

Deriving state-and-transition models from an image series of grassland pattern dynamics

ABSTRACT

We present how state-and-transition models (STMs) may be derived from image data, providing a graphical means of understanding how ecological dynamics are driven by complex interactions among ecosystem events. A temporal sequence of imagery of fine scale vegetation patterning was acquired from close range photogrammetry (CRP) of 1m quadrats, in a long term monitoring project of *Themeda triandra* (Forsskal) grasslands in north western Australia. A principal components scaling of image metrics calculated on the imagery defined the state space of the STM, and thereby characterised the different patterns found in the imagery. Using the state space, we were able to relate key events (i.e. fire and rainfall) to both the image data and aboveground biomass, and identified distinct ecological 'phases' and 'transitions' of the system. The methodology objectively constructs a STM from imagery and, in principle, may be applied to any temporal sequence of imagery captured in any event-driven system. Our approach, by integrating image data, addresses the labour constraint limiting the extensive use of STMs in managing vegetation change in arid and semiarid rangelands.

Keywords: Pattern dynamics, Close range photogrammetry, *Themeda triandra* grasslands, Adaptive management, Vegetation monitoring, Image metrics, Pilbara

1. Introduction

The state-and-transition model (STM) is a conceptual tool to organize our understanding of how the dynamics (or temporal change) of vegetation communities are driven by complex interactions among events (e.g. fire, grazing and flooding), processes (e.g. mineralisation and drainage) and biological factors (e.g. invasive species). STMs were originally developed to explain limitations in Clementsian linear succession models in predicting the consequences of land management on rangeland vegetation dynamics, including: irreversible vegetation change; grazing catastrophe; episodic plant recruitment; and, alternative stable vegetation states (Westoby et al., 1989). As a management tool, the STM illustrates the potential impacts, nonlinearities and uncertainties inherent under different environmental conditions or management scenarios (Bestelmeyer et al., 2004). Consequently, STMs are a key component of the proposed framework for science-based land management (SBLM) of rangelands in the western United States and elsewhere (Herrick et al., 2006). However, a science-based, quantitative implementation that goes beyond current qualitative methods of deriving STMs is constrained by: (i) a lack of long term data encompassing the breadth of different possible ecosystem behaviours; (ii) the need to derive STMs indirectly by collating a range of often disparate data sources; and, (iii) system and geographical specificity, thereby creating an inequity in terms of which systems and localities will be investigated first (Bestelmeyer et al., 2003, 2004). Data in the form of aerial imagery would ameliorate the above difficulties in STM modelling, as imagery is spatially extensive and may be reliably captured, georeferenced and processed at a range of geographical scales. A methodology to arrive directly at a STM from temporal sequences of imagery that capture vegetation pattern dynamics is therefore proposed.

In its simplest form the STM identifies the possible vegetation communities that can occur in a particular system or locale, labelling them as different meta-stable phases (i.e. these vegetation communities may vary from being transient to persistent; Westoby et al., 1989; Stringham et al., 2003). The system may then experience a transition (or pathway) from one phase to another that was triggered (or driven) by the interacting effect of different events and processes. Hence, some of the nonlinearities observed in the field that confound linear succession models may be described, a list including: ecological thresholds and irreversible transitions; multiple stable states; historical contingency; cross-scale interactions, and lagged effects (Friedel, 1991; Rietkerk and van de Koppel, 1997; Peters et al., 2004). The dynamics of a

system are therefore an observed sequence of nonlinear transitions between different phases (Fig. 1).

A key concept of the STM is the ecological threshold, a boundary in space and time between different phases of the system (or domains of stable system behaviour; May, 1977; Friedel, 1991). Transitions across ecological thresholds can occur along a continuum from irreversible to immediately reversible pathways (Stringham et al., 2003; Briske et al., 2005). Irreversible transitions are triggered by events that are typically unprecedented and which lead to a degraded state that is constrained by a different set of biological and soil processes (e.g. soil erosion, overgrazing, introduction of exotic plant species, altered fire regimes or a combination of such events). The typology of Stringham et al. (2003) uses irreversible transitions to define 'states' of the system, whereas reversible transitions separate 'phases' or 'communities'. Collections of phases, connected by a network of transitions are visualized as the phase dynamics nested within each system state (Fig. 1; Stringham et al., 2003).

Science-based land management (SBLM) is defined by four core elements: (1) a method of land classification to describe 'ecological' sites; (2) a data storage and management facility; (3) conceptual models of ecosystem dynamics, including site specific STMs; and (4) a methodology to evaluate the status of the ecosystem, be it qualitative or quantitative (Herrick et al., 2006). To date, SBLM has been applied to arid lands of western United States, where episodic events such as rainfall drive often dramatic shifts in ecological processes. In arid lands, intensive assessment is necessary to match and capture the frequency of driving events, otherwise important factors leading to vegetation change will likely be missed and long term trends will be confounded by temporary responses to recent events. Consequently, the need for long term monitoring is implicit in SBLM, and is akin to adaptive management (Holling, 1978): data from long term monitoring is used to update models and knowledge that assists management decision making in dynamic and uncertain ecological systems.

