

Forest management decisions for wildlife objectives: system resolution and optimality

Clinton T. Moore ^{a,*}, Michael J. Conroy ^b, Kevin Boston ^c

^a Georgia Cooperative Fish and Wildlife Research Unit, Warnell School of Forest Resources,
University of Georgia, Athens, GA 30602, USA

^b US Geological Survey, Biological Resources Division,
Georgia Cooperative Fish and Wildlife Research Unit, Warnell School of Forest Resources,
University of Georgia, Athens, GA 30602, USA

^c Warnell School of Forest Resources, University of Georgia, Athens, GA 30602, USA

Abstract

Managers of forest wildlife populations make recurring management decisions based on incomplete knowledge of system states. For example, animal population estimates may ignore spatial structure that may influence population viability. We built a spatially-explicit model for a population of birds in a forested landscape. Rates of bird population growth within forest compartments and rates of bird dispersal among compartments were functions of stand age and basal area, compartment population size, and inter-compartment distance. Stand characteristics were imbedded in a dynamic model and assumed perfectly observable and under the complete control of managers. We constructed a genetic algorithm to search for the schedule and spatial distribution of silviculture to maximize total bird abundance at the end of a fixed planning horizon, under combinations of initial habitat and population distribution. We also found policies for a smaller set of population distributions that a manager may only presume to occur (e.g. birds equally distributed among stands), as when managers are only able to observe abundance and not spatial distribution. We investigated the effect of this loss of system resolution on optimality by examining differences in projected population sizes under the two types of policies. That is, we used the set of ‘presumed-state’ policies to project population size from each true initial system state, then we compared these to projections under the best policy for that state. For the planning horizon that we considered, loss in optimality was highly dependent on initial habitat state and on choice of presumed population distribution. Generally, loss in optimality and species extinction rate were both greater for habitat states that were initially poor than initially favorable. For some

* Corresponding author.

E-mail address: cmoore@smokey.forestry.uga.edu (C.T. Moore)

initial habitat states, population projections based on policies for presumed states often exceeded objective function values for known-state policies, suggesting that the genetic algorithm frequently fell short of finding bona fide optima. © 2000 Elsevier Science B.V. All rights reserved.

Keywords: Genetic algorithm; Optimal scheduling; Forest management; Wildlife management; Dynamic system; Spatially-explicit population model

1. Introduction

Management for many forest animal species, particularly forest birds, centers on silvicultural actions carried out over space and time. Optimal scheduling and optimal control techniques are finding greater application in silvicultural planning where objectives and constraints for wildlife populations or their habitats are explicitly considered (Williams, 1989; Bettinger et al., 1997; Hof and Raphael, 1997; Hughell, 1997).

Wildlife populations, however, are notoriously difficult to measure. Populations are rarely completely enumerated. In many cases, population size estimates derived from a sampling program must suffice in lieu of real measurements, and in many more cases, all that may be available are indicators of presence or absence of a species (Conroy and Noon, 1996). Furthermore, many populations follow birth, death, and movement dynamics so rapid that implementation of the action occurs at a state substantially different than the one at the time of measurement.

A specific problem of wildlife population measurement is that of measuring population parameters at spatial scales that are relevant to population dynamics. For example, all that a manager may be able to discern is that two populations comprise 1000 animals each. Population dynamics and expected persistence for these populations may be considerably different if one population is distributed over 20 proximate habitat patches and the other is distributed over two isolated patches. Consequently, management approaches for each population, given a common management objective, might be entirely different if the full state of the system was known; however, this is often not the case.

Thus, managers are frequently required to make management decisions based on aggregated system measurements, even as these systems follow finer resolutions of spatial dynamics. Our objective was to estimate the loss in optimality incurred, if any, when one is forced to manage in this way, and to determine whether losses vary as a function of initial states of the system.

2. Model development

We built a deterministic model of a forest ecosystem comprised of two submodels. The forest submodel projected stand characteristics through time and in response to management actions. The population submodel projected abundance of a forest bird species in response to forest stand characteristics. Our model most

closely resembled population dynamics for bird species that prefer mature, low-density southern US pine (*Pinus* spp.) forests, e.g. the endangered Red-cockaded woodpecker (*Picoides borealis*) (Ligon et al., 1986; Walters, 1991) and the potentially threatened Bachman's sparrow (*Aimophila aestivalis*) (Dunning and Watts, 1990; Pulliam et al., 1992). However, our purpose was to investigate suboptimal management in a general, non-specific forest setting. Therefore, we chose model forms and parameterizations that we found convenient to work with yet reasonably characteristic of a forest system: they are not necessarily representative of any known system.

2.1. Forest growth and regeneration model

We modeled a square forest landscape as a grid of four equal-sized (1 square land unit) square compartments or stands. We numbered the compartments $s = 1, 2, 3, 4$ in a clockwise direction starting from the northwest compartment. We assumed that each compartment is managed on an even-aged basis. At any 5-year decision stage, the manager chooses to either (0) do nothing, (1) thin to a fixed basal area, or (2) clearcut the stand. In any event, the stand is clearcut when it reaches age 120 years.

