

Modelling spatial and temporal changes with GIS and Spatial and Dynamic Bayesian Networks



Yung En Chee ^{a, b, *}, Lauchlin Wilkinson ^c, Ann E. Nicholson ^c,
Pedro F. Quintana-Ascencio ^d, John E. Fauth ^d, Dianne Hall ^e, Kimberli J. Ponzio ^e,
Libby Rumpff ^a

^a School of BioSciences, The University of Melbourne, Parkville, VIC 3010, Australia

^b School of Ecosystem and Forest Sciences, The University of Melbourne, 500 Yarra Boulevard, Burnley, VIC 3121, Australia

^c Faculty of Information Technology, Monash University, PO Box 197, Caulfield East, VIC 3145, Australia

^d Department of Biology, University of Central Florida, 4000 Central Florida Blvd, Orlando, FL 32816-2368, USA

^e Division of Water Resources, St. Johns River Water Management District, PO Box 1429, Palatka, FL 32178-1429, USA

ARTICLE INFO

Article history:

Received 28 January 2016

Received in revised form

11 April 2016

Accepted 12 April 2016

Keywords:

Probabilistic graphical models

State-and-transition models

Object-oriented

Adaptive management

Willow

ABSTRACT

State-and-transition models (STMs) have been successfully combined with Dynamic Bayesian Networks (DBNs) to model temporal changes in managed ecosystems. Such models are useful for exploring *when* and *how* to intervene to achieve the desired management outcomes. However, knowing *where* to intervene is often equally critical. We describe an approach to extend state-and-transition dynamic Bayesian networks (ST-DBNs) — incorporating spatial context via GIS data and explicitly modelling spatial processes using spatial Bayesian networks (SBNs). Our approach uses object-oriented (OO) concepts and exploits the fact that ecological systems are hierarchically structured. This allows key phenomena and ecological processes to be represented by hierarchies of components that include similar, repetitive structures. We demonstrate the generality and power of our approach using two models — one developed for adaptive management of eucalypt woodland restoration in south-eastern Australia, and another developed to manage the encroachment of invasive willows into marsh ecosystems in east-central Florida.

© 2016 Elsevier Ltd. All rights reserved.

1. Introduction

Bayesian networks (Pearl, 1988) are increasingly popular for ecological and environmental modelling, decision support and adaptive management (Nyberg et al., 2006; Korb and Nicholson, 2010; Aguilera et al., 2011). Ecosystem management problems characteristically involve variable, complex and imperfectly understood biophysical, social and economic interactions. The iterative knowledge-engineering process of developing BNs is invaluable for: a) clarifying objectives; b) identifying and articulating alternatives; c) synthesising available knowledge; d) quantifying uncertainties and d) pinpointing critical assumptions to be tested by purposeful monitoring. When fully parameterised, such models help us explore and (where possible) resolve uncertainty about the consequences of management decisions. This is integral

to adaptive management (*sensu* Holling, 1978; Walters and Hilborn, 1978) which supplies the broader framework for evaluating the performance of decision actions and updating our knowledge base to improve future management (Nichols and Williams, 2006; Duncan and Wintle, 2008).

Despite the obvious value of using BNs to support learning over time for adaptive management (see e.g., Ames et al., 2005; Chee et al., 2005), most published examples of BNs for environmental applications have focused on formalising static conceptual models of the system in question, and do not explicitly represent ongoing dynamics (e.g. multiple time steps and sequential decisions) (Barton et al., 2012). Examples that incorporate spatiality explicitly are even rarer. Yet it is critical to address these gaps because the ability to understand change over time, and to account for spatial context and interactions is often necessary for meaningful decision support.

For instance, in our eucalypt woodlands case study, restoring species composition, ecosystem structure and function is a long-

* Corresponding author.

term undertaking that needs to effectively manage threats like weed establishment, so that the recovery process can build upon successive gains. In our invasive willows management case study, control efforts are long-term because adult willows have become firmly established within the catchment. In both cases, spatial considerations are crucial because the encroachment of weeds (in woodlands) and willow seedlings (in marsh ecosystems) depends on seed production and dispersal from surrounding areas, and spatial characteristics also determine the applicability and effectiveness of management actions.

State-and-transition dynamic Bayesian networks (ST-DBNs) as described by Nicholson and Flores (2011) provide a viable approach for explicitly modelling change over time. Here, we extend the capabilities of ST-DBNs – first, coupling them to GIS data so we can harness spatially relevant data, and then explicitly modelling key spatial processes using spatial Bayesian networks (SBNs). Our approach makes use of object-oriented (OO) concepts and exploits the fact that ecological systems are hierarchically structured such that key phenomena and processes of interest can be represented by nesting components that include similar, repetitive structures.

First, we explain the ‘buildings blocks’ and concepts of the tools we use for modelling spatial and temporal changes with BNs. We then present and illustrate our approach using two models—one developed for adaptive management of eucalypt woodland restoration in south-eastern Australia (‘Woodlands weed’ model, Rumpff et al. (2011)), and another developed to manage willow spread into marsh ecosystems in east-central Florida, USA (‘Willows’ model, Wilkinson et al. (2013)). Of course, incorporating spatial context and processes can lead to a massive increase in the size and complexity of the networks, which in turn generates computational issues and difficulties with the probabilistic updating—we discuss our approach to handling these challenges and provide a generic system architecture, templates and algorithms for combining GIS, object-oriented spatial BNs and object-oriented state-transition DBNs.

To our knowledge, this is the first demonstration of the integration of these three tools. This novel and powerful approach allows the incorporation of spatial context where it is critical for decision-making.

2. Background: building blocks and OO concepts

State-and-transition models (STMs) are management-focused, qualitative conceptual models that synthesise knowledge about an ecological system, in the form of observed and/or hypothesised system states and transitions that are of management interest (Westoby et al., 1989; Jackson et al., 2002). STMs are a popular tool for modelling changes over time in ecological systems that have clear transitions between distinct states. They combine graphical depiction of transitions and their causal factors with tables of qualitative descriptions of the transitions. They have been widely applied both to understand and help manage vegetation change in ecosystems such as rangelands (e.g., Westoby et al., 1989; Bestelmeyer et al., 2003; Bashari et al., 2009), grasslands (e.g., Sadler et al., 2010) and woodlands (e.g., Yates and Hobbs, 1997b; Rumpff et al., 2011).

Bayesian networks (BNs) are graphical models of cause-effect relationships used for reasoning under uncertainty. More formally, a Bayesian network (Pearl, 1988) is a directed, acyclic graph whose nodes represent the random variables in the problem. A set of directed arcs connect pairs of vertices, representing the direct dependencies of variables. The set of nodes pointing to X are called its parents and is denoted $pa(X)$. BNs display key variables in the system succinctly, show which variables are linked and how the causal chain or argument links events to outcomes of interest. The

relationship between variables is quantified by conditional probability tables (CPTs) associated with each node, namely $P(X|pa(X))$. The CPTs together compactly represent the full joint distribution. Users can set the values of any combination of nodes in the network that they have observed. This evidence, e , propagates through the network, producing a new posterior probability distribution $P(X|e)$ for each variable in the network. There are a number of efficient exact and approximate inference algorithms for performing this probabilistic updating, providing a powerful combination of predictive, diagnostic and explanatory reasoning.

Dynamic Bayesian Networks (DBNs) are a variant of ordinary BNs (Dean and Kanazawa, 1989; Kjærulff, 1992; Nicholson, 1992) that explicitly model changes over time and can be used to model feedback functions in problem contexts where this is important. A typical DBN has nodes for N variables of interest and for each domain variable X_i , there is one copy for each *time slice* for interest: $X_i^T, X_i^{T+1}, X_i^{T+2}$ etc. Links in a DBN include those between nodes in the same time slice, and those in the next time slice. Of the latter, temporal arcs may link the same variable over time, $X_i^T \rightarrow X_i^{T+1}$, and different variables over time, $X_i^T \rightarrow X_j^{T+1}$. Environmental applications employing DBNs are scarce (e.g., Shihab and Chalabi, 2007; Dawsey et al., 2007; Shihab, 2008). This may be because they are perceived to be “very tedious” (Uusitalo, 2007), or because DBN algorithms are available only in software resulting from research projects¹, with DBN functionality less well supported in popular commercial products.

State-and-transition Dynamic Bayesian Networks (ST-DBNs) combine the advantages of graphical visualisation of transitions and their influencing factors with quantitative representation of dependencies and uncertainty, along with explicit representation of time. Our example models are based on Nicholson and Flores (2011)’s template.

S^T represents the state of the system, has n possible values $s_1 \dots s_n$, and may directly influence any of the environmental and management factors, which are divided into m main factors, F_1, \dots, F_m (which directly influence transitions) and other sub-factors, X_1, \dots, X_r (which influence the main factors).

Transition nodes, ST_1, \dots, ST_n , represent the transitions from each state s_i . Each has at most $n + 1$ values (usually fewer), one for each “next” state plus “impossible”, giving explicit modelling of impossible transitions. Like ordinary DBNs, there is an implied δT , which can be included explicitly as a parent of all the ST nodes, if the time step varies. Each transition node ST has only some of the causal factors as parents. The CPT for the ST node is just a partition of the corresponding CPT if the problem was represented as an ordinary DBN, without the transition nodes. The next state node, S^{T+1} , has to combine the results of all the different transition nodes, given the starting state S , and thus has $n + 1$ parents. However, the relationship between the transition nodes and S^{T+1} is deterministic, so the CPT can be generated from a straightforward equation.