The main constraint to SBLM is the ability to capture data. Currently, SBLM defines a protocol of 17 qualitative indicators for the evaluation of three key ecosystem attributes: soil and site stability, hydrologic function and biotic integrity (e.g. Pellant et al., 2005). If more precise information is required then a further protocol composed of quantitative indicators may be applied (Herrick et al., 2005). Both protocols are similar to landscape functional analysis (LFA) in the type of indicators they employ (Ludwig et al., 1997; Tongway and Hindley, 2005). However

acquiring data by both LFA and SBLM is labour intensive, resulting in limited spatial sampling (i.e. a sampling bias, Watso and Novelly, 2004), and risking significant observer bias (different observers may provide different assessments, and the same observer may provide different assessments under different conditions, Burrough and McDonnell, 1998; Hunt et al., 2003; Boot et al., 2006). The limitations inherent in labour-intensive, qualitative assessments is exacerbated by the need for repeated long term monitoring in capturing ecological dynamics, and is likely to limit the application of STMs beyond intensively studied sites.

Ecosystem assessments that serve multiple management goals, such as that promoted by SBLM, employ a large number of indicators, as individual indicators represent only singular aspects of ecosystem behaviour. For example, a number of metrics that quantify pattern in imagery may be applied, be they structural (e.g. patch based) or textural (i.e. pixel based). These image metrics frequently describe similar aspects of image pattern, such as average patch radius and the perimeter–area ratio (Riitters et al., 1995). Dimension reduction techniques, such as principal components, can be used to simplify the multivariate information of a large number of candidate metrics into a small number of summary variables (Riitters et al., 1995). In effect, the summary variables define a ‘state space’ (or blank canvas) on to which individual images can be plotted.

An ordination of image metrics calculated on time sequences of imagery will generate a trajectory of a system’s dynamics in the ordination defined state space. Time trajectories of a system’s behaviour have previously been applied to studying changes in plant community composition, with ordination plots used to elucidate the difference in dynamics between good and poor condition shrubby grasslands of arid Australia (Friedel, 1991). Nonmetric multidimensional scaling has also described the regeneration of grassland at polluted sites in comparison to semi-natural meadows in East Germany (Voigt and Perner, 2004). A partitioning of the state space into various ecological phases can then be linked to ecological events that drive transitions between the ecological phases, thus deriving a STM.

The flexibility of imagery in terms of scale, geographical extent and relative ease of capture makes imagery an ideal data source, with potential to alleviate both the sampling and observer bias of manual methods, and to reduce costs associated with long term monitoring. However, it remains uncertain how to best incorporate imagery within existing assessment

frameworks, despite remote sensed imagery being viewed as an important facet of ecosystem monitoring in spatially extensive systems (Ludwig et al., 2004; Herrick et al., 2006). Our challenge is to derive models of ecosystem behaviour directly from imagery sourced data, a step critical in furthering the utility of both imagery and the STM in the SBLM framework (Bestelmeyer et al., 2004). Consequently, our objectives were to: (i) generate a state space by applying image metrics to an image series; (ii) partition the state space to represent different 'phases' of the system's dynamics; (iii) associate sets of system events to the different phases, and thus define the ecological triggers for transitions between phases; and, (iv) construct a state-and-transition model by integrating the above information. As discussed below our approach extends that of (Jackson and Bartolome, 2002), who were first to realize the feasibility of applying data driven methods in developing STMs. Further, our approach uses only off-the-shelf statistical tools, readily available to both managers and applied scientists, in fulfilling our stated objectives.




Fig. 1. State-and-transition model.

The state space contains ecological states (or vegetation communities) separated by transitions across irreversible ecological thresholds. Each state may be described by its own characteristic set of phase dynamics. Transitions between these phases are reversible, either directly or indirectly. Transitions, whether reversible or irreversible, are triggered by interacting ecological events.

Modified from Stringham et al. (2003).




Fig. 2. Monthly and annual rainfall at Hamersley Station (October 1996–October 2005). (a) The monthly rainfall record shows a seasonal summer rain pattern, and dates of fires and image capture; (b) variability in monthly rainfall is illustrated by the boxplots, with median monthly rainfall given by a thick horizontal bar. Average monthly maximum temperature and average monthly minimum temperature are given by the solid and dashed red lines, respectively. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

2. Method

The methodology for deriving a state-and-transition model from an image series was developed on a grazing exclosure experiment that focuses on the long-term biomass and nutrient cycling of *Themeda triandra* (kangaroo grass) tussock (bunch) grasslands in the sub-tropical and semi-arid Pilbara region of north western Australia. The data here represent a case study of how the methodology may be applied to other systems and image series.

2.1. Study sites

The study sites were located on Hamersley Station, a cattle property ~1700km north of Perth, where *T. triandra* grasslands are confined to deep self-mulching red Vertisol soils on very gently sloping alluvial plains. Soils are extremely low in organic matter (less than 1%) and nutrients (Bentley et al., 1999). The *T. triandra* grasslands studied are monospecific swards, with less than 3% annuals and forbs by biomass (Bennett et al., 2002). On drier slopes adjacent to the plains, vegetation is generally a mix of open mulga (*Acacia anuera* complex) woodland and spinifex (*Triodia* spp.) hummock grasslands.

The dynamics of tussock grassland biomass in the Pilbara region are largely driven by episodic events such as flood and fire. In particular, growth of *T. triandra* is highly responsive to seasonal deluges associated with summer cyclonic activity, and senesces under drought conditions that may last for several years. Mean annual rainfall is 350mm y^{-1} , but varies from 50 to 800mm y^{-1} (Fig. 2a). Fire generally consumes all standing biomass and several good growing seasons (i.e. high rainfall) are required before pre-fire biomass is regained. Consequently, fuel accumulation after fire is inextricably linked to the unpredictable timing and extent of drought and large rainfall events. The fire return interval for *T. triandra* grasslands in the Pilbara is in the order of 4–10 years, depending on the occurrence of rainfall, compared to 1–3 years for *T. triandra* in temperate and tropical regions elsewhere (Bennett et al., 2002 Lunt and Morgan, 2002).