At any year t and compartment s , three state variables describe the system: (1) stand age a_{ts} , (2) stand age h_{ts} at last thinning, and (3) population size n_{ts} of a resident forest bird. However, we assume that a manager measures the 12 state variables only at times $t \in \{0, 5, 10, \dots\}$. A stand younger than a fixed threshold age $a_{TT} = 30$ years cannot be thinned. A logistic function relates stand basal area b_{ts} to age of stand and to age of stand at last thinning. Thinning reduces stand basal area to that corresponding to a_{TT} , and we represented an unthinned stand by setting $h_{ts} = a_{TT}$. Because management interventions over the life of a stand may abruptly change h_{ts} , the basal area growth function for a stand is discontinuous in a_{ts} :

$$b_{ts} = \begin{cases} K_0 / (1 + \exp(-r_0[a_{TT} - h_{ts} + a_{ts} - a_1])), & h_{ts} > a_{TT} \\ K_0 / (1 + \exp(-r_0[a_{ts} - a_1])), & h_{ts} \leq a_{TT} \end{cases}$$

Here, $K_0 = 27.55 \text{ m}^2 \text{ ha}^{-1}$ ($120 \text{ ft}^2 \text{ acre}^{-1}$) is the asymptotic basal area, $r_0 = 0.1 \text{ year}^{-1}$ is the basal area maximum growth rate, and $a_1 = 30$ years is age of stand at curve inflection point.

2.2. Bird population model

We modeled a compartment-specific bird population growth rate λ_{ts} as a quadratic function of stand age and basal area

$$\lambda_{ts} = \lambda_{\max} + a_0(a_{ts} - a_x)^2 + b_0(b_{ts} - b_x)^2,$$

where $\lambda_{\max} = 1.15 \text{ year}^{-1}$ is population maximum growth rate, $a_0 = -3.0 \times 10^{-5} \text{ year}^{-3}$ is the curvature coefficient for stand age, $a_x = 120$ years is stand age at λ_{\max} , $b_0 = -1.318 \times 10^{-7} \text{ year}^{-1} \text{ m}^{-4} \text{ ha}^2$ ($-1.0 \times 10^{-4} \text{ year}^{-1} \text{ ft}^{-4} \text{ acre}^2$) is the

curvature coefficient for basal area, and $b_x = 13.77 \text{ m}^2 \text{ ha}^{-1}$ ($60 \text{ ft}^2 \text{ acre}^{-1}$) is basal area at λ_{\max} . These parameter settings define a surface that ascends broadly with stand age to a maximum at age 120 years and that more rapidly ascends with basal area (from either direction) to a maximum at $13.77 \text{ m}^2 \text{ ha}^{-1}$.

We projected population growth, decline, and movement under the model of source-sink population regulation (Pulliam, 1988). We considered compartments where $\lambda_{ts} \geq 1.0$ as population source compartments and those where $\lambda_{ts} < 1.0$ as population sinks. Source habitats may accommodate individuals to a ceiling of $n_{\max} = 200$ birds; beyond this ceiling, surplus individuals become candidates for export to other compartments eligible to receive them. All neighboring (any adjacent compartment sharing a common boundary or corner) sink compartments may receive unlimited numbers of birds, and neighboring source compartments may receive birds up to the limit of n_{\max} . In no instance does a sink compartment export birds to other compartments. Surplus animals to be moved are allocated to recipient compartments in inverse proportion to inter-compartment distances, measured center-to-center. Compartments sharing a common boundary have an inter-compartment distance of 1.0, and those sharing a common corner have a distance of $\sqrt{2}$. Furthermore, we assumed that the area surrounding this forest landscape comprised only sink habitat, arranged in similarly-configured compartments.

2.3. Decision model

For an initial set of habitat and population states $\{a_{0s}, h_{0s}, n_{0s}\}$, the manager seeks the schedule of decisions $d_{ts} \in \{\text{no action, thin, clearcut}\}$, $t \in \{0, 5, 10, \dots, 45\}$ that maximizes $\sum_s n_{50,s}$, the total population size at the end of a 50-year (ten-decision period) planning horizon. The decision policy must satisfy the constraint that ≥ 4 clearcut actions are taken during decision periods 0–3 (years 0–15) and during periods 4–7 (years 20–35). One extreme approach to meeting this constraint is to clearcut all four stands in one decision period; another is to repeatedly clearcut one compartment over four decision periods. A constraint of this sort might be enforced to provide for minimum levels of certain secondary forest objectives including wood volume extraction, openings for other wildlife species, or recreation opportunity.

Following a set of decisions at time t , the new forest habitat state is immediately established. From this state, both the forest stands and the bird populations are grown annually until the next decision period 5 years hence.

3. Optimization heuristic

We used a genetic algorithm (GA) (Goldberg, 1989; Davis, 1991) to find the schedule of decisions that maximize the bird population objective. Alternative approaches that rely on some degree of decision-state enumeration, such as dynamic programming, would have to consider 10^{13} or more decision-state combinations per decision stage, even with population states coarsely discretized. We

believed that a heuristic procedure, such as the GA, provided a reasonable search alternative given our desire to derive temporally and spatially-explicit decision schedules. Though decision schedules provided by the GA are not guaranteed to be optimal, we nevertheless found their use reasonable for our purposes of making relative comparisons in objective function values between competing decision schedules.

