Computational issues surrounding the dynamic optimisation of management of an ecological food web

William J M Probert [1 2 3] and **Eve McDonald-Madden**[1] and **Nathalie Peyrard**[2] and **Régis Sabbadin**[2]

Abstract. We discuss computational issues surrounding current research that investigates the relevance of graph centrality metrics to the management of ecological food webs. Ecological food webs can be viewed as directed acyclic graphs and we use Markov decision processes to model management. Using dynamic programming we optimally solve the management of an Alaskan food web through time so as to maximise the expected number of species surviving. To generalise our results we investigate policies on generated food webs of varying size. For large food webs the state and action spaces are too large for dynamic programming to be computationally feasible and we use heuristical methods to approximate the optimal policy.

1 INTRODUCTION

Artificial intelligence has been proven to be a useful field for tackling the highly complex and uncertain nature of ecological systems [4, 20]. Management of such ecological systems, with the aim of preserving the world's natural resources in all forms, is the concern of conservation biology.

One area of ecological research in which artificial intelligence can contribute is in the application of graphical models to ecological food webs. In this context, ecological systems are represented mathematically as directed acyclic graphs. Nodes represent 'trophic species', which may be groups of plant or animal species, macroscopic or microscopic, that have a common set of predators and prey [26]. The term 'trophic' relates to the consumption of another species for energy. We will refer to these nodes of 'trophic species' as just 'species'. The basal species in a graph usually represent a ubiquitous food source such as zooplankton or kelp. Edges represent the trophic interactions between species. The terminating node of an edge represents the predator in the interaction (that is, the terminating node eats the source node).

In a reasonably general form, it is possible to frame the problem of managing a food web for conservation purposes as a Markov decision process (MDP). An MDP is an appealing choice because management is then explicitly part of the framework and solution algorithms exist [22]. The MDP framework has been used extensively within ecology and conservation [5, 21]. Solution policies, that prescribe an optimal action given a state and time step for a given reward structure, are specific to the structure of each food web in question and, due to the complexity of food webs, may only be solved using generous assumptions to simplify the system and available actions. Even then, solutions are only possible for moderately sized food webs of less than about 15 species. Such a constraint on finding exact solutions is considerable given that some documented food webs in the ecological literature have over 200 species groups (nodes), such as the Caribbean food web described in Bascompte et al (2005) [2]. It is therefore of interest to investigate heuristical policies that approximate the exact solution in small food webs so that approximate policies may be applied to manage larger food webs of more realistic sizes.

A myriad graph metrics have been applied to study ecological food webs and these provide a suitable starting point for investigating heuristic policies. Previous ecological and conservation research on food webs has mainly focused on measures of how the structure (trophic interactions between species) of a food web persists through species extinctions [15]. This 'stability' has numerous definitions, generally (and traditionally) referring to the stability of the characteristic polynomial from the square, real interaction matrix that defines the species interactions (as defined below) [9, 18].

Typically, the suggested metrics are calculated for each node and, thus, a ranking scheme of nodes is implicit. Research has yet to investigate the use and relevance of said metrics for the purposes of managing ecological systems [6, 1, 15, 13]. Our research compares management of a simple model of an ecological system using the optimal solution with management that is guided by ranking implied by graph metrics. If it is possible to identify metrics that provide management benefits (for a given reward function) that are close to the optimal solution for small scale food webs then such metrics may provide use in guiding management of larger networks.

Initial research has investigated the use of Bayesian networks to model this problem over one time step [19]. Bayesian networks are probability distributions with a structure defined on them [14, 16]. By 'structure', we mean that transition probabilities between different states can be factorised into a product of 'local' transition probabilities defined by the edges between nodes in a graph (and rewards may similarly be decomposed into a summation of local rewards). We extend the previous approach to the multiple time-step case using a finite-horizon Markov decision process (MDP).

An MDP is not just a Bayesian network that is extrapolated to the multiple time-step case and that includes modelling of decisions and rewards; this would usually be called an 'influence diagram' [12]. Figure 1 illustrates the relationship between different optimisation and modelling frameworks and the commonly used names for these techniques. The arrows indicate axes that add additional components to the framework: either time, structure or decisions and rewards. For instance, a Markov chain may be thought of as a one-dimensional random variable that includes a temporal component, similarly if one wishes to add a time component to a Bayesian network, the resultant

---

[1] ARC Centre for Excellence in Environmental Decisions (CEED), The University of Queensland, St Lucia, Queensland, 4072 Australia

[2] INRA-Toulouse, Unité de Biométrie et Intelligence Artificielle, Auzeville, Castanet-Tolosan, FRANCE

[3] email: wprobert AT maths DOT uq DOT edu DOT au; william DOT probert AT toulouse DOT inra DOT fr

model is typically called a dynamic Bayesian network, or if one has a one-dimensional random variable to which they wish to add a framework for decisions and rewards then the resultant framework is generally called stochastic optimisation. This diagram is only concerned with finite time horizons and the structure axis is not strictly one-way, in that, some of those techniques with structure may be modelled using those without formally defined structure (for instance, an influence diagram may be modelled as a Markov decision process).