2.2. Data acquisition

We acquired data from two fenced grazing exclosures located 12km apart ('Ridge' and 'Cattlewell' paddocks). Exclosures were established in 1995 to assess the productivity and dynamics of *T. triandra* grasslands in the absence of cattle grazing. Changes in cover and biomass in response to climate, fire and nutrient additions have been measured at these sites since October 1996 (Bennett et al., 2002).

Imagery captured through close range photogrammetry (CRP) have spatial resolutions as fine as 1mm, and includes nadir photography captured by on-ground frame mounted cameras (e.g. Cooper, 1924; Bennett et al., 2000; Laliberte et al., 2007), or more recently digital sensors mounted on ultralight airplanes (Hunt et al., 2003; Booth et al., 2006). CRP methods have small fields of view, typically ranging from 1 to 100 m, and do not provide continuous photographic coverage over extensive landscapes. Instead, images are sampled intermittently across the landscape (Booth et al., 2006). Combined with image processing, CRP in rangelands is at least as accurate as non-image assessment methods in estimating vegetation variables such as percentage foliage cover (Bennett et al., 2000; Booth et al., 2006). Here, we use CRP to regularly capture the fast system dynamics of biomass turnover in *T. triandra* grasslands. As biomass will likely be highly correlated with vegetation structure or "texture", we propose a STM mapping of the "texture" dynamics. In principle, the same methodology may be applied to image series captured at different scales for event driven vegetation systems exhibiting slower dynamical cycles and/or shifts.

At each site a sub-plot measuring $1\text{m} \times 1\text{m}$ central to a larger $5\text{m} \times 5\text{m}$ plot was photographed from above using a camera mounted on a 2m high tripod (Bennett et al., 2000). Initially, from October 1996 to October 2002, a film SLR camera was used, but was replaced with a digital SLR from March 2003 to October 2005. Sites were sampled twice a year, generally corresponding to before (October) and after (near April) the summer growing season, although at times major rainfall events that resulted in extensive sheet flooding of the sites limited accessibility to the sites. Consequently, in 2000 and 2001 plots were sampled only once. Half of the plots were experimentally burnt and photographed in November 1996. The complete data set therefore consists of four replications of two time series (experimentally burnt and unburnt) of

18–19 images at each of two sites, giving a total of 295 images (with one image treated as a missing value).

Covariates assigned to each image were: site ('Ridge', 'Cattlewell'); biomass; time since last fire (months); and rainfall (mm). Total aboveground biomass was harvested from half of the replicated plots at any one sampling date, from a sub-plot on the perimeter of each 5m × 5m plot to minimise disturbance to the central unharvested photo/monitoring sub-plot. The sub-plot sampled for biomass therefore differed at each sampling date. Total aboveground biomass was oven dried (75 ° C for 48 h) and weighed (Bennett et al., 2000, 2002).

Wildfire, ignited by lightning strikes, was documented for each of the two sites (Fig. 2a). However, an arbitrary value of 100 months for time since last fire was assigned to the first sampling date in October 1996, as fire had not been previously observed by the station managers at either site for at least a decade. Rainfall data for the period were acquired from Hamersley Station, with temperature data from the nearest town (Tom Price, 35kmdistant). Annual rainfall was defined as two variables, summed over the early wet season (November–January), and middle wet season (February–March). The early season and middle season rainfall variables correspond respectively to rising and peak ambient temperatures and evapotranspiration demand over the summer season (Fig. 2b).

2.3. Image processing and metrics

A complete sequence of digital RGB (red–green–blue) images was assembled by scanning photographs acquired by the SLR camera during the first part of the monitoring program at high resolution (2000 dpi). Images were processed in Adobe® Photoshop® 7.0 (ADOBE®, 2002) and were firstly color adjusted using a black and white color tile placed in each image. The color tile permitted standardisation of color between sampling periods, as the sub-plots were shaded with a purpose built shade cloth when photographed. Georeferencing was enabled by placing a 1m × 1m metal sub-plot frame over marker pegs, and the camera tripod inserted onto a pole, both permanently situated in the ground at the start of the monitoring program. Images were warped and clipped to a 1000 × 1000 pixel template to compensate for any errors in camera angle. To increase separation between foliage and red soil, the blue

brightness of the green color channel was increased by 100%. Images were then reduced to 2×250 pixels, the coarsest resolution that permitted detection of fine scale vegetation features by the naked eye. To further simplify computation and data management images were thresholded (i.e. transformed) at a brightness value of 80 to produce black and white (or binary) images. Data were then imported into an R software environment (version 2.7; R Development Core Team, 2008), coupled to the Geographical Information System GRASS (version 6.0.4 GRASS Development Team, 2006; Neteler and Mitasova, 2004; Bivand et al., 2008).

The final processing step applied a smoothing filter to remove noise from the images. All patches of nine pixels size or smaller were removed to produce an image series (Fig. 3). Table 1 summarizes the image metrics calculated for each image in the GRASS-R computing environment. Our choice of metrics depended on qualitative review of their mathematical properties (such as scale, rotation and translation invariance; see for example Frohn, 1998).




Fig. 3. Image series of Cattlewell plot. The image series are of a single $1\text{m} \times 1\text{m}$ control treatment plot at the Cattlewell site (i.e. not burnt in November 1996). Two main types of pattern seem to occur (above vs. below rows), the separation of which coincides with the wildfire of April 2001.

