Impact of climate change uncertainty on optimal forest management policies at stand level

S. Backéus[1], L.O. Eriksson[1] and F. Garcia[2]

[1]SLU, Umea, Sweden, [2]INRA UBIA, Toulouse, France, E-Mail: fgarcia@toulouse.inra.fr

Keywords: *climate change; forest management planning under uncertainty; reinforcement learning*

EXTENDED ABSTRACT

Forest management is classically modeled as a deterministic planning problem where decisions are determined in advance unconditionally to future events. This approach has led to the development of efficient optimization methods based on linear programming which can solve large forest management problems as encountered in Scandinavian countries. Recent research results have nevertheless established that it could be essential from an economical point of view to consider stochastic phenomena in the definition of long term forest management problems such as natural hazards and market uncertainties.

Climate change will have important effects on forest ecosystems in the long run. It is a new fundamental reason to go beyond deterministic planning approaches in forest management. According to the last scenarios published by IPCC, the average temperature increase in the world might be in the range of 1.4 to 5.8° by year 2100. At regional level, a model developed by SWECLIM for Sweden predicts an increase between 2.5°C to 4.5°C by 2100, but there is also a considerable uncertainty regarding the future climate trajectory.

In the study presented in this paper, we investigated the impact of such climate change uncertainty on the solutions of forest management problems for typical Swedish stands with several species. Our main objective was to determine whether taking into account uncertainty can improve the best deterministic solutions obtained by considering average temperature change scenarios. Such stochastic planning problems are theoretically solved by dynamic programming or related methods. However, these approaches can rarely be employed when complex growth and yield models are used. The second methodological objective of this study was thus to assess the effectiveness of new alternative stochastic simulation methods like simulation-based optimization or reinforcement learning for solving forest management planning problems with stochastic features.

In the work described in this paper, we analyzed approximate optimal policies obtained for forest management problems under climate variability and change. These problems were defined on the basis of a forest management model called GAYA and a climate change model developed for the purpose of this study. Approximate optimal policies were obtained with Linear-Q-learning(λ), a reinforcement learning algorithm. They were compared with approximate optimal plans obtained by considering average climate change scenarios.

We studied fictive stands with one, two and three species. The stands were located in southern and northern Sweden and contained the species pine, spruce and birch. The stands were simulated for 20 five-year long periods years. Reinforcement learning converged toward approximate optimal policies after few simulations, even for complex stands with several species, either for problems with deterministic or stochastic scenarios.

Ours results showed that there was a gain in considering stochastic models, but also that the magnitude of this gain was small. Considering climate change uncertainty improved the value of approximate optimal programs, but deterministic plans was sometime still optimal. This important conclusion can be explained by the relatively small estimated value of the growth effect, that lead to a quasi-deterministic dynamics of the stand state, and by the symmetry of the temperature change distribution around the average scenario.

Our analysis lies entirely on the stochastic climate change model we developed, based on the SWECLIM regional climate modeling for northern Europe. Its main limitation is its stationary assumption on probabilities of the future climate change scenarios. Modeling today the probable fall of uncertainty in the future may be a complicated task. Note however that such a decreasing uncertainty model should strengthen our present conclusions.

1. INTRODUCTION

Forest management is classically modeled as a deterministic planning problem where decisions are determined in advance unconditionally to future events. This approach has led to the development of efficient optimization methods based on linear programming which can solve large forest management problems as encountered in Scandinavian countries. Recent research results (see for instance Valsta (1992) or Lohmander (2001) have nevertheless established that it could be essential from an economical point of view to consider stochastic phenomena in the definition of long term forest management problems such as natural hazards and market uncertainties.
Climate change will have important effects on forest ecosystems in the long run. It is a new fundamental reason to go beyond deterministic planning approaches in forest management. According to the last scenarios published by IPCC (2002), the average temperature increase in the world might be in the range of 1.4 to 5.8° by year 2100. At regional level, a model developed by SWECLIM for Sweden predicts an increase between 2.5°C to 4.5°C by 2100 (Rummukainen, 2003), but there is also a considerable uncertainty regarding the future climate trajectory.