We used the simple GA described by Goldberg (1989) with some fitness-scaling modifications suggested by Davis (1991) (Fig. 1). Chromosomes in our GA population comprised 40 genes corresponding to ten decision periods \times four compartments. Alleles for each gene were the set $\{0, 1, 2\}$, representing the no action, thin, and clearcut decisions respectively. Thus, one chromosome with a specific configuration of alleles represented a candidate decision schedule for the planning horizon.

Because some decisions were nullified depending on system state, two chromosomes with different sets of allele values could encode the same decision schedule. For example, $a_{ts} = 20$ years represents a stand younger than the minimum allowable thinning age. Therefore, an allele value of 1 for the gene corresponding to this decision period and compartment is acted upon as if its allele value was 0. We allowed these schedules to freely enter and reproduce in the population. The only drawback of doing so was a loss of efficiency caused by the GA having to search unnecessarily over certain regions of the decision space. However, we felt that this loss was more than compensated by the computational savings in not having to verify each schedule.

Each chromosome encoded a candidate schedule that was either feasible (clearcut constraint met) or infeasible. Each schedule provided a non-negative real value for the population objective function. For feasible schedules, we added 1000 to the population result. For infeasible schedules, we assigned the value 0.0001. By doing so, we segregated the list of objective function values arising from a population of chromosomes into two distinct ranges: ≥ 1000 and 0.0001. We applied a linear ranking to the set of feasible values (Davis, 1991, pp. 31–34), with 1.0 assigned to the smallest value ≥ 1000 , 100.0 assigned to the largest value, and proportionate rankings assigned to intermediate values. To construct the mating pool, we drew chromosomes from the population with selection probability in proportion to transformed objective function values. In this way, individuals encoding infeasible solutions were highly unlikely to contribute to the next generation. Individuals representing feasible solutions were likely to reproduce, but the use of linear ranking helped to ensure that the GA population did not too quickly favor a super individual and converge on a local optimum.

We conducted all GA runs with a population size of 200 chromosomes reproducing over 200 generations. We used elite chromosome selection to ensure that the top-ranked individual in one generation was transferred intact to the next generation; the other 199 chromosomes were constructed from the previous generation through random crossover and mutation (Davis, 1991, p. 34). We used two-site crossover with crossover probability of 0.8. We chose genes for mutation with probability 0.15, then we chose an allele with equal probability from the allele set.

Therefore, the effective gene mutation probability (overall probability of allele change) was 0.10. We chose these values subjectively, but using guidance from Schaffer et al. (1989) and from our own experience with related problems.

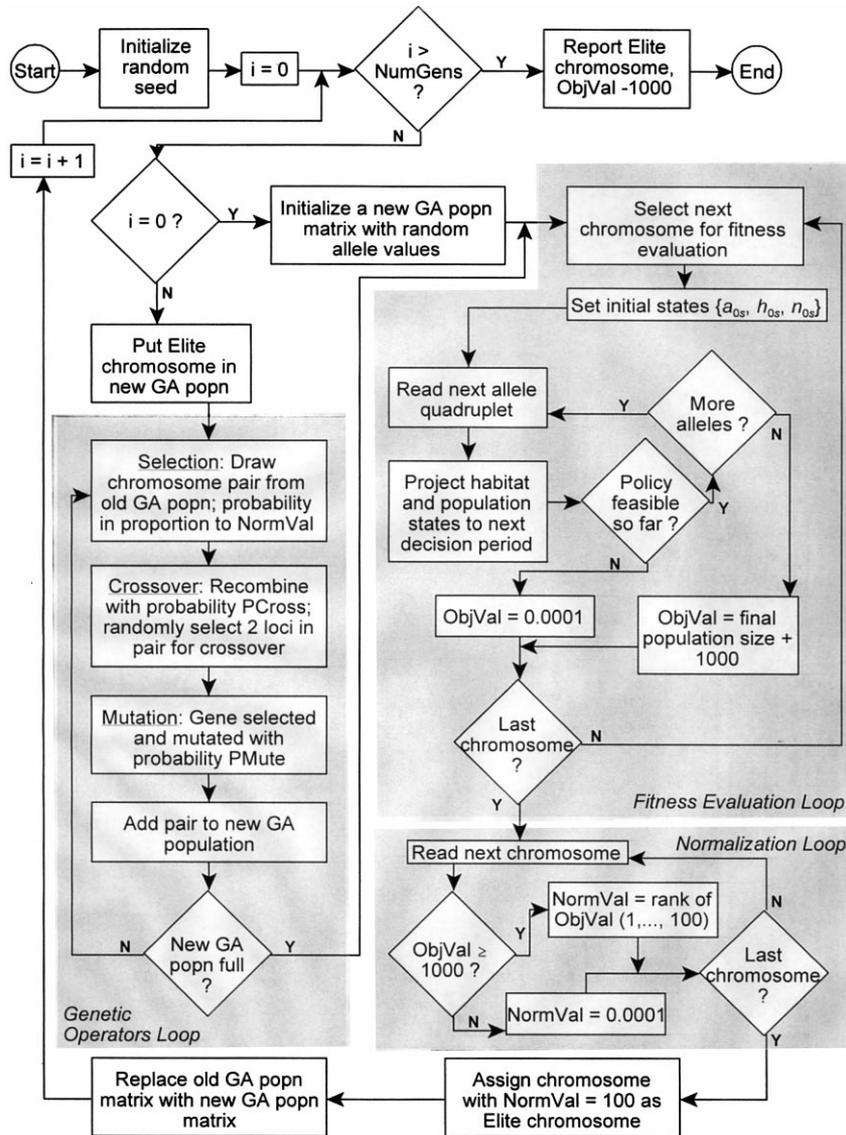


Fig. 1. Flow diagram of genetic algorithm used for finding optimal decision schedules for forest bird population objective. Parameter values were set to each of the following: NumGens (number of generations) = 200, GA population size = 200, PCross (probability of crossover) = 0.80, and PMute (effective probability of gene mutation) = 0.10.