Table 1

Image metrics

Metric	Description	Reference
White Fraction	The proportion of the image occupied by the ‘white’ fraction of a binary black and white image.	
Number of Patches	Counts the number of discrete “white” patches.	See Frohn (1998) for a discussion.
Twist Number	A measure of overall shape complexity. Calculated by counting the number of twists in the perimeter of a shape.	Bogaert et al. (1999)
Effective Mesh Size	Associated with the probability of two locations belonging to the same patch, and determines the size of uniform patches that corresponds to that probability.	Jaeger (2000)
Box Counting Slope	Complex systems may display some form of multi-fractality, i.e. differently scaled regions may display different space filling properties. Slopes of the box counting plot are therefore used as separate metrics using box dimensions of 1, 2, 5, 10, 25, and 50 pixels in length.	Li (2000)
Contrast	Measures local variability of pixel values.	Haralick et al. (1973)
Contagion	Based on the relative frequency of finding a pixel of one type next to a pixel of another type.	O’Neill et al. (1988); Li and Reynolds (1993)
Recursivity	Measures uniformity of pixel pair combinations. Designed to be independent of contrast and scaled to between 0 and 1.	Baraldi and Parmiggiani (1995)
Compactness	Counts the number of internal pixel edges contained in a shape, and is expressed as a fraction of the maximum number of internal edges possible for a shape of the same number of pixels.	Bribiesca (1997)

2.4. Statistical methodologies

A two-dimensional (2D) state space was defined by the first two axes produced from a principal component analysis of the image metrics, with data points corresponding to individual images. A key concern of the analysis is how to classify the state space into different ecological phases. A classification could be elicited through a combination of expert opinion, on-ground evidence and knowledge of site history. However, such a classification should be compared to an unsupervised classification that makes few assumptions a priori regarding the possible structure of the STM. Here we apply model-based clustering to the image metric scores, grouping the four replicates for each treatment \times site \times capture date combination into a single multivariate observation through concatenation. This grouping of observations assumes that all replicates are situated within the same ecological phase given they have the same history of fire and rainfall. The model-based clustering identifies clusters of observations based on estimating finite mixture models that vary in both their location and spread, while choosing an optimal number of clusters through applying a Bayesian Information Criterion (BIC; implemented in the ‘Mclust’ package in R; Fraley and Raftery, 2002). The original data were then classified into one of the inferred clusters by their treatment \times site \times capture date information, with each identified cluster representing a hypothesized ecological phase. In addition to the initial principal components ordination and model-based clustering, our approach included a further three statistical procedures: (i) kernel density estimation; (ii) classification trees; (iii) penalized spline geoaddivitive regression.

A ‘porous’ phase boundary for each cluster may be drawn in the principal component state space through kernel density estimation. Kernel density estimation superimposes a density curve (or surface), known as the kernel, over each data point and sums the resulting collection of density curves, and may be interpreted simply as a smoothed histogram (Bowman and Azzalini, 1997). phase boundary can therefore be represented by choosing a specific contour (e.g. the contour given by the 85th percentile defines all points on the density surface of equal height that contains 85% of the volume under the density surface). A Gaussian kernel was employed, with the smoothing parameter controlling the kernel’s dispersion chosen automatically as the asymptotic normal smoothing parameter estimate (Wand and Jones, 1995). Kernel density estimation was implemented using the ‘sm’ package of Bowman and Azzalini (1997) in R.

A classification tree analysis identified which covariates, such as time since fire and rainfall, were associated with transitions between different hypothesized phases. Some hypothesized phases may be different in terms of structural patterning as captured by the imagery, but are not ecologically different in terms of what drivers or covariates explain them. The operation of classification trees is relatively straightforward, choosing splits in a single variable through some optimality criterion (in this case a deviance measure based on the conditional likelihood) before proceeding to further splits (Breiman et al., 1984). A split occurs at a 'node', and terminal nodes are known as 'leaves'. Classification trees have been used frequently in the ecological literature (e.g. De'ath and Fabricius, 2000; Jackson and Bartolome, 2002) and were implemented using the 'rpart' package in R (Venables and Ripley, 2002), with confusion matrix statistics estimated with the 'caret' package.

A geoaddivitive model was used to understand how biomass varied across the state space after taking into account such variables as time since fire and rainfall, and thereby give some ecological meaning to the state space (in a manner similar to surface fitting in ordinated spaces; Dixon, 2003). Geoaddivitive models combine both semiparametric regression (i.e. including linear and smooth functions such as splines of regressor variables) and universal kriging (i.e. covariate regression and autocorrelated error) within a shared mixed model framework (Kammann and Wand, 2003), facilitating model fitting and selection. Smoothing parameters for fitting trend lines were selected automatically using generalised crossvalidation (Ruppert et al., 2003), whilst the number of knots used to define the splines was constrained to five knots for univariate regressors and 20 knots for bivariate regressors to highlight any main trends in the data. Standard F-statistics were used in choosing an optimal model through combined backward and forward selection of both smooth and linear functions of covariates. Although there is no explicit form for the distribution of these statistics within semiparametric regression, some guidance is provided by comparison to a corresponding F distribution in the absence of computationally intensive simulation (Hastie and Tibshirani, 1990). Thus a test with significance levels set to 0.05 in this context provides only some evidence (as opposed to strong evidence) of a trend or effect in response to explanatory variables. Implementation of geoaddivitive models used the 'SemiPar' package in R (Ruppert et al., 2003).




Fig. 4. Density estimation of hypothesized phases. The data were allocated to one of five hypothesized phases identified through model-based clustering of the entire set of principal components. The contour for each hypothesized phase represents the 85th percentile of a 2D kernel density estimate of points assigned to the hypothesized phase. Points are plotted on the first two principal component axes.