It is important to note that ST-DBNs that explicitly model all the transitions, only remain tractable when there are natural constraints in the domain; that is, if the number of transitions from each state is limited and only influenced by a small number of causal factors such that the underlying state transition matrix for S is sparse (Nicholson and Flores, 2011).

2.1. How does object-oriented (OO) thinking help?

The complexity of ecological systems is such that representing even a moderate degree of ecological realism tends to lead to large networks. The resulting visual ‘clutter’ of large networks makes

¹ e.g. BNT, code.google.com/p/bnt.

working on submodels cumbersome and less readily communicable to stakeholders. This problem of network size is compounded when temporal and spatial dynamics are explicitly integrated.

Although examples in environmental modelling are scant (Molina et al., 2010; Carmona et al., 2011; Johnson and Mengersen, 2012), OO modelling techniques can help: a) manage BN complexity via abstraction and encapsulation, b) facilitate the construction of classes of objects that are internally coherent and potentially more reusable, and c) formalise interfaces so that information flow between OOBNs is properly defined prior to integration (Koller and Pfeffer, 1997; Neil et al., 2000; Kjærulff and Madsen, 2008; Korb and Nicholson, 2010; Molina et al., 2010).

We follow the definition of OOBNs used in Kjærulff and Madsen (2008), and implemented in the Hugin BN software package. A standard BN is made up of ordinary nodes, representing random variables. An OOBN class is made up of both nodes, and objects, which are instances of other classes. Thus an object may encapsulate multiple sub-networks (i.e. OOBNs can contain other OOBNs), giving a composite and hierarchical structure.

Objects are connected to other nodes via some of its own ordinary nodes, called *interface* nodes. Interface nodes specify how other objects may interact with it, and are divided into *input* nodes and *output* nodes. Input nodes are the root nodes within an OOBN class, and when an object (instance) of that class becomes part of another class, each input node may be mapped to a single node (with the same state space) in the encapsulating class. The output nodes are the only nodes that may become parents of nodes in the encapsulating class. Non-interface nodes are not visible to the “outside world”, thus hiding information detail. An OOBN class can therefore be thought of as a self-contained template for an OOBN object, described by its name, its interface and its hidden part. Potentially, this allows classes to be used as libraries, to be reused and combined into a model in different contexts (Koller and Pfeffer, 1997).

Since ecological systems are hierarchically structured and key phenomena and processes of interest can be represented by nested components of similar, repetitive structures, we exploit this and convert our example ST-DBNs and SBNs into submodels represented by objects. Being able to access and display submodels at varying levels of abstraction within a time slice is useful for communicating complicated nested model structure to different audiences. A single time slice of a ST-OODB can also be thought of as an object, allowing model structure to be handily replicated over multiple linked time periods.

3. Two case studies

Here we describe the problem domain and the spatial and dynamic processes that need to be modelled for each case study.

3.1. Woodlands case study

Temperate eucalypt woodlands were formerly widespread throughout southern Australia. Extensive clearing and grazing has dramatically reduced their extent to remnants of varying size, quality and isolation (Yates and Hobbs, 1997a). These changes have led to regional losses of woodland species diversity, community structure and ecosystem function. Site and landscape level protection and restoration is necessary to halt and reverse degrading processes such as further clearing, grazing, nutrient enrichment, altered fire regimes and invasion of exotic species (Yates and Hobbs, 1997a).

Assessing the restoration requirements and recovery trajectories of woodland sites with different histories, starting states and variable environmental settings is complex and involves many

uncertainties. The Woodlands model (Rumpff et al., 2011, 2012) set up the adaptive management conceptual framework in the form of a ST-BN developed in collaboration with domain experts. A ST-OODB version of this model, representing change from time t to $t + 1$ at location (x, y) , is shown in Fig. 1. Distinct woodland states (e.g. “Reference”, “Oldfield”, “Native pasture”, “Thicket”; Fig. 1) are characterised by different combinations of structural and compositional attributes (called state variables). These include attributes such as native species cover and richness, woody recruitment and weed cover. State variables may change in response to environmental factors, process variables and management actions. In the context of this model, environmental variables refer to ‘given’, non-modifiable conditions such as the site’s land-use history (e.g. time since cropping). Process variables, on the other hand, represent site conditions that are alterable (e.g. grazing pressure and soil phosphorus (P) enrichment), and these process variables are affected to different degrees by the chosen type and level of intensity of available management actions (e.g. stocking rate, weed control and soil treatment). Management actions such as direct seeding and the planting of tube stock can also directly influence state variables. Threshold changes to state variables result in transitions between woodland states.

Here we concentrate on modelling the spatial processes that influence the weed cover state variable. Since weed cover at a site is strongly influenced by whether there are weed sources in the surrounding neighbourhood, spatially explicit modelling of this is necessary for effective management.

3.2. Willows case study

The Upper St. Johns River in east-central Florida has been the focus of considerable restoration investment (Quintana-Ascencio et al., 2013). However, the original herbaceous marshlands are increasingly being invaded by woody shrubs such as Carolina willow (*Salix caroliniana* Michx.) (Kinser et al., 1997). This change is undesirable because extensive willow thickets obscure viewsheds, reduces local vegetation heterogeneity, habitat diversity, navigable access and scope for recreation activities such as wildlife viewing, fishing and hunting. The evapotranspiration rate of *S. caroliniana* is much greater than that of herbaceous marsh vegetation and so, large-scale vegetation changes in this headwater catchment may reduce the amount of water available to both the riparian ecosystem and humans. Managers seek to control the overall extent of willows, their rate of expansion into other extant wetland types and encroachment into recently restored floodplain habitats. Spatial context matters because areas differ in biodiversity, aesthetic and recreational value, “invasibility” and applicable interventions; intact vegetation communities are less susceptible to invasion and prescribed fire depends on the type of vegetation present and its “burnability”.

The Willows model captures current knowledge and uncertainties about how key life-history stages and attributes of *S. caroliniana* respond to environmental factors and management actions (Fig. 2). Its purpose is to serve as a decision support tool for managers to test and learn about the landscape-scale consequences of different management strategies (Nicholson et al., 2012; Wilkinson et al., 2013).

Modelling willow spread requires spatially explicit data on willow occupancy, an understanding of seed production, dispersal, germination and survival, and how the key life-history stages respond to environmental factors and management actions.

Model development drew upon knowledge derived from ecological and physiological theory, field observations, field and greenhouse experiments and experts (e.g., Kinser et al., 1997; Pezeshki et al., 1998; Lee et al., 2005a, b; Ponzio et al., 2006;

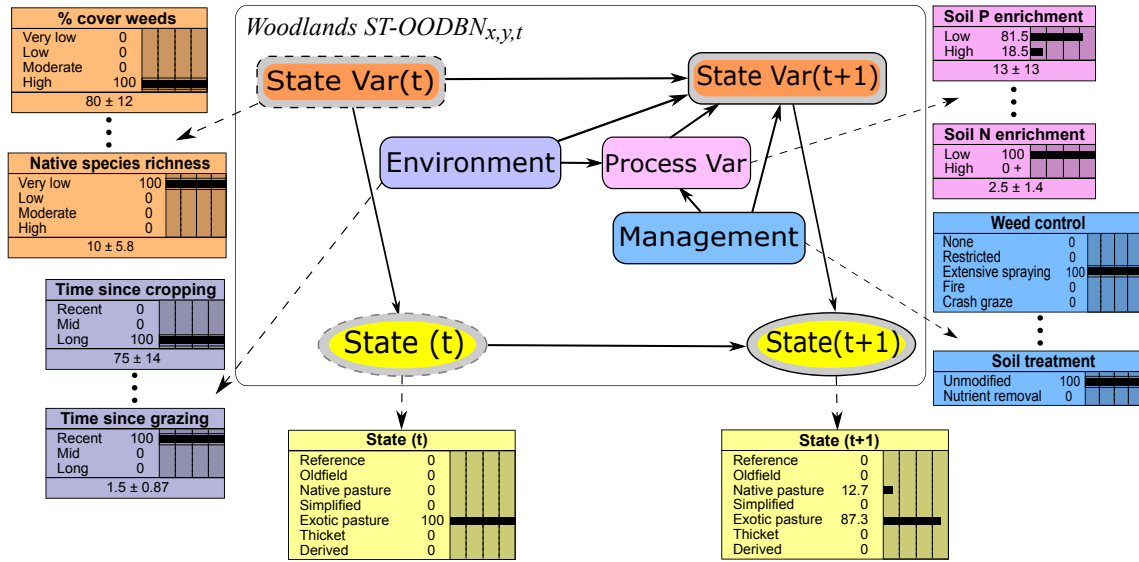


Fig. 1. ST-OODB of eucalypt woodland restoration showing the key input nodes (dotted and shaded outlines), and output nodes (bold and shaded outline). Round cornered boxes represent OOBN classes of the main types of variables that define and ultimately drive changes in woodland state (in yellow). These include: state variables (in orange) which represent sets of structural and compositional attributes that collectively define distinct woodland states; environmental variables (in purple) which describe non-modifiable site environmental conditions such as land use history; process variables (in pink) that represent modifiable environmental factors (e.g. soil nitrogen (N) enrichment); and management actions (in blue) that represent discrete actions (e.g. weed control) that can be implemented to influence state variables directly, or indirectly via process variables. The full model has eight state variables, nine environmental, five process and ten management variables. For brevity, we show just two illustrative examples of each variable type (belief-bar boxes indicated by dotted arrows). The reader is referred to Rumpff et al. (2011) for details of the complete set of model variables. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