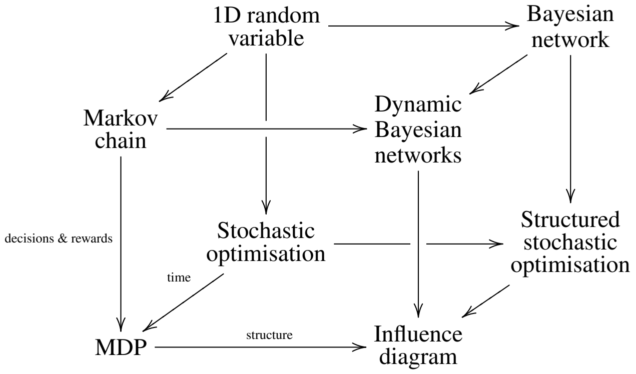

Figure 1. Relationships between modelling techniques with a finite time horizon

By modelling this decision making problem as an MDP, we risk facing the curse of dimensionality when trying to find an optimal solution; the state and action spaces can become combinatorially large as the number of species in the food web increases. Hence, it is necessary to investigate heuristical approaches to approximate exact solutions if results are to be extrapolated to food webs of realistic size. The exact solution and the heuristical solutions are compared via total expected rewards and using statistical analyses to search for patterns in the optimal policy. This is the first research to investigate the use of graph metrics for conservation management purposes over multiple time steps. The extension to multiple time steps also allows comparison of adaptive policies with the myopic policies of the graph heuristics. This manuscript focuses on the computational issues in calculating the exact solution algorithm and therefore some details of the current analysis are omitted.

2 FRAMEWORK & METHODS

We briefly outline the model used and then the computational issues and suggested work-arounds.

We have a directed acyclic graph, $G = \langle V, E\rangle$, representing a food web (figure 2). V respresents a set of species (vertices) and E represents a set of trophic interactions (oriented edges) between those species. There are n species in our food web, $\mid V \mid= n$. Species i is a prey of (is eaten by) species j if there is an edge from i to j, $(i, j) \in E$. We denote the adjacency matrix defining the food web with $\mathbf{G}$, where $\mathbf{G}_{i,j} = 1$ if $(i, j) \in E$, otherwise 0.

Markov decision process. A Markov decision process (MDP) may be used to model the decision making problem of managing an ecological food web. An MDP is a common framework for sequential decision making problems and is composed of a tuple, $\langle \mathcal{X}, \mathcal{A}, P, R\rangle$, consisting of the state space, action space, state transition probabilities and rewards, respectively [22]. We model this system for a set of time steps, $t \in \{1, \ldots, T\}$, where we have a finite time horizon, $T < \infty$. A superscript of t represents the state or action at time step t. The state of each species can take one of two values, extant (1) or extinct (0), $x_i^t \in \{0, 1\},\ \forall i \in V$. The set x^t then represents the state of all species in the food web at time t, $x^t \in \{0, 1\}^n$. Management actions that represent a species' protection status are available for each species. Simply, a species may or may not be protected, that is, $a_i^t \in \{0, 1\},\ \forall i \in V$, where protection takes the value 1. The term a^t then represents the set of actions chosen for all species in the food web at time t, $a^t \in \{0, 1\}^n$. To recognise limitations in conservation funding, we introduce a budget constraint on the choice of actions at each time step, $c(a^t) = \sum c_i a_i^t \leq B^t$, for a per-time-step budget B^t, summation is over $i \in V$, and for a cost of protecting species i of c_i ($i \in V$).

To incorporate local interactions and dependencies in the food web we define neighbourhood functions. The assumption of a neighbourhood allows the factorisation of state transition probabilities into the product of species' local state transitions (and the respective decomposition of rewards into a sum of local rewards) [23]. We define a species' neighbourhood to include all prey species and the species itself, $N(i) = \{j \in V \mid (j, i) \in E\} \cup \{i\}$. The set $x_{N(i)}^t$ then denotes the states of all species in the set $N(i)$ at time t.