Table 1

Initial age of stand (a , years), initial age of stand at last thinning (h , years), and estimated optimal population sizes at end of planning horizon (median over 50 samples of initial population state) for eight configurations of initial forest state^a

Configuration	Compartment								Abundance
	1		2		3		4		
	a	h	a	h	a	h	a	h	
1	100	95	100	95	100	95	100	95	2.2
2	100	30	100	30	100	30	100	30	2.2
3	60	55	60	55	60	55	60	55	617.5
4	20	30	20	30	20	30	20	30	27.0
5	100	95	70	65	10	30	40	30	483.6
6	100	30	70	65	10	30	40	30	480.3
7	100	95	100	30	10	30	10	30	17.7
8	100	95	10	30	100	30	10	30	20.0

^a Stands never thinned are denoted by $h = 30$.

4. Experimental design

Our aim was to evaluate performance of decision schedules corresponding to presumed initial population distributions on a sample of true population distributions. We randomly generated a sample of 1000 vectors from a four-component simplex volume (all vector components sum to 1.0), then we picked the 50 most dispersed within the volume (Steuer, 1986, pp. 311–332). We used these vectors to allocate an initial total population size of 200 birds into the four forest compartments, permitting population states to be represented as fractions of organisms.

We selected eight initial configurations of forest habitat (Table 1). Configurations 1–4 are four states of a homogeneous forest: mature/thinned, mature/unthinned, mid-rotation/thinned, and young/unthinned, respectively. The other four configurations represent heterogeneous forests. In configurations 5 and 6, stand ages vary in a gradient across the forest. In configurations 7 and 8, stand ages vary dichotomously between young and mature states. Similar stands are contiguous in configuration 7 and are isolated in configuration 8.

We conducted 20 replications of our GA for each of the $50 \times 8 = 400$ initial habitat and population states. For each set of replications, we saved the decision schedule for the replication that provided the greatest value of the objective function. We refer to this set of 400 decision schedules as the ‘known-state’ policy or KSP set, and an individual policy from that set is called a KSP. The corresponding objective value for a KSP is called the KSPV.

We also carried out the same optimization exercise for a much smaller set of initial population states. We considered five typical distributions of birds that a manager, having knowledge of only the total population size of 200, may presume to occur among the forest compartments: (1) 50 birds in each compartment, (2) all 200 birds in compartment 1, (3) all 200 in compartment 2, (4) all 200 in

compartment 3, and (5) all 200 in compartment 4. We considered the same set of habitat configurations because we believed that a manager would be able to observe the habitat state of the system at each decision period with relative ease. We refer to this set of $5 \times 8 = 40$ decision schedules as the ‘presumed-state’ policy or PSP set.

For each of the 400 true population and habitat configurations, we projected bird populations through time and management interventions under five corresponding PSPs. We thus obtained a set of five population projections under presumed conditions for each combination of true conditions, and each was comparable to a corresponding KSPV. We refer to the population projection of a true state under a PSP schedule as a PSPV.

We calculated several statistics to aid our comparisons. By computing averages or medians for these statistics over the sample of true initial population states, we obtained ‘expected values’ of decision schedule performance when initial population states must be presumed rather than known.

PCTDIFF was the difference between KSPV and PSPV, scaled by KSPV and expressed on a percentage basis. We classified the membership of PCTDIFF in four range classes: $\text{PCTDIFF} \geq 95\%$, $5\% < \text{PCTDIFF} \leq 95\%$, $-5\% < \text{PCTDIFF} \leq 5\%$, and $\text{PCTDIFF} \leq -5\%$. These classes qualitatively describe degree of performance loss under PSPs relative to KSPs as either high, moderate, negligible, or negative (i.e. performance gain), respectively. Note that if the KSP set constitutes a set of bona fide optimal decision schedules, then PCTDIFF should never take on negative values. However, we allow this possibility as our heuristic procedure provides a set of ‘good’ decision schedules that are not guaranteed to be optimal.

For any KSP or PSP schedule, we considered the population at $t = 50$ to be extinct if $n_{50,s} < 2.0$ for all s . The indicator variable EXTINCT takes on the value 1 if the population becomes extinct under a PSP but not under the corresponding KSP, and it takes on the value 0 otherwise. Thus EXTINCT measures the tendency for the PSP to carry the population to extinction conditional on the fact that the KSP does not. Another indicator variable, DISAME, takes on the value 1 when the first-period decision is identical in the PSP and the KSP and 0 otherwise.