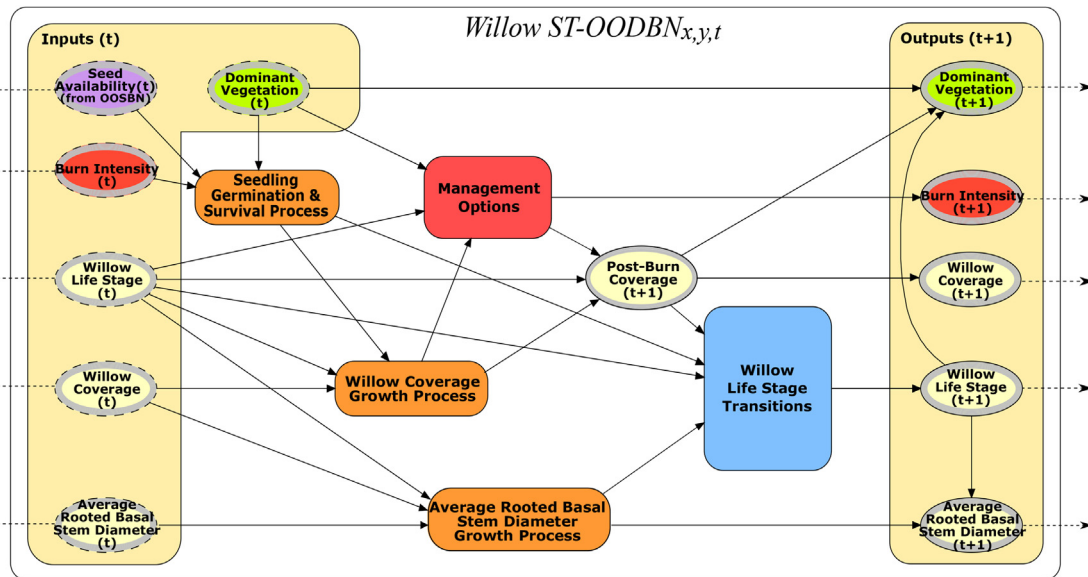


Fig. 2. ST-OODB for willow management showing the key input nodes (dotted and shaded outlines) and output nodes (bold and shaded outlines). Note that *Seed Availability (t)* for germination (in purple) depends on willow seed production and dispersal from a cell's surrounding region and this process is modelled by a separate OO network described in Section 4. Round cornered boxes represent OOBN classes for the key processes we model. They include seedling germination and survival, annual growth of willow coverage and stem diameter (in orange), management options (in red) and transitions between willow life stages (in blue). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Quintana-Ascencio and Fauth, 2010).

We used a modelling unit of 100 m × 100 m (1 hectare) in size and a time step of one year. The stages of management interest modelled in the Willow Life Stage node are: unoccupied, yearling, sapling (non-reproductive juvenile) and adult. The key points of interest are whether willows are present in a cell or not. If present, what life stage are they in, what is the average size (Rooted Basal Stem Diameter) and what is the coverage (i.e. proportion of the cell

that is occupied) (see outputs in Fig. 2).

For each cell (spatial unit), data on environmental attributes such as water availability, soil, vegetation type and landscape position and context is supplied from GIS data. This data provides inputs to parameterise the ST-OODB and seed dispersal model. The chosen cell size reflects the resolution of available spatial data for environmental attributes, makes the computational demand associated with seed dispersal modelling feasible, and is a

reasonable scale with respect to candidate management actions. A time step of one year was considered appropriate given the willow's growth and seed production cycle (Nicholson et al., 2012).

4. Modelling spatial processes with object-oriented spatial BNs (OOSBNs)

Scale is an inherent consideration in modelling spatial processes because the interactions of interest may vary from highly localised to spatially extensive. These contrasting situations are exemplified by each of our case studies (described below). We therefore developed a generic template with the flexibility to accommodate different scales of interest.

In this generic OOSBN (Fig. 3), the *Process of Interest* takes into account any salient OOSBN inputs and can be represented by a single node or modelled in greater detail by an object. The scale of the *Process of Interest* is determined by what we call the spatial region of influence (*RoI*) and this in turn, is defined by the *Distance Factor*. Each cell (x, y) is influenced by all cells within its *RoI* (i.e. (x'_1, y'_1) to (x'_n, y'_n)) and the contributions of these cells is depicted by the multiple *SpatialProcess* objects shown within the OOSBN (Fig. 3). These contributions are accumulated to provide inputs to the *Process of Interest* and ultimately, to compute the required output for the OOSBN.

We present a simple application of the generic OOSBN using the woodlands study and a more sophisticated instance with the willows study.

4.1. Woodlands weed model

We are concerned here with clumping weed species whose spread is localised, proceeding at a maximum rate of one cell unit

per time step. We used a modelling unit of $25\text{ m} \times 25\text{ m}$ and a time step of one year. The goal is to model expected weed cover at a target cell after taking into account weed spread from its neighbouring spatial region of influence (*RoI*). In this case, the *Distance Factor* equals one and the *RoI* includes just the eight immediately adjacent cells (Fig. 4, left). The process of weed spread is represented by a single node parameterised by expert elicitation (*Weed Spread Process* in Fig. 4).

Because neighbouring weed contribution is a simple additive function, we create a *Cumulative Cover* node and successively add the weed contribution from each adjacent cell to the *Cumulative Cover* (Fig. 4). This is equivalent to repeatedly “divorcing” parents to reduce the size of the combinatorial state space (Jensen, 1996). The modelling of localised weed contribution is thus achieved by running the OOSBN (Fig. 4) over the spatial dimension of the eight adjacent cells in a similar way to rolling out a DBN over time.

4.2. Willow seed production and dispersal

The goal here is to model to total amount of seed available for germination at a target cell, after taking into account seed production and dispersal from surrounding cells at a given distance. Here, the *Process of Interest* incorporates separate, explicit models for seed production and dispersal over a distance. The number of seeds arriving at the target cell is then accumulated to compute the total quantity of seed available for germination. The sub-models are described in detail below.

4.2.1. Seed production

S. caroliniana flowers in early spring and produces vast numbers of small seeds (a median of ~165,000 seeds annually; Quintana-Ascencio et al., unpublished data). We model seed production

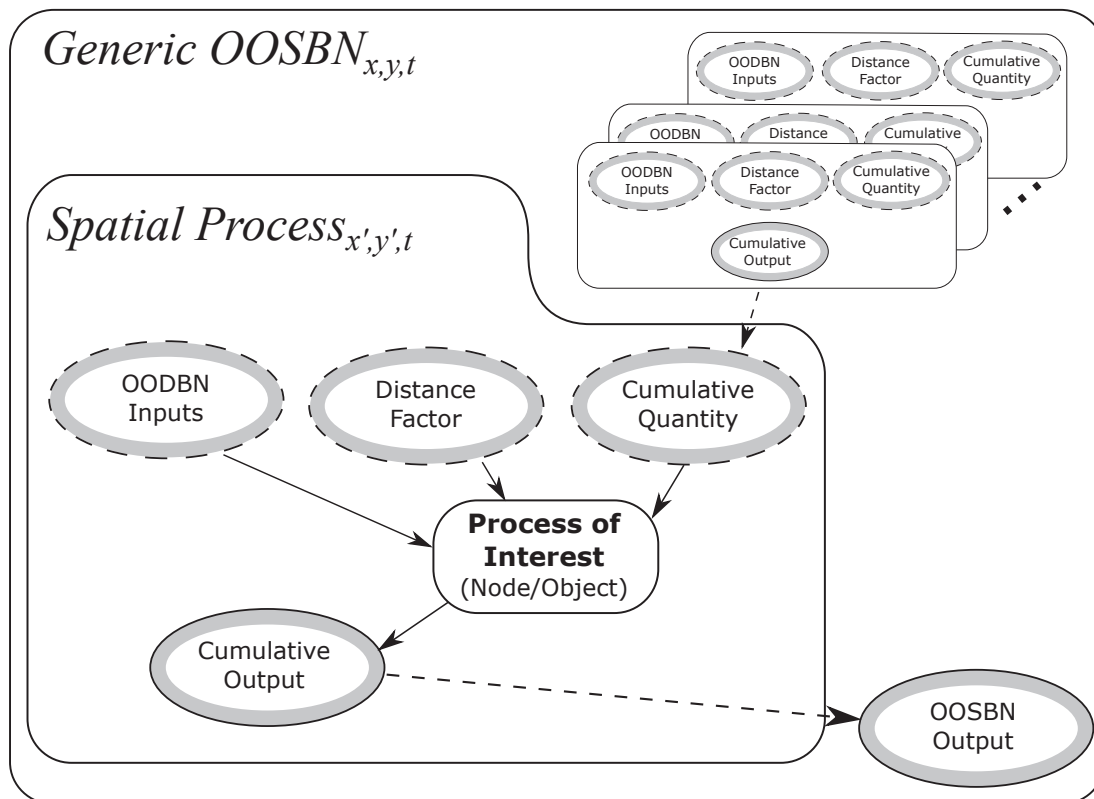


Fig. 3. Generic OOSBN template for modelling any spatial *Process of Interest* at variable spatial scales (via *Distance Factor*). The contribution of individual cells within the region of influence (*RoI*) is denoted by the multiple *Spatial Process* objects.

