

University of Warwick institutional repository: <http://go.warwick.ac.uk/wrap>

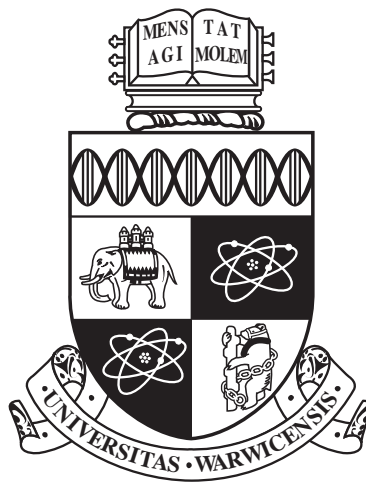
A Thesis Submitted for the Degree of PhD at the University of Warwick

<http://go.warwick.ac.uk/wrap/67166>

This thesis is made available online and is protected by original copyright.

Please scroll down to view the document itself.

Please refer to the repository record for this item for information to help you to cite it. Our policy information is available from the repository home page.



Pattern Formation and Persistence in Spatial Plant Ecology

by

Michael Alastair Irvine

Thesis

Submitted to the University of Warwick

for the degree of Mathematical Ecology and Complexity

Science

Doctor of Philosophy

Complexity Science

October 2014

THE UNIVERSITY OF
WARWICK

Contents

List of Tables	vi
List of Figures	vii
Acknowledgments	x
Declarations	xii
Abstract	xiii
Chapter 1 Introduction	1
1.1 Motivation	2
1.2 Related efforts	3
1.3 Summary of main development of thesis	6
1.4 Outline of remainder of thesis	7
Chapter 2 Background	15
2.1 Introduction	16
2.2 Pattern formation	17
2.2.1 Semi-arid ecosystems	20
2.3 Critical phenomena	22
2.4 Scaling	24
2.5 Conclusion	30
Chapter 3 Seagrass	38
3.1 Introduction	39
3.2 Biology & ecology	39
3.3 Isles of Scilly	42
3.4 Mapping	43
3.5 Survey	45

3.5.1	Banding survey	47
3.6	Modelling	49
3.7	Conclusion	50
Chapter 4	Spatial analysis	57
4.1	Introduction	58
4.2	Monofractal	58
4.2.1	Measuring monofractality from data	61
4.3	Fractional Brownian motion	62
4.4	Lacunarity	65
4.5	Multifractality	65
4.5.1	The multifractal spectrum: density against occupancy data	69
4.5.2	Pattern anisotropy	70
4.6	Results	73
4.6.1	Box-counting dimension	73
4.6.2	Lacunarity	74
4.6.3	Anisotropy	75
4.6.4	Hölder exponent	77
4.7	Discussion	78
Chapter 5	Modelling	86
5.1	Introduction	87
5.2	Reaction-diffusion equations	88
5.2.1	A Short deviation into the literature	95
5.2.2	Conclusion	96
5.3	Integro-differential equations	97
5.3.1	Derivation	98
5.3.2	Kernels	100
5.3.3	Anisotropy	101
5.4	Probabilistic cellular automata	101
5.5	Simulation	103
5.6	Results	103
5.6.1	Geometry	106
5.7	Conclusion	106
Chapter 6	Fractal heuristics of return rate	110
6.1	Introduction	111
6.2	Theoretical background	113

6.2.1	Fractal growth and boundaries	114
6.2.2	Relative patchiness	120
6.3	Models	121
6.3.1	Introduction	121
6.3.2	Two-state model	121
6.3.3	Three-state model	122
6.4	Method	125
6.4.1	Calculation of fractal dimensions	125
6.4.2	Calculation of the return rate	127
6.4.3	Comparison to data	133
6.4.4	Simulations	134
6.5	Results	135
6.5.1	Disturbance & recovery in a homogeneous environment . . .	135
6.5.2	Invading cluster with environmental noise	137
6.5.3	Properties of a vegetation boundary along an environmental gradient	138
6.6	Conclusion	140

Chapter 7 Parameter estimation from a snapshot of regular vegetation spatial patterns 154

7.1	Introduction	155
7.2	Density as an observable	156
7.2.1	Method	156
7.2.2	Minimising for one parameter	157
7.2.3	Minimising for two parameters	158
7.2.4	Searching over parameter space	159
7.3	Spatial correlations as an observable	160
7.3.1	Derivation of the variational correlation for spatial pattern inference	161
7.3.2	Non-spatial birth death process	166
7.4	Inference	169
7.4.1	Simulated annealing	169
7.4.2	Approximate Bayesian computation	173
7.4.3	Synthetic likelihood	174
7.4.4	Derivation of the likelihood	175
7.4.5	Priors	178
7.4.6	Fitting to real data	182

7.5	Results	182
7.5.1	Simulated data	182
7.5.2	Real data	185
7.6	Conclusion	185
Chapter 8 Aggregation		195
8.1	Introduction	196
8.2	Introduction to equations	197
8.3	An ecological interpretation of the aggregation kernel	199
8.4	An analogue for the return rate	200
8.5	Constant aggregation with linear death	202
8.6	Comparison to explicit spatial modelling of vegetation growth	204
8.7	Empirical analysis of patch-size distribution using the aggregation- fragmentation equation	206
8.7.1	Developing the likelihood	207
8.7.2	Results	207
8.8	Conclusion	208
Chapter 9 Vegetation disease on spatial pattern formation		214
9.1	Introduction	215
9.1.1	Modelling the spread of a species in a disordered environment	216
9.1.2	The Multiplicative cascade model	217
9.1.3	Methods	221
9.1.4	Observables	221
9.1.5	Results	222
9.1.6	Discussion	225
9.2	Disease diffusion in a changing environment	225
9.2.1	Disease model	226
9.2.2	Competition in regulating disease spread	227
9.3	Evolutionary model for pattern formation in reaction to disease . . .	230
9.3.1	Evolutionary model for competition that is transmitted . . .	232
9.4	Results	236
9.4.1	Altruistic competition	236
9.4.2	Selfish competition	239
9.5	Conclusion	239

Chapter 10 Conclusion	245
10.1 Future work	249
10.2 Conclusion	251

List of Tables

4.1	Density, box-counting dimension and range for surveyed sites.	73
7.1	Measurement of correlation distributions for various system sizes. . .	178
8.1	Bayes factor comparison for all three models over the five sites. . . .	208
8.2	Inferred parameters and 95% confidence intervals for all three models and five sites.	211

List of Figures

2.1	Reaction-Diffusion mechanism	16
2.2	Examples of regular pattern formation	17
2.3	Semi-arid patch-size distribution	26
3.1	Global seagrass distribution	40
3.2	Old Grimsby Harbour (2013)	42
3.3	Isles of Scilly survey sites	44
3.4	Seagrass surveying methods	46
3.5	Time-series of quadrat-based survey	47
3.6	Seagrass banding sites	48
3.7	grey seal	51
4.1	Simulating fractional Brownian motion	62
4.2	Clipping in box-counting	66
4.3	Lacunarity of survey sites	74
4.4	Comparison of lacunarity to null model	76
4.5	Anisotropy analysis	77
4.6	Anisotropy analysis	80
4.7	Boundary dimension of survey sites	81
4.8	local scaling of survey sites	82
4.9	correlations between density and local scaling	83
5.1	Hydrological scouring	87
5.2	spectral analysis of offset parameter	104
5.3	spectral analysis of ambient birth parameter	105
5.4	impact of geometry on banding	106
6.1	Eden model	115
6.2	Example realisation of three-state vegetation model	128

6.3	Calculation of return rate for the birth-death process	131
6.5	Growth rate and Hurst exponent in three-state model.	137
6.9	Korcak - return rate relationship for surveyed and simulated data. .	140
6.4	Comparison between dynamic heuristics and average time to equilib- rium in two-state model.	143
6.6	Comparison of boundary dimension and return rate for three-state model.	144
6.7	Comparison of boundary and Korcak dimension for three-state model.	145
6.8	Comparison of Korcak dimension and return rate for three-state model.	146
6.10	Varying environmental parameters in two-state model.	147
6.11	Varying dynamic parameters in two-state model.	148
7.1	Expected change in density varying offset.	157
7.2	Expected change in density varying two parameters.	158
7.3	Expected change in density varying three parameters.	159
7.4	Change in expected correlation for birth-death process.	168
7.5	Change in correlation for simulated data.	170
7.6	Calculating ζ^2 over a two parameter subset of parameter space. . . .	171
7.7	Simulated annealing for four parameters	172
7.8	Scaling of expected number of XY pairs.	177
7.9	Distribution of errors for increasing system size.	179
7.10	Example posterior for simulated data.	186
7.11	Marginals of posterior for simulated data	187
7.12	Change in correlation structure for changing competition k	188
7.13	Posterior for seagrass data.	191
7.14	Posterior for tiger bush data.	192
8.1	Plausible Korcak-return rate relationship.	201
8.2	Exponents of patch-size distribution compared to simulations.	206
9.1	Equilibrium distribution for Poisson point process.	217
9.2	Summary of diffusion for Poisson point process.	222
9.3	Summary of growth rate for multiplicative cascade model.	224
9.4	Schematic diagram for disease model.	226
9.5	Mean and variance for population in disease model.	227
9.6	Example of dynamics of disease model with no banding.	229
9.7	Example of dynamics of disease model with strong banding.	229

9.8	Realisations of altruistic competition model with no mutation and disease.	232
9.9	Effect of infectivity on competition in selfish and altruistic model. .	237
9.10	Realisations of altruistic competition model.	238

Acknowledgments

I am firstly grateful to my two supervisors Matt Keeling and Jim Bull for their constant support throughout the whole process, keeping me pointing in the right direction and for interesting discussions throughout. I would like to thank the Engineering and Physical Sciences Research Council (EPSRC) for funding me and allowing me the opportunity to work on this project. I am grateful to everyone involved with the Complexity Science Doctoral Training Centre (DTC) for the support, training and generally being my home for the past four years. The people of the Warwick Infectious Disease and Epidemiology Research (WIDER) group, for their ideas and feedback. I am also thankful to Emma Jackson and Natural England for providing the seagrass mapping data. Also to the Isles of Scilly *Zostera* group for their data and allowing me the opportunity to visit the islands; conduct field work and swim with grey seals. I'd like to thank my thesis proof-readers, Jim Bull, Matt Keeling, Liz Buckingham-Jeffrey, Peter Dawson, Tom Rafferty and Merve Alanyali for getting through my sometimes obfuscated writing.

Finally, I'd like to thank my family for their unwavering support, encouragement and reminding me to keep my feet on the ground. To my awesome partner Jenny Lo, whose stuck with me through all the peaks and troughs during the PhD. Also to all the friends and colleagues I have made during my time at Warwick. Including, but not limited to: Ben Collyer, Daniel Sprague, Mike Maitland, Peter Dawson, Jenny Webb, Jonathan Skipp, Arran Tamsett, Quentin Caudron, Sergio Morales, Marcus Ong, Martine Barons, Chris Oates, Jamie Harris, Anthony Woolcock and Adam Miller, whose wit and appreciation of tea kept me going through the long

months when results were not forthcoming.

Declarations

This thesis is submitted to the University of Warwick in support of my application for the degree of Doctor of Philosophy. It has been composed by myself and has not been submitted in any previous application for any degree.

The survey work discussed in chapter 3 was carried out as part of an on-going 18 year study into the health of *zostera* beds in the Isles of Scilly with the volunteer group IOS *zostera*. I personally took part in the study in the summer of 2013, where I was involved with the collection of data and assistance in the survey work. I also devised and participated in the study of regular spatial banding of the seagrass beds mentioned. The calculation of the return rate on the Isles of Scilly *Zostera* data conducted in chapter 6 was performed by James C. Bull. Every effort has been made to credit the authors of the peer-reviewed literature which lay the foundations of this work.

Abstract

The main aim of the thesis is to explore the interaction between pattern and process in vegetation ecology using a variety of mathematical and statistical methods. Of particular interest is what information about the dynamics of the underlying system can be gained through a single spatial snapshot, such as an aerial photograph or satellite image. The hypotheses are related to seagrass ecology, whose growth is primarily clonal and broadly exists as a monoculture and thus makes it an ideal candidate to study these interactions. The thesis broadly concerns two forms of spatial pattern and the underlying dynamics that give rise to them. The first concerns regular pattern formation, where the pattern has a characteristic length scale. Examples are abundant in natural systems, such as mussel beds, semi-arid ecosystems as well as seagrass. The developments concerned with regular pattern formation include methods of detection in a large spatial dataset, a novel stochastic model of vegetation that produces regular pattern with plausible mechanisms, the development of a new methodology to fit regular spatial pattern data to the model and the impact as well as evolutionary mechanisms of regular patterning in the presence of disease.

The second form of spatial pattern exhibited in a wide variety of sessile species is scale-free or fractal patterning. Certain scaling heuristics, such as the boundary dimension of a vegetation cluster or the power-law exponent of the patch-size distribution have been used to infer properties of the dynamics. We explore these heuristics using a variety of plausible models of vegetation growth and find the circumstances under which there is a clear relationship between the spatial heuristics and the dynamics. These are then supplemented by viewing vegetation growth as an aggregation process. A novel model of vegetation aggregation with death is produced to find the origin of the ubiquitous power-law patch-size distribution found in nature. Finally the impact of scaling on the spread of disease is explored.

Chapter 1

Introduction

[W]e must find ways to quantify patterns of variability in space and time, to understand how patterns change with scale. . . , and to understand the causes and consequences of pattern. . .

(Simon A. Levin)

1.1 Motivation

Traditional survey techniques of a marine vegetation systems involve longitudinal studies with local sampling techniques such as quadrat sampling [Duarte and Kirkman, 2001; Pringle, 1984; Shears and Babcock, 2002]. These techniques rely on the assumption that the relevant scale of reference is of the order of the size of the quadrat (normally 1m^2) and will often miss processes occurring at larger scales, such as the scale of whole ecosystems. The time-scale over which longitudinal, quadrat sampling occurs is also an issue. For a perennial plant species that has a turn-over rate greater than 2 years a study could typically last decades in order to collect sufficient data to answer questions on the dynamics of the ecosystem [Bull et al., 2012].

Often the assumption of spatial homogeneity is used to simplify mathematical analysis. Vegetation ecosystems are rarely completely homogeneous, often forming patchy landscapes with non-smooth boundaries [Adler et al., 2001; Wiens, 1989]. Processes such as dispersal of organisms, disease spread, persistence and colonisation depend on the form of these spatial patterns that the vegetation exhibits [Bertness and Callaway, 1994; Jolles et al., 2002; Nathan and Muller-Landau, 2000; Tilman, 1994]. For the management of ecosystems, it is therefore important to have techniques that can describe vegetation spatial patterns as well as techniques to understand their formation and their underlying dynamics. In order to probe the relationship between pattern and process, ecosystems where a single vegetation species dominates can be explored. One such species is seagrass [Hemminga and Duarte, 2000].

Seagrass exists largely as a monoculture exhibiting both clonal and sexual reproduction in the form of shooting and seeding respectively [Kendrick et al., 2012]. It exhibits long and short-range feedback processes in the form of sedimentation; retardation of currents; crowding and anchoring from the growth of clonal mats [Abdelrhman, 2003; Larkum et al., 2006]. It forms a number of different spatial patterns at a range of scales including banding and can display fractal meadow boundaries as well as scale-free patch-size distributions [Cunha and Santos, 2009; Van Der Heide et al., 2010]. As such it is an ideal species candidate to test the relationship between spatial pattern and dynamics.

Seagrass provides a number of ecosystem services that make it of significant economic and ecological relevance [Costanza et al., 1998]. It is involved in coastline

stabilisation, by raising the seabed relative to the baseline level [Wright and Jones, 2006]. This breaks the wave before the coastline, thus reducing erosion from wave action. It can engineer its own environment, meaning it's able to reduce currents significantly whilst also increasing deposition of sand giving protection and more access to light. The rhizome structure is an important carbon sink and recycles organic nitrogen [Duarte et al., 2005; Touchette and Burkholder, 2000]. The meadows themselves provide an ecosystem for a large number of organisms and biodiversity is typically high with species of many different phyla including molluscs, fish and waterfowl [Beck et al., 2001]. Due to the economic and ecological importance and their rapid loss in recent years it is of pressing concern to study and understand the extent of seagrass losses, their ability to persist as meadows and their general efficacy as ecosystem service providers.

In recent years there has been a huge increase in the amount of spatial data obtained from remote sensing techniques available for ecological analysis [Kerr and Ostrovsky, 2003]. Data in the form of satellite imagery and aerial photographs have been used to identify species habitat, spatial extent and variability and how cofactors such as anthropogenic influences can change the resulting distribution [Buermann et al., 2008; Hansen et al., 2010; Lefsky et al., 2002]. This relatively new form of data in ecology requires new methodologies that will allow it to supplement more traditional data sources such as quadrat-based surveys. In particular the main hypothesis is that information on the dynamics of a vegetation system can be gained through analysing a single spatial snapshot. The gain in knowledge from the spatial pattern alone would represent a significant improvement over current surveying techniques.

1.2 Related efforts

Objects in nature such as boundaries between water and land (coastlines, lakes rivers etc.), vegetation boundaries and clouds are difficult to describe using classical geometry [Mandelbrot, 1983]. This is due to their apparent lack of scale when viewed at increasing levels of resolution. The classic question raised by Mandelbrot, "How long is a coastline?" highlights the issue [Mandelbrot, 1967]. If a ten meter ruler were used to measure a section of coastline a different answer may be obtained than if a one meter ruler were used to measure it. This is due to undulations in the coastline that exist on all scales, while some may be too small for the ten meter ruler to measure, the one meter ruler could measure them, thus the coastline appears to

grow in length the smaller you make the measurements. The development of this is to observe how the length of the object being studied grows as the length of the measurements decreases. This produces a power law, with an exponent that defines the *Fractal Dimension* of an object [Falconer, 2013].

Fractal Theory has had a number of applications in vegetation ecology [Sugihara and M May, 1990]. Questions arising from studying patchy landscapes such as the size of habitat that can sustain a species population and the length of interfaces between two forms of vegetation have a natural setting in fractal theory. Often processes can dominate over a finite range of scales leading to processes that dominate for a different range of scales. This is why the term scaling theory has become more popular in recent years, circumventing some of the strict formalism laid out by Fractal theory [Halley et al., 2004]. The local scaling dimension of vegetation has been shown to affect species distributions, but relatively little work has been done on discerning the underlying mechanisms of what gives rise to these species distributions [Green et al., 2003; Wilson et al., 2004].

How the dynamics of a vegetative process gives rise to its shape has been discussed and originally proposed by Sugihara and M May [1990]. Vegetation is viewed broadly as a diffusion process with both dynamic and spatial scaling that are inter-related. This is quantified using a modified Brownian diffusion process with a tunable parameter representing the persistence or anti-persistence of the process [Mandelbrot, 1979]. This parameter, known as the *Hurst Exponent* has then been used to characterise different vegetative spatial patterns and used as a heuristic for dynamics such as succession [Hastings et al., 1982]. Being a phenomenological model however, it still remains unclear where the mechanistic link is between the dynamics and the resulting spatial pattern.

Other spatial statistics that have been explored include the size and shape of patches [Seuront, 2009]. These can take a number of forms including estimating the patch perimeter, its area and its radius of gyration [Li, 2000]. Of particular importance is the area of a patch. The distribution of which has been well studied and under certain conditions can have a power law tail [Kéfi et al., 2011]. Power laws in patch-size distributions were originally studied by Korcak, where he observed them in size distributions of Aegean islands. Since then there has been a lot of research into using these distributions to define an exponent known as the Korcak exponent after its discoverer.

The Korcak exponent has been applied to several ecological processes including: correlating with grazing pressure on a landscape [Xin et al., 1999]; providing an estimate of patchiness and re-forestation [Imre et al., 2011]; relating to the cover between two species [Erlandsson et al., 2011] and providing a measure of persistence [Pascual et al., 2002]. There has been tantalising glimpses that varying Korcak exponents can give different vegetative dynamics and hence the Korcak exponent can be used as a measure of the vegetative dynamics and in particular its persistence [Cunha and Santos, 2009]. However, up to now there has been no research to demonstrate this relationship between dynamics and the Korcak exponent in situ.

A patch-size distribution displaying scale-free behaviour intuitively has fractal properties. As the distribution is a power law there is no relevant scale at which to observe the patches and hence there is a self-similar property over certain length scales. This has lead to discussion over how the patch-size scale free property can be related to the self-affine properties of the boundary of the patches. Various authors have looked at this relationship, but as yet there is no clear picture as to how they do relate if at all in a vegetative process [Hastings and Sugihara, 1993; Imre et al., 2012; Sugihara and M May, 1990].

The origin of such vegetation distributions has also gained much interest recently. Pattern formation concerns the emergence of regular patterns from homogeneous starting conditions and has broadly been studied using reaction-diffusion type systems to characterise the necessary symmetry-breaking bifurcations needed for pattern formation to occur. In sessile ecology there are numerous examples of pattern formation in semi-arid ecosystems [Lejeune and Tlidi, 1999], mussel beds [van de Koppel et al., 2005], coral reefs [Mistr and Bercovici, 2003] etc.

As a vegetative system evolves, certain parameters have a stronger influence over the nature of the dynamics than others. In semi-arid ecosystems, parameters such as rainfall determine whether the system is driven to a steady-state of high vegetation density or a steady-state of low vegetation density [Schwinning et al., 2004]. As a system approaches this point through variation of its parameters it undergoes a tipping point or bifurcation, where there is a qualitative change in the system dynamics. Often vegetative systems will display hysteresis, with the barren state persisting even if rainfall is increased [Von Hardenberg et al., 2001].

As a dynamic system approaches a bifurcation a number of predictable changes occur that can be seen in an observable of the system such as density or variance

[Scheffer et al., 2009]. These measures can then be used as an indicator of whether a system is close to a catastrophe or not and hence can inform ecosystem managers to prevent a complete collapse of the system [Guttal and Jayaprakash, 2009]. These measures are applicable when the system is close to a bifurcation point, however if a system is far from equilibrium or not close to a tipping point then it is not clear what spatial indicators are suitable in order to establish the persistence of the ecosystem.

Indicators of bifurcation in vegetative spatial patterns have recently received much attention [Dakos et al., 2010, 2011; Kéfi et al., 2014]. If the underlying mechanisms of a vegetative process are well understood then the system undergoes a predictable series of changes to the spatial pattern as it moves from the vegetative state to the barren state [Rietkerk et al., 2004]. These spatial indicators could then provide a way of rapidly assessing how close an ecosystem is to a bifurcation. It remains unclear, however, whether spatial indicators can be detected in a real system and what the nature of them would be for a system with many scales of interaction.

The related efforts naturally lead to a number of hypotheses that shall be explored in this thesis. The primary of which is what the relationship is between temporal and spatial persistence. This can be broken down into a number of further hypotheses. Specifically, can the Korcak exponent, boundary dimension and other fractal measures of a spatial vegetation distribution be used as an effective heuristic in measuring the dynamic persistence in a real or simulated ecological system. What is the theoretical origin of the power-law scaling of the patch-size distribution and how does this exponent relate to underlying physical and ecological processes. Does the Korcak exponent change continuously with parameters of a system or is there an abrupt transition. Further hypotheses that shall be explored shall be whether a mechanistic model can both reproduce regular pattern formation and scale-free pattern formation. Is it possible to fit parameters to a single spatial snapshot using this model. Finally, how does the heterogeneity and pattern of a vegetation landscape impact other processes that may be occurring on said landscape, such as species dispersion or disease spread.

1.3 Summary of main development of thesis

The main developments of this thesis have been the following:

A key finding has been the exploration of the relationship between persistence in the form of the return rate and spatial pattern in the form of the Korcak exponent. These results are a combination of analysis of seagrass datasets as well as from numerical simulation.

Current modelling of regular spatial pattern in vegetation is done using continuous deterministic models such as reaction diffusion. These models provide a good qualitative description of how certain underlying mechanisms lead to pattern formation, however are difficult to apply directly to data. A novel model of regular pattern formation in vegetation is explored where the vegetation is modelled as individual units occupying a lattice. This allows direct comparison between the simulated spatial pattern and data. A novel technique of performing Bayesian model inference on vegetation spatial pattern data has also been constructed and explored in this thesis. This technique provides a way of taking a single snapshot of a spatial pattern and inferring a plausible model of vegetation growth with competition to it.

The spatial work has been supplemented by viewing vegetation as an aggregation process. This leads to a theoretical origin for the power law tail of the patch-size distribution as well as an explanation for the onset of an exponential tail observed in multiple spatial ecological systems such as mussel beds and semi-arid ecosystems.

Finally the role of pattern on process was explored by analysing how fractal and regular pattern properties of a vegetation affect the spread of disease through the vegetation. Heterogeneity of the distribution was found to slow down the progression of the disease. While regular patterns in the form of banding was also found to regulate the progression and persistence of the disease and vegetation. The hypothesis explored was then whether regular patterns could have evolved in the presence of disease. It was found that under a wide-range of conditions, in particular when the disease has a high level of virulence, then the regular pattern trait dominates.

1.4 Outline of remainder of thesis

The structure of the thesis is as follows. The next chapter covers the background of related efforts in the field of vegetation spatial pattern analysis. In particular, the origin of certain types of pattern commonly found in vegetation ecosystems as well as efforts related to interpreting these spatial patterns in relation to the extinction risk of the community and other dynamical indicators. The next chapter

discusses seagrass, which is used as a particular vegetation community where ideas developed in the thesis can be explored. Then spatial analysis is conducted on the seagrass dataset to ascertain the types of spatial pattern present in the data that will then be interpreted in coming chapters. Particular forms of pattern are discussed and various models that are both based on reasonable assumptions about the underlying biology in the next chapter. Following on, specific ways of relating the spatial pattern of a vegetation to its persistence are discussed including a number of modelling approaches and specifically applying to the seagrass data in order to ascertain its validity. The next chapter deals with model fitting to a spatial pattern, when the spatial pattern has a regular structure. A way of interpreting the patch-size distribution using aggregation is then explored. Finally, the thesis ends with an exploration of disease on a vegetation pattern, including how the fractal properties of the spatial pattern impact the spread of disease as well as how regular banding can regulate disease spread.

Bibliography

- Mohamed A Abdelrhman. Effect of eelgrass *zostera marina* canopies on flow and transport. *Marine Ecology Progress Series*, 248:67–83, 2003.
- P Adler, D Raff, and W Lauenroth. The effect of grazing on the spatial heterogeneity of vegetation. *Oecologia*, 128(4):465–479, 2001.
- Michael W Beck, Kenneth L Heck Jr, Kenneth W Able, Daniel L Childers, David B Eggleston, Bronwyn M Gillanders, Benjamin Halpern, Cynthia G Hays, Kaho Hoshino, Thomas J Minello, et al. The identification, conservation, and management of estuarine and marine nurseries for fish and invertebrates: A better understanding of the habitats that serve as nurseries for marine species and the factors that create site-specific variability in nursery quality will improve conservation and management of these areas. *Bioscience*, 51(8):633–641, 2001.
- Mark D Bertness and Ragan Callaway. Positive interactions in communities. *Trends in Ecology & Evolution*, 9(5):191–193, 1994.
- Wolfgang Buermann, Sassan Saatchi, Thomas B Smith, Brian R Zutta, Jaime A Chaves, Borja Milá, and Catherine H Graham. Predicting species distributions across the amazonian and andean regions using remote sensing data. *Journal of Biogeography*, 35(7):1160–1176, 2008.
- James C Bull, Emma J Kenyon, and Kevan J Cook. Wasting disease regulates long-term population dynamics in a threatened seagrass. *Oecologia*, 169(1):135–142, 2012.
- Robert Costanza, Ralph d’Arge, Rudolf de Groot, Stephen Farber, Monica Grasso, Bruce Hannon, Karin Limburg, Shahid Naeem, Robert V O’Neill, Jose Paruelo, et al. The value of the world’s ecosystem services and natural capital. *nature*, (387):253–260, 1998.

- AH Cunha and RP Santos. The use of fractal geometry to determine the impact of inlet migration on the dynamics of a seagrass landscape. *Estuarine, Coastal and Shelf Science*, 84(4):584–590, 2009.
- Vasilis Dakos, Egbert H van Nes, Raúl Donangelo, Hugo Fort, and Marten Scheffer. Spatial correlation as leading indicator of catastrophic shifts. *Theoretical Ecology*, 3(3):163–174, 2010.
- Vasilis Dakos, Sonia Kéfi, Max Rietkerk, Egbert H van Nes, and Marten Scheffer. Slowing down in spatially patterned ecosystems at the brink of collapse. *The American Naturalist*, 177(6):E153–E166, 2011.
- C Middelburg Duarte, Jack J Middelburg, N Caraco, et al. Major role of marine vegetation on the oceanic carbon cycle. *Biogeosciences*, 2(1):1–8, 2005.
- Carlos M Duarte and Hugh Kirkman. Methods for the measurement of seagrass abundance and depth distribution. *Global seagrass research methods*, pages 141–153, 2001.
- Johan Erlandsson, Christopher D McQuaid, and Martin Sköld. Patchiness and co-existence of indigenous and invasive mussels at small spatial scales: the interaction of facilitation and competition. *PloS one*, 6(11):e26958, 2011.
- Kenneth Falconer. *Fractal geometry: mathematical foundations and applications*. John Wiley & Sons, 2013.
- Jessica L Green, John Harte, and Annette Ostling. Species richness, endemism, and abundance patterns: tests of two fractal models in a serpentine grassland. *Ecology Letters*, 6(10):919–928, 2003.
- Vishweshha Guttal and Ciriya Jayaprakash. Spatial variance and spatial skewness: leading indicators of regime shifts in spatial ecological systems. *Theoretical Ecology*, 2(1):3–12, 2009.
- JM Halley, S Hartley, AS Kallimanis, WE Kunin, JJ Lennon, and SP Sgardelis. Uses and abuses of fractal methodology in ecology. *Ecology Letters*, 7(3):254–271, 2004.
- Matthew C Hansen, Stephen V Stehman, and Peter V Potapov. Quantification of global gross forest cover loss. *Proceedings of the National Academy of Sciences*, 107(19):8650–8655, 2010.

- Harold M Hastings and George Sugihara. Fractals. a user's guide for the natural sciences. *Oxford Science Publications, Oxford, New York: Oxford University Press*, — c1993, 1, 1993.
- Harold M Hastings, Richard Pielou, Richard Monticciolo, David Vun Kannon, and Diane Del Monte. Time scales, persistence and patchiness. *BioSystems*, 15 (4):281–289, 1982.
- Marten A Hemminga and Carlos M Duarte. *Seagrass ecology*. Cambridge University Press, 2000.
- Attila R Imre, Dániel Cseh, Markus Neteler, and Duccio Rocchini. Korcak dimension as a novel indicator of landscape fragmentation and re-forestation. *Ecological indicators*, 11(5):1134–1138, 2011.
- Attila R Imre, Josef Novotný, and Duccio Rocchini. The korcak-exponent: a non-fractal descriptor for landscape patchiness. *Ecological complexity*, 12:70–74, 2012.
- Anna E Jolles, Patrick Sullivan, Alisa P Alker, and C Drew Harvell. Disease transmission of aspergillosis in sea fans: inferring process from spatial pattern. *Ecology*, 83(9):2373–2378, 2002.
- Sonia Kéfi, Max Rietkerk, Manojit Roy, Alain Franc, Peter C De Ruiter, and Mercedes Pascual. Robust scaling in ecosystems and the meltdown of patch size distributions before extinction. *Ecology letters*, 14(1):29–35, 2011.
- Sonia Kéfi, Vishweshia Guttal, William A Brock, Stephen R Carpenter, Aaron M Ellison, Valerie N Livina, David A Seekell, Marten Scheffer, Egbert H van Nes, and Vasilis Dakos. Early warning signals of ecological transitions: Methods for spatial patterns. *PloS one*, 9(3):e92097, 2014.
- Gary A Kendrick, Michelle Waycott, Tim JB Carruthers, Marion L Cambridge, Renae Hovey, Siegfried L Krauss, Paul S Lavery, Donald H Les, Ryan J Lowe, Oriol Mascaró i Vidal, et al. The central role of dispersal in the maintenance and persistence of seagrass populations. *BioScience*, 62(1):56–65, 2012.
- Jeremy T Kerr and Marsha Ostrovsky. From space to species: ecological applications for remote sensing. *Trends in Ecology & Evolution*, 18(6):299–305, 2003.
- Anthony WD Larkum, Robert Joseph Orth, Carlos M Duarte, et al. *Seagrasses: biology, ecology, and conservation*. Springer, 2006.

- Michael A Lefsky, Warren B Cohen, Geoffrey G Parker, and David J Harding. Lidar remote sensing for ecosystem studies lidar, an emerging remote sensing technology that directly measures the three-dimensional distribution of plant canopies, can accurately estimate vegetation structural attributes and should be of particular interest to forest, landscape, and global ecologists. *BioScience*, 52(1):19–30, 2002.
- Olivier Lejeune and Mustapha Tlidi. A model for the explanation of vegetation stripes (tiger bush). *Journal of Vegetation Science*, 10(2):201–208, 1999.
- Bai-Lian Li. Fractal geometry applications in description and analysis of patch patterns and patch dynamics. *Ecological Modelling*, 132(1):33–50, 2000.
- BB Mandelbrot. *Fractals: form, chance and dimension*. WH Freeman & Co., 1979.
- Benoit B Mandelbrot. How long is the coast of britain. *Science*, 156(3775):636–638, 1967.
- Benoit B Mandelbrot. *The fractal geometry of nature*, volume 173. Macmillan, 1983.
- Susannah Mistr and David Bercovici. A theoretical model of pattern formation in coral reefs. *Ecosystems*, 6(1):0061–0074, 2003.
- Ran Nathan and Helene C Muller-Landau. Spatial patterns of seed dispersal, their determinants and consequences for recruitment. *Trends in ecology & evolution*, 15(7):278–285, 2000.
- Mercedes Pascual, Manojit Roy, Frédéric Guichard, and Glenn Flierl. Cluster size distributions: signatures of self-organization in spatial ecologies. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 357(1421):657–666, 2002.
- JD Pringle. Efficiency estimates for various quadrat sizes used in benthic sampling. *Canadian Journal of Fisheries and Aquatic Sciences*, 41(10):1485–1489, 1984.
- Max Rietkerk, Stefan C Dekker, Peter C de Ruiter, and Johan van de Koppel. Self-organized patchiness and catastrophic shifts in ecosystems. *Science*, 305(5692):1926–1929, 2004.
- Marten Scheffer, Jordi Bascompte, William A Brock, Victor Brovkin, Stephen R Carpenter, Vasilis Dakos, Hermann Held, Egbert H van Nes, Max Rietkerk, and George Sugihara. Early-warning signals for critical transitions. *Nature*, 461(7260):53–59, 2009.

- Susan Schwinning, Osvaldo E Sala, Michael E Loik, and James R Ehleringer. Thresholds, memory, and seasonality: understanding pulse dynamics in arid/semi-arid ecosystems. *Oecologia*, 141(2):191–193, 2004.
- Laurent Seuront. *Fractals and multifractals in ecology and aquatic science*. CRC Press, 2009.
- Nick T Shears and Russell C Babcock. Marine reserves demonstrate top-down control of community structure on temperate reefs. *Oecologia*, 132(1):131–142, 2002.
- George Sugihara and Robert M May. Applications of fractals in ecology. *Trends in Ecology & Evolution*, 5(3):79–86, 1990.
- David Tilman. Competition and biodiversity in spatially structured habitats. *Ecology*, 75(1):2–16, 1994.
- Brant W Touchette and JoAnn M Burkholder. Review of nitrogen and phosphorus metabolism in seagrasses. *Journal of Experimental Marine Biology and Ecology*, 250(1):133–167, 2000.
- Johan van de Koppel, Max Rietkerk, Norbert Dankers, and Peter MJ Herman. Scale-dependent feedback and regular spatial patterns in young mussel beds. *The American Naturalist*, 165(3):E66–E77, 2005.
- Tjisse Van Der Heide, Tjeerd J Bouma, Egbert H Van Nes, Johan Van De Koppel, Marten Scheffer, Jan GM Roelofs, Marieke M Van Katwijk, and Alfons JP Smolders. Spatial self-organized patterning in seagrasses along a depth gradient of an intertidal ecosystem. *Ecology*, 91(2):362–369, 2010.
- J Von Hardenberg, E Meron, M Shachak, and Y Zarmi. Diversity of vegetation patterns and desertification. *Physical Review Letters*, 87(19):198101, 2001.
- John A Wiens. Spatial scaling in ecology. *Functional ecology*, 3(4):385–397, 1989.
- Robert J Wilson, Chris D Thomas, Richard Fox, David B Roy, and William E Kunin. Spatial patterns in species distributions reveal biodiversity change. *Nature*, 432(7015):393–396, 2004.
- Justin P Wright and Clive G Jones. The concept of organisms as ecosystem engineers ten years on: progress, limitations, and challenges. *BioScience*, 56(3):203–209, 2006.

Xiao-ping Xin, Qiong Gao, Yi-yin Li, and ZY Yang. Fractal analysis of grass patches under grazing and flood disturbance in an alkaline grassland [j]. *Acta Botanica Sinica*, 41(3):307–313, 1999.

Chapter 2

Background

... when he sees in a snail, or nautilus, or tiny foraminiferal or radiolarian shell a close approach to sphere or spiral, he is prone of old habit to believe that after all it is something more than a spiral or a sphere, and that in this “something more” there lies what neither mathematics nor physics can explain.

(D’Arcy W. Thompson - *On Growth and Form*)

2.1 Introduction

Ecosystems such as marine, semi-arid and tropical are predominantly defined via the composition and distribution of their vegetation. Vegetation can colonise barren land and through successional species lead to a composition of species that appear in equilibrium with their environment. The vegetation species of many ecosystems therefore are pivotal in the sustainability and resilience of the overall system and an understanding of their dynamics is vital.

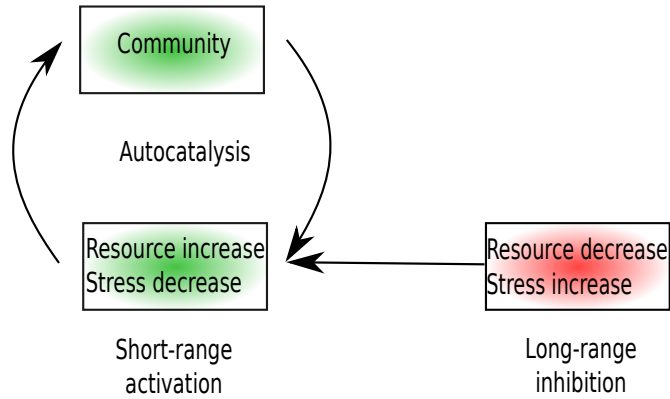


Figure 2.1: General mechanism for an activator-inhibitor system in vegetation. Locally positive interaction occurs due to increase in resources and decrease stress. On larger spatial scales inhibition occurs through resource depletion and increase in stress (Adapted from [Rietkerk and Van de Koppel, 2008]).

Many processes affect the persistence and pattern of the vegetation including: grazing, disease, environment and fire. These processes will often be mediated at a range of scales: on larger scales composition of the vegetation is predominantly determined by the environment, which determines the viability of where vegetation can occupy. On smaller scales feedback processes such as through environmental factors such as soil composition and ground water determine the finer structure of the vegetation community. At even finer scales on the level of individuals stochastic processes dominate: these processes include growth of new plants from the underlying seed bank; clonal growth from rhizomal propagation and shooting; and death. This highlights the importance of scale in the vegetation community. If observing at a particular scale, identifying what processes dominate at those spatial ranges is vital. It is also important to identify if such a process is acting at a particular

characteristic length or if it is operating over a range of scales.

This chapter contains an outline of the modelling and theoretical research carried out to address the hypotheses outlined here. An outline of regular pattern formation in vegetation is carried out in the next section. This is where processes between vegetation and its environment create spatial patterns at a characteristic length scale. Patterns where there is no characteristic length scale, such as power law patch-size distributions are then reviewed. This leads on to stochastic patterns of vegetation with rough boundaries that are indicative of fractal phenomena. A review of the explanations of these phenomena are then given.

2.2 Pattern formation

In order to understand how persistence influences spatial pattern and vice versa, we must first understand the broad range of spatial patterns that exist in vegetation ecosystems. Rietkerk and Van de Koppel [2008] provides a review of well-established regular spatial patterns that exist in ecology. The leading explanation of how patterns of characteristic length scale arise from vegetation dynamics is that of Turing patterns [Murray, 2001]. In the seminal paper Turing [1952] proposes a mechanism of symmetry-breaking that leads to stable patterns in space. The mechanism under which this occurs is based on the interaction between two chemical species: an activator and an inhibitor. The activator is in a positive feedback loop with itself, while it also produces an inhibitor. The inhibitor inhibits the further production of the activator (Fig. 2.1). For a spatially homogeneous system, these dynamics lead

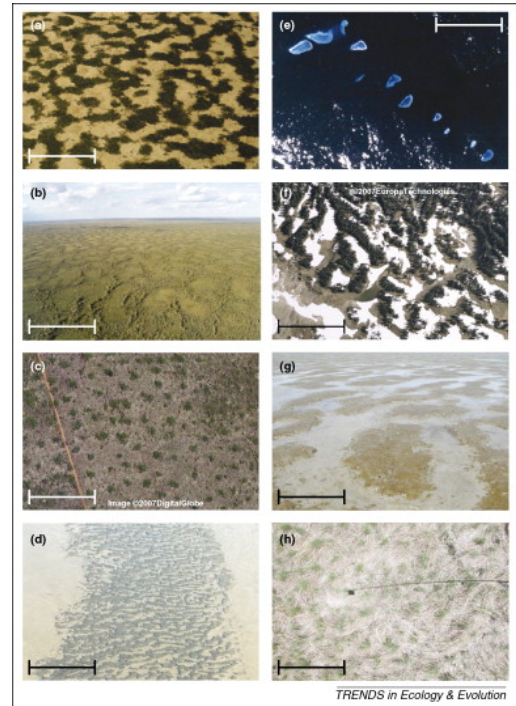


Figure 2.2: Various examples of regular pattern formation in vegetation. Examples are (a) semi-arid vegetation (b) shrubs & trees in Siberia (c) spotted vegetation in Niger (d) mussel beds (e) reef islands (f) ribbon forests (g) marine benthic diatoms and (h) sedge (Reproduced from [Rietkerk and Van de Koppel, 2008]).

to a stable fixed point with a non-zero concentration of activator and inhibitor. However, when space is introduced, via a diffusion operator on the inhibitor and activator, small spatial perturbations can lead to the system moving away from the homogeneous fixed point to one of stable spatial patterns (when the diffusion constant of the inhibitor is sufficiently larger than the diffusion operator of the activator).

The resulting patterns are due to the separation of scales between activator and inhibitor. On small scales, the activator is able to diffuse to reinforce the local high concentration of activator. On larger scales the inhibitor prevents the further diffusion of the activator leading to patterns of a fixed length scale such as spots, stripes and labyrinths. This activator-inhibitor mechanism provides an allegory for vegetation and its environment, where vegetation is dominated by positive local feedback on small scales, but negative feedback on larger spatial scales.

These regular patterns have been found in a number of ecosystems including wetlands [Foster et al., 1983; Sakaguchi, 1980; Swanson and Grigal, 1988], Savanna [Lejeune et al., 2002], Mussel beds [Bertness and Grosholz, 1985; Okamura, 1986; van de Koppel et al., 2005], coral reefs [Mistr and Bercovici, 2003] and intertidal mudflats [Blanchard et al., 2000; De Brouwer et al., 2000] (See Fig. 2.2). For marine systems such as mussel beds, intertidal mudflats and coral reefs mechanisms that induce spatial patterning are associated with the interplay between currents; the density of vegetation and other species; and the environment by way of nutrients, sediment and algae concentrations. The resulting spatial patterns give an indication as to the underlying model parameters that gave rise to such patterns.

Turing bifurcations are not the only pattern-inducing transitions known to occur in vegetation ecosystems. Shnerb et al. [2003] identified a model where formation of patches in an irregular pattern is due to a non-Turing mechanism. Growth was divided into two seasons: summer and winter. Although growth in the first season was modelled by equations with a stable homogeneous fixed point, the dynamics of the second season, where vegetation is depleted depending on a fixed size cut-off, keeps the system away from the equilibrium state. Thus the system develops characteristics of a glassy system more commonly associated with lattice dynamics in physics. These are systems where there is some fixed random structure in the assemblage of the individual units that are held in place due to the system being trapped in an energy well far from equilibrium.

Another type of non-Turing pattern was introduced by D’Odorico et al. [2007]. They considered a model of vegetation in the presence of random fire events where the vegetation has logarithmic growth and spatial diffusion. In the presence of no stochastic term, this is the regular Fisher equation, which produces no Turing bifurcation. An added stochastic term in the form of a Poisson process was used to model disturbances due to fires. The resulting spatial patterns produced are transient in nature and do not have a regular length-scale. This provides caution that although if a particular spatial pattern is observed, it is not necessarily the case that the pattern forming process can be immediately identified from the spatial pattern alone and the underlying mechanisms must also be discerned.

Regular spatial patterns are one class of patterning that is ubiquitous in vegetation ecosystems. Another type are scale-free spatial patterns, which in contrast to regular patterns do not have a characteristic length-scale. Usually there is some property of the pattern, such as the roughness of cluster boundaries or the size distribution of clusters that has a scale free or power-law distribution where the cluster size s is related to its frequency via the relationship

$$N(s) \sim s^{-\gamma}, \quad (2.1)$$

where γ is the exponent of the power-law distribution. An explanation for the formation of these scale-free patterns comes from the the idea put forward by Bak et al. [1987], known as self-organised criticality. This is where, without fine-tuning of parameters or starting from a particular initial state, an extended spatial system with many degrees of freedom evolves towards a critical state. In this critical state, the system is barely stable and small perturbations can lead to dissipations on all length-scales. A concrete, relevant example of a system that displays self-organised criticality is the forest fire model [Drossel and Schwabl, 1992]. The model is an abstraction of a forest that can be colonised with a single species and is subject to random disturbances in the form of fire that can burn through the vegetation leading to barren states that can be occupied again. The model is specified as a probabilistic cellular automata on a square lattice, where each cell can either be empty, occupied or burning. The model is updated at each time-step using four rules

1. A burning site becomes an empty site with probability 1.
2. An occupied site becomes burning if one of its neighbours is in a burning state.

3. An empty site becomes occupied with probability p .
4. An occupied site can become burning with probability f .

Time can be rescaled so there is a single parameter on which the dynamics depends p/f . For a suitable parameter range, there is a cascade of fire events, where the total size of a forest fire has a power-law distribution. This critical state is reached regardless of initial conditions and for a large range of the parameter. The model dynamics evolve the system to a state where disturbance cascades of any size can occur (if the system size is suitably larger than the range of scales considered).

2.2.1 Semi-arid ecosystems

Semi-arid ecosystems have produced a large amount of research into regular pattern formation, where the competition for water dominates growth of vegetation. Modelling has taken two forms. Thiery et al. [1995] introduces a model of pattern formation in arid ecosystems via a deterministic cellular automata model. Space is divided up into a lattice and each lattice site contains the quantity $S_{i,j,t}$ representing the quantity of vegetation at site i, j at time t . The deterministic dynamics are then updated according to a convolution, which mediates the local and long-range cooperation and competition effects. The dynamics are then summarised as

$$S_{i,j,t+1} = S_{i,j,t} + \max\{-1, \min\{1, (A * S)_{i,j}\}\}, \quad (2.2)$$

$*$ is the convolution operator and A is a matrix representing the cooperation and competition between plants. Thiery considers a simplified interaction matrix with two parameters representing interactions in the presence of a gentle slope and the resulting spatial patterns are qualitatively consistent with observations from aerial photographs. The interaction being locally positive and negative at long-range is again important for vegetation bands to occur. Although the spatial patterns produced are striking, the model is purely phenomenological in its description and no quantitative parameter fitting was performed.

Lefever and Lejeune [1997] also consider a deterministic model of vegetation growth with long-range spatial competition that induces regular spatial patterns. The kinetic equation has three components F_1 represents the growth of vegetation density c due to facilitation and local dispersal. F_2 is a competition term that limits the growth of the c term. The final term, F_3 represents mortality and leads to a decrease in c . Lejeune and Tlidi [1999] further analysed this model to find that the length-scale of reproduction must be shorter than the length-scale of competition

in order for pattern formation to occur. Increasing of the death rate leads to a predictable change in the resulting spatial pattern: from vegetation with regular gaps, to a banded labyrinth structure, to vegetation spots until finally reaching the barren state for high mortality. The deterministic nature of the kinetic model however leads to regular patterns, which are difficult to observe in the field due to heterogeneity of the environment and other stochastic processes associated with vegetation growth and death. Klausmeier [1999] introduces a model of vegetation growth where the concentration of ground water is explicitly modelled via a reaction-diffusion equation. The model produces stripes that are perpendicular to the direction of flow of the ground water. The resulting regular patterns are dynamic and move uphill. The wavelength of the bands are connected with the parameters of the system; the wavelength of stripes decrease for increasing water input rate and plant loss rate. Stripes occur as a boundary between the barren state, where there is no vegetation present and the homogeneous state, where there are no regular spatial patterns. This model indicates the applicability of inferring parameters from spatial pattern, although if it were purely done from wavelength a family of parameter values would be applicable, hence other parameters such as water input would need to be inferred by other means. Sherratt [2005] extended the analysis of the model to the other model parameters. They found a square root relationship between the wavelength of the stripes and the gradient of the slope. This again shows the applicability of inferring parameters from spatial snapshots if other parameters are known. HilleRisLambers et al. [2001] extended the model again to include dynamics of surface water without the presence of a slope (no ground water current term). The resulting model is a three-dimensional reaction-diffusion equation that has a large region of parameter space where banding exists. This is contrasted with Klausmeier [1999] where an environmental gradient was necessary to induce banding. Kéfi et al. [2010] further extends the model of HilleRisLambers et al. [2001]; Rietkerk et al. [2002] to include new non-linearities in the interaction between vegetation density and surface water. The resulting equations admit a bistable region, where both vegetation and the barren state can coexist. Both Turing and non-Turing spatial patterns are observed and the resulting spatial patterns occur in the region where the spatially-homogeneous system is bistable. This indicates in this model at least that certain spatial patterns can be associated with a bistable region and hence imminent desertification. These conclusions do rely on the underlying model being valid as we have seen that other models produce similar spatial patterns without the need to be bistable.

The interaction between vegetation and its environment producing regular patterns is particularly well-studied in the field of ecohydrology, where there is competition between vegetation units for ground water inducing spatial pattern. Borgogno et al. [2009] reviews the main mechanisms proposed to pattern formation in ecohydrology including both deterministic and stochastic models. They conclude with a set of open challenges in the field. These include under what conditions pattern can form in a noisy environment and what form of model validation can occur purely on observations of spatial pattern when many mechanisms are known to generate similar spatial patterns.

2.3 Critical phenomena

Criticality is the phenomena exhibited in many systems where small changes in the underlying parameters of the system lead to large-scale qualitative system changes. The classically studied forms of criticality have been those of thermodynamics such as the liquid-gas phase transition. For a slight temperature increase there is a large-scale observable change in the qualitative properties of the matter being heated. At the critical point, both phases exist simultaneously and spatial correlations decay algebraically. This introduces the idea of using spatial statistics, such as spatial correlation functions to detect the presence of a critical point. These spatial statistics are known as order parameters due to their characterisation of order near phase transitions, where there is a sudden change in the order of the system. Many ecological systems exhibit patchiness and scale-free behaviour in their spatial distribution near a critical point. Pascual and Guichard [2005] identify three types of criticality: classical, self-organised and robust. Classical have been discussed previously as with the liquid-gas phase transition, when a system has an order parameter that varies dramatically with an external parameter over a single point, the order parameter is considered to not affect the external parameter. Self-organised criticality (SOC) on the other hand occurs when a system evolves freely towards its critical point without any external fine-tuning [Bak, 1997; Pascual et al., 2002]. The prototypical example of SOC is the sand-pile model introduced by Bak et al. [1987]. The idea is to imagine a flat surface where sand is dropped at constant rate but in random points causing sand-piles to emerge. The abstraction of this process is a cellular automata where a one-dimensional grid of infinite length is set-up with initially empty sites. Sand is modelled as single units that are placed randomly onto the length of the grid. Relaxation occurs when a sand unit lands on a site; the neighbouring sites values are checked against the centre one, if the difference between the site and its

neighbours is greater than one then the sand particle jumps into the lower site. The new site, where the sand has jumped to, is again checked to see if the difference in particles between itself and its neighbours is greater than one and the relaxation dynamics continue until this inequality is satisfied. After a certain relaxation period it is found that the length of these relaxation events or avalanches that occur can be arbitrary in size; the distribution of avalanche sizes is in fact a power law. The system is thus one that has freely evolved into a state of criticality. Other models that exhibit SOC have also been produced such as an evolutionary model, where extinction events have a power law distribution [Sneppen et al., 1995]. The final form of criticality proposed is robust criticality, where there is scale-invariance and long-range correlations for a wide-range of parameters in the system (for an example see [Guichard et al., 2003]). The review also states that three or more states are required for robust criticality to occur. Obvious questions arise about why this number of states is necessary for robust criticality to emerge and how the power-law observations depend upon the parameters in the critical region.