5. Results

Decision schedules derived by the GA replications provided objective values that were generally consistent with our intuition of how initial habitat states would influence bird abundance over time (Table 1). Habitat configurations 3, 5, and 6 provided initial excellent bird habitat and provided numerous opportunities over time for management to perpetuate population growth. Configurations 1 and 2, however, provided little opportunity for improvement through management (the entire forest had to be clearcut within 20 years) despite the initial provision of favorable bird habitat. Configurations 4, 7, and 8 provided relatively poor initial habitat but provided plenty of opportunity for habitat improvement through management. With a longer planning horizon, we would expect large population sizes for these latter configurations.

Average values of DISAME ranged from 0.0 to 0.34 (Fig. 2). Least similarity occurred for habitat configurations 3, 5, and 6; policies for these configurations allowed the greatest degree of management latitude over the planning horizon. In such cases, the first-period decision appears to be highly sensitive to how well one's presumption matches the true initial population state. As management options become more limited (e.g. configurations 1 and 2), choice of an initial decision seems to become less dependent on the accuracy of this presumption.

The median of PCTDIFF varied strongly with initial habitat configuration and with initial presumed population state (Fig. 3). For habitat configurations 3, 5, and 6, PSPVs generally were in agreement with KSPVs. We found important exceptions, however, for initial presumed population state 4 in habitat configurations 5 and 6. This configuration of states presumed that all 200 birds initially occurred in a long-term sink habitat, and not surprisingly, PSPVs fell significantly short of KSPVs. In five cases, median PCTDIFF was negative (one value of -0.03 was too small to display in Fig. 3).

In habitat configurations 1 and 2, PSPVs were generally in agreement with KSPVs (Fig. 3), however, both tended to fall near 0 (Table 1). Median PCTDIFF was typically significant for habitat configurations 4, 7, and 8 (Fig. 3). Populations under these initial habitat configurations would certainly grow large given a

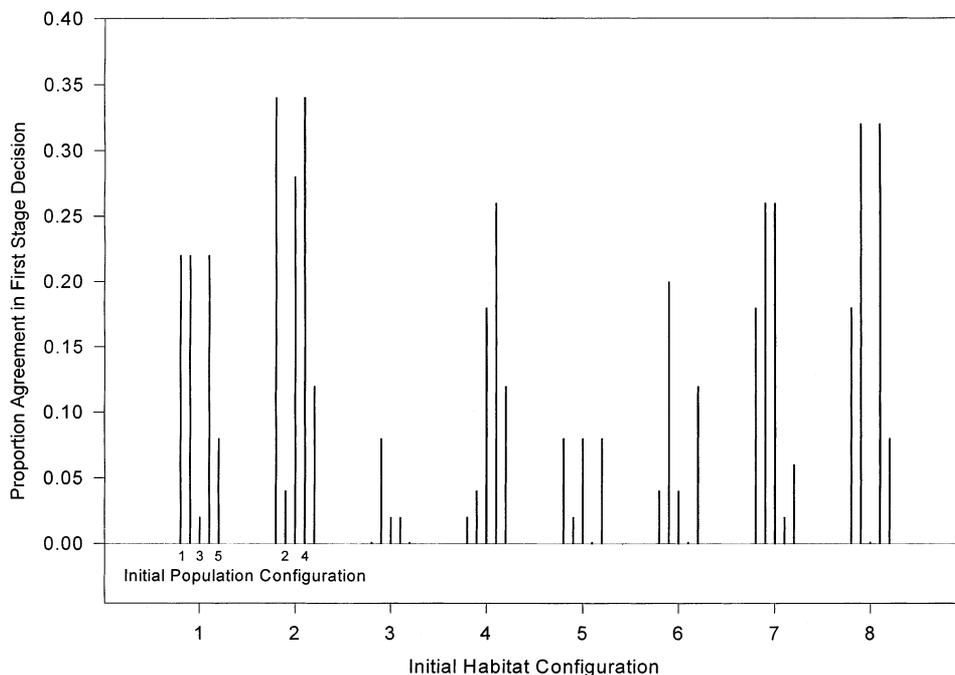


Fig. 2. Proportion of known-state decision schedules ($n = 50$) where first-period decision agreed with that of the presumed-state schedule. Values are displayed for five presumed initial population configurations at each of eight initial habitat configurations.