The transition probability function is defined as follows:

$$P^t(x^{t+1} \mid x^t, a^t) = \prod_{i=1}^{n} P_i^t(x_i^{t+1} \mid x_{N(i)}^t, a_i^t). \tag{1}$$

For simplicity, we assume transition probabilities are the same across time, $P_i^t(\cdot) = P_i(\cdot) \quad \forall t$. We define individual species' transition probabilities from one state to the next, for a given action, as some baseline probability of survival, p_i^0, times the ratio of the number of prey species that are extant ($f^{\star,t}$) to the total number of prey species (f),

$$P_i(x_i^{t+1} = 1 \mid x_i^t = 1, x_{N(i)\setminus i}, a_i^t = 0) = p_i^0 \left(\frac{f^{\star,t}}{f}\right)$$

$$P_i(x_i^{t+1} = 0 \mid x_i^t = 1, x_{N(i)\setminus i}, a_i^t = 0) = 1 - p_i^0 \left(\frac{f^{\star,t}}{f}\right)$$

Such a probability is subject to the following conditions:

- For basal species, the probability of transitioning from extant to extinct is solely equal to some baseline probability of survival, p_i^0.
- For any species, extinction (death) is an absorbing state,

$$P_i(x_i^{t+1} = 0 \mid x_i^t = 0, a_i^t) = 1, \forall a_i^t$$

- To survive, a species must have at least one prey species extant,

$$P_i(x_i^{t+1} = 0 \mid \sum_{k \in N(i)} x_k^t < 2, a_i^t) = 1, \forall a_i^t$$

- A species' is guaranteed to survive if it is protected, $a_i^t = 1$, and the above conditions hold,

$$P_i(x_i^{t+1} = 1 \mid x_i^t = 1, \sum_{k \in N(i)} x_k^t \geq 2, a_i^t = 1) = 1.$$

Transition probabilities are simplistic in this research project to avoid complexity in computation. Species demographics are not directly taken into account but they may be used in calculations to contribute to an overall probability of species extinction. Note that although basal species may be secured indefinitely by protecting them, this does not guarantee that species further up the food chain will not go extinct and thus may not lead to the most desirable outcome for a given choice of reward function.

We denote the array of transition probabilities as $\mathbf{P}$, where $\mathbf{P}_{i,i',a}$ is the probability of transitioning from state i to state i' when action a is taken.

Our objective is to maximise the number of species surviving at the end of the project. Rewards are assumed stationary through time (with the exception of rewards in the final time step). We use a final time-step reward function of the number of extant species,

$$R^T(x^T) = \sum_{i=1}^{n} x_i^T \tag{2}$$

and per-time-step rewards are zero, $R^t(x^t) = 0,\ t < T$. Various different reward functions may be investigated, including rewards that acknowledge the presence of ecologically meaningful structure in the state of the food web, but for conciseness we only mention one basic reward function here.

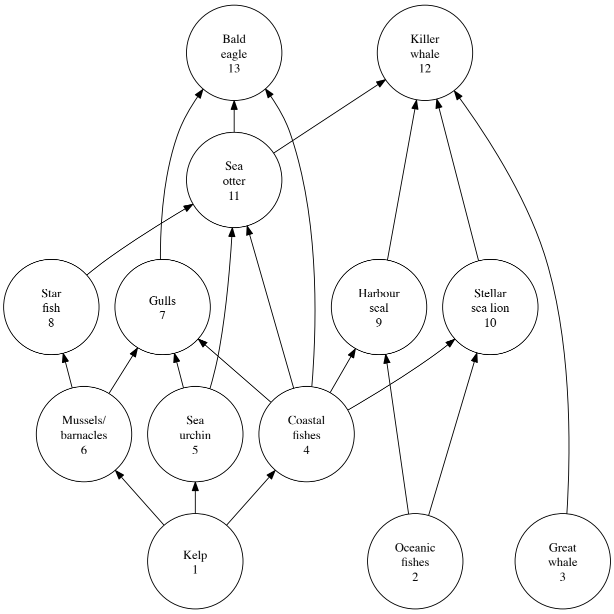

Figure 2. Example food web based on an Alaskan trophic network with 13 species and 21 trophic edges. The numbers underneath the names of the species represent the species index, i, which are referred to in the subsequent tables.

Policies. Let $\delta_r = (d^1, \ldots, d^{T-1})$ represent a policy according to rule r, that designates an action, $a^t = d^t(x^t)$, to take at each time step according to the current state, x^t, and decision rule $d^t(x^t)$. The total expected reward of any policy, δ, is defined as

$$\upsilon_\delta^T(x^1) = \mathbf{E}\left[\sum_{t=1}^{T} R^t(x^t) \,\middle|\, x^1, \delta\right] \tag{3}$$

Optimal policy. Solving the above MDP involves finding the optimal policy, δ_*, that provides the highest total expected reward. This maximum total expected reward is called the 'value'. For food webs with 13 species or less, we may solve the above MDP using the backwards induction algorithm as follows [22]:

1. Set the current time-step to $t = T$ and the value in the final time-step to $\upsilon_*^T(x^T) = R^T(x^T)\ \forall x^T \in \mathcal{X}$
2. Set $t = t - 1$ and calculate $\upsilon_*^t(x^t)$ for each state using

$$\upsilon_*^t(x^t) = \max_{a^t \in \mathcal{A}} Q^t(x^t, a^t)$$
$$a_*^t = \arg\max_{a^t \in \mathcal{A}} Q^t(x^t, a^t)$$

where

$$Q^t(x^t, a^t) = R^t(x^t) + \sum_{x^{t+1}} P(x^{t+1} \mid x^t, a^t)\ \upsilon_*^{t+1}(x^{t+1})$$

3. If $t = 1$ then stop, otherwise return to step 2.

For our problem, this is initialised by setting the optimal value in the final time step equal to the final rewards in equation (3) and setting per-time step rewards to zero.

Graph based policies. We generate policies guided by ranking schemes that are defined by several graph metrics, outlined below. Graph metrics map nodes to integer or real values. Policies that are defined on graph metrics will manage nodes in descending order of the graph metric. Ties in graph metrics are determined by randomisation and isolated species, for instance if the current state causes the food web to become disconnected, have metric values of zero (which means they are managed last).

The total expected rewards of a graph based policy is evaluated using the following finite horizon metric-policy evaluation algorithm with inputs $\delta = (d^1, \cdots, d^{T-1})$ and the interaction matrix, $\mathbf{G}$ [22]:

1. Set the current time-step to $t = T$ and the terminal rewards in the final time-step to $\upsilon_\delta^T(x^T) = R^T(x^T)\ \forall\ x^T \in \mathcal{X}$
2. Set $t = t - 1$ and calculate $\upsilon_\delta^t(x^t)$ for each state using

$$\upsilon_\delta^t(x^t) = Q^t(x^t, d(x^t))$$

where

$$Q^t(x^t, d(x^t)) = R^t(x^t) \quad + \sum_{x^{t+1}} P(x^{t+1} \mid x^t, d^t(x^t))\ \upsilon_\delta^{t+1}(x^{t+1}) \tag{4}$$

3. If $t = 1$ then stop, otherwise return to step 2.

This algorithm is also initialised by setting final time step rewards equal to equation (3) and setting per-time step rewards to zero.

Below we define several graph metrics. From the social or mathematical sciences we define degree centrality, betweenness centrality and closeness centrality. From the ecological literature we define prey degree, predator degree, keystone index, bottom-up prioritisation and trophic level. Metric values for each species in the 13 species Alaskan food web of figure 2 are presented in tables 3 and 4.

Degree centrality

The degree of a species is the number of connections (in any direction) to a particular species [25]. This can be normalised by the size of the food web so that degree centrality can be compared between graphs of different sizes (note that $n - 1$ is the maximum number of connections a species can have in any graph). The degree of species i is the sum of the number of prey and predators that it has, normalised by the size of the food web,

$$D_i = \frac{D_i^{\leftarrow} + D_i^{\rightarrow}}{\mid V \mid -1} \tag{5}$$

where $D_i^{\leftarrow}$ and $D_i^{\rightarrow}$ are respectively the prey degree and predator degree of species i. More specifically, $D_i^{\leftarrow}$ is the size of the set $V_i^{\leftarrow} = \{j \in V : (j,i) \in E\}$ of all prey of species i, and $D_i^{\rightarrow}$ is the size of the set $V_i^{\rightarrow} = \{j \in V : (i,j) \in E\}$, the set of all predators of species i.

Betweenness centrality

The betweenness centrality of a species reflects how central a species is in the transmission of energy that links other species in the food web. It is the proportion of shortest paths between any two species that pass through a particular species, normalised by the size of a network [25]. Betweenness centrality is calculated as

$$BC_i = \frac{\sum\limits_{j<k} \frac{g_{jk}(i)}{g_{jk}}}{(\mid V \mid -1)(\mid V \mid -2)}, \tag{6}$$

where g_{jk} denotes the number of shortest paths (geodesics) between species j and k, and $g_{jk}(i)$ denotes the number of shortest paths between species j and species k which pass through species i.

Closeness centrality

Closeness centrality is a function of the sum of distances between a node and all other nodes in the graph [3, 7, 8]. Specifically it is the ratio of 1) the sum of distances from a particular species to every other species in the network to 2) the minimum possible value that this sum may take (which is $n-1$ for a network of size n).