The theory of systems near criticality is well-established for physical systems, where criticality is associated with the presence of a second order phase transition [Stanley, 1987]. For more biological systems there is little general theory of systems near criticality, Scheffer et al. [2009] provides some examples of indicators that occur as a system approaches a threshold. An example of this general theory comes from imagining a system with a smooth potential. The potential has two local minima with a potential barrier separating the two. For the ecosystem case these two minima could represent a vegetated and barren state. We imagine that there is some parameter (e.g. rainfall) that is able to change the height of the potential barrier. When the barrier is high, perturbations in the system quickly decay as the local minima is distinct and it would require a large perturbation to leave the local minima. When the barrier is lowered, the gradient of the potential well also reduces, thus perturbations decay at a slower rate. At the critical point, perturbations can take an arbitrarily long time to decay. These dynamics translate into population time-series statistics as an increasing variance and skewness in the time-series as a system approaches a critical point. Flickering phenomena is also seen to occur where, if the potential barrier is low enough, small perturbations can allow a system to jump between local minima and thus the time-series appears to have two stable points with random switching between the two. Thus the three hallmarks of a system at criticality are increased variance; spiking in skewness and flickering phenomena. For a spatially-explicit system where regular spatial patterns occur, predictable

spatial patterns can form near a critical bifurcation [Rietkerk et al., 2004]. However, [Scheffer et al., 2009] notes that these change in spatial pattern are model-specific and thus the conclusions of how the spatial pattern predicts the critical transition is only true insofar as the underlying model is an adequate description of the system in question.

The work on detection of criticality in the observations of natural systems has mainly focused on the temporal characteristics of the system. Guttal and Jayaprakash [2009] proposes a methodology where the spatial characteristics of a signal spatial snapshot can be used to detect when an ecological system is close to criticality. They consider a vegetation with logistic growth, various dispersal mechanisms including local diffusion and heavy-tailed spatial kernels as well as a stochastic grazing pressure that leads to the vegetation being bistable between the barren and vegetated states. Spatial variance was shown to increase and spatial skewness spikes near the critical point. They also demonstrate that intervention based on use of these spatial statistics can lead to recovery of the ecosystem that would otherwise collapse under constant grazing pressure. Although a wide variety of models were proposed where variance and skewness were good indicators of critical transitions in the sense that acting on them was able to avert the collapse of a population, Hastings and Wysham [2010] shows this is not always the case. They constructed two models based on the Ricker equation of population growth in order to demonstrate the limitations of these indicators. Whereas Guttal and Jayaprakash [2009] had model equations with smooth potentials around the critical transition points, the Hastings and Wysham model does not and thus the general argument of population fluctuations near a transition point do not hold. This leads to no warning of a dramatic population shift from the fluctuation statistics. Thus model-specific assumptions are vital to understand when applying general theory to real ecological data.

2.4 Scaling

Scaling in spatial ecology has been recognised as an important link between pattern and process [Levin, 1995]. The idea that an underlying generative process can lead to patterns that hold a particular scaling that are invariant began with fractals.

Fractal theory began as a purely descriptive science for characterising the patchy and irregular patterns in nature. Since then it has had an increasing number of applications in Ecology. Sugihara and May [1990] lay out the claim that there is a

connection between the patterns generated by an underlying process and the process itself. They give an example of the fractional Brownian motion (fBm) to highlight this claim. Standard Brownian motion was originally proposed as a model for how grains of pollen diffuse in a liquid being buffeted by random forces. The motion of the particle is described in one dimension by a displacement $X(t)$. For standard Brownian motion the root mean squared of the particle distance scales as

$$\sqrt{\mathbb{E}[(X(t) - X(0))^2]} \sim t^{1/2}. \quad (2.3)$$

In fractional Brownian motion this relationship is generalised where the scaling of root mean squared is replaced by an exponent H , known as the Hurst exponent. i.e.

$$\sqrt{\mathbb{E}[(X(t) - X(0))^2]} \sim t^H. \quad (2.4)$$

There are three main domains for the Hurst parameter: $0 < H < 1/2$, $H = 1/2$ and $1/2 < H < 1$. When $0 < H < 1/2$ the trajectory is known as anti-persistent, where the future trajectory of the particle is anti-correlated with its current displacement. For $1/2 < H < 1$ the fBm is described as persistent; the trajectory is correlated with its past displacements and a realisation of the process is smoother than for lower values of the Hurst exponent. $H = 1/2$ is the classical Brownian motion where the process is memoryless and the future trajectory is uncorrelated with its past. This model is traditionally used to detect the presence of memory in time-series analysis. The process can also be generalised to more than one dimension, in three dimensions level-sets of the process have been employed as neutral models of patchy landscapes [Hastings et al., 1982; Keitt, 2000]. These patches have a resulting fractal dimension defined by the parameter H . This is a tantalising clue that spatial pattern can be analysed to determine the underlying process.

There are however issues when applying fractal theory to real data. Fractal definitions are defined as limits that can go arbitrarily small [Falconer, 2013]. In other words, in order to perform fractal analysis the data set in question would have to have an infinite number of scales. This is clearly not possible and hence instead fractal analysis is usually performed by considering a range of scales over which some power law relationship holds [Seuront, 2009]. This leads to issues such as biases in the linear regression performed or if the range over which the scaling occurs is too narrow. Indeed, there is confusion and dispute over the definitions of the various fractal measures and their relationship to one another [Halley et al., 2004].

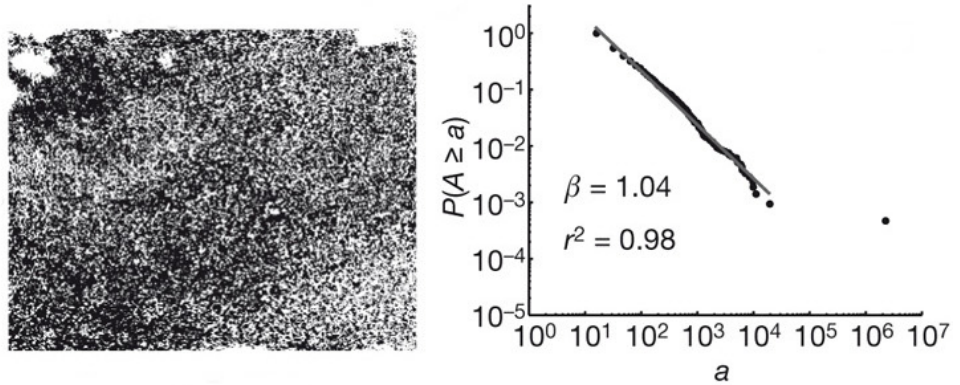


Figure 2.3: Measuring the Korcak exponent k from an aerial photograph of vegetation clusters in the Kalahari. The image is converted into counts of patch sizes, which is used to estimate the patch-size distribution, $P(A \geq a)$, the probability that a randomly chosen patch is greater than size a . On a log-log scale the distribution is approximately linear and can be fitted using linear regression. The gradient of the fitted line is the estimate of the Korcak exponent. Reproduced from Scanlon et al. [2007].

Another widely used fractal measure used in spatial ecology to characterise the patchiness of landscapes is the Korcak dimension. Originally, it was introduced by Korčák [1938] studying the size distribution of Aegean islands. Mandelbrot [1983] further developed the theory of the dimension and how it could be applied to analyse certain fractal models. If the sizes of patches follows the distribution

$$N(A \geq a) = na^{-k}. \quad (2.5)$$

Then the Korcak dimension D_k is defined to be

$$D_k = 2k. \quad (2.6)$$

Mandelbrot [1983] explored the possibilities of using this as a measure of fragmentation in a vegetation landscape. The Korcak dimension was proposed as a first degree measure of succession for a vegetation species [Hastings et al., 1982]. More recently it has been used to measure patchiness and reforestation [Imre et al., 2011, 2012].

Scaling of power-law distributions has been observed in the semi-arid vegetation over a wide range of environmental conditions [Scanlon et al., 2007]. There was found to be variation in the power-law exponent of the patch-size distribution as

rainfall varied in different locations. A binary cellular automata model was used to explain this distribution where there was a global term based on overall rainfall and a local positive feedback term. It was found that the positive feedback term was necessary to induce power-law behaviour, but the behaviour was robust to a range of local feedback parameters. Although the model does reproduce the distribution well, the relationship between the scaling exponent of the power-law distribution and the ability of the system to persist still remains an open problem.

Kéfi et al. [2007a] analysed patch-size distributions in the presence of various grazing pressures. Under high grazing pressure the patch-size distribution deviated from a power law characterised by a distribution with exponential cut-off. They propose a model of vegetation growth where growth has a local positive feedback mechanism and grazing is modelled as a spatially homogeneous rate at which cells die. The model was found to reproduce the broad features of the observed vegetation patches including the power-law meltdown for increased grazing pressure. Maestre and Escudero [2009] tested whether deviations from a power-law patch-size distribution are indicative of immediate desertification. They found that deviations as measured by fitting a truncated power law were not positively correlated with desertification variables, but vegetation cover was. Kéfi et al. [2011] provides further restrictions on when the meltdown of power-law structures in robust critical systems occurs when the system is under an increased amount of stress. They analysed four models that display robust criticality from an arid ecosystem model [Kéfi et al., 2007b], a mussel bed model [Guichard et al., 2003], predator-prey model [Durrett and Levin, 1994; Pascual et al., 2002] and a null model where there is no spatial dependency. Although there is a certain universality to the meltdown of the power-law patch size distribution near a critical transition, no single indicator provides an overall picture of the distance to a critical transition and measurements must be taken relevant to each other. Also no experimental work on real ecosystems has been used to verify these claims, hence there are still gaps in the knowledge of how the patch-size distribution can be used in the practical application of measuring extinction risk.

Guichard et al. [2003] explored patch-size dynamics in a mussel bed system. Mussels attach to rock and conspecifics, thus facilitating local positive interaction. However, they are also vulnerable to disturbance from wave action, which can buffet mussels leading to uprooting and disturbance. This system was again modelled as a lattice where each site could be in one of three states: empty, occupied or disturbed. The occupied sites spread through into empty sites through positive density depen-

dence. The disturbed sites are also assumed to have a density dependent feedback, where the disturbance can spread locally. The resulting dynamics leads to robust criticality, where there is a large region of parameter space where power-law scaling of the patch-size distribution is observed. The exponent of the distribution was found to almost be independent of the parameters however. This raises the question as to what conditions the exponent is indicative of the underlying parameters and ultimately the dynamics.

One particular application of fractal theory is in invasion fronts of vegetation. Cannas et al. [2006] introduce a model of vegetation growth with long-range dispersal and show that the fractal boundary of the front increases with increasing dispersal length under certain model assumptions. This indicates that it is theoretically possible to capture some broad features of an invasion front from the underlying dynamics. How general this relationship is in the presence of environmental feedback is still an open problem, however. Recent evidence suggests that fractal roughening of a vegetation invasion front can be detected in an experimental system [Allstadt et al., 2014].

The concept of scale and spatial variability in vegetation spatial patterns leads naturally to the idea of multifractals. Intuitively we can consider multifractals as a description of a spatial pattern with scale-invariance like in the standard fractal case, but the scale-invariance changes locally in space. This extends the standard definition of a fractal dimension to include a range of fractal dimensions or Rényi Dimensions with parameter q . Multifractal methodology has had some application in vegetation patterns, including the use of it to classify the varying dispersal length-scales for two competing vegetation species Scheuring and Riedi [1994]. A multifractal distribution of vegetation has also been interpreted as evidence for self-organised criticality in the system [Solé and Manrubia, 1995a]. Theoretically, the multifractal measure seems to be an elegant solution to the problem of characterising both scale and variability in a vegetation landscape. However, there has been some criticism surrounding how the measure can be applied to noisy ecological data and the validity of the results based on these measurements [Zamir, 2003].

Kubo et al. [1996] introduced the idea of using the Ising model to characterise the observed spatial features of the canopy dynamics. The Ising model, although sharing features with other cellular automata such as being based on a finite lattice with a finite state space differs in dynamics as it is a continuous-time process,

whereas cellular automata are discrete. The Ising model considers a state space where each lattice site is in one of two states $\{+1, -1\}$. The total lattice S therefore has $2^{N \times N}$ configurations, where each configuration is denoted as σ . The probability of a configuration being in state x is given by

$$P(\sigma = x) = \frac{1}{Z} \exp \left(\beta \sum_{i \sim j} x_i x_j + h \sum_i x_i \right), \quad (2.7)$$

where $i \sim j$ denotes the neighbours of j . β is the interaction strength between neighbours and h is a forcing term that represents the tendency to favour the occupied (+1) state. When the parameters h and β are fine-tuned, critical phenomena occurs such as algebraic decay of spatial correlation and power-law size distributions of the +1 and -1 state, which is indicative of a phase transition.

It was found that although the Ising model can recover some of the properties of the observed gap dynamics [Kizaki and Katori, 1999], in general the Ising model does not provide a good fit to the observed data due to forest gap dynamics being a far from equilibrium process [Schlicht and Iwasa, 2004]. One example of where multifractal methodologies have been applied is to forest gap dynamics. Forests have been observed to have power-law distributions in their gap sizes. A cellular automata model of growth with gap formation was considered where births and deaths are random with independent probabilities. Each individual grows according to the light resources available in its immediate surroundings, this is determined by the total height of the immediate neighbours of the individual and an interaction parameter γ . The resulting multifractal spectrum was then computed and compared to data of forest gap distributions taken from an aerial photographic survey of Barro Colorado Island, Panama. A region of parameter space shows power-law distribution in the gap sizes as well as a broad multifractal spectrum indicative of a self-organised critical process [Manrubia and Solé, 1996; Solé and Manrubia, 1995a,b]. As the multifractal spectrum varies throughout the parameter space, there would be a possibility of using it for parameter estimation of a single spatial snapshot, although estimations of probability of death and birth were taken from studies and only the interaction parameter was varied to fit to data, this does at least give the plausibility of performing some form of parameter estimation on a single spatial snapshot.

2.5 Conclusion

In this review a wide variety of spatial mechanisms that induce pattern formation in vegetation ecology have been explored, with the focus on the relationship between spatial pattern and persistence. The types of spatial pattern belong roughly to two classes: scale-free and regular. Although there have been many models proposed to explain these phenomena in various ecosystems there are some similarities between them. Regular spatial patterns are often seen as a sign of local positive feedback processes along with long-range inhibition mediated by the environment either through competition for resources or some other long-range mechanism. Regular spatial patterns can also be viewed as transient states as a system moves towards either a barren or vegetated equilibrium. This can also occur where the system is somehow kept far from equilibrium, hence although it appears the system is moving towards an equilibrium point it may take an arbitrarily long time in order to reach it.

Scale-free patterning in the form of power law patch-size distributions are also a ubiquitous characteristic of vegetation ecosystems. The meltdown of the power-law distribution has been found to be a robust indicator as to whether a system is close to a tipping point. Although questions remain over what systems this form of behaviour is exhibited and indeed if there are other indicators that can be used on single spatial snapshots to determine if a system is close to criticality. The other indicators proposed include the exponent of the patch-size distribution itself, which is often referred to as the Korcak exponent. It is not clear what the origin of this scale-free behaviour is in these systems or whether the exponent is related to the dynamic persistence of the system (although it does provide a measure of the spatial persistence).

Scale-free behaviour in the form of fractal boundaries of a growing vegetation cluster can also give insight into the underlying dispersal mechanisms that produce the spatial pattern, although the effect is conflated by a heterogeneous environment. Therefore what mechanisms lead to the roughening of a vegetation boundary in a heterogeneous environment and how this roughening relates to dynamics remains an open problem. A scale-free patch-size distribution has also been observed in many ecosystems and has been numerically confirmed using cellular automata models of vegetation growth. The origin of these power-law structures in vegetation growth is still a mystery as is the truncation of the power-law to an exponential distribution under increased environmental stress.

Both scale-free and regular patterns are exhibited in seagrass ecosystems. Seagrass exists as a monoculture and has strong feedback loops with its environment. As such it is an ideal candidate to study the interaction between pattern and process. Although a certain amount of modelling research has been performed on seagrass ecosystems, models have not been produced that exhibit both types of pattern phenomena and can be used to determine how spatial pattern impacts persistence.

Ecosystems that are capable of producing regular spatial pattern are also capable of producing scale-free patterns. Indeed the underlying mechanisms such as local positive-feedback are similar for inducing both scale-free and characteristic length-scale patterns. If there are similar mechanisms for both, then are models able to produce both types of behaviour depending on model parameters. If this is the case is there a sharp transition between the two states or is there a smooth transition leading to a region of parameter space where there is a coexistence of the two states. Also in such a model what do both spatial patterns indicate about the temporal persistence. Do regular patterns only form when the system is under a large amount of stress and equally is a scale-free pattern indicative of a system far from some tipping point. These are the questions that shall be addressed in the coming chapters.

Bibliography

- Andrew J Allstadt, Jonathan A Newman, Jonathan A Walter, G Korniss, and Thomas Caraco. Dispersal limitation and roughening of the ecological interface. *arXiv preprint arXiv:1403.4574*, 2014.
- Per Bak. *How nature works*. Oxford university press Oxford, 1997.
- Per Bak, Chao Tang, and Kurt Wiesenfeld. Self-organized criticality: An explanation of the $1/f$ noise. *Physical review letters*, 59(4):381, 1987.
- Mark D Bertness and Edwin Grosholz. Population dynamics of the ribbed mussel, *geukensia demissa*: the costs and benefits of an aggregated distribution. *Oecologia*, 67(2):192–204, 1985.
- GF Blanchard, DM Paterson, LJ Stal, P Richard, R Galois, V Huet, J Kelly, C Honeywill, J De Brouwer, K Dyer, et al. The effect of geomorphological structures on potential biostabilisation by microphytobenthos on intertidal mudflats. *Continental Shelf Research*, 20(10):1243–1256, 2000.
- Fabio Borgogno, P D’Odorico, Francesco Laio, and Luca Ridolfi. Mathematical models of vegetation pattern formation in ecohydrology. *Reviews of Geophysics*, 47(1), 2009.
- Sergio A Cannas, Diana E Marco, and Marcelo A Montemurro. Long range dispersal and spatial pattern formation in biological invasions. *Mathematical biosciences*, 203(2):155–170, 2006.
- JFC De Brouwer, S Bjelic, EMGT De Deckere, and LJ Stal. Interplay between biology and sedimentology in a mudflat (biezelingse ham, westerschelde, the netherlands). *Continental shelf research*, 20(10):1159–1177, 2000.
- Paolo D’Odorico, Francesco Laio, Amilcare Porporato, Luca Ridolfi, and Nicolas Barbier. Noise-induced vegetation patterns in fire-prone savannas. *Journal of Geophysical Research: Biogeosciences (2005–2012)*, 112(G2), 2007.

- Barbara Drossel and Franz Schwabl. Self-organized critical forest-fire model. *Physical Review Letters*, 69(11):1629, 1992.
- Richard Durrett and Simon Levin. The importance of being discrete (and spatial). *Theoretical population biology*, 46(3):363–394, 1994.
- Kenneth Falconer. *Fractal geometry: mathematical foundations and applications*. John Wiley & Sons, 2013.
- DR Foster, GA King, PH Glaser, and HE Wright. Origin of string patterns in boreal peatlands. 1983.
- Frédéric Guichard, Patti M Halpin, Gary W Allison, Jane Lubchenco, and Bruce A Menge. Mussel disturbance dynamics: signatures of oceanographic forcing from local interactions. *The American Naturalist*, 161(6):889–904, 2003.
- Vishwesh Guttal and Ciriya Jayaprakash. Spatial variance and spatial skewness: leading indicators of regime shifts in spatial ecological systems. *Theoretical Ecology*, 2(1):3–12, 2009.
- JM Halley, S Hartley, AS Kallimanis, WE Kunin, JJ Lennon, and SP Sgardelis. Uses and abuses of fractal methodology in ecology. *Ecology Letters*, 7(3):254–271, 2004.
- Alan Hastings and Derin B Wysham. Regime shifts in ecological systems can occur with no warning. *Ecology Letters*, 13(4):464–472, 2010.
- Harold M Hastings, Richard Pielou, Richard Monticciolo, David Vun Kannon, and Diane Del Monte. Time scales, persistence and patchiness. *BioSystems*, 15(4):281–289, 1982.
- Reinier HilleRisLambers, Max Rietkerk, Frank van den Bosch, Herbert HT Prins, and Hans de Kroon. Vegetation pattern formation in semi-arid grazing systems. *Ecology*, 82(1):50–61, 2001.
- Attila R Imre, Dániel Cseh, Markus Neteler, and Duccio Rocchini. Korcak dimension as a novel indicator of landscape fragmentation and re-forestation. *Ecological indicators*, 11(5):1134–1138, 2011.
- Attila R Imre, Josef Novotný, and Duccio Rocchini. The korcak-exponent: a non-fractal descriptor for landscape patchiness. *Ecological complexity*, 12:70–74, 2012.

- Sonia Kéfi, Max Rietkerk, Concepción L Alados, Yolanda Pueyo, Vasilios P Papanastasis, Ahmed ElAich, and Peter C De Ruiter. Spatial vegetation patterns and imminent desertification in mediterranean arid ecosystems. *Nature*, 449(7159):213–217, 2007a.
- Sonia Kéfi, Max Rietkerk, Minus van Baalen, and Michel Loreau. Local facilitation, bistability and transitions in arid ecosystems. *Theoretical Population Biology*, 71(3):367–379, 2007b.
- Sonia Kéfi, Maarten B Eppinga, Peter C de Ruiter, and Max Rietkerk. Bistability and regular spatial patterns in arid ecosystems. *Theoretical Ecology*, 3(4):257–269, 2010.
- Sonia Kéfi, Max Rietkerk, Manojit Roy, Alain Franc, Peter C De Ruiter, and Mercedes Pascual. Robust scaling in ecosystems and the meltdown of patch size distributions before extinction. *Ecology letters*, 14(1):29–35, 2011.
- Timothy H Keitt. Spectral representation of neutral landscapes. *Landscape Ecology*, 15(5):479–494, 2000.
- Shinya Kizaki and Makoto Katori. Analysis of canopy-gap structures of forests by ising-gibbs states-equilibrium and scaling property of real forests. *Journal of the Physical Society of Japan*, 68(8):2553–2560, 1999.
- Christopher A Klausmeier. Regular and irregular patterns in semiarid vegetation. *Science*, 284(5421):1826–1828, 1999.
- Jaromír Korčák. *Deux types fondamentaux de distribution statistique*. Comité d’organisation, 1938.
- Takuya Kubo, Yoh Iwasa, and Naoki Furumoto. Forest spatial dynamics with gap expansion: total gap area and gap size distribution. *Journal of theoretical Biology*, 180(3):229–246, 1996.
- René Lefever and Olivier Lejeune. On the origin of tiger bush. *Bulletin of Mathematical Biology*, 59(2):263–294, 1997.
- Olivier Lejeune and Mustapha Tlidi. A model for the explanation of vegetation stripes (tiger bush). *Journal of Vegetation Science*, 10(2):201–208, 1999.
- Olivier Lejeune, Mustapha Tlidi, and P Couteron. Localized vegetation patches: a self-organized response to resource scarcity. *Physical Review E*, 66(1):010901, 2002.

- Simon A Levin. *The problem of pattern and scale in ecology*. Springer, 1995.
- Fernando T Maestre and Adrián Escudero. Is the patch size distribution of vegetation a suitable indicator of desertification processes? *Ecology*, 90(7):1729–1735, 2009.
- Benoit B Mandelbrot. *The fractal geometry of nature*, volume 173. Macmillan, 1983.
- Susanna C Manrubia and Ricard V Solé. Self-organized criticality in rainforest dynamics. *Chaos, Solitons & Fractals*, 7(4):523–541, 1996.
- Susannah Mistr and David Bercovici. A theoretical model of pattern formation in coral reefs. *Ecosystems*, 6(1):0061–0074, 2003.
- JD Murray. *Mathematical biology: I. An introduction*. Springer, 2001.
- B Okamura. Group living and the effects of spatial position in aggregations of *mytilus edulis*. *Oecologia*, 69(3):341–347, 1986.
- Mercedes Pascual and Frédéric Guichard. Criticality and disturbance in spatial ecological systems. *Trends in ecology & evolution*, 20(2):88–95, 2005.
- Mercedes Pascual, Manojit Roy, Frédéric Guichard, and Glenn Flierl. Cluster size distributions: signatures of self-organization in spatial ecologies. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 357(1421):657–666, 2002.
- Max Rietkerk and Johan Van de Koppel. Regular pattern formation in real ecosystems. *Trends in Ecology & Evolution*, 23(3):169–175, 2008.
- Max Rietkerk, Maarten C Boerlijst, Frank van Langevelde, Reinier HilleRisLambers, Johan van de Koppel, Lalit Kumar, Herbert HT Prins, and André M de Roos. Self-organization of vegetation in arid ecosystems. *The American Naturalist*, 160(4):524–530, 2002.
- Max Rietkerk, Stefan C Dekker, Peter C de Ruiter, and Johan van de Koppel. Self-organized patchiness and catastrophic shifts in ecosystems. *Science*, 305(5692):1926–1929, 2004.
- Y Sakaguchi. On the genesis of banks and hollows in peat bogs: an explanation by a thatch line theory. *Bulletin of the department of Geography University of Tokyo*, 12:35–58, 1980.

- Todd M Scanlon, Kelly K Caylor, Simon A Levin, and Ignacio Rodriguez-Iturbe. Positive feedbacks promote power-law clustering of kalahari vegetation. *Nature*, 449(7159):209–212, 2007.
- Marten Scheffer, Jordi Bascompte, William A Brock, Victor Brovkin, Stephen R Carpenter, Vasilis Dakos, Hermann Held, Egbert H van Nes, Max Rietkerk, and George Sugihara. Early-warning signals for critical transitions. *Nature*, 461(7260):53–59, 2009.
- István Scheuring and Rudolf H Riedi. Application of multifractals to the analysis of vegetation pattern. *Journal of Vegetation Science*, 5(4):489–496, 1994.
- Robert Schlicht and Yoh Iwasa. Forest gap dynamics and the ising model. *Journal of theoretical biology*, 230(1):65–75, 2004.
- Laurent Seuront. *Fractals and multifractals in ecology and aquatic science*. CRC Press, 2009.
- Jonathan A Sherratt. An analysis of vegetation stripe formation in semi-arid landscapes. *Journal of mathematical biology*, 51(2):183–197, 2005.
- NM Shnerb, P Sarah, H Lavee, and S Solomon. Reactive glass and vegetation patterns. *Physical review letters*, 90(3):038101, 2003.
- Kim Sneppen, Per Bak, Henrik Flyvbjerg, and Mogens H Jensen. Evolution as a self-organized critical phenomenon. *Proceedings of the National Academy of Sciences*, 92(11):5209–5213, 1995.
- Ricard V Solé and Susanna C Manrubia. Are rainforests self-organized in a critical state? *Journal of Theoretical Biology*, 173(1):31–40, 1995a.
- Ricard V Solé and Susanna C Manrubia. Self-similarity in rain forests: evidence for a critical state. *Physical review E*, 51(6):6250, 1995b.
- H Eugene Stanley. Introduction to phase transitions and critical phenomena. *Introduction to Phase Transitions and Critical Phenomena, by H Eugene Stanley, pp. 336. Foreword by H Eugene Stanley. Oxford University Press, Jul 1987. ISBN-10: 0195053168. ISBN-13: 9780195053166*, 1, 1987.
- George Sugihara and Robert May. Applications of fractals in ecology. *Trends in Ecology & Evolution*, 5(3):79–86, 1990.

- David K Swanson and David F Grigal. A simulation model of mire patterning. *Oikos*, pages 309–314, 1988.
- JM Thiery, J-M d’Herbes, and Christian Valentin. A model simulating the genesis of banded vegetation patterns in niger. *Journal of Ecology*, pages 497–507, 1995.
- Alan Mathison Turing. The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 237(641): 37–72, 1952.
- Johan van de Koppel, Max Rietkerk, Norbert Dankers, and Peter MJ Herman. Scale-dependent feedback and regular spatial patterns in young mussel beds. *The American Naturalist*, 165(3):E66–E77, 2005.
- M Zamir. Critique of the test of multifractality as applied to biological data. *Journal of theoretical biology*, 225(3):407–412, 2003.

Chapter 3

Seagrass

Stretch'd on her mossy couch, in trackless deeps,
Queen of the coral groves, *Zostera* sleeps;
The silvery sea-weed matted round her bed,
And distant surges murmuring o'er her head.

(Erasmus Darwin)

3.1 Introduction

Vegetation-based ecosystems, where an ecosystem is described by a single or assemblage of sessile species, represent a large class of vital communities that are found globally. Examples of these ecosystems dominated by an assemblage include wet woodland [Peterken, 1996], temperate bogs [Scheffer et al., 2001], coral reefs [Huston, 1985] and tropical mangrove swamps [Lugo and Snedaker, 1974]. In chapter 2, we considered the variety of pattern formation processes that have been observed and the modelling efforts that have been performed in order to understand the origins of pattern formation in vegetation-based ecosystems. The efforts shall be to understand how pattern and dynamics are linked through modelling work that will be introduced in chapter 5 and further developed in chapters 6 & 7. In order to assess the validity of this modelling work, we shall refer to a test case of a marine vegetation ecosystem. This test case allows us to apply hypotheses directly to a real ecosystem for which we have data in order to test their validity.

The system that shall be considered is Seagrass, which is a class of marine flowering plants that are found globally in coastal areas and contribute a significant number of ecosystem services. In this chapter their biology and ecology will be described with specific reference to how they relate to pattern formation and persistence. A set of field studies taken from seagrass habitats in the Isles of Scilly, UK will then be described along with specific field work carried out to confirm the nature of the observed spatial patterns. The author personally participated in the field study in the summer of 2013 and was involved with data collection, assisting survey work as well as devising and participating the the banding study. These field studies shall be used to test key hypotheses laid out in the introduction on the nature of dynamic and temporal persistence in a vegetation community as well as the various observed spatial patterns taken from an aerial photographic survey conducted in conjunction with the longitudinal study.

3.2 Biology & ecology

Seagrasses represent a broad class of sixty marine flowering plants that are divided into four separate families: *Posidoniaceae*, *Zosteraceae*, *Hydrocharitaceae* and *Cymodoceaceae*. They evolved from terrestrial grasses approximately 100 million years ago to have an entirely aquatic life history and have number of adaptations that make them suitable for aquatic life [Hemminga and Duarte, 2000]. They are both a pioneer and climax species, being able to colonise soft substrate such as sand or mud

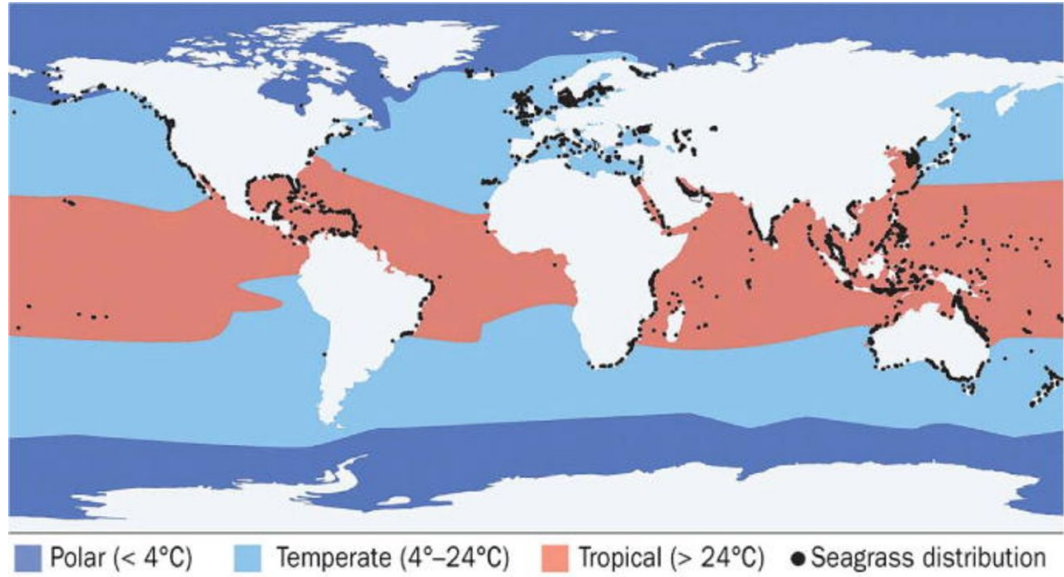


Figure 3.1: The global distribution of Seagrass species. Reproduced from Orth et al. [2006].

and producing a large-scale habitat often referred to as a meadow due to their resemblance to terrestrial grassland. They are environmental engineers, by producing an interwoven root and rhizome structure that is able to retain sediment. Shoots reduce water current allowing the capture of suspended sediment and decreasing the risk of damage due to hydrological scouring. Their presence is also able to modify nutrient concentrations and fundamentally alter the local food web. Although originally thought not to be a food source for many species due to their low nutritional content, it has been discovered that they are a vital food source for tropical herbivores such as sea turtles, dugongs, manatees as well as wild fowl [Thayer et al., 1984]. Further, they provide a habitat for a large number and diversity of species including both recreationally and commercially important fish species [Beck et al., 2001]. The rhizome layer also represents an important form of carbon sequestration, where atmospheric carbon is converted into organic carbon that remains in the seabed [Fourqurean et al., 2012].

Seagrasses are found globally in sheltered coastlines and estuaries from as far south as the southern tip of New Zealand (46°S) to as far north as Iceland and Greenland (66°N) occupying tropical, temperate and polar climates (Fig. 3.1). Their life-cycle is dependent upon local adaptations to climate, though they are often perennial with a stronger growing season in the summer for temperate species and significant die-back of shoots in the winter. Flowering occurs throughout the year and can be

triggered by temperature. Pollination then occurs through water action, but there is evidence to suggest that certain crustacean and fish species could also be involved in the pollination process [Sumoski and Orth, 2012; van Tussenbroek et al., 2012]. This is followed by production and dispersion of seeds that occurs both locally or long-range due to uprooted leaves being suspended in the water column [Kendrick et al., 2012].

Seagrass is under direct threat from anthropogenic factors through industry coastal development projects and land reclamation. These influences have accelerated the losses of seagrass globally and there are many areas where they are under direct threat. For instance, coastline development in Cockburn Sands, Australia lead to large scale losses in the 1990s and density remains low today [Walker et al., 2006]. Losses were also identified in Gibraltar [Bull et al., 2010]. There was a dramatic decline in numbers between 1993 and 2009 to the extent that two species are locally extinct in the area. The cause of this dramatic decline was attributed to high anthropogenic stress due to large developments and land reclamation. It caused a vast amount of silt to be suspended in the surrounding waters, thus destroying the seagrasses' potential to colonise.

Seagrass is also under threat from wasting disease, whose causative agent is a slime-mould protist *Labyrinthula zosterae* [Muehlstein et al., 1991]. A large-scale epidemic in the 1930s severely reduced the number of Atlantic seagrass leading to the collapse of the scallop industry; a dramatic decrease in waterfowl numbers and the extinction of at least one species of marine gastropod [Rasmussen, 1977]. Numbers may not have fully recovered from the epidemic and many open questions remain as to how the spread of the disease occurred and what conditions lead to the rapid spread of the disease [Den Hartog, 1987]. The protist infects leaves of the plant producing dark patches of necrotic tissue. The infection then spreads along the leaf eventually destroying its photosynthetic ability, which can lead to the death of the whole plant. Experiments have suggested that an infection event occurs due to direct leaf-to-leaf contact [Muehlstein et al., 1991], but infected leaves can be severed from the main plant and float, thus facilitating more long-range infection events [Vergeer and Den Hartog, 1991]. It is believed that infection spreads when seagrass is under stress and a healthy plant is not susceptible to infection [Orth et al., 2006]. Temperature is also believed to be associated with an increase in the number of infected plants [Rasmussen, 1977].

Current understanding of seagrass vegetation dynamics comes from broad-scale monitoring of extent often having to span several years and large scales [Waycott et al., 2009]. A number of studies also only consider small plots [Cambridge and McComb, 1984; Deslous-Paoli et al., 2001] or through experimental work by manipulating certain environmental variables that affect growth [Tomasko and Lapointe, 1991]. What makes the Isles of Scilly study unique is the combination of long time scale measurements spanning several sites having their own unique set of conditions with a large scale, high-resolution mapping of the distribution in the whole archipelago. These combine together to form a powerful set of data suitable to test how seagrass dynamics relate to its overall distribution.

3.3 Isles of Scilly



Figure 3.2: Old Grimsby Harbour (2013). A seagrass habitat highly stressed due to anthropogenic disturbance.

The Isles of Scilly is an archipelago of 200 islands approximately 40km of the south west coast of the UK mainland. The cluster of islands are mainly formed of granite with shallow seas in between the islands. The sea bed itself is composed of

a variety of pebbles and sand with a size distribution that changes from the more sheltered waters in the East to the more exposed waters in the West. The Isles are composed of five main islands: St. Mary's the largest island, which is also the location of the largest settlement in the archipelago, St. Martin's, Treco, Bryher and St. Agnes. The islands are home to various bays used for recreational boat mooring and there are also farming, fishing and tourist activities on the islands. Despite this there is relatively little agriculture and industry on the islands. This combined with a strong tidal action means that there is relatively little pollution making the site an ideal place to study the vegetation in isolation of other effects. Local exposure varies throughout the area, with the most extreme areas being along the Western islands, where there is a rocky shoreline and strong wave action from the open ocean. The eastern islands, which have a large amount of sediment deposition, are more protected.

The isles are the habitat of a wide-variety of flora and fauna with some unique species not observed on the UK mainland and are also home to a large colony of Atlantic grey seal (Fig. 3.7). There are three species of seagrass found in the waters around the Isles of Scilly: *Zostera noltii*, *Zostera angustifolia* and *Zostera marina* also known as eelgrass due to the elongated shape of its leaves. Annual surveying of eelgrass in Scilly goes as far back as 1984. However, it was only until 1992 when large scale surveying financed by Natural England began. This was in response to the appearance of wasting disease in the early 1990s, which had previously been absent in the waters [Fowler, 1992]. Initially, only two sites were chosen to be surveyed: Higher Town Bay and Old Grimsby Harbour. However, since then the sites of West Broad Ledges, Broad Ledges Treco and then Little Arthur in the Eastern Isles were added to the list of sites annually surveyed. Apart from two points, all five sites have been monitored every year from 1992 to the latest survey, which occurred in July 2014.

3.4 Mapping

The spatial section of the data comes from an aerial survey conducted on behalf of Natural England [Jackson et al., 2011]. The main purpose of the project was to use aerial photography and Geographic Information System analysis to map the distribution of eelgrass in the Isles of Scilly. The study was split into two phases. The first being the acquisition of both RGB and infra-red digital aerial photographs to a $10 \times 10\text{cm}^2$ resolution. Unsupervised learning was then conducted on the resulting

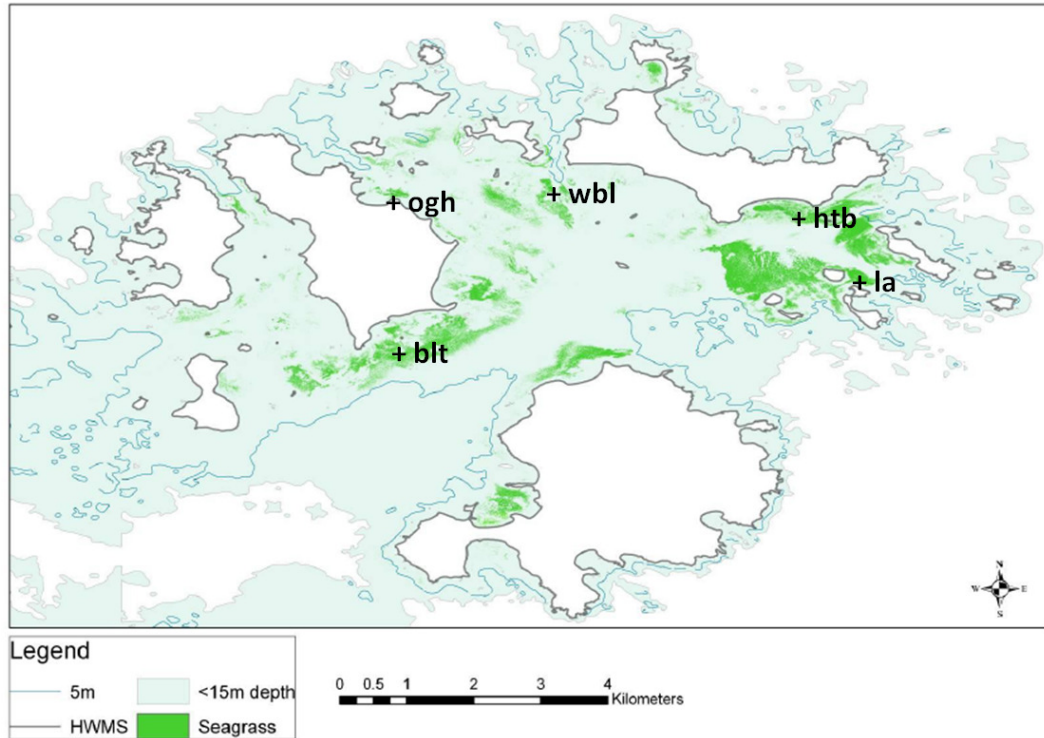


Figure 3.3: Final distribution of eelgrass produced from the 2008 mapping study. The five sites that have been regularly surveyed by dive teams are also highlighted with key: *blt*- Broad Ledges Tresco; *htb*- Higher Town Bay; *la*- Little Arthur; *ogh*- Old Grimsby Harbour; *wbl*- West Broad Ledges. Reproduced from Jackson et al. [2011].

photographs to identify a number of classes that could then undergo ground truthing for the second phase. An aerial photographic survey using a 90 Mega-pixel camera took place on 26th September between 0925 and 1155 BST. The survey took place when there was little cloud cover and eelgrass coverage was at its highest (eelgrass shoot density has strong seasonality and there is a significant die-back in winter). Since eelgrass is sub-tidal and typically resides at depths greater than one metre, the infra-red data could not be utilised as infra-red is only able to penetrate a few cm of water.

The resulting RGB images form a number of overlapping mosaics that were spliced together to form one single continuous image. The Iterative Self-Organising Data Analysis Technique (ISODATA), which is an unsupervised classification algorithm, was then performed on the image set to classify each pixel into one of thirty categories. With these classifications, ground truthing occurred in the summer of 2009

to provide a data set for supervised learning to determine which pixels were occupied with eelgrass. Each chosen site was ground truthed either visually for shallow waters or by using an underwater video camera for deeper sites. It was noted where seagrass was sighted along with the percentage cover. The information from the ground truth survey together with records of the distribution of eelgrass was then used to find training areas that were suitable for classification. Adopting a maximum likelihood technique, each pixel was assigned a classification using a probability density function based on the training data set. This results in a distribution of the presence or absence of eelgrass based upon the supervised learning algorithm.

The final stage was to validate whether the pixels were correctly classified and changing the classification of pixels that were incorrect. Some places, such as the north coasts, were removed as eelgrass is unlikely to exist there due to the presence of rocky substrate and exposure. Pixels that were below the 5m in sea depth were also removed as local knowledge and expert opinion consented that eelgrass is unable to survive in the waters around the Isles of Scilly below 5m due to the lack of light. A Kappa statistic was then used to assess the agreement between classifications. The kappa statistic is defined as

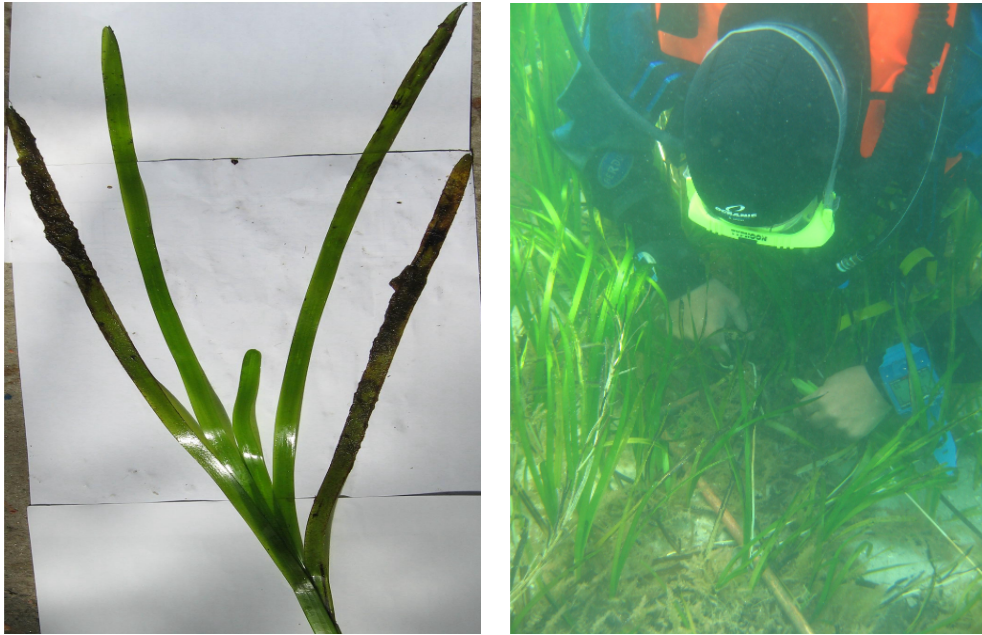
$$\kappa = \frac{P(A) - P(E)}{1 - P(E)}, \quad (3.1)$$

where $P(A)$ is the proportion of correctly classified pixels and $P(E)$ is the proportion of correctly classified pixels when classification occurred at random. The lowest kappa statistics were found on the north and south coasts of St. Martin's. The highest were observed at St. Mary's harbour, the area between Tresco and St. Martin's and the Eastern Isles.

The final stage was to apply a high-pass filter to the digital image to remove isolated pixels classified as eelgrass, based on the fact that it is highly unlikely to have an isolated $10 \times 10\text{cm}^2$ patch of eelgrass. The resulting smoothed map was then converted to a polygon table.

3.5 Survey

The five sites surveyed are Old Grimsby Harbour, Tresco ($49^\circ 57.611'N$ $06^\circ 19.784'W$) ; Higher Town Bay, St. Martin's ($49^\circ 57.428'N$ $06^\circ 16.448'W$) ; Broad Ledge, Tresco ($49^\circ 56.327'N$ $06^\circ 19.773'W$) ; West Broad Ledge, St. Martin's ($49^\circ 57.418'N$ $06^\circ 18.264'W$)



(a) An example of eelgrass shoots taken from survey. Both healthy and diseased leaves are present along with epiphyte coverage on several leaves. Reproduced from Cook and Paver [2007]

(b) Surveying of eelgrass using quadrat central datum line method. Reproduced from Cook and Paver [2007]

Figure 3.4: Isles of Scilly surveying methodology

and Little Arthur, Eastern Isles ($49^{\circ}56.961'N$ $06^{\circ}15.932'W$). Surveys performed annually during summer either in the last week of July or the first week of August using a central datum marker method. At each site, a central line is lowered to the sea floor to be used as a central marker. The central marker varied from year to year due to the boat's anchorage point being located over barren sand so as not to disturb the seagrass habitat. Pairs of divers then took random bearings and distances between 0 and 30m from the central line. These bearings are measured using a compass and the distance measured using a tape measure. Having moved to the required position, a $0.5 \times 0.5m^2$ quadrat is placed onto the seabed. After the quadrat is properly placed, all leaves within the bottom left of the quadrat were cut above the rhizome in order for the plant to be able to survive with the rhizomal mass intact (Fig. 3.4b). These leaves were then bagged and marked for identification once ashore.

For each chosen quadrat site the bearing and distance, number of leaves, maximum leaf length, amount of wasting disease cover, epiphyte cover, number of flowering

plants and number of eggs on leaves were all recorded on land. The wasting disease coverage (Fig. 3.4a) was measured by eye by observing the amount of blackened tissue on the leaf and noting the percentage cover [Burdick et al., 1993]. The resulting density estimates for the five meadows was produced by combining the total number of leaves taken for all the randomly selected quadrat sites. Disease prevalence at each of the meadows was also calculated by taking the average disease leaf coverage over all quadrat sites. The resulting eelgrass densities and disease prevalence are shown in Fig. 3.5.

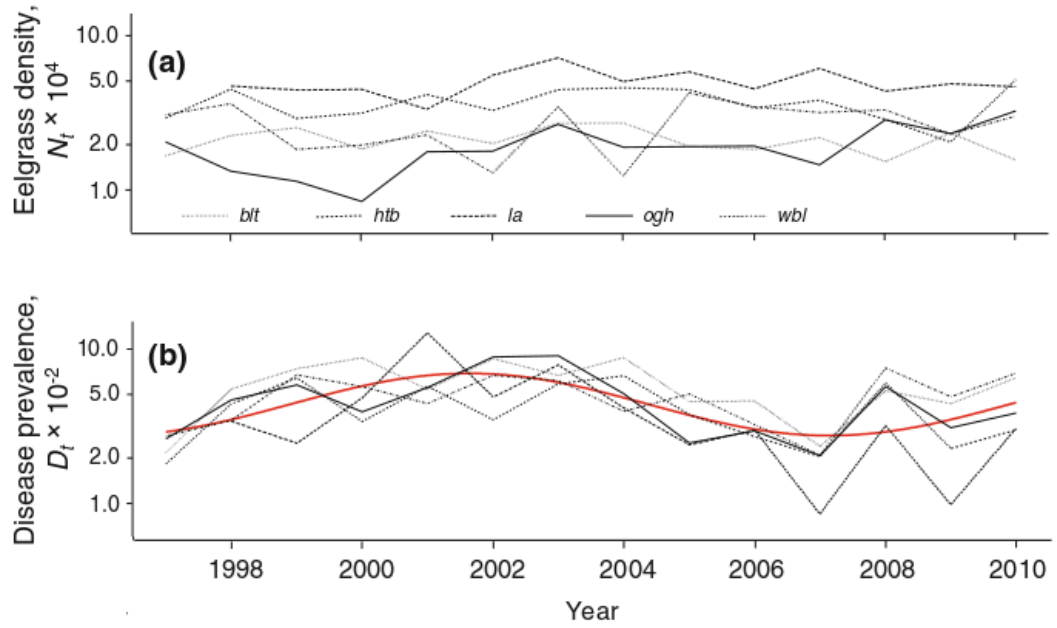
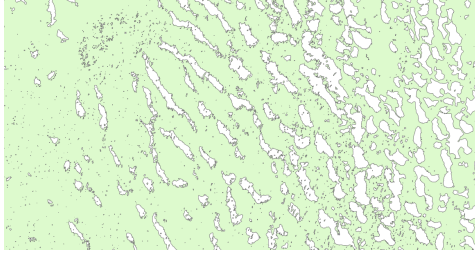


Figure 3.5: Data produced from dive survey for 1998-2008. (a) The overall densities of eelgrass for each site. (b) Disease coverage with a fitted spline shown in red. Reproduced from Bull et al. [2012]

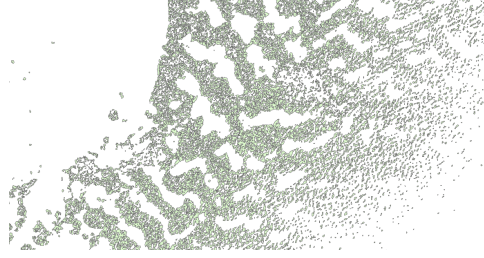
3.5.1 Banding survey

During the survey in 2013, a study was conducted first to determine the validity of banding patterns observed in the data. Banding sites were also assessed for their vegetation distribution, including whether the edges were sharply defined going from high density vegetation to none and also how the substrate co-occurred with the presence of the bands.

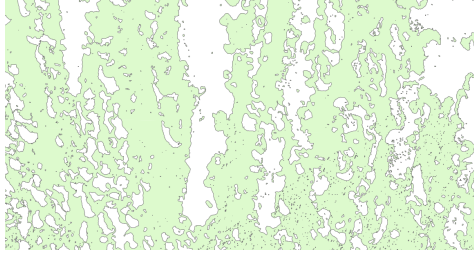
The study was performed by isolating three separate areas where banding was present (Fig. 3.6). These sites were chosen by inspection such that they contained



(a) Location one: wbl. In Latitude Longitude degrees decimal minutes $49^{\circ}57.5'N, 06^{\circ}18.4'W$. Scale is 1:1,000. width of band is approximately 8 metres



(b) Location two: blt. In Lat Long, Degrees minutes and seconds $49^{\circ}56.1'N 6^{\circ}19.79'W$. Scale is 1:1,000. width of band is approximately 16 metres



(c) Location three: South West of htb In Lat Long, Degrees minutes and seconds $49^{\circ}57.0'N 6^{\circ}16.6'W$. Scale is 1:1,000. width of band spacing is approximately 16 metres and width of band approximately 24 metres.