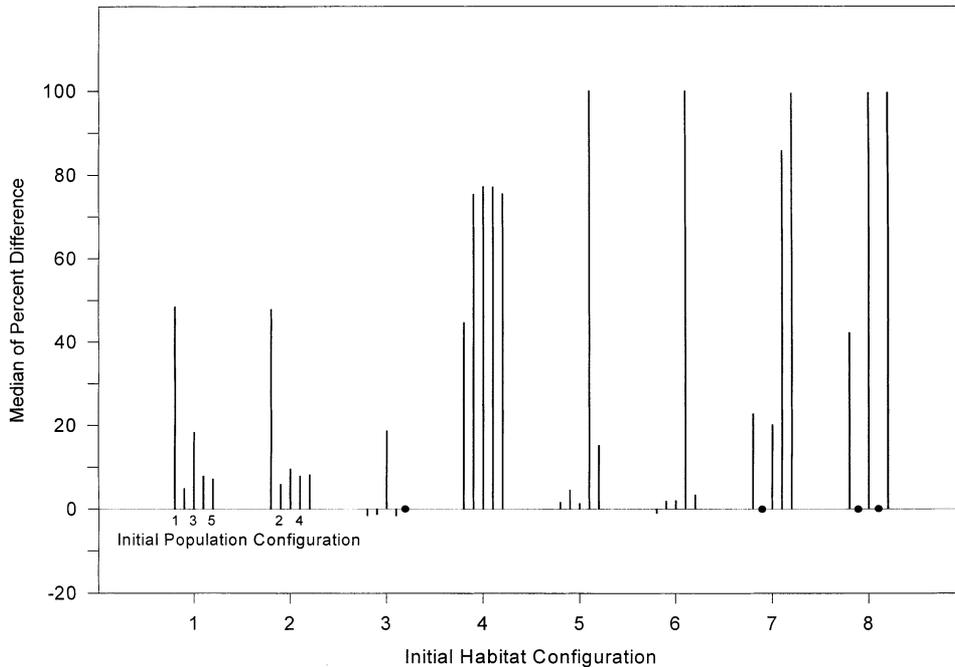


Fig. 3. Median difference between objective value from known-state decision schedules ($n = 50$) and population projection from presumed-state decision schedule, where difference is expressed as percentage of objective value. Values are displayed for five presumed initial population configurations at each of eight initial habitat configurations. Round symbols denote values < 0.03 in absolute value.

sufficiently long planning horizon: whether these large differences between PSPV and KSPV would persist is unknown.

The distribution of PCTDIFF in range classes provides a better illustration of how often PSPs were highly, moderately, or negligibly inferior to KSPs (Fig. 4). In fact, the 'optimal' KSP frequently performed poorly relative to the PSP, particularly for initial habitat configuration 3, although the typical difference in performance was not great (Fig. 3). We found that performance was most often worse for initial habitat configurations 4, 7, and 8, though exceptions did occur for particular combinations of habitat and population configurations.

Average values of EXTINCT followed patterns similar to those described above (Fig. 5). For habitat configurations 3, 5, and 6, extinctions did not appear to occur more frequently except in two cases. As noted in the summary of PCTDIFF, presumed population configuration 4 placed all birds in the long-term sink habitats of configurations 5 and 6. Therefore, these particular PSPs often led to extinction whereas the KSPs did not.

More frequent extinctions were also noted for habitat configurations 4, 7, and 8. As in the previous case, extinctions greatly increase under policies that presume birds are initially concentrated in very poor habitat (see population configurations 4 and 5 for habitat configuration 7 and population configurations 3 and 5 for

habitat configuration 8; Fig. 5). We also note that extinctions were more frequent for all PSPs in habitat configurations 1 and 2, but these results may have little meaning as many KSPs also carried the population to extinction.

6. Discussion and recommendations

We have shown that in many instances under a spatially-explicit model of population dynamics, decision schedules based on presumed initial distributions of populations perform quite poorly relative to schedules based on known distributions. In other instances, the presumed initial population distribution mattered little. Despite their basis on a very simple model, our analyses suggest that common presumptions about animal distribution may indicate decision policies that are suboptimal, or even irreversibly harmful, with respect to the managed resource. Decision policies for spatially-explicit systems should therefore address uncertainty regarding the distribution of individuals. One approach is adaptive resource management (Walters, 1986).

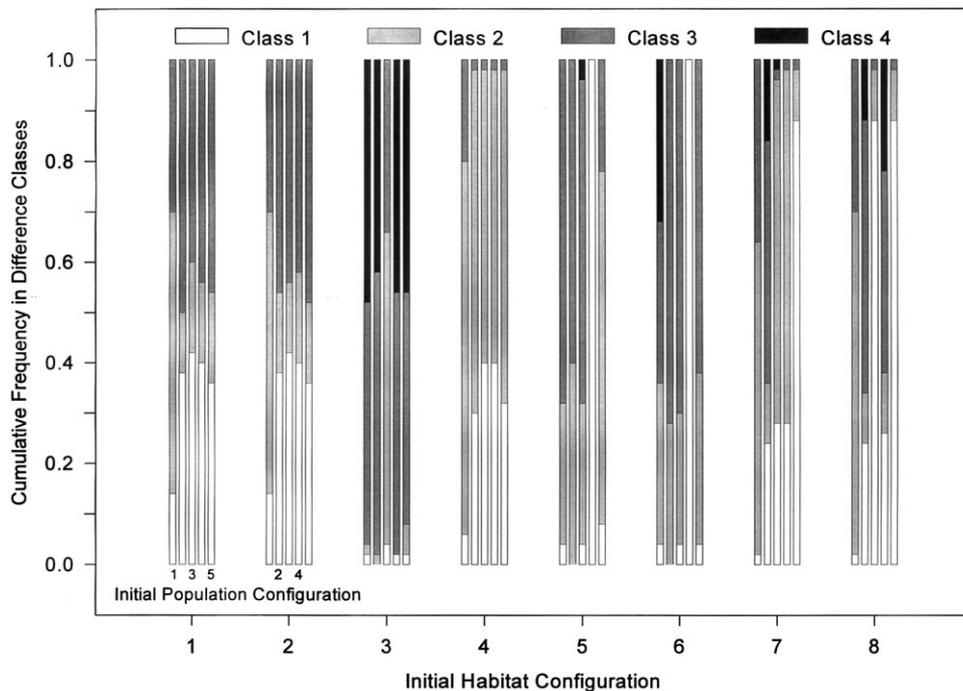


Fig. 4. Distribution of difference between objective value from known-state decision schedules ($n = 50$) and population projection from presumed-state decision schedule, where difference is expressed as percentage of objective value. Differences are assigned to range classes [95%, 100%], [5%, 95%), [-5%, 5%), and $(-\infty, -5\%)$ numbered 1–4, respectively. Values are displayed for five presumed initial population configurations at each of eight initial habitat configurations.