In the study presented in this paper, we investigated the impact of such climate change uncertainty on the solutions of forest management problems for typical Swedish stands with several species. Our main objective was to determine whether taking into account uncertainty can improve the best deterministic solutions obtained by considering average temperature change scenarios. Such stochastic planning problems are theoretically solved by dynamic programming or related methods. However, these approaches can rarely be employed when complex growth and yield models are used. The second methodological objective of this study was thus to assess the effectiveness of new alternative stochastic simulation methods like simulation-based optimization or reinforcement learning for solving forest management planning problems with stochastic features.

2. OPTIMIZING FOREST MANAGEMENT UNDER CLIMATE VARIABILITY

The forest management problem we considered in this paper is defined at the stand level. We assume here that the dynamics of the forest stand is influenced by thinning actions over a finite horizon H, given some evolution of average temperature due to climate change.

Let us denote by s_t the state of the stand at time t, a_t the action applied on this stand at time t, and T_t the average observed temperature change over the period $[t,t+1]$. The general equation of the state dynamics is :

$$\forall t \leq H, \quad s_{t+1} = f\left(s_t, a_t, T_t\right)$$

where s_1 is the initial state of the stand. At each time period, depending on the state and the action that has been applied, a financial reward $r_t(s_t,a_t,s_{t+1})$ that corresponds to wood production gains minus operation costs is received.
The objective of the planning problem is then to maximize over $t = 1,\ldots,H$ the value R_H of the discounted sum of future rewards

$$R_H = r_1 + \gamma r_2 + \cdots + \gamma^{H-1} r_H .$$

2.1. Optimal plans

When one assumes that the sequence $T=(T_t)_{t=1,\ldots,H}$ is known in advance, the optimization problem consists in maximizing the sum R_H over the set of all possible plans $(a_1,\ldots,a_H)$:

$$R_H^* = \max_{(a_1,\ldots,a_H)} r_1 + \gamma r_2 + \cdots + \gamma^{H-1} r_H .$$

In that case, a solution is a deterministic plan $P^* = (a^*_1,\ldots,a^*_H)$.

When the scenario T is not known in advance, as it is the case with climate change, the value of such deterministic plans P must be defined by considering the expected value V_H of the sum of discounted rewards, with regard to the probability distribution of T:

$$V_H(P) = E\left[r_1 + \gamma r_2 + \cdots + \gamma^{H-1} r_H \middle| P\right].$$

Optimizing this expected value is generally very costly. Approximate optimal plans $\overline{P}^*$ can be obtained by solving the transformed optimization problem:

$$\max_{(a_1,\ldots,a_H)} r_1 + \gamma r_2 + \cdots + \gamma^{H-1} r_H$$

for the average temperature scenario $\overline{T}$:

$$\left(\overline{T}_t\right)_{t=1,\ldots,H} = E\left[\left(T_t\right)_{t=1,\ldots,H}\right].$$

The idea here is thus to look for the sequence of decisions $\overline{P}^*$ that maximizes the global revenue, assuming an average temperature scenario $\overline{T}$ occurs, instead of looking for the plan P^* that maximizes the average value of the global revenue, assuming the temperature scenario T is uncertain. This approach is fully justified when fluctuations of T around $\overline{T}$ are small, with thus $\overline{P}^* \approx P^*$.

2.2. Optimal policies

These approaches for solving planning problems in an uncertain context correspond to the classical modelling of forest management, where management decisions like plantation, thinning, etc. can be determined in advance independently of future events. However, forest management planning under uncertainty can be solved more efficiently by exploiting the fact that decision makers adapt their behavior over time as new information is revealed about the current state of the forest ecosystem and the economical context (Eriksson and Backéus, 2003).