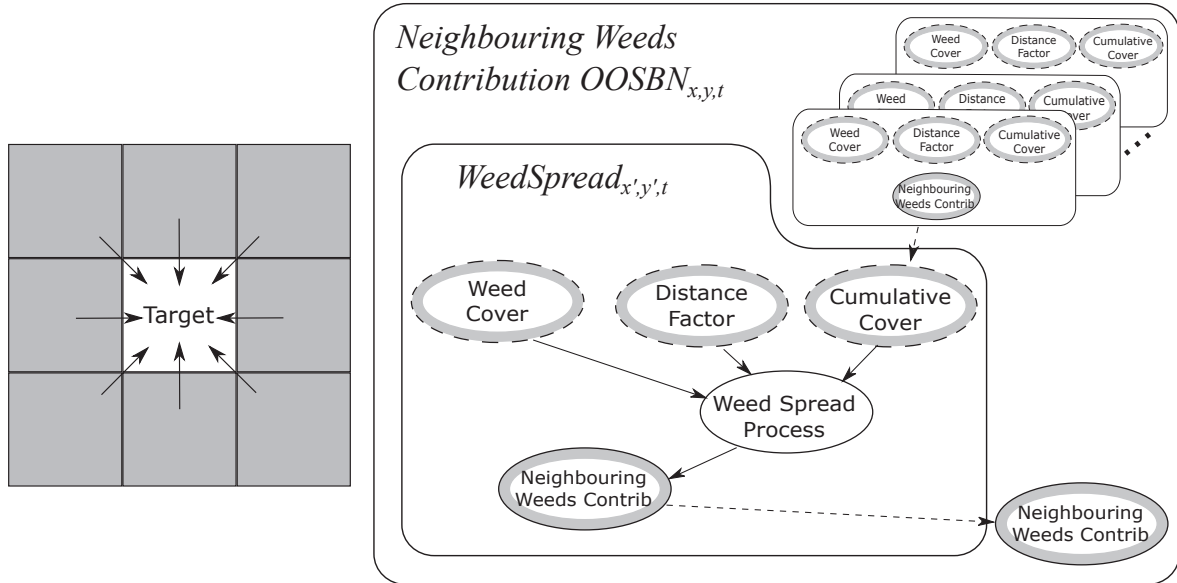


Fig. 4. Neighbouring Weeds Contribution OOSBN used in Woodlands model. As weed spread is localised, the *Distance Factor* is set at one cell unit and the *RoI* is restricted to the eight immediately adjacent cells (i.e. x'_1, y'_1 to x'_8, y'_8) of the target cell (x, y) (left). The contributions from these cells are shown by the multiple *WeedSpread* objects.

with a *WillowSeedProduction* OOBN, which is embedded in the broader *Seed* OOSBN shown in Fig. 5. The number of seeds produced by an adult is given by the product of the number of *Inflorescences*, the number of *Fruits per inflorescence* and the number of *Seeds per fruit*. *Fruits per inflorescence* and *Seeds per fruit* are defined by distributions estimated from empirical data. The number of *Inflorescences* increases as a function of adult size (represented by *Rooted Basal Stem Diameter*) and this relationship is estimated from empirical data.

Coverage is the percentage of a 1 hectare cell that is occupied by willows and *Average Canopy Area* is modelled as a function of *Rooted Basal Stem Diameter*. Together these two variables provide an estimate of the number of reproductive stems. Overall seed

production within a cell, *Seeds per Hectare*, is then simply the product of the seed production per stem, by the number of reproductive stems.

4.2.2. Seed dispersal

We model seed dispersal phenomenologically rather than mechanistically. Wind-mediated seed dispersal is calculated using the Clark et al. (1999) dispersal kernel:

$$SD_{x,y}^{x',y'} = SP_{x',y'} \times \frac{1}{2\pi\alpha^2} e^{-\left(\frac{d}{\alpha}\right)^2} \quad (1)$$

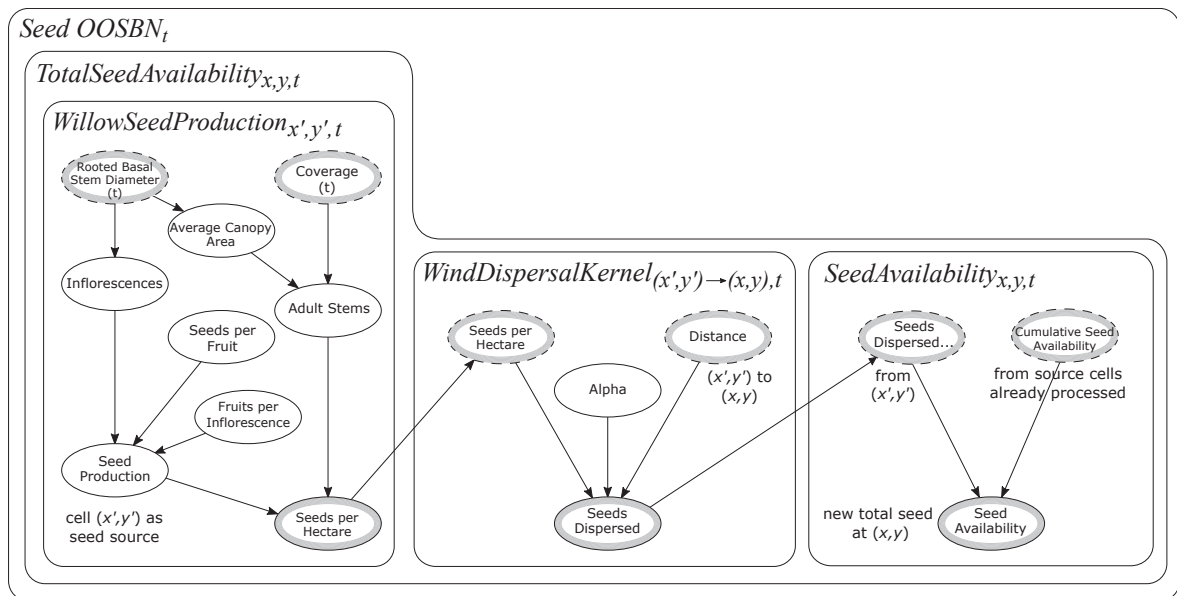


Fig. 5. Willow Seed OOSBN architecture showing the embedded OOBNs used to compute seed production, dispersal by wind and total seed availability for germination. Inputs for *Routed Basal Stem Diameter* and *Coverage* are used to estimate the *Seeds per Hectare* produced within source cell (x', y') at time t . The amount dispersed to each destination cell (x, y) depends on *Seeds per Hectare* from (x', y') and the *Distance* between cell (x, y) and (x', y') . Final *Seed Availability* for germination at cell (x, y) combines the amount of seed that has arrived from (x', y') with *Cumulative Seed Availability*—the seed arriving from source cells processed previously).

where $SD_{x,y}^{x',y'}$ is the number of seeds arriving at cell (x, y) from those produced at a cell (x', y') ; it is the product of seed produced $SP_{x',y'}$ and an exponential kernel where d is the distance between cells (x, y) and (x', y') , and α is a distance parameter.

To simulate stochasticity in dispersal events, α is a random variable that can be sampled from distributions designed to reflect the expected nature of dispersal, such as short versus long distance dispersal (Fox et al., 2009). This seed dispersal model is captured within the $WindDispersalKernel_{(x',y') \rightarrow (x,y),t}$ OOBN (Fig. 5), where the input *Distance* node is set as the distance between cells (x, y) and (x', y') , and α is set as a discretised normal distribution with a mean of 1 and a variance of 0.25.

4.2.3. Seed availability for germination

The number of seeds dispersed from a seed-producing cell declines exponentially with increasing distance from that cell such that after a certain distance, the number of seeds dispersed is effectively negligible. We model this by specifying a *RoI*. Seed availability $SA_{x,y}$ for a target cell (x, y) , is then the sum of the seeds dispersed to it from every seed-producing cell in its *RoI*:

$$SA_{x,y} = \sum_{x',y' \in RoI} SD_{x,y}^{x',y'} \quad (2)$$

The *RoI* can be designed to take on different shapes and sizes to reflect potentially important influences on wind dispersal such wind direction, wind strength and terrain characteristics. For simplicity in this example however, we assume a circular *RoI* with a radius of 800 m (eight cells). This implies $\pi 8^2$, or ~ 201 cells providing parents to the final *Seed Availability* node. If the *Seed Availability* node is discretised to n states and the *Seeds Dispersed* node discretised to m states, the CPT for *Seed Availability* would include $n \times m^{\pi 8^2}$ probabilities – a massively infeasible number for any realistic m !

However, because *Seed Availability* (like neighbouring weed contribution) is a simple additive function, we can add *Seed Dispersed* from each source cell, using the *Cumulative Seed Availability* node to store the running total (Fig. 5). So in this example, we

run the seed OOSBN $\pi 8^2$ times for each cell (x', y') in the study area to compute seed production, dispersal and availability for germination at every cell.

4.3. Belief updating

In both case studies, the OOSBNs for every cell within the defined spatial *RoI* can be linked together to produce a single very large OOBN. However, in the BN software, the exact belief updating algorithm used to compute the posterior probability distributions first compiles the BN into a different underlying tree structure (Jensen et al., 1990), which for such a complex BN produces huge probability tables that far exceed even what a high-end computer can store in memory, thus rendering belief updating impossible.

We therefore developed an algorithm (Algorithm 1) to provide a scalable and computationally feasible solution for belief updating of the weed spread and seed dispersal OOSBNs. The *PROPAGATE* function in Algorithm 1 starts by taking the OOSBN, a list of *PTLayers* whose cells correspond to the study area, and the user-defined *RoI*.