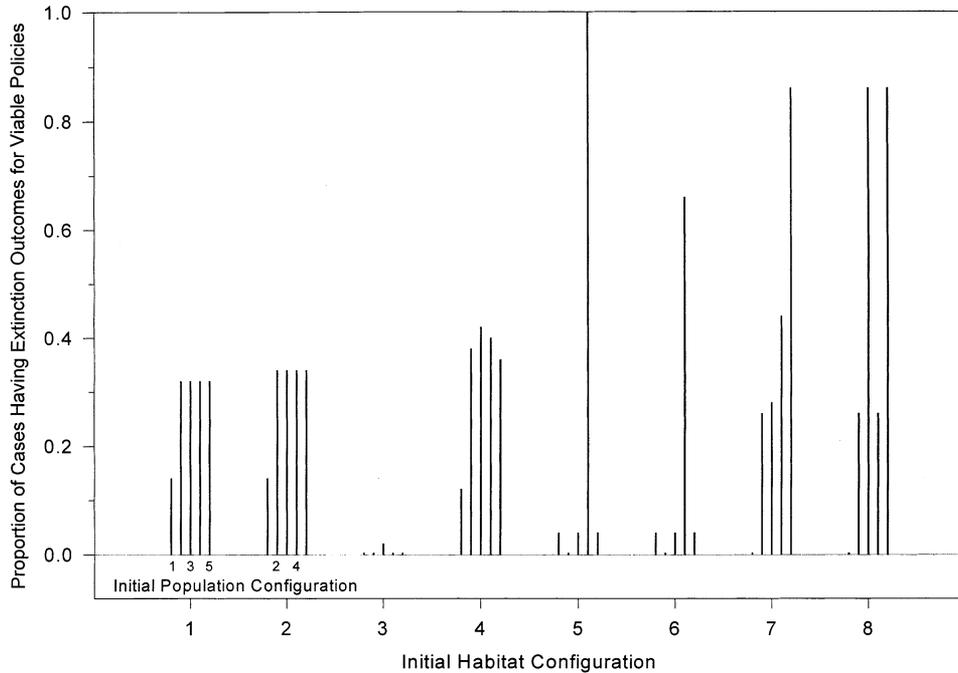


Fig. 5. Proportion of known initial population states ($n = 50$) that are carried to extinction under a presumed-state policy but not under the corresponding known-state policy. Values are displayed for five presumed initial population configurations at each of eight initial habitat configurations.

The average of the PCTDIFF variable is related to the expected value of perfect information (EVPI) (Lindley, 1985) for the unknown bird system state. The EVPI is a measure of what one should be willing to pay, in units of the utility, to obtain perfect information in making a decision (Lindley, 1985). EVPI compares the expected utility under perfect information (EUPI) against the maximum expected utility for any decision in the absence of information. In contrast, our analyses compared the EUPI against expected utility for that optimal decision corresponding to a single presumed system state. In a sense, our analysis measured the value of information against specific 'worst-case' scenarios, any of which might be reasonable presumptions for a habitat manager. Our statistic may most closely approximate EVPI in case 1 of the presumed initial state, in which birds were distributed equally among compartments.

We suspect that observed outcomes were highly dependent on the size of our forest system, on the dynamics and initial size of our bird population, and on length of planning horizon. Had we considered a much longer planning horizon, these differences may have been of smaller magnitude as we would have expected the system to become less dependent on initial state over time. On the other hand, had we considered smaller initial population sizes, more forest compartments, or stochastic population dynamics, we might have made system performance even more sensitive to initial conditions.

Even for such small and simplistic modeled systems as ours, it is apparent that obtaining temporally and spatially-explicit decision policies for resource managers will require alternatives to the classical mathematical programming techniques (linear programming, dynamic programming, integer programming, etc.). We see that heuristic procedures including GA, Tabu search (Glover, 1986), neural networks (Narendra, 1996), and simulated annealing (Kirkpatrick et al., 1983), especially hybridized in some manner with the classical techniques, hold great promise for optimal management of wildlife resources.

Our analyses suggested that a considerable number of decision policies were not optimal. Without comparison to results from an enumerative optimization procedure, we do not know the extent or degree to which this occurred. However, for a related optimal control problem of another dynamic system (harvest of waterfowl populations), comparisons between our GA solutions and those provided by dynamic programming were not substantially different (C.T. Moore, unpublished data). Furthermore, we are confident that several options we did not pursue could have improved the performance of our GA. For example, a ‘GA-within-GA’ approach could have been used to adapt the GA parameters over the course of the GA run.