This leads to formalizing the problem as an optimal control problem, where optimal solution is no more a priori a deterministic plan P, but instead adaptive policy $\pi = (\pi_t)_{t=1,\ldots,H}$ that maps the state of the system to possible actions at time t:

$$\forall t \leq H, \quad a_t = \pi_t(s_t).$$

Depending on the scenarios T, a policy will generate different sequences of states, actions and rewards. Like for deterministic plans, the value of a policy is then defined as the expected value of the sum of discounted rewards

$$V_H(\pi) = E\left[r_1 + \gamma r_2 + \cdots + \gamma^{H-1} r_H \,\middle|\, \pi\right].$$

The optimal control problem thus consists in determining the policy π^* that maximizes $V_H(\pi)$.

Although this planning under uncertainty approach is theoretically optimal, its complexity is very high and the standard exact optimization methods for solving stochastic problems – stochastic optimization, stochastic programming (Birge and Louveaux, 1997), Markov decision process models (Buongiorno, 2001) – can only be applied to single stand or small forest problems. For larger and more complex problems, approximate methods like simulation-based optimisation methods (Gosavi, 2001) or reinforcement learning methods (Bertseklas and Tsitsiklis, 1996; Sutton, 1998), can be applied. These approaches generate approximate optimal policies $\tilde{\pi}^*$, with $V_H(\tilde{\pi}^*) \leq V_H(\pi^*)$.

2.3. Plans or policies?

Like deterministic plans are just special instances of policies, we know that the optimal policies are always at least as good as optimal plans:

$$\forall P, \quad V_H(\pi^*) \geq V_H(P),$$

with in particular $V_H(\pi^*) \geq V_H(P^*)$ and $V_H(\pi^*) \geq V_H(\overline{P}^*)$. For deterministic problems, with no uncertain factors, optimal policies are optimal plans. For near-deterministic problems, value of optimal plans and optimal policies can have very close values. It is more difficult to compare a priori $V_H(\tilde{\pi}^*)$ with $V_H(P^*)$ and $V_H(\overline{P}^*)$: in the general case, it can occurs that optimal plans perform better than approximate optimal policies, depending on the approximations that are made.

In the work described in this paper, we analyzed approximate optimal policies obtained for forest management problems under climate variability and change. These problems were defined on the basis of a forest management model called GAYA and a climate change model developed for the purpose of this study. Approximate optimal policies were obtained by reinforcement learning. They were compared with approximate optimal plans obtained by considering average climate change scenarios.

3. FOREST STAND SIMULATOR

Management programs were generated by the GAYA stand simulation system (Eriksson, 1983). The system confers a realization of the state dynamics for growth periods with a length of five years based on forest state, action, and temperature (see the implementation of temperature below). Up to three species in the same stand can be handled. The system also yields revenue as a function of state and action.

Basal area increment and natural mortality is estimated with functions by Agestam (1985). Based on this data timber volume is calculated (Agestam, 1985). Since projections are performed only for established stands, i.e. stands with a average height of about 8 meters, projection functions for young stands are not needed. Timber revenues are derived with bucking functions by Ollas (1980) and with the timber and pulpwood price list of the SCA forest company of year 2001 (SCA SKOG AB, 2002). Costs are estimated with functions from SLU (1990), where unit costs of harvester is set to 800 and 700 SEK per hour in final felling and thinning, respectively, and of forwarder to 600 SEK and 500 SEK per hour in final felling and thinning, respectively.

Four different actions can be considered in each time period. They are: do nothing, in which case the net revenue is zero, clear felling, in which case all trees are removed, and light and heavy thinning. The light thinning involves the removal of 20 percent of the volume and the heavy thinning the removal of 30 percent. The thinning is performed such that the same proportion of basal area and no of stems is removed. Do nothing is always possible, clear felling can be done when the stand has reached 80 years, and thinning is permissible as long as the number of stems after thinning is at least 400.