Table 2

Classification of phases by events

Node 1	Node 2	Node 3	Class (leaf)	Positive prediction rate ^a	Number of observations
Rain over last 2 months > 106mm	High November–January rainfall > 215mm		5	75%	16
	Low November–January rainfall < 106mm		4	71%	59
Rain over last 2 months < 106mm	High November–January rainfall > 62mm	Time since fire \leq 6.5 months	3	25%	16
		Rain over last 6 months < 36mm & time since last fire between 20 and 85 months	5	75%	16
		Otherwise	1 and 2	100%	104
	Low November–January rainfall < 62mm	Time since fire > 5 years	1 and 2	83%	24
		Rain in last 12 months < 99mm and time since fire < 5 years	4	83%	12
		Otherwise	5	75%	48

^a The positive prediction rate is the difference between the total number of predicted positives and number of falsely predicted positives, as a proportion of total number of positives for each leaf of the classification tree.

3. Results

3.1. Defining the state space and ecological phases

A state space was constructed using the first two principal components of the sphered image metric data (where sphering divides the data by its correlation matrix), explaining 84.5% of the variation in the image metric data (66% principal component 1; 18.5% principal component 2). Each point in the state space corresponds to an observed image with observations concentrated to the right hand side of the state space (Fig. 4a). A user defined mask covered all points, illustrating that not all parts of the state space contain observations.

Five phases (clusters) were hypothesized by the model-based algorithm: phases 1 and 2 shared the same location, but with phase 2 more widely dispersed across the state space than phase 1 (Fig. 4b and c). In order of proximity to phases 1 and 2 the phases were defined as phase 5, phase 4 and phase 3. Both phase 5 and phase 4 were more broadly dispersed than phases 1 and 2, whereas phase 3 was the least frequently occurring phase containing just eight of the 295 observations in total. Note that at the 85th percentile there is significant overlap between the different hypothesized phases in the state space.

3.2. Interpretation of transition pathways

A partition of the state space using a classification tree shows the intensities and types of event combinations most associated with the hypothesized phases (Table 2). For example, phase 4 was associated with two differing sets of conditions defined by two separate branches of the classification tree: (i) under dry November–January conditions (combined seasonal rainfall < 215 mm), but when rainfall in the previous 2 months was greater than 106mm; or, (ii) under more extreme drought conditions (previous yearly rainfall < 99 mm), but within 5 years of the most recent fire. Phase 1 and Phase 2 were largely congruous, but with 36% of Phase 2 observations distinguished from Phase 1 observations by a time since fire of greater than 7 years. Phases 1 and 2 were then combined to lift the overall classification accuracy rate up from 73% to 81% (a statistically significant increase in accuracy at the $\alpha = 0.05$ significance level), suggesting that in terms of the ecological drivers examined here the two phases differed little. Together, phases 1 and 2 were predicted well by low recent rainfall (< 106mm) but with a high combined rainfall in the most recent November–January period (> 62 mm), or when time since fire was greater than 8

years. Phase 3 was the most poorly predicted (a sensitivity rate of 50% in Table 3), but was associated with recent fire within the previous 6 months. Significantly, the previous phase of the system was not found to be an important predictor of the current system phase.

Exploratory data analysis (EDA) may be employed to support or contradict the model derived from the model-based clustering and tree classification. The event histories at the two sites in the timing of fire in relation to a drought over 2001 and 2002. The fire at Cattlewell in April 2001 preceded the drought and led to phases 4 and 5 type behaviour (Fig. 5). The fire at Ridge in December 2002 was at the tail of the drought: images captured during the drought are tightly clustered within the phases 1 and 2 region (Fig. 6). Pronounced behaviour such as this response to the interaction of fire and drought is encapsulated in the classification tree (Table 2). More transitory behaviour, such as a sheet flow event in early 2004 that was observed to strip the *T. triandra* grasslands of much of its senesced biomass, in part explains why phase 4 and phase 5 were associated with both drought and extreme wet events, both being biomass removing ecological events. Here is an example of similar structural configurations, as captured by the imagery, resulting from different ecological stimuli. Further, where events or hypothesized phases were infrequently observed then the model-based clustering and classification tree performs less reliably, as when the November 1996 images for Ridge being classed as phase 4 whereas ideally they would be classed as phase 3 as they were captured within 2 months of an experimental burn. The low positive prediction rate for phase 3 (Table 3) arises from the allocation of June 1998 images from Cattlewell into phase 3 through the location/spread form of the mixture models estimated in the model-based clustering. This result supports an assertion that distances between images in the state space should not be automatically assumed to represent 'ecological distance', and effort should be invested in EDA to query the underlying drivers of state space movements and transitions.

Not all state space mappings will have the advantage of these data in having available auxiliary state variables such as biomass to verify or ground-truth the dynamics captured by the image series. Biomass data for the *T. triandra* grassland, harvested at the same time as when the imagery were captured, ranged from 0.2 to 10 t ha⁻¹. In general, high scores on the first principal component represented a 2 t ha⁻¹ higher biomass than low scores, once site, time since fire and rainfall (summed over the previous 6 months) were taken into account in a geoaddivitive semi

parametric (spline) regression of sub-plot biomass (Fig. 7). The single regressors illustrated the following trends: (i) biomass first increased with rainfall but showed a relatively constant response thereafter (Fig. 7a); (ii) biomass increased linearly with time since last fire (Fig. 7b); and, (iii) the ‘Cattlewell’ site had 0.45 t ha^{-1} higher biomass than ‘Ridge’, after taking into account the other regressors ($p\text{-value} = 0.023$). The inclusion of a bivariate smoothing interaction between fire and rainfall did not lead to a strong improvement in model fit. Thus phases 1 and 2 may be interpreted as a high biomass phase, and phase 3 as the low biomass phase in association with recent fire.