PTLayers are simply a type of internal data structure that combines the spatial structure of a GIS, with distributions for the (discretised) nodes in the networks. In the first instance, GIS data representing initial conditions such as starting *Weed Cover* for the woodlands weed study or *Coverage* for the willows study is stored in the *PTLayer* for each node. Each *PTLayer* contains a number of fields, one for each of the node states of the linked input (*I*) and output (*O*) nodes. Each field stores the probability mass of the corresponding node state. In our scheme, *PTLayers* are used to provide intermediate storage across the spatial grid. More specifically, we use *PTLayers* to store and pass the spatially referenced prior distributions of input (*I*) nodes and posterior distributions of output (*O*) nodes for both the OOSBNs (Figs. 4 and 5) and *ST-OOSBNs* (described in Section 5).

Algorithm 1 works in an analogous manner in both case studies, but here we describe how it works with respect to the more complex willow seed production and dispersal OOSBN.

Algorithm 1 An algorithm for weed spread or seed production and dispersal across a user-defined study area and *RoI* using an OOSBN

```

1: function PROPAGATE(OOSBN, PTLayers, Area, RoI)
2:   I ← I(OOSBN)
3:   O ← O(OOSBN) // note: in the Willows case study, SA is the sole output node
4:   for all (x', y') ∈ Area do
5:     L = getLayer(PTLayers, CSA)
6:     p(Lx',y' = none) ← 1 // cumulative seed availability initialised to zero
7:   end for
8:   for all (x', y') ∈ Area do // seed source cell (x', y')
9:     for all Ii ∈ Ix',y'(OOSBN) do
10:      L ← getLayer(PTLayers, Ii)
11:      p(Ii) ← L(x', y')
12:    end for
13:    for all (x, y) ∈ RoI do // seed destination cell (x, y)
14:      L = getLayer(PTLayers, SA)
15:      p(CSAx,y) ← L(x, y) // CSA input set to previous SA output
16:      d ← getDistance({x, y}, {x', y'}, RoI)
17:      p(Distance = d) ← 1
18:      update beliefs in OOSBN
19:      L ← getLayer(PTLayers, SA)
20:      L(x,y) ← Bel(SAx,y) // pushes seed availability out to PTLayers
21:    end for
22:  end for
23:  return PTLayers
24: end function

```

In Algorithm 1, $I(OOSBN)$ and $O(OOSBN)$ denote functions that return the OOSBN interface input and output nodes respectively. $getLayer(PTLayers, V)$ denotes a function that returns the $PTLayer$ corresponding to a node V . At commencement, a $PTLayer$ that tracks seed contributions (cf. *Cumulative Seed Availability* in Fig. 5) is initialised to 0 at all (x', y') co-ordinates. Algorithm 1 then loops through each cell (x', y') in the study area, setting the distribution for the other input nodes $I_{1...M}$ (e.g., *Rooted Basal Stem Diameter* and *Coverage*, Fig. 5) from their corresponding $PTLayer$ cell. The algorithm then enters a second loop for every cell (x, y) that is a possible seed destination based on the RoI . The *Cumulative Seed Availability* input node is set (via the $PTLayer$) to the previous seed availability output, as the model maintains a running total of seed that has arrived (x, y) at from source cells already processed. The *Distance* node is set using a distance function ($getDistance$). In the wind dispersal kernel, this is the Euclidean distance between the current cell and the target co-ordinates. Finally, belief updating is done within the OOSBN (Fig. 5) and the beliefs (posterior probability distributions) from the output (*Seed Availability*) node at each cell are saved back to the appropriate $PTLayer$. In our Willows example, the $SA_{x,y}$ $PTLayer$ will contain the overall RoI -derived total seed available for germination in each cell. Given a study area of $N \times N$ cells, with a region of interest with radius r cells, the upper bound on the computational complexity of Algorithm 1 is $O(N^2r^2)$.

one ST-OOSBN. In practice, we do not require multiple individual copies of OOSBNs and ST-OOSBNs, but rather re-use network structures, whose input nodes are re-parameterised for each cell, at each time step.

With reference to Fig. 6 and Algorithm 2, the main steps are as follows:

1. At time t , GIS data for initial conditions at each cell (x, y) in the study area are read into the relevant $PTLayers$. In situations where no GIS layers are available for input, a user-defined (and properly justified) prior distribution can be used.
2. $PTLayers_{1...l}$ are used to initialise the priors of nodes $I_{1...m}$ and $I_{1...M}$ of $ST-OOSBN_{x,y,t}$ and $OOSBN_t$ respectively.
3. After propagation within $OOSBN_t$ using Algorithm 1, beliefs from the output node O are stored in $PTLayer_{OOSBN}$, and used to update the priors of the *Seed Availability* input node, S , of the $ST-OOSBN_{x,y,t}$.
4. Belief updating is then done with the $ST-OOSBN_{x,y,t}$ and the beliefs from output nodes $O_{1...n}$ copied back to the relevant corresponding $PTLayers_{1...b}$ ready to provide updated priors for nodes $I_{1...m}$ and $I_{1...M}$ at time $t + 1$.

When running the model for T time steps, the computational complexity of Algorithm 2, which calls Algorithm 1, is $O(N^2r^2T)$.

Algorithm 2 An algorithm for propagating a GIS-coupled ST-OOSBN with OOSBN sub-networks

```

1: function PROPAGATE(ST-OOSBN, Area, RoI, PTLayers, t)
2:    $I \leftarrow I(ST-OOSBN)$ 
3:    $O \leftarrow O(ST-OOSBN)$ 
4:    $S \leftarrow S(I)$  // Seed availability input node in the ST-OOSBN
5:   for  $t := 0$  to  $t$  do
6:     PROPAGATE(OOSBN, PTLayers, Area, RoI)
7:     // Produces new priors for seed availability  $S$ 
8:     for all  $(x, y) \in Area$  do
9:       for all  $I_i \in I$  do
10:         $L_j \leftarrow getLayer(PTLayers, I_i)$ 
11:         $p(I_i) \leftarrow L_j(x, y)$ 
12:       end for
13:       update beliefs in ST-OOSBN
14:       for all  $O_i \in O$  do
15:         $L_j \leftarrow getLayer(PTLayers, O_i)$ 
16:         $L_j(x, y) \leftarrow Bel(O_i)$ 
17:       end for
18:     end for
19:   end for
20:   return PTLayers
21: end function

```

This approach of iteratively rolling out a single OOSBN instance across space is equivalent to the standard “roll-out” followed by “roll-up” operation done with two time slice DBNs (Boyen and Koller, 1998) to avoid the computational complexity of rolling out a DBN all at once over a large number of time steps. This mitigates the problem of infeasibly large tables and converts the problem to one of computation time. This opens up the possibility of improving computation efficiency via parallel computing.

5. System architecture and algorithm for integrating GIS data, the ST-OOSBN and the OOSBN

Fig. 6 illustrates our system architecture, showing the interactions between GIS data and the input (I) and the output (O) nodes of the spatial (OOSBN) and temporal (ST-OOSBN) networks. For each cell in the study area, there is conceptually one OOSBN and

In our scheme, $PTLayers$ are used to provide intermediate storage across the spatial grid, and inputs for the next time step, $t + 1$, come not only from the network for the same cell, but from outputs of networks for *other* cells. This is necessary for modelling the processes of seed production, dispersal and availability for germination, as described in Section 4.

In effect, the $PTLayers$ replace both the *spatial* arcs between the networks for different cells (i.e. the cross-network arcs from seed production in one place to seed availability in another), and the *temporal* arcs if the network was rolled-out over many time steps. This method is limited to prediction only; we cannot reason backwards from some given state to identify the starting states and management actions required to achieve a preferred end-state. The specific tools used to implement the software architecture are described in Appendix A.

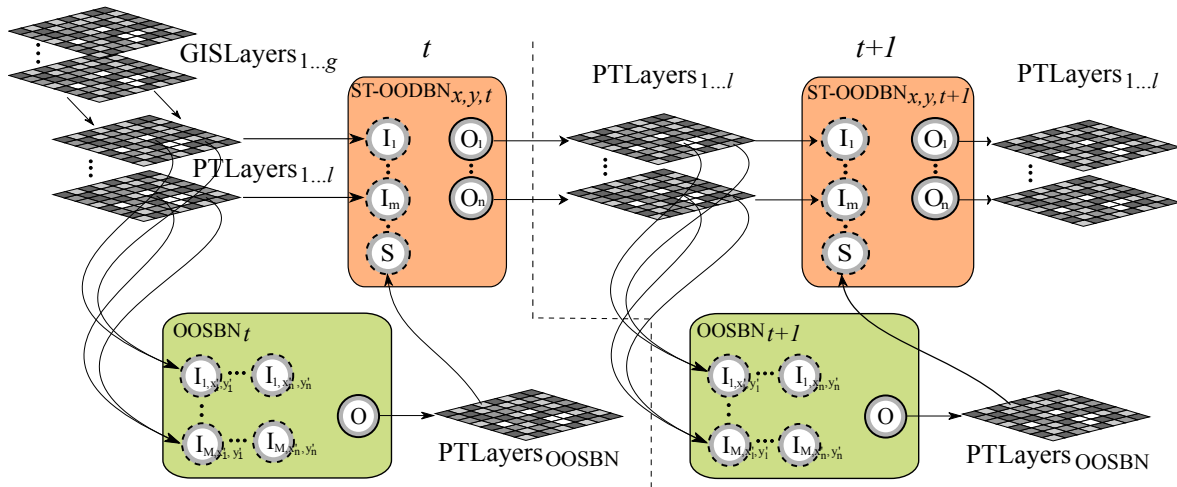


Fig. 6. An abstract representation of our system architecture for integrating GIS data, and spatial and dynamic OOBNs (refer to text for detailed description).