Figure 3.6: Sites identified from aerial survey data used to investigate banding phenomena.

banding and were close to the locations of the dive sites to allow easy investigation of the site. The first location is at West Broad Ledges ($49^{\circ}57.5'N, 06^{\circ}18.4'W$). Vegetation covers the extent of the site with regular interspersed gaps, with the width of the banding being approximately 8 metres. The gaps however, are not continuous, and have smaller regular gaps towards the eastern edge of the site. The second location is at Broad Ledges Tresco ($49^{\circ}56.1'N 6^{\circ}19.79'W$). Bands are larger compared to the first site, with width approximately 16 metres. The vegetation site is more patchy (being composed of several patches that interlock together to form the global banded structure). There is also a gradient in the vegetation coverage from west to east. The third location is south west of Higher Town Bay ($49^{\circ}57.0'N 6^{\circ}16.6'W$). Banding is less apparent in this site, but there is still an anisotropic distribution of vegetation running west-east. Bands are not regular in size even though gap spacing of the larger band is approximately 16 metres and band width is approximately 24 metres. It is apparent from these three sites that anisotropic pattern observed in

the eelgrass distribution is not dominated by a single scale, but rather changes from site to site.

Due to time-constraints, only one site was studied for the presence of banding. The study was conducted by starting from the centre co-ordinates of the site and moving in a transect perpendicular to the banding. The eelgrass was observed both by visual inspection on the side of the boat and through the use of underwater cameras. Further to the transect observations, diving took place to observe the banding edges.

Regular elongated patches of eelgrass, interrupted by barren sand, were seen along the transect. The edges of the bands were seen to be distinct and vegetation within the band appeared continuous at a similar density throughout. Diving confirmed the distinct transition between bands of vegetation and barren sand. A raise in the sea bed from where vegetation was present was also noted. This confirms the pattern detected in the digital image and provides evidence towards the validity of observed banding patterns in sub-tidal seagrass.

3.6 Modelling

Seagrass is a marine flowering plant that can either exist inter-tidally as in the case of *Z. noltii* or sub-tidally as is the case with *Z. marina*. Species are found globally across coastlines and provide an ecosystem vital to a large number of important species [Larkum et al., 2006]. It is estimated that seagrass ecosystems are comparable with coral reefs in their ecosystem services [Costanza et al., 1997], however their numbers have been rapidly declining in recent years [Orth et al., 2006]. Due to the fact that seagrass exists as a monoculture, it is an ideal candidate to study the interaction between process and pattern. It can also be used as a system to test hypotheses associated with inferring dynamics from spatial pattern.

Regular pattern formation has been exhibited in seagrass vegetation. There is a strong association between the shear velocity profile of current and the resulting localised density of seagrass [Fonseca et al., 1983]. Meadows in the presence of low tidal regimes were found to be homogeneous, whereas in the presence of high currents, seagrass is spatially organised into elongated patches. The interaction between sediment and seagrass vegetation was found to also be strongly linked to the presence of bands in strong currents, where strong currents raised the surface bed

around seagrass vegetation. Frederiksen et al. [2004] performed a study by analysing a series of aerial photographs of the Wadden Sea with different physical factors. They found that the presence of elongated patches or bands was associated with a strong presence of wave dynamics. Marba and Duarte [1995] studied the interplay of sand dunes and seagrasses and found there was a strong cross-correlation between the height of sediment and dynamics of the elongated patch. More recently, the interaction between seagrass and soft sediment in the presence of wave action has been explicitly modelled. The modelling produces bands of seagrass, whose spatial structure are in agreement with data collected from an inter-tidal species of seagrass whose environment is dominated by wave action [Van Der Heide et al., 2010].

More explicit modelling work has taken place to identify the characteristic growth properties of a growing cluster of seagrass [Sintes et al., 2005, 2006]. Seagrass propagates as rhizomes that branch off from shoots, where further shoots form from the rhizomal layer. The model explicitly involving these two interactions found that there were two phases to the growth pattern. The geometry of the boundary of the growing cluster changed from being strongly asymmetric with rough boundaries in the early time period, to a smooth disc like growth phase. The model does not explicitly take into account interaction with the environment and other individuals of the same species, hence it is difficult to conclude how these growth properties might be interpreted in situ.

Finally fractal theory has been used to some extent in seagrass ecology [Cunha and Santos, 2009]. The impact of changes of a coastal inlet on the distribution of seagrass was assessed by estimating the Korcak dimension for a number of time-points. Cunha and Santos [2009] used the Hurst exponent to characterise the persistence of the spatial pattern, however there was no accompanying work performed on estimating the temporal persistence of the seagrass meadow and only the spatial persistence is mentioned. This leaves the question as to whether or not temporal persistence is measurable from the fractal characteristics of a spatial distribution.

3.7 Conclusion

Seagrass represents a diverse class of marine vegetation involved with a number of important coastal and ecological processes. These include acting as a nursery habitat for fish; carbon sequestration and reducing coastal erosion. There are a number of interesting aspects to seagrass dynamics that make it an interesting test case to



Figure 3.7: A grey seal belonging to a colony that inhabit the Isles of Scilly. Photo courtesy of T. Reid

study. Firstly, it acts as an ecosystem engineer, both raising the seabed by capturing suspended sediment and reducing the water flow around the bed. Thus on a local-scale there are a number of positive feedback loops present. Seagrass reproduces both sexually and asexually. Asexual reproduction occurs through growth of rhizomes and is primarily local, whilst sexual reproduction occurs through production of seeds, which can be carried on currents and thus can be dispersed over large ranges.

The Isles of Scilly eelgrass represents a distinct monoculture that exists in isolation of the mainland. That combined with the lack of large-scale agriculture and industry on the islands makes the eelgrass dataset an ideal case to study the interaction vegetation and its environment in isolation. Further the study combines a large-scale mapping project of seagrass vegetation to very high resolution ($0.1 \times 0.1\text{m}^2$) with a longitudinal survey spanning over two decades and across several distinct meadows each with a unique combination of exposure, wave action, aspect and anthropogenic influences. These two studies combined represent a unique dataset that where both spatial scale and temporal scale can be investigated in terms of the underlying persistence of the vegetation. Of particular interest is the relationship between dynamic and spatial persistence, which shall be investigated using these dataset combined with theory and modelling work in chapter 6.

Banding field work was carried out based upon results shown in the mapping data, which identified areas of vegetation displaying a high amount of anisotropy. The field work confirmed the presence of banding for one of these sites. The bands were shown to be distinct patches of continuous vegetation with a sharp contrast to uncolonised sand, the raising of the seabed was also observed at the edge of the vegetation patch. This banding phenomena forms the focus of chapters 4 & 5 where both the spatial properties and the origin of the banding shall be investigated.

Bibliography

- Michael W Beck, Kenneth L Heck Jr, Kenneth W Able, Daniel L Childers, David B Eggleston, Bronwyn M Gillanders, Benjamin Halpern, Cynthia G Hays, Kaho Hoshino, Thomas J Minello, et al. The identification, conservation, and management of estuarine and marine nurseries for fish and invertebrates: A better understanding of the habitats that serve as nurseries for marine species and the factors that create site-specific variability in nursery quality will improve conservation and management of these areas. *Bioscience*, 51(8):633–641, 2001.
- James C Bull, Emma J Kenyon, Danny Edmunds, and Kevan J Cook. Recent loss of gibraltar seagrasses. *Botanica Marina*, 53(1):89–91, 2010.
- James C Bull, Emma J Kenyon, and Kevan J Cook. Wasting disease regulates long-term population dynamics in a threatened seagrass. *Oecologia*, 169(1):135–142, 2012.
- David M Burdick, Frederick T Short, and Jaimie Wolf. An index to assess and monitor the progression of wasting disease in eelgrass *zostera marina*. *Marine Ecology-Progress Series*, 94:83–83, 1993.
- ML Cambridge and AJ McComb. The loss of seagrasses in cockburn sound, western australia. i. the time course and magnitude of seagrass decline in relation to industrial development. *Aquatic Botany*, 20(3):229–243, 1984.
- K Cook and LFC Paver. Report on 2007 isles of scilly *zostera marina* survey. Technical report, 2007.
- R Costanza, R dArge, R deGroot, S Farber, M Grasso, B Hannon, K Limburg, S Naeem, RV ONeill, J Paruelo, RG Raskin, P Sutton, and M vandenBelt. The value of the world’s ecosystem services and natural capital. *Nature*, 387(6630):253–260, 1997.

- AH Cunha and RP Santos. The use of fractal geometry to determine the impact of inlet migration on the dynamics of a seagrass landscape. *Estuarine, Coastal and Shelf Science*, 84(4):584–590, 2009.
- C Den Hartog. wasting disease and other dynamic phenomena in *Zostera* beds. *Aquatic Botany*, 27(1):3–14, 1987.
- Jean-Marc Deslous-Paoli, Isabelle Auby, and Françoise Dagault. Factors influencing primary production of seagrass beds (*Zostera noltii* hornem.) in the thau lagoon (french mediterranean coast). *Journal of Experimental Marine Biology and Ecology*, 259(1):63–84, 2001.
- Mark S Fonseca, Joseph C Zieman, Gordon W Thayer, and John S Fisher. The role of current velocity in structuring eelgrass (*Zostera marina* l.) meadows. *Estuarine, Coastal and Shelf Science*, 17(4):367–380, 1983.
- James W Fourqurean, Carlos M Duarte, Hilary Kennedy, Núria Marbà, Marianne Holmer, Miguel Angel Mateo, Eugenia T Apostolaki, Gary A Kendrick, Dorte Krause-Jensen, Karen J McGlathery, et al. Seagrass ecosystems as a globally significant carbon stock. *Nature Geoscience*, 5(7):505–509, 2012.
- SL Fowler. Marine monitoring in the isles of scilly. report to natural england, truro. Technical report, Natural England, 1992.
- Morten Frederiksen, Dorte Krause-Jensen, Marianne Holmer, and Jens Sund Laursen. Spatial and temporal variation in eelgrass (*Zostera marina*) landscapes: influence of physical setting. *Aquatic Botany*, 78(2):147–165, 2004.
- Marten A Hemminga and Carlos M Duarte. *Seagrass ecology*. Cambridge University Press, 2000.
- MA Huston. Patterns of species diversity on coral reefs. *Annual Review of Ecology and Systematics*, pages 149–177, 1985.
- EL Jackson, S Higgs, T Allsop, A Cathray, J Evans, et al. Isles of scilly seagrass mapping. *Natural England Commissioned Reports*, (087), 2011.
- Gary A Kendrick, Michelle Waycott, Tim JB Carruthers, Marion L Cambridge, Renae Hovey, Siegfried L Krauss, Paul S Lavery, Donald H Les, Ryan J Lowe, Oriol Mascaró i Vidal, et al. The central role of dispersal in the maintenance and persistence of seagrass populations. *BioScience*, 62(1):56–65, 2012.

- Anthony WD Larkum, Robert Joseph Orth, Carlos M Duarte, et al. *Seagrasses: biology, ecology, and conservation*. Springer, 2006.
- Ariel E Lugo and Samuel C Snedaker. The ecology of mangroves. *Annual review of ecology and systematics*, pages 39–64, 1974.
- Nuria Marba and Carlos M Duarte. Coupling of seagrass (*cymodocea nodosa*) patch dynamics to subaqueous dune migration. *Journal of Ecology*, pages 381–389, 1995.
- Lisa K Muehlstein, David Porter, and Frederick T Short. *Labyrinthula zosterae* sp. nov., the causative agent of wasting disease of eelgrass, *zostera marina*. *Mycologia*, pages 180–191, 1991.
- Robert J Orth, Tim JB Carruthers, William C Dennison, Carlos M Duarte, James W Fourqurean, Kenneth L Heck, A Randall Hughes, Gary A Kendrick, W Judson Kenworthy, Suzanne Olyarnik, et al. A global crisis for seagrass ecosystems. *Bioscience*, 56(12):987–996, 2006.
- George F Peterken. *Natural woodland: ecology and conservation in northern temperate regions*. Cambridge University Press, 1996.
- Erik Rasmussen. The wasting disease of eelgrass (*zostera marina*) and its effects on environmental factors and fauna. *Seagrass ecosystems—a scientific perspective*, pages 1–15, 1977.
- RA Scheffer, RS Van Logtestijn, and JTA Verhoeven. Decomposition of *carex* and *sphagnum* litter in two mesotrophic fens differing in dominant plant species. *Oikos*, 92(1):44–54, 2001.
- Tomas Sintes, Núria Marba, Carlos M Duarte, and Gary A Kendrick. Nonlinear processes in seagrass colonisation explained by simple clonal growth rules. *Oikos*, 108(1):165–175, 2005.
- Tomàs Sintes, Núria Marbà, and Carlos M Duarte. Modeling nonlinear seagrass clonal growth: assessing the efficiency of space occupation across the seagrass flora. *Estuaries and coasts*, 29(1):72–80, 2006.
- Sarah E Sumoski and Robert J Orth. Biotic dispersal in eelgrass *zostera marina*. *Mar. Ecol. Prog. Ser.*, 471:1–10, 2012.
- Gordon W Thayer, Karen A Bjorndal, John C Ogden, Susan L Williams, and Joseph C Zieman. Role of larger herbivores in seagrass communities. *Estuaries*, 7(4):351–376, 1984.

- DA Tomasko and BE Lapointe. Productivity and biomass of *Thalassia testudinum* as related to water column nutrient availability and epiphyte levels: field observations and experimental studies. *Marine ecology progress series. Oldendorf*, 75 (1):9–17, 1991.
- Tjisse Van Der Heide, Tjeerd J Bouma, Egbert H Van Nes, Johan Van De Koppel, Marten Scheffer, Jan GM Roelofs, Marieke M Van Katwijk, and Alfons JP Smolders. Spatial self-organized patterning in seagrasses along a depth gradient of an intertidal ecosystem. *Ecology*, 91(2):362–369, 2010.
- Brigitta I van Tussenbroek, L Veronica Monroy-Velazquez, and Vivianne Solis-Weiss. Meso-fauna foraging on seagrass pollen may serve in marine zoophilous pollination. *Marine Ecology Progress Series*, 469:1, 2012.
- LHT Vergeer and C Den Hartog. Occurrence of wasting disease in *zostera noltii*. *Aquatic Botany*, 40(2):155–163, 1991.
- Diana I Walker, Gary A Kendrick, and Arthur J McComb. Decline and recovery of seagrass ecosystemsthe dynamics of change. In *Seagrasses: biology, ecology and conservation*, pages 551–565. Springer, 2006.
- Michelle Waycott, Carlos M Duarte, Tim JB Carruthers, Robert J Orth, William C Dennison, Suzanne Olyarnik, Ainsley Calladine, James W Fourqurean, Kenneth L Heck, A Randall Hughes, et al. Accelerating loss of seagrasses across the globe threatens coastal ecosystems. *Proceedings of the National Academy of Sciences*, 106(30):12377–12381, 2009.

Chapter 4

Spatial analysis

He had bought a large map representing the sea,
Without the least vestige of land:
And the crew were much pleased when they found it to be
A map they could all understand.

(Lewis Carroll - *The Hunting of the Snark*)

4.1 Introduction

In order to determine the spatial characteristics that are relevant for understanding pattern formation and persistence in plant Ecology we must first consider the number of techniques used to characterise spatial patterns in Ecology and other Life Sciences. This includes determining their applicability, various merits and disadvantages for use on our dataset.

To infer properties of spatial pattern and how they relate to underlying dynamics, which we shall come on to in later chapters, we must first determine the spatial properties of the vegetative patterns in a statistically robust manner. We will then concentrate our analysis on the seagrass dataset collected from the Isles of Scilly, UK, although the analysis would be equally applicable to other forms of vegetative pattern. The Scilly dataset was collected from an aerial photographic survey conducted in August, 2008 and prepared to produce a binary dataset of Seagrass patches. Using standard techniques, locations of seagrass occupation were resolved to a 10cm^2 grid size (See chapter 3). Five distinct seagrass meadows were sampled from this dataset in order to compare to the time-series data. GPS positions of these meadows are: Broad Ledges, Tresco (*blt*: $49^\circ 56.4'N$, $06^\circ 19.6'W$); Higher Town Bay, St. Martin's (*htb*: $49^\circ 57.2'N$, $06^\circ 16.6'W$); Little Arthur (*la*: $49^\circ 56.9'N$, $06^\circ 15.9'W$, depth: 1.0m); Old Grimsby Harbour, Tresco (*ogh*: $49^\circ 57.6'N$, $06^\circ 19.8'W$); and West Broad Ledges (*wbl*: $49^\circ 57.5'N$, $06^\circ 18.4'W$). We used a 700m bounding box around the centre of each of these sites to produce five datasets. These datasets are explored in this chapter using a variety of scaling and anisotropy statistics. Initially, we consider the monofractal scaling properties of the datasets. This then leads on to a discussion of the multifractal properties as well as the lacunarity. Finally, we consider the anisotropic properties of the whole Scilly dataset.

4.2 Monofractal

To study the fractal properties of a set we must first define what a fractal is. This is done by considering the Hausdorff dimension. Although the definition is rather technical, we shall only consider it in a theoretical context and introduce the box-counting dimension afterwards, which gives a much more intuitive sense of dimension. We should briefly here discuss an intuitive idea of what dimension is. For simple Euclidean objects the idea of dimension seems obvious and we can appeal to the linear algebra definition for a dimension as the size of the basis set i.e. how many co-ordinates you would need to describe a position in a set. For example a circle is

a one-dimensional object as only one co-ordinate, namely the angle θ is needed to describe the position.

To make the ideas of dimension more explicit consider the prototypical example of fractal: the Cantor set. First, consider the unit interval $[0, 1]$. The construction is iterative; in the first step the middle third is removed leaving a disjoint union of two intervals $[0, \frac{1}{3}] \cup [\frac{2}{3}, 1]$. At this stage we have two identical copies of the original set, the procedure continues by removing the middle third of each of the intervals, leaving $[0, 1/9] \cup [2/9, 1/3] \cup [2/3, 7/9] \cup [8/9, 1]$. This iteration continues to infinity leaving the resulting Cantor set. Now we may ask the question how many co-ordinates are needed in order to describe the set, this can be done by considering how it was constructed. Consider a point in the Cantor set, at each stage of the construction there is a choice over whether the point is in the left third or the right third of the set being divided, which can be represented as a co-ordinate that takes the value 0 or 1. Hence a point that has the co-ordinates $(0, 1, 1 \dots)$ will be in the interval $[8/27, 1/3]$. The problem then comes as to describe a point in the Cantor set it appears you need an infinite number of co-ordinates. On the other hand, we know that the Cantor set is a subset of the unit interval which has dimension one. For a definition of dimension to hold we would require that if the sets $E \subset F \subset G$ then the dimension of each satisfies $\dim(E) \leq \dim(F) \leq \dim(G)$. We would also desire a definition of dimension to coincide with the dimension of an Euclidean space. In the next section we will explore two such dimensions: The Hausdorff dimension and the Box-Counting dimension.

The idea behind the Hausdorff Dimension is to consider various coverings for a set and consider the sum of the size of these coverings raised to a power s . The sum is minimised over all coverings of size δ . As the size of the coverings reduces as the sum increases, they reach a supremum as $\delta \rightarrow 0$. This provides the Hausdorff measure. More formally, let F be a set with a δ -covering $\{U_i\}$ i.e. $|U_i| < \delta$ for all i and $F \subset \cup_{i=0}^{\infty} U_i$, then the s -dimensional Hausdorff measure is defined as

$$\mathcal{H}^s(F) = \liminf_{\delta \rightarrow 0} \left\{ \sum_{i=0}^{\infty} |U_i|^s : \{U_i\} \text{ is a } \delta\text{-cover of } F \right\}. \quad (4.1)$$

The Hausdorff dimension is defined from the s -dimensional Hausdorff measure as

$$\dim_H F = \inf \{s \geq 0 : \mathcal{H}^s(F) = 0\}. \quad (4.2)$$

As an example of a calculation of the Hausdorff dimension consider the unit interval $[0, 1]$. A minimal δ -covering for this set would be $n = \lceil 1/\delta \rceil$. Hence,

$$\begin{aligned} \sum_{i=0}^n |U_i|^s &= \sum_{i=0}^{\lceil 1/\delta \rceil} |U_i|^s \\ &= \sum_{i=0}^{\lceil 1/\delta \rceil} \delta^s = \delta^{s-1}, \end{aligned} \quad (4.3)$$

where we assume that $1/\delta$ has an integer value. As $\delta \rightarrow 0$ we can see that

$$\mathcal{H}^s(F) = \begin{cases} \infty & : s < 1 \\ 0 & : s > 1 \end{cases}. \quad (4.4)$$

The corresponding Hausdorff dimension is calculated directly from definition as $\dim_H[0, 1] = 1$. In general, the Hausdorff dimension need not take an integer value, as for the Cantor set the Hausdorff dimension is $\log(3)/\log(2)$ [Falconer, 2013].

Although the Hausdorff dimension is appealing from a theoretical point of view, it is in general intractable to calculate for real data. The box-counting or Minkowski-Bouligand dimension is defined on a set F embedded in a Euclidean space \mathbb{R}^n . For our purposes we should only concern ourselves with sets embedded in \mathbb{R}^2 , since we are considering vegetation occupancy patterns. The box-counting dimension is calculated via the box-counting algorithm. Boxes of length ϵ are placed in a regular grid fashion over the space. The number of boxes of length ϵ that intersect the set F is denoted $N(\epsilon)$. The box-counting dimension is then defined as

$$\dim_{\text{BC}}(F) := \lim_{\epsilon \rightarrow 0} \frac{\log N(\epsilon)}{\log(1/\epsilon)}. \quad (4.5)$$

For instance for the Cantor set, if we take a box size of length $(1/3)^n$, then the number of boxes that are filled are 2^n . It is then straightforward to calculate $\dim_{\text{BC}}(F)$ as

$$\begin{aligned} \dim_{\text{BC}}(\text{Cantor set}) &= \lim_{n \rightarrow \infty} \frac{\log(2^n)}{\log(3^n)} \\ &= \lim_{n \rightarrow \infty} \frac{n \log(2)}{n \log(3)} = \frac{\log(2)}{\log(3)}, \end{aligned} \quad (4.6)$$

which is the same value as the Hausdorff dimension. The Hausdorff dimension often gives an equivalent value as the box-counting dimension, however they are not

the same and in general the box-counting dimension is greater than or equal to the Hausdorff dimension.

There are various definitions of a fractal set. It is often defined as a set that roughly repeats itself on finer scales [Gisiger, 2001; Hastings and Sugihara, 1993]. Mandelbrot [1983] gave the definition of a fractal as a set whose Hausdorff Dimension strictly exceeds its Topological Dimension. For our purposes, we define a set to be fractal if the number of non-empty boxes of length ϵ scales as a power-law over some range of ϵ . This does mean that certain trivial sets such as the empty set or the spatial Poisson Process would be classed as fractals under this definition. We shouldn't necessarily be concerned with the definition of fractal and non-fractal too much, as we are more focused on being able to robustly measure scaling properties such as the box-counting dimension. A set that we class as non-fractal is then one where the developed fractal analyses such as the box-counting dimension are not applicable.

4.2.1 Measuring monofractality from data

In the previous section, we have concerned ourselves with what theoretical measurements we wish to apply to data in order to detect the underlying scaling properties. This lead to establishing the box-counting dimension as a measure that can be efficiently calculated and easily applied to several different types of data including time-series and occupancy. The standard method of calculating the box-counting dimension is to measure the number of boxes of length ϵ that are occupied and plot it against $1/\epsilon$ on a log-log plot. The relationship between the two for a fractal is linear on the plot and the gradient is taken to be the box-counting dimension. For real data, there is an issue that the same type of scaling might not be present over all ranges being considered. As an example, over smaller spatial scales seagrass is buffeted by small currents and wave action and is also undergoing turnover in the form of clonal and seed spreading. Over larger ranges, patches form and these coagulate into meadows. On even larger scales, the vegetation is affected by large-scale currents and geographic features such as coastline and coastal shelf. It is therefore not expected that there should be a similar scaling throughout the entirety of the spatial scales. In order to deal with this problem, a number of ideas can be proposed. The first is to use splines to fit to the log-log plot in order to detect different forms of scaling. There are a few issues with this: firstly, there may be a smooth transition between one form of scaling and another making it difficult to find the optimal scale to place the spline; secondly, it may not be obvious how many forms

of scaling and thus how many splines should be used to fit to the data.

The second method is to find an appropriate range of scales over which the power law relationship does hold for a single dimension. It is this method that shall be employed to analyse the Isles of Scilly data. The method we use is outlined in Seuront [2009]; for all ranges of $\log(\delta)$ where the number of δ points is 5, fit a line by linear regression and calculate the coefficient of determination r^2 and the sum of squared residuals (SSR). Then, find the range that minimises the SSR and see if this range corresponds with the range that maximises the r^2 . If it does, then use the gradient calculated by linear regression from this range to determine the box-counting dimension. If the ranges do not match up, there is no best-fit for any range and we say that the data is non-fractal as it has no range over which the scaling is constant.

4.3 Fractional Brownian motion

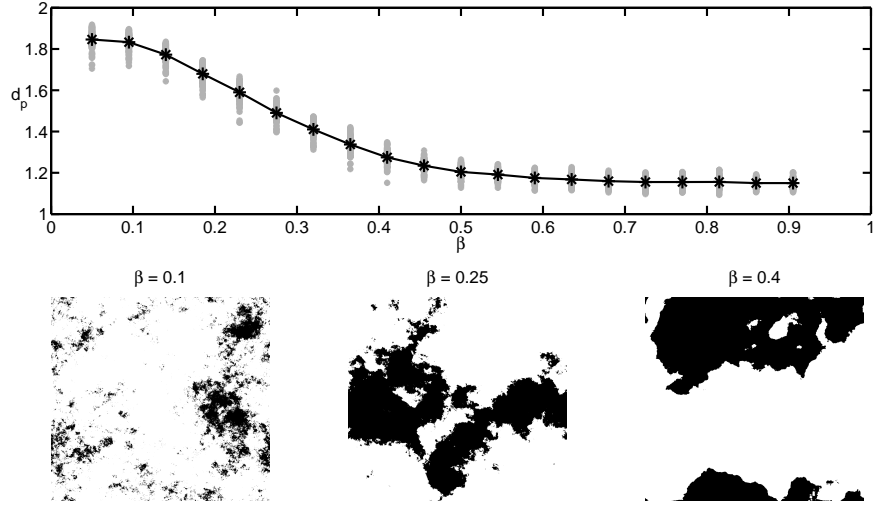


Figure 4.1: Realisations of fBm for a range of power spectrum scaling $\beta = 1 + 2H$ as defined in Eq. 4.8. The calculated boundary dimension (d_b) for a number of realisations is given in the top figure, whilst the bottom figure gives example realisations at different β levels. d_b scales linearly with β over a small range due to finite size effects.

In studies of spatial analysis, it is useful to be able to construct a neutral landscape model where certain statistical properties of the landscape are at specific

values, whilst other properties are allowed to vary randomly. Fractional Brownian Motion is one of the canonical examples of a neutral landscape model [Keitt, 2000]. The idea is to construct a two-dimensional random field X , where level sets of the random field form boundaries that have the desired fractal properties. Fractional Brownian motion (fBm) is a unique probability distribution with independent increments, stationarity and finite variance and is controlled by the Hurst exponent H . For $(x, y) \in \mathbb{R}^2$, we define a random field $X(x, y)$ such that the following conditions are satisfied:

1. $X(0, 0) = 0$ with probability 1.
2. $X(x, y)$ is a continuous function.
3. The height increments $X(x + h, y + k) - X(x, y)$ have a normal distribution with zero mean and variance $(h^2 + k^2)^H$ for $(x, y), (h, k) \in \mathbb{R}^2$.

The model then has one parameter H and this can be interpreted as the persistence of the process. If $H = \frac{1}{2}$, the increments become uncorrelated. With $H > \frac{1}{2}$, correlations are positive and realisations of the surface are smoother or more persistent. Conversely for $H < \frac{1}{2}$ points are anti-correlated and realisations become rougher or anti-persistent. Fractal Brownian patches are then defined as a level set of this process, i.e. $X^{-1}(c) = \{(x, y) : X(x, y) = c\}$, which defines the boundaries of the patches. Note that, in this context when talking about persistence or anti-persistence, it is about the spatial process only and has nothing to say about the underlying dynamics that caused it. fBm can be easily generalised to n -dimensions and thus can produce spatio-temporal models of vegetation patches, however each dimension may have its own Hurst exponent and hence spatial persistence may be independent of dynamic persistence for this general statistical model. However, fBm does have some desirable properties. The graph of fBm (referring to the set $\{(x, y, z) : x, y \in \mathbb{R}, z \in X(x, y)\}$ as opposed to the trail of fBm which only refers to the set $\{z : z \in X(x, y)\}$) can be shown to have both Hausdorff and box-counting dimension $3 - H$. This has led to the assumption that the box-counting dimensions and the Hurst exponent are the same, even though this only applies to fBm and other derivatives of this model. In general the Hurst exponent may be different or even independent of the box-counting dimension [Gneiting and Schlather, 2004]. The level-set of a fBm $X^{-1}(c)$ has box-counting dimension one less than the box-counting dimension of the corresponding fBm field. This then provides a relationship between the Hurst exponent for the underlying fBm and the box-counting dimension as

$$\dim_{\text{BC}}(X^{-1}(c)) = 2 - H. \quad (4.7)$$

Note that this only refers to the collection of boundaries of the level set that has this fractal property. A single patch boundary may have a different dimension due to finite-size effects. It also only refers to the boundary as opposed to the patch itself, which will necessarily have a box-counting dimension 2.

A realisation of fBm can be simulated by considering the spectral properties of the field. It can be shown via the WienerKhinchin theorem that a fBm generated with Hurst exponent H has the power spectrum

$$S(f) \sim 1/f^{1+2H} \quad (4.8)$$

This formula gives insight into the scaling properties of fBm, for example when $H = 0$, the process has $1/f$ noise. Also note that white noise corresponds to a constant power spectrum or $1/f^0$, which means that the corresponding Hurst exponent would be $H = -1/2$. Although this is not defined for the model it does give insight into how the model changes smoothly from Brownian motion $H = 1/2$ to Gaussian noise $H = -1/2$, where there is a transition from the path being continuous to discontinuous at $H = 0$.

An approximate realisation of fBm can be constructed via a spectral method [Hastings and Sugihara, 1993; Peitgen et al., 1988]. For an $N \times N$ grid an i.i.d. Gaussian white noise process is simulated for each point. A discrete two-dimensional Fourier transform can then be taken on this grid. The Fourier transformed Gaussian noise can then be multiplied by $1/(f_x^2 + f_y^2)^{\beta/2}$, where f_x, f_y are the wavelengths of the x and y component respectively. The corresponding power spectrum then has power law scaling of β . The inverse Fourier can be taken and the absolute value used as an approximation to the fBm process. Examples of level set of this can be seen in Fig. 4.1. Notice that due to finite size effects, there is a range of corresponding Hurst values for each value of β , this can be used to characterise the error in the estimation of a Hurst exponent using the box-counting dimension of the boundary. Since the simulated realisations of fBm can have a level set at any value, we may vary the constant continuously until the desired density is reached (for Fig. 4.1 all outputs are held constant with density at a half). We may, therefore, use this to simulate landscapes with the desired density and Hurst exponent, but with other properties allowed to vary randomly.

4.4 Lacunarity

Fractal Dimension can give insight into the complexity of a geometric object. However, this is one particular aspect of shape relevant to ecology. Another important concept is *lacunarity*, derived from the Latin for "lake" or "gap", which was first proposed by Mandelbrot [1983] as a complementary way of analysing surfaces and textures other than the fractal dimension. It has since received interest in Ecology, as a way of characterising the heterogeneity of a spatial pattern on a variety of length-scales. [Plotnick et al., 1996].

Lacunarity is calculated in a similar fashion to the box-counting dimension. A mesh of boxes of size ϵ is laid over the spatial data and points are counted up inside each box. $s_\epsilon(i, j)$ shall denote the point count in the box of size ϵ in the (i, j) position of the mesh. The spatial moments can then be taken and the lacunarity Λ at size ϵ is given by

$$\Lambda(\epsilon) = \frac{\langle s_\epsilon^2 \rangle}{\langle s_\epsilon \rangle^2}. \quad (4.9)$$

Essentially it is the second moment normalised by the mean squared. $\Lambda(\epsilon)$ has dependence on three variables: The density ρ of the spatial data, the size ϵ of the boxes and the complexity of the spatial data. If a pattern is highly heterogeneous on a scale ϵ then it would be expected that the second moment of the s is high compared to the first moment. It is also worth noting that in a pattern that is monofractal we would expect $\Lambda(\epsilon)$ to scale as a power law for increasing ϵ . It has been shown [Mandelbrot, 1983] that the slope of a lacunarity curve on a log-log plot for a monofractal is equal to $D - E$, where D is the fractal dimension and E is the Euclidean dimension (in our case this is 2). Notice that the slope is related to the Hurst Exponent by $D - E = -H$. Hence, if the curve deviates from a straight line, we may also use the measure to detect multifractal scaling and identify under which scales the relevant processes become dominant.

4.5 Multifractality

Monofractal techniques such as the box-counting dimension give insight into the scaling processes occurring in a vegetative landscape, however they do not provide information about the heterogeneity of the landscape. In this section, the multifractal formalism is introduced by first defining the local scaling of a spatial pattern and then using these local scaling laws to define sets of points that have the same local scaling. These sets also develop fractal scaling and thus have a Hausdorff or

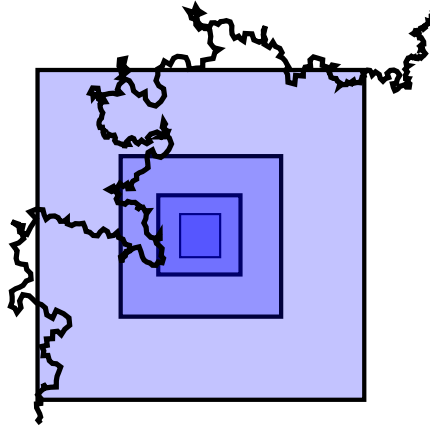


Figure 4.2: An example of clipping for a boundary. When box-counting is not centred on the mass of a fractal, spurious relationships form between the length-scale and the box-count due to boxes below a certain cut-off being empty.

box-counting dimension. The resulting analysis produces a spectrum of dimensions for each local scaling dimension and hence has the name *multifractal spectrum*. This technique is relatively easy to understand from a theoretical standpoint, despite being difficult to implement for reasons that shall be briefly discussed. Instead of using the local scaling properties to get at the multifractal spectrum, it is possible to study the multifractal properties by analysing the moments of the occupancy of boxes as they scale with box size. The scaling of these generalised dimensions define the Rényi dimensions D_q . These generalised dimensions can then be related to the multifractal spectrum via a Legendre transform. How this new multifractal formalism differs from the previously defined one shall be briefly discussed as will its limitations and problems on a binary dataset.

The Hölder exponent provides a description of the local scaling of a spatial pattern. The idea is very similar to the box-counting dimension, however with the Hölder exponent a dimension is given for each point in the set as opposed to the set itself. Let $N(\epsilon, x)$ be the number of points contained in the box of width ϵ centred at x . Then the corresponding Hölder exponent is defined as

$$\dim_{\text{loc}}(F, x) = \lim_{\epsilon \rightarrow 0} \frac{\log N(\epsilon, x)}{\log(1/\epsilon)}. \quad (4.10)$$

Due to the finite scale cut-off of data, the exponent is calculated by plot $N(\epsilon, x)$ against $1/\epsilon$ on a log-log plot and calculating the gradient via linear regression similarly to the box-counting dimension.

For a large dataset the number of linear regressions needed to be taken is large making computation slow. A novel way of increasing the efficiency of the algorithm, therefore, is to compute the corresponding regression matrix over all data points and use this matrix to perform the regression calculation. Using matrix notation, linear regression can be defined using the design matrix X defined as

$$\mathbf{X} := \begin{pmatrix} 1 & -\log(\epsilon_1) \\ 1 & -\log(\epsilon_2) \\ \vdots & \vdots \\ 1 & -\log(\epsilon_n) \end{pmatrix} \quad (4.11)$$

for box sizes $\epsilon_1, \epsilon_2, \dots, \epsilon_n$. The response variable \mathbf{y} is $\mathbf{y} = (N(\epsilon_1, x), N(\epsilon_2, x), \dots, N(\epsilon_n, x))^T$ for an arbitrary x . The parameter vector can then be estimated using ordinary least squares regression as

$$\hat{\beta} := (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y}. \quad (4.12)$$

Instead of performing this calculation for every single point in the dataset $x \in F$, we observed that the design matrix for each of the points is the same. Thus by indexing all the points in the dataset as x_1, x_2, \dots, x_m , we construct a matrix \mathbf{Y} as

$$\mathbf{Y} := \begin{pmatrix} N(\epsilon_1, x_1) & N(\epsilon_1, x_2) & \cdots & N(\epsilon_1, x_m) \\ N(\epsilon_2, x_1) & N(\epsilon_2, x_2) & \cdots & N(\epsilon_2, x_m) \\ \vdots & \vdots & \ddots & \vdots \\ N(\epsilon_n, x_1) & N(\epsilon_n, x_2) & \cdots & N(\epsilon_n, x_m) \end{pmatrix}. \quad (4.13)$$

The parameter matrix can then be similarly defined using ordinary least squares as

$$\hat{\mathbf{B}} := (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{Y}. \quad (4.14)$$

This produces a $2 \times m$ matrix where the Hölder exponent can be read off from the second row. Now we have a single matrix calculation, which can provide a significant speed-up compared with performing linear regression on each point separately.

Having defined the Hölder exponent, we can now introduce the multifractal formalism also known as fine multifractal analysis [Falconer, 2013]. The idea is to compute the Hölder exponent for all points in a set F and then define sets that have the same Hölder exponent. Let the set E_α be

$$E_\alpha := \{x \in F : \dim_{\text{loc}}(F, c) = \alpha\}. \quad (4.15)$$

This produces sets covering a range of values from α_{\min} to α_{\max} . As these are also binary sets, the box-counting dimension can be taken on each to give the fine multifractal spectrum

$$f(\alpha) = \dim_{\text{BC}}(E_\alpha). \quad (4.16)$$

The multifractal spectrum gives an indication of the heterogeneity in the scaling of the spatial pattern. For a monofractal spatial pattern the scaling is homogeneous, hence a single set E_α is defined and the multifractal spectrum is a single point. A pattern that is more multifractal will generally have a wider spectrum. There are, however, issues with this calculation for data. Firstly the resolution may not be fine enough to produce sets E_α that can have their box-counting dimension accurately inferred. Secondly, it is not necessarily clear how the sets E_α should be defined. A range of α values may be close together due to error in regression or sampling technique etc. and so it is not clear whether a point should be included in one set or should form an independent set with that value of the Hölder exponent. We can circumvent these issues by introducing another type of multifractal approach.

The idea is to construct a measure for the spatial pattern F and calculate how the moments of the measure scale for increasing length. We define the measure to be the probability of occupancy. This is the ratio of the number of points in a box divided by the total number of points. It can be seen that this satisfies the properties of a measure. the generalised moments of $P(x, \epsilon)$ can then be calculated as

$$M_q(\epsilon) = \sum_{x \in F} P(x, \epsilon)^q. \quad (4.17)$$

The scaling of these moments is assumed to be a power-law i.e. $M_q \sim \epsilon^{-\beta(q)}$, where $\beta(q)$ is defined as

$$\beta(q) := \lim_{\epsilon \rightarrow 0} \frac{\log M_q(\epsilon)}{\log(1/\epsilon)}. \quad (4.18)$$

$\beta(q)$ can also be used to define the Rényi dimensions D_q via the following formula

$$D_q := \frac{1}{1-q} \beta(q) = \frac{1}{1-q} \lim_{\epsilon \rightarrow 0} \frac{\log M_q(\epsilon)}{\log(1/\epsilon)} \quad (4.19)$$

where D_0 is the familiar box-counting dimension. D_1 can be derived in the limit as $q \rightarrow 1$ using L'hôpital's rule as

$$D_q := \lim_{q \rightarrow 1} \frac{1}{1-q} \beta(q) = \lim_{\epsilon \rightarrow 0} \frac{-\sum_x P(x, \epsilon) \log P(x, \epsilon)}{\log(1/\epsilon)}, \quad (4.20)$$

which is the definition of entropy for a probability measure P divided by the log of the box length. D_1 is, therefore, referred to as the information dimension. Higher orders of q can be taken until the limit as $q \rightarrow \infty$ where the contribution to the sum comes only from those boxes which are maximally filled, hence D_∞ provides information on how the densest regions of the set F scales. D_q can also be defined for $q < 0$, with boxes that contribute towards the sum in Eq. 4.17, which are sparse. Thus, in the limit $D_{-\infty}$ gives insight into how the sparsest parts of the set F scale. We can see intuitively how the concept of the Rényi dimension relates to the fine multifractal spectrum, however in order to produce something comparable the Legendre transform needs to be taken.

The Legendre transform takes the extensive pairing $(q, \beta(q))$ and converts it into the intensive pairing $(\alpha, f(\alpha))$ using the following form

$$f(\alpha) = \inf_{-\infty < q < \infty} \{\beta(q) + \alpha q\}. \quad (4.21)$$

Assuming differentiability of $\beta(q)$ the multifractal spectrum $(\alpha, f(\alpha))$ can be calculated by minimising $\beta(q) + \alpha q$ over q , i.e.

$$\alpha(q) = -\beta'(q), \quad (4.22)$$

$$f(\alpha(q)) = \beta(q) + q\beta'(q). \quad (4.23)$$

The multifractal spectrum then produces a concave curve with a single peak at $\alpha(0)$ corresponding to the box-counting dimension. It can also be noted by differentiating f with respect to α using the chain rule that $\frac{df}{d\alpha} = q$. The multifractal spectrum therefore admits a single stationary point at $\alpha(0)$, which is a maximum.

4.5.1 The multifractal spectrum: density against occupancy data

For the purposes of calculating the multifractal spectrum, care must be taken over the form of the data being used and whether the underlying data is indeed multifractal in the sense that it has moments M_q defined in Eq. 4.17 that scale as a power-law. If not, then the multifractal spectrum is not an adequate measure of the heterogeneity in the spatial pattern and other techniques should be considered instead.

For density data, the form of the measure used to compute the multifractal spectrum can be taken by normalising the density values over the whole space. The

measure of a box of length ϵ then becomes a sum of the normalised density of that box. This can then lead to a calculation of the spectrum by calculating the moments and then following the Legendre approach.

For occupancy data the situation is less clear. A probability measure can be constructed once again by dividing each occupancy point by the total number of occupied lattice points. The probability measure is then calculated by summing the total number of occupied lattice points in a box of size ϵ and dividing by the total number of points. The issue with this lie with how the support of measure should be formed.

The support is defined as having a non-zero measure everywhere. At the boundary of a measure however, boxes are not centred on a point and hence will give spurious values when the negative moments are calculated. This phenomena is known as clipping (Fig. 4.2). This can be circumvented by only taking boxes centred over a measure and ignoring all others. The second drawback of the multifractal measure is having a more obscure interpretation than other scaling methods.

4.5.2 Pattern anisotropy

In the previous sections, we have concentrated on the scaling of spatial pattern, that is how certain quantities such as patch-size or box-occupancy grow or decay with increasing size. Another method of analysing spatial pattern in vegetation is to consider how certain quantities change with angle. Anisotropy for spatial patterns is defined as the variation in density for changing angle. If the spatial pattern is statistically invariant under rotation then the pattern is described to be *isotropic*.

Anisotropy in vegetation has been of great interest [Haase, 2001; Purves and Law, 2002; Watt, 1947]. However, few statistically robust methods exist for characterising the anisotropy if the spatial pattern is regular, where the vegetation is formed into regular stripes or bands. Two-dimensional Fourier spectral analysis can then be used to detect the dominant direction of the spatial pattern. The discrete two-dimensional Fourier transform is similar to its one-dimensional counter-part and conceptually, compares the spatial pattern to sine and cosine functions of varying scale. A term where there is a strong overlap between the function and the data leads to a larger contribution than when there is little overlap. For spatially-discrete two-dimensional data $x_{i,j}$ with spatial co-ordinates i, j , the discrete Fourier transform is

given as

$$X_{k,l} = \sum_{n_1=0}^{N-1} e^{-2\pi i l n_1 / N} \sum_{n_2=0}^{N-1} e^{-2\pi i k n_2 / N} x_{i,j}, \quad (4.24)$$

where the size of the space is $N \times N$. The transform takes the data $x_{i,j}$ in the spatial domain and converts it into $X_{k,l}$ in the frequency domain. Since the Discrete Fourier Transform is complex, the absolute value can be taken by multiplying it with its conjugate in order to obtain the two-dimensional power spectrum. This can then be used to analyse the dominant frequencies in the spatial pattern. There are a number of problems with this approach for two-dimensional occupancy data. The test functions that are being compared to the data $e^{2\pi i k}$, although discrete in space, take on continuous values whereas the data has only binary values 0 or 1 if the site is occupied or not. The test functions also assume underlying periodicity in the data, which may not be the case. The basis for the function is in the x- and y-direction only and as such it is not clear how the measure of anisotropy would change if the data was rotated. These problems can be overcome by using wavelet analysis, which has had a number of applications in Ecology in recent years [Bradshaw and McIntosh, 1994; Bradshaw and Spies, 1992; Cazelles et al., 2008; Dale and Mah, 1998; Nakken, 1999; Perry et al., 2002]

Wavelet Analysis is similar to Fourier Analysis in the sense that the data is compared to test functions of varying scale. In contrast to the Fourier Transform, wavelets have compact support, which means strength at location as well as scale can be measured. There is also larger freedom with wavelet analysis by being able to choose a certain form of a wavelet for the task, whichever is more appropriate.

For the purposes of detecting regular spatial pattern in the Isles of Scilly data, we employ a method along the lines of Rosenberg [2004]. Wavelet functions used in wavelet analysis often satisfy a number of technical conditions, but in general they can have variety and a large number of families of wavelets exist. The wavelet used in the analysis is known as the French Top Hat, although other wavelet families have been used in the analysis of vegetation spatial pattern such as the Mexican Hat wavelet [Dale and Mah, 1998]. The French Top Hat wavelet is given by the following formula:

$$g(x) = \begin{cases} -1 & \text{if } 1/2 < |x| < 3/2, \\ 2 & \text{if } |x| < 1/2, \\ 0 & \text{otherwise.} \end{cases} \quad (4.25)$$

Produces a single discrete wave shape with a flat top. This creates a shape that is ideal for fitting to occupancy data. The wavelet function can be translated via a parameter a and stretched by a parameter b by taking $x \rightarrow \frac{x-a}{b}$. These parameters provide a way of finding both the scale and position of a strong banded pattern. The wavelet transform at position a and scale b is defined as

$$W(b, a) = \frac{1}{b} \sum_{j=1}^n y(\theta_j) g\left(\frac{\theta_j - a}{b}\right), \quad (4.26)$$

where n is the number of transects taken and $y(\theta_j)$ is the data at transect point θ_j . The overall variance over all scales b can then be easily calculated at a given transect position

$$P(\theta_i) = \frac{1}{m} \sum_{k=1}^m W^2(b_k, \theta_i), \quad (4.27)$$

where m is the largest length scale considered (half the total distance of the transect). $P(x_i)$ is then maximised when the shape of the wavelet $g(x)$ best fits the data for that transect, in other words, when there is strong banding present. This is due to occupied sites being sparse if a transect runs along a gap, whereas sites are dense if a transect runs perpendicular to the banding. This produces a distinct change in variation around certain angle transects and can also determine the angle width over which the banding occurs.

The analysis in Rosenberg [2004] has been adapted from point process data to site occupancy data. This is done by approximating the occupancy data as a point pattern if a box is in the occupied state taking its centre as the point. This is an approximation as on the smallest scale we have no knowledge over whether the occupied box represents a single individual, a cluster of vegetation or even a continuous mass of vegetation. For larger scales, this approximation does not bias the data as a transect would in general be much larger than a single box, thus where the box is occupied is inconsequential. The problem as to whether a box is occupied with a single point or a cluster of points can be overcome by taking the box size small enough such that no more than a single clonal unit could occupy the box. It is still not clear for a continuous mass of vegetation where the distinction between individual or cluster should be made. Our analysis, however, is simply to look at the overall variation in spatial pattern and as such this distinction should not affect the overall results. This approximation therefore seems an appropriate one to make. The analysis can then proceed in the same way as outlined by Rosenberg [2004]. Each site contains a number of occupied sites, each occupied site can have the same

wavelet analysis performed on it in order to produce the statistic $P(\theta_i)$. The $P(\theta_i)$ can then be averaged over all occupied sites producing an overall variation $P(\theta_i)$ for the whole site. The maximum peak can then be recorded and used to compare to the other sites in order to detect where the vegetation is strongly anisotropic.

4.6 Results

4.6.1 Box-counting dimension

Site	ρ	d_b	range (m)
blt	0.4612	1.9916	25.6-409.6
htb	0.1463	1.5949	6.4-409.6
la	0.2013	1.6827	3.2-51.2
ogh	0.0679	1.6101	1.6-51.2
wbl	0.1832	1.9275	12.8-204.8

Table 4.1: Computed values for the density of occupied sites ρ , the box-counting d_b dimension and the range over which scaling was computed to be constant.

All sites passed the $R^2 - SSR$ criterion for fractality, except for the site wbl where the estimated box dimension failed the test. However, this site just failed, in the sense that the set which maximised the r^2 had an almost equivalent value of the SSR as for the set that minimised the SSR . A possible explanation for why this site would fail would be the strong anisotropy introduced from the observed banding. Banding introduces a characteristic length scale that violates the scale-free definition of fractality. The test also gives estimates for the range of scaling that approximates a fractal closest. These are: blt: 25.6-409.6m , ogh: 1.6-51.2m, wbl: 12.8-204.8m, htb: 6.4-409.6m, la: 3.2-51.2m. The two lowest fractal values also interestingly have the smallest range reaching a maximum of only 51.2m. The smaller ranges of la and ogh is due to the size of the meadows being smaller than in the other meadow sites. The estimated box-counting dimension values are given in Table 4.1. Sites blt and wbl have the highest box-counting dimensions, whereas ogh and htb have the lowest box-counting dimensions. Visually comparing the spatial maps of the vegetation distribution (Fig. 4.7), it is difficult to ascertain the general nature of the boundaries of the patches. Hence, the box-counting dimension provides a description of the spatial properties of the system beyond what can be done through eye-balling the data.

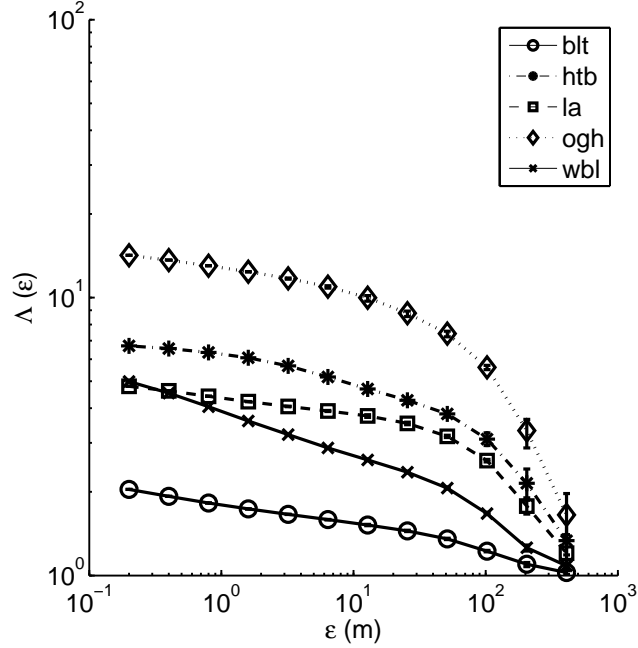


Figure 4.3: The lacunarity calculated from box-counting for each study site.

4.6.2 Lacunarity

Lacunarity analysis was performed on each of the five sites over a range of approximately 500m^2 . The lacunarity was calculated using a box-counting method with periodic boundary conditions to avoid clipping [Feagin et al., 2007]. Estimates of the error were calculated by repeating the lacunarity analysis over images that were transformed with a random offset. The results can be seen in Fig. 4.3. All the study sites have a change in scaling around $\epsilon \approx 10^2\text{m}$ that is indicative of a clustered distribution with cluster size around 10^2m as it can be observed from visual analysis of the images. The scaling change is less pronounced in blt and wbl, where the bounding box of the image is contained in the patch, thus leading to less of a change in scale.

