

Adaptive Management of Migratory Birds Under Sea Level Rise

Samuel Nicol

CSIRO Ecosystem Sciences
sam.nicol@csiro.au

Takuya Iwamura

Stanford University
takuya@stanford.edu

Olivier Buffet

INRIA/
Université de Lorraine
olivier.buffet@loria.fr

Iadine Chadès

CSIRO Ecosystem Sciences
iadine.chades@csiro.au

Abstract

The best practice method for managing ecological systems under uncertainty is adaptive management (AM), an iterative process of reducing uncertainty while simultaneously optimizing a management objective. Existing solution methods used for AM problems assume that the system dynamics are stationary, i.e., described by one of a set of pre-defined models. In reality ecological systems are rarely stationary and evolve over time. Importantly, the effects of climate change on populations are unlikely to be captured by stationary models. Practitioners need efficient algorithms to implement AM on real-world problems. AM can be formulated as a hidden model Markov Decision Process (hmMDP), which allows the state space to be factored and shows promise for the rapid resolution of large problems. We provide an ecological dataset and performance metrics for the AM of a network of shorebird species utilizing the East Asian-Australasian flyway given uncertainty about the rate of sea level rise. The non-stationary system is modelled as a stationary POMDP containing hidden alternative models with known probabilities of transition between them. We challenge the POMDP community to exploit the simplifications allowed by structuring the AM problem as an hmMDP and improve our benchmark solutions.

1 Introduction

The best practice method for managing ecological systems under uncertainty is adaptive management (AM), an iterative process of reducing uncertainty by learning the true system model from management outcomes [Walters, 1986]. The true model is assumed to be one of a suite of alternative models. The AM problem of finding the best management action given model uncertainty can be formulated as a hidden model Markov decision process (hmMDP) [Chades *et al.*, 2012], that is typically solved as a POMDP.

Solution methods currently used for AM suffer from two key issues. Firstly, previous implementations of AM have as-

sumed that the true model is stationary [Martin *et al.*, 2011]. In reality, ecological systems are rarely stationary and evolve over time. For example, many populations will be affected by climate change, which induces non-stationarity. Incorrectly assuming that the environment is stationary means that AM may learn the wrong model, leading to poor management performance. AM problems can be modelled with either online reinforcement learning or POMDP planning, however both methods require some form of workaround to adapt to non-stationary environments. The second issue affecting implementation of AM is that exact solution methods used for AM problems are slow. It remains a challenge to find ever more efficient and user-friendly solution techniques.

In this paper we provide an ecological dataset, data files formatted for the point-based POMDP solver symbolic Perseus [Poupart, 2005] and performance metrics for the solution of an AM problem. The dataset represents networks corresponding to migratory routes used by 10 shorebird species utilizing the East Asian-Australasian (EAA) flyway given uncertainty about the rate of sea level rise and its effect on shorebird populations. The problem is interesting because sea level rise is non-stationary. We model the non-stationary system as a stationary POMDP containing alternative models with known probabilities of transition between them. We provide benchmark solution times for the problem in the hope that more efficient algorithms and implementations will improve performance.

1.1 Problem Description

The EAA flyway is one of the world's major migratory shorebird routes. After breeding in eastern Siberia, shorebirds undertake a return annual migration through Asia to Australasia. Along the way, shorebirds depend on intertidal mudflats [Myers *et al.*, 1987]. At these staging sites shorebirds feed on invertebrates that provide them with energy to complete the migration.

Globally, sea levels are predicted to rise between 0.5–1.4m by 2100 [Rahmstorf, 2007]. As sea levels rise, low-lying intertidal areas that are critical habitat for migrating shorebirds will be inundated and lost. Development of urban coastal areas is rapidly removing the flexibility for intertidal mudflats to migrate inland as the sea level rises. Development plans

need to ensure that staging sites are protected against an uncertain future sea level rise. We use AM to predict when and where actions will be required to protect shorebird populations against the predicted effects of sea level rise.

2 Methods

2.1 Sea Level Rise Scenarios

We allow 3 alternative models for the extent of future sea level rise (0m, 1m, 2m) covering the predicted range of sea level rise for the next century. The extent of staging sites inundated by each model was predicted using digital elevation models, bathymetry maps and tidal information. Full details are contained in [Iwamura, 2011].

Intertidal areas may advance inland if development does not prevent advancement. In our model, if management was implemented then intertidal areas were allowed to advance inland where existing impermeable surfaces did not inhibit movement [Iwamura, 2011]. No inland encroachment was possible if management was not implemented.

2.2 Network Flow Model

We represent the flyway of each species as a weighted directed graph, which is used to predict the population at the breeding site from one year to the next. Each node represents the number of birds using a regional group of shorebird sites given a sea level rise scenario. Edges represent the maximum flow of birds that can migrate between nodes. The structure of the network was obtained using empirical data and expert knowledge [Iwamura, 2011]. Birds migrate deterministically through the graph from the breeding node. After returning to the breeding node the surviving population undergoes stochastic breeding, which re-sets the population for the next annual migration.