The closeness centrality for species k is

$$CC_k = \left[\frac{\sum\limits_{i=1}^{n} d(i,k)}{n-1} \right]^{-1} \tag{7}$$

This is a relative measure and can be compared between networks of different sizes. The distance between two species, $d(i,j)$, is the smallest number of trophic connections between species i and j (geodesic). Note that the sum of distances from a species to all other species will grow with the distance between species (hence the inverse is taken) and distances ignore direction of the trophic interaction.

Keystone index

The keystone index is based on the idea of ‘status’ and ‘contrastatus’ of an organisation by Hararay [10, 11, 15]. This measure is only for directed, acyclic graphs. Using the definitions of $V_i^{\rightarrow}$ and $V_i^{\leftarrow}$ from the degree centrality description the keystone index may be composed of two sets of two components:

$$\begin{aligned} K_i &= K_i^{\downarrow} + K_i^{\uparrow} \\ K_i^{\downarrow} &= \sum_{c \in V_i^{\leftarrow}} \frac{1}{D_c^{\leftarrow}} (1 + K_c^{\downarrow}) \\ K_i^{\uparrow} &= \sum_{e \in V_i^{\rightarrow}} \frac{1}{D_e^{\rightarrow}} (1 + K_e^{\uparrow}) \end{aligned} \tag{8}$$

where $K_i^{\downarrow}$ and $K_i^{\uparrow}$ represent the top-down and bottom-up keystone indices of node i respectively. Alternatively, equation (8) can be rearranged to express the keystone index as a sum of direct and indirect effects on node i,

$$\begin{aligned} K_i &= K_i^{\text{dir}} + K_i^{\text{undir}} \\ K_i^{\text{dir}} &= \sum_{c \in V_i^{\leftarrow}} \frac{1}{D_c^{\leftarrow}} + \sum_{e \in V_i^{\rightarrow}} \frac{1}{D_e^{\rightarrow}} \\ K_i^{\text{undir}} &= \sum_{c \in V_i^{\leftarrow}} \frac{K_c^{\uparrow}}{D_c^{\leftarrow}} + \sum_{e \in V_i^{\rightarrow}} \frac{K_e^{\downarrow}}{D_e^{\rightarrow}} \end{aligned}$$

i	D_i	$D_i^{\leftarrow}$	$D_i^{\rightarrow}$	BC_i	CC_i
1	3	0	3	0	0.48
2	2	0	2	0	0.364
3	1	0	1	0	0.353
4	6	1	5	0.035	0.632
5	3	1	2	0.008	0.462
6	3	1	2	0.01	0.429
7	4	3	1	0.011	0.5
8	2	1	1	0.015	0.444
9	3	2	1	0.008	0.5
10	3	2	1	0.008	0.5
11	5	3	2	0.04	0.6
12	4	4	0	0	0.522
13	3	3	0	0	0.522
Mean	3.231	1.615	1.615	0.01	0.485
SD	1.25	1.273	1.273	0.013	0.076
Min	1	0	0	0	0.353
Max	6	4	5	0.04	0.632

Figure 3. Graph metrics for each species in the full Alaskan food web of figure 2. Metrics are respectively (L to R in columns 2 to 6) degree, prey degree, predator degree, betweenness centrality, and closeness centrality. Index numbers, i, for each species are labelled in figure 2.

Bottom-up prioritisation

Bottom-up prioritisation (BUP) ranks species firstly according to trophic level, L_i, which is calculated from the complete food web, and then secondly by the number of extant predators in the current state. Let L_1 define the set of basal species in a food web with a trophic level of 1, that is, all those species which do not have prey in the complete food web,

$$L_1 = \{k \in V : D_k^{\leftarrow} = 0\} \tag{9}$$

We then define the subgraph food web, $G_i = (V_i, E_i)$, as that which excludes species in trophic level i and all trophic levels below i, where

$$V_i = V \setminus \bigcup_{k<i} L_k \tag{10}$$

and E_i is the corresponding edges between these vertices from the original graph. All species with trophic level i are then those which belong to the set

$$L_i = \left\{ k \in V_i : D_k^{\leftarrow, i} = 0 \right\} \tag{11}$$

i	K_i	K_i^{dir}	K_i^{indir}	K_i^{bu}	K_i^{td}	L_i
1	8.5	3	5.5	8.5	0	1
2	1.25	1	0.25	1.25	0	1
3	0.25	0.25	0	0.25	0	1
4	2.889	2.333	0.556	2.556	0.333	2
5	1.306	1	0.306	0.972	0.333	2
6	2.306	1.667	0.639	1.972	0.333	2
7	1.933	1.533	0.4	0.333	1.6	3
8	1.194	0.833	0.361	0.528	0.667	3
9	1.017	0.95	0.067	0.25	0.767	3
10	1.017	0.95	0.067	0.25	0.767	3
11	3.183	2.283	0.9	0.583	2.6	4
12	6.333	3.5	2.833	0	6.333	5
13	4.667	1.7	2.967	0	4.667	5
Mean	2.757	1.615	1.142	1.342	1.415	2.692
SD	2.315	0.902	1.572	2.196	1.904	1.323
Min	0.25	0.25	0	0	0	1
Max	8.5	3.5	5.5	8.5	6.333	5

