Optimal decision rules for marked point process models

Marie-Colette van Lieshout

CWI / University of Twente

Motivating example: tree harvesting

Many forests are harvested each year for timber. Common strategies include

- felling trees whose diameter at breast height (dbh) exceeds some threshold (French thinning) – supposed to stimulate rejuvenation;
- fell trees whose dbh is smaller than some threshold (German thinning) – supposed to enhance natural selection.

Goal: optimise discounted total expected timber value over a given time horizon.

Markov decision theory

When the system is in state $x \in \mathcal{X}$, the decision maker selects action $a \in A(x)$ earning reward $r(x, a)$. The next state is governed by $p(\cdot|x, a)$.

A policy $\Phi = (\phi_i)_{i=0}^{\infty}$ is a procedure for selecting actions at each decision time $i = 0, 1, 2, \cdots$. If $\phi_i \equiv \phi$, then Φ is stationary.

Write $(X_i, Y_i)_{i=0}^{\infty}$ for states and actions. An optimal policy maximises the α-discounted total expected reward

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

Note: when r is bounded and $\mathcal{X}$ and all $A(x)$ are finite, it suffices to consider only Markov policies (in which actions chosen depend only on current state) that are stationary and deterministic.

Policy iteration and dynamic programming

For current policy $\Phi = (\phi, \phi, \dots)$, find value function v by solving

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

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. Repeat until no further improvement.

Dynamic programming improves the current value function v by

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

until a precision threshold is met. This procedure is amenable to non-finite state and action spaces.

Poisson model with logistic growth – state and action spaces

State space: finite simple marked point patterns $\mathbf{x} = \{(x_i, m_i)_i\}$ on compact set $W \subset \mathbb{R}^2$ with marks in $[0, K]$, $K > 0$ (e.g. wood content for harvesting).

Action spaces: thinnings/subsets of $\mathbf{x}$. Denote retained points by $\phi(\mathbf{x}) \subseteq \mathbf{x}$.

Reward:

$$R \sum_{(x_i, m_i) \in \mathbf{x} \setminus \phi(\mathbf{x})} m_i, \quad R > 0.$$

Poisson model with logistic growth – dynamics

Dynamics: in current state $\mathbf{x}$ under action $\mathbf{a}$,

- delete $(x_i, m_i) \in \mathbf{x} \setminus \mathbf{a}$;
- independently of other points, let $(x_i, m_i) \in \mathbf{a}$ die with probability $p_d \in (0, 1)$ and otherwise grow to

$$\left(x_i, \frac{K}{1 + e^{-\lambda} \left(\frac{K}{m_i} - 1 \right)} \right), \quad \lambda > 0;$$

- add a Poisson process on W with intensity $\beta > 0$ and marked i.i.d. according to probability measure ν on $[0, K]$.

Poisson model with logistic growth – discounted rewards

Theorem: For $0 \le \alpha < 1$, French thinning with threshold

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

(with 0/0=0) is optimal and has α-discounted total expected reward

$$v^*_\alpha(\mathbf{x}) = R \sum_{(x_i, m_i) \in \mathbf{x}} s(m_i) + \frac{\alpha R \beta |W|}{1 - \alpha} \int_0^K s(m) d\nu(m)$$

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

Note: modifications to finite horizons and location dependent β exist.

A comparison with German thinning

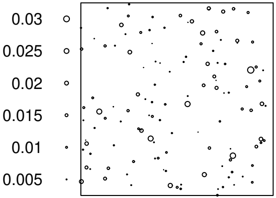

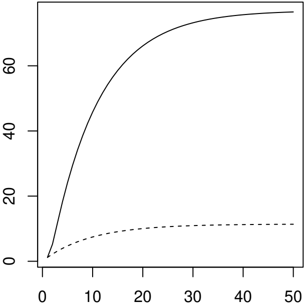

Left: sample $\mathbf{x}$ from a Poisson($\beta = 5$, $W = [0, 5]^2$) process marked by $\nu =$Beta($\lambda_1 = 2$, $\lambda_2 = 20$) on $[0, 0.1]$.

Right: graphs of the finite horizon 0.9-discounted total expected reward $v_n(\mathbf{x})$ against n for optimal French (solid line) and best German (dotted line) thinning. Here $R = 1$, $\lambda = 2$ and $p_d = 0.05$.

Hard core models with interaction

Modification: allow only trees that can grow to their full size.

Dynamics: in current state $\mathbf{x}$ under action $\mathbf{a}$,

- delete $(x_i, m_i) \in \mathbf{x} \setminus \mathbf{a}$;
- independently of other points, let $(x_i, m_i) \in \mathbf{a}$ die with probability $p_d \in (0, 1)$ and otherwise grow to

$$\left(x_i, \frac{K}{1 + e^{-\lambda \left(\frac{K}{m_i} - 1 \right)}} \right), \quad \lambda > 0;$$

- add a hard core(K) process on W with intensity $\beta > 0$ and marked i.i.d. according to probability measure ν on $[0, K]$ and remove all points that fall within distance K to a point in $\mathbf{a}$.

Hard core model with logistic growth – discounted rewards

Theorem: For $0 \leq \alpha < 1$ and $\mathbf{x}$ respecting hard core K, the optimal α-discounted total expected reward $v_\alpha(\mathbf{x})$ satisfies

$$\tilde{v}_\alpha(\mathbf{x}) \leq v_\alpha(\mathbf{x}) \leq \hat{v}_\alpha(\mathbf{x}),$$

where $\tilde{v}_\alpha$ and $\hat{v}_\alpha$ take the form

$$R \sum_{(x,m)\in\mathbf{x}} s(x,m) + R \sum_{k=1}^{n-1} \alpha^k \int_{W\times[0,K]} s(w,l)\beta dw d\nu(l)$$

with integrand s given by, respectively,

$$\tilde{s}(x,m) = \sup_{n\in\mathbb{N}_0} \left\{ \frac{K\alpha^n(1-p_d)^n}{1+e^{-\lambda n}\left(\frac{K}{m}-1\right)} - \alpha K \beta |W \cap b(x,K)| \sum_{i=0}^{n-1} \alpha^i (1-p_d)^i \right\}$$

and

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

Sketch of proof

Use dynamic programming and bound integrals of the form

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

from below by

$$\int_{W\times[0,K]} s(w,l)\beta dw d\nu(l) - K \sum_{(x,m)\in\phi(\mathbf{x})} \beta|W\cap(b(x,K)|$$

and from above by $\int_{W\times[0,K]} s(w,l)\beta dw d\nu(l)$.

Note: modifications to finite horizons and location dependent β exist.

Tightness of bounds

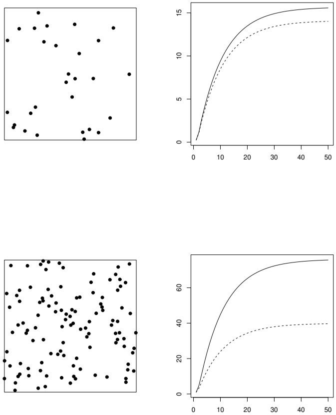

Samples **x** from a hard core ($\beta = 1$ resp. $\beta = 4.3$, $W = [0,5]^2$) process marked by $\nu =$Beta($\lambda_1 = 2$, $\lambda_2 = 20$) on $[0, 0.1]$ and upper and lower bounds of the finite horizon 0.9-discounted total expected reward ($R = 1$, $\lambda = 2$ and $p_d = 0.05$).

Thank you for your attention!

M.N.M. van Lieshout (2024)
Optimal decision rules for marked point process models.
Stochastic Envirnomental Research and Risk Assessment.