Sea level rise will reduce the habitat at each node. We assume that for a fixed sea level rise, the maximum population that can be supported by a node declines deterministically in proportion to the habitat inundated at the node. Declines in maximum population at a node may have a flow-on effect to other nodes as fewer birds can pass through the node. Given a population at the breeding node and a sea level rise model, we compute the number of birds returning to the breeding node after migration using a maximum flow algorithm.

Shorebirds breed between annual migrations. We estimate the annual population change after breeding for our 10 taxa using a stochastic Gompertz model. Population is assumed to be a function of growth rate, a density dependence term and a normally distributed process error term. We have no time series data for the breeding node of each population, so we use population data from Moreton Bay, Australia [Wilson *et al.*, 2010], assuming that fluctuations in the breeding population were proportional to fluctuations elsewhere in the flyway.

2.3 POMDP Formulation

Managers need to select nodes to protect given stochastic population fluctuations and uncertainty about the effect of sea level rise on bird populations. We allow managers to perform an action at one node during each time step. By taking an action, a node is protected against a fixed amount of sea level

rise described by one of the three models (0m, 1m or 2m rise). If a node is protected against the realized sea level rise then mortality is computed using only the proportion of the potential expansion area covered with impermeable surfaces. If sea level rise is greater than the protection level then mortality is computed assuming that no intertidal area expansion is possible. Actions are not taken at the breeding node, so that for a network with N nodes, actions can be applied to $N - 1$ nodes. We formulate the hmMDP as a POMDP tuple $(S, A, O, T, Z, r, H, b_0, \gamma)$ [Chades *et al.*, 2012] where:

- $S = X \times V_1 \times \dots \times V_{N-1} \times M$ is the factored state space. X is the completely observable population state at the breeding node, discretized into 4 population sizes so that $|X| = 4$. V_n represents the completely observable protection states of node n ($n \in \{1, \dots, N - 1\}$). M is the set of models representing the flyway population dynamics under different extents of sea level rise. We include a set $M = \{m_1, m_2, m_3\}$ of models representing 0m, 1m and 2m sea level rise. The size of V_n depends on the number of models (because each protection state protects against the sea level rise represented by one of the models), so that $|V_n| = |M| = 3$. The total state space size is $|X||M|^N$.
- A is the finite set of actions. Protection level can be increased at one node per time step ($(N - 1)$ actions).
- $O = X \times V_1 \times \dots \times V_{N-1}$ is the set of observations. The size of the observation space is $|X||M|^{N-1}$.
- T is the transition matrix $P(x', \vec{v}', m' | x, \vec{v}, m, a) = P(x' | x, \vec{v}, m, a)P(\vec{v}' | \vec{v}, a)P(m' | m)$ which gives the probability to end up in state (x', \vec{v}', m') given that action a is performed in state (x, \vec{v}, m) . $P(x' | x, \vec{v}, m, a)$ uses the Gompertz model fit to determine the population leaving the breeding site. If the population leaving the breeding node x' is less than the maximum flow through the network given the protection states of the nodes and the model then all birds complete the migration, otherwise the number of returning migrants is the maximum flow. The population is rounded up to the nearest discrete state to determine x' . The protection levels of each node are independent, i.e. $P(\vec{v}' | \vec{v}, a) = P(v'_1 | v_1, a) \dots P(v'_{N-1} | v_{N-1}, a)$. For a node $n \in \{1, \dots, N - 1\}$ going from protection level i to protection level j (where $i, j \in \{1, \dots, |M|\}$), $P(v'_n | v_n, a)$ is determined by the probability matrices of decline $P_{decline}(n)$ and action $P_{action}(n)$. The elements of $P_{decline}(n)$ are 1, if $i = j = 1$; 0.075, if $j = i - 1$; 0.925, if $j = i \neq 1$; and 0 otherwise. The elements of $P_{action}(n)$ are all 1 if no action is performed at n . If action a is performed at n , the elements of $P_{action}(n)$ are: 1, if $i = j = |M|$; 0.01, if $j = i \neq |M|$; 0.99, if $j = i + 1$; and 0 otherwise. The transition matrix for node n is $P(v'_n | v_n, a) = P_{decline}(n)P_{action}(n)$. $P(m' | m)$ is the probability of transition between models. Sea level rise is expected to increase over time. Consequently for $i < M$ we set $P(m_j | m_i)$ to: 0.05 if $j = i + 1$; 0.95 if $i = j \neq |M|$; 1 if $i = j = |M|$ (max-

imum sea level rise achieved); and 0 otherwise. These probabilities correspond to a sea level rise of one model every 20 years on average. Our model differs from the previous hmMDP formulation of AM because it allows the true model to change over time (i.e. we relaxed assumption 3 in [Chades *et al.*, 2012]). We do this to mimic the behaviour of a non-stationary system.

