, Schwartz A (2002) Handbook of Markov decision processes. Springer, Berlin

Fransson P, Franklin O, Lindroos O, Nilsson U, Brännström Å (2020) A simulation-based approach to a near optimal thinning strategy: allowing for individual harvesting times for individual trees. Can J For Res 50:320–331

Ghorbani M, Cronie O, Mateu J, Yu J (2021) Functional marked point processes: a natural structure to unify spatio-temporal frameworks and to analyse dependent functional data. TEST 30:529–568

Kelly FP, Ripley BD (1976) On Strauss's model for clustering. Biometrika 63:357–360

Khloussy E, Gelabert X, Jiang Y (2015) Investigation on MDP-based radio access technology selection in heterogeneous wireless networks. Comput Netw 91:57–67

Lee W, Jung BC, Lee H (2020) DeCoNet: density clustering-based base station control for energy-efficient cellular IoT networks. IEEE Access 8:120881

Lieshout MNM van (2009) Sequential spatial processes for image analysis. In: Capasso V (ed) Stereology and image analysis. ECS10-Proceedings of the 10th European congress of ISS

Loosmore NB, Ford ED (2006) Statistical inference using the G or K point pattern statistics. Ecology 87:1925–1931

Lu X, Salehi M, Haenggi M, Hossain E (2021) Stochastic geometry analysis of spatial-temporal performance in wireless networks: a tutorial. IEEE Commun Surv Tutor 23:2753–2801

Matérn B (1986) Spatial variation. Springer, Berlin

Myllimäki M (2009) Statistical models and inference for spatial point patterns with intensity-dependent marks. Ph.D. thesis, University of Jyväskylä

Pretzch H (2009) Forest dynamics, growth and yield. Springer, Berlin

Pukkala T, Miina J (1998) Tree-selection algorithms for optimizing thinning using a distance-dependent growth model. Can J For Res 28:693–702

Pukkala T, Lähde E, Laiho O (2015) Which trees should be removed in thinning treatments? For Ecosyst 2:1–12

Puterman ML (1994) Markov decision processes. Wiley, New York

Renshaw E, Särkkä A (2001) Gibbs point processes for studying the development of spatial-temporal stochastic processes. Comput Stat Data Anal 36:85–105

Renshaw E, Comas C, Mateu J (2009) Analysis of forest thinning strategies through the development of space-time growth-interaction simulation models. Stoch Environ Res Risk Assess 23:275–288

Richards FJ (1959) A flexible growth function for empirical use. J Exp Bot 10:290–300

Rönnqvist M (2003) Optimization in forestry. Math Program Ser B 97:267–284

van Lieshout MNM (2008) Depth map calculation for a variable number of moving objects using Markov sequential object processes. IEEE Trans Pattern Anal Mach Intell 30:1308–1312

Wiegand T, Huth A, Getzin S, Wang X, Hao Z, Savitri Gunatilleke CV, Nimal Gunatilleke IAU (2012) Testing the independent species' arrangement assertion made by theories of stochastic geometry of biodiversity. Proc R Soc B 279:3312–3320