Figure 4. Graph metrics for each species in the full Alaskan food web of figure 2. Metrics are respectively (L to R in columns 2 to 7) keystone index, directed keystone index, indirected keystone index, bottom-up keystone index, top-down keystone index and trophic level. Index numbers, i, for each species are drawn in figure 2.

where $D_k^{\leftarrow,i}$ is the prey degree of species k in the food web defined by G_i. In other words, a species is in trophic level 3, for example, if it has no prey species present after removing species in trophic levels 1 and 2. The bottom-up prioritisation ranks species first by ascending trophic level, and then by descending number of predators.

Policies and comparisons. In addition to the policies defined by graph metrics and the optimal policy, we also include the policy that manages nothing and a policy that chooses extant species to protect at random. The total expected reward of policies, calculated from the metric-policy evaluation algorithm, will be compared with the optimal value (the maximum total expected reward) from the exact solution, using the backwards induction algorithm, to find which approximation methods are the best.

Computational issues and remedies. To utilise the backwards induction algorithm, transition probabilities must be calculated for the current problem. The number of probabilities to calculate will be $\mid S \mid \times \mid S \mid \times \mid A \mid$. For a small real food web with 13 species and a budget with the capacity to protect 4 species at each time step, we have $2^{13}.2^{13}.C_4^{13} = 47,982,837,760$ transition probabilities to calculate (over 40 billion). With 25 species and a budget capacity to protect 8 species, this is more than 1.2×10^{21} transition probabilities, illustrating the need to use approximations to the optimal dynamic programming solution when dealing with larger food webs. Our modelling framework already uses generous assumptions to simplify the model and uses transition probabilities that do not change through time. Previously studied food webs have included over 150 different groups of organisms [17, 2] which is out of reach for the optimal solution in the current investigation.

Figure 5 plots the average CPU time over 10 different food webs for calculating four metrics against the food web size, of up to 25 species. Computations were performed using Matlab version 7.7.0.471 (Mathworks, 2009). A selection of 10 food webs of each size were randomly generated using a published method for generating food webs, the cascade model, and using a 'connectance' value, $C = |E|/|V|^2$, of 0.1 [24, 26]. A cascade model is one which is constructed by first assigning each species a uniform random number and secondly setting the probability that a species predates on species with a random uniform value less than its own to $2.C.n/(n-1)$ [24, 26]. Connectance values in real, observed food webs range between 0.1 and 0.2 [17, 26]. Calculations were performed under MacOS 10.6.8 on a 2.53GhZ, mid-2009 Macbook Pro with 4Gb of 1067MHz DDR3 RAM.

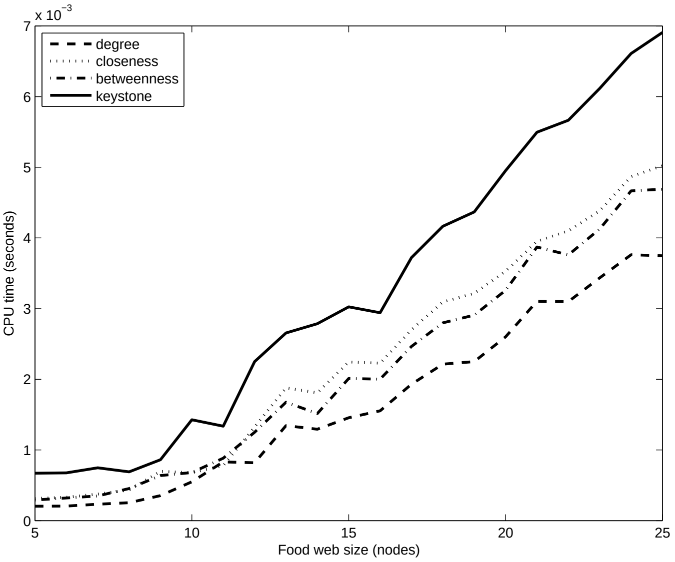

Figure 5. Average CPU time for calculation of four different metrics as a function of food web size.