- Z is the observation matrix. As all variables are observable except the hidden model, Z is the identity matrix.
- r is a state-dependent reward function. If a proportion of habitat is lost then a cost equal to the number of birds that could have used the habitat is incurred. The reward for a state is the population size at the breeding node minus the cost.
- H is the infinite time horizon for the problem. b_0 is the initial belief in each model— we assumed that all models had the same initial belief. $\gamma = 0.95$ is a discount factor to facilitate convergence of the algorithm.

3 Benchmark Solution Times

We used symbolic Perseus [Poupart, 2005] to solve the AM problem for 10 species. Default parameter values were used for all species (50 iterations, 300 belief points retained, maximum alpha values retained= 200, maxnorm threshold= 0.001; see symbolic Perseus documentation for full description of parameters). Time to compute the approximate POMDP solution and average reward after 500 simulations are reported in Table 1. Solution time increases with the size of the network. Required solution times are viable, however experts need to experiment and compute solutions multiple times.

Table 1: Benchmark times and rewards for 10 flyway species. Results generated on a dual 3.46GHz Intel Xeon X5690 computer with 96GB of memory.

Species	Nodes	States	Solution time (s)	Average Reward
Lesser sand plover	3	108	10	4675
Bar-tailed godwit <i>b.</i>	5	972	48	18217
Terek sandpiper	6	2916	48	7263
Bar-tailed godwit <i>m.</i>	6	2916	58	24583
Grey-tailed tattler	6	2916	378	4520
Red knot <i>pearsonii</i>	8	26244	289	5510
Red knot <i>rogersi</i>	8	26244	3323	6314
Great knot	9	78732	3479	35193
Far eastern curlew	9	78732	29489	5086
Curlew sandpiper	9	78732	4442	22994

The range of network solution times means that our AM problems are suitable for different types of POMDP solvers that exploit the hmMDP problem structure. For the smaller networks, exact solutions may be possible. For larger networks, faster approximation techniques can be implemented.

4 Conclusion/ Data Challenge

Natural resource managers strive to control populations in space and time, however the curse of dimensionality has prevented optimization beyond trivially small examples in sta-

tionary systems. Advances developed in the AI community (i.e. ADDs, POMDP algorithms) show promise for application in ecological networks, but solution times remain high and uptake by ecologists is low. The hmMDP formulation allows AM problems to be factored, i.e., the corresponding POMDP can be solved more efficiently because several state variables are fully observable. The input files we provide are formatted for POMDP, but existing POMDP solvers could be adapted to exploit the mixed observability. We provide a real ecological dataset and instructions to formulate it as a hmMDP to challenge POMDP researchers keen to develop new methods for ecologists. The networks we present are small but planning algorithms should be scaleable to larger networks. Low runtime, high reward algorithms that can accommodate non-stationary models are needed to find good management solutions in ecological network problems.

5 Appendix

Data sets and symbolic Perseus input files are available at <https://sites.google.com/site/adaptivemanagementeaa>. We also provide files in Cassandra’s format [Cassandra, 1998] for networks with less than 8 nodes. S.N. was supported by the National Environmental Research Program and the CSIRO Climate Adaptation Flagship.

References

- [Cassandra, 1998] A. Cassandra. *Exact and Approximate Algorithms for Partially Observable Markov Decision Processes*. PhD thesis, Brown University, 1998.
- [Chades *et al.*, 2012] I. Chades, R. Sabbadin, J. Carwardine, T.G. Martin, S. Nicol, and O. Buffet. MOMDPs: a solution for modelling adaptive management problems. In *AAAI-12*, 2012.
- [Iwamura, 2011] T. Iwamura. *Spatial conservation prioritisation under global threats*. PhD thesis, University of Queensland, 2011.
- [Martin *et al.*, 2011] J. Martin, P.L. Fackler, J.D. Nichols, B.C. Lubow, M.J. Eaton, M.C. Runge, B.M. Stith, and C.A. Langtimm. Structured decision making as a proactive approach to dealing with sea level rise in Florida. *Climatic Change*, 107:185–202, 2011.
- [Myers *et al.*, 1987] J.P. Myers, R.I.G. Morrison, P.Z. Antas, B.A. Harrington, T.E. Lovejoy, M. Salaberry, S.E. Senner, and A. Tarak. Conservation strategy for migratory species. *American Scientist*, 75:19–26, 1987.
- [Poupart, 2005] P. Poupart. *Exploiting Structure to Efficiently Solve Large Scale Partially Observable Markov Decision Processes*. PhD thesis, University of Toronto, 2005.
- [Rahmstorf, 2007] S. Rahmstorf. A semi-empirical approach to projecting future sea-level rise. *Science*, 315(5810):368–370, 2007.
- [Walters, 1986] C.J. Walters. *Adaptive management of renewable resources*. Biological Resource Management in Agriculture. Macmillan, 1986.
- [Wilson *et al.*, 2010] H. Wilson, B.E. Kendall, and H.P. Possingham. Variability in population abundance and the classification of extinction risk. *Conservation Biology*, 25:747–757, 2010.