4. CLIMATE CHANGE MODEL

A stochastic climate model was created to predict the temperature increase in every five-year period for the forest simulator. The temperature increase 100 year ahead was assumed to be between 2.4°C and 4.5°C according to the regional climate modeling for northern Europe done by SWECLIM (Rummukainen, 2003). The inter periodic variation was assumed to have the same pattern as past climate in Sweden. The temperature in time t was

$$T_t = Tr + \rho T_{t-1} + \sigma\sqrt{1-\rho^2}\, z_t$$

where Tr was the temperature trend, z_t was a independent normally distributed random number, ρ was the autocorrelation and σ the standard deviation. The standard deviation and autocorrelation were calculated ($\rho = 0.8$, $\sigma = 0.55$) from temperature data from 1860 to 2003 based on average data from 37 fairly evenly distributed weather stations in four regions in Sweden (Alexandersson, 2002 and Alexandersson pers. com.). We then simulated an accumulated temperature trend as

$$Tr = Tr_{\min}(1-U) + Tr_{\max}U$$

where Tr_{min} (0.12) and Tr_{max} (0.22) was the assumed minimum and maximum temperature change per period according to SWECLIM, and U was a uniform random number [0, 1].

Temperature changes are assumed to affect basal area growth only. Thus, the effect mediated by the temperature change was expressed as

$$BI_{ti} = B0_{ti} \cdot (1 + T_t \cdot GE_i)$$

where $B0_t$ is the calculated basal area growth over five year period t of species i without consideration to temperature change and GE_i is the growth effect per degree of temperature change. GE_i was set according to Bergh et al. (2003) to 0.065 for pine and birch and to 0.05 for spruce.

5. REINFORCEMENT LEARNING

Since the dynamics of the stand is complex and can only be simulated, we decided to solve that planning problem under uncertainty by using reinforcement learning, a simulation-based control approach (Sutton, 1998).

The idea is to learn by simulation a close approximation of the optimal value function of the problem $Q_t^*(s_t,a_t)$, $t=1,\ldots,H$, from which a optimal policy can be deduced:

$$\pi_t^*(s_t) = \arg\max_a Q_t^*(s_t, a).$$

The well known reinforcement learning algorithm Q-learning consists in simulating N trajectories and updating the $Q_t(s,a)$ estimates of $Q_t^*(s_t,a_t)$ after each observed transition of the system from state s_t to state s_{t+1} given that action a_t has been applied and that the reward r_t has been received:

$$Q_t(s_t, a_t) \longleftarrow Q_t(s_t, a_t) + \alpha\Delta_t$$

with

$$\Delta_t = r_t + \gamma \max_a Q_{t+1}(s_{t+1}, a) - Q_t(s_t, a_t),$$

and $Q_{H+1} = 0$. In that equation, Δ_t is called the temporal difference (TD) error, γ is the discount factor of the problem and α is the learning rate. This learning rate can vary with k and t, and decreasing with the number of simulated transitions.

For large or continuous domains, these $Q_t(s,a)$ values cannot be individually learned. A classical approach for estimating Q_t^* is to then use a linear approximation structure

$$Q_t(s_t, a_t) = \sum_{k=1}^{K} \omega_t^k \phi_t^k(s_t, a_t)$$

where the functions ϕ_t^k represent special features of the problem that are important for defining a good quality control policy.

The convergence of Q-learning can also be accelerated by considering a new TD error:

$$\Delta_t^{\lambda} = \gamma\lambda\Delta_{t+1}^{\lambda} + \Delta_t ,$$

where λ is a parameter of the method, $\lambda \in [0,1]$ and $\Delta^{\lambda}_{H+1} = 0$.

A natural implementation of the algorithm, called Linear-Q-learning(λ) (Garcia, 1999), is based on a backward iteration from $t = H$ to $t = 1$, after the H transitions have been observed:

Linear-Q-learning(λ)

1. Initialize weights $\omega = 0$
2. For $n = 1$ to N
   - a. Initialize s_1
   - b. For $t = 1$ to H
     - Choose a_t
     - $(s_{t+1}, r_t) = \text{Simulate}(s_t, a_t)$
   - c. For $t = H$ to 1
     - $\omega_t^k \longleftarrow \omega_t^k + \alpha\Delta_t^{\lambda}\phi_t^k(s_t, a_t)$
     - $k = 1,\ldots,K$
3. Return ω

The choice of actions a_t during the learning process is important since it is the only way of controlling on-line the algorithm. The *greedy* action $a_t = \text{argmax}_a\, Q_t(s_t,a)$ is a good choice but other actions need also to be visited in order to explore the action domain. A classical approach is to choose the greedy action with a probability 1-ε, and a random action with probability ε, where ε is a parameter of the algorithm.