Comparing to the null model

The lacunarity plots were compared to a null model where spatial correlation is ignored. This is done by considering the same lattice size as the data and modelling each point as a Bernoulli trial with probability p equals to the density of the original data. Note that for this model and for an infinite lattice, the lacunarity at the smallest box size (the size of the lattice) can be calculated explicitly since the

variance and the mean of the boxes is just the variance and mean of the Bernoulli distribution. Therefore

$$\Lambda(1) = 1 + \frac{\text{Var}(X)}{\text{E}(X)^2} = 1 + \frac{p(1-p)}{p^2} = p^{-1}.$$

Hence $\Lambda(1)$ is the inverse of the density. For box-sizes larger than one, the lacunarity can also be calculated as a box of length ϵ contains $\epsilon \times \epsilon$ Bernoulli trials and the count is a Binomial distribution with parameters $\epsilon \times \epsilon$ and p . Hence in general for an infinite lattice size the lacunarity calculation is for the box-count $X \sim B(\epsilon^2, p)$

$$\Lambda(\epsilon) = 1 + \frac{\text{Var}(X)}{\text{E}(X)^2} = 1 + \frac{np(1-p)}{n^2p^2} = 1 + \frac{1}{\epsilon^2}(p^{-1} - 1). \quad (4.28)$$

The lacunarity decays to 1 for increasing box-size ϵ at $O(\epsilon^{-2})$. For a large occupation probability the resulting distribution would be relatively uniform, leading to $(p^{-1} - 1) \ll 1$, thus a lacunarity that is close to one for all box-sizes. A small probability of occupation would lead to some boxes being occupied with many not being occupied. This would lead to a more heterogeneous spatial pattern and as such a lacunarity that is large (although still scaling as ϵ^{-2}). A high lacunarity then may not indicate structure anymore interesting than just a low occupation probability. The simple occupation model therefore, can be used as a null model to test for statistical significance on the lacunarity plots. Each site's probability of occupation p is used to generate 10^3 realisations of the model with the same lattice size. The lacunarity analysis is then performed on each realisation and the upper and lower percentiles are compared with the lacunarity of the original plots to check for statistical significance. Note that we perform the statistical significance test in this manner as opposed to using the theoretically derived results in Eq. 4.28 as the lattice sizes are finite and needs to be accounted for. The results show (Fig. 4.4) significant lacunarity from the null model for each of the five sites. Blt shows the lowest lacunarity of the five sites, although its scaling can be seen to be significantly different to the null model indicating the presence of spatial correlation.

4.6.3 Anisotropy

The anisotropy of the seagrass distribution was measured for the entire data set of the Isles of Scilly by partitioning the data into a 322×322 grid where each grid represented an area of approximately $15\text{m} \times 15\text{m}$. The wavelet analysis was performed on each box and the maximum point was taken as a measure of the strength and direction of the anisotropy. The results were then overlayed on a

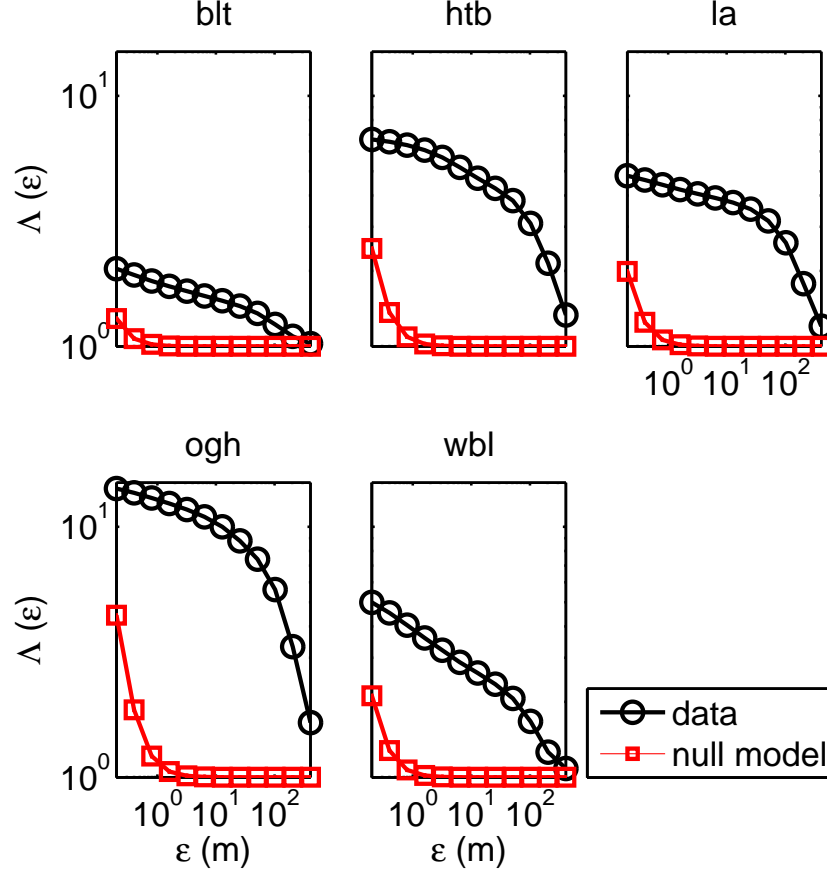


Figure 4.4: Comparison of the lacunarity plots to the null model where spatial correlations are ignored. All sites show a significance divergence from the null model (95% confidence intervals were plotted, but are completely contained in data markers). The null model scaling on the log-log scale is exponential with exponent -2 as expected from the analytic result in Eq. 4.28

map of the Isles (See Fig. 4.5). Areas where the strength of the anisotropy were highest were found to have a strong banded structure in the occupation of vegetation. Regular bands of similar width, approximately $1 - 5\text{m}$, were found everywhere there was a strong region of anisotropy. The method appears to be able to pick out areas where vegetation banding is strong. Fig. 4.6 displays the anisotropy analysis as directors (headless vectors) with a length proportional to the magnitude of the dominant directionality. Each director was averaged around a box size of 10×10 of the original analyses so that it can be displayed clearly in the figure. There is a clear East-West directionality for the wavelet anisotropy running perpendicular to the dominant currents in the Isles of Scilly (which run North to South). Fig. 4.6

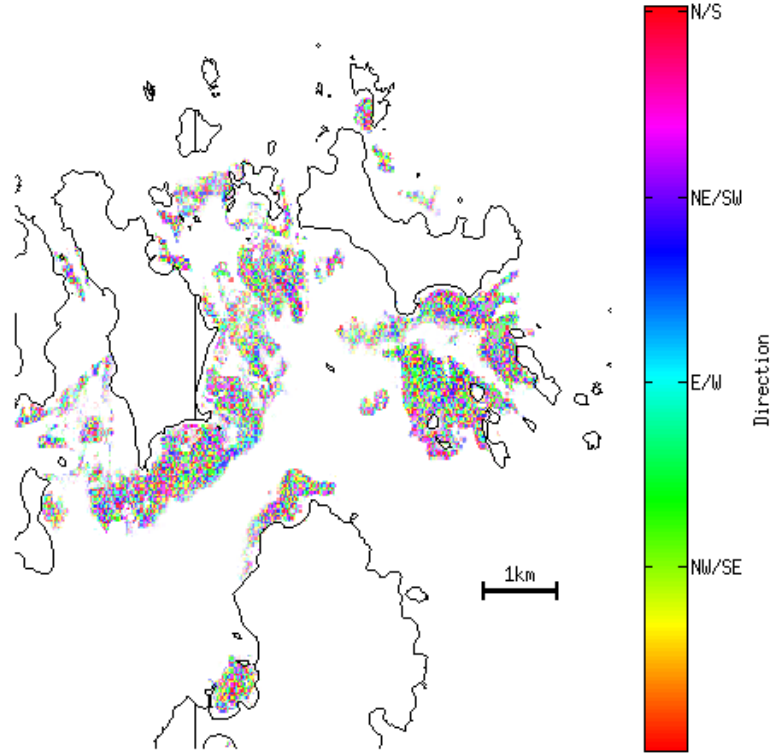


Figure 4.5: Anisotropy analysis of the seagrass spatial pattern. Contains Ordnance Survey data © Crown copyright and database right 2011

is derived by averaging over a number of sites where the analysis was performed. The original distribution of dominant angles in the wavelet analysis also reveals the strong East-West direction.

4.6.4 Hölder exponent

The Hölder exponent was used to test how scaling varies with space across the five sites. This was calculated by taking the occupancy data Ω , where Ω is an $m \times n$ matrix with values 0 or 1 depending on whether the site was occupied or not. The Hölder exponent was then calculated according to the method outlined in section 4.5. Each of the five meadows surveyed displayed a large range of values for the Hölder exponent, between 1 and 2.5 (Fig. 4.8). Each site has a unimodal distribution of exponents with a peak of approximately 2, representing the interior of the meadows. All four sites have very similar distributions with the exception of *wbl*, where the

distribution contained more mass in the left tail (between 2 and 1) than for the other sites. Whereas in the other sites that contain large patches with an interior, *wbl* is more fragmented, thus introducing areas where the Hölder exponent is lower. All sites have values greater than 2, although theoretically this is not possible as the maximum a box of length ϵ can be occupied is ϵ^2 , where there may be a sharp change in the scaling around an area (such as by a patch border) this can lead to a sharp change in the scaling. As linear regression is being performed over all the scales this leads to some sites that have unexpectedly high Hölder exponents.

In order to compare the Hölder exponent with the local density where the exponent was calculated the box of length half of that over which regression was taken was used to determine if there is a strong correlation between the exponent and the average density (Fig. 4.9). All five meadows sites' exponents have a medium correlation coefficients ($0 < r < 0.5$) with the average density indicating the exponent is providing more information than just density alone. Although there is a degree of correlation with the average density, by comparing to the plots between the two, there is no clear trend.

4.7 Discussion

In this chapter we have laid out the significant body of work in descriptive statistics for spatial ecology. The work includes scaling statistics, such as the box-counting dimension and Hölder exponent, as well as measures of anisotropy and heterogeneity in the distribution of vegetation. The merits of each have been considered for use in measuring disturbance and patchiness in the vegetation ecosystem. Later chapters will compare these measurements to the underlying dynamics of the vegetation community using both modelling techniques as well as direct comparison with a quadrat-based survey of the five sites.

The five sites used in the seagrass study for the Isles of Scilly, UK were found to have a number of interesting spatial features. Their heterogeneity as defined by the lacunarity was found to be significantly higher than the lacunarity of a spatially uncorrelated landscape. Each site also showed a degree of power law scaling in both the patch-size distribution and the boundary dimension defined through box-counting. These results show there is a significant degree of fractality in the vegetative community and hence the hypothesis that there is fractal scaling in seagrass vegetation is valid.

It was found that strong anisotropic patterns exist in a number of locations in the Isles of Scilly dataset. The patterning was found to be bands of vegetation regularly spaced in a vegetation cluster. This banding was found to be prevalent in most of the sites in the study. Banding was also found to be perpendicular to the dominant currents. The origin of these bands and their impact on the dynamics of the Seagrass and vegetation in general is an interesting problem and one that will be the focus of the next chapter as well as Chapter 7 and Chapter 9.

The next chapter begins the discussion of modelling vegetation communities with strong spatial structure. The modelling shall try to consider in particular two different form of spatial pattern exhibited: the power-law scaling structure in the patch-size distribution and boundary; and the anisotropy of the spatial pattern characterised by banding perpendicular to the main flow of currents.

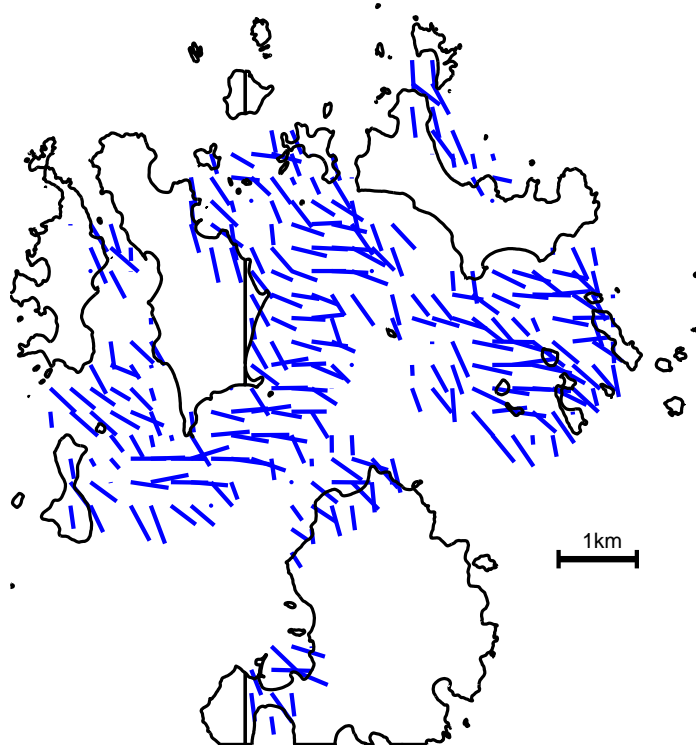


Figure 4.6: Anisotropy analysis of the vegetation spatial pattern using the wavelet method. The directors in the graph represent the direction of the dominant variation in the spatial pattern with it's length proportional to the strength of the spatial pattern. Results were averaged for box sizes and a vector of the Scilly Isles was approximately overlayed on top. Contains Ordnance Survey data © Crown copyright and database right 2011

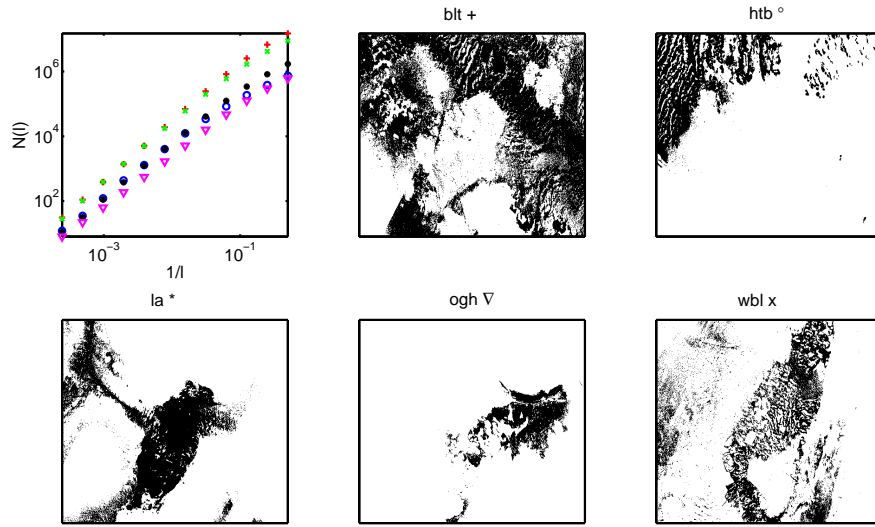


Figure 4.7: Estimating the Box Dimension from the five locations. The spatial data for patches A-E are displayed around the central figure. Values of 1 representing the presence of Seagrass are shown in green and values of 0 representing no Seagrass are shown in blue. The main figure shows how the number of occupied boxes N scales inversely with the length of the box L on a log-log plot. We find that the power law scaling is remarkably clear over a wide range of L , with only slight deviations as L becomes small for sites A and D. We also note a strong difference for site B. For a large range of L the values of N deviate greatly from the other sites.

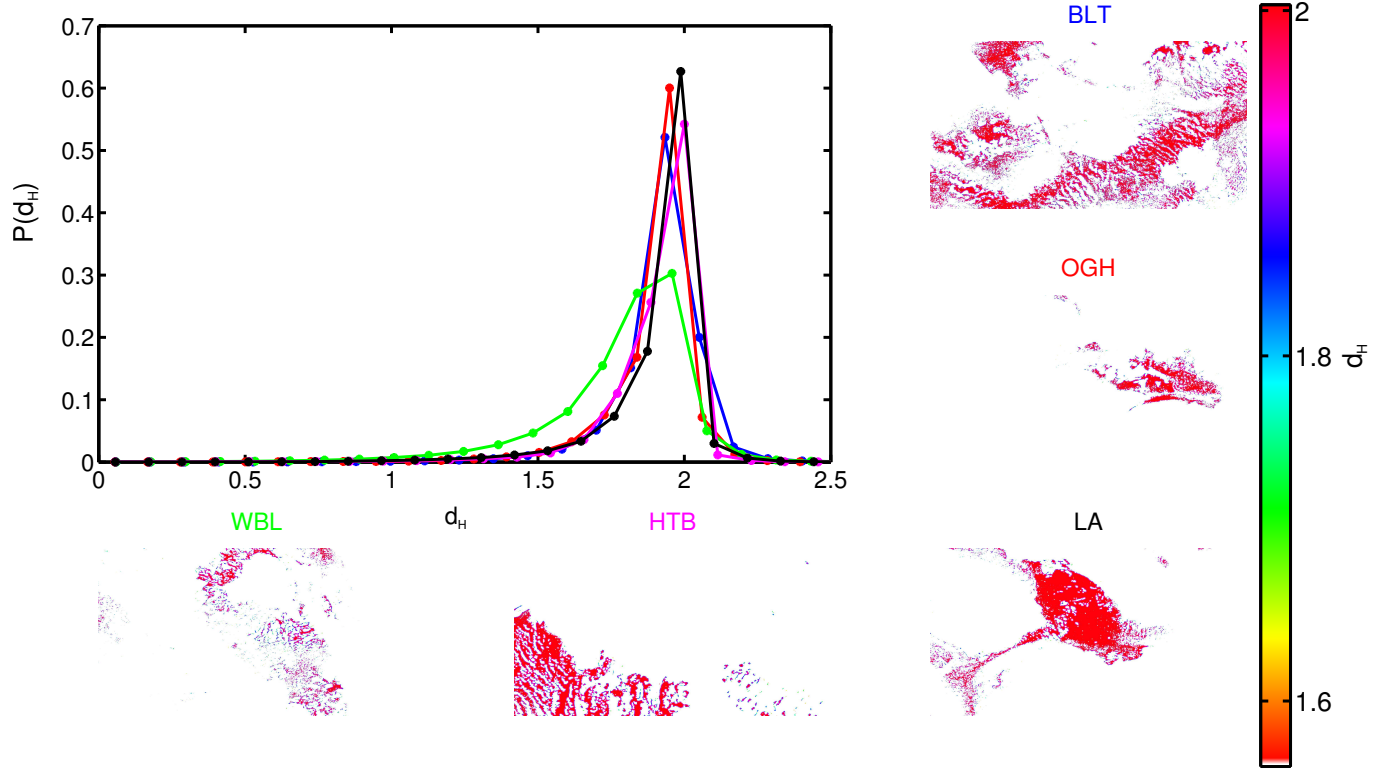


Figure 4.8: The Hölder exponents (d_H) calculated for the five meadows *blt*, *htb*, *la*, *ogh* and *wbl*. Distribution of the exponents is given in the top left figure and their locations are given in the surrounding figures for each meadow. All meadows were found to have broadly similar unimodal distributions of exponents. *wbl* however, is an outlier with a heavier left-tail due to the more fragmented structure of the spatial distribution of vegetation.

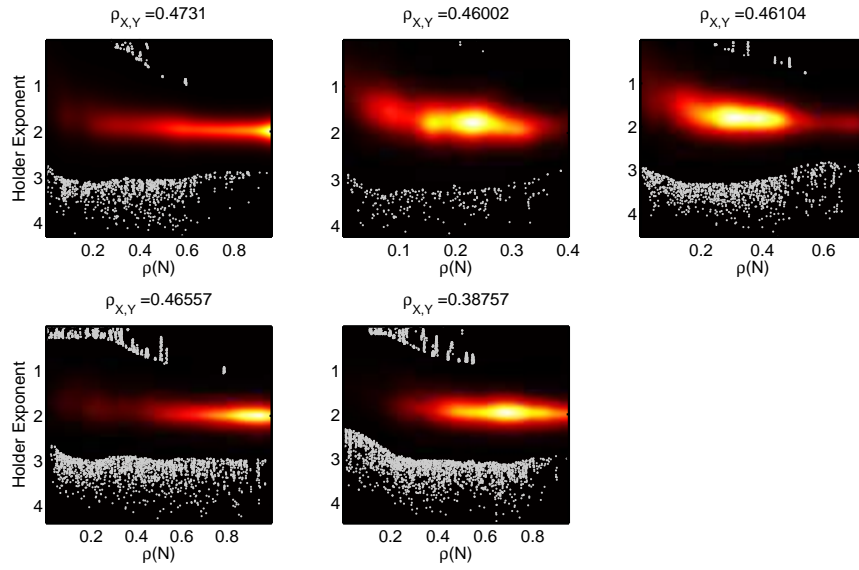


Figure 4.9: correlation plots of Hölder exponents against the density of box with the average length for each of the five meadows. Each site explored had a medium strength correlation (less than 0.5) between the box density and the Hölder exponent providing evidence that the exponent is providing more information than just the average density alone. The density of points is given as a heat map with outliers shown in white.

Bibliography

- GA Bradshaw and BA McIntosh. Detecting climate-induced patterns using wavelet analysis. *Environmental Pollution*, 83(1):135–142, 1994.
- GA Bradshaw and Thomas A Spies. Characterizing canopy gap structure in forests using wavelet analysis. *Journal of ecology*, pages 205–215, 1992.
- Bernard Cazelles, Mario Chavez, Dominique Berteaux, Frédéric Ménard, Jon Olav Vik, Stéphanie Jenouvrier, and Nils C Stenseth. Wavelet analysis of ecological time series. *Oecologia*, 156(2):287–304, 2008.
- MRT Dale and M Mah. The use of wavelets for spatial pattern analysis in ecology. *Journal of Vegetation Science*, 9(6):805–814, 1998.
- Kenneth Falconer. *Fractal geometry: mathematical foundations and applications*. John Wiley & Sons, 2013.
- RA Feagin, XB Wu, and T Feagin. Edge effects in lacunarity analysis. *Ecological modelling*, 201(3):262–268, 2007.
- T Gisiger. Scale invariance in biology: coincidence or footprint of a universal mechanism? *Biological Reviews of the Cambridge Philosophical Society*, 76(02):161–209, 2001.
- Tilman Gneiting and Martin Schlather. Stochastic models that separate fractal dimension and the hurst effect. *SIAM review*, 46(2):269–282, 2004.
- Peter Haase. Can isotropy vs. anisotropy in the spatial association of plant species reveal physical vs. biotic facilitation? *Journal of Vegetation Science*, 12(1):127–136, 2001.
- Harold M Hastings and George Sugihara. Fractals. a user’s guide for the natural sciences. *Oxford Science Publications, Oxford, New York: Oxford University Press*, — c1993, 1, 1993.

- Timothy H Keitt. Spectral representation of neutral landscapes. *Landscape Ecology*, 15(5):479–494, 2000.
- B.B. Mandelbrot. *The fractal geometry of nature*. Wh Freeman, 1983.
- Margriet Nakken. Wavelet analysis of rainfall–runoff variability isolating climatic from anthropogenic patterns. *Environmental Modelling & Software*, 14(4):283–295, 1999.
- Heinz-Otto Peitgen, Dietmar Saupe, Michael Fielding Barnsley, Yuval Fisher, and Michael McGuire. *The science of fractal images*. Springer New York etc., 1988.
- JN Perry, AM Liebhold, MS Rosenberg, J Dungan, M Miriti, A Jakomulska, and S Citron-Pousty. Illustrations and guidelines for selecting statistical methods for quantifying spatial pattern in ecological data. *Ecography*, 25(5):578–600, 2002.
- R.E. Plotnick, R.H. Gardner, W.W. Hargrove, K. Prestegard, and M. Perlmutter. Lacunarity analysis: a general technique for the analysis of spatial patterns. *Physical Review E*, 53(5):5461, 1996.
- DW Purves and R Law. Fine-scale spatial structure in a grassland community: quantifying the plants-eye view. *Journal of Ecology*, 90(1):121–129, 2002.
- Michael S Rosenberg. Wavelet analysis for detecting anisotropy in point patterns. *Journal of Vegetation Science*, 15(2):277–284, 2004.
- Laurent Seuront. *Fractals and multifractals in ecology and aquatic science*. CRC Press, 2009.
- Alex S Watt. Pattern and process in the plant community. *The Journal of Ecology*, pages 1–22, 1947.

Chapter 5

Modelling

Chaos is found in greatest abundance wherever order is being sought.
It always defeats order, because it is better organized.

(Terry Pratchett -*Interesting Times*)

5.1 Introduction

In Chapter 4 a number of spatial features were discovered for the eelgrass dataset. We found that seagrass exhibits strong regular patterns in the form of bands that, in general, are perpendicular to the dominant currents. The focus of this chapter will be on exploring the mechanisms that underpin these spatial patterns and what set of parameters control the strength and scale of the banding pattern. In Section 5.2 the modelling is constructed from reaction-diffusion type equations. In Section 5.4 the role of stochasticity is explored to see what contributes to the strength and maintenance of regular spatial patterns in vegetative communities. Stochasticity is further explored in Chapter 6

Eelgrass is subject to a number of different environmental factors that need to be considered in a modelling context [Moore and Short, 2006]. Propagation occurs through local clonal growth in the form of rhizomes that grow throughout the year. It is a sub-tidal species of seagrass and is found in estuaries and along coastlines in shallow waters. It is characterised by long, broad leaves that can grow to 30-60cm in length [Kuo and Den Hartog, 2001]. The resulting morphology of the plant structure is affected by substrate type [Short, 1983], temperature [Moore et al., 1996], light and nutrient availability [Marba et al., 1996], and tide and wave regimes [Fonseca and Bell, 1998]. Dense clusters of foliage can act as a filter, trapping and binding sediment and dampening wave and current energy. This also facilitates the accumulation of organic material in the sediments leading to an increased concentration of nutrients.

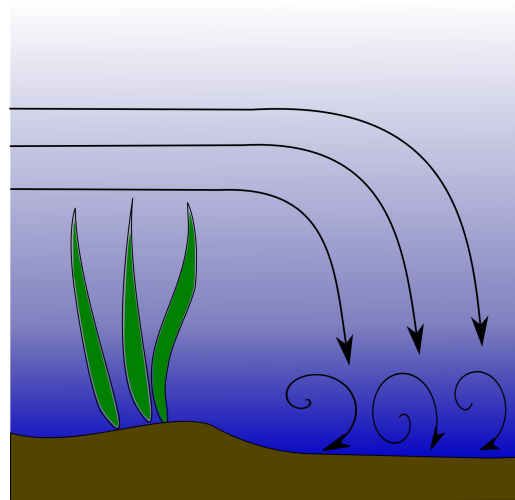


Figure 5.1: A caricature of the process of hydrological scouring in a seagrass bed. current flow (shown as black arrows) over the vegetation (in green) causing turbulent flow down-current of the vegetation that scours the bed lowering the probability of new growth.

There are potentially a huge number of effects that can affect growth and would potentially need to be modelled. We are therefore presented with a choice between realism and abstraction. A realistic model would explicitly model all factors associ-

ated with growth such as nutrient cycles, individual shoot and rhizome propagation, current and wave action, as well as sediment interaction and seasonal cycles. An abstract model would contain only a small number of features that would be associated with eelgrass growth laid out in very general terms, such as a functional response to local density and general environmental conditions. A realistic model suffers from analytic and computational intractability as well as difficulty in parameter fitting, whereas an abstract model suffers from lack of predicting power and realism. Therefore, we have proceeded with a compromise between the two, by taking only the salient aspects of colonisation and interaction and ignoring other possible interactions. This also allows the model to be more general and thus applicable to a wider number of ecosystems. For now we will also ignore distributed disturbance effects, such as disease or grazing, and focus on seagrass and its interaction with the environment. There are clear positive local interactions through sediment deposition; nutrient production and deposition, and mutual sheltering. Negative interaction between vegetation occurs through competition for nutrients and hydrological scouring; where a raised dense bed of vegetation induces turbulent flow that scours the seabed down current of the vegetation patch (Fig. 5.1). Interaction is therefore mediated by the environment through sediment, nutrient concentrations, and through density of vegetation itself. These mechanisms are explored using a variety of modelling frameworks.

This chapter begins with an introduction to reaction-diffusion equations applied to vegetation-environment interaction. A novel attempt to characterise this system is explored as well as discussion of similar efforts that have been made in the literature. The next section explores continuous models, where long-range interaction occurs via integral terms in the kinetic equation. A model of tiger bush [Lefever and Lejeune, 1997] is adapted to our system and a novel derivation of this model is given. In order to explore how stochasticity interacts with pattern formation in later chapters, the continuous kinetic equation is adapted into a probabilistic cellular automata, where the rates of the system relate to the rates given in the continuous model. Pattern formation is then explored in this model via simulation.

5.2 Reaction-diffusion equations

The main purpose of this section shall be to derive a parsimonious set of reaction-diffusion equations that will explicitly model the vegetation's interaction with itself in the form of spatial diffusion and competition, and the environment in the form of

a substrate density that provides protection to the vegetation and aides its growth. The vegetation density $u(t, x, y)$ and the environmental substrate $v(t, x, y)$ shall therefore be explicitly modelled in the domain $[0, T] \times \mathbb{R}^2$ for some time T .

The general approach to reaction-diffusion equations in ecological systems is outlined in Dieckmann et al. [2000], although a brief summary is given here. The strategy is to find stable homogeneous fixed points of the equation. Small spatial perturbations are added to the homogeneous solution and the set of wavelengths that are unstable i.e. the wavelengths of a perturbation that grow in time are calculated. If the unstable wavelengths occupy a finite range and are greater than zero then the reaction-diffusion system is said to have a Turing bifurcation. We begin with a general model of a reaction-diffusion equation with 2 species each with concentration u_i (note that u_i depends both on space x and time t , but these dependencies are not shown to save on notation) in one spatial dimension. The general equation is then

$$\frac{\partial u_1}{\partial t} = u_1 f_1(u) + \mu_1 \frac{\partial^2 u_1}{\partial x^2}, \quad (5.1a)$$

$$\frac{\partial u_2}{\partial t} = u_2 f_2(u) + \mu_2 \frac{\partial^2 u_2}{\partial x^2}, \quad (5.1b)$$

where f_1 and f_2 are the density dependent growth rates for species 1 and 2. μ_1 and μ_2 are the diffusion rates for both species respectively. A homogeneous solution $\bar{u} = (\bar{u}_1, \bar{u}_2)$ is one with no time or spatial dependency and satisfies the set of simultaneous equations

$$f_1(\bar{u}) = 0, \quad (5.2a)$$

$$f_2(\bar{u}) = 0. \quad (5.2b)$$

The conditions under which this homogeneous solution is stable can be found by linearising around \bar{u} . This transforms Eq. 5.1 into

$$\frac{\partial u_1}{\partial t} = (Au)_1 + \mu_1 \frac{\partial^2 u_1}{\partial x^2}, \quad (5.3a)$$

$$\frac{\partial u_2}{\partial t} = (Au)_2 + \mu_2 \frac{\partial^2 u_2}{\partial x^2}, \quad (5.3b)$$

where Au is the linearised form of $f(u)$ around the fixed point \bar{u} . In order for this fixed point to be stable it is required that the eigenvalues of A have negative real

part. This occurs when

$$\text{Tr}(A) = a_{11} + a_{22} < 0, \quad (5.4a)$$

$$\text{Det}(A) = a_{11}a_{22} - a_{12}a_{21} > 0. \quad (5.4b)$$

In order to test whether the solution is stable in the presence of spatial fluctuations a test function of the form $u = \bar{u} + \exp(\lambda t) \cos(kx)$ is applied to Eq. 5.3 and calculating the real part of λ in order to determine if the spatial fluctuations increase in time. This leads to the following condition

$$\mu_1 a_{22} + \mu_2 a_{11} > 2\sqrt{\mu_1 \mu_2 \text{Det}(A)}. \quad (5.5)$$

Hence if the system satisfies Eq 5.4a, Eq 5.4b and Eq 5.5 then the system sustains regular spatial patterns. It is important to note that Turing patterns are not the only path to pattern formation in reaction-diffusion equations and so called global patterns can also exist [Kéfi et al., 2010].

We begin by explicitly modelling the vegetation with its interaction on the environment. The vegetation propagates clonally and hence has a local dispersion process characterised by diffusion. The growth rate of the vegetation is assumed to be directly proportional to the concentration of substrate v . Local growth of the vegetation is logistic with carrying capacity $K(v) = 1 + \beta v$. The vegetation is therefore assumed to be able to be sustained when there is no substrate present ($v = 0$), this also avoids a singularity when no substrate is present. β controls the dependency of the vegetation carrying capacity on the substrate. The substrate dynamics are dominated by currents that are dependent on both the vegetation and the substrate and a diffusion term that again both depends on the vegetation and substrate. The kinetic equations are therefore

$$\frac{\partial u}{\partial t} = f(u, v) + D_1 \nabla^2 u, \quad (5.6a)$$

$$\frac{\partial v}{\partial t} = g(u, v) + \nabla \cdot (-\mathbf{j}(u, v) + D_2(u, v) \nabla v), \quad (5.6b)$$

where f has the form

$$f(u, v) = ruv \left(1 - \frac{u}{1 + \beta v} \right), \quad (5.7)$$

where r is the intrinsic growth rate, ruv is the actual growth rate in the presence of the substrate, j is the current of the substrate, D_1 is the diffusion coefficient of vegetation and D_2 is the diffusion coefficient of the substrate that may depend on the local concentration on substrate and vegetation. $g(u, v)$ is harder to understand.

We can classify it as having its own intrinsic growth rate μ and carrying capacity K that are both affected by the presence of vegetation only. Hence

$$g(u, v) = \mu u \left(1 - \frac{u}{K}\right). \quad (5.8)$$

Before considering the form of the current term $\mathbf{j}(u, v)$ and the environmental diffusion term $D_2(u, v)$, we may consider the spatially homogeneous fixed point solution of this system (u^*, v^*) by setting $f, g = 0$. Eq. 5.7 gives

$$\begin{aligned} f(u^*, v^*) &= r u v \left(1 - \frac{u}{1 + \beta v}\right) = 0 \\ \implies u^* &= 0 \text{ or } v^* = 0 \text{ or } v^* = \beta^{-1}(u^* - 1), \end{aligned} \quad (5.9)$$

using the non-zero solution and setting g to zero we obtain from Eq. 5.8

$$\begin{aligned} g(u^*, v^*) &= \mu u \left(1 - \frac{u}{K}\right) = 0, \\ \implies u^* &= 0 \text{ or } u^* = K. \end{aligned}$$

Hence the spatially homogeneous solutions are $(0, 0)$, $(1, 0)$, $(0, -\beta)$ and $(K, \beta^{-1}(K - 1))$. There is immediately a problem with these solutions. Although the non-zero solution is positive (if $K > 1$), there is a pathological solution $(0, -\beta)$, although in the absence of vegetation $g \geq 0$ and hence this fixed point is never reached if the initial conditions are in the positive cone ($u_0 > 0, v_0 > 0$). The fixed point $(1, 0)$ is also concerning as this suggests that the vegetation can sustain itself without the presence of environmental substrate.

We perform stability analysis on the fixed point $(K, \beta^{-1}(K - 1))$ by calculating the Jacobian

$$\begin{pmatrix} \frac{\partial f}{\partial u} & \frac{\partial f}{\partial v} \\ \frac{\partial g}{\partial u} & \frac{\partial g}{\partial v} \end{pmatrix} = \begin{pmatrix} r v \left(1 - \frac{u}{1 + \beta v}\right) - r u v & r u \left(1 - \frac{u}{1 + \beta v}\right) + \frac{r \beta u v}{(1 + \beta v)^2} \\ \mu \left(1 - \frac{u}{K}\right) - \frac{\mu}{K} u & 0 \end{pmatrix}. \quad (5.10)$$

Taking the determinant of the Jacobian at the fixed points $(u^*, v^*) = (K, \beta^{-1}(K - 1))$

$$\begin{aligned} \det J &= 0 - r \mu u \left(1 - 2 \frac{u^*}{K}\right) \left(1 - \frac{u^*}{1 + \beta v} + \beta \frac{v^*}{(1 + \beta v^*)^2}\right), \\ &= \mu r K (1 - 2) \left(1 - \frac{K}{1 + K - 1} + \frac{K - 1}{(1 + K - 1)^2}\right), \\ &= -\mu r K \left(1 - \frac{1}{K^2}\right). \end{aligned} \quad (5.11)$$

Since it is assumed that $\mu, r, K > 0$, the condition that the fixed point is stable when the determinant of the Jacobian is positive occurs when

$$\begin{aligned} 1 - \frac{1}{K^2} &< 0, \\ \implies 1 &< \frac{1}{K^2}, \\ \implies K &< 1. \end{aligned} \tag{5.12}$$

However, it is always assumed that the carrying capacity K is greater than one to avoid the environmental growth rate g being negative. Also note that the trace of the Jacobian should be negative if the fixed point is stable and hence

$$\text{tr}J = -\frac{r}{\beta}K(K-1) < 0. \tag{5.13}$$

As such for the fixed point to be stable $K > 1$, however for the determinant to be positive $K < 1$. Hence both conditions cannot be satisfied and the fixed point is never stable. We seek an instability in the spatially inhomogeneous solution in order to detect a Turing bifurcation that would lead to pattern formation. First the spatial components of Eq. 5.6 shall be simplified.

In order to simplify the spatial components we define new spatial co-ordinates x', y' that are rotations of the co-ordinates x, y such that the current \mathbf{j} lies in the x' direction and has no y' component. We may also reduce the number of parameters by rescaling the new co-ordinates x', y' by $D_1^{-1/2}$ and introducing the new diffusion parameter $d(u)$ and current term $j(u)$ for the environmental substrate. In the new co-ordinate system Eq. 5.6 becomes

$$\frac{\partial u}{\partial t} = f(u, v) + \nabla^2 u, \tag{5.14a}$$

$$\frac{\partial v}{\partial t} = g(u, v) + \nabla \cdot (-j(u)(1, 0)^T + d(u)\nabla v). \tag{5.14b}$$

If the diffusion and flux dependence on u is linear we may take the functional forms of these to be

$$d(u) = d \left(1 - \frac{u}{K} \right), \tag{5.15a}$$

$$j(u) = K - u. \tag{5.15b}$$

This ensures that when u is maximised, both the diffusion and flux of the environmental quality v is 0. Inserting these forms of d and j into Eq. 5.14 gives

$$\frac{\partial u}{\partial t} = f(u, v) + \nabla^2 u, \quad (5.16a)$$

$$\frac{\partial v}{\partial t} = g(u, v) - \frac{\partial u}{\partial x} - \frac{d}{K}(\nabla u) \cdot (\nabla v) + d \left(1 - \frac{u}{K}\right) \nabla^2 v. \quad (5.16b)$$

As the previous form of f and g did not produce a stable homogeneous solution in the interior ($u, v > 0$) we must consider a more general form of g . Consider g as a polynomial of order n ;

$$g(u) = a_0 + a_1 u + a_2 u^2 + \dots + a_n u^n. \quad (5.17)$$

The environmental quality should decrease in the absence of vegetation, hence $g(0) < 0$ and so $a_0 < 0$. We also desire the environmental quality to saturate at a certain carrying capacity, thus limiting the growth of the vegetation. Hence we desire $g(K) = 0$. As there are only two conditions that are imposed the minimal degree of g is 2. Take $a = -a_0$, so a is positive and $b = a_1$. Then a and b can be determined by considering the equation $g(K) = 0$

$$\begin{aligned} K^2 + bK - a &= 0, \\ K^2 + bK &= a, \\ K(K + b) &= a, \\ b &= \frac{a}{K} - K. \end{aligned} \quad (5.18)$$

Hence g has the form

$$g(u) = \mu \left[u^2 + \left(\frac{a}{K} - K \right) u - a \right], \quad (5.19)$$

and the interior fixed point is now $(u^*, v^*) = (K, \beta^{-1}(K - 1))$. The Jacobian at this fixed point is calculated to be

$$J = \begin{pmatrix} -\frac{r}{\beta}(K - 1) & \frac{r}{\beta}(K - 1) \\ \mu(K + \frac{a}{K}) & 0 \end{pmatrix}. \quad (5.20)$$

In order to deduce the conditions for the fixed point to be stable we desire $\det J > 0$ and $\text{tr } J < 0$. As such the conditions for stability are

$$-r\beta\mu(K-1)\left(K + \frac{a}{K}\right) > 0, \quad (5.21a)$$

$$-\frac{r}{\beta}(K-1) < 0. \quad (5.21b)$$

The condition in Eq. 5.21a is never satisfied as all parameters are assumed positive. Eq. 5.21b is satisfied if $K > 1$. Hence we may conclude that the fixed point $(K, \beta^{-1}(K-1))$ is never stable.

With the spatially homogeneous case dealt with we now seek a Turing bifurcation by analysing a small spatial perturbation around the fixed point. If the spatial perturbation increases over time then the spatially homogeneous solution is unstable in the sense spatial perturbations of certain characteristic wavelength will lead to the solution moving away from the homogeneous state. We therefore introduce the ansatz perturbation solution

$$u = u^* + u_\epsilon, \quad u_\epsilon = a(t) \cos(kx) \cos(ly), \quad (5.22a)$$

$$v = v^* + v_\epsilon, \quad v_\epsilon = b(t) \cos(kx) \cos(ly). \quad (5.22b)$$

Linearising around (a, b) we obtain the linear equation

$$\frac{d}{dt} \begin{pmatrix} a \\ b \end{pmatrix} = A \begin{pmatrix} a \\ b \end{pmatrix}, \quad (5.23)$$

where the matrix A is

$$A = \begin{pmatrix} k^2 + l^2 + r\beta^{-1}(K-1)(1-2K) & 0 \\ 2\mu K + \frac{a}{K} - K & d(k^2 + l^2) \end{pmatrix}. \quad (5.24)$$

Note there is no dependency on the flux j as the term $(\nabla u) \cdot (\nabla v)$ only has a non-linear part. We again consider the determinant and trace of this matrix to determine the stability of this solution, where we also for simplicity write the sum squares of the wave numbers as $k^2 + l^2 = x^2$

$$\det A = d[r\beta(K-1)(1-2K) + x^2]x^2, \quad (5.25a)$$

$$\text{tr } A = r\beta^{-1}(K-1)(1-2K) + (1+d)x^2. \quad (5.25b)$$

To determine the eigenvalues of A we solve the equation $\lambda^2 - (\text{tr } A)\lambda + \det A = 0$. These have positive real values if either $\det A < 0$ or $\text{tr } A > 0$. For the determinant

to be negative we would require that

$$\begin{aligned} r\beta(K-1)(1-2K) + x^2 &< 0, \\ \implies x^2 &< -r\beta(K-1)(1-2K). \end{aligned}$$

Since $-r\beta(K-1)(1-2K) > 0$ this is possible for certain values. For $\text{tr } A > 0$ we require that

$$\begin{aligned} r\beta^{-1}(K-1)(1-2K) + (1+d)x^2 &> 0, \\ \implies \frac{-r\beta^{-1}(K-1)(1-2K)}{1+d} &< x^2. \end{aligned}$$

We therefore have conditions under which small scale spatial perturbations would be unstable leading to pattern formation. Unfortunately, however, due to the fact that the interior fixed point is never stable, there are no pattern forming Turing bifurcations in this system.

5.2.1 A Short deviation into the literature

For the types of systems we have considered there has been no Turing bifurcation due to the interior fixed point never being stable. It is not obvious what functional form the reaction terms need to take for a bifurcation to exist. There is an issue with choosing the reaction terms without specific knowledge of what form they should take as this would lead to a model that may have similar pattern characteristics, but would have poor predictive power. Many such model mechanisms do exist in the literature, as an example from semi-arid ecosystems [Sherratt, 2005] where the environment is explicitly modelled via groundwater, an equation that does lead to stripe pattern formation is

$$\frac{\partial u}{\partial t} = u^2v - au + \nabla^2 u, \tag{5.26a}$$

$$\frac{\partial v}{\partial t} = b - v - u^2v + c\frac{\partial v}{\partial x} + d\nabla^2 v, \tag{5.26b}$$

where the environmental parameter has a flux controlled by the parameter c in the x direction. In Sherratt [2005] the authors do not consider diffusion of the environment, however this can be added in as has been done in Hille Ris Lambers et al. [2001]. The term u^2v represents plant growth subject to the environmental variable v . The term $-au$ represents vegetation mortality and is assumed constant. There is also a diffusion term (where the diffusion coefficient has been absorbed

by the spatial variables to reduce the number of parameters). The environmental parameter has a constant rate of improvement b ; for the semi-arid ecosystem this is interpreted as constant rainfall, but we can consider this more generally as any independent process that improves the environment such as influx of nutrients, soft-sediment, organic matter etc. The environment is also subject to depletion, $-v$, and a transfer to vegetation, $-u^2v$. This term represents the absorption of water, nutrients or other aspects of environmental quality in the process of vegetative growth. The system is of an activator-inhibitor type, where the environment inhibits the further growth of vegetation.

We may consider the vegetation growth term u^2v and assume that there is only linear dependency on the density of vegetation so the system of equations would be

$$\frac{\partial u}{\partial t} = uv - au + \nabla^2 u, \quad (5.27a)$$

$$\frac{\partial v}{\partial t} = b - v - uv + c \frac{\partial v}{\partial x} + d \nabla^2 v. \quad (5.27b)$$

However, there is an issue with this system. There does exist an interior stable point at $(u^*, v^*) = (\frac{b-a}{a}, a)$ when $b > a$. However, by performing the standard spatial perturbation analysis to detect a Turing phase transition the trace of the dynamic matrix linearised around the stable fixed point is always negative and hence for there to be a Turing phase transition we would require the determinant to be negative. This can only happen if $a > b$, hence contradicting the stability condition and so no Turing phase transition can occur. This shows that the non-linearity of the vegetation growth term is a requirement to induce a phase transition that would lead to pattern formation.

The stripe patterns that result from Eq. 5.26 are perpendicular to the environmental gradient induced by the $c \frac{\partial v}{\partial x}$ term. The bands also move up the environmental gradient with a velocity that is increasing with the strength of the environmental gradient c . In the case of seagrass, the environmental gradient is induced by the strength of the prevailing current. Hence c will be a function of the local current strength.

5.2.2 Conclusion

Although reaction-diffusion equations can produce regular spatial patterns in the form of banding and these bands can align perpendicular to an environmental flux,

we have in the process had to be very general about the form of the environmental reaction term g in order for banding to occur. It is difficult to, from first principles, deduce what the nature of this function g should be without proper experimentation. If the environment was simply due to the density of soft sediment then an experiment could potentially be devised. However, the term v includes other aspects of the environment that allow the vegetation to persist, which may be difficult to capture. This means that the form of g would always be a phenomenological one and would lack a mechanistic underpinning. The conclusion is that pure reaction-diffusion equations provide a good step in explaining the origin of regular pattern formation in a vegetative community with environmental feedback, however for a full mechanistic underpinning further modelling work must be sought.

5.3 Integro-differential equations

Partial Differential Equations have been effective in capturing salient details of vegetative systems [Holmes et al., 1994]. Vegetation is modelled as a density function $\rho(x, t)$ that is dependent on space and time. It is assumed that this function remains positive everywhere and furthermore it has an upper bound K denoting its carrying capacity. Spatial patterning such as patches and banding have been observed in such details and been successfully applied to vegetative patterns, namely Tiger Bush [Lefever and Lejeune, 1997].

Here a kinetic equation is introduced for the evolution of $\rho(x, t)$ taking into account processes which contribute to growth and death of a vegetation patch. The equation is given by

$$\partial_t \rho(\mathbf{x}, t) = F_1 \times F_2 - F_3. \quad (5.28)$$

The kinetic equation encapsulates aspects of growth and death that are dependent not only on the local density, ρ , at \mathbf{x} , but on long-range densities. This is a departure from normal reaction diffusion dynamics, where only local interactions are included via the diffusion operator.

F_1 represents growth due to reproduction (with intrinsic growth rate λ), this can include sexual reproduction via the production of gametes or clonal growth via propagules such as rhizomes, tubers or cuttings. In the case of local clonal growth this shall mainly be a local interaction term, however the extent of rhizome networks in vegetation such as seagrass is largely unknown and hence could involve longer range dynamics. Sexual reproduction such as via seed dispersal can be very long-

range, this would be a classic example of where nearest neighbour dynamics would break down. There is also mutual interaction, with rate Ω , due to sheltering and attachment that increases the growth rate when density is high. Spatial asymmetry can also be a factor due to prevailing winds, currents or animal migration patterns directing seed dispersal. These spatial processes are captured by a spatial kernel, $k_1(\mathbf{x})$.

F_2 encapsulates interaction between vegetation and its environment. The environment can limit the growth of a plant in a variety of ways such as competition between plants for nutrients, water and sunlight leading to a carrying capacity K . However the environment can additionally affect growth in more subtle ways. High density of vegetation can funnel and change wind and wave action, inhibiting growth in other surrounding areas by creating barriers that can change the dynamics of the environment. Again there is a local range to this interaction as well as a long range one in this term that must be taken into account with a spatial kernel $k_2(\mathbf{x})$.

F_3 represents ambient death, which depends on a number of factors related to the environment as well as local density. The simplest form of this is to assume death occurs everywhere at a constant rate η , with a spatial kernel that is the Dirac delta function $w_3(\mathbf{x}) = \delta(\mathbf{x})$. Although we wish for this equation to have the fullest generality possible, for now we shall concentrate on how they relate to the modelling of seagrass dynamics. The resulting forms of F_1 , F_2 and F_3 are

$$F_1(\mathbf{x}, t) = \int d\mathbf{y} \lambda k_1(\mathbf{y}) \rho(\mathbf{x} + \mathbf{y}, t) (1 + \Omega \rho(\mathbf{x} + \mathbf{y}, t)), \quad (5.29a)$$

$$F_2(\mathbf{x}, t) = \int d\mathbf{y} k_2(\mathbf{y}) (1 - \rho(\mathbf{x} + \mathbf{y}, t)/K), \quad (5.29b)$$

$$F_3(\mathbf{x}, t) = \int d\mathbf{y} \eta k_3(\mathbf{y}) \rho(\mathbf{x} + \mathbf{y}, t) = \eta \rho(\mathbf{x}, t). \quad (5.29c)$$

5.3.1 Derivation

We begin by modelling the environment as a separate variable $E(\mathbf{x}, t)$, with its own kinetic equation dependent on vegetation density $S(\mathbf{x}, t)$. We assume that the dynamics of this environmental variable (soft matter and nutrients) is dominated by the density of vegetation. In full generality we have two processes affecting the quality of the environment at a certain location. The first term in the rate equation for the environment variable E is the diffusion of soft sediment etc. due to the presence of seagrass, this is controlled by the kernel k_3 . The second term is the process by which sediment is prevented from piling up, this is controlled by the

absence of seagrass and the kernel k_4 . The whole rate of the process is controlled by the parameter ϵ , which represents the rate at which this process occurs relative to the seagrass dynamics. Hence the resulting equations are

$$\begin{aligned}\partial_t E(\mathbf{x}, t) &= \frac{1}{\epsilon} \left(-E \int d\mathbf{y} k_3(\mathbf{y}) S(\mathbf{x} + \mathbf{y}, t)/K + (1 - E) \int d\mathbf{y} k_4(\mathbf{y}) (1 - S(\mathbf{x} + \mathbf{y}, t)/K) \right), \\ \partial_t S(\mathbf{x}, t) &= \beta \int d\mathbf{y} k_1(\mathbf{y}) S(\mathbf{x} + \mathbf{y}, t) \times \int d\mathbf{y} k_2(\mathbf{y}) E(\mathbf{x} + \mathbf{y}, t) - \eta S(\mathbf{x}, t).\end{aligned}\quad (5.30)$$

We assume that the environmental dynamics are fast relative to the dynamics of the seagrass. Hence the leading order parameter ϵ can be assumed to be small and in the limit $\epsilon \rightarrow 0$ we have

$$E = \frac{\int d\mathbf{y} k_4(\mathbf{y}) (1 - S(\mathbf{x} + \mathbf{y}, t))}{\int d\mathbf{y} (k_3(\mathbf{y}) - k_4(\mathbf{y})) S(\mathbf{x} + \mathbf{y}, t) + \int d\mathbf{y} k_4(\mathbf{y})}. \quad (5.31)$$

We must consider the denominator of this equation (5.31). Note it is of the form $(a + x)^{-1}$, where x is small. We may hence apply a Taylor expansion $(a + x)^{-1} = a^{-1}(1 - \frac{x}{a} + O(x^2))$. Denote the integral $\int d\mathbf{y} k_4(\mathbf{y}) = \alpha$. Hence

$$E = \alpha^{-1} \left(\int d\mathbf{y} k_4(\mathbf{y}) (1 - S(\mathbf{x} + \mathbf{y}, t)) \right) \left(1 - \alpha^{-1} \int d\mathbf{y} (k_3(\mathbf{y}) - k_4(\mathbf{y})) S(\mathbf{x} + \mathbf{y}, t) \right). \quad (5.32)$$

Let us return to the discussion of the kernels k_3 and k_4 for a moment. The two kernels relate to the process of sand diffusion and sand deposition respectively. We have assumed until now that these two processes have different kernels. However, it is likely that they are on the same length scale, hence we use the simplifying assumption $k_3 = k_4$. Substituting this into equation (5.30),

$$\begin{aligned}\partial_t S(\mathbf{x}, t) &= \alpha^{-1} \beta \int d\mathbf{y} k_1(\mathbf{y}) S(\mathbf{x} + \mathbf{y}, t) \times \int d\mathbf{y} \int d\mathbf{z} k_2(\mathbf{y}) k_4(\mathbf{z}) (1 - S(\mathbf{x} + \mathbf{y} + \mathbf{z}, t)/K) \\ &\quad - \eta S(\mathbf{x}, t).\end{aligned}\quad (5.33)$$

Concentrating on the double integral term, which can be simplified using a convolution $*$ and noting their associativity

$$\begin{aligned}
P(\mathbf{x}) &= \int d\mathbf{y} \int d\mathbf{z} k_2(\mathbf{y}) k_4(\mathbf{z}) (1 - S(\mathbf{x} + \mathbf{y} + \mathbf{z}, t)/K), \\
&= \int d\mathbf{y} \int d\mathbf{z} k_2(\mathbf{z} - \mathbf{y}) k_4 * (1 - S/K)(\mathbf{x} + \mathbf{y}), \\
&= k_2 * (k_4 * (1 - S/K))(\mathbf{x}) = (k_2 * k_4) * (1 - S/K)(\mathbf{x}), \\
&= \int d\mathbf{\zeta} (k_2 * k_4)(\mathbf{\zeta}) (1 - S(\mathbf{x} + \mathbf{\zeta}, t)/K). \tag{5.34}
\end{aligned}$$

There remains a convoluted kernel $k_2 * k_4$. It is desirable for these kernels to have the usual properties of a probability distribution, namely being non-negative everywhere and integrating to 1. Assuming all kernels are Gaussian, then the convolution of the two is another Gaussian with a transformed variance. This is denoted as $h(\mathbf{y})$. Also denote the leading rate as $\lambda = \alpha^{-1}\beta$, where this can be interpreted as the birth rate. Finally denoting $k_1 = k$, gives the final form as

$$\partial_t S(\mathbf{x}, t) = \lambda \int d\mathbf{y} k(\mathbf{y}) S(\mathbf{x} + \mathbf{y}, t) \times \int h(\mathbf{y}) (1 - S(\mathbf{x} + \mathbf{y}, t)/K) - \eta S(\mathbf{x}, t). \tag{5.35}$$

This equation coincides with Eq. 5.28 when the death kernel is a delta-function. This equation therefore models the interaction between vegetation and environment when there is a separation of time scales between the environmental processes and the vegetation turnover.