3.3. Mapping the state-and-transition model

A STM was mapped onto the 2D state space derived from the principal component ordination of image metrics applied to the CRP imagery (Fig. 8). The STM was composed of the hypothesized ecological phases that partition the state space. The overlapping of phase boundaries was permitted as there was a priori no reason to assume different phases should be entirely discrete (for these data the 60th percentile of the 2D kernel density estimate of the data points in each ecological phase was considered an acceptable heuristic for a relatively discrete representation of phases in Fig. 8). The ecological phases were simply connected as a graph by transitions defined by specific sets of ecological events deduced from the tree classifier.

In summary, driving events associated with transitions of the system from one phase to another are learned by: (i) hypothesizing phases either through an unsupervised classifier such as model-based clustering, through expert opinion, or through an EDA (such as breaking down the system’s dynamics in the image state space into individual sampling dates and searching for similar clustering behaviour); (ii) by associating threshold intensities of different ecological events to each of the hypothesized phases through a supervised classifier such as a classification tree; (iii) assigning the state space some ecological meaning through querying the performance of the classifier through EDA techniques, or if available the geospatial regression of auxiliary state variables over the state space. Ecological meaning could further be validated by applying the classifier to test image series captured at novel locations.

Table 3

Confusion matrix for predicting hypothesized phases

Predicted	Hypothesized phase				
	Phase 1 and 2	Phase 3	Phase 4	Phase 5	Specificity ^a
Phases 1 and 2	124	0	0	4	97.4%
Phase 3	4	4	4	4	95.8%
Phase 4	0	4	52	15	91.9%
Phase 5	16	0	4	60	90.6%
Sensitivity	86.1%	50.0%	86.7%	72.3%	n = 295




Fig. 5. Cattlewell vegetation pattern dynamics, 1996–2005. Each plot maps the images defined as points on the first two principal components for individual sampling dates at the Cattlewell site. Control (○) and treatment plots experimentally burnt in November 1996 (△) are plotted, with the size of each observation dependent on time since last fire. A large plotting symbol indicates a recent fire, with a wildfire occurring in April 2001. Rainfall (mm) summed over the 3 months previous to the sampling date is represented by the height of the bar at the bottom of each image. Dark shading of the state space corresponds to high biomass whilst light shading corresponds to low biomass, after factors such as site, time since fire and rainfall have been taken into account in a geoaddivitive regression.




Fig. 6. Ridge vegetation pattern dynamics, 1996–2005. Each plot maps the images defined as points on the first two principal components for individual sampling dates at the Ridge site. Control (\circ) and treatment plots experimentally burnt in November 1996 (\triangle) are plotted, with the size of each observation dependent on time since last fire. A large plotting symbol indicates a recent fire, with a wildfire occurring in December 2002. Rainfall (mm) summed over the 3 months previous to the sampling date is represented by the height of the bar at the bottom of each image. Dark shading of the state space corresponds to high biomass whilst light shading corresponds to low biomass, after factors such as site, time since fire and rainfall have been taken into account in a geoaddivitive regression.




Fig. 7. Semi-parametric regression of biomass. Total biomass (gm^{-2} ; square-root transformed for variance stabilization) was regressed on: (a) the natural logarithm of the sum of the previous 6months rainfall; (b) the natural logarithm of time since last fire; and (c) the two PCA axes. Grey regions in plots (a) and (b) indicate ± 2 standard errors. The PCA component shows high biomass in red and low biomass in blue. Total estimated biomass is the additive sum of the component smooth regressions. The regressions used either 5 knots (rainfall and fire) or 20 knots (bivariate PCA). The plotting mask in (c) was user defined to cover all observed points.

4. Discussion

We have demonstrated the direct construction of a state-and-transition model from a time series of image data. Thus far, state-and-transition models have been used to cogently summarize the existing state of knowledge of an ecosystem's complex dynamics, to identify the transitions that are most important for management, to direct research towards conditions that lead to those transitions, or as qualitative tools to evaluate the relative benefits and risks of different management actions (Westoby et al., 1989; Stringham et al., 2003). High spatial and temporal variability in driving event processes were previously thought to preclude a quantitative approach (Bestelmeyer et al., 2004). Where quantitative transition rules have been assigned to a STM, as in the modelling of long term carbon dynamics in Australian savannas, parameter values for these rules were sourced from other studies (Hill et al., 2005). Our approach builds on that of Jackson and Bartolome (2002) by employing an unsupervised classifier for hypothesizing different system phases and states, and a supervised classifier to associate ecological covariates and events to the hypothesized phases. The methods proposed here differ from Jackson and Bartolome (2002), however, in a number of ways:

- The classifiers are applied to patterning in imagery (as characterised by a set of image metrics) as opposed to community composition data.
- The phases/states are mapped directly onto a state space defined by an ordination of the data through kernel density percentiles, rather than deduced by the expert from the classification. The plotting of dynamic trajectories onto an ordinated state space follows Friedel (1991) and Voigt and Perner (2004).
- The model-based clustering provides an objective method of choosing the number of hypothesized phases/clusters, though exploring alternative numbers of clusters should be an important component of the EDA. The supervised classifier should 'whittle' down the number of hypothesized classes when hypothesized classes may be explained by similar event histories (as with phases 1 and 2 in the *T. triandra* data set).
- We employed the tree classifier to predict the occurrence of system phases and not observed transitions (as in Jackson and Bartolome (2002)). In the *T. triandra* system, previous system phase was not found to be competitive with other ecological covariates in predicting the current system phase. While the previous system phase is likely to be an important predictor

in other ecological systems, the evidence presented here suggests that the dynamics of *T. triandra* grasslands are highly event dependent: processes such as fire occurrence and drought mask any possible relationship between the current and previous system states (supported by the high classification rate > 80% overall). Our formulation of including previous system phase as a predictor also allows for the possible identification of events that sustain the persistence of a system over time in a given phase.