6. Model demonstration and results

6.1. Example scenarios for model demonstration

To demonstrate our working implementation, we ran the woodlands weed model for an area of grassy eucalypt woodlands near Wollert, Victoria (86 by 80 cells). Weed cover in the target area was initialised using GIS data (White et al., unpublished data) and the timeframe of interest was 15 years. We investigated two scenarios — a) no management intervention at all, and b) intervention when weed cover is low.

For the willows study, we ran the model for a portion of the Blue Cypress Marsh Conservation Area (44 × 45 cells) within the Upper St. Johns River basin, for a management horizon of 20 years. Willows coverage was initialised using GIS data from 2008 to 2009 (SJRWMD, unpublished data). At commencement (t = 0), adult willows occupied ~15% of the total study area. We investigated three scenarios — a) no intervention, b) burn management of cells when the probability that willow stage = yearling is >10%, and c) burn management of cells when the probability that stage = sapling is >10%.

6.2. Results

In the woodlands weed study, weed cover across the target area is initially low, with a scattering of small, spatially disjunct patches

of moderate weed cover (Fig. 7). In the absence of management intervention however, the patches of weed cover increase from moderate to high and then very high cover over time — eventually coalescing into large swathes of very high weed cover that dominate the target area (Fig. 7a).

The strategy of applying management whenever weed cover in a cell is expected to be low does not stop patches with initially moderate cover increasing to very high cover over time. But it helps to “contain” and isolate such patches so that relatively weed-free areas can persist in spaces within a matrix of heavily infested patches (Fig. 7b).

In the willows case study we present the distribution of three model outputs — the probability of each willow stage, the probability of willow coverage and the most likely seed availability level — for each management option (Fig. 8). We first consider the ‘no intervention’ option (Fig. 8a): at t = 5, there is a clump of adult willows at the bottom-centre of the study area with very high probability of very high cover. Immediately bordering this clump is a band of cells where willows are most likely to be at the sapling stage, and to be at a low level of cover. A substantial proportion of the remaining study area north of the clump of adults is most likely unoccupied by willows. However, due to the productivity of willows and the effectiveness of dispersal, seed availability is high to very high in the areas surrounding the adults, but also moderate to high in patches where there are likely to be saplings. In the absence of intervention, the bottom clump of adult willows expands over

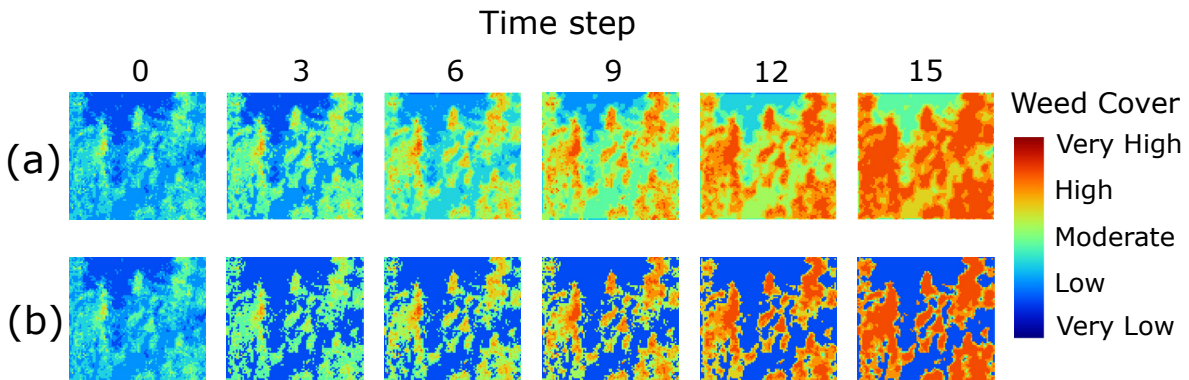


Fig. 7. Woodlands weed model results: maps of (most likely value of) predicted weed cover at three yearly time slices when there is (a) no management intervention, and (b) intervention in cells where weed cover is expected to be low.

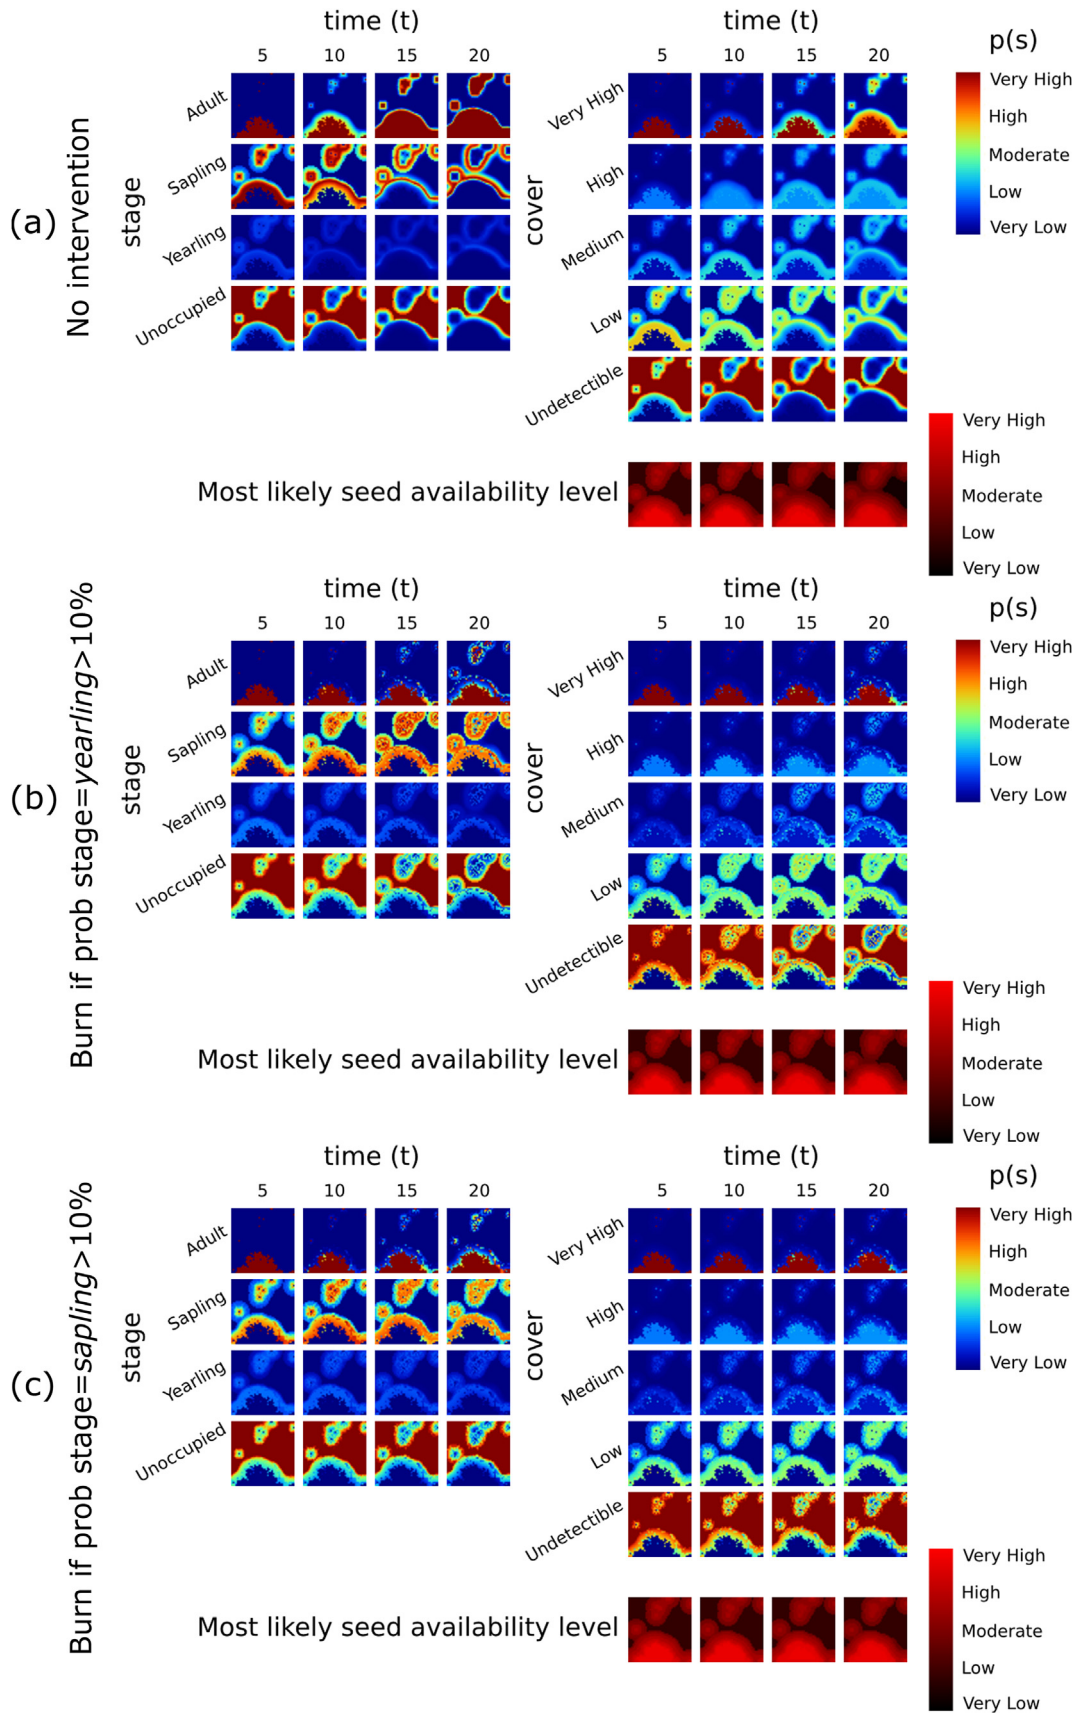


Fig. 8. Willows model results: maps of probability of willow stage and coverage, and most likely seed availability level at five-yearly time slices. The scenarios are (a) no management intervention, (b) burn management of cells when probability that stage = *yearling* > 10% and (c) burn management of cells when probability that stage = *sapling* > 10%.

time, eventually occupying just under a third of the study area (~32.5%). The 'northern' patch of saplings from $t = 5$ also matures over time turning into a separate expanding clump of adults (Fig. 8a, $t = 15$ and $t = 20$). By $t = 20$, adult willows occupy ~41% of the total area and ~20% of the total area now has high to very high seed availability.