5.3.2 Kernels

Let us briefly discuss the interaction kernels and their interpretation. The kernels describe the range of the processes for reproduction and competition. The simplest form of these would be a Gaussian kernel

$$k(\mathbf{x}) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{1}{2\sigma^2} (x^2 + y^2)\right) \tag{5.36}$$

We may extend this to include processes, that are anisotropic which may be induced by currents, wind and other environmental factors. The simplest way of doing this is via the skewed Gaussian, which has skewness parameter α ;

$$k(\mathbf{x}) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{1}{2\sigma^2} (x^2 + y^2)\right) \left(1 + \operatorname{erf}\left(\alpha \frac{x}{\sigma\sqrt{2}}\right)\right). \tag{5.37}$$

Longer range, sub-exponential processes may also be considered, such as power law scaling. This means the kernel would decrease proportional to some power α . The kernel would then be a Pareto distribution of the form

$$k(\mathbf{x}) = \begin{cases} \alpha x_m^\alpha x^{-(\alpha+1)} & \text{for } x \geq x_m \\ 0 & \text{if } x < x_m \end{cases}.$$

5.3.3 Anisotropy

The previous section was concerned with kernels whose domain of interaction was located over the vegetation. Environmental forces such as currents, tides and wave action introduce anisotropy into the system. We may analyse this anisotropy by considering when the competition kernel is offset i.e. the competition of the vegetation is felt by the competing vegetation at distance o away. For simplicity of analysis assume that this offset is always in the x -direction. The resulting non-dimensionalised kinetic equations are

$$\begin{aligned} \frac{\partial}{\partial t} S(\mathbf{x}, t) = & \left[\int e^{-[x_1'^2 + x_2'^2]/2L^2} S(\mathbf{x} + \mathbf{x}') d\mathbf{x}' \right] \left[1 - \int e^{-[(x_1' - o)^2 + x_2'^2]/2} S(\mathbf{x} + \mathbf{x}') d\mathbf{x}' \right] \\ & - \mu S(\mathbf{x}, t) \end{aligned} \quad (5.38)$$

Assuming Gaussian kernels, Lefever and Lejeune [1997] showed that the homogeneous solution matches the isotropic case, and the corresponding eigenvalues for a small perturbation of wavelength $\mathbf{k} = (k_x, k_y)^T$ is

$$\begin{aligned} \omega_{\mathbf{k}} = & -\mu + \mu e^{-L^2|\mathbf{k}|^2/2} - (1 - \mu) \cos(ok_x) e^{-|\mathbf{k}|^2/2} \\ & - i(1 - \mu) \sin(ok_x) e^{-|\mathbf{k}|^2/2}. \end{aligned} \quad (5.39)$$

Using this dispersion relationship, the unstable wavelengths can be estimated. unstable wavelengths appear for $L < 1$

5.4 Probabilistic cellular automata

In the previous section we defined a continuous time continuous space model that describes the kinetic evolution of seagrass interacting with its possibly inhomogeneous environment. Although the continuous limit model can be appealing, we should be aware of the plethora of interesting spatial patterning due to finite size effects. In particular, we wish to focus our attention on fractals, banding and clus-

tering patterns. This can be done by introducing a probabilistic cellular automata [Balzter et al., 1998; Ermentrout and Edelstein-Keshet, 1993; Hogeweg, 1988] based on the rates described in Eq. 5.35. A probabilistic cellular automata is a stochastic lattice model where the probability of transitioning to a new state of the lattice is dependent on the current state only, hence the model can be viewed of as a Markov chain.

The domain of the model shall be the square lattice of size N . Each lattice site shall either be occupied or not and hence the total possible number of states is $2^{N \times N}$. The system is described at time t by

$$S = \{s_{ij} : s_{ij} \in \{0, 1\}, \quad i, j \in \mathbb{Z}\}. \quad (5.40)$$

A master equation based upon the previous continuous model can be written down

$$P(S, t + 1) = P(S, t) + \sum_{S'} [w(S|S')P(S', t) - w(S'|S)P(S, t)] \quad (5.41)$$

where $w(S|S')$ are the transition rates from state S' to state S and can be defined using the birth and death rates B, D

$$w(S'|S) = \int B(\mathbf{x}, S)S'(\mathbf{x})\Delta(S + \delta_{\mathbf{x}} - S')d\mathbf{x} + \int D(\mathbf{x}, S)S(\mathbf{x})\Delta(S - \delta_{\mathbf{x}} - S')d\mathbf{x}, \quad (5.42)$$

where Δ is the delta-function defined as

$$\Delta(A) = \begin{cases} 1 & \text{if } A=0 \\ 0 & \text{otherwise} \end{cases}, \quad (5.43)$$

the birth and death rates are equivalent to the continuous model case with an added density-independent birth rate r_0 ,

$$B(\mathbf{x}, S) = \int k(\mathbf{y})S(\mathbf{x} + \mathbf{y})d\mathbf{y} \times \int h(\mathbf{y})(1 - S(\mathbf{x} + \mathbf{y})/K)d\mathbf{y} + r_0, \quad (5.44a)$$

$$D(\mathbf{x}, S) = \mu. \quad (5.44b)$$

Due to the large state space, analytical treatment of the preceding model is difficult. We therefore proceed using simulation.

5.5 Simulation

In order to perform fast simulations on the stochastic process defined in the previous section a number of approximations can be made. The model system at time t is described by the matrix in Eq. 5.40, where each lattice site s_{ij} is either occupied (1) or empty (0). The sites are updated synchronously where at each time-step, the transition probabilities for each of the birth and death events are calculated from Eq. 5.44. To calculate the birth probabilities two convolutions need to be performed for each site. If the kernels are the same size as the system, then for large system sizes this computation becomes infeasible. Since both kernels are Gaussian, their support is over the whole system. Both Gaussian kernels are approximated by having a finite cut-off where the probability is less than 10^{-4} . All the probabilities, once calculated, were multiplied by an ϵ term, where $\epsilon = 0.01$. This was used to reduce the number of sites updating per time step, in order to reduce the number of correlations occurring in the updating process. First the probabilities for transitions on the empty sites were computed, then there is a birth at each site according to those probabilities. Then the probabilities of sites transitioning to the empty state were computed. Deaths were then assigned randomly according to these probabilities and the whole lattice is updated producing the system S_{t+1} .

5.6 Results

Simulations were performed on a 256×1024 grid with fixed boundary conditions. The East boundary was fixed with occupied sites and the North, West and South boundary were fixed with empty sites. The rectangular lattice sized and fixed boundary were implemented in order to increase the strength of the banded pattern in order to study it using spectral methods. Each simulation was allowed to run until sufficient time where as to reach statistical stationarity. Each time-step is recorded and the average over the previous 100 time-steps is then taken, this provides a smoother data sequence to which spectral methods can be performed. A Fourier transform implemented using the Fast Fourier Transform (FFT) was performed for each i on the series $\mathbf{y}_i = (x_{ij})_{j=1, \dots, 1024}$ to produce the Fourier transformed series $\mathcal{F}(\mathbf{y}_i)$. The Fourier transform is complex-valued representing both sine and cosine components of the Fourier series decomposition of the data \mathbf{y}_i . Taking the absolute value of $\mathcal{F}(\mathbf{y}_i)$ gives the power spectrum, which provides an estimation of the power at varying wavelengths in the signal. These power spectra can be averaged so a smooth estimation of the power spectrum for the spatial distribution is found.

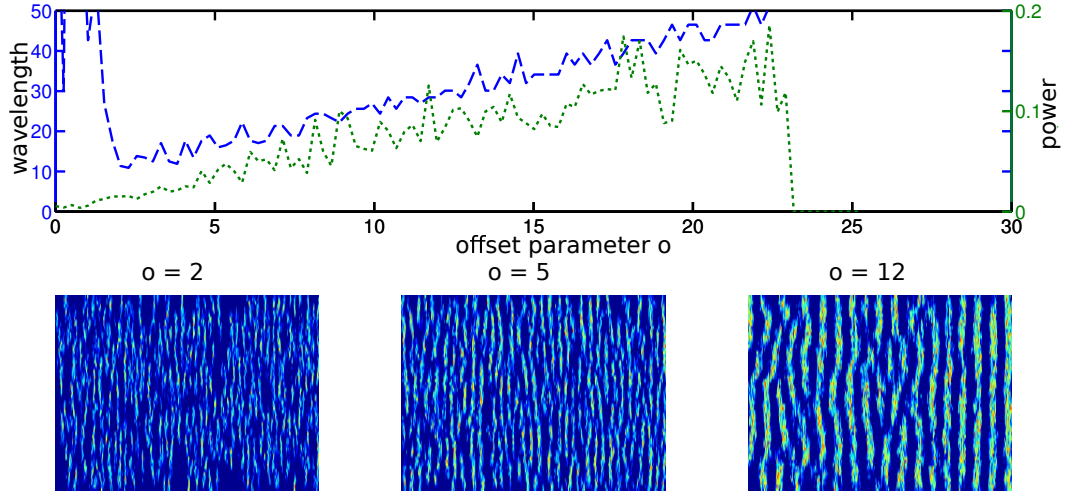


Figure 5.2: The effect of the offset parameter o on the strength and wavelength of vegetation banding. Top: wavelength (blue dashed line) and power (green dotted line) of banding for increasing offset parameters. Bottom: spatial snapshots of the resulting banding for three offset values. Colour represents the proportion of time spent occupied in 100 time-steps.

It was found that for all simulations where strong banding was visually present produce a single-peak power spectrum, therefore the wavelength and the power of the spectrum of the maximal peak were recorded to represent the banding of the spatial pattern.

For fixed kernel and demographic parameters, the offset parameter o was varied in the x -direction (Fig. 5.2). For small offset values ($o < 1$) the power of the largest peak is low, leading to the wavelength of the largest peak being subject to stochastic effects. In contrast as the offset parameter increases, both the wavelength and the power of the spatial pattern increase in an approximately linear fashion. For offsets greater than 23 the power of the banding dramatically decreases. The larger bands are therefore unable to be sustained when the offset reaches a size that is on the same length-scale of the whole system.

r_0 represents the ambient birth probability in the environment; that is the probability of a birth event on an empty site not including births due to the local spreading term. This term can therefore be interpreted as births due to seeding if the seed-dispersal process is sufficiently long-range on the scale of the system size. The effect of ambient reproduction on banding was studied by performing a number of simula-

tions for various values of r_0 , keeping other model parameters fixed at constants that lead to banding when ambient reproduction is not present ($K = 1.6, o = 15, \alpha = 1$). r_0 strongly affects the power of the banding pattern. In the absence of any ambient births, the power of the pattern is around 0.1, this increases to around 0.4 for $r_0 > 0.05$. Even for small values of r_0 there is still a significant impact on the spatial pattern. The number of topological defects in the banding pattern resulting from bands not being fully aligned is also reduced for increasing r_0 values. This leads to strong uniform bands of vegetation that are static at the centre and have small boundaries where the probability of remaining occupied is lower than in the case of lower ambient reproduction, where the probability of persisting varies more continuously through the band (Fig. 5.3). It was found that varying the α parameter had little or no affect on the resulting strength or length scale of banding.

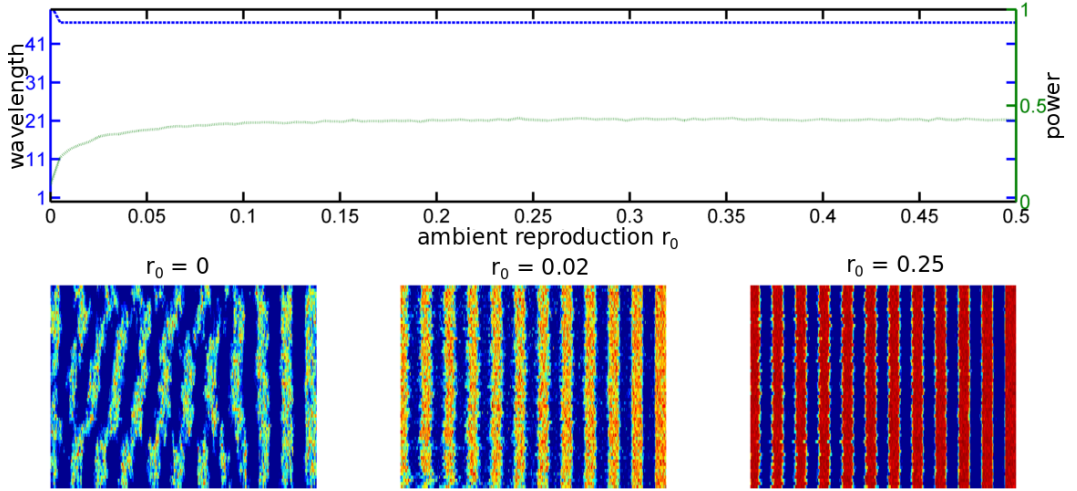


Figure 5.3: The strength of banding for simulations with varying r_0 . Increasing r_0 leads to strong uniform bands with fewer topological defects. Example snapshots are given for varying r_0 parameters and colours represent the proportion of being occupied over 100 time-steps and can be interpreted as the survival probability at each lattice site.

5.6.1 Geometry

In the previous section a number of numerical experiments were ran where the boundary of the domain was fixed with occupied vegetation sites along the East boundary. For each simulation the offset of the competition kernel was perpendicular to the line of constant vegetation. However, there can of course be situations where the offset of the competition kernel is not perpendicular to the geometry of the boundary, which can lead to frustration between the band of constant vegetation and the direction of the offset (Fig. 5.4). This can provide some insight into expected spatial patterns along coastlines, where there is a strong tendency for bands of vegetation to align along the coast, but dominant currents may flow in a different direction.

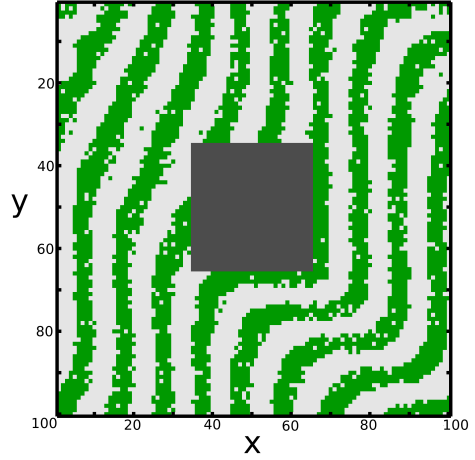


Figure 5.4: Banding in the presence of a geographic feature that does not allow vegetation to grow. Although here the competition kernel is offset along the diagonal of the system, the direction of the vegetation (shown in green) locally is more strongly affected by the geometry of the geographic feature (shown in grey). The other parameters of the system are $K = 1, r_0 = 0, \mu = 0.1, l_1 = 1, l_2 = 2$.

5.7 Conclusion

A variety of models of spatial vegetative processes have been presented and analysed. Of specific interest is the interaction between environment and vegetation that leads to regular spatial patterns such as banding. Whilst investigating these specific features, we have also tried to keep the resulting equations at the fullest generality possible. This has included keeping the number of parameters in the model to a minimum and also proposing a general model of vegetation with interaction from its environment.

The interaction between environment and vegetation is in particular a spatial one. As such, the first step in the analysis of mechanisms that generate spatial patterns was to generate a number of reaction-diffusion models that explicitly model the interaction between vegetation and its environment. Marine vegetation is affected by a number of non-local processes such as currents and wave actions, however these were ignored in the initial approach. Instead soft sediment was modelled via

a transport-diffusion process, where both the transport term and diffusion term are affected by the density of vegetation. The growth of vegetation itself depends on both the current density of vegetation and the presence of soft sediment that allows the vegetation to root and provides protection from other environmental forces. With these considerations and for specific growth terms regular patterns can form via a symmetry breaking Turing bifurcation. Although the reaction-diffusion approach has some appeal, it lacks a mechanistic underpinning for seagrass ecology in particular.

The limitations of the reaction-diffusion approach can be overcome by considering a form of vegetative growth involving an integro-differential equation, where kernels mediate both the growth and competition felt by the vegetation species. We have shown that this equation is the limiting case of a system where both environment and vegetation are explicitly modelled via integro-differential equations. The resulting model overcomes the previous limitations by allowing arbitrary choice over the spatial extent of the growth and competition terms.

The integro-differential equation provides the necessary broad spatial distribution as observed, however fails to capture the boundaries between vegetation and bare sea floor. These are due to small numbers of vegetative units near the boundaries leading to the continuous assumption breaking down. The integro-differential model was therefore converted into a discrete-time Markov process. This model still has the regular pattern formation associated with the integro-differential equation and also had the stochastic properties of the boundaries between vegetation and empty states. This model then qualitatively fits the properties of the seagrass-environment system with plausible underlying mechanisms. We may therefore use this model to explore the relationship between spatial pattern and persistence (Chapter 6); how we may validate the model by fitting to spatial pattern data (Chapter 7) and how disease impacts the dynamics and spatial pattern of vegetation (Chapter 9).

Bibliography

- Heiko Balzter, Paul W Braun, and Wolfgang Köhler. Cellular automata models for vegetation dynamics. *Ecological modelling*, 107(2):113–125, 1998.
- Ulf Dieckmann, Richard Law, and Johan AJ Metz. *The geometry of ecological interactions: simplifying spatial complexity*. Cambridge University Press, 2000.
- G Bard Ermentrout and Leah Edelstein-Keshet. Cellular automata approaches to biological modeling. *Journal of theoretical Biology*, 160(1):97–133, 1993.
- Mark S Fonseca and Susan S Bell. Influence of physical setting on seagrass landscapes near beaufort, north carolina, usa. *Marine Ecology-Progress Series*, 171:109, 1998.
- Reinier Hille Ris Lambers, Max Rietkerk, Frank van den Bosch, Herbert HT Prins, and Hans de Kroon. Vegetation pattern formation in semi-arid grazing systems. *Ecology*, 82(1):50–61, 2001.
- P Hogeweg. Cellular automata as a paradigm for ecological modeling. *Applied mathematics and computation*, 27(1):81–100, 1988.
- Elizabeth E Holmes, Mark A Lewis, JE Banks, and RR Veit. Partial differential equations in ecology: spatial interactions and population dynamics. *Ecology*, pages 17–29, 1994.
- Sonia Kéfi, Maarten B Eppinga, Peter C de Ruiter, and Max Rietkerk. Bistability and regular spatial patterns in arid ecosystems. *Theoretical Ecology*, 3(4):257–269, 2010.
- John Kuo and C Den Hartog. Seagrass taxonomy and identification key. In *Global seagrass research methods*, volume 33, pages 31–58. Elsevier, 2001.
- René Lefever and Olivier Lejeune. On the origin of tiger bush. *Bulletin of Mathematical Biology*, 59(2):263–294, 1997.

- Nuria Marba, Just Cebrian, Susana Enriquez, and Carlos M Duarte. Growth patterns of western mediterranean seagrasses: species-specific responses to seasonal forcing. *Marine ecology progress series. Oldendorf*, 133(1):203–215, 1996.
- Kenneth A Moore and Frederick T Short. *Zostera*: biology, ecology, and management. In *Seagrasses: Biology, Ecology and Conservation*, pages 361–386. Springer, 2006.
- Kenneth A Moore, Hilary A Neckles, and Robert J Orth. *Zostera marina*(eelgrass) growth and survival along a gradient of nutrients and turbidity in the lower chesapeake bay. *Marine Ecology Progress Series*, 142(1):247–259, 1996.
- Jonathan A Sherratt. An analysis of vegetation stripe formation in semi-arid landscapes. *Journal of mathematical biology*, 51(2):183–197, 2005.
- Frederick T Short. The seagrass, *Zostera Marina* l.: Plant morphology and bed structure in relation to sediment ammonium in izembek lagoon, alaska. *Aquatic Botany*, 16(2):149–161, 1983.

Chapter 6

Fractal heuristics of return rate

...once you know what the question actually is, you'll know what the answer means.

(Douglas Adams - *The Hitch-hikers Guide to the Galaxy*)

6.1 Introduction

Ecology produces a wealth of spatial patterns, including regular [Rietkerk and Van de Koppel, 2008] and scale-free [Pascual and Guichard, 2005]. Perhaps surprisingly a wide-range of environmental and underlying processes can give rise to very similar spatial patterns [Valentin et al., 1999]. Broadly speaking, a spatio-temporal system (such as a spatial vegetative system) will typically have many degrees of freedom and a model to predict individual locations would quickly become intractable for all, but trivial system sizes. This introduces the idea of using spatial observables or summary statistics to encapsulate information about the underlying dynamics of an ecological growth process [Dieckmann and Law, 2000]. The idea is that for an initial spatial pattern, a large number of the degrees of freedom will decay and the system will evolve on a sub-manifold of slow-moving parameters. This manifold is referred to as the *Relaxation Manifold* as the system can be described by a subset of spatial statistics. We can then apply the assumption that if the transient dynamics are quick, the dynamics depend entirely upon this subset of statistics. This assumption is known as a *relaxation projection*. An example of this is the mean field dynamics in a spatial birth-death process. The mean field assumes that the system depends entirely the mean density at a particular time $\rho(t)$ and ignores all spatial correlations. The relaxation manifold for this system is a one-dimensional interval on the positive real line. Nearest neighbour pair correlations can also be considered and hence the relaxation manifold would be four-dimensional, made up of the density and the three pair-correlations. Higher order correlations (triplets, quadruplets etc.) may also be considered should the dynamics still not be sufficiently explained. However, as more correlations are considered, less information on the global dynamics is generally gained for each new observable. There then exists an optimum number of statistics where the dynamics of the spatial system are almost entirely described by the list of statistics. This viewpoint shall be more thoroughly explored in Chapter 7. Another class of spatial statistics that can be considered are scaling statistics, such as the patch-size distribution or box-occupancy for increasing box size. Again, initially these distributions may have many degrees of freedom, however as the system evolves they relax onto a low-dimensional manifold where their distribution can be described by a few or even one exponent. The purpose here shall be to ascertain how these exponents relate back to the model parameters that gave rise to them and ultimately to the dynamical persistence of the system.

The theory of fractals has had a number of applications in the past thirty years in Ecology. The theory originally proposed by Mandelbrot was used to explain certain

seemingly ubiquitous patterns in nature Mandelbrot [1983]. Since, there has been tantalising speculation over using fractal theory to elucidate certain ecologically meaningful parameters from spatial patterns. The Hurst exponent (H) has been hypothesised as a first order approximation to succession in a vegetative community Hastings et al. [1982]. These approaches were summarised by Sugihara and May [1990]. The link between static and dynamic scaling has also been explored for an invading percolation cluster [Cannas et al., 2004, 2006].

Understanding the processes underlying the production of fractal structure has been another large area of research. Certain mechanisms have been identified to produce fractals including: self-organised criticality, where a system evolves into a critical state without fine-tuning of the parameters [Bak et al., 1988] ; robust criticality, where critical scaling is observed over a range of parameters in phase space [Pascual and Guichard, 2005] ; multi-scaled randomness [Halley and Kunin, 1999] ; Iterated maps/ successive branching rules [Turcotte, 1997] ; diffusion-limited aggregation [Witten Jr and Sander, 1981] ; power-law dispersal of species such as Lévy Flights [Harnos et al., 2000] ; birth-death processes, where birth is random, but death is spatially aggregated or vice versa [Shapir et al., 2000] ; and multiplicative with additive noise, this is where a combination of multiplicative and additive noise in the system produces intermittency and anomalous scaling in the time-series [Sornette, 1998], as well as in the spatial field [Benzi et al., 1993]. It is important to note that a spatial pattern without knowledge of the underlying system does not give a clear indication of what process caused it. Therefore a strong mechanistic explanation is essential.

This chapter begins with an overview of fractal growth in systems that are out of equilibrium with particular reference to how scaling in spatial statistics relates to the dynamic scaling properties of the system. The models considered have traditionally been applied to physical problems such as crystal growth and polymer formation, but it is emphasised throughout how these models can be applied to the context of vegetation growth. The next section describes several possible measures for dynamic persistence and measuring of spatial scaling for a single snapshot of vegetation based upon previous theoretical work. These measures are then applied to various plausible models of vegetation growth in various settings where the spatial environment is homogeneous, has a static noise term or is composed of a gradient and noise term. As a way of testing these measures, the simulation studies are compared to the Isles of Scilly seagrass dataset discussed in chapter 3. The applicability of using these

static measures to gain insight into the underlying persistence of the vegetation system are then discussed.

6.2 Theoretical background

Perhaps the simplest model of spatial plant growth is percolation. There are two types of Percolation: bond percolation and site percolation- the latter of the two shall be concentrated on here due to its relevance to vegetation spatial pattern. Percolation is defined on a discrete lattice of infinite size, where each site on the lattice can be occupied with probability p and empty with probability $1 - p$. For the purposes here an occupied site represents a location of a vegetative species. Each site has an independent probability of being occupied conditional on its neighbouring sites. This is obviously a simplification of a vegetative process, but it can still provide insights into more complex situations (epidemics, forest fires etc. [Bak et al., 1990; Sander et al., 2003]).

For low p the lattice is sparsely occupied and clusters, defined as sites connected by neighbours, are generally small and finite. At a certain critical point p_c , there is a phase transition, where a cluster of infinite size has a probability of forming. P_{inf} is defined to be the probability of a randomly chosen site being connected to the cluster of infinite size known as the infinite spanning cluster. For the probability of occupancy p less than the critical probability p_c , the probability of a site being in the infinite spanning cluster is zero. however for $p > p_c$, P_{inf} is non-zero and increases to one when $p = 1$. Hence P_{inf} defines an order parameter for the percolation system.

The percolation cluster has a number of interesting scaling properties, the most salient of which is the fractal property. A percolation cluster at criticality takes on the form of a fractal cluster, where the cluster is invariant under a certain scaling. This can be defined by considering the number of occupied sites M in a box of length l . For certain length scales and range of p there is the following scaling relation

$$M(l) \sim l^{d_f}, \quad (6.1)$$

where d_f is known as the mass fractal dimension of the percolation cluster. Although there are many other dimension characteristics of the percolation cluster that exist. d_f defines a static exponent of the percolation cluster, however dynamic exponents may also be defined.

The simplest dynamic model in a disordered system is the random diffusion of particles. A random walk can be defined on a percolation cluster, where the walker hops randomly around nearest neighbour sites only updating and jumping to a new site if it randomly selects a site that is occupied. If several random walkers are started from the same position and evolve over time, then the statistical properties of the random walker can be measured. A generalisation of Fick's law of diffusion for the expected mean squared displacement $r^2(t)$ can be found

$$\langle r^2(t) \rangle \sim t^{2/d_w}, \quad (6.2)$$

where d_w is the generalised diffusion exponent and describes the scaling at which a growth front of diffusing particles would grow. Remarkably, there is a connection between this dimension and the previously defined Mass fractal dimension d_f according to the Alexander-Orbach conjecture (Alexander and Orbach [1982])

$$d_w = \frac{3}{2}d_f. \quad (6.3)$$

Applied to the context of a growing vegetative system this implies that a higher mass fractal dimension d_f leads to a slower diffusion rate. To conclude, for the percolation model there is a strong link between static scaling properties of the cluster and the underlying dynamics that lead to such a pattern. In fact, it has more recently been shown that the Alexander-Orbach conjecture only holds for certain cases [Kozma and Nachmias, 2009] and in general d_w is bounded by d_f as opposed to attaining equality, [Barlow, 2004]. Nevertheless a correlation between the static fractal scaling properties and the dynamic scaling does exist.

6.2.1 Fractal growth and boundaries

Fractal growth phenomena is a very general and well-studied concept in the realms of Physics, Chemistry and Biology [Barabási, 1995]. Fractal growth deals with out of equilibrium systems where there is an irreversible growth process leading to a rough front that follows some scaling law. An example of this process is given by the Eden model, which is defined on a binary lattice with an initial random seed. From the initial seed there is a constant probability of growth p in any of the neighbouring sites that are occupied. At each time-step a site neighbouring the cluster is chosen and becomes part of the cluster. This model is out of equilibrium in the sense that the number of particles is always increasing at each time-step. Note also that there is a constant probability of a site becoming occu-

pied along the whole edge of the growing cluster, this means that holes and fjords tend to be filled in and hence the dimension of the whole cluster is equal to the embedding dimension of the space, which would be two in the case of a 2D lattice. The natural scaling in the system to study then is the boundary. Consider the model when the initial seed is a strip running along the horizontal length of the lattice. For each point along the horizontal line a height can be defined as the maximum occupied site in the vertical direction. An average height can then be defined as

$$h = \frac{1}{N} \sum_{i=1}^N h_i, \quad (6.4)$$

where N is the length of the system size being considered and h_i is the height for each strip of the lattice i (Fig. 6.1). Note the particular scaling of interest is the boundary as opposed to the bulk or mass scaling in percolation in section 6.2. A relationship that connects the dynamic scaling to the static scaling can be found by firstly defining the width of the growing boundary

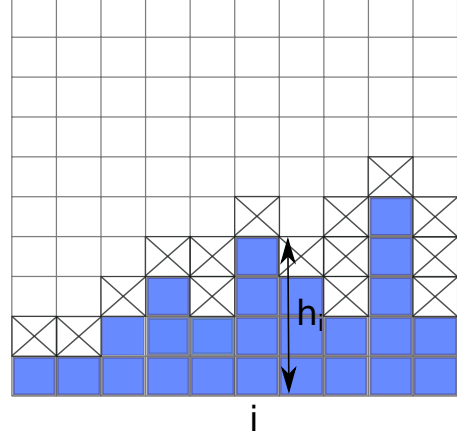


Figure 6.1: Eden model for a particular configuration at time t . The occupied sites are in blue (grey) and the potential growth sites at time $t + 1$ are highlighted with an x. The i th column is highlighted with its current height h_i . For this configuration the height at i would be 4.

$$\sigma(L, t) = \left[\frac{1}{L} \sum_{i=1}^L (\bar{h} - h_i)^2 \right]^{1/2}, \quad (6.5)$$

where the system is measured over the length L . The Family-Viscek scaling law [Family and Vicsek, 1985] gives the following form of the boundary width

$$\sigma(L, t) = \begin{cases} L^\alpha (t/L^z)^\beta & \text{if } t < L^z \\ L^\alpha & \text{if } t > L^z. \end{cases}$$

The parameter α is the roughness exponent of the system and represents the static scaling of the boundary. This can be related to the boundary dimension by $\alpha = 2 - d_b$. β is the early-time growth rate of the boundary and z is the dynamic exponent, which sets the time at which the scaling of the width saturates. These

exponents are not unrelated and for the Eden process has a simple relationship, which defines its universality class [Kardar et al., 1986].

$$z = \alpha/\beta, \tag{6.6a}$$

$$\alpha + z = 2, \tag{6.6b}$$

$$\alpha = 1/2. \tag{6.6c}$$

The system has a single exponent ($\alpha = 1/2$) that defines the scaling of the dynamic and growth exponent. This phenomena is known as universality, where the large-scale properties of the system are not dependent on the underlying neighbourhood interactions [Feigenbaum, 1983]. From the context of a growing vegetation front, the Eden model captures a homogeneous environment with a single growth rate and zero death rate. Although this may be an oversimplification, it does give some interesting results notably an exponent for a null model for comparing against growth in a heterogeneous environment.

Analytical treatment of the Eden model and other similar probabilistic cellular automata models came in the form of deriving a Langevin equation to describe the evolution of the surface growth. The Langevin approach is to model the surface growth as a continuous-time variable u , subject to a force term F and a stochastic term η that describes the microscopic noise. The first such approach, known as the Edwards-Wilkinson (EW) equation [Edwards and Wilkinson, 1982] modelled the growth of a surface $u(x, t)$ where stochastic growth is uncorrelated in space and time, but the surface is able to smooth out roughness. The resulting EW Langevin equation is

$$\frac{\partial u}{\partial t} = \nabla^2 u + \eta(x, t), \tag{6.7}$$

where u represents a growing front under the action of diffusion (∇^2) and a white noise term η with the correlation structure $\mathbb{E}[\eta(x', t')\eta(x, t)] = \delta(x - x')\delta(t - t')$. The model was extended by Kardar et al. [1986] to include a surface relaxation term

$$\frac{\partial u}{\partial t} = \nu \nabla^2 u + (\lambda/2)(\nabla u)^2 + \eta(x, t). \tag{6.8}$$

This is often referred to as the Kardar-Parisi-Zhang equation (KPZ equation). The first term represents surface relaxation or diffusion as in Eq. 6.7. The second term is growth normal to the surface and the third term is a white noise term. The inclusion of the white noise term is the same as in Eq. 6.7. For a growing vegetative species, the second term represents shooting into unoccupied substrate and the first

term represents clonal spreading. There has been success in modelling the growth of bacterial colonies with this class of equation [Vicsek et al., 1990]. The KPZ equation has the same scaling laws as the ones for the Eden model

$$z = \alpha/\beta, \tag{6.9a}$$

$$\alpha + \alpha/\beta = 2, \tag{6.9b}$$

$$\alpha = 1/2. \tag{6.9c}$$

In fact these scaling laws hold for a wide variety of models [Barabási, 1995], hence any that have the same scaling exponents are referred to as belonging to the KPZ universality class. The KPZ also assumes a homogeneous environment with each individual growth unit being indistinguishable from one another. Hence the only source of variation is from the stochasticity is due to random effects in space and time.

At a basic level, models of vegetative growth can broadly be described as diffusion processes. These can be realised as a number of particles performing a random walk in two dimensional space. Each single random walker is defined as a process $X(t)$ where X is a location of vegetative species at time t . A walker waits for a certain time τ according to a distribution $P(\tau)$ and then jumps with a magnitude of k according to a distribution $Q(k)$. A number of these non-interacting random walkers can be considered together and their trail produces a diffusion growth front. This model of vegetative growth is rather skeletal however, as certain processes in growth can be strongly directional such as in the presence of a resource gradient. The production of ramets as well produces a directed jump, where the parent plant is stationary and the daughter plant appears at the end of the ramet. Nevertheless, this model is instructive in providing insight between dynamic and static scaling for a situation that is more complex than a simple random walk. If we define a density of particles of a species at time t to be $X(t)$, then the standard deviation of these particles obeys a generalised Fick's law (Metzler and Klafter [2004])

$$\langle X^2(t) \rangle = D_\alpha t^\alpha. \tag{6.10}$$

As Fick's law is a consequence of the Central Limit Theorem, a break down of this would lead to deviating diffusion power laws. If we were to unpick the parts of the criteria for the central limit theorem to hold then we find there are two main plausible explanations as to why a natural system might violate it. These are 1. broad distributions in the waiting time or jump size and 2. the existence of long-

range correlations. In the context of a vegetative growth process, a long waiting time relates to a growth process that might be constrained by environmental factors and life history of the species. A long jump size can occur with a process such as current-distributed seeding, where new colonies can form long distances from their original source. Finally, the existence of long-range correlations could relate to a number of possible factors, such a correlation in the underlying environment attributing to growth rates of specific locations as well as external factors, such as weather, leading to a correlated disturbance across a large spatial range.

Anomalous diffusion has been widely studied in recent years [Klafler and Sokolov, 2005]. This is where the mean squared displacement of a growth front is non-linear in time. For example, a heterogeneous environment leads to a form of diffusion where the increase in variance is slower than expected in a homogeneous environment. The growth process of the interface can then be described by the following Langevin equation [Metzler and Klafter, 2000]

$$\frac{\partial}{\partial t} h(x, t) = \frac{\partial^\nu}{\partial \|x\|^\nu} h(x, t) + \eta(x, t). \quad (6.11)$$

Here $h(x, t)$ is the height of the growth front compared to some reference point (such as the centre of a patch), x is the spatial reference of the growth interface. The diffusion part of the process is given by the Riesz fractional derivative [Agrawal, 2007]. The parameter ν essentially determines the correlations in the diffusion process. When ν is a non-integer value the process becomes non-markovian i.e. the system has memory of its past states. η is a white noise process, which can be uncorrelated or correlated, and drives the growth of the boundary.

Anomalous diffusion processes as described in the previous paragraph can be formulated as fractional diffusion processes. This is where the standard diffusion operator ∇^2 is replaced by a more generalised version of the differential operator. A fractal growth process with general waiting time distribution characterised by a power law tail with exponent γ and jump size distribution with a power law tail of exponent ν can be formulated as a Langevin equation of the following type

$$\frac{\partial^\gamma h}{\partial t^\gamma} = k \frac{\partial^\nu h}{\partial |x|^\nu} + \eta(x, t), \quad (6.12)$$

where k is a diffusion constant and η is a white noise term with 0 mean and covariance of the form $\langle \eta(x, t) \eta(x', t') \rangle = 2D \delta(x - x') \delta(t' - t)$. Where D is a constant

and $\delta(\cdot)$ is a delta function. The details of these pseudo-differential operators shall not be discussed here (See Leith [2003] for a review). The important property to note is that the differential operators are non-local and hence take into account longer range correlations than in the standard diffusion process. This equation is divided into two cases to review its scaling properties; the first is when $\gamma = 1$. This recovers the usual time derivative of a diffusion equation and hence the anomalous aspect come from the spatial operator characterised by ν , where $0 < \nu \leq 2$. The second case is when $\nu = 2$; this is where there is a standard spatial diffusion operator and the anomalous behaviour is purely from the time derivative. Each case produces characteristically different behaviour, the first is known as super-diffusion, where there exists longer range jumps than normal diffusion and hence the diffusion rate is greater. The second case is known as sub-diffusion where there exists longer waiting time between jumps than standard diffusion leading to a slower diffusion rate. The results can be summarised by the roughness exponent α

$$0 < \alpha < 1/2 \text{ sub-diffusion,} \quad (6.13a)$$

$$\alpha = 1/2 \text{ normal diffusion,} \quad (6.13b)$$

$$1/2 < \alpha < 1 \text{ super-diffusion.} \quad (6.13c)$$

$$(6.13d)$$

Sub-diffusion leads to a rougher boundary than expected and super-diffusion leads to a smoother boundary than expected from standard diffusion.

Scaling analysis may be performed on Eq. 6.12 in order to determine the relationship between the dynamic and static parameters. This is where an ansatz solution of the form $h(x, t) \sim t^\beta f(x/t^{1/z})$ for some function f is assumed. The exponents α, β and z completely determine the fractal growth process of the boundary. β is the growth exponent as $h(t) \sim t^\beta$ for early time. z is the dynamic scaling exponent, and controls the time to saturation of the growth interface. Finally, α is the roughness exponent and characterises how the boundary fluctuates in the spatial component.

Scale analysis can be carried out on Eq. 6.11. This leads to the following exponents

$$\alpha = \frac{\nu - 1}{2}, \quad \beta = \frac{\nu - 1}{2\nu}, \quad z = \nu. \quad (6.14)$$

ν completely determines all the scaling parameters. Hence once one scale parameter is found the rest are necessarily determined. α is related to the fractal box-counting

dimension of the boundary d_b via the simple relationship

$$\alpha = 2 - d_b. \quad (6.15)$$

A similar scaling analysis can be performed on Eq. 6.12 with $\nu = 2$ to give the following exponents [Leith, 2003]

$$\alpha = \frac{3}{2} - \frac{1}{\gamma}, \quad \beta = \frac{3\gamma}{4} - \frac{1}{2}, \quad z = \frac{2}{\gamma}. \quad (6.16)$$

Anomalous diffusion can lead to a one- or many-parameter class that defines the relationships between the static and dynamic exponents of a growing front. The continuous parameters ν and γ can be used to explain the apparent variation of a growth front for different environmental conditions. Note that this is in contrast to the EW and KPZ equation ($\nu = 2$, $\gamma = 1$) where there is no dependency on a free parameter and hence the roughness of the boundary is constant in the limit for any realisation of the system. The key idea here is that if the underlying mechanism is the same then relationships such as those in Eq. (6.16) and Eq. (6.14) should be resolvable with enough data.

6.2.2 Relative patchiness

Vegetation dynamics due to demographic and environmental factors produce spatial patterns comprised of patches of varying sizes. It has often been observed that these patterns can produce power laws in both patch size distributions, patch-perimeter distributions and the patch radius of gyration distributions. Power law patch size distributions have been observed in Mussel beds [Guichard et al., 2003], power law gap size distributions in wind-disturbed forests [Kizaki and Katori, 1999] and power law patch size of fire-disturbed forests [Malamud et al., 1998]. The Korcak exponent is defined by considering the patch-size distribution. If this is of the form of a power law i.e.

$$N(A \geq a) = ka^{-K}, \quad (6.17)$$

K is defined to be the Korcak exponent, where $0 \leq K \leq 1$. The Korcak relationship Mandelbrot [1983] gives the following relationship to the Hausdoff dimension

$$K = \frac{D_H}{D_E}, \quad (6.18)$$

where D_E is the Euclidean dimension of the embedding space, hence would be 2 in this context.

Recent studies have shown that in general the Korcak relationship does not hold [Imre et al., 2012]. Although studies indicated previously that there is a class of self-iterating maps for which the relationship does exist [Hastings and Sugihara, 1993], and this claim has been made in several other articles [Seuront, 2009]. This leads to the conclusion that although the Korcak exponent does relate to the Hausdorff dimensions in certain cases, there is no universal relationship.

6.3 Models

6.3.1 Introduction

Two plausible models of vegetation growth are introduced in this section in order to ascertain how more realistic vegetation growth than simply percolation and other models discussed in the background impacts the relationship between persistence characteristics and scaling heuristics such as the boundary dimension and the patch-size distribution. The first model is equivalent to the PCA model discussed in Chapter 5, where the probability of a birth at an unoccupied site is dependent on a growth kernel and a competition kernel. For brevity, an offset to the competition kernel is not considered here.

The second model includes a third state where the vegetation is in a degraded or recovering state. The first state is an active phase where the plant is established and able to seed or shoot in order to reproduce. The second is a rested or dead phase, where the plant has lost the ability to reproduce, but still occupies the site preventing invasion from surrounding vegetation. An example of this in coral, where a coral can undergo an acute disturbance known as bleaching. When in the bleached state a coral can recover, but suffers from higher mortality during this process [Brown, 1997].

6.3.2 Two-state model

The model is derived in chapter 5 and based upon similar assumptions to the hydrodynamic model of vegetative growth given in Lefever and Lejeune [1997], however a summary of it shall also be given here. The model is defined on a square $N \times N$ lattice denoted Ω , where each site can either be occupied (1 for short-hand) or unoccupied (0). Dirichlet boundary conditions are implemented such that the boundary is held constant at 0 for any site outside of Ω . The model is updated synchronously, hence at each time step there is a certain probability of each site flipping states.

Death is considered to be constant i.e. for each site that is occupied there is a constant probability of death μ . For a state to become occupied there are two factors that are included in the model: local reproduction, which is mediated by a Gaussian kernel with 0 mean and variance σ_B ; and long-range competition, which is mediated by another Gaussian kernel with mean 0 and variance σ_C . For a site that is unoccupied; the probability of transitioning to an occupied state at the next time step is

$$(k_B * \Omega)(1 - k(k_C * \Omega)), \quad (6.19)$$

where the two-dimensional discrete convolution term $f * g$ is defined as

$$(f * g)[x, y] = \sum_{n_x=0}^{N-1} \sum_{n_y=0}^{N-1} f[n_x, n_y] \times g[x - n_x, y - n_y]. \quad (6.20)$$

The dynamics can then be summarised as follows

$$P(0 \rightarrow 1) = (k_B * \Omega)(1 - k(k_C * \Omega)), \quad (6.21a)$$

$$P(1 \rightarrow 0) = \mu. \quad (6.21b)$$

6.3.3 Three-state model

Many growth processes involving sessile units such as mussels, coral and semi-arid ecosystems [Guichard et al., 2003; Hoegh-Guldberg, 1999; Kéfi et al., 2007] exhibit three-stages of colonisation: empty, occupied and resting/disturbed. An empty site is where no vegetation exists, but is susceptible to becoming colonised. The occupied stage is where vegetation has shoots and taken root and represents the entire active stages of the vegetation. The resting state is where vegetation has died back, but new shooting or growth cannot take place due to conditions left by previously occupied vegetation. In the context of seagrass clonal growth the three states may therefore be interpreted as

1. Unoccupied soft substrate state, where seagrass has the ability to invade (S).
2. Seagrass has invaded the soft substrate with rhizomes and shooting (I).
3. Shooting has died back leaving dead or dormant rhizomes layer that is not able to invade by new seagrass shooting (R).

The three states correspond similarly to the three states in the Susceptible-Infected-Recovered epidemic modelling [Anderson and May, 1991]. Here the modelling is

conducted on a lattice $L = N \times N$, with a state space $\Omega = \{S, I, R\}^{2^N}$ and a set of probabilities describing the transitions of the system as follows

$$P(S(x) \rightarrow I(x)) = \lambda(x) \times \sum_{y \in n(x)} I(y), \quad (6.22a)$$

$$P(I(x) \rightarrow R(x)) = \mu_1, \quad (6.22b)$$

$$P(R(x) \rightarrow S(x)) = \mu_2, \quad (6.22c)$$

where $n(x)$ are the neighbours of x using the Von Neumann neighbourhood structure. The fecundity rate $\lambda(x)$ is a static environment variable representing the variability of the environment. This is taken as the inverse of a waiting time distribution that is heavy-tailed. i.e.

$$\lambda(x) = \frac{1}{w(x)} \quad \text{where} \quad w(x) \sim \text{Pareto}(\theta, w_0 = 1). \quad (6.23)$$

The heavy-tailedness of the distribution is a salient detail, we have chosen the pareto distribution here for convenience. Some properties of the distribution are

$$\langle w \rangle = \frac{\theta}{\theta - 1}, \quad (6.24a)$$

$$\text{Var}(w) = \frac{\theta}{(\theta - 1)^2(\theta - 2)}. \quad (6.24b)$$

It is simple to show that for $\theta < 1$ the mean does not exist and for $\theta < 2$ the second moment and hence the variance does not exist. To calculate the moment properties of λ then, proceed using the substitution $w = \langle w \rangle - z$. Using a geometric expansion

$$\begin{aligned} \frac{1}{w} &= \frac{1}{\langle w \rangle - z}, \\ &= \frac{1}{\langle w \rangle} \frac{1}{1 - z/\langle w \rangle}, \\ &= \frac{1}{\langle w \rangle} \left(1 + \frac{z}{\langle w \rangle} + \left(\frac{z}{\langle w \rangle} \right)^2 + O \left(\left(\frac{z}{\langle w \rangle} \right)^3 \right) \right). \end{aligned}$$

Hence the mean of the inverse can be approximated by

$$\langle \lambda \rangle = \langle w^{-1} \rangle = \frac{1}{\langle w \rangle} + O(\text{Var}(w)/\langle w \rangle^3), \quad (6.25)$$

hence $\langle \lambda \rangle \geq 1/\langle w \rangle$, which may be obtained from Jensen's inequality. Similarly the variance is approximated by

$$\text{Var}(\lambda) \approx \text{Var}(w)/\langle w \rangle^4 = \frac{(\theta - 1)^2}{(\theta - 2)\theta^3}. \quad (6.26)$$

An environmental gradient was also considered for the model described in Eq. 6.22. This is achieved by assuming the fecundity rate $\lambda(x)$ has the form

$$\lambda(x) = \gamma(1 - x) + \xi \frac{\gamma}{2} \zeta(x, y), \quad (6.27)$$

where γ is the gradient of the environment, ξ controls the strength of the noise term and $\zeta(x, y)$ is an i.i.d random variable drawn from the standard normal distribution. The $\gamma/2$ factor is used as it is the mean of λ in the presence of no noise. The master equation representing the probability of observing a state Ω at time t .

$$\begin{aligned} \frac{dP(\Omega, t)}{dt} = & \sum_{x' \in \Omega} [r(S(x') \rightarrow I(x'))P(\Omega(S(x')), t) + r(I(x') \rightarrow R(x'))P(\Omega(I(x')), t) \\ & + r(R(x') \rightarrow S(x'))P(\Omega(R(x')), t)] - \\ & - \sum_{x' \in \Omega} [r(S(x) \rightarrow I(x))P(\Omega(S(x)), t) + r(I(x) \rightarrow R(x))P(\Omega(I(x)), t) \\ & + r(R(x) \rightarrow S(x))P(\Omega(R(x)), t)]. \end{aligned}$$

Rate equation

From the three rates the expected change in density of occupied sites is for a given density ρ may be calculated. This is done by first considering what the expected rate of change is for an individual site x .

$$\begin{aligned} \frac{\partial}{\partial t} \mathbb{E}[I(x)] &= (-1) \times r(I \rightarrow R) + (+1) \times r(S \rightarrow I) \\ &= -\mu_2 + \lambda(x) \sum_{y \in n(x)} I(y). \end{aligned}$$

The expected rate of change is locally correlated by the number of neighbours that are infected. In order to proceed analytically the mean field assumption is applied. First note that since all $\lambda(x)$ are independent, the expectations with the sum of the infected neighbourhoods can be separated and the expectation of the sum is simply

four times the population density. Adding this together produces

$$\begin{aligned}
\frac{\partial}{\partial t} \mathbb{E}[\rho_I] &= \frac{1}{N^2} \sum_{x \in \Omega} \frac{\partial}{\partial t} \mathbb{E}[I(x)] \\
&= \sum_{x=0}^{N^2} -\mu_2 + \mathbb{E}[\lambda(x)](4\rho) \\
&= -\rho\mu_2 + 4\bar{\lambda}\rho_S.
\end{aligned}$$

If the rate of going from recovered to susceptible is high i.e. $\mu_2 \gg \mu_1$ then the density of the recovereds can be ignored and $\rho_S = (1 - \rho_I)$ and hence,

$$\frac{\partial}{\partial t} \mathbb{E}[\rho_I] = \rho_I(\bar{\lambda} - (\mu_2 + \bar{\lambda})\rho_I). \quad (6.28)$$

This forms a binomial function in terms of the density. By setting the rate of change to zero, the positive equilibrium density ρ_I^* is then calculated as

$$\rho_I^* = \frac{\bar{\lambda}}{\mu_2 + \bar{\lambda}}. \quad (6.29)$$

The gradient at this point then gives the return rate

$$\begin{aligned}
\frac{\partial}{\partial t} \mathbb{E}[\rho_I^*] &= \bar{\lambda} - 2(\mu_2 + \bar{\lambda})\rho_I^* \\
&= \bar{\lambda} - 2\bar{\lambda} \\
&= -\bar{\lambda}.
\end{aligned}$$

Under the mean field assumption the expectation of the environmentally-determined growth rate $\bar{\lambda}$ completely determines the return rate and there is no dependency on the variance or higher-order moments of the environmental variable. This approximation shall be compared to simulation in order to determine its accuracy.

6.4 Method

6.4.1 Calculation of fractal dimensions

Boundary dimension

The boundary dimension of a growing cluster was calculated as follows. Firstly a flood-fill algorithm is performed from the base of the growing cluster in order to determine the connected lattice sites (this ignores small islands that may be growing

away from the main growing cluster). Flood-fill was used again in order to fill in the gaps of the main cluster. This leaves a distinct edge between the cluster and the unoccupied sites. From this a boundary can be traced out as the lattice sites that border both an occupied and unoccupied site.

Box-counting was then performed on this boundary object. A series of convolutions with unit matrices of varying size $l \times l$ were run over the lattice. All non-zero entries of the convolution were summed together in order to calculate $N(l)$ the number of occupied boxes of length l . Linear regression can then be performed on the graph $(\log(1/l), \log N(l))$. The gradient of this linear regression is then used as the estimator for the boundary dimension d_b .