Our study identified four phases in the fine scaled vegetation pattern dynamics of *T. triandra* grasslands in the Pilbara: a 'high' biomass ecological phase (combined phases 1 and 2) associated either with greater early season rain in November–January or when there is an extended absence of fire from the system (> 7 years); a transitory 'low' biomass phase occurring immediately after infrequent fire (phase 3); and two intermediate biomass phases spanning the remainder of the state space (4 and 5). There was no evidence of irreversible transitions in the observed dynamics of the system, and so what is reported here represents 'phase' dynamics within the *T. triandra* grassland state as opposed to transitions between system states.

Identification of multiple phases can be interpreted as a response to a highly variable and unpredictable environment of fires and droughts. Transitions between meta-stable phases in vegetation patterning brought about by shifts in standing plant biomass and plant numbers are considered to confer resilience in arid ecosystems to drought and grazing (van de Koppel and Rietkerk, 2004), and concurs with previous observations of northern *T. triandra* grasslands as being amongst the most resilient grazing systems in Australia (Tongway and Ludwig, 1994). The efficacy by which image metrics of fine-scale grassland patterning were able to distinguish distinct ecological phases also lends weight to the close coupling of vegetation pattern and ecological process in arid systems (Noy-Meir, 1973; Ludwig et al., 1997). The role of pattern dynamics as a useful surrogate for the underlying dynamics of ecological processes is to an extent supported. Two key caveats in utilizing pattern dynamics as a surrogate for process dynamics remain: (i) similar vegetation pattern configurations as captured in the imagery may result from different ecological stimuli (as in phases 4 and 5 that may occur after sheet flow or interacting drought/fire events); and (ii) distance between points in the state space is not to be confused with 'ecological' distance, in that infrequently observed phases (e.g. phase 3) may have both poor sensitivity and specificity in their prediction, and are thereby not well identified. Both these issues result from data constraints in the number and length of image series that can be

generated through more extensive monitoring in both time and space. Some of the misclassification may result from the difficulty of monitoring appropriate or sufficiently many covariates. The practical consequence of these data constraints is that there is still a need for rigorous querying by the ecological expert of outputs from the clustering and classification methods, and of how they are to be interpreted.

The derivation of STMs from image time series complements the adaptive management paradigm that underpins the SBLM framework for arid and semiarid ecosystems. For example, the hypothesis that an early season rainfall of greater than 62mm allows a ‘high’ biomass phase to persist can be refined and updated by reapplying the STM methodology when further monitoring data are acquired. Thus, the STM methodology provides a facility to learn the dynamics of a largely uncertain system – to date not all intensities, timing and interactions between ecological events possible under the current event regime have been observed, with their consequent impacts on vegetation dynamics largely unknown (Holling and Allen, 2002). However, caution is needed in interpreting model output as exact values (such as 62mm) are an automatic product of the classification tree methodology. Further analyses, utilizing random forests (Breiman, 2004), may be used to provide a measure of variability about such estimates. Threshold intensities in events associated with transitions between phases should therefore be viewed as provisional, but valid, starting points in predicting ecological phase, and in providing an increasingly precise understanding of system behaviour over time.

A quantitative basis to STM construction also provides new possibilities for understanding and thus managing vegetation dynamics, in addition to detecting potential deviations from the domain of normally observed dynamics, including: (i) comprising an ‘early warning’ system that detects potential deviations away from the domain of normally observed dynamics, observed as a trajectory moving away from commonly occurring phases in the state space; (ii) classifying landscapes and landscape change according to ecological phases; (iii) determining how different ecological phases underpin other processes such as the likelihood of fire; and (iv) prediction of future transitions, although this last application will be dependent on the ability to predict the occurrence and intensity of ecosystem events.

Overlapping boundaries of the polygons defining different ecological phases in the constructed STM (Fig. 8) may be compared to higher order phase transitions in physical systems. In physical systems, a system parameter may be increased so one phase becomes dominant over

another (i.e. that phase becomes the meta-stable phase) but that dominance is not complete, or not realized immediately. For example, shrub dominance is in a constant state of flux when there is no sustained overgrazing in semi-arid wooded grasslands in eastern Australia, with changes in shrub dominance driven by fire interacting with drought events (Westoby et al., 1989). The high variability in timing and intensity of fire and rainfall ensures that shrub or grass dominance in wooded grasslands is never complete at landscape scales. An awareness of the lack of complete dominance (or lack of discreteness of phase boundaries in the state space) of any single phase should be maintained when interpreting STMs generated from image series, and multivariate data in general.

Our approach of deriving STMs of ecosystem dynamics from CRP imagery has the potential to reduce the labour required in extensive long term monitoring by ground based methods. However, non-trivial issues in advancing the implementation of the methodology include calculating how much ground-truthing is required to attribute some form of ecological meaning to phases of the state space (exemplified in this study by the linkage of the state space to labour intensive biomass data), and the initial cost of developing image processing protocols. Without auxiliary state variables interpretation will have to rely on judgements inferred from the EDA and knowledge of event histories at different sites. An auxiliary state variable or other interpretation is not strictly necessary for purposes of detecting novel system dynamics, but is likely essential in communicating the event driven STM behaviour from an abstract state space to on-ground managers.