The strategy of burning cells when the probability that stage = *yearling* is >10% is effective at containing the bottom-centre patch of adult willows and inhibiting its expansion to surrounding cells. However, by $t = 20$, the small 'northern' clump of saplings seen at $t = 5$, has most likely matured and become established as adults, though the level of cover is mostly low (Fig. 8b). The prevalence of adult willows at $t = 20$ is much lower for this strategy than that under no intervention (~20% of the total area vs ~41%). Seed availability however, remains high to very high across large portions of the study area (Fig. 8b).

Strategy (c) of burning cells when the probability that stage = *sapling* is >10% produces a similar outcome to that of burning when probability that stage = *yearling* is >10%, but is arguably slightly more effective (Fig. 8c). This seems unintuitive and one would expect strategy (b) to be superior to strategy (c) because burning is more effective at killing yearlings than saplings. However, because yearlings are a transient stage, the probability mass of stage = *yearling* is typically very low, which means the burn intervention is rarely triggered (see stage = *yearling* in Fig. 8(a)–(c)). In contrast, burn intervention is triggered more frequently with the rule probability that stage = *sapling* is >10%.

7. Discussion and conclusions

The adoption of an OO approach allowed us to design an architecture for using Bayesian networks to reason about change over both time and space. The use of abstraction and encapsulation via components with formalised input and output interfaces helped manage the complexity. We have demonstrated the generality and power of our approach through two environmental management case studies.

We note, however, that throughout the research and integration process we encountered challenges with the development, management and use of OOBNs. The tools we used to design and implement our OOBNs still lack powerful refactoring, making the management of object interface changes a time-consuming and error-prone task. Integrated version control is non-existent and documentation tools rudimentary. Improvements in these areas, which are now standard in modern software engineering IDEs, would make working with OOBNs easier and more robust.

As mentioned earlier, incorporating spatial and temporal processes can result in networks with large numbers of parents and child nodes with state space sizes so massive that they are effectively "uncompilable" because the total size of the compiled network exceeds the memory capacities of current hardware. We get around this problem by iteratively "rolling-out" a single OOSBN instance across the.

RoI, and using a "roll-out, roll-up" approach for the ST-OOSBN from one time slice to the next. The example scenarios we present provide a proof-of-concept demonstration of how our approach can be used to model spatial and temporal changes in real world case studies. An important caveat however, is that the "roll-out, roll-up" approach is an approximation – errors can (and do) accumulate and compound over space and time. This means that model estimates of weed cover in the woodlands weed study are probably over-estimated. Similarly, the progression of willow to later life stages, higher coverage and greater seed availability levels is also exaggerated, and worse than might be expected over the given timeframe.

We recognise this problem, and future work should focus on quantifying the magnitude of errors, and understanding their implications for decision making. One possible approach, is to compile spatially explicit input data at known points in time, matched with detailed knowledge of management actions that were undertaken. We could then calibrate our models for each of the case studies using these historical settings, and track how convergent or divergent our model outputs are relative to what occurred historically. Compiling and verifying the requisite data for this exercise however, is a non-trivial undertaking, and is beyond the resources and scope of the current paper. This important step remains as future work.

A potential solution for avoiding "roll-out, roll-up" approximation errors is to generate the complete model, but then use stochastic simulation to generate approximate belief distributions, instead of compiling it to perform exact belief updating. Stochastic simulation (Shachter and Peot, 1989) uses the network to generate a large number of cases from the network distribution, which are then used to estimate the posterior probabilities of the target nodes. By the Law of Large Numbers from statistics, as more cases are generated, the estimate converges on the exact probability.

As with exact inference, there is a computational complexity issue with approximate updating (Dagum and Luby, 1993). However, in practice, stochastic simulation approaches for approximate inference converge fairly quickly (Korb and Nicholson, 2010). For example, Woodberry et al. (2014) used stochastic simulation updating for an OOSBN grasslands management model that involved up to 30 nodes per time slice, run in prediction mode for 4 seasons per year for a total of 20 years. The simulation took approximately 8 hours of CPU run-time on a powerful desktop computer. This suggests that scaling up the stochastic simulation updating across a spatial grid should be computationally feasible on high-performance computing clusters (such as those now available on the Cloud). However, this remains as future work.

7.1. Conclusion

For coherent, coordinated and effective landscape-scale decision support, managers need the capability to predict state changes across *space* and *time*. We have tackled these challenges by synthesising ideas and techniques from object-oriented knowledge engineering, dynamic BNs, GIS-coupled BNs and dispersal modelling. To our knowledge, the resultant general framework, using OOBNs to model spatially-explicit process interactions, is the first of its kind.

We have demonstrated the generality and power of the framework via models for two real-world case studies. Whilst these methods are computationally demanding, this modelling is valuable for adaptive management because it allows us to examine the consequences of spatially explicit alternative management actions. Uncertainty in predicted consequences is also explicit; together these outputs can be combined with cost and utility data to evaluate and select among alternative actions using formal decision analysis.

There are challenges still to be resolved when building and using such complex spatial-temporal models, including the need for better development tools and improved approximate computation. However, we hope that our proof-of-concept demonstration, model templates and algorithms for spatial and temporal scenario simulations will pave the way for future improvements.

Acknowledgements

This work was supported by ARC Linkage Projects LP110100304 and LP110100321, the Australian Centre of Excellence for Risk

Analysis, and the ARC Centre of Excellence for Environmental Decisions. Steven R Miller, Jo Anna Emanuel, Ken Snyder, Chris Oman provided critical information about the St Johns River catchment, and knowledge about willow ecology and management. We thank Steve Sinclair for assistance with developing the woodlands weed model, and Owen Woodberry for valuable discussions about approximation errors and stochastic simulation for approximate inference.

Appendix A. Implementation tools

We used Hugin Researcher 7.7 (Hugin Expert A/S, 2013) to develop the ST-OODBs and OOSBNs, the Hugin Researcher Java API 7.7 (2013) to provide programmatic access to the developed networks, the Image-IO-ext (GeoSolutions, 2013) Java library to provide access to GIS raster layer formats, and the Java programming language to implement the algorithms tying the components together. Hugin was chosen as the OOB development platform as it currently has one of the most complete OOB implementations. Java was chosen as the implementing language as it is platform independent and provides for a well established and understood OO development environment. We implemented the tool as a standalone program allowing pre-processing of GIS data and post-processing of outputs to be performed in whatever program the end user was most familiar with. In our case we used a combination of ArcGIS (ESRI, 2013), Quantum GIS (Quantum GIS Development Team, 2013) and SAGA GIS (SAGA Development Team, 2013).