Korcak dimension

For a given lattice configuration Ω in one of two states 0 or 1 for each site the Korcak dimension is calculated as follows. Firstly a flood-fill algorithm is performed over the lattice to determine the size of each of the clusters. This produces the sample $x = \{x_1 \dots x_n\}$. The strategy is to fit a Pareto distribution of the form

$$P(X = x) = \begin{cases} \frac{\alpha x_m^\alpha}{x^{\alpha+1}} & \text{if } x \geq x_m \\ 0 & \text{if } x < x_m \end{cases} \quad (6.30)$$

α represents the scaling exponent and is the Korcak exponent in this context. Hence the distribution is fitted to the data x such that the estimated value $\hat{\alpha}$ is found. A maximum likelihood estimator approach was implemented [Clauset et al., 2009] by first estimating \hat{x}_m to be $\min\{x_1, \dots, x_N\}$. The maximum likelihood can then be calculated by setting the partial derivative of the likelihood with respect to α to zero. This then produces the estimate for α as

$$\hat{\alpha} = \frac{N}{\sum_{i=1}^N (\log x_i - \log \hat{x}_m)}. \quad (6.31)$$

The Korcak dimension was then estimated to be twice this. i.e.

$$K = 2\hat{\alpha}. \quad (6.32)$$

6.4.2 Calculation of the return rate

Introduction

A concept of a return rate on a lattice system can be defined in a number of ways. Two methods were explored for measuring the rate by measuring the time to return to equilibrium following a perturbation or by measuring the fluctuations around equilibrium. A naive method was developed based on the time to return to an equilibrium population number following a disturbance. This method was extended by considering the stochastic fluctuations of density around the population equilibrium.

Return rate following perturbation

A lattice simulation is initialised with a random distribution of a certain population density ρ_0 . The system then evolves until equilibrium is reached. The equilibrium point is measured as the first point where the averaged population derivative becomes negative. The system is allowed to evolve for another 1000 time-steps. It is then hit with a disturbance where 20% of the population is randomly killed. The system will then evolve back to equilibrium again and the time taken for it to reach within 1% of the equilibrium population is recorded as τ_e . the return rate following a disturbance rr_d is then defined as

$$rr_d = \frac{1}{\tau_e}. \quad (6.33)$$

There are a number of issues with this method. If the population is low then a disturbance of 20% will affect a small number of sites. This also means that the return time will be much shorter than a system with an equilibrium at a much higher population count. This can be a problem for any definition of return rate. For a population with a small but stable equilibrium, the population density is bounded from below and hence its fluctuations around equilibrium will encompass more of the viable population space than a system at a higher equilibrium.

Return rate at equilibrium

The return rate for a lattice system at equilibrium is defined to be the expected change in density around the equilibrium density. Suppose there is a fixed spatial pattern at time t , denoted Ω^t . A single realisation of the system one step ahead is denoted as the configuration Ω^{t+1} . The dash here is used to denote the fact that this is only one of multiple possible configurations that would result from the the dynamics being run on the same configuration Ω^t . A simple method of measuring

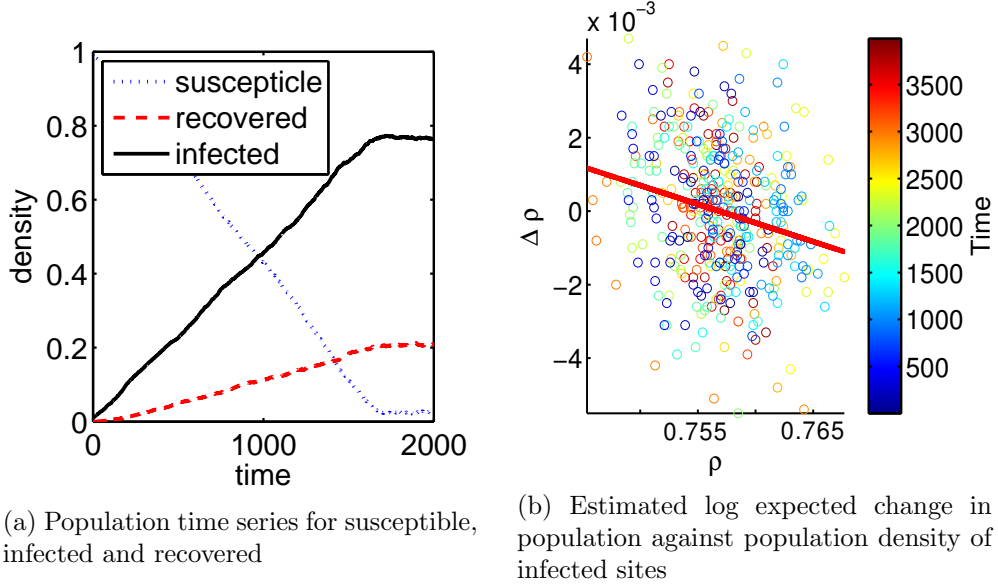


Figure 6.2: Example output simulation of three-state lattice model with $\sigma = 1$, $\mu_1 = 0.1$, $\mu_2 = 0.1$ over a lattice of size $N = 100$. The system reaches statistical equilibrium shortly after 1000 time-steps. However as can be seen in Fig. 6.2b there is still large variation in the expected change in population around the equilibrium level. The return rate at equilibrium is estimated to be $rr = 0.006$.

the change in density would be to calculate the difference in density between two consecutive time-steps $\Delta\rho_{t+1} = \rho_{t+1} - \rho_t$. The stochastic nature of the system, however, implies that $\Delta\rho_{t+1}$ is itself a random variable of unknown variance. For time-series data, this is the best estimator we can use to determine the return rate, as there is only one realisation of the process (Fig. 6.2). For simulation data there are methods that can reduce the variance of this estimate. One method would be to simulate multiple Ω^{t+1} from a fixed configuration Ω_t in order to determine the expected change in density $\mathbb{E}[\Delta\rho_{t+1}]$, however this would be computationally costly as the configuration Ω^{t+1} would have to be drawn many times to achieve a good quality estimate.

$\mathbb{E}[\Delta\rho|\Omega^t]$ can be determined directly from the simulation for a given configuration Ω^t . The probability of a birth event at lattice site (i, j) is $P_{ij}(E_B|\{\Omega^t\})$ and a death event is $P_{ij}(E_D|\{\Omega^t\})$, the expected change in density can be calculated exactly as

$$\mathbb{E}[\Delta\rho|\Omega^t] = \sum_{i=0}^{n-1} \sum_{j=0}^{n-1} [P_{ij}(E_B|\{\Omega^t\}) - P_{ij}(E_D|\{\Omega^t\})]. \quad (6.34)$$

For a fixed density ρ there are a number of corresponding configurations Ω with this density. In particular for a two-state system on an m by n lattice, the number of states would be given by

$$\binom{mn}{mn\rho}. \quad (6.35)$$

In general it would be computationally intractable to calculate the appropriate event probabilities for every configuration of a given ρ . Instead it is assumed that around the equilibrium point ρ^* , the function $\mathbb{E}[\Delta\rho|\rho]$ is assumed to be approximately linear. A simulation given an initial configuration Ω^0 is then run until it reaches statistical equilibrium. Once it has, the datum $(\rho(\Omega^t), \mathbb{E}[\Delta\rho|\Omega^t])$ is recorded for a given number of generations N . Linear regression is then performed on the data set and the gradient of regression is taken to be the return rate.

Example of return rate calculation on two-state non-spatial model

In order to determine the accuracy of the return rate calculation discussed in section (6.4.2) the method is performed on a non-spatial birth-death model where the return rate can be calculated exactly and compared with the return rate calculated from the simulation.

A population of N sites is taken. A site has a probability of a birth event occurring b and a probability of a death event d . A birth event can only occur at a site if it is unoccupied and a death event can only occur if the site is occupied. Hence the model is interpreted as N Bernoulli processes running synchronously. The density ρ is defined to be the proportion of sites occupied i.e. $N\rho$ is the number of occupied sites. Hence the probability of n births is

$$P(B = n) = \binom{N(1-\rho)}{n} b^n (1-b)^{N(1-\rho)-n}, \quad (6.36)$$

and the probability of m deaths is

$$P(D = m) = \binom{N\rho}{m} d^m (1-d)^{N\rho-m}. \quad (6.37)$$

The expected number of births is then derived as $\mathbb{E}[B] = N(1-\rho)b$ and the expected number of deaths is $\mathbb{E}[D] = N\rho d$. The expected change in density is therefore

$$\begin{aligned}\mathbb{E}[\Delta\rho] &= \mathbb{E}[B - D], \\ &= \mathbb{E}[B] - \mathbb{E}[D], \\ &= b(1 - \rho) - d\rho.\end{aligned}\tag{6.38}$$

This function shows that the expected rate of change is linearly dependent on density for this model. Setting the expected rate of change to zero determines the expected density at equilibrium to be

$$\rho^* = \frac{b}{b + d}.\tag{6.39}$$

The expected change in density can be differentiated with respect to the density in order to determine the return rate

$$\frac{\partial}{\partial\rho}\mathbb{E}[\Delta\rho] = -(b + d).\tag{6.40}$$

The expected change in density can also be calculated at $\rho = 0$, here the change in density is always positive and is equal to the birth probability b , since at this density the dynamics reduces to a pure birth process. The return rate is dependent purely on the magnitude of both birth and death at equilibrium. Hence a system can have the same equilibrium value but two different return rate values. For example if $(b_1, d_1) = (0.2, 0.1)$ and $(b_2, d_2) = (1, 0.5)$ then $\rho_1^* = \rho_2^* = 0.3$. However, $r_1 = -0.3$ and $r_2 = -1.5$.

A birth-death process was simulated in order to determine the accuracy of the calculated return rate. This was done by initialising the simulation with a random configuration of occupied sites. the model dynamics were then run forward until the equilibrium density had been established and the density was fluctuating around the equilibrium. The expected rate of change in the density as calculated from Eq. 6.38 along with the density was recorded for 2000 time-steps. The theoretical expected change in density was then compared to the change in density as calculated from the formula $\Delta\rho_{t+1} = \rho_{t+1} - \rho_t$. These two estimates of $\Delta\rho$ were then used to calculate the return rate via linear regression. There was found to be a close agreement between the two calculated return rates for this example (Fig. 6.3)

The next step would be to look at a three state non-spatial model with density-dependence. This would be similar to the final model, with the main difference being that the rate of infection is not spatial and as such simply depends upon the global density of infecteds. The infection kernel clearly introduces local correlations that would need to be corrected for should a full analytical treatment of the model take place. For now we shall ignore this dilemma and return to calculating the general dynamics of the non-spatial model in similar fashion as was done at the beginning of this chapter.

Example of return rate calculation on three-state non-spatial model

Similarly to section (6.4.2), the return rate for a three-state birth-death-recovery model with no spatial interaction was calculated analytically. This is equivalent to a mean field of the SIR model introduced in section (6.3.3).

A population of N sites is considered. Each site can be in one of three states $\{S, I, R\}$. For a site that is in the susceptible S state, there is a probability of a birth event b at the next time-step. Equivalently, for an occupied site there is a probability of a death event d . For the dead state R , there is a probability of a recovery event r . there are then N Bernoulli processes at each time-step each with a probability that is dependent upon the current state of the individual site. The density is now defined to be the density of the the occupied sites i.e. ρ_I , similarly the density for dead sites is ρ_R and empty sites as ρ_S . The number of occupied, dead and empty are $N\rho_I, N\rho_R, N\rho_S$ respectively. In the next time step the birth,

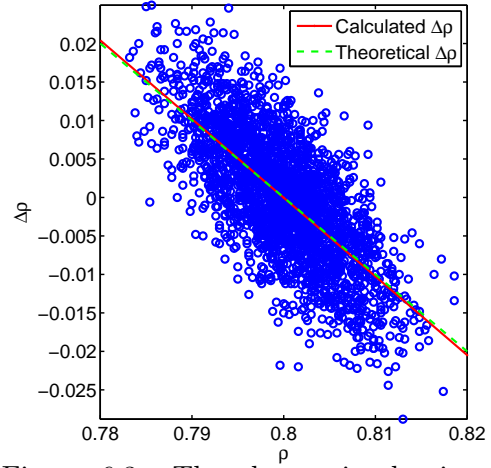


Figure 6.3: The change in density $\Delta\rho$ is plotted against the density ρ for 2000 time-points after equilibrium is established. The theoretical relationship between the expected $\Delta\rho$ is plotted in green and the calculated expected $\Delta\rho$ (in red) is calculated using linear regression on the data points. There is a close agreement between the calculated and theoretical values, indeed the calculated return rate for this example was -1.0225 , close to the theoretical value of -1 . (Here parameters are $b = 0.8, d = 0.2$, hence $\rho^* = 0.8$).

death and recovery events are all independent and thus can be written as

$$P(B = k) = \binom{N\rho_S}{k} b^k (1 - b)^{N\rho_S - k}, \quad (6.41a)$$

$$P(D = l) = \binom{N\rho_I}{l} b^l (1 - b)^{N\rho_I - l}, \quad (6.41b)$$

$$P(R = m) = \binom{N\rho_R}{m} b^m (1 - b)^{N\rho_R - m}. \quad (6.41c)$$

There are three occupation states and hence there are three expected changes in density that can be calculated

$$\mathbb{E}[\Delta\rho_I] = \mathbb{E}[B] - \mathbb{E}[D] = b\rho_S - d\rho_I,$$

$$\mathbb{E}[\Delta\rho_R] = \mathbb{E}[D] - \mathbb{E}[R] = d\rho_I - r\rho_R,$$

$$\mathbb{E}[\Delta\rho_S] = \mathbb{E}[R] - \mathbb{E}[B] = r\rho_R - b\rho_S.$$

By setting the expected change in density of each state to zero, the equilibrium point can be calculated as

$$\rho_I^* = \frac{rd}{bd + r(b + d)}, \quad (6.42a)$$

$$\rho_R^* = \frac{br}{bd + r(b + d)}, \quad (6.42b)$$

$$\rho_S^* = \frac{bd}{bd + r(b + d)}. \quad (6.42c)$$

For this system the rate of change around the equilibrium value of the occupied density ρ_I can then be calculated. In the three-state system this is now a multi-valued function dependent on both ρ_I and ρ_S . Again, this function can be assumed to be approximately linear around the fixed point and so find its gradient in the ρ_I co-ordinate. ρ_S can be rewritten as $\rho_S = 1 - \rho_R - \rho_I$. ρ_R can also be found at equilibrium in terms of ρ_I . Hence the return rate is calculated as

$$\begin{aligned} \frac{\partial \mathbb{E}[\Delta\rho_I]}{\partial \rho_I} &= \frac{\partial}{\partial \rho_I} (b(1 - \rho_R - \rho_I) - d\rho_I), \\ &= \frac{\partial}{\partial \rho_I} (b(1 - \rho_I - \frac{d}{r}\rho_I) - d\rho_I), \\ &= - \left(b + d + \frac{bd}{r} \right). \end{aligned}$$

Compare this with the return rate in Eq. 6.40 derived for the non-spatial birth-death model. There is now an extra factor bd/r , strictly greater than 0 implying that a three-state model has a higher return rate than the equivalent two-state model. Note that for a small recovery rate, the return rate is unbounded and decreasing. We may interpret this as a system entirely populated by dead sites is a quasi-absorbing state of the system i.e. as r decreases to zero the probability of exiting the quasi-absorbing state also decreases to zero.

In order to measure the equilibrium return rate for a simulation there needs to be an estimate for when the simulation first reaches statistical stationarity. This is estimated by recording the density at a number of time-steps as

$$\mathbf{R} = \{\rho_t : t \in \{0 \dots N\}\}.$$

The time at which the maximum occurs for this set is t_{\max} and the value ρ_{\min} is defined as

$$\rho_{\min} := \min\{\rho_t \in \mathbf{R} : \rho_t \geq \rho_{t_{\max}}\}.$$

The corresponding time at which the value ρ_{\min} is first achieved is denoted as t_{\min} . The set of points that are assumed to be at stationarity from which inference can be performed are then defined as all points in the set \mathbf{R} that are at time greater than t_{\min} .

6.4.3 Comparison to data

In order to compare the derived static and scaling properties derived from the models the eelgrass dataset from the Isles of Scilly, UK (IoS) as discussed in chapter 3 is used. The time-series data consists of five surveyed sites, whilst the spatial data consists of a single snapshot for the whole of the Isles of Scilly. In order to be able to make a direct comparison of both of the time-series and spatial data, the spatial data was sub-sampled where a $500\text{m} \times 500\text{m}$ bounding box was placed over each of the five sites centred on the 2008 survey for each location.

Time series modelling

Initially, a series of autoregressive (AR) models were fitted to time series of both mean quadrat shoot densities and patch occupancy (proportion of occupied quadrats), in order to explore long-term trends in density and extent, as well as evidence of dependence. Based on this preliminary analysis, it was found that it was sufficient

to fit a relatively simple population dynamic model to the time series, allowing us the main temporal processes to be quantified for comparison with our spatial data modelling. The population dynamics of shoot density, X , for year t , were modelled using a discrete time model:

$$X_{t+1} = X_t \exp(r - b \ln(X_t)),$$

where r is the intrinsic growth rate and population growth is regulated by a density-dependent process governed by the parameter, b .

Zostera marina leaves typically survive for less than 100 days over the summer months [Larkum et al., 2006]; however, rhizomes under the sand are something of an unknown but presumably persist for longer. The statistical population model was fitted to spatially replicated time series data in a mixed-effects framework [Pinheiro and Bates, 2000]. Spatial heterogeneity (within survey heteroscedasticity and between survey correlation) was modelled with an empirical variance-covariance matrix. Multiplicative (log-scale) Gaussian noise was assumed, as this has been shown to be an appropriate descriptor of stochastic processes in spatially explicit systems [Bonsall and Hastings, 2004]. The fitted models for each survey area were used to calculate return rates from perturbations around the equilibrium point; b provides an estimate for the return rate and are associated with persistence, perturbations in a system with high b quickly decay. This calculation of the return rate along with estimation of the boundary dimension and Korcak exponent as discussed in chapter 4 can be compared to the values obtained through simulation of the vegetation models.

6.4.4 Simulations

Introduction

Simulations were carried out as a probabilistic cellular automata on an $N \times N$ lattice denoted Ω . The simulations were synchronously updated, such that at each time-step each site has a probability of changing state depending on the model specifics. Any number of states can change at each time-step. Random numbers are generated from the *Matlab R2013a* rand function for each possible event at each time-step. A proposed event for each site was then accepted if the random number was below the probability for that site. The specific probabilities for each of the models shall be outlined in the following sections.

Two-state model

For a given lattice configuration Ω , The probabilities of a switch for each state are calculated according to Eq. (6.21a) and Eq. (6.21b). Each site is then assigned a uniformly distributed random number on the unit interval and flipped if this number falls below the probability at the site. The probabilities are multiplied by a factor ϵ , which was taken to be 0.01. This was used to slow down the rate at which state transitions occur and hence prevent pathological dynamics, such as complete extinction at all sites for a large population.

Three-state model

A configuration Ω now contains one of three states at each site denoted $\{0, 1, 2\}$. The state transitions flow according to the following schematic diagram

$$0 \rightarrow 1 \rightarrow 2 \rightarrow 0. \quad (6.43)$$

Hence for each site, there is one of two possibilities for the next time step: transitioning to the next state or remaining in the state in the previous time-step. This then allows the simulation to be carried out in the same procedure as the two-state model. Now the probabilities are calculated according to Eq. (6.22a), Eq. (6.22b) and Eq. (6.22c) for states 0,1 and 2 respectively. Again a factor ϵ was multiplied by the probabilities to slow the speed of the simulation down and prevent pathological dynamics.

6.5 Results

The simulation results are divided into three main sections. The initial numerical investigation was concerned with the relationship between the fractal statistics and the time to equilibrate following a disturbance. The second numerical investigation is concerned with an invading cluster along an environment with no gradient, but with a noise term characterised by its variance. The third numerical investigation is concerned with studying the boundary of a vegetation cluster along an environmental gradient.

6.5.1 Disturbance & recovery in a homogeneous environment

The first numerical investigation is concerned with how the observed fractal statistics are related to the ability of a vegetative system to recover following a disturbance.

This was applied to the two-state model with Gaussian and power law competition kernels. The persistence following a disturbance was measured by first allowing the system to equilibrate and then randomly removing 25% of the occupied sites. The time taken for the density to return to within 1% of the equilibrium density was then recorded. At the final distribution both the Korcak exponent and the box-counting dimension of the mass cluster(as opposed to the boundary) were recorded.

A binary toroidal lattice Ω is initiated with a uniform random density. Birth events occur stochastically with rate

$$k_1 * \Omega(1 - k_2 * \Omega/K), \quad (6.44)$$

where $*$ represents a two-dimensional discrete convolution, k_1, k_2 are the kernels for growth and competition respectively and K is the carrying capacity. Death occurs uniformly at rate μ . For the growth kernel a Gaussian centred at zero with variance l_1 was used while for the competition kernel both a power law distribution was used with power law exponent α and a Gaussian kernel with variance l_2 was used. Hence the parameters of the model are $\mu, k, \alpha/l_2, l_1$.

In order to measure persistence the simulation was initialised with randomly with a density of 0.2 lattice sites occupied. The simulation was then allowed to equilibrate at which point the temporal variance and skewness were measured. After 100 time-steps a shock is introduced to the system by randomly removing 20% of the population. The number of time-steps between the shock and the time it takes to return to within 0.01 standard deviations from the equilibrium density is taken. This was then repeated for 10^4 separate simulation runs for lattice sizes 100×100 and 200×200 (See Fig. 6.4 for the main results of these simulations). The average time to equilibrium relates well to the box-counting dimension, however the variance of the dimension at a given resilience value is high. The relationship also displays heteroscedasticity, where lower values of resilience have a far larger variance in the mass fractal dimension than at higher values of persistence. The relationship is however, expected from theory as a "rougher", more point-like spatial pattern is indicative of a system with lower resilience, at least from the perspective of this heuristic.

The Korcak dimension as calculated from the patch size distribution also has a positive relationship with the persistence heuristic. The variance of the dimension is lower than the equivalent with the mass fractal dimension. However, the relationship

is more convex, thus the higher values of persistence make the dimension more indistinguishable.

The measure of persistence was also compared against known statistics of criticality and disturbance from time-series statistics, Guttal and Jayaprakash [2009]; Scheffer et al. [2009], namely the variance and skewness of the population density. Both follow a convex relationship similar to the Korcak dimension and have small variance. It was found that there were no strong differences between Gaussian and power law competition kernels as both gave very similar relationships to the ones expected from theory.

6.5.2 Invading cluster with environmental noise

The second numerical experiment focuses on the impact of environmental

noise on both the growth rate of an invading cluster and the dimension of the boundary. According to the theory of fractional diffusion, a waiting time distribution characterised by an exponent θ leads to a functional relationship with the static properties of the boundary as well as the dynamic properties characterising the width of the growth front. As time is discrete, in place of the waiting time the constant growth rate λ in Eq. (6.22a) is replaced with $\lambda(x)$, where $\lambda(x)$ is a random variable for each site that is drawn from an inverse Pareto distribution of the form

$$\lambda(x) = \frac{1}{w(x)}, \quad \text{where } w(x) \sim \text{Pareto}(\sigma, w_0 = 1), \quad (6.45)$$

$\lambda(x)$ is defined as the probability of growth and hence its inverse is the expected number of time-steps until a site becomes occupied in the absence of other growth factors, such as competition.

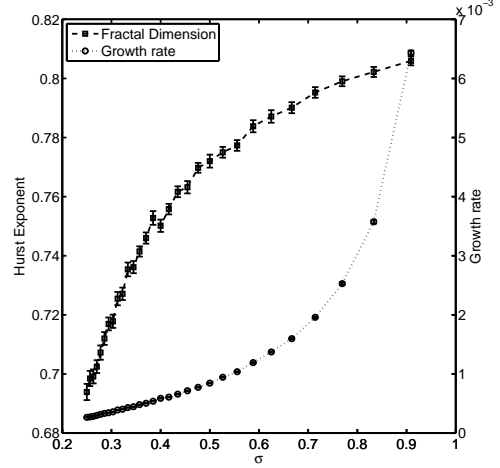


Figure 6.5: Measuring rates of growth from the three-state model with no environmental gradient and a simple environmental variation noise term characterised by the variance σ . Growth rate was calculated as the inverse of the time to reach the halfway point of the total population of sites.

In order to ascertain how the parameter σ affects the invasion rate of vegetation in an unoccupied landscape numerical simulation is performed. The system was initialised with an empty 100×100 lattice. The boundary of the lattice were held constant as unoccupied sites. A line of occupied was started at one end of a $N \times N$ lattice and the time until reaching half the distance of the lattice $N/2$ was taken. The growth rate was then calculated as the inverse of this time. The results are shown along with the fractal dimension of the boundary in Fig. 6.5

The exponent of the waiting time σ is varied for constant probability of death and recovery of the sites μ_1 and μ_2 , where recovery is high and death is low ($\mu_1 = 0.2$, $\mu_2 = 0.8$). The boundary dimension is monotonically decreasing (Hurst exponent increasing) in σ . The growth rate is also monotonically increasing, however is convex, whereas the Hurst exponent relationship is concave.

6.5.3 Properties of a vegetation boundary along an environmental gradient

The relationship between the boundary dimension of a growing cluster in the presence of an environmental gradient imposed on the birth rate has been explored. There was found to be a strong relationship between the boundary dimension and the return rate, where from theory we would expect an inverse relationship between the boundary dimension and the return rate. There were however, certain regions of parameter space where this relationship does not hold, moreover the data obtained from remote sensing and survey work conducted in Scilly, UK did not seem to match up. This lead to the idea of exploring the Korcak dimension. The Korcak dimension is defined as twice the exponent obtained by fitting a power law distribution to the patch size distribution. Mandelbrot hypothesised a simple relationship between the Korcak dimension and the boundary dimension for fracturing process, but this has since been overturned by recent studies [Imre et al., 2012].

Parameter space results

The hypothesis explored is that the Korcak and boundary dimension form a complete set of summary statistics for phase space. i.e. if we ascertained the Korcak and the boundary dimension we would be able to determine the return rate regardless of other parameters. Parameter space for Eq. 6.22 was divided up into four regions:

1. $\mu_1 = 0.3, \mu_2 = 0.8$;
2. $\mu_1 = 0.8, \mu_2 = 0.8$;

3. $\mu_1 = 0.3, \mu_2 = 0.3$;

4. $\mu_1 = 0.8, \mu_2 = 0.3$.

For each region, the environmental gradient described in Eq. 6.27 was used. Simulations were run over a range of environmental gradient γ and noise term ξ . As with previous results the environmental gradient γ gave a larger variance of boundary, Korcak dimension and return rate than the noise term ξ . Density also varies across both μ values and environmental values. The relationship between the Korcak dimension and the boundary dimension is distinct for each region in parameter space (Fig. 6.7). All regions do not obey the relationship proposed by Mandelbrot and in fact, there is broadly a negative relationship between the boundary and Korcak dimension. The seagrass data also violates this prediction, with a broad negative relationship structure.

Fig. 6.6 displays the relationship between the boundary dimension and the return rate. In each region there is a separate distinct relationship produced. For low disturbance rate μ_1 , an inverse relationship is produced, whilst when disturbance is high, either a positive ($\mu_2 = 0.3$) or flat relationship ($\mu_2 = 0.8$) occurs.

Finally the Korcak dimension was compared to the return rate (Fig. 6.8). The relationship here is less clear than for the boundary dimension. All regions have distinct relationships, with either broadly constant Korcak dimension, or positively or negatively correlated with the return rate.

Two state model

To elucidate the previous relationships discussed the two-state vegetation model is also considered in the presence of an environmental gradient. Initially, dynamic and spatial parameters of the model were fixed and the environmental parameters were varied (ξ and γ) on the unit interval. The death probability was kept constant at $\eta = 0.2$, the spatial growth and competition scales were kept constant at $\sigma_1 = 0.5, \sigma_2 = 1$ (where the total length of the system is 20). High and low competition factor k were investigated ($k = 0.2$ and $k = 0.8$). The Korcak exponent gives a strong linear trend with the return rate; there is also a high likelihood indicating a good fit of the exponent (See Fig. 6.10a and Fig. 6.10a). The same relationship does not hold when the environmental parameters are held constant and the demographic parameters are altered (Fig. 6.11), instead there is no strong correlation between the either the return rate and Korcak or boundary dimension.

Fig. (6.9) gives a comparison between simulations in the two-state model and the empirical seagrass data for the five sites between the Korcak Dimension and the return rate for the two-state model. With the assumption that the demographic parameters remain constant between sites, but the environmental parameters vary the resulting inverse relationship between the return rate and the Korcak dimension is reproduced. The simulation results were fitted to the data by using known values of the recruitment and death rate in *Zostera marina* [Larkum et al., 2006]. The resulting confidence intervals for the simulations fit within the bounds of all five sites, with la being marginal.

6.6 Conclusion

A variety of spatial lattice-based models with varying environmental and demographic parameters were constructed in order to determine the relationship between the return rate of a system at equilibrium or following a disturbance and the dimensionality of a cluster boundary (Minkowski-Bouligand), mass fractal and the patch-size distribution (Korcak dimension).

For a spatially homogeneous environment the two-state model was implemented with a variety of spatial kernels to determine a relationship between the mass fractal dimension and the time to equilibrium following a disturbance. The Korcak and the mass dimension both had a positive relationship with the average time to equilibrium, however the variation in the Korcak dimension for a particular equilibrium time was lower than that of the mass dimension. This suggests that the Korcak dimension of the entire vegetative spatial pattern provides a more accurate estimate of the equilibrium time than that of the mass fractal. As the Isles of Scilly, UK data was taken when each of the five sites were at statistical stationarity no comparison can be made.

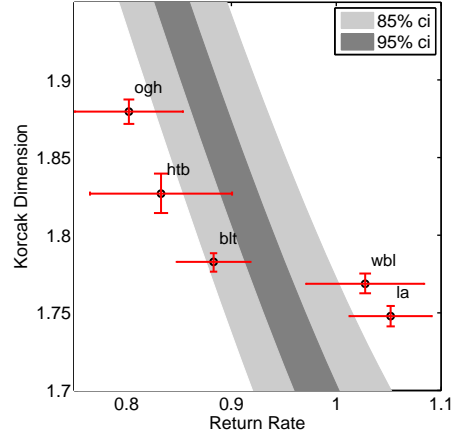


Figure 6.9: Korcak return rate relationship for surveyed and simulated vegetation. Empirical seagrass data for the five sites surveyed is shown as labelled points with \pm SE error bars. The diagonal solid band describes the inverse relationship reproduced by the two-state PCA model over a range of environmental parameters. The simulation rate is set using known parameters of death and recruitment for *Zostera marina* [Larkum et al., 2006]

For a growing cluster out of equilibrium there was found to be a good correspondence between the growth rate and the boundary dimension of the invading cluster over a range of environmental noise parameters. These values were taken as ensemble averages and hence would provide a useful link between growth rate and spatial properties if the system could be observed for a number of time-steps (assuming ergodicity).

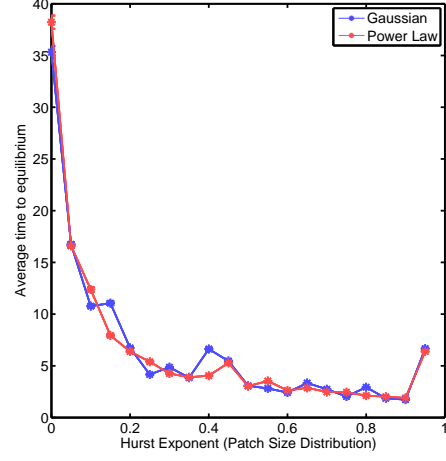
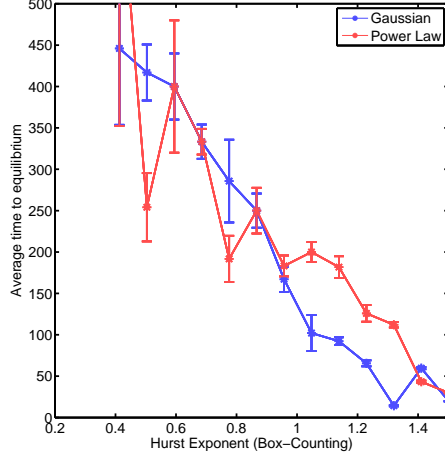
For a growth front at equilibrium that is being constrained by an environmental gradient, the return rate and Boundary Dimension as well as the Korcak Dimension were compared for both the two-state and the three-state model. It was found that for the two-state model both the boundary and Korcak dimension correlate well with the return rate at equilibrium for both high and low competition. This relationship is also hinted at by the Isles of Scilly data. Here the environmental variables representing the noise and slope of the gradient were varied and found to reproduce the negative dimension-return rate relationship well. However, if the environmental parameters were held constant and the dynamic parameters (namely the death rate and the competition factor) were varied (see Fig. 6.11), then this produced no strong dimension-return rate relationship. This shows the limitations of this method, where comparing between sites that may have strong demographic differences would be difficult. This also demonstrates that comparisons between spatial patterns of different species in order to elucidate the return rate of each would not be possible given this method.

For the three-state model the inverse relationship holds for a high recovery probability. However, when the probability of site recovery is low the relationship is inverted and boundary dimension is positively correlated with return rate. The region of parameter space where the return rate-dimension relationship is most pronounced is in the low death-high recovery region. This state corresponds most with the previous model as a recovery probability of one leads to an infinite recovery rate and hence the time spent in the third state for each site is negligible compared with the other two states.

For a probability of low recovery the recovered state (R) begins to play a non-negligible role in the dynamics. As it is assumed this state cannot be directly observed and hence the dimensionality statistics are calculated just from the occupied states. This can lead to a situation where the relative return rate is high, (in the case of no recovery the return rate is infinite), however the roughness of

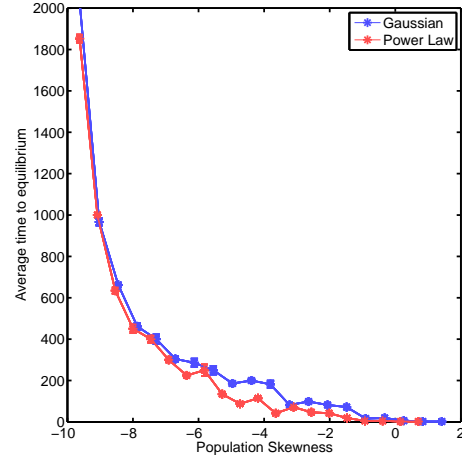
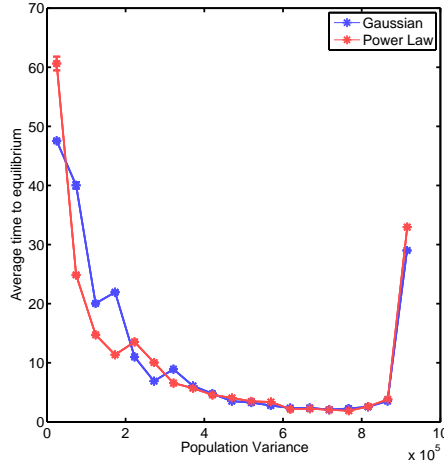
the boundary as well as the patchiness would increase due to the recovered state blocking the invasion of the occupied state. This leads to the inverse relationship that is observed as the environmental statistics are varied, but can also explain why there is no strong relationship as the demographic parameters μ_1 and μ_2 are varied and the environmental parameters are held constant.

The purpose of this chapter was to elucidate the relationship between the spatial scaling parameters that could be ascertained from a single spatial snapshot of vegetation and the dynamics that underpin the spatial distribution. The theory behind fractal growth phenomena has focused on situations where there is a certain amount of experimental control over the growth, such as initialising from a flat surface or single seed as well as taking regular spatial snapshots of the growth over time in order to calculate the scaling exponents of the system. In the context of spatial ecological modelling often this form of data does not exist and the initial stages of growth are not observed. It is necessary therefore to be able to assess vegetation spatial pattern that is well-established, although it may be in a constant state of flux. This work then demonstrates the conditions under which dynamics can be interpreted from the spatial snapshots alone and when this cannot be done. It also highlights that the particular mechanisms of the model, such as the inclusion of a recovering state or not can rapidly change the type of static-dynamic scaling relationships observed.



(a) Hurst Exponent (H) calculated from box-counting method versus Average time to equilibrium.

(b) Hurst Exponent (H) calculated from patch size distribution method versus Average time to equilibrium.



(c) Variance of time-series at equilibrium versus Average time to equilibrium.

(d) Skewness at equilibrium versus Average time to equilibrium.

Figure 6.4: Comparison between dynamic heuristics and average time to equilibrium in the two-state model. There is a linear relationship between the Hurst Exponent and the measure of persistence. The lower the value of H the longer the system takes on average to return to equilibrium and the higher the value of H the quicker it takes

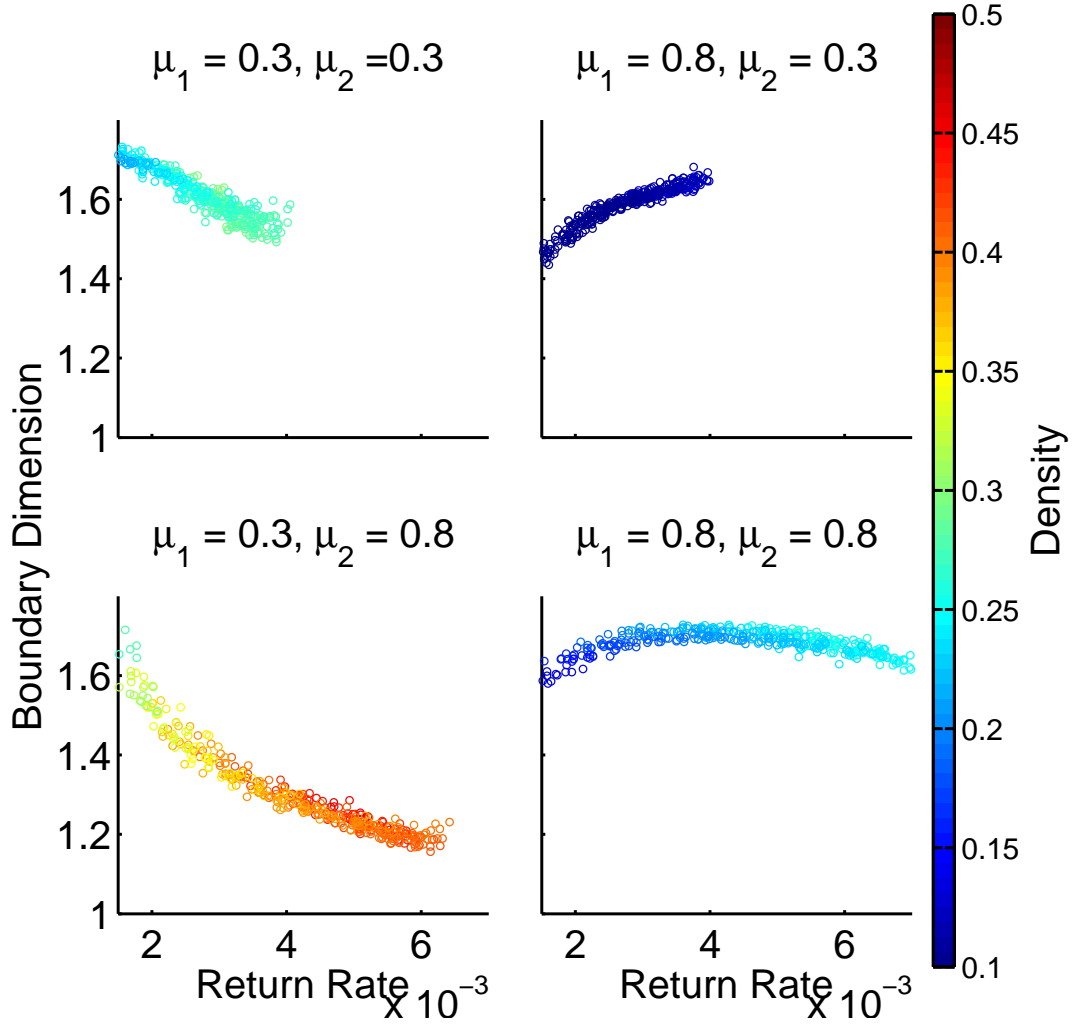


Figure 6.6: Comparison of boundary dimension and return rate for four regions of phase space in the three-state model. Shading represents total density of final simulation size. $\mu_1 = 0.3, \mu_2 = 0.8$ gives the strongest relationship in concordance with the original hypothesis and also has notably higher density than the other regions. $\mu_1 = 0.3, \mu_2 = 0.3$ also follows a similar relationship, although for low densities this does not hold. $\mu_1 = 0.8, \mu_2 = 0.8$ gives a less clear pattern as there is a positive correlation between return rate and boundary dimension for return rate less than 4×10^{-3}

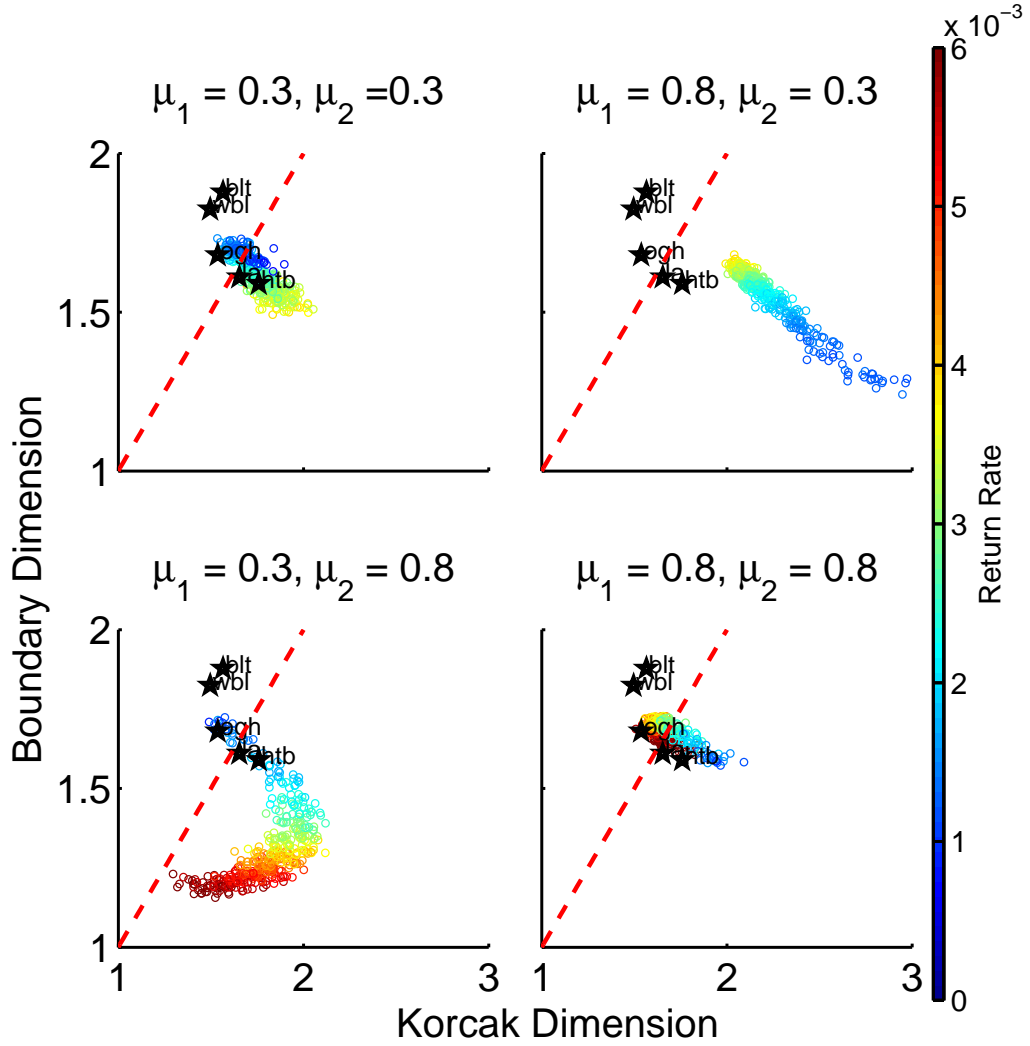


Figure 6.7: Comparison between the Korcak and boundary dimension for the four regions in parameter space in the three-state model. The Mandelbrot hypothesis is shown as the red dotted line on each of the plots. Also displayed on each of the plots are the five sites obtained from Scilly, UK. There appears to be a general negative correlation between the Korcak and Boundary dimension. However, the relationship to the return rate is less clear. For $\mu_1 = 0.3$, $\mu_2 = 0.3$ and $\mu_1 = 0.3$, $\mu_2 = 0.8$ a high boundary dimension and a low Korcak dimension corresponds to a low return rate. however for $\mu_1 = 0.8$, $\mu_2 = 0.8$ the relationship is reversed and a high boundary and low Korcak dimension correspond to a high return rate. For low death rate μ_1 the boundary dimension acts as a predictor of the return rate, as an increasing boundary leads to a decreasing return rate. From fractal growth models with a pure birth process this would be expected. For a high death rate, the boundary dimension relationship is reversed and the Korcak dimension now becomes an indicator of return rate, which would be expected from models of fracture cluster formation.

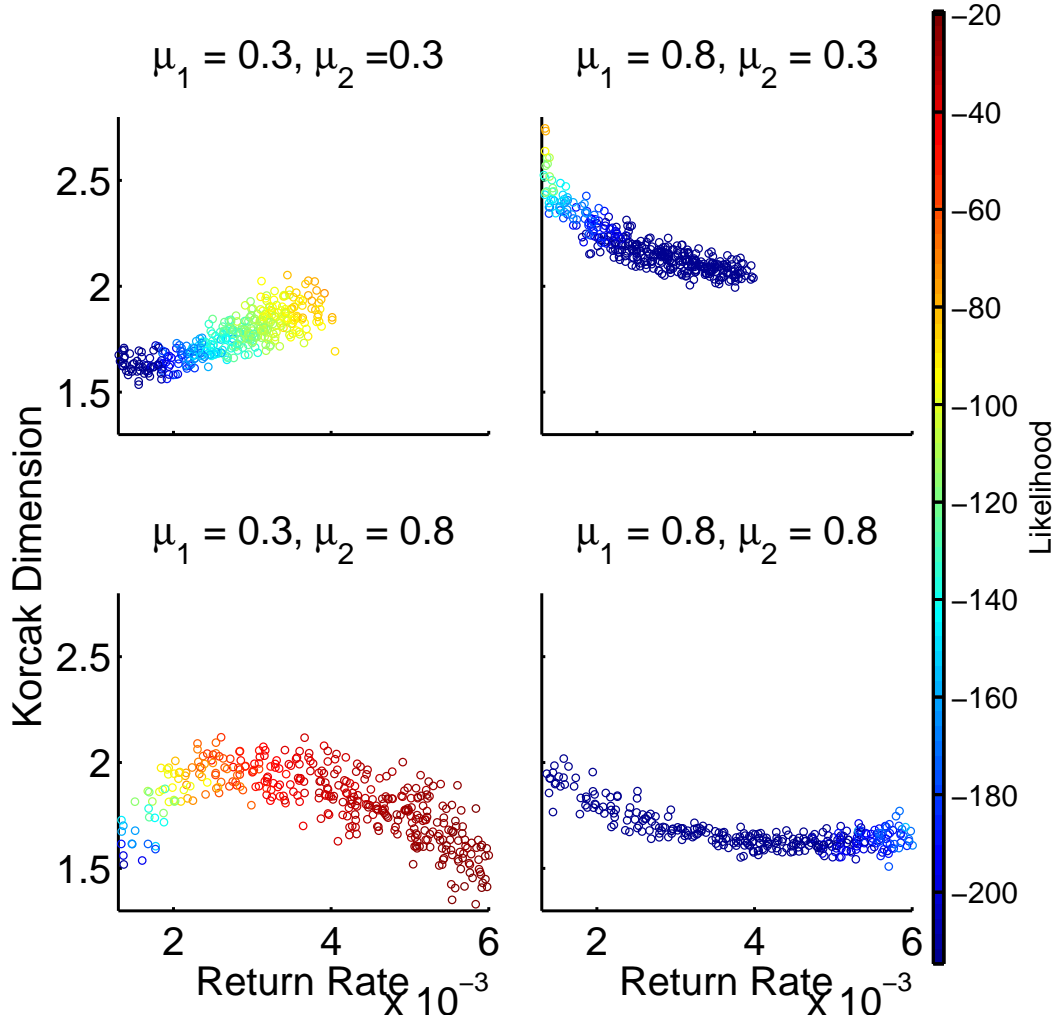
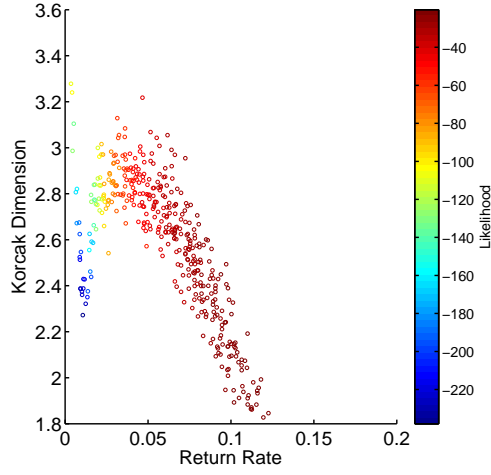
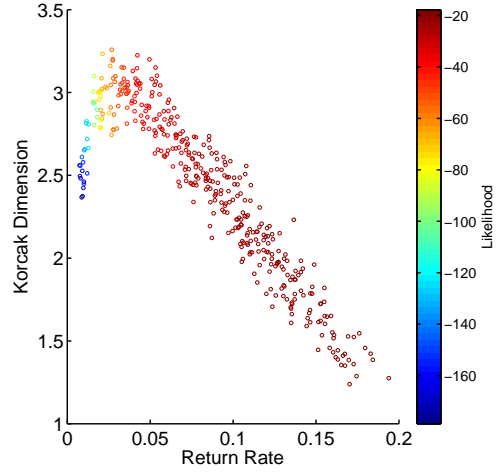


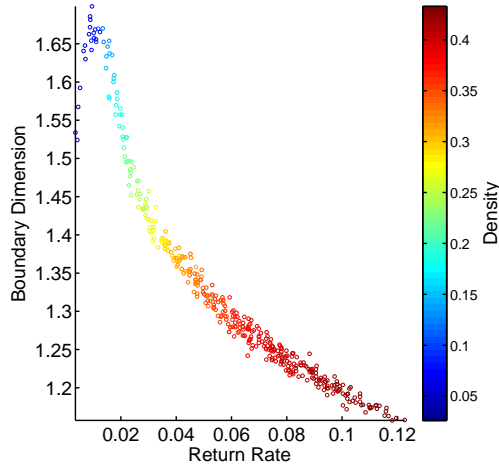
Figure 6.8: Comparison of the Korcak Dimension and the Return rate for four regions of phase space in the three-state model. Points are shaded according to their log likelihood, where red indicates a high likelihood and blue a low likelihood of the patch-size distribution fitting a power law. The only case that fits particularly well is the low death high recovery rate model indicative of a simple birth-death process. When the death rate is high the likelihood is lowest. Interestingly this seems to oppose the idea that a high death rate leads to a more fractured boundary and thus a larger distribution of patch-sizes.



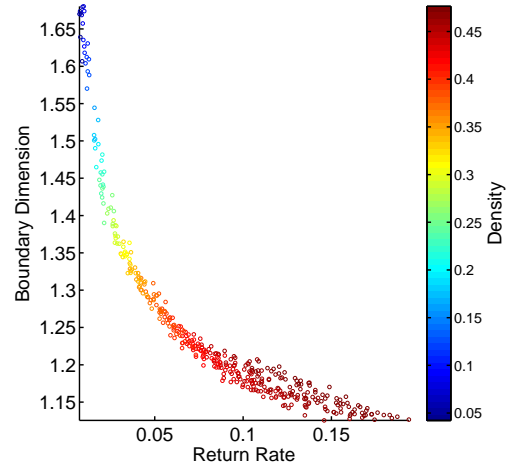
(a) Korcak dimension. High k



(b) Korcak dimension. Low k



(c) Boundary dimension. High k



(d) Boundary dimension. Low k

Figure 6.10: Relationship between return rate and boundary dimension and Korcak dimension for high and low competition in two-state model. Other parameters were kept constant at $\sigma_1 = 0.5, \sigma_2 = 1, \eta = 0.2$

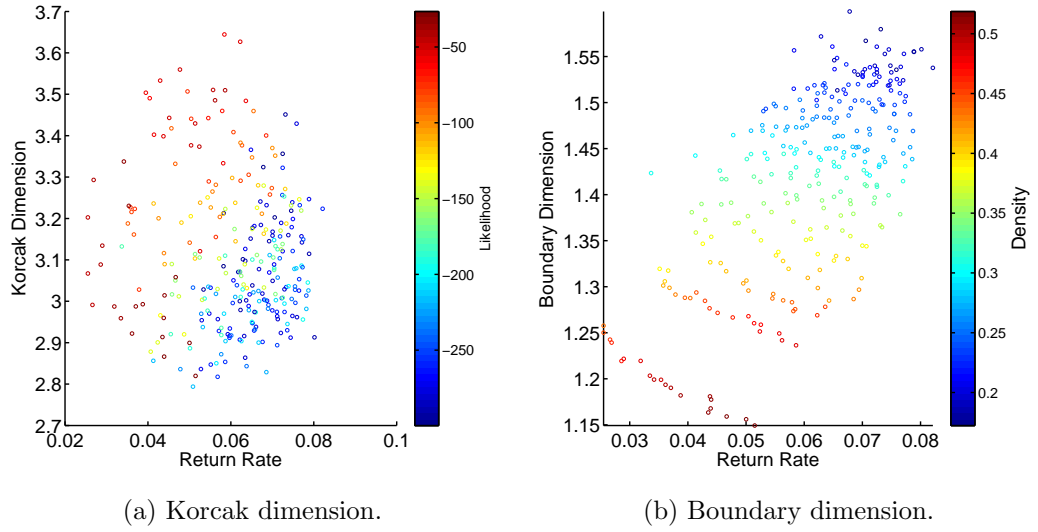


Figure 6.11: Varying the dynamic parameters k and η for fixed spatial and environment parameters $\sigma_1 = 0.5, \sigma_2 = 1, \gamma = 0.5, \xi = 0.1$ in two-state model. There is a lack of strong correlation between both the Korcak Dimension and the Boundary Dimension compared to the return rate.