When developing protocols for deriving STMs from image series, four criticisms of applying image metrics should be answered: (i) the image processing procedure has the potential to alter what processes are actually being monitored, and the ecological interpretation of metrics applied to the imagery; (ii) the same metrics may have potentially different meanings at different scales and in different ecosystems, therefore representing possibly unrelated processes (Li and Wu, 2004); (iii) image metrics are subject to criticism similar to that of on-ground ecological indicators in that not any one metric may fully represent a process; and (iv) image metrics also possess mathematical properties that may impact on the utility of a metric (e.g. 'energy' based measures will in general have a higher sensitivity to deviations from landscape uniformity than from landscape heterogeneity in comparison to other textural measures; Baraldi and Parmiggiani, 1995). In consequence, several image metrics are typically used to represent a process, analogous

to what occurs for on-ground monitoring in SBLM and Landscape Functional Analysis. However, while on-ground indicators are often designed for measuring specific ecological processes (e.g. soil electrical conductivity for salinity) image metrics measure only pattern, with little a priori knowledge of which image metrics will usefully capture ecological processes (Riitters et al., 1995; Fortin et al., 2003; Li and Wu, 2004). For example, there is not an automatic correlation between textural metrics of image pattern and the scale at which plants may utilize resources in a landscape. A further task in constructing STM state spaces from image data is therefore the selection of the subset of image metrics that will be appropriate to a specific application. However, once established image processing would be at least semi-automatic for the repeated capture of images and of benefit to any long term monitoring program.

The need of land managers to distinguish alternative ecological phases at broader scales (i.e. paddock to landscape) than the 1m^2 quadrats captured here by CRP can in part be addressed by multiple image capture over a study site. In this study, for example, eight quadrats were photographed at each of the two sites at each sampling date. Contours of kernel density estimates defining boundaries of the hypothesized ecological phases were then one way of representing the mean field behaviour of the *T. triandra* grasslands over the paddock scale sites where the imagery were captured. However, land managers in other contexts will likely employ imagery of different resolution at larger spatial scales. Differently scaled vegetation patterns will therefore be detected (e.g. a grass-woodland matrix), dominated by differently scaled ecological processes (e.g. landscape flows of water and plant invasion). This diversity of applications may be furthered by managers choosing other metrics, processing methods and driver variables that are directed more towards management needs, rather than the approach presented here which aims at detecting differences in image pattern with as few a priori assumptions as possible (e.g. managers may be more interested in shifts in basal area of the grasslands for assessing stocking rates, as opposed to overall changes in patterning of above ground biomass). However, the deployment of alternative imagery and metrics will likely affect the generation of STMs, potentially resulting in different, hypothesized ecological phases correlated to somewhat different system events. This specificity of the derived STM in relation to the image source will constrain the capacity of managers to “mix-and-match” STMs generated from different image sources and at varying spatial scales for the one vegetation system. Despite the limitations of

moving beyond a specific spatial scale and imagery source, the salient point here is that STMs can be generated from an image series. The methodology may therefore be applied more generally to larger scale imagery in future.

4.1. Alternative applications

Our study of *T. triandra* grasslands of north western Australia represents a relatively simple system: a monoculture that responds rapidly to driving events, with experimental exclusion of grazing. Where grazing is a factor, changes in grazing intensity may be incorporated as a driver of a system's dynamics by including grazing intensity as a further covariate in the tree classification, assuming the data are available. In addition to semi-arid and arid systems, the methodology may be applied to the pattern dynamics of other event-driven, multi-phase systems such as seagrass-macroalga communities (Fourqurean and Rutten, 2004), where strong changes in patterning may be captured by aerial imagery over relatively small time scales. More complex vegetation systems, such as coniferous forests whose event-driven dynamics evolve over decadal scales, will require monitoring of multiple sites in a space-for-time substitution, with the divergence in event histories between sites to be maximised (Pickett, 1989). For example, 'Ridge' and 'Cattlewell' were sufficiently distant spatially to be subject to separate wildfires, permitting the interaction between fire and drought to be better identified in the *T. triandra* image series (Figs. 5 and 6). An image derived dynamics has then a potentially wide application to a range of vegetation systems.

More generally, our approach consists simply of defining a state space, hypothesizing phases through an unsupervised classifier, and then assigning thresholds in driving events to those phases using a supervised classifier. As demonstrated by (Jackson and Bartolome, 2002), the methodology may be extended to other multivariate, long-term monitoring data and not just to pattern metrics derived from image data. For example, restoration projects using plant community data in comparing phase dynamics between disturbed and benchmark sites could be used to evaluate project completion targets (Grant, 2006). Remote sensing technologies quantifying aspects of the system elements other than vegetation pattern (e.g. biomass and net primary productivity; 'greenness' indices; water yield and heat fluxes) would simply extend the number of indicators to be incorporated into the state space, or be treated as auxiliary state variables. Existing rangeland monitoring data comprising SBLM type indices to detect changes in ecosystem processes may also be integrated within this framework.




Fig. 8. State-and-transition model of *T. triandra* grassland dynamics. Boundaries of ecological phases are represented by the 60th percentile of 2D kernel density estimates of all observations placed into each hypothesized phase, and were collapsed onto the one plot. Note that the different phases overlap but in this plot have been discretised to an extent by the choice of percentile. Together the phases and transitions form a phase dynamics that characterise an ecological state of the system (i.e. *T. triandra* grassland).