A couple of steps can be used to remedy the curse of dimensionality. As mentioned, we assumed transition probabilities may be factorised into a product of species' local transitions. This means it is only necessary to calculate the probability of all possible local transitions for any species. Furthermore, various transitions in states can be set to zero based on the conditions of the transition probabilities. Concretely, consider the matrix $\mathbf{M} = \mathbf{SG}$, where $\mathbf{S} : \mathcal{X} \to 2^{\mathcal{X}}$, is a $2^n \times n$ Boolean matrix that indicates for each possible state which species is extant (one row for each state, columns index species) and $\mathbf{G}$ is the $n \times n$ adjacency matrix of the food web in question. That is, $\mathbf{G}_{ij} = 1$ if species i is a prey of species j and otherwise 0; cannibalism is not allowed $\mathbf{G}_{ii} = 0 \; \forall \; i$. The elements of the $2^n \times n$ matrix $\mathbf{M}$ are then

$$\mathbf{M}_{i,j} = \text{Number of extant prey of species } j \text{ when the state is } \mathbf{S}_{i,:}$$

where $\mathbf{S}_{i,:}$ is the ith row of $\mathbf{S}$. It is not necessary for basal species to have prey to survive so we set $M_{i,j} = 1$ if species j is a basal species. Further, defining $\mathbf{Q} = \mathbf{M} \odot \mathbf{S}$, where $\odot$ defines element-wise multiplication (the Hadamard product), we have a matrix that has elements

$$\mathbf{Q}_{i,j} = \begin{cases} 0 & \text{if species } j \text{ is extinct or has no prey in state } i \\ & \text{(and is not a basal species)} \\ \{1, 2, \cdots\} & \text{otherwise} \end{cases}$$

and thus $\mathbf{P}_{i,i',a} = 0$ if $\mathbf{Q}_{i,j} = 0 \; \forall j$ s.t. $\mathbf{S}_{i',j} > 0$.

3 Results

Despite the methods described above for speeding up computation time, the calculation of transition probabilities still takes a considerable amount of time. For illustrative purposes, we present initial results (table 6) for the Alaskan food web with only 10 species (by removing the Great whale, star fish and mussels/barnacles from the 13 species food web). It is possible to solve the exact solution to the 10 species web in several minutes on the computer described above.

Policy	$\upsilon_\delta(x^1)$
Optimal	5.92
$K^{\uparrow}$	5.52
BUP	5.52
$D^{\rightarrow}$	5.51
K	4.99
K^{indir}	4.97
K^{dir}	4.76
BC	4.00
D	3.84
CC	3.72
Random	3.66
$K^{\downarrow}$	3.49
$D^{\leftarrow}$	3.45
None	1.10

Figure 6. Preliminary results for the 10 species Alaskan food web (with Great whale, star fish and mussels/barnacles removed) over 10 time steps with $p_i^0 = 0.9,\ \forall i$, and budget of 4 protected species at each time step. Column 2 is the total expected number of species surviving in the final time step when using each policy to manage the full 10 species Alaskan food web for 10 years (higher is better). Rows are arranged in descending order of their values in column 2.

We present the total expected reward (and respectively the value in the exact policy) in the first time step for the 10 species food web for a project of 10 years, with an underlying probability of survival of $p_i^0 = 0.9, \forall i$, and budget capacity to protect 4 species at each time step. Terminal rewards are the final number of species surviving and per-time-step rewards are zero. The policies presented are based on the eleven metrics described above, the random policy, the policy of protecting no species, and the optimal policy.

Results from the 10 species Alaskan food web suggest that the heuristic policies that perform well compared to the exact solution are those which are based on metrics that acknowledge the number of extant predators that a species has (the bottom-up keystone index, predator degree, and bottom-up prioritisation). On the other hand, heuristic policies that prioritise management of species based on the number of extant prey perform worse than a random strategy (prey degree, top-down keystone index).

4 Discussion

We have discussed both exact methods for solving a Markov decision processes with an underlying graphical structure and the potential of heuristics, based on graph metrics, for guiding decision making when exact methods are not computationally feasible. Preliminary results have been presented for a 10 species food web and suggest that heuristic policies that prioritise species according to the number of extant predators perform the best. Results suggest management preference for species of a lower trophic level which may be associated with the assumption of guaranteed survival for protected basal species.

Ten species is relatively small for an ecological food web and, thus, findings should be taken cautiously at this early stage of research. Future research will compare the exact solution with heuristic policies for additional food webs in the ecological literature and use statistical analyses to search for patterns in the optimal policies that may be predicted by either graph metrics or local features of the food web. Heuristic policies that are found to approximate the optimal solution consistently will be used to simulate management of larger food webs. Future research will also investigate the potential of more sophisticated methods of approximating the exact method to ecological food-web management, similar to those used in graph-based MDPs [23].

ACKNOWLEDGEMENTS

Thanks to INRA-BIA in Toulouse for WP's travel and accommodation funding for the ECAI workshop, and thanks to the Australian Research Council's Centre of Excellence for Environmental Decisions for WP's funding to attend the workshop. We wish to acknowledge two reviewers for pertinent comments that improved this manuscript.