There exist some theoretical results (Bertseklas and Tsitsiklis, 1996) that prove the convergence of Linear-Q-Learning to the closest linear approximation to $Q_t^*(s_t,a_t)$. In practice, the parameters α, ε, λ and the features ϕ_t^k have to be fixed correctly in order to define efficient algorithms.

6. RESULTS

We studied fictive stands with one, two and three species. The stands were located in southern and northern Sweden and contained the species pine, spruce and birch (Table 1). The productivity was medium. The stands were simulated for 20 five-year long periods years and the interest rate γ was set to three percent.

Table 1. Initial state of the three stands used in the study

		Stand 1	Stand 2	Stand 3
Number of species		1	2	3
Lat., Alt.		64°200m	58° 0m	58° 0m
Age (years)		27	28	29
Pine	Nb. of stems	2800	1400	1400
	Basal area (m^2)	24.3	13	8.2
Spruce	Nb. of stems	-	1400	700
	Basal area (m^2)	-	5.1	5.5
Birch	Nb. of stems	-	-	700
	Basal area (m^2)	-	-	5.5

In these stand planning problems, the state s_t was correctly described by two state-variables ns_t^i (number of stems) and ba_t^i (basal area) for the three species $i = 1,\ldots,3$. Actions were also decomposed into three action variables, one per each species: $a_t = (a_t^1, a_t^2, a_t^3)$, with $a_t^i \in \{1,2,3,4\}$ (1 = do nothing, 2 = clear felling, 3 and 4 = light and heavy thinning as given in section 3).

For each stand, we calculated and compared three approximate solutions. Approximate optimal plans $\overline{P}^*$ were first calculated for the average temperature scenario defined by $\overline{T}_t = tTr_{mean}$, where $Tr_{mean} = 1/2(Tr_{\min} + Tr_{\max})$ is the average rate of temperature change. We used a systematic optimization approach. Due to the large size of the search space (64^{20} possible plans for the 3-species stand), we were not able to evaluate all the candidates and we optimized plans where thinning and clear felling occur at the same time for all species ($a_t^1 = a_t^2 = a_t^3$).

We then applied Linear-Q-Learning(λ) for calculating approximate optimal policies $\overline{\pi}^*$ and $\tilde{\pi}^*$, considering respectively the deterministic scenario $\overline{T}$ and the stochastic scenarios T. Best results were obtained from a linear representation of the Q-value function

$$Q_t(s_t, a_t) = \sum_{i=1}^{3} Q_t^i\left(ns_t^i, ba_t^i, a_t^i\right),$$

$Q_t^i(ns, ba, a) = \omega_t^{i,1,a} ns + \omega_t^{i,2,a} ba + \omega_t^{i,3,a}$ (for a time horizon set to $H = 20$, the number of weights to learn was thus equal to 3x20x3x4=720).

Linear-Q-Learning(λ) was applied with N = 50,000 simulated trajectories. Parameters were experimentally tuned and set to α = 0.1, ε = 0.2 and λ = 1.0. The value of these approximate solutions $\overline{P}^*$, $\overline{\pi}^*$ and $\tilde{\pi}^*$ was estimated on N = 1000 simulated stochastic climate change scenarios (Table 2).