Management under a coarse resolution of system scale takes on many forms, and we addressed only one specific instance. Another common form is the identification of population states through species presence/absence data. Here, one obtains spatially-explicit information on the state of the population, but the dichotomous nature of the data provides only the absolute minimum in information content, namely, whether a species is truly present or possibly absent in the area. A Bayesian analysis might be used to identify a spatially-explicit probability distribution of animal abundance conditional on a particular configuration of presence/absence indicators among forest compartments. By sampling population values from this distribution and deriving a known-state policy for each sample, one could begin to search for policies that return the greatest expected value of the population objective. Because many sampling programs that provide an estimate of aggregated population size also provide species indicator data at some finer resolution of scale, managers might couple both pieces of information to build a probability distribution on abundance.

Another form of coarse-scale management that deserves study is the inability to identify the spatially-explicit model operative in the population. A reasonable null model is that of no spatially-explicit processes whatsoever, but other (incorrect) spatially-explicit models could be proposed as well. An adaptive approach (Walters, 1986; Williams, 1997) might serve this management situation well, where measurements of the system state over time are matched to each model’s predictions. These comparisons update relative confidence in each of the proposed models, and in turn, the revised model credibility weights affect future decision policies. However, to distinguish among spatially-explicit models through adaptive management will generally require knowledge of animal distributions.

A final typical form of management under coarse system resolution is that of carrying out management actions at scales that are large relative to the scale of

population dynamics of the target organism. Here, a degree of control over management is lost in that actions are difficult to fine-tune to the precise level desired for optimal population response. For each management decision and system state, one must specify a probability distribution for the system's response to management. The expected objective value of the optimal policy under this scenario would likely be smaller than that under a scenario in which the spatial scale of actions more closely agreed with that of population dynamics.

Acknowledgements

We thank an anonymous reviewer and R.E. Plant for comments which improved our manuscript. Work of CTM was supported by the U.S. Environmental Protection Agency, STAR Graduate Fellowship Program, and by the U.S. Geological Survey, Biological Resources Division (Work Order # 34, Coop. Agreement # 14-16-0009-1551). The Georgia Cooperative Fish and Wildlife Research Unit is jointly sponsored by the Biological Resources Division, the Georgia Department of Natural Resources, the University of Georgia, and the Wildlife Management Institute.

References

- Bettinger, P., Sessions, J., Boston, K., 1997. Using Tabu search to schedule timber harvests subject to spatial wildlife goals for big game. *Ecol. Model.* 94, 111–123.
- Conroy, M.J., Noon, B.R., 1996. Mapping of species richness for conservation of biological diversity: conceptual and methodological issues. *Ecol. Appl.* 6, 763–773.
- Davis, L. (Ed.), 1991. *Handbook of Genetic Algorithms*. Van Nostrand Reinhold, New York.
- Dunning, J.B. Jr, Watts, B.D., 1990. Regional differences in habitat occupancy by Bachman's Sparrow. *Auk* 107, 463–472.
- Glover, F., 1986. Future paths for integer programming and links to artificial intelligence. *Comput. Oper. Res.* 13, 533–549.
- Goldberg, D.E., 1989. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, Reading, MA.
- Hof, J., Raphael, M.G., 1997. Optimization of habitat placement: a case study of the Northern Spotted Owl in the Olympic Peninsula. *Ecol. Appl.* 7, 1160–1169.
- Hughell, D.A., 1997. *Simulated adaptive management for timber and wildlife under uncertainty*. Ph.D. Thesis, North Carolina State University, Raleigh, USA.
- Kirkpatrick, S., Gelatt, C.D., Vecchi, M.P., 1983. Optimization by simulated annealing. *Science* 220, 671–680.
- Ligon, J.D., Stacey, P.B., Conner, R.N., Bock, C.E., Adkisson, C.S., 1986. Report of the American Ornithologists' Union Committee for the Conservation of the Red-cockaded Woodpecker. *Auk* 103, 848–855.
- Lindley, D.V., 1985. *Making Decisions*, second ed. Wiley, London.
- Narendra, K.S., 1996. Neural networks for control: theory and practice. *Proc. IEEE* 84, 1385–1406.
- Pulliam, H.R., 1988. Sources, sinks, and population regulation. *Am. Nat.* 132, 652–661.
- Pulliam, H.R., Dunning, J.B. Jr, Liu, J., 1992. Population dynamics in complex landscapes: a case study. *Ecol. Appl.* 2, 165–177.

- Schaffer, J.D., Caruana, R.A., Eshelman, L.J., Das, R., 1989. A study of control parameters affecting online performance of genetic algorithms for function optimization. In: Schaffer, J.D. (Ed.), *Proceedings of the Third International Conference on Genetic Algorithms*. Morgan Kaufmann, San Mateo, CA.
- Steuer, R.E., 1986. *Multiple Criteria Optimization*. Wiley, New York.
- Walters, C.J., 1986. *Adaptive Management of Renewable Resources*. MacMillan, New York.
- Walters, J.R., 1991. Application of ecological principles to the management of endangered species: the case of the red-cockaded woodpecker. *Ann. Rev. Ecol. Syst.* 22, 505–523.
- Williams, B.K., 1989. Review of dynamic optimization methods in renewable natural resource management. *Nat. Resour. Model.* 3, 137–216.
- Williams, B.K., 1997. Approaches to the management of waterfowl under uncertainty. *Wildl. Soc. Bull.* 25, 714–720.