References

- Aguilera, P.A., Fernández, A., Fernández, R., Rumí, R., Salmerón, A., 2011. Bayesian networks in environmental modelling. *Environ. Model. Softw.* 26, 1376–1388.
- Ames, D.P., Neilson, B.T., Stevens, D.K., Lall, U., 2005. Using Bayesian networks to model watershed management decisions: an East Canyon creek case study. *J. Hydroinformatics* 7, 267–282.
- Barton, D.N., Kuikka, S., Varis, O., Uusitalo, L., Henriksen, H.J., Borsuk, M., de la Hera, A., Farmani, R., Johnson, S., Linnell, J.D.C., 2012. Bayesian networks in environmental and resource management. *Integr. Environ. Assess. Manag.* 8, 418–429.
- Bashari, H., Smith, C., Bosch, O., 2009. Developing decision support tools for rangeland management by combining state and transition models and Bayesian belief networks. *Agric. Syst.* 99, 23–34.
- Bestelmeyer, B.T., Brown, J.R., Havstad, K.M., Alexander, R., Chavez, G., Herrick, J.E., 2003. Development and use of state-and-transition models for rangelands. *J. Range Manag.* 56, 114–126.
- Boyer, X., Koller, D., 1998. Tractable inference for complex stochastic processes. In: *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, pp. 33–42.
- Carmona, G., Molina, J., Bromley, J., 2011. Object-Oriented Bayesian networks for participatory water management: two case studies in Spain. *J. Water Resour. Plan. Manag.* 137, 366–376.
- Chee, Y., Burgman, M., Carey, J., 2005. Use of a Bayesian Network Decision Tool to Manage Environmental Flows in the Wimmera River, Victoria. Report No. 4, LWA/MDBC Project UMO43: Delivering Sustainability through Risk Management. University of Melbourne, Australia, p. 84.
- Clark, J.D., Silman, M., Kern, R., Macklin, E., Hille Ris Lambers, J., 1999. Seed dispersal near and far: patterns across temperate and tropical forests. *Ecology* 80, 1475–1494.
- Dagum, P., Luby, M., 1993. Approximating probabilistic inference in belief networks is NP-hard. *Artif. Intell.* 60, 141–153.
- Dawsey, W.J., Minsker, B.S., Amir, E., 2007. Real time assessment of drinking water systems using a Dynamic Bayesian network. In: *World Environmental and Water Resources Congress*. [http://dx.doi.org/10.1061/40927\(243\)507](http://dx.doi.org/10.1061/40927(243)507).
- Dean, T., Kanazawa, K., 1989. A model for reasoning about persistence and causation. *Comput. Intell.* 5, 142–150.
- Duncan, D., Wintle, B., 2008. Towards adaptive management of native vegetation in regional landscapes. In: *Pettit, C., Cartwright, W., Bishop, I., Lowell, K., Pullar, D., Duncan, D. (Eds.), Landscape Analysis and Visualisation: Spatial Models for Natural Resource Management and Planning*. Springer, Berlin, pp. 159–182.
- ESRI, 2013. ArcGIS (Version 10.1) [Software]. <http://www.esri.com/software/arcgis>.
- Fox, J.C., Buckley, Y.M., Panetta, F.D., Bourgoin, J., Pullar, D., 2009. Surveillance protocols for management of invasive plants: modelling Chilean needle grass (*Nassella neesiana*) in Australia. *Divers. Distrib.* 15, 577–589.
- GeoSolutions, 2013. Image-IO-ext Java Library (Version 1.1.7) [Software]. <http://github.com/geosolutions-it/>.
- Holling, C., 1978. *Adaptive Environmental Assessment and Management*. Wiley, Chichester, UK.
- Hugin Expert A/S, 2013. Hugin Researcher (Version 7.7) [Software]. <http://www.hugin.com>.
- Jackson, R., Bartolome, J., Allen-Diaz, B., 2002. State and transition models: response to an ESA symposium. *Bull. Ecol. Soc. Am.* 83, 194–196.
- Jensen, F., Lauritzen, S., Olesen, K., 1990. Bayesian updating in causal probabilistic networks by local computations. *Comput. Stat. Q.* 4, 269–282.
- Jensen, F.V., 1996. *An Introduction to Bayesian Networks*. Springer Verlag, New York.
- Johnson, S., Mengersen, K., 2012. Integrated Bayesian network framework for modeling complex ecological issues. *Integr. Environ. Assess. Manag.* 8, 480–490.
- Kinsler, P., Lee, M.A., Dambek, G., Williams, M., Ponzio, K., Adamus, C., 1997. Expansion of Willow in the Blue Cypress Marsh Conservation Area, Upper St. Johns River Basin. Technical Report. St. Johns River Water Management District, Palatka, Florida.
- Kjærulff, U., 1992. A computational scheme for reasoning in dynamic probabilistic networks. In: *Proceedings of the Eighth Conference on Uncertainty in Artificial Intelligence*, San Mateo, CA, pp. 121–129.
- Kjærulff, U.B., Madsen, A., 2008. *Bayesian Networks and Influence Diagrams: a Guide to Construction and Analysis*. Springer Verlag, New York.
- Koller, D., Pfeffer, A., 1997. Object-oriented Bayesian networks. In: *Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, pp. 302–313.
- Korb, K.B., Nicholson, A., 2010. *Bayesian Artificial Intelligence*, second ed. Chapman&Hall/CRC, Boca Raton, FL.
- Lee, M.A.B., Ponzio, K.J., Miller, S.J., 2005a. Response of willow (*Salix caroliniana* Michx.) in a floodplain marsh to a growing season prescribed fire. *Nat. Areas J.* 25, 239–245.
- Lee, M.A.B., Synder, K.L., Valentine-Darby, P., Miller, S.J., Ponzio, K.J., 2005b. Dormant season prescribed fire as a management tool for the control of *Salix caroliniana* Michx. in a floodplain marsh. *Wetl. Ecol. Manag.* 13, 479–487.
- Molina, J., Bromley, J., García-Aróstegui, J., Sullivan, C., Benavente, J., 2010. Integrated water resources management of overexploited hydrogeological systems using Object-Oriented Bayesian Networks. *Environ. Model. Softw.* 25, 383–397.
- Neil, M., Fenton, N., Nielson, L., 2000. Building large-scale Bayesian networks. *Knowl. Eng. Rev.* 15, 1–33.
- Nichols, J., Williams, B., 2006. Monitoring for conservation. *Trends Ecol. Evol.* 21, 668–673.
- Nicholson, A., Chee, Y., Quintana-Ascencio, P., 2012. A state-transition DBN for management of willows in an american heritage river catchment. In: *Nicholson, A., Agosta, J.M., Flores, J. (Eds.), Ninth Bayesian Modeling Applications Workshop at the Conference of Uncertainty in Artificial Intelligence*, Catalina Island, CA, USA. <http://ceur-ws.org/Vol-962/paper07.pdf>.
- Nicholson, A.E., 1992. *Monitoring Discrete Environments Using Dynamic Belief Networks*. Ph.D. thesis. Department of Engineering Sciences, University of Oxford.
- Nicholson, A.E., Flores, M.J., 2011. Combining state and transition models with dynamic Bayesian networks. *Ecol. Model.* 222, 555–566.
- Nyberg, J., Marcot, B., Sulyma, R., 2006. Using Bayesian belief networks in adaptive management. *Can. J. For. Res.* 36, 3104–3116.
- Pearl, J., 1988. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, San Mateo, CA.
- Pezeshki, S.R., Anderson, P.H., Shields, F.D., 1998. Effects of soil moisture regimes on growth and survival of black willow (*Salix nigra*) posts (cuttings). *Wetlands* 18, 460–470.
- Ponzio, K.J., Miller, S.J., Underwood, E.C., Rowe, S.P., Voltolina, D.J., Miller, T.D., 2006. Response of a willow (*Salix caroliniana* Michx.) community to roller-chopping. *Nat. Areas J.* 26, 53–60.
- Quantum GIS Development Team, 2013. Quantum GIS (Version 1.8.0) [Software]. <http://www.qgis.org/>.
- Quintana-Ascencio, P., Fauth, J.E., 2010. *Ecological Studies of Willow (Salix caroliniana)*. Final Report - Year 2. Department of Biology, University of Central Florida, Orlando, Florida, p. 60.
- Quintana-Ascencio, P.F., Fauth, J.E., Castro Morales, L.M., Ponzio, K.J., Hall, D., Snyder, K., 2013. Taming the beast: managing hydrology to control Carolina Willow (*Salix caroliniana*) seedlings and cuttings. *Restor. Ecol.* 21, 639–647.
- Rumpff, L., Duncan, D., Vesk, P., Keith, D., Wintle, B., 2011. State-and-transition modelling for adaptive management of native woodlands. *Biol. Conserv.* 144, 1224–1236.
- Rumpff, L., Duncan, D., Vesk, P., Wintle, B., 2012. Development of a state-and-transition model to guide investment in woodland vegetation condition. In: *Lefroy, T., Curtis, A., Jakeman, A.J., McKee, J. (Eds.), Landscape Logic: Pattern, People and Process in Landscape Management*. CSIRO Publishing, Melbourne, pp. 157–172.
- Sadler, R.J., Hazelton, M., Boer, M.M., Grierson, P.F., 2010. Deriving state-and-transition models from an image series of grassland pattern dynamics. *Ecol. Model.* 221, 433–444.
- SAGA Development Team, 2013. SAGA GIS (Version 2.0.8) [Software]. <http://www.saga-gis.org/>.
- Shachter, R., Peot, M., 1989. Simulation approaches to general probabilistic inference on belief networks. In: *Proceedings of the Fifth Workshop on Uncertainty in Artificial Intelligence*, pp. 311–318.
- Shihab, K., 2008. Dynamic modeling of groundwater pollutants with Bayesian networks. *Appl. Artif. Intell.* 22, 352–376.

- Shihab, K., Chalabi, N., 2007. Dynamic modeling of ground-water quality using Bayesian techniques. *J. Am. Water Resour. Assoc.* 43, 664–674.
- Uusitalo, L., 2007. Advantages and challenges of Bayesian networks in environmental modelling. *Ecol. Model.* 203, 312–318.
- Walters, C., Hilborn, R., 1978. Ecological optimization and adaptive management. *Annu. Rev. Ecol. Syst.* 9, 157–188.
- Westoby, M., Walker, B., Meir, I., 1989. Opportunistic management for rangelands not at equilibrium. *J. Range Manag.* 42, 266–274.
- Wilkinson, L., Chee, Y., Nicholson, A., Quintana-Ascencio, P., 2013. An object-oriented spatial and temporal Bayesian network for managing willows in an American heritage river catchment. In: Almond, R., Mengshoel, O. (Eds.), *Proceedings of the 2013 UAI Application Workshop: Models for Spatial, Temporal and Network Data at the Conference of Uncertainty in Artificial Intelligence*, 15 July, Bellevue, Washington, USA, CEUR-WS.org, pp. 77–86. <http://ceur-ws.org/Vol-1024/paper-10.pdf>.
- Woodberry, O., Millett-Riley, J., Sinclair, S., Nicholson, A., 2014. An object-oriented dynamic Bayesian decision network model for grasslands adaptive management (abstract). In: *Proceedings of the Eleventh UAI Bayesian Modeling Application Workshop*, 27 July, Quebec, Canada, CEUR-WS.org, p. 101. http://ceur-ws.org/Vol-1218/bmaw2014_abstract_2.pdf.
- Yates, C.J., Hobbs, R.J., 1997a. Temperate eucalypt woodlands: a review of their status, processes threatening their persistence and techniques for restoration. *Aust. J. Bot.* 45, 949–973.
- Yates, C.J., Hobbs, R.J., 1997b. Woodland restoration in the western Australian wheatbelt: a conceptual framework using a state and transition model. *Restor. Ecol.* 5, 28–35.