Table 2. Expected value of approximate optimal solutions for the stochastic temperature scenario

	Expected net revenue (SEK)		
Solution	Stand 1	Stand 2	Stand 3
$\overline{P}^*$	14688	12185	8644
$\overline{\pi}^*$	14606	15059	11189
$\tilde{\pi}^*$	14711	15291	11515

For the three stands we observe that $V_H(\tilde{\pi}^*) \geq V_H(\overline{\pi}^*) \geq V_H(\overline{P}^*)$, except for Stand 1 where $V_H(\overline{\pi}^*) \leq V_H(\overline{P}^*)$. These results show that i) reinforcement learning allowed us to generate efficiently (50000 simulations are run in few minutes) approximate optimal solutions; ii) taking into account climate change uncertainty in the optimization process improved slightly the quality of the solutions.

Best results of the reinforcement learning method were obtained for Stands 2 and 3. One explanation could be that for these two cases, Linear-Q-Learning(λ) was able to learn different policies for each species, which was not the case for the systematic optimization approach.

We analysed the structure of the solutions obtained with Linear-Q-Learning(λ). For the one species stand, $\overline{\pi}^*$ and $\tilde{\pi}^*$ were adaptive policies (Figure 1) but $\overline{\pi}^*$ lead to the same sequence of action than $\overline{P}^*$ for the deterministic scenario $\overline{T}$.

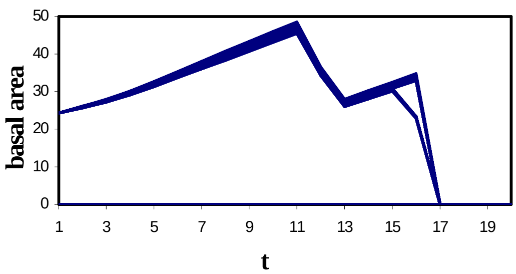

Figure 1. Basal area (Stand 1) under $\tilde{\pi}^*$.

For Stand 2, $\overline{\pi}^*$ and $\tilde{\pi}^*$ were adaptive policies, $\tilde{\pi}^*$ being quasi-deterministic (Table 3). For Stand 3, $\overline{\pi}^*$ was an adaptive policy but $\tilde{\pi}^*$ was a deterministic plan.

Table 3. Structures of the approximate optimal plans and policies plans for Stand 2. The percent figure in adaptive policies indicates how often the subsequence was chosen on 100 trajectories.

Solution	(a_t^1, a_t^2)		
$\overline{P}^*$	11111111144141311121		
$\overline{\pi}^*$	1111114444111	12 (23%)	
		21 (77%)	
	1111111111111	313412 (2%)	
		14	331 → 12 (27%)
			331 → 21 (19%)
			11112 (50%)
$\tilde{\pi}^*$	111113444	44444341112 (97%)	
		31444341112 (3%)	
	111111111111414	11112 (97%)	
		31112 (3%)	

7. DISCUSSION

We applied reinforcement learning for computing solutions to forest planning problems within an adaptive and stochastic framework. This optimization method converged toward approximate optimal policies and required few simulations, even for complex stands with several species. The quality of both $\overline{\pi}^*$ and $\tilde{\pi}^*$ solutions also indicates that reinforcement learning was quite efficient either for solving problems with deterministic or stochastic scenarios.

There exist many options for adapting forest management to climate change. Our main objective in this paper was to assess the importance of considering uncertainty in climate change when one designs management programs at stand level. Ours results, based on three fictive stands, showed that there was a gain in considering stochastic models, but also that the magnitude of this gain was small. Considering climate change uncertainty improved the value of approximate optimal programs, but deterministic plans were sometime still optimal. This important conclusion can be explained by the relatively small estimated value of the growth effect (GE = 0.065 for pine), that lead to a quasi-deterministic dynamics of the stand state (See Figure 1), and by the symmetry of the temperature change distribution T around $\overline{T}$.

Our analysis lies entirely on the stochastic climate change model we developed, based on the SWECLIM regional climate modeling for northern Europe by (Rummukainen, 2003). Its main limitation is its stationary assumption, as it assigns constant probabilities to future climate change scenarios. Modeling today the probable fall of uncertainty in the future may be a complicated task. Note however that such a decreasing uncertainty model should strengthen our present conclusions.