REFERENCES

[1] M. Almeida-Neto, P. Puimarães, P. R. Puimarães Jr, R. D. Loyola, and W. Ulrich, 'A consistent metric for nestedness analysis in ecological systems: reconciling concept and measurement', *Oikos*, **117**, 1227–1239, (2008).

[2] J. Bascompte, C. J. Melián, and E. Sala, 'Interaction strength combinations and the overfishing of a marine food web', *Proceedings of the National Academy of Sciences of the United States of America*, **102**, 5443–5447, (2005).

[3] M. A. Beauchamp, 'An improved index of centrality', *Behavioral Science*, **10**(2), 161–163, (1965).

[4] I. Chadès, T.G. Martin, S. Nicol, M.A. Burgman, H.P. Possingham, and Y.M. Buckley, 'General rules for managing and surveying networks of pests, diseases, and endangered species', *Proceedings of the National Academy of Sciences of the United States of America*, **108**, 8323–8328, (2011).

[5] I. Chadès, E. McDonald-Madden, M.A. McCarthy, B. Wintle, M. Linkie, and H.P. Possingham, 'When to stop managing or surveying cryptic threatened species', *PNAS*, **105**(37), 13936–13940, (2008).

[6] E. Estrada, 'Characterization of topological keystone species: local, global, and 'meso-scale' centralities in food webs', *Ecological Complexity*, **4**, 48–57, (2007).

[7] L. C. Freeman, 'Centrality in social networks: conceptual clarification', *Social Networks*, **1**, 215–239, (1978).

[8] L. C. Freeman, 'The gatekeeper, pair-dependency and structural centrality', *Quality and Quantity*, **14**, 585–592, (1980).

[9] M. R. Gardner and W. R. Ashby, 'Connectance of large dynamic (cybernetic) systems: critical values for stability', *Nature*, **228**, 784, (1970).

[10] F. Harary, 'Status and contrastatus', *Sociometry*, **22**(1), 23–43, (1959).

[11] F. Harary, 'A structural analysis of the situation in the Middle East in 1956', *The Journal of conflict resolution*, **5**(2), 167–178, (1961).

[12] R.A. Howard and J.E. Matheson, 'Influence diagrams', in *Readings on the Principles and Applications of Decision Analysis, Vol. II*, eds., R.A. Howard and J.E. Matheson, Menlo Park CA: Strategic Decisions Group, (1984).

[13] A. James, J. W. Pitchford, and M. J. Plank, 'Disentangling nestedness from models of ecological complexity', *Nature*, **487**, 227–230, (2012).

[14] F. V. Jensen and T. D. Nielsen, *Bayesian Networks and Decision Graphs*, Springer Verlag, 2001.

[15] F. Jordán, 'Keystone species and food webs', *Philosophical Transactions of The Royal Society, B*, **364**, 1733–1741, (2009).

[16] D. Koller and N. Friedman, *Probabilistic graphical models: principles and techniques*, MIT Press, 2009.

[17] N. D. Martinez, 'Artifacts or attributes? effects of resolution on the Little Rock Lake food web', *Ecological Monographs*, **61**, 367–392, (1991).

[18] R. May, 'Will a large complex system be stable?', *Nature*, **238**, 413–414, (1972).

[19] R. McDonald-Madden, E.and Sabbadin, P. W. J. Baxter, I. Chadès, and H. P. Possingham, 'Making food web theory useful for conservation decisions', *Submitted*, (2013).

[20] S.C. Nicol, I. Chadès, S. Linke, and H.P. Possingham, 'Conservation decision-making in large state spaces', *Ecological Modelling*, **221**, 2531–2536, (2010).

[21] W. J. M. Probert, C. E. Hauser, E. McDonald-Madden, M. C. Runge, P. W. J. Baxter, and H. P. Possingham, 'Managing and learning with multiple models: objectives and optimization algorithms', *Biological Conservation*, **144**, 1237–1245, (2011).

[22] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, John Wiley and Sons, New York, NY, 1994.

[23] R. Sabbadin, N. Peyrard, and N. Forsell, 'A framework and a mean-field algorithm for the local control of spatial processes', *International Journal of Approximate Reasoning*, **53**, 66–86, (2012).

[24] A. R. Solow and A. R. Beet, 'On lumping species in food webs', *Ecology*, **79**(6), 2013–2018, (1998).

[25] S. Wasserman and K. Faust, *Social Network Analysis*, Cambridge University Press, Cambridge, 1994.

[26] R. J. Williams and N. D. Martinez, 'Simple rules yield complex food webs', *Nature*, **404**, 180–183, (2000).