Bibliography

- Om P Agrawal. Fractional variational calculus in terms of riesz fractional derivatives. *Journal of Physics A: Mathematical and Theoretical*, 40(24):6287, 2007.
- S Alexander and R Orbach. Density of states on fractals: fractons. *Journal de Physique Lettres*, 43(17):625–631, 1982.
- Roy M Anderson and Robert McCredie May. *Infectious diseases of humans*, volume 1. Oxford university press Oxford, 1991.
- Per Bak, Chao Tang, and Kurt Wiesenfeld. Self-organized criticality. *Physical review A*, 38(1):364, 1988.
- Per Bak, Kan Chen, and Chao Tang. A forest-fire model and some thoughts on turbulence. *Physics letters A*, 147(5):297–300, 1990.
- A-L Barabási. *Fractal concepts in surface growth*. Cambridge university press, 1995.
- Martin T Barlow. Which values of the volume growth and escape time exponent are possible for a graph? *Revista Matemática Iberoamericana*, 20(1):1–31, 2004.
- R Benzi, L Biferale, A Crisanti, G Paladin, M Vergassola, and A Vulpiani. A random process for the construction of multiaffine fields. *Physica D: Nonlinear Phenomena*, 65(4):352–358, 1993.
- Michael B Bonsall and Alan Hastings. Demographic and environmental stochasticity in predator–prey metapopulation dynamics. *Journal of Animal Ecology*, 73(6):1043–1055, 2004.
- BE Brown. Coral bleaching: causes and consequences. *Coral reefs*, 16(1):S129–S138, 1997.
- Sergio A Cannas, Diana E Marco, Sergio A Paez, and Marcelo A Montemurro. Plant spread dynamics and spatial patterns in forest ecology. *Nonextensive Entropy: Interdisciplinary Applications: Interdisciplinary Applications*, page 243, 2004.

- Sergio A Cannas, Diana E Marco, and Marcelo A Montemurro. Long range dispersal and spatial pattern formation in biological invasions. *Mathematical biosciences*, 203(2):155–170, 2006.
- Aaron Clauset, Cosma Rohilla Shalizi, and Mark EJ Newman. Power-law distributions in empirical data. *SIAM review*, 51(4):661–703, 2009.
- Ulf Dieckmann and Richard Law. Relaxation projections and the method of moments. *The Geometry of Ecological Interactions: Simplifying Spatial Complexity* (U Dieckmann, R. Law, JAJ Metz, editors). Cambridge University Press, Cambridge, pages 412–455, 2000.
- Samuel F Edwards and DR Wilkinson. The surface statistics of a granular aggregate. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 381(1780):17–31, 1982.
- Fereydoon Family and Tamás Vicsek. Scaling of the active zone in the eden process on percolation networks and the ballistic deposition model. *Journal of Physics A: Mathematical and General*, 18(2):L75, 1985.
- Mitchell J Feigenbaum. Universal behavior in nonlinear systems. *Physica D: Nonlinear Phenomena*, 7(1):16–39, 1983.
- Frédéric Guichard, Patti M Halpin, Gary W Allison, Jane Lubchenco, and Bruce A Menge. Mussel disturbance dynamics: signatures of oceanographic forcing from local interactions. *The American Naturalist*, 161(6):889–904, 2003.
- Vishwesha Guttal and Ciriya Jayaprakash. Spatial variance and spatial skewness: leading indicators of regime shifts in spatial ecological systems. *Theoretical Ecology*, 2(1):3–12, 2009.
- John M Halley and William E Kunin. Extinction risk and the $1/f$ family of noise models. *Theoretical Population Biology*, 56(3):215–230, 1999.
- A Harnos, G Horvath, AB Lawrence, and G Vattay. Scaling and intermittency in animal behaviour. *Physica A: Statistical Mechanics and its Applications*, 286(1):312–320, 2000.
- Harold M Hastings, Richard Pekelney, Richard Monticciolo, David Vun Kannon, and Diane Del Monte. Time scales, persistence and patchiness. *Biosystems*, 15(4):281–289, 1982.

- HM. Hastings and G Sugihara. *Fractals: A User's Guide for the Natural Sciences*. Oxford University Press, 1993.
- Ove Hoegh-Guldberg. Climate change, coral bleaching and the future of the world's coral reefs. *Marine and freshwater research*, 50(8):839–866, 1999.
- Attila R Imre, Josef Novotný, and Duccio Rocchini. The korcak-exponent: A non-fractal descriptor for landscape patchiness. *Ecological complexity*, 2012.
- Mehran Kardar, Giorgio Parisi, and Yi-Cheng Zhang. Dynamic scaling of growing interfaces. *Physical Review Letters*, 56(9):889, 1986.
- Sonia Kéfi, Max Rietkerk, Concepción L Alados, Yolanda Pueyo, Vasilios P Papanastasis, Ahmed ElAich, and Peter C De Ruiter. Spatial vegetation patterns and imminent desertification in mediterranean arid ecosystems. *Nature*, 449(7159):213–217, 2007.
- Shinya Kizaki and Makoto Katori. Analysis of canopy-gap structures of forests by ising-gibbs states-equilibrium and scaling property of real forests. *Journal of the Physics Society Japan*, 68(8):2553–2560, 1999.
- J Klafler and Igor M Sokolov. Anomalous diffusion spreads its wings. *Physics world*, 18(8):29, 2005.
- Gady Kozma and Asaf Nachmias. The alexander-orbach conjecture holds in high dimensions. *Inventiones mathematicae*, 178(3):635–654, 2009.
- Anthony WD Larkum, Robert J Orth, and Carlos Duarte. *Seagrasses: biology, ecology and conservation*. Springer, 2006.
- René Lefever and Olivier Lejeune. On the origin of tiger bush. *Bulletin of Mathematical Biology*, 59(2):263–294, 1997.
- JR Leith. Fractal scaling of fractional diffusion processes. *Signal processing*, 83(11):2397–2409, 2003.
- Bruce D Malamud, Gleb Morein, and Donald L Turcotte. Forest fires: an example of self-organized critical behavior. *Science*, 281(5384):1840–1842, 1998.
- Benoit B Mandelbrot. *The fractal geometry of nature*. Macmillan, 1983.
- Ralf Metzler and Joseph Klafter. The random walk's guide to anomalous diffusion: a fractional dynamics approach. *Physics reports*, 339(1):1–77, 2000.

- Ralf Metzler and Joseph Klafter. The restaurant at the end of the random walk: recent developments in the description of anomalous transport by fractional dynamics. *Journal of Physics A: Mathematical and General*, 37(31):R161, 2004.
- Mercedes Pascual and Frédéric Guichard. Criticality and disturbance in spatial ecological systems. *Trends in ecology & evolution*, 20(2):88–95, 2005.
- José C Pinheiro and Douglas M Bates. *Mixed-effects models in S and S-PLUS*. Springer Verlag, 2000.
- Max Rietkerk and Johan Van de Koppel. Regular pattern formation in real ecosystems. *Trends in Ecology & Evolution*, 23(3):169–175, 2008.
- LM Sander, CP Warren, and IM Sokolov. Epidemics, disorder, and percolation. *Physica A: Statistical Mechanics and its Applications*, 325(1):1–8, 2003.
- Marten Scheffer, Jordi Bascompte, William A Brock, Victor Brovkin, Stephen R Carpenter, Vasilis Dakos, Hermann Held, Egbert H Van Nes, Max Rietkerk, and George Sugihara. Early-warning signals for critical transitions. *Nature*, 461(7260):53–59, 2009.
- Laurent Seuront. *Fractals and multifractals in ecology and aquatic science*. CRC Press, 2009.
- Yonathan Shapir, Subhadip Raychaudhuri, David G Foster, and Jacob Jorne. Scaling behavior of cyclical surface growth. *Physical Review Letters*, 84(14):3029, 2000.
- Didier Sornette. Linear stochastic dynamics with nonlinear fractal properties. *Physica A: Statistical Mechanics and its Applications*, 250(1):295–314, 1998.
- George Sugihara and Robert May. Applications of fractals in ecology. *Trends in Ecology & Evolution*, 5(3):79–86, 1990.
- Donald L Turcotte. *Fractals and chaos in geology and geophysics*. Cambridge university press, 1997.
- Christian Valentin, Jean-Marc d’Herbès, and Jean Poesen. Soil and water components of banded vegetation patterns. *Catena*, 37(1):1–24, 1999.
- Tamás Vicsek, Miklós Cserző, and Viktor K Horváth. Self-affine growth of bacterial colonies. *Physica A: Statistical Mechanics and its Applications*, 167(2):315–321, 1990.

TA Witten Jr and Leonard M Sander. Diffusion-limited aggregation, a kinetic critical phenomenon. *Physical review letters*, 47(19):1400, 1981.

Chapter 7

Parameter estimation from a snapshot of regular vegetation spatial patterns

Data do not give up their secrets easily.
They must be tortured to confess.

(Jeffrey Hooper - Bell Labs)

7.1 Introduction

For a particular vegetation ecosystem we wish to estimate parameters that affect the dynamic persistence of the ecosystem such as strength of competition and extent of local positive and negative spatial interactions. For time series data, this type of analysis is well-established and often a full Bayesian likelihood methodology is used to fit a stochastic model to spatio-temporal data [Gibson and Austin, 1996; Wikle, 2003]. Methodology is less well-established when only a single spatial snapshot is present, this can occur due to the large costs of measuring vegetation distributions on a regular basis or if rapid inference is required, such as when a vegetation community is threatened. Keeling et al. [2004] uses the conservation of certain spatial statistics to produce an error function that can be minimised in order to estimate parameters for an epidemiological model. The procedure is to calculate the expected rate of change of spatial correlation functions given a certain parameter set. This can then be considered as a cost function, which can be minimised over several variables.

Previous analyses of vegetation spatial snapshots have used windowed Fourier transform Penny et al. [2013] or Wavelet analysis Rosenberg [2004], however these techniques are based purely on spatial statistics alone without using prior knowledge of vegetation growth dynamics. The advantage of model inference over these techniques is that they directly obtain biological mechanisms that underpin the pattern formation.

In this chapter we explore several methods of model fitting to single snapshots of vegetation. Throughout, the model of banding formation developed in chapter 5 shall be used to fit to both generated data and real data obtained from aerial photographic surveys such as the one outlined in chapter 3 as well as from other ecosystems where banding is present such as mussel beds and semi-arid ecosystems [Rietkerk and Van de Koppel, 2008]. The first section explores using the expected change in density of a spatial pattern as a cost function in order to fit single model parameters to data when other parameters are known. The limitations of this method shall then be discussed. In the next section we develop the methodology of Keeling et al. [2004] into a Bayesian framework. A Likelihood is developed that can be then used in order to calculate the posterior distribution over the parameters in the model. This methodology is tested using simulated data and then applied to a variety of real datasets.

7.2 Density as an observable

7.2.1 Method

For a particular site where strong banding occurs, it is desirable to estimate certain parameters from the model from a single spatial snapshot. To recap, the model is as follows. There is an $N \times N$ lattice of states that can either be occupied (1) or unoccupied (0) i.e. $S = \{0, 1\}^{N \times N}$. At each time step there is a probability of a birth event and a death event at each lattice site. The probabilities are computed as follows

$$P_B = k_1 * S \times (1 - k k_2 * S), \quad (7.1a)$$

$$P_D = \eta S, \quad (7.1b)$$

the $*$ operator represents a convolution of the discrete lattice space. k_1 and k_2 are a Gaussian and offset Gaussian representing clonal growth and long range competition,

$$k_1 = \exp\left(\frac{1}{2\sigma_1^2}[x^2 + y^2]\right), \quad (7.2a)$$

$$k_2 = \exp\left(\frac{1}{2\sigma_2^2}[(x - o)^2 + y^2]\right). \quad (7.2b)$$

In order to perform inference, the usual method is to calculate the posterior distribution of the model and estimate the maximum a posteriori to find the parameter set that closely fits the data. For a single spatial snapshot this may be an intractable problem due to the large amount of unobserved data. However, we may consider the rate of change of leaving the state $\frac{d}{dt}P(\Omega)$. Again, this would be dependent upon other probabilities of state that we do not have access to. We can then consider the density observable ρ , which maps the set of states onto the closed unit interval $[0, 1]$ and then use probability of observing a density ρ conditioned on the current spatial pattern to calculate the expected value $E[\rho]$. Hence the rate of change of the expected value of the density can be defined $\frac{d}{dt}E[\rho]$, this can then be calculated directly from the observed state. The Expected change in density (*EDD*) for a spatial pattern Ω is

$$\frac{d}{dt}E[\rho] = \int_{\Omega} P_B(\omega) - P_D(\omega) d\omega. \quad (7.3)$$

This then means that in order for the rate of change of density to be zero, the probability of birth events summed over all sites must be equal to the probability of death events summed over all sites.

It is assumed that transient dynamics of the system are short-lived and hence the fluctuations of observables such as density and spatial correlations are around an equilibrium value. More rigorously, we assume at equilibrium that

$$\frac{d}{dt}E[\rho] = 0. \quad (7.4)$$

Since $\frac{d}{dt}E[\rho]$ can be calculated from a single spatial snapshot, the procedure is to minimise this quantity subject to the parameters of the model. In order to test this procedure we begin by minimising one parameter over synthetic data whilst keeping all other parameters fixed.

7.2.2 Minimising for one parameter

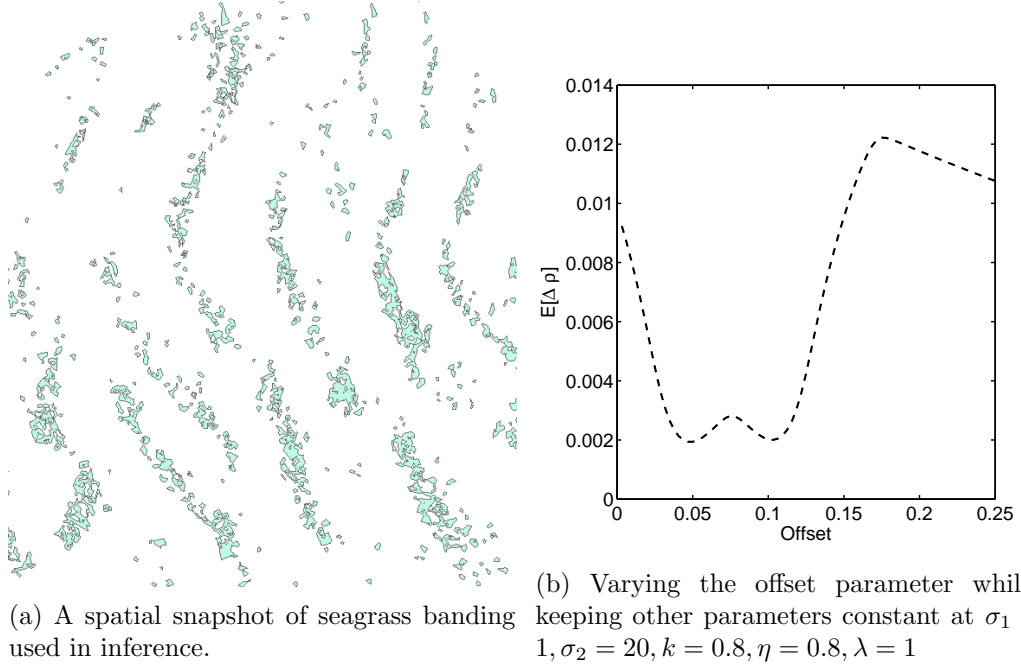
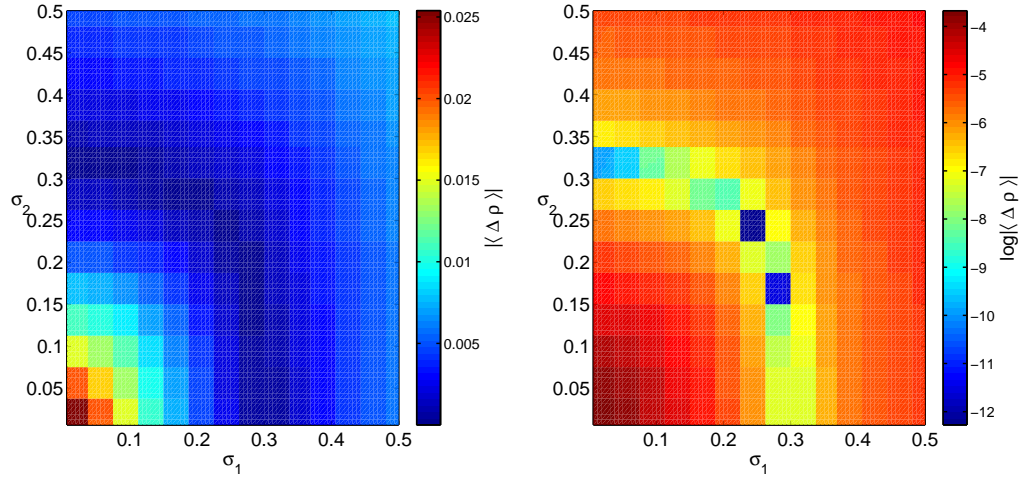


Figure 7.1: The expected change in density of an example spatial snapshot of banding over a range of a single parameter. Two local minima are observed, one at $o = 0.05$ and another at $o = 0.1$. o is measured in the total length of the bounding box of the data, hence an offset $o = 1$ would represent an offset of the total system size.

Any deviation from Eq. 7.4 is assumed to be a small fluctuation around an equilibrium. The strategy then is to minimise the function on the left-hand side over parameter space. For an initial investigation most of the state parameters shall be kept constant and a single variable will be varied, to see where the local minima

(if any) exist. Inference was performed on a spatial snapshot of real data of seagrass banding. In Fig. 7.1b the offset parameter is varied whilst other parameters are kept fixed. Two local minima exist, one at $o = 0.05$ and the other $o = 0.1$. The development of two local minima may be due to the intrinsic periodicity of the system. If there is a strong wavelength at λ then one would also expect due to the periodic nature of the system a peak at $\lambda/2$, albeit a weaker one than the primary wavelength. Interestingly here, there appears to be no distinction between the minima of either, possibly suggesting that two wavelengths are observed for this data, or else the inference procedure cannot properly distinguish between the two o for this region of parameter space.

7.2.3 Minimising for two parameters



(a) absolute expected change in density. (b) Logarithm of absolute expected change in density for varying the variance of the growth and competition kernel. There appears to be a symmetry where the minimum is achieved for a continuous range of parameters σ_1 and σ_2

Figure 7.2: Expected change in density over a range of of parameter space for σ_1 and σ_2 . Other parameters were kept fixed for $k = 0.8, o = 0.1, \eta = 0.8, \lambda = 1$. The linear scale reveals a broad range of parameters where the expected change in density is minimum. There is an apparent symmetry between the two parameters conferring positive feedback and negative feedback

An exhaustive search over parameter space for two parameters may also be performed with relatively little computation. As an example the parameters σ_1 and σ_2 were varied whilst the other parameters were kept constant (See Fig. 7.2). On a linear scale a minima valley is revealed, representing a range of σ_1 and σ_2

minimise the expected change in density. There appears to be a symmetry between the two parameters, however on a log scale an asymmetry is revealed and there are local minima isolated in certain regions. This apparent asymmetry however may just be due to the sampling performed over the parameter space. If there is some functional relationship between σ_1 and σ_2 over where the expected change in density is minimised, then the minimisation procedure cannot uniquely identify a plausible set of parameters for the model.

7.2.4 Searching over parameter space

In order to search over all of parameter space an exhaustive search would not be applicable due to the high dimensionality of the problem. Hence we may implement a search of parameter space using a *Markov Chain Monte Carlo* (MCMC) method such as the Metropolis-Hastings algorithm (MH-MCMC) [Hastings, 1970; Metropolis et al., 1953]. The expected change in density can then be used as a cost function to be minimised. This procedure would not be able to determine both the dynamic parameters

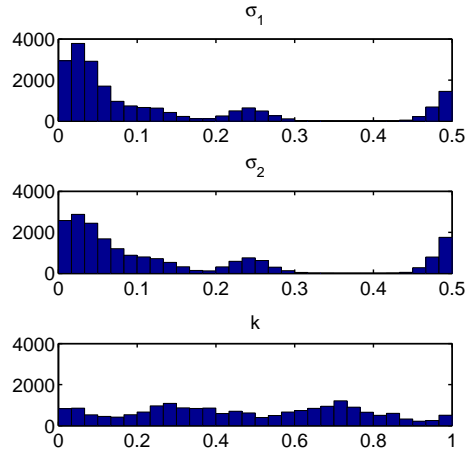


Figure 7.3: M-H MCMC to minimise the expected change in density over the parameters σ_1 , σ_2 and k for the spatial snapshot shown in Fig. 7.1a.

(η, λ) , since if both parameters are set to zero, the expected change in density is trivially zero. These could however be estimated from the time-series data or else approximated and then sensitivity analysis could be conducted to determine if the "guess" strongly biases the other parameters. We begin by performing a small Metropolis-Hastings sampler over a small subregion of parameter space to test for convergence and sensitivity of the expected change in density.

The Metropolis-Hastings Markov chain Monte Carlo scheme to take into account the fact that the absolute expected change in density is to be minimised (as opposed to a log posterior, which is maximised in the scheme). The algorithm is as follows, Parameters λ, η and o were held constant at 1, 0.8, 0.1 respectively (these were chosen in keeping with the previous results where an offset of 0.1 and competition value k gave a minima of the EDD). The other parameters were initialised by drawing from a uniform distribution. The data is then convoluted with the parametrised kernels

and from this the EDD was calculated (see Eq. 7.3) for the proposal parameters, which we denote as τ_p and the proposal EDD, which we shall denote as EDD_p . The ratio of $EDD(t)$ is compared to the proposal EDD_p and the value α is taken according to the following

$$\alpha = \min \left\{ 1, \frac{EDD(t)}{EDD_p} \right\}. \quad (7.5)$$

If $\alpha > 1$, the proposal parameters τ_p are accepted and $EDD(t+1) = EDD_p$. If $\alpha < 1$ then the proposal parameters τ_p are accepted with probability α , else they are rejected, so $EDD(t+1) = EDD(t)$ and the parameters remain unchanged. This gives a way of sliding down the valleys and troughs of the EDD landscape, but also gives the possibility of sometimes climbing, thus lowering the chance of ended up trapped in a local minima. Further, after an adequate number of Markov chain steps (known as the *burn-in time*) the current parameters may be recorded for each time step and hence a histogram of values may be recorded. This gives an estimate of the robustness of different parameters and can also indicate if there are strong correlations between certain parameters that would make a minimisation process difficult.

The results of the three parameter inference procedure with a burn-in time of 100 and an inference time of 10^3 can be seen in Fig. 7.3. The histogram of values for σ_1 and σ_2 are very similar, confirming that the scheme cannot identify between the two parameters. k also has a broad distribution. As the actual underlying parameters are not known, it would be a better first step to perform this inference procedure on generated data where the parameters are known. This also shows that certain parameter sets are inherently unidentifiable with the current procedure.

7.3 Spatial correlations as an observable

Density can only give very basic information on the dynamics of the model for a single spatial pattern. In order to identify parameters, more informative observables of the spatial pattern are required. The idea then is to use spatial correlation functions in place of density to determine parameters by minimising the expected change of correlations.

Before we consider the calculation of the expected rate of change of this observable, we shall slightly alter the model. As the transition of states are synchronous, there

is an issue using the previous method where the rates are directly converted into probabilities. In particular there is an issue with the logistic growth term $P_1(1-kP_2)$. Although P_1 and P_2 are probabilities, k is a scalar and as such the resulting term is not a probability as it will not in general sum to one. Instead the following rates were implemented

$$P(0 \rightarrow 1) = k_1 * \Omega, \quad (7.6a)$$

$$P(1 \rightarrow 0) = (k_2 * \Omega)^k, \quad (7.6b)$$

where Ω is the current state of the lattice and k_1 and k_2 are kernels mediating the growth and death terms. Note that in doing so we have removed the ambient death parameter, this can be added in later if needed. k is a competition factor, but is now an exponent of the death term to guarantee the probability of a transition is bounded by one. k can be interpreted as a mortality factor due to spatial crowding, as $k \rightarrow \infty$ nearly all terms apart from those at unity will go to 0 and hence there is no death due to spatial factors. As $k \rightarrow 0$ all terms go to one except those at 0 and hence competition is spatially very strong. k then acts like an inverse of the strength of competition.

To generalise the spatial competition death kernel in 2D it is assumed there is an offset to the kernel due to an environmental factor such as wind or currents. This is described by a two parameter offset (o_x, o_y) . k_1, k_2 are Gaussian, and have a parameter each to measure their variance, σ_1, σ_2 . To summarise there are two spatial parameters σ_1 & σ_2 , the spatial competition parameter k , and two offset parameters o_x and o_y making a total of five parameters that would need to be fitted.

7.3.1 Derivation of the variational correlation for spatial pattern inference

We wish to derive the full term of the estimator used in the spatial correlation inference. We begin by assuming the Kolmogorov forward equation for the probability of a system to be in state $S \in \Omega$. Given a rate of moving from state $S \rightarrow S'$ as $w_\Theta(S'|S)$, where $\Theta = \{\theta_1, \theta_2, \dots, \theta_N\}$ are the parameters for the system,

$$\frac{\partial}{\partial t} P_t(S) = \int_{\Omega} [w_\Theta(S|S')P_t(S') - w_\Theta(S'|S)P_t(S)] dS'. \quad (7.7)$$

We have some observable of the system state $\phi : \Omega \rightarrow \mathbb{R}$. For our particular example we have used spatial correlation as the observable, however this method is general

and would apply to any observable of the system. The expectation of an observable is defined as

$$\mathbb{E}[\phi(S)] = \int_{\Omega} \phi(S) P(S) dS. \quad (7.8)$$

The rate of change of this expectation can then be calculated using Eq. 7.7

$$\frac{\partial}{\partial t} \mathbb{E}[\phi(S)] = \frac{\partial}{\partial t} \int_{\Omega} \phi(S) P_t(S) dS \quad (7.9)$$

$$= \int_{\Omega} \phi(S) \frac{\partial}{\partial t} P_t(S) dS \quad (7.10)$$

$$= \iint_{\Omega \times \Omega} \phi(S) [w_{\Theta}(S|S') P_t(S') - w_{\Theta}(S'|S) P_t(S)] dS' dS. \quad (7.11)$$

Observe that

$$\iint_{\Omega \times \Omega} \phi(S) [w_{\Theta}(S|S') P_t(S') - w_{\Theta}(S'|S) P_t(S)] dS' dS \quad (7.12)$$

$$= \iint_{\Omega \times \Omega} \phi(S) w_{\Theta}(S|S') P_t(S') dS' dS - \iint_{\Omega \times \Omega} \phi(S) w_{\Theta}(S'|S) P_t(S) dS' dS \quad (7.13)$$

$$= \iint_{\Omega \times \Omega} \phi(S) w_{\Theta}(S|S') P_t(S') dS' dS - \iint_{\Omega \times \Omega} \phi(S') w_{\Theta}(S|S') P_t(S') dS' dS \quad (7.14)$$

$$= \iint_{\Omega \times \Omega} [\phi(S) - \phi(S')] w_{\Theta}(S|S') P_t(S') dS' dS. \quad (7.15)$$

We define the change in an observable due to the transition to a new state to be $\Delta\phi(S|S') = [\phi(S) - \phi(S')]$. Hence we have

$$\frac{\partial}{\partial t} \mathbb{E}[\phi(S)] = \iint_{\Omega \times \Omega} \Delta\phi(S|S') w_{\Theta}(S|S') P_t(S') dS' dS. \quad (7.16)$$

We assume that the system is at statistical stationarity and so the expectation of the observable is time invariant. There is also a single observation of the system at equilibrium denoted E . The probability function of the state can then be estimated by a single delta function $\hat{P}(S) = \delta_E(S)$. With the following assumptions we can

then apply them to Eq. 7.16

$$0 = \iint_{\Omega \times \Omega} \Delta\phi(S|S') w_{\Theta}(S|S') \hat{P}(S') dS' dS \quad (7.17)$$

$$= \iint_{\Omega \times \Omega} \Delta\phi(S|S') w_{\Theta}(S|S') \delta_E(S') dS' dS \quad (7.18)$$

$$= \int_{\Omega} \Delta\phi(S|E) w_{\Theta}(S|E) dS. \quad (7.19)$$

The estimator is therefore defined to be

$$\zeta_E(\Theta) = \int_{\Omega} \Delta\phi(S|E) w_{\Theta}(S|E) dS. \quad (7.20)$$

The integral currently goes over all possible states of the lattice $S \in \Omega$. For an $N \times N$ lattice size, this implies there are $2^{N \times N}$ possible combinations to sum over. In order to reduce the size of the integral to be performed we may approximate it by assuming that the transition rates of each site are small. It can then be assumed at each time step that the transitioning states are uncorrelated with one another and hence each state that E can transition to is $E(s)$ for all lattice sites s . $E(s)$ denotes the state E with site s switched. The integral now sums over all lattice sites s instead of all possible states and hence the size of the sum is $N \times N$ as opposed to $2^{N \times N}$. Hence, if I is the index set for the lattice S , $\zeta_E(\Theta)$ is approximated as

$$\zeta_E(\Theta) = \sum_{s \in I} \Delta\phi(E(s)|E) w_{\Theta}(E(s)|E). \quad (7.21)$$

Note that the transition rate $w_{\Theta}(E(s)|E)$ is the rate of site s transitioning, hence is either the rate of a birth at site s or the rate of a death at site s depending on the current state, hence we may simplify notation by defining $w_{\Theta}(s) := w_{\Theta}(E(s)|E)$.

Note that the estimator $\zeta_E(\Theta)$ has no dependence upon the probability of states P . $\Delta\phi(S|E)$ and $w_{\Theta}(S|E)$ are calculated by considering all the ways in which the system can leave the state E . This can be calculated with complete knowledge of state E for any parameter set Θ . By using the equilibrium assumption there should be a set of parameters $\zeta_E(\hat{\Theta}) = 0$. $\hat{\Theta}$ is defined to be the point estimator of the parameters after minimising $\zeta_E(\Theta)$.

The observables being considered are the pair-wise correlations at distance d . As seen in Eq. 7.21, each site contributes an amount towards the whole expected change in the observable. for any birth-death system there are three correlations to consider: 00, 01 and 11. The correlation of 00 therefore is the probability that given a current site in state 0, a randomly selected site at distance d away is also in state 0 with probability $P_{00}(d)$. Eq. 7.21 is then used to calculate the expected rate of change of the observable $P_{00}(d)$. The increase or decrease of the observable for each site depends on whether that site is occupied, in which case the only way in which the observable can change is if there is a death event at that site or when the site is unoccupied, in which case the only way the observable can change is if there is a birth event at that site. Let $I(s)$ be the indicator function which is 1 if the site is occupied or 0 if it is unoccupied. The calculation of the rate of change of the observable $P_{00}(d)$ is therefore

$$\begin{aligned}
\frac{d}{dt}\mathbb{E}[\Delta P_{00}(d)] &= \sum_{s \in I} \Delta P_{00}(d)(E(s)|E)w_{\Theta}(s) \\
&= \sum_{s \in \mathbb{I}: I(s)=1} \Delta P_{00}(d)(E(s)|E)w_{\Theta}(s) + \sum_{s \in \mathbb{I}: I(s)=0} \Delta P_{00}(d)(E(s)|E)w_{\Theta}(s) \\
&= \sum_{s \in \mathbb{I}: I(s)=1} \Delta P_{00}(d)(E(s)|E)d_{\Theta}(s) + \sum_{s \in \mathbb{I}: I(s)=0} \Delta P_{00}(d)(E(s)|E)b_{\Theta}(s) \\
&= \sum_{s \in \mathbb{I}: I(s)=1} N_{10}^s(d)d_{\Theta}(s) - \sum_{s \in \mathbb{I}: I(s)=0} N_{00}^s(d)b_{\Theta}(s),
\end{aligned}$$

where $N_{XY}^s(d)$ represents the number of pairs where site s is in state X and the sites distance d away are in state Y . As an example, if site s is in state 0 and there is a birth event at site s then the number of 00 pairs that are destroyed at distance d apart due to this event is $N_{00}^s(d)$.

The expected rate of change for the pairs 01 and 11 at distance d apart may similarly be calculated by considering how each event at a site contributes to the creation or the destruction of a XY pair. The calculation for 01 is

$$\begin{aligned}
\frac{d}{dt}\mathbb{E}[\Delta P_{01}(d)] &= \sum_{s \in I} \Delta P_{01}(d)(E(s)|E)w_{\Theta}(s) \\
&= \sum_{s \in \mathbb{I}: I(s)=1} \Delta P_{01}(d)(E(s)|E)d_{\Theta}(s) + \sum_{s \in \mathbb{I}: I(s)=0} \Delta P_{01}(d)(E(s)|E)b_{\Theta}(s) \\
&= \sum_{s \in \mathbb{I}: I(s)=1} (N_{11}^s(d) - N_{10}^s(d))d_{\Theta}(s) + \sum_{s \in \mathbb{I}: I(s)=0} (N_{00}^s(d) - N_{01}^s(d))b_{\Theta}(s).
\end{aligned}$$

Similarly the calculation for the 11 pair is

$$\begin{aligned}
\frac{d}{dt}\mathbb{E}[\Delta P_{11}(d)] &= \sum_{s \in I} \Delta P_{11}(d)(E(s)|E)w_{\Theta}(s) \\
&= \sum_{s \in \mathbb{I}: I(s)=1} \Delta P_{11}(d)(E(s)|E)d_{\Theta}(s) + \sum_{s \in \mathbb{I}: I(s)=0} \Delta P_{11}(d)(E(s)|E)b_{\Theta}(s) \\
&= - \sum_{s \in \mathbb{I}: I(s)=1} N_{11}^s(d)d_{\Theta}(s) + \sum_{s \in \mathbb{I}: I(s)=0} N_{01}^s(d)b_{\Theta}(s).
\end{aligned}$$

We wish to calculate the total rate of change for a correlation pair XY for all distances. This is accomplished by constructing a weighted sum over all distances d , where the weight of each sum corresponds to the inverse of the number of sites at distance d from the site s . Define

$$N^s(d) = \#\{k \in \mathbb{I} : |s - k| = d\}. \quad (7.22)$$

For a toroidal lattice this is the same for all sites s and hence $N^s(d) = N(d)$, however if the boundary conditions are not toroidal, then the number of neighbours at distance d will vary throughout the lattice. The general calculation of the total rate of change of a XY with general boundary conditions is therefore

$$\frac{d}{dt}\mathbb{E}[\Delta P_{XY}] = \sum_{d=1}^{d_{\max}} \sum_{s \in \mathbb{I}: I(s)=1} \frac{\Delta P_{XY}(d)(E(s)|E)d_{\Theta}(s)}{N^s(d)} + \sum_{s \in \mathbb{I}: I(s)=0} \frac{\Delta P_{XY}(d)(E(s)|E)b_{\Theta}(s)}{N^s(d)}. \quad (7.23)$$

Once each observable has been established an estimator for the total rate of change for all observables can be constructed. Index the set of observables with j , so that the set of observables is denoted $\{\phi_j : j = 1 \dots K\}$. For each observable ϕ_j denote the corresponding estimator as ζ_j such that

$$\zeta_j(\Theta, E) = \int_{\Omega} \Delta \phi_j(S|E)w_{\Theta}(S|E)dS. \quad (7.24)$$

A full estimator that minimises each corresponding estimator ζ_j can then be constructed as

$$\zeta_{\Theta}(E) = \sum_{j=1}^N a_j \zeta_j^2 \quad (7.25)$$

$$= \sum_{j=1}^N a_j \left(\int_{\Omega} \Delta \phi_j(S|E) w_{\Theta}(S|E) dS \right)^2. \quad (7.26)$$

Note that for any $a_j > 0$ Eq. (7.25) equates to 0 if and only if each estimator ζ_j is 0. There is hence a choice over how the a_j should be constructed.

7.3.2 Non-spatial birth death process

In order to determine the accuracy of the method, a simple birth-death process was constructed where the correlations could be calculated exactly and hence compared against the numerical results. For a lattice Ω there $N \times N$ sites that can occupy one of two states: Alive (1) and dead (0). The system can be modelled as a finite Markov chain where each site has the transition probabilities

$$\pi_{01} = d, \quad (7.27)$$

$$\pi_{10} = b. \quad (7.28)$$

Assuming detailed balance we may calculate the equilibrium probabilities of being in states 1, P_1 and 0, P_0 at each site,

$$\pi_{01}P_1 = \pi_{10}P_0, \quad (7.29)$$

$$dP_1 = bP_0. \quad (7.30)$$

$$(7.31)$$

Observing that $P_0 + P_1 = 1$,

$$P_0 = \frac{d}{b} (1 - P_0) \quad (7.32)$$

$$\implies P_0 = \frac{d}{b+d} \quad P_1 = \frac{b}{b+d}. \quad (7.33)$$

The correlation distance functions can be easily calculated owing to the fact that there is no spatial correlation. Hence for $d > 0$,

$$P_{00}(d) = P_{00} = P_0 P_0 = \frac{d^2}{(b+d)^2}, \quad (7.34)$$

$$P_{01}(d) = P_{01} = P_0 P_1 = \frac{bd}{(b+d)^2}, \quad (7.35)$$

$$P_{11}(d) = P_{11} = P_1 P_1 = \frac{b^2}{(b+d)^2}. \quad (7.36)$$

These values were found to match up well to numerical results. The set of observables is hence defined to be $\{P_{00}, P_{01}, P_{11}\}$. Thus the estimator $\zeta_\Theta(E)$ is

$$\begin{aligned} \zeta_\Theta(E) = & \left(\sum_i \Delta P_{00}(E_i|E) w_\Theta(E_i|E) \right)^2 \\ & + \left(\sum_i \Delta P_{01}(E_i|E) w_\Theta(E_i|E) \right)^2 \\ & + \left(\sum_i \Delta P_{11}(E_i|E) w_\Theta(E_i|E) \right)^2. \end{aligned} \quad (7.37)$$

Where E_i denotes the state E with the i th site swapped. Consider the first term of the sum. there are two ways in which a 00 pair can change. Either there is a death in a 01 pair, thus contributing to the creation of a 00 pair or there is a birth in a 00 pair leading to its loss. Hence

$$\begin{aligned} & \sum_i \Delta P_{00}(E_i|E) w_\Theta(E_i|E) \\ = & \sum_i \Delta P_{00}(0_i \rightarrow 1_i|E) w_\Theta(0_i \rightarrow 1_i|E) + \Delta P_{00}(1_i \rightarrow 0_i|E) w_\Theta(1_i \rightarrow 0_i|E) \\ = & \sum_i -\frac{d}{(b+d)} b \delta_i + d(1 - \delta_i) \frac{d}{(b+d)} \\ = & \frac{bd}{(b+d)} (1 - \rho) - \frac{d^2}{(b+d)} \rho. \end{aligned}$$

Denoting the density of state E as ρ and equating to 0 to find estimators of parameters \hat{b} and \hat{d} we find

$$\begin{aligned} -\frac{d^2}{(b+d)}\rho + \frac{bd}{(b+d)}(1-\rho) &= 0, \\ \implies b(1-\rho) &= d\rho, \\ \implies \frac{d}{b} &= \frac{1-\rho}{\rho}. \end{aligned}$$

Thus $\rho = b/(b+d)$ as is expected from the equilibrium distribution. It can also be shown by similar calculation that the two other squared terms in the estimator corresponding to the 11 and 01 pairs give the same result when set to 0. This was compared to numerical results in Fig. (7.4) and was found to be in close agreement. In particular the difference in where the respective estimates of ζ_{00}^2 were minimised were in close agreement with the true value of the parameter.

With the confirmation made for the simple birth-death model we may extend to look at numerical solutions for the spatially-aggregated birth and spatial-competition death model. A simulation was initialised randomly on a 100×100 grid and ran until reaching stationarity (approximately 2000 time-steps). The resultant configuration E was then used to estimate each parameter in turn. The true parameter values are $\tau_1 = 1, \tau_2 = 0.5, k = 0.5, o_x = 10, o_y = 10$. Keeping all parameters fixed and measuring each in turn produces a convex functional form that has a unique minimum for the precisions τ_1, τ_2 and the competition factor k . The offset parameters are not convex, however have global minima that match well with the true parameters of the simulation (Fig. 7.5).

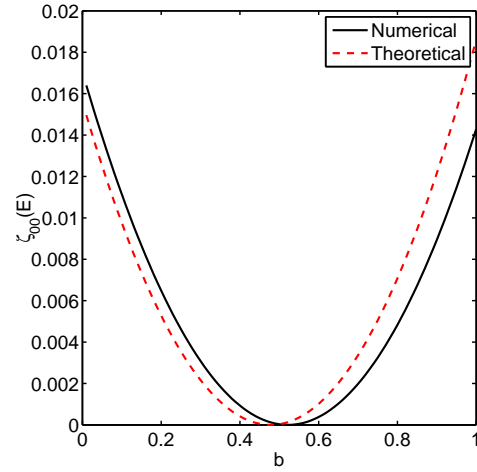


Figure 7.4: Comparison of the variational parameter $\zeta_{00}(E)$ for a single realisation of the non-spatial birth-death process. Here the parameter b is varied whilst d is held constant at the true value 0.5. The function calculated numerically is minimised around 0.52, close to the theoretically calculated value from the density 0.48 and the true value 0.5. The system was ran for 6000 time-steps on a 50×50 grid.

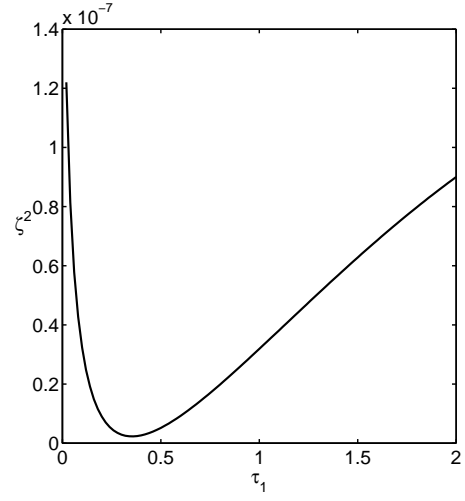
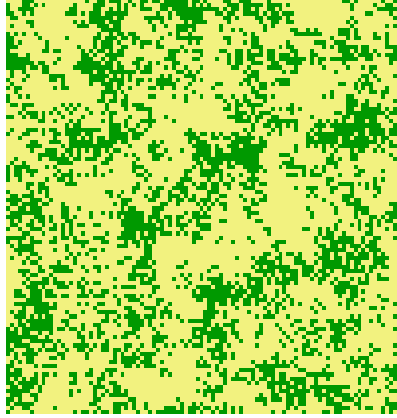
With the method established for estimation over one parameter, attention can now turn to varying two parameters. The parameters used for the simulation are $\tau_1 = 0.8, \tau_2 = 0.1, k = 0.1, o_x = 10, o_y = 10$. Two parameters were varied whilst all others were fixed at their true values. Varying the offset parameters o_x, o_y produces a complex multi-modal landscape, however the global minimum is close to the true values of the parameters (Fig. 7.6a) . Varying of the two precision values τ_1 and τ_2 produces a smoother cost function landscape where the minima occupy a valley (Fig. 7.6b). Again, the true value of the parameters is close to the global minima. Hence for two parameters at least the method is able to identify parameters that fit well to the spatial pattern. The next step is to consider how to perform inference when no parameters are known.

7.4 Inference

The expected rate of change of spatial correlations can be used as a cost function in a minimisation procedure in order to detect the most likely model parameters for a single spatial snapshot. The procedure can be done using an exhaustive search when fitting a single parameter or two parameters as was done in Section 7.3. However, for a high dimensional parameter space, other methods need to be employed due to an exhaustive search being computationally intractable. Further, the minimisation procedure only gives a single estimate for the likely parameter and does not indicate the confidence in the parameter. It is desirable, therefore, to have a fitting procedure that also gives a probability distribution associated with the parameters. We consider three methods to achieve this: simulated annealing, Approximate Bayesian Computation and a synthetic-likelihood approach. In this section each method shall be described in how it relates to the fitting problem and its relative merits shall be assessed.

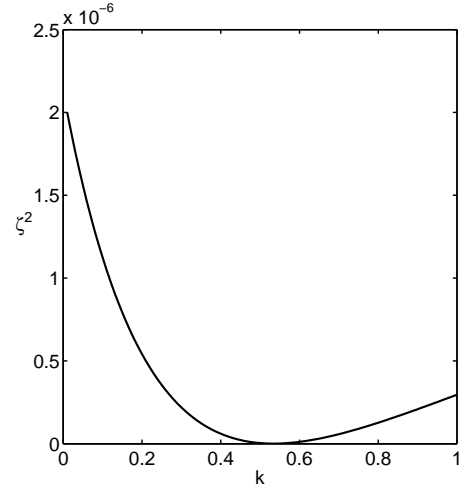
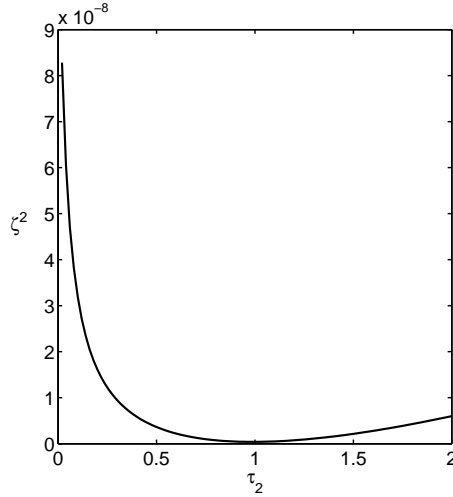
7.4.1 Simulated annealing

One approach towards parameter estimation is simulated annealing [Kirkpatrick et al., 1983]. The main idea of this scheme is to find a global minima of an energy function by mimicking the physical process of annealing, where a structure is able to configure to its global minimum energy state through a cooling schedule. First an energy function is defined with a parameter T , which corresponds to the temperature of the system. For larger values of T , the energy function appears flat and there is weak selection among the parameters, for a lower T the energy function is highly dependent on the function which is being minimised. If temperature is fixed, then



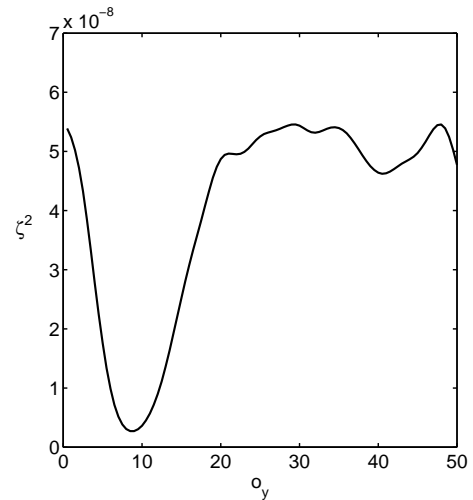
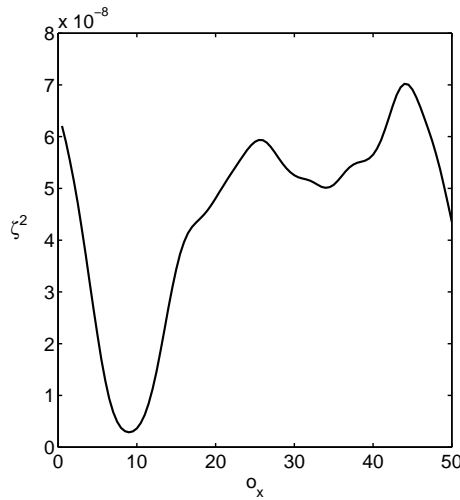
(a) simulated spatial snapshot over which parameters are estimated.

(b) Precision of growth kernel τ_1 . Minimises around 0.1



(c) The precision fro death τ_2 minimises around 1.

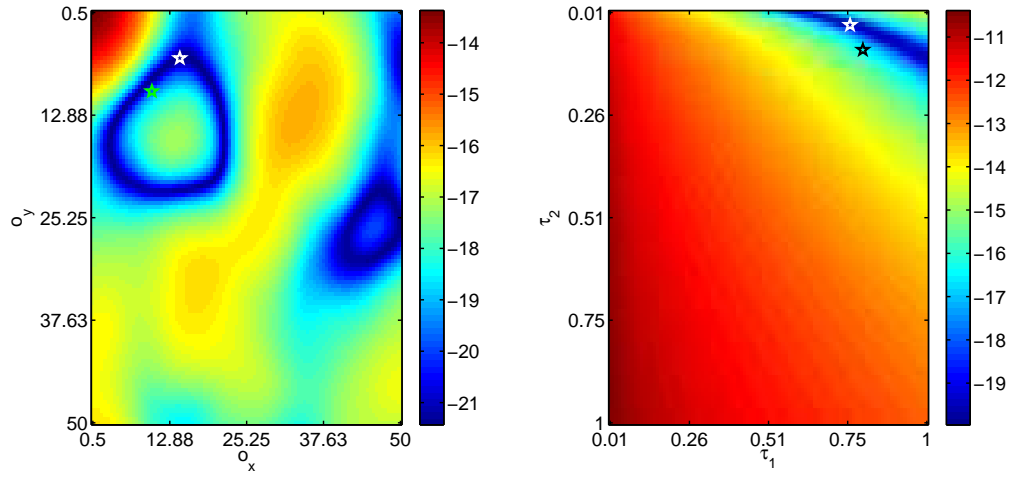
(d) Competition factor k minimises around 0.5 close to the true value.



(e) Offset o_x .

(f) Offset o_y .

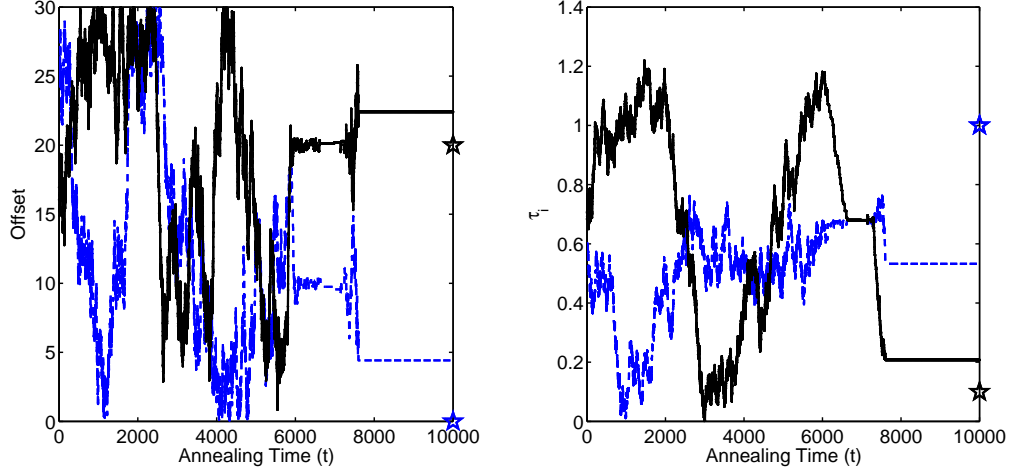
Figure 7.5: Calculating ζ^2 for each parameter of the model whilst holding other parameters constant. The true parameters of the system are $\sigma_1 = 6, \sigma_2 = 0.1, k = 1, o_x = 10, o_y = 10$



(a) Function $\log \zeta^2$ for $\{o_x, o_y\}$ slice of parameter space. True parameter values displayed in white. Minimum of ζ^2 displayed in green.

(b) Function $\log \zeta^2$ for $\{\tau_1, \tau_2\}$ slice of parameter space. True parameter values displayed in white. Minimum of ζ^2 displayed in black.

Figure 7.6: Calculating ζ^2 over a two parameter subset of parameter space. The true parameters of the system are $\tau_1 = 0.8, \tau_2 = 0.1, k = 0.1, o_x = 10, o_y = 10$. There is close agreement between the predicted and the actual parameters indicating the method is stable for a single spatial snapshot.



(a) Annealing over parameters o_x and o_y . (b) Annealing over parameters τ_1 and τ_2

Figure 7.7: Simulated Annealing for four parameters $\{\tau_1, \tau_2, o_x, o_y\}$, with $k = 0.2$. true values of the system are displayed as stars and the temperature is lowered every 100 steps. All parameters converge close to their true values, however there is a large amount of error in each of the final parameters.

for a large temperature no single state is strongly selected for and the scheme would not converge. If temperature is low then the scheme is likely to get stuck in only local minima and not find the global minima. The idea then is to vary the temperature, such that the scheme can move between minima eventually finding the lowest one. To summarise, a biased random walk is performed over parameter space, where lower energy values are proportionally selected according to a temperature T . The system is cooled or annealed, by decreasing the temperature until all the parameters are fixed and the random walk no longer makes new jumps. This method provides a way of finding a global minimum in a complex energy landscape that may have many local minima. Where banding is weak, the ζ^2 function can have many local minima where a simple scheme with no annealing may become stuck (Fig. 7.6a gives an example of such an energy landscape).