8. ACKNOWLEDGMENTS

We thank Sören Holm at the department of Forest Resource Management and Geomatics, SLU, for assistance with statistical analyses of climate data, and Hans Alexandersson (SMHI; Sweden) for providing climate data.

9. REFERENCES

Agestam, E. 1985. En produktionsmodell för blandbestånd av tall, gran och björk i Sverige (A growth simulator for mixed stands of pine, spruce and birch in Sweden). Dept. of Forest Production, Swed. Univ. of Agri. Sciences, Umeå. Report No. 15. 150 pp. Eng summary.

Alexandersson, H. (2002), Temperatur och nederbörd i Sverige 1860-2001, Temperature and precipitation in Sweden 1860-2001. SMHI, Norrköping (in Swedish with English summary).

Bergh, J., Freeman, M., Sigurdsson, B., Kellomäki, S., Laitinen, K., Niinisto, S., Peltola, H., Linder, S. (2003), Modeling the short-term effects of climate change on the productivity of selected tree species in Nordic countries. *Forest Ecology and Management* 183, 327-340.

Bertsekas, D. P. and Tsitsiklis, J. N. 1996. Neuro-Dynamic Programming. Athena Scientific. Belmont (MA)

Birge, J.R. and Louveaux, F. 1997. Introduction to Stochastic Programming. Springer. New York.

Buongiorno, J. 2001. Generalization of Faustmann's Formula for Stochastic Forest Growth and Prices with Markov Decision Process Models. For. Sci. 47(4), 466-474.

Eriksson, L.O., 1983. Long range forestry planning—A case study of a forest in a transition period (Långsiktig skoglig planering—En studie av ett skogsinnehav i förändring). Inst för skogsteknik. Swed. univ. agr. sci., section of forest technology, Garpenberg, Sweden. Report 154.

Eriksson, L.O. and Backéus, S. 2003. Stochasic modelling at the forest level. In: Proceedings of EFITA2003 - Conference Papers 5-9 July 2003, Debrecen-Budapest, Hungary

Garcia, F. 1999. Use of reinforcement learning and simulation to optimize wheat crop technical management. In Proc. of the international congress on modeling and simulation (MODSIM'99), 801-806, Hamilton, New-Zealand.

Gosavi, A. 2003. Simulation-Based Optimisation: Parametric Optimisation Techniques and Reinforcement Learning. Kluwer Academic Publishers.

IPCC, 2002. Climate Change and Biodiversity. (Ed) Gitay, H., Suárez, A., Watson, R.T. and Dokken, D.J. IPCC Technical Paper (www.ipcc.ch/pub/techrep.htm).

Lohmander 2001. On Risk and Uncertainty in Forest Management Planning Systems. In Proceedings of SNS Meeting: Nordic Trends in Forest Inventory, Management Planning and Modelling, Solvalla, Finland.

Ollas, R. 1980. Nya utbytesfunktioner för träd och bestånd (New yield functions for trees and stands). Forskningsstiftelsen Skogsarbeten, Ekonomi no 5. 4 pp. Stockholm. In Swedish.

Rummukainen, M., 2003. The Swedish regional climate modeling program, SWECLIM, 1996-2003. Final Report. Rep. No. 104, Norrköping, Sweden.

SCA SKOG AB, 2002, "Virkesprislista—AA0301", (In Swedish). <http://www.skog.sca.com/se/koporganisation/prislistor/prislista_virkesyd.asp>, (030825)

SLU, 1990. Områdeskalkyler - skog :mellersta Sverige 1990/91. Sveriges lantbruksuniversitet, SLU info rapporter, Allmänt nr.165. Uppsala.

Sutton, R.S. and Barto, A.G., 2001, *Reinforcement Learning: an introduction*, MIT Press, Cambridge.

Valsta, L. 1992. A scenario approach to stochastic anticipatory optimization in stand management. Forest Sci. 38: 430-447.