More formally there is a range of temperatures denoted $\mathbb{T} = \{T_0, T_1, \dots\}$ where typically T_0 is large and $\lim_{i \rightarrow \infty} T_i = 0$. \mathbb{T} is known as the cooling schedule and needs to be defined carefully in order for the random walk to not become stuck in a local minima. An initial parameter state Θ_0 is drawn from a prior distribution. The probability is then calculated as

$$\pi(\Theta_0) = \exp(-\zeta^2/T_0) \quad (7.38)$$

This is then fed into a standard Metropolis-Hastings sampler. After a certain number of steps t_{inf} , the temperature is updated so that $T_i \rightarrow T_{i+1}$ and the calculation of the probability for the current state is repeated. This process continues until the chain makes no more updates. An example of this scheme is given in Fig. (7.7) on a 100×100 snapshot that has reached statistical equilibrium. Here the cooling schedule is $\{1000, 990, 980, \dots, 1\}$, with 100 time-steps between each new temperature regime. The scheme was performed on four parameters: the offsets o_x and o_y and the variance parameters τ_1 and τ_2 . Although the scheme was approximately able to find the offsets to within an accuracy of 5 spatial units, the precisions were not accurately determined. This is also coupled with the problem that there is no estimate to the errors of the fitted parameters. For these reasons simulated annealing is not considered further.

7.4.2 Approximate Bayesian computation

The second method considered was Approximate Bayesian Computation (ABC) [Beaumont et al., 2002]. In standard Bayesian inference a likelihood is developed based on a statistical model of the data. For simple statistical models, this may be a simple analytic function and hence the posterior distribution can be calculated either directly or through the use of MCMC. For more complex statistical models on the other hand, there may be no simple way of writing down the likelihood, it can either be difficult to establish or it becomes computationally very expensive to compute. When such cases arise, it may be appropriate to approximate the likelihood. Instead of computing the likelihood directly, data drawn from the statistical model may be compared to the data being assessed using appropriate summary statistics. The probability of a set of parameters given that the simulated data and the underlying data is less than a certain distance according to some metric is then used as an approximation to the full likelihood. This then approximates the full likelihood i.e.

$$P(D|\theta) \approx P(\rho(D, \hat{D}) < \epsilon|\theta) \quad (7.39)$$

where ρ is a metric measuring the distance between the simulated data \hat{D} and the true data D for some tolerance threshold ϵ . The simplest way of computing $P(\rho(D, \hat{D}) < \epsilon|\theta)$ is via a rejection scheme [Pritchard et al., 1999; Tavaré et al., 1997].

For a pre-defined tolerance threshold ϵ , the rejection scheme begins by drawing a set of parameters θ from the prior distribution. Simulated data is then produced

from the model given the current parameter set, which is compared to the true data using a set of summary statistics and a distance function ρ to compare between the data. If the distance between the simulated data and the model data is less than ϵ then the set of parameters is accepted and recorded, otherwise it is rejected. The resulting set of parameters is an approximate sample of the posterior distribution. This sample can then be used to estimate the maximum a posteriori parameters along with their uncertainty.

For the model of spatial competition in vegetation, the summary statistic to be chosen would be the spatial correlation function for the pairs 00, 01 and 11. The distance between the two correlation distributions can then be any that compares two distributions, for instance the dot product. The tolerance can then be set in order for the rejection scheme to be performed.

Although theoretically it is possible for the approximate Bayesian computation scheme to work, there are issues. The greatest issue is that for a model with a large number of dimensions, such as the spatial banding model the number of samples that need to be taken in order to well-represent the likelihood space is large. This combined with the spatial correlation function being costly to calculate makes this procedure a large computational affair. This approach is therefore also not considered further.

7.4.3 Synthetic likelihood

The likelihood function $P(D|\theta)$ is defined as the probability of seeing the observed data D given the parameters θ . If instead of considering the full data, we choose an appropriate function of the data ρ , the resulting synthetic likelihood may be written as $P(\zeta(D)|\theta)$ [Wood, 2010]. This synthetic likelihood has several advantages. Firstly, although a full likelihood of the data may be analytically intractable, the statistic ζ may have comparatively small dimension and hence an analytic form of the likelihood may be produced. The likelihood can also be calculated directly without the need to simulate for each chosen parameter set θ and hence has a computational advantage over ABC. Although a certain amount of information is lost with summarising the data using ζ , if ζ is properly specified then the resulting probability distribution will closely match the full likelihood. For the banding model, we consider the expected rate of change in the pair-wise correlations as the summary statistic. This produces a synthetic likelihood with a five-dimensional parameter

space corresponding to the parameters of the model $\{r, \theta, l_1, l_2, k\}$ i.e.

$$P\left(\frac{d}{dt}\mathbb{E}[\phi(D)]|\{r, \theta, l_1, l_2, k\}\right) \quad (7.40)$$

This likelihood may then be used in a MH-MCMC scheme. In order to proceed, a functional form of this likelihood needs to be derived.

7.4.4 Derivation of the likelihood

In order for a Likelihood to be determined for the summary statistic ζ , the exact probability distribution of this statistic must be discovered. The likelihood $L(\Theta|D) = P(D|\Theta)$ is approximated by $P(\zeta(D)|\theta)$, where ζ is the sum of squares of correlation rates, i.e.

$$\zeta = \sum_{i=0}^{N-1} \zeta_i^2. \quad (7.41)$$

Assuming that each statistic is drawn independently from a uniform distribution with the same variance $\zeta_i \sim N(0, \sigma)$. The distribution for ζ defined in Eq. 7.41 can be calculated using the change of variables formula. Take F is the cumulative distribution for ζ_i , then the cumulative distribution for ζ_i^2 is

$$G(y) = F(r^{-1}(y)),$$

where $r(y) = y^2$. Differentiating and substituting in the pdf for ζ_i we find

$$\begin{aligned} \frac{d}{dy}G(y) &= \frac{d}{dy}F(r^{-1}(y)), \\ g(y) &= f(r^{-1}(y))\frac{d}{dy}r^{-1}(y), \\ g(y) &= \frac{1}{\sqrt{2\pi\sigma^2}}\exp\left(\frac{-r^{-1}(y)^2}{2\sigma^2}\right)\frac{d}{dy}r^{-1}(y), \\ g(y) &= \frac{1}{\sqrt{2\pi\sigma^2}}\exp\left(\frac{-\sqrt{y}^2}{2\sigma^2}\right)\frac{d}{dy}\sqrt{y}, \\ g(y) &= \frac{1}{2\sqrt{2\pi\sigma^2}}\exp\left(\frac{-y}{2\sigma^2}\right)y^{-1/2}. \end{aligned}$$

g is therefore gamma distribution with parameters $k = 1/2, \theta = 2\sigma^2$. It follows either from the moment-generating function or from the properties of the gamma function that the resulting distribution for ζ is $\text{Gamma}(\sum_{i=0}^{N-1} 1/2, 2\sigma^2)$. Hence the

resulting pdf h is

$$h(\zeta) = \frac{1}{(2\sigma^2)^{N/2}\Gamma(N)} \exp\left(\frac{-y}{2\sigma^2}\right) y^{N/2-1}. \quad (7.42)$$

However, for the summary statistic considered in the Metropolis-Hastings inference scheme, the values for each $P_{XY}(d)$ are summed together for all d and spatial location and then squared. This means that the actual statistic is $\zeta = P_{00}^2 + P_{01}^2 + P_{11}^2$, if we assume each statistic is approximately normally distributed then the total distribution is gamma with $k = 3$, this is approximately equal to the Exponential distribution and hence we shall use the Exponential distribution as the Likelihood. The exponential distribution has a single parameter λ that is the inverse of the mean of the distribution. In order to calculate this unknown parameter a simulation with parameter values that produce a spatial pattern qualitatively similar to that of the data is run to statistical stationarity. The expected rate of change of the correlations is then calculated and the total error ζ^2 is recorded every 100 time-steps. After sufficient time the mean of these errors at statistical stationarity can then be used as an estimate for λ . The sensitivity of the estimated λ can then be assessed by varying the model parameters and re-running the calculation to check if the resulting λ is significantly different.

Table 7.1 shows the estimated distributions of the values P_{00} , P_{01} and P_{11} for a range of system sizes. For our assumption of the distribution of ζ to work we would need the distributions of P_{00} , P_{01} and P_{11} to all have a normal distribution. A Kolmogorov-Smirnov test was performed on each sample of P_{XY} for system size $50 \times 50, 70 \times 70, 90 \times 90, 110 \times 110, 130 \times 130$. Each test did not reject the null hypothesis that the samples were drawn from a normal distribution, with level 0.05 for any of the correlations or system sizes providing evidence that the correlation rates can be considered normal. This then provides support for the argument that the appropriate choice of the synthetic likelihood is the exponential distribution.

To test the statistical stationarity assumption, spatial correlations were recorded for single model runs. Transient dynamics quickly decay (less than 50 time-steps) to a stationary state that is approximately drawn from a normal distribution (Fig. 7.9a). The total rate of change across all correlations is also calculated from simulation runs of various system sizes. It was found that the total expected rate of change in the correlation structure scales approximately with the system size (Fig. 7.10).

The issue with using the gamma form of the likelihood is that it is based on the sum squared of the observables ζ_i without considering each observable individually. This means that if one observable dominates the sum-squared value, then it is only this observable that is being selected for in the likelihood and the other observables are redundant. One way to overcome this would be to add weights to each term in the sum in Eq. 7.41, however it is not clear what values these weights should take. Another approach is to assume each statistic ζ_i has normal error with variance σ_i and hence the resulting likelihood would be

$$l(\zeta) = \prod_{i=0}^{N-1} \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp\left(-\frac{1}{2\sigma_i^2}\zeta_i^2\right), \quad (7.43)$$

where the log-likelihood function is up to some constant

$$\log l(\zeta) = -\sum_{i=0}^{N-1} \frac{\zeta_i^2}{2\sigma_i^2}. \quad (7.44)$$

Note the similarity between this form of the likelihood and the exponential of the sum-squared errors. There is still an issue over the choice of σ_i for each observable i . It was found for typical simulation runs that the variance of the rate of change of the number of pairs scales with the distance as a square (Fig. 7.8). Hence, the functional form of $\sigma_i(d) = \sigma N(d)^2$ for both 00 and 11 pairs, 01 does not obey a simple scaling relationship however. The functional form of the variance then leaves a single free parameter σ for the likelihood, which can then be found through simulation. For example, a 100×100 lattice gives a variance of $\sigma = 10^{-6}$

Both the gamma likelihood and the normal likelihood were used in model fitting, however most of the analysis uses the gamma form of the likelihood.

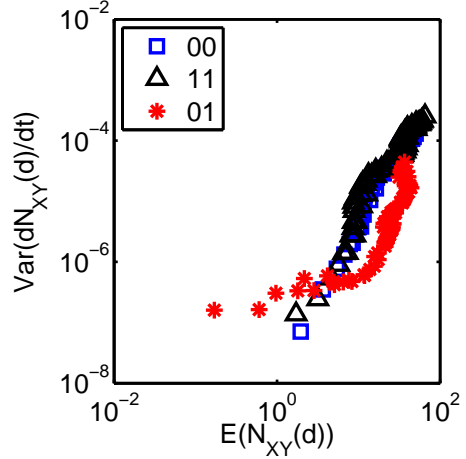


Figure 7.8: Expected number of XY pairs at a given distance d against the variance of the rate of change of the number of XY pairs. The scaling for both 00 and 11 is approximately square.

Size (L)	Error	mean	standard deviation	p-value
50	P_{00}	-19.6746	373.64	0.14304
	P_{01}	5.6085	176.8485	0.09033
	P_{11}	14.0662	197.3716	0.25076
70	P_{00}	-21.0737	1211.4098	0.40162
	P_{01}	42.3027	497.5222	0.30539
	P_{11}	-21.2289	715.431	0.49747
90	P_{00}	-235.2413	2500.6328	0.77123
	P_{01}	61.3701	1059.2863	0.61697
	P_{11}	173.8711	1444.6568	0.97663
110	P_{00}	-42.1362	4429.6086	0.86911
	P_{01}	19.2043	1922.7386	0.97634
	P_{11}	22.9318	2512.4194	0.77743
130	P_{00}	31.6279	7355.5629	0.39338
	P_{01}	-32.1154	3160.5081	0.21348
	P_{11}	0.48749	4199.9252	0.54093

Table 7.1: Measurement of errors for different system sizes and the corresponding p-value for the corresponding one-sample Kolmogorov-Smirnov test. A p-value of 0.05 or lower would indicate that the distribution is not Gaussian.

7.4.5 Priors

For Bayesian inference we wish to calculate the posterior $P(\theta|D)$ from the likelihood. This is accomplished using Bayes theorem

$$P(\theta|D) = \frac{P(D|\theta)P(\theta)}{P(D)}. \quad (7.45)$$

The posterior $P(\theta|D)$ is the probability of observing the parameters θ given the data D . This is calculated as the probability of the data given the parameters θ , which is the likelihood, multiplied by the probability of the parameters, divided by the normalisation factor $P(D)$. The normalisation factor $P(D)$ can be calculated using the marginal probabilities

$$P(D) = \sum_{\theta} P(D|\theta)P(\theta), \quad (7.46)$$

this is however, computationally intractable for all but trivial statistical models and hence methods such as MCMC are used instead. The other factor that needs to be determined is $P(\theta)$. This is known as the prior and represents the belief in the parameters before the data [Ellison, 2004]. The banding model has five parameters representing the displacement and direction of the competition r, θ , the length scales

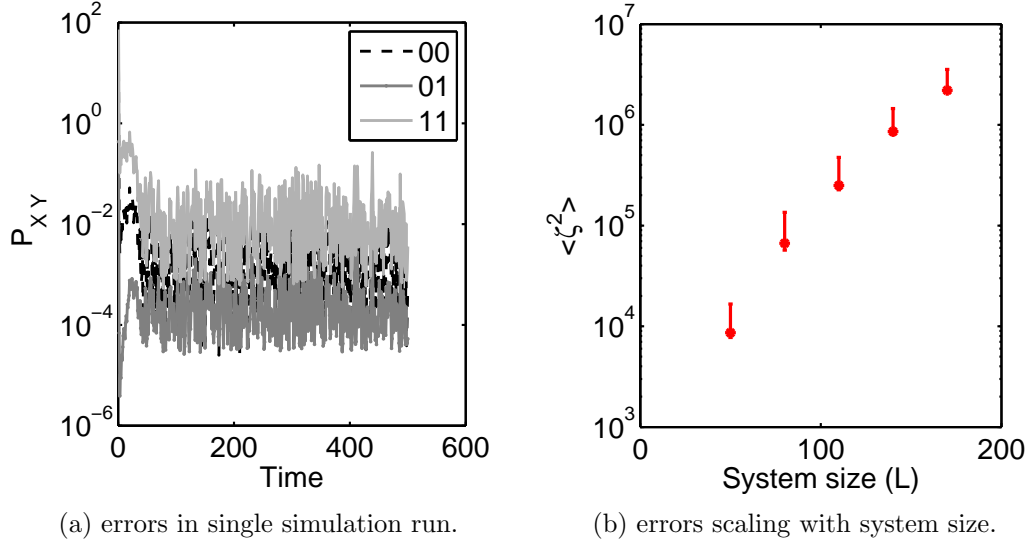


Figure 7.9: How the distribution of errors ζ^2 grows for increasing system size. Simulations were run for each system size with fixed parameters for a suitable time until enough values of ζ^2 had been calculated from the snapshots in time. The errors grow almost linearly on a semi-log scale. The bars around the points represent the lower and upper quartile respectively. Fig. 7.9a shows how the individual errors for each of the correlation observables change as the system equilibrates.

of the growth and competition l_1, l_2 and the competition factor k . It is assumed that the parameters are independent and so the prior has the following multiplicative structure

$$P(\Theta) = P(\theta)P(r)P(l_1)P(l_2)P(k). \quad (7.47)$$

Maximum entropy distributions

For a distribution $p(x)$ with support $[a, b]$, the entropy of a distribution is a scalar defined as

$$S = - \int_a^b p(x) \log p(x) dx. \quad (7.48)$$

This quantity gives a measure of uncertainty in the distribution. The maximum entropy approach is to then find a distribution $p(x)$ that maximises S with certain constraints that are known about the distribution, such as the mean or variance [Jaynes, 2003].

Assume that the mean of a distribution is known and denote it μ . Using the calculus of variations a Lagrangian can be constructed

$$L = S + \lambda_0 \left(\int_a^b p(x) dx - 1 \right) + \lambda_1 \left(\int_a^b xp(x) dx - \mu \right). \quad (7.49)$$

The fixed point of L then gives the probability distribution that maximises S for a fixed mean. Consider the derivative with respect to $p(x)$ for a fixed x .

$$\begin{aligned} \frac{\partial L}{\partial p(x)} &= -1 - \log p(x) + \lambda_0 + \lambda_1 x = 0 \\ \implies \log p(x) &= \lambda_0 + \lambda_1 x \\ \implies p(x) &= \exp(\lambda_0 + \lambda_1 x). \end{aligned}$$

This can be simplified to

$$p(x) = c \exp(\lambda_1 x). \quad (7.50)$$

Differentiating the Lagrangian with respect to the Lagrangian multiplier λ_0 then gives the normalisation condition for the distribution $p(x)$. For our purposes all parameters of the model exist on the interval $[0, \infty)$ (with the exception of the directional parameter θ). Inputting this into the normalisation condition

$$\begin{aligned} \int_0^\infty c \exp(\lambda_1 x) dx &= -\frac{c}{\lambda_1} = 1 \\ c &= -\lambda_1 \end{aligned}$$

Differentiating with respect to the second Lagrange multiplier gives

$$\begin{aligned} \int_0^\infty -x \lambda_1 \exp(\lambda_1 x) dx &= -\frac{1}{\lambda_1} = \mu, \\ \lambda_1 &= -\frac{1}{\mu}. \end{aligned}$$

Hence the probability distribution that maximises entropy on the interval $[0, \infty)$ with a known expectation μ is none other than the exponential distribution. This technique can also be applied to distributions on other intervals. For instance, we may know that the distance of a band cannot be over a certain height, and so r would lie on an interval $[0, r_{\max}]$, this reduces to adding a normalisation factor to the standard exponential distribution so it integrates to one on the defined interval.

Directional statistics

The prior for θ also needs to be considered. θ is defined on the wrapped interval $[0, \pi)$ i.e.

$$\theta = x \bmod \pi, \quad (7.51)$$

where x is a corresponding linear random variable. If x is defined on the interval $[0, \infty)$ and has the corresponding exponential distribution $p(x) = \frac{1}{\lambda} \exp(-\lambda x)$ then the distribution for θ is calculated as

$$\begin{aligned} p_w(\theta) &= \sum_{k=0}^{\infty} p(\theta + \pi k) \\ &= \sum_{k=0}^{\infty} \frac{1}{\lambda} \exp(-\lambda(\theta + \pi k)) \\ &= \frac{1}{\lambda} \exp(-\lambda\theta) \sum_{k=0}^{\infty} \exp(-\lambda\pi k) \\ &= \frac{\exp(-\lambda\theta)}{\lambda(1 - \exp(-\lambda\pi))}. \end{aligned}$$

This produces an exponential with an added normalisation factor to take into account the fact that the distribution is defined on a wrapped interval. It is straightforward to show that, on a wrapped interval, a distribution with known expectation that maximises the entropy is $p_w(\theta)$, known as the wrapped exponential distribution.

There is an issue however, as slightly more is known about the distribution $p_w(\theta)$ than the expectation alone. In particular, if the expectation of θ is close to 0 then all of the probability density will be concentrated on the positive plane and none near π . As the interval is wrapped so π is matched up with 0 this is clearly erroneous. To fix this we need a distribution that is truly circular i.e. will be continuous on the whole wrapped interval. One such distribution is the von Mises distribution defined as

$$p(\theta) = \frac{\exp(\kappa \cos(\theta - \mu))}{2\pi I_0(\kappa)}, \quad (7.52)$$

where κ is analogous to the precision and μ to the expectation in a normal distribution. The one issue with this distribution is finding a way of estimating the variance of the prior. This could be done by considering the variation of the spatial pattern for all θ and taking the variation of this as an estimate for κ .

The simplest form of a probability distribution defined on C^1 such that it is continuous would be the uniform distribution $U[0, 1]$. This can be used when there is an absence of prior knowledge about the direction of the banding. However, where banding is present, a two-dimensional Fourier transform can be used to detect the direction and strength of banding. The largest Fourier mode is then used as an estimate for the direction of the banding. This is taken as the mean of an empirical circular distribution, where the parameter $x = \sin \theta$ has a peak at the maximal direction of the Fourier modes.

The prior distributions of l_1 and l_2 can be similarly estimated from the spatial snapshot by measuring the length-scale of the banding. This is again done using the spatial Fourier transform of the banding snapshot, where the displacement of the largest Fourier mode is used as the mean for the exponential distribution of l_1 and l_2 . Finally the exponential distribution for k is also used with mean taken to be the power of the largest Fourier mode. This completely defines the prior for the model.

7.4.6 Fitting to real data

Fitting to real data is problematic due to the boundary conditions. The simulations up until now have been performed on a toroidal lattice. However, this does not apply to a real-world dataset, where the form of the boundary conditions for a single spatial snapshot are not explicitly known.

If D is the data in the form of a binary image of size $n \times n$, then take $E \subset D$ to be the centered box of length m (hence $m < n$). The set $D - E$ represents the border around the inner-box. The rate of change at each point in E is considered, however when the change in correlation due to an event in E is calculated, the whole of D is taken into account. As long as n is sufficiently greater than m and the parameters l_2 and r do not get too large, this provides a true description of the boundary without having to consider occupancy beyond the outer boundary. The downside to this approach is that the rate of change in the correlation structure for lattice sites around the edge is not explicitly considered.

7.5 Results

7.5.1 Simulated data

With the machinery in place to perform likelihood-based inference on the banding model, we first consider a simulated spatial snapshot with toroidal boundary condi-

tions and model parameters $r = 10, \theta = \pi/4, l_1 = 0.5, l_2 = 1, k = 0.1$. The toroidal boundary conditions are used so that boundary conditions can be ignored while the method is tested. Visual inspection of the likelihood as well as the sampled parameters indicates that the chain after a sufficient burn-in time is well-mixed and the chain is sufficiently exploring the posterior. The first attempt at inferring parameters from the single snapshot was performed on a synthetic likelihood that is the negative of the total error ζ^2 . This likelihood was not strongly informative as the posterior differed only marginally from the prior distribution. There was however some indication that the method could detect both the competition factor k and the angle of the competition offset θ . The same inference was performed using the exponential likelihood function $l(\zeta^2) = \exp(-\lambda\zeta^2)$, where λ was calculated using the mean total error from a single simulation run at statistical stationarity. The chain was run for 10^6 iterations with a burn-in time of 0.5×10^6 . The resulting posterior differed significantly from the prior distribution (Fig. 7.10), with all parameters correct within the 95% confidence interval. Whilst there is low correlation between most parameters, there is a significant correlation structure exhibited between r and θ . The resulting structure represents a pay-off between the angle and the length between bands. If the angle is perpendicular to the bands, then the offset of the competition is minimised in order for the centre of the competition kernel to be in the centre of the gap, as the angle θ changes, r must compensate in order for the centre of the kernel to remain in the gap. This produces a bent curve where r achieves its minima approximately by r 's true value. l_1 and l_2 have similar profiles for their marginal posterior distributions, although the prior is the same for both, their posteriors differ, with the mean of the distribution close to the true underlying value. k is the only parameter where the mean of the marginal posterior differs from the true underlying value.

k represents the strength of competition and is implemented via the death term, where it is the exponent of the competition kernel convolved with the lattice state S . a large exponent increases the probability of a death event when the convolution is greater than one and decreases it when it is less than one, this means that for a large competition factor, the competition kernel has highly platykurtic. This means that k is only strongly affecting the tails of the competition kernel. This can be elucidated by running a simulation of banding with a low competition factor to statistical stationarity and then observing the expected rate of change for each of the sites when k is larger (Fig. 7.12a). The sites with the largest rate of death occur along the borders of the band where the rate of change of sites is already

highest, thus a change in k may not be detectable from the background stochastic fluctuation of the system. If the simulation is then continued, this time with a new competition factor k , then we may measure the change in the correlation structure as the system evolves using the Kullback-Liebler divergence. For a change in k from 0.1 to 1, there is little change in the correlation structure over small time-scales, although there is over larger time-scales although there is significant change over larger time-scales (Fig. 7.12b). For larger differences in k (0.1 to 20) there is however an immediate change in the correlation structure of the snapshots at each time-step even on relatively small time-scales (Fig. 7.12c). The method relies on there being a change in the correlation in a single time-step, for small changes in k , this change in correlation occurs slowly and hence k ends up with a broad distribution in the posterior.

As the competition factor is weakly fitted for simulated data, another form of the banding model is considered. Instead of k being a competition factor that exponentiates the competition kernel, a rate of competition (β) is considered instead. Hence the rate of transition for a site x is

$$r_x(0 \rightarrow 1) = (k_1 * S)_x, \quad (7.53a)$$

$$r_x(1 \rightarrow 0) = \beta(k_2 * S)_x, \quad (7.53b)$$

where k_1, k_2 are defined as before in Eq. 7.6. As the model is now in term of rates, there needs to be a conversion to probabilities. This is achieved by assuming each individual site is a Poisson process with rates equal to the rate of transitions. Assuming a small time interval ϵ , the resulting probability of transitions are

$$P_x(0 \rightarrow 1) = 1 - \exp(-\epsilon(k_1 * S)_x), \quad (7.54a)$$

$$P_x(1 \rightarrow 0) = 1 - \exp(-\epsilon\beta(k_2 * S)_x). \quad (7.54b)$$

The previous exponential form of the likelihood was compared to the normal likelihood defined in Eq. 7.44, with the new form of the model defined in Eq. 7.54. Similar parameter values were used for the simulation $r = 10, \theta = \pi/4, l_1 = 1, l_2 = 2, \beta = 2$ as for the exponential likelihood case. The resulting posterior distribution matches closely with the true underlying parameter values (Fig. 7.11). The θ marginal posterior has a bimodal structure indicating the invariance under reflection in the spatial snapshot. All other marginals are centred on the true parameters, while l_1 and l_2 have discrete cut-offs due to the discreteness of the lattice. As there is evidence the normal form of the likelihood provides a better tool for inference on simulated data,

this form of the likelihood shall be used to perform inference on real data.

7.5.2 Real data

With the validity of the scheme for simulated data established we can now turn our attention to real data in the form of three different vegetative ecosystems that all exhibit banding phenomena: seagrass, tiger bush and mussel beds.

The normal form of the likelihood (Eq. 7.44) was applied to two different spatial snapshots: seagrass and tiger bush. The first snapshot was taken from an area banding observed in Isles of Scilly seagrass dataset. The binary image was rescaled so that it fit on a 100×100 lattice, with each pixel representing the presence of vegetation. The resulting marginal posteriors have small variance and are unimodal (Fig. 7.13). The maximum a posteriori for each parameter was used in a simulation with random initial conditions and ran for 500 time-steps. The resulting spatial distribution closely matches the original data, with the direction of the banding, inter and intra band length correctly reproduced.

A similar procedure was conducted on a snapshot of tiger bush in South West Niger ($13^{\circ}30' \text{ N}$; $2^{\circ}40' \text{ E}$), taken from Thiery et al. [1995]. The posterior distribution is flat for r , θ and l_2 (Fig. 7.14). The marginal posterior for l_1 has a maximum at 1 and over 95% of its mass less than 1. β has an extremely sharp distribution at 0. The poor fit of parameters is reflected in the simulation based on the maximum a posteriori, where the simulated spatial pattern matches poorly to the original data. In particular there is no banding present in the simulated spatial pattern as the low competition rate leads to a low death rate and hence the pattern is uniformly occupied.

7.6 Conclusion

In this chapter a method for performing Bayesian inference for a model of vegetation banding has been developed. First the expected rate of change in density was used as a measure of how a chosen parameter set deviates from the true parameters that produced a spatial snapshot. This provided a way of fitting single spatial parameters when other parameters were known, however was unable to differentiate cases with two or more unknown parameters.

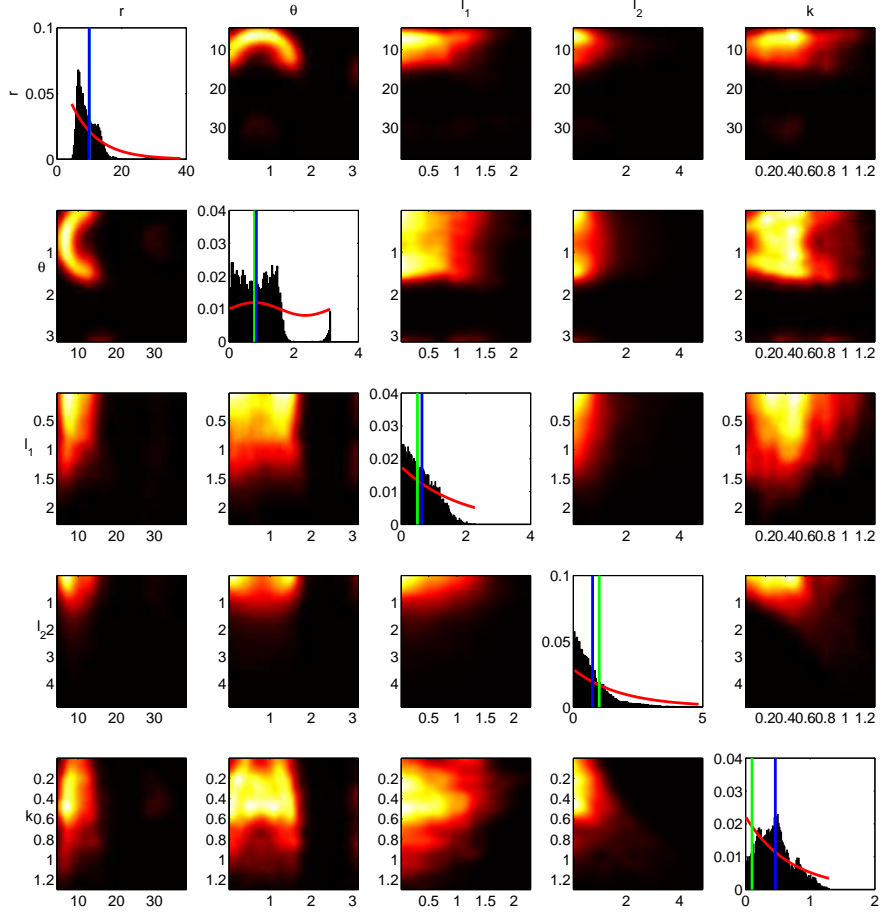


Figure 7.10: Example of posterior for a 150×150 spatial snapshot ran until equilibrium with parameters $r = 10, \theta = \pi/4, l_1 = 0.5, l_2 = 1, k = 0.1$. The exponential likelihood $\exp(-\lambda\zeta^2)$ is applied, where λ is calculated from the errors of the simulation run. True parameters are highlighted in green, while the mean of the posterior is shown in blue. The crescent shape indicates the play-off between r and θ in a single band.

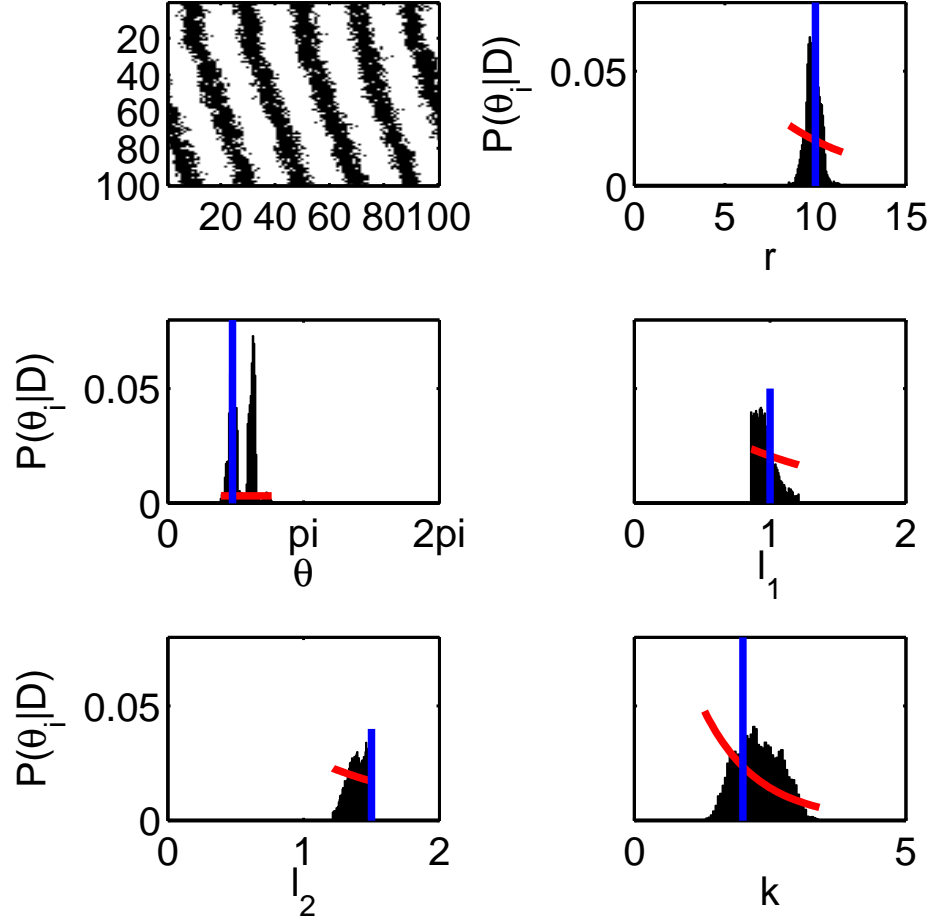
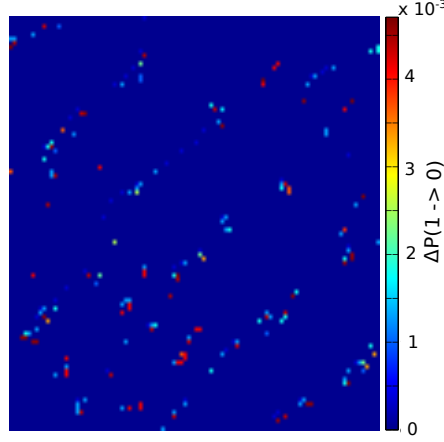
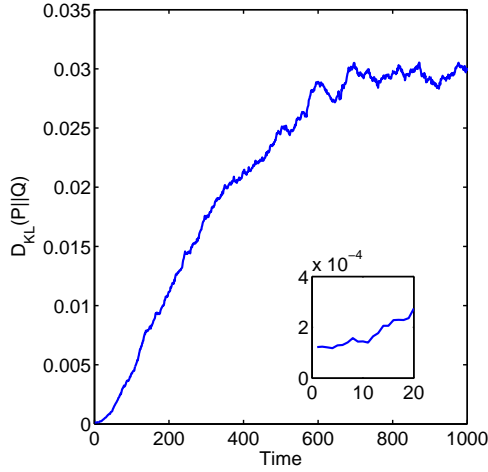


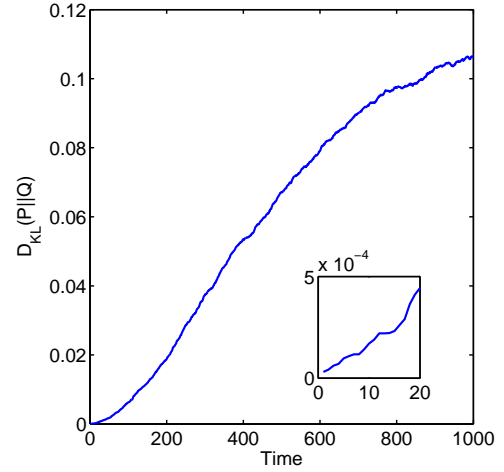
Figure 7.11: Marginals of posterior for a 100×100 spatial snapshot ran until stationarity with parameters $r = 10, \theta = \pi/4, l_1 = 1, l_2 = 2, k = 2$. The full normal likelihood defined in Eq. 7.44 is applied. The top left is the simulation snapshot on which inference was performed. Each of the marginal posteriors (black) are compared to the priors (red) and the true parameter values (blue).



(a) change in rate of death



(b) $k = 0.1$ to $k = 1$.



(c) $k = 0.1$ to $k = 20$.

Figure 7.12: Change in the Kullback-Liebler Divergence of the P_{11} distribution. The evolves at $k = 0.1$ until there is only small statistical fluctuations in the correlation structure. The system parameter k is then switched and the KL-divergence is recorded compared against the new correlation distribution P_{11} , with the reference distribution as P_{11} at stationarity. For a smaller change in k there is little change in the P_{11} distribution for the first few number of time-steps, but increases for longer time-steps. This provides evidence as to why k is not strongly selected for in the inference procedure, as the correlation structure only changes dramatically for large values of k . Fig. 7.12a shows the change in the death rate ($1 \rightarrow 0$) from a simulation with $k = 0$ switching to $k = 1.5$ after equilibrium time, where all other parameters are kept constant. Due to the nature of k , the Gaussian death kernel only strongly changes at it's edge for increasing k . At the edge of the kernel there are few sites in the alive state, hence a change in k does not change the rate of these events to a huge degree. Indeed, for this case the average change in probability is approximately 0.01.

The next step was to consider the expected rate of change in pair-wise correlations for a spatial snapshot. It was found that this method was able to resolve all parameters of the model with other parameters kept fixed. This was then applied to two unknown parameters on test data, with other parameters fixed. The displacement parameter r and angle parameter θ both minimised the expected change in correlation close to the true values of the parameters.

For fitting all five parameters to a spatial snapshot an exhaustive approach would be too computationally costly. Therefore the approach was adapted into the framework of Bayesian inference. This was achieved by deriving a synthetic likelihood that is based upon the sum squared errors of the correlation functions. The resulting likelihood was then combined with empirical priors based on spatial estimates using a two-dimensional Fourier transform. Further discussion of the likelihood lead to considering each correlation term individually, where each correlation was assumed to have normal error with an individual variance.

The method with sum of squares likelihood, was first tested on simulated data where it was found to be able to recover all parameters within 95% error margins except for the competition factor k . k was found to be only weakly resolved, as a change in k relates mainly to the boundaries of the competition kernel, for a spatial pattern at statistical equilibrium this would only strongly affect occupied sites at the edge of bands. These sites are naturally subject to higher death rate, due to fluctuations in the banding and hence do not contribute a significant amount to the correlation structure. This was further investigated by observing how the correlation structure changes in time as a simulation is equilibrated at a small value of k and then observing the changing correlation structure when k is altered to a large value. For significant changes in k little change in the correlation structure is observed on small time-scales, although on larger time-scales this is observed.

Due to the problematic nature of resolving the k parameter the model was altered such that k was replaced by β , which represents the competition rate and the transition probabilities were calculated by assuming each site in the lattice is a Poisson process. This was combined with the normal form of the likelihood and the method was again tested on simulated data. It was found that the method was able to identify all five parameters of the model within 95% confidence intervals of the posterior and hence this form of the likelihood and model was applied to real data.

The method was then applied to a variety of spatial patterns from various ecosystems that exhibit banding: tiger bush and seagrass. The method was able to resolve the direction and distance of the banding for the seagrass snapshot, however was unable to for the tiger bush snapshot. There are many reasons as to why the method failed to capture appropriate parameters. In order to represent the data on a binary lattice, several image processing steps were performed. These methods include having a cut-off for when a site is occupied. This cut-off implies sites near the boundary of banding would be subject to noise and hence may not be properly categorised. An extension would be to consider instead the probability of a site being occupied by using a greyscale image of the vegetation. The method can then be adapted to consider probabilities of a site's occupation. The particular form of the competition kernel may also be misspecified for the tiger bush system.

Banding is but one of a plethora of spatial patterns observed in vegetation ecosystems. These patterns include spots, labyrinths and barren patches known as “fairy circles” [Van Rooyen et al., 2004]. Models have been proposed to explain the mechanisms of these patterns, but none have been explicitly fitted to spatial data [Barbier et al., 2008; Fernandez-Oto et al., 2013; Tlidi et al., 2008]. The method proposed in this chapter gives a way of fitting parameters to data for a general probabilistic cellular automata model. An extension would then be to apply this method to other models of vegetation pattern formation in order to fit to more general spatial patterns.

The main spatial observable used has been the pair-wise spatial correlation function. A further extension would be to not only consider pairs of sites, but also triples of sites. This would produce a significant increase in the computation of the expected change in the observable, but may be able to more accurately distinguish between spatial patterns, especially where more than one direction to the pattern is observed. An alternative approach could be to split the spatial data into overlapping boxes of a certain length and fit the model to each of these boxes. The results would then indicate how local anisotropy and competition effects vary throughout the space.

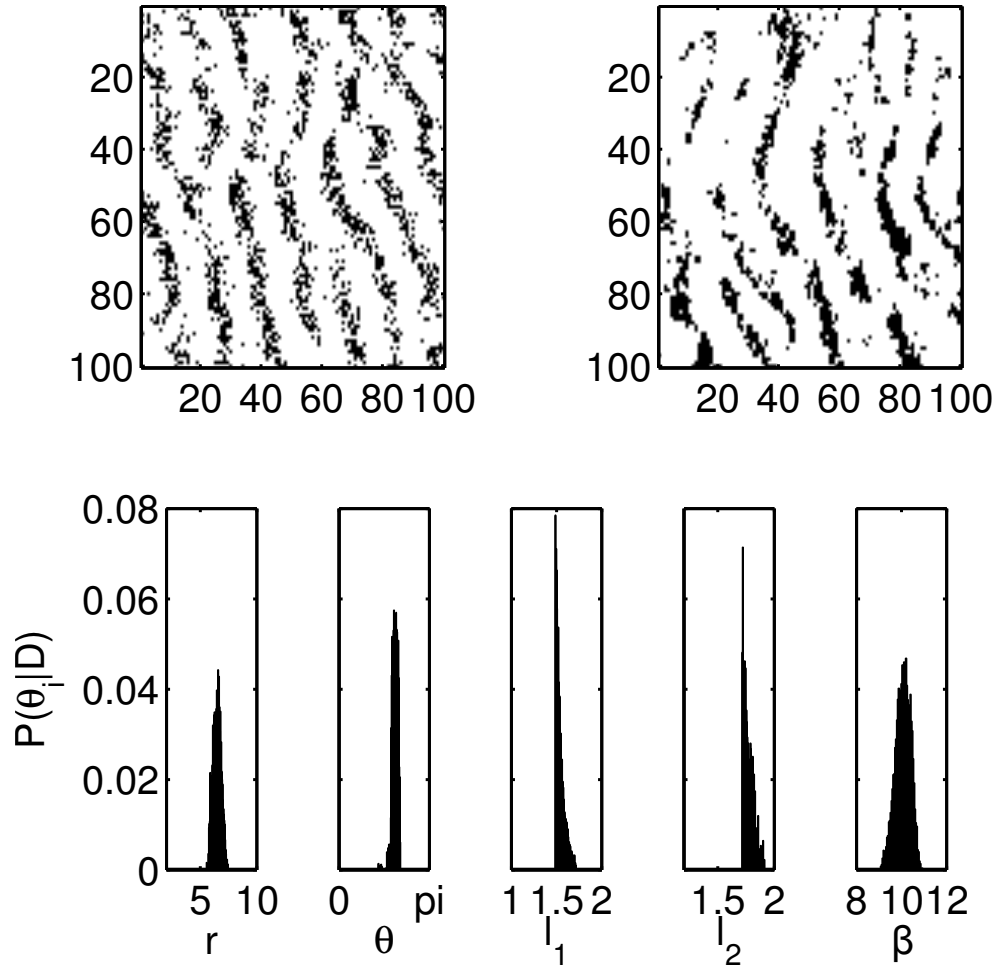


Figure 7.13: Top-right: Banding pattern observed in seagrass. Top-left: An example simulation with maximum a posteriori parameters using the normal likelihood (Eq. 7.44) and the rate model (Eq. 7.54). Bottom: marginal probability distribution for each parameter.

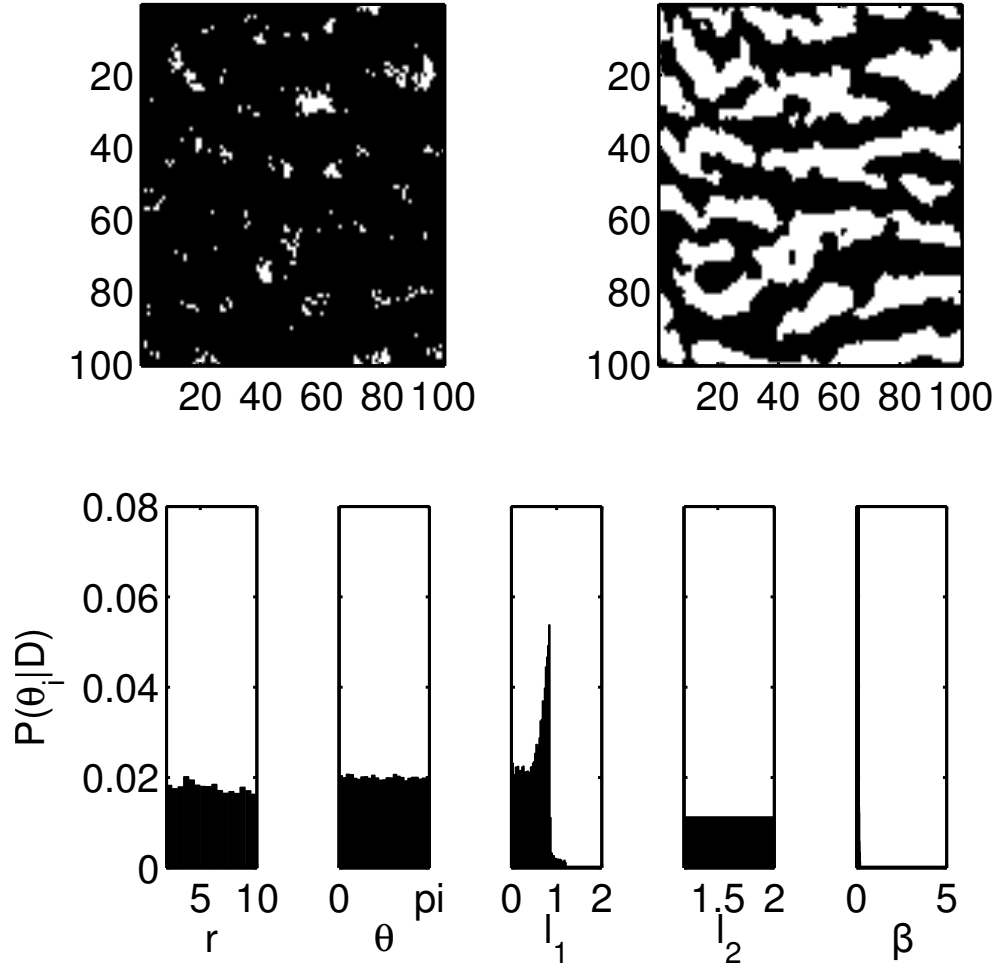


Figure 7.14: Top-right: Banding pattern observed in tiger bush (reproduced from Thiery et al. [1995]). Top-left: An example simulation with maximum a posteriori parameters using the normal likelihood (Eq. 7.44) and the rate model (Eq. 7.54). Bottom: marginal probability distribution for each parameter using the normal likelihood.

Bibliography

- Nicolas Barbier, Pierre Couteron, René Lefever, Vincent Deblauwe, and Olivier Lejeune. Spatial decoupling of facilitation and competition at the origin of gapped vegetation patterns. *Ecology*, 89(6):1521–1531, 2008.
- Mark A Beaumont, Wenyang Zhang, and David J Balding. Approximate bayesian computation in population genetics. *Genetics*, 162(4):2025–2035, 2002.
- Aaron M Ellison. Bayesian inference in ecology. *Ecology letters*, 7(6):509–520, 2004.
- C Fernandez-Oto, M Tlidi, D Escaff, and MG Clerc. Strong interaction between plants induces circular barren patches: fairy circles. *arXiv preprint arXiv:1306.4848*, 2013.
- GJ Gibson and EJ Austin. Fitting and testing spatio-temporal stochastic models with application in plant epidemiology. *Plant Pathology*, 45(2):172–184, 1996.
- W Keith Hastings. Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.
- Edwin T Jaynes. *Probability theory: the logic of science*. Cambridge university press, 2003.
- Matt J Keeling, Stephen P Brooks, and Christopher A Gilligan. Using conservation of pattern to estimate spatial parameters from a single snapshot. *Proceedings of the National Academy of Sciences of the United States of America*, 101(24):9155–9160, 2004.
- Scott Kirkpatrick, C Daniel Gelatt, Mario P Vecchi, et al. Optimization by simulated annealing. *science*, 220(4598):671–680, 1983.
- Nicholas Metropolis, Arianna W Rosenbluth, Marshall N Rosenbluth, Augusta H Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6):1087–1092, 1953.

- Gopal G Penny, Karen E Daniels, and Sally E Thompson. Local properties of patterned vegetation: quantifying endogenous and exogenous effects. *arXiv preprint arXiv:1303.4360*, 2013.
- Jonathan K Pritchard, Mark T Seielstad, Anna Perez-Lezaun, and Marcus W Feldman. Population growth of human y chromosomes: a study of y chromosome microsatellites. *Molecular Biology and Evolution*, 16(12):1791–1798, 1999.
- Max Rietkerk and Johan Van de Koppel. Regular pattern formation in real ecosystems. *Trends in Ecology & Evolution*, 23(3):169–175, 2008.
- Michael S Rosenberg. Wavelet analysis for detecting anisotropy in point patterns. *Journal of Vegetation Science*, 15(2):277–284, 2004.
- Simon Tavaré, David J Balding, Robert C Griffiths, and Peter Donnelly. Inferring coalescence times from dna sequence data. *Genetics*, 145(2):505–518, 1997.
- JM Thiery, J-M d’Herbes, and Christian Valentin. A model simulating the genesis of banded vegetation patterns in niger. *Journal of Ecology*, pages 497–507, 1995.
- Mustapha Tlidi, René Lefever, and A Vladimirov. On vegetation clustering, localized bare soil spots and fairy circles. In *Dissipative Solitons: From Optics to Biology and Medicine*, pages 1–22. Springer, 2008.
- MW Van Rooyen, GK Theron, N Van Rooyen, WJ Jankowitz, and WS Matthews. Mysterious circles in the namib desert: review of hypotheses on their origin. *Journal of Arid Environments*, 57(4):467–485, 2004.
- Christopher K Wikle. Hierarchical bayesian models for predicting the spread of ecological processes. *Ecology*, 84(6):1382–1394, 2003.
- Simon N Wood. Statistical inference for noisy nonlinear ecological dynamic systems. *Nature*, 466(7310):1102–1104, 2010.

Chapter 8

Aggregation

In this house, we obey the laws of thermodynamics!

(Homer J. Simpson)

8.1 Introduction

In Chapter 6 a variety of spatially explicit models were explored to elucidate the relationship between the scaling of the patch-size distribution and the dynamic persistence of the vegetation system. It was found that a power-law distribution provides a good fit to the patch-size distribution under a variety of conditions, however there are marginal cases to this. Kéfi et al. [2007] analysed patch-size distributions in semi-arid vegetation in the Kalahari and found that there was not only a power-law distribution evident in the patch-size distribution, but also a truncated exponential term, when the system was under increased grazing pressure. Similar phenomena have also been detected in a number of other ecosystems including mussel beds [Guichard et al., 2003] and marine benthic diatoms [Weerman et al., 2012].

The leading explanation of this power-law pattern formation is due to local interactions driving the large-scale behaviour [Scanlon et al., 2007b]. This has been supported through the use of numerical simulation of spatially-explicit models of vegetation growth combined with a global effect on the population density interpreted as the amount of rainfall or other global processes. This explanation does not answer how a power-law should form on a more fundamental level, whether it is due to an aggregation of smaller clusters or a competition effect between larger clusters dominating the landscape.

In this chapter spatially implicit models of vegetation clusters are investigated by considering how patches form and aggregate, instead of modelling vegetation as a probabilistic cellular automata as was done in chapter 5 and 6. The general conditions under which a power-law distribution is expected to emerge are explored as well as when there is a meltdown of the power law distribution due to an exponential truncation.

Models of aggregation and fragmentation have been considered in other areas in ecology such as the size of fish schools [Niwa, 1998] and marine diatoms [Jackson, 1990]. Aggregation phenomena has been more generally studied in the Physical sciences [Aldous, 1999]. This chapter gives an overview of these models and their solutions as well as using these solutions to give insight into vegetation ecosystems.

Further, the connection between the power-law exponent and the persistence of the distribution in this model is explored. We begin with defining a novel model of aggregation with linear death and then deriving an asymptotic solution when

the death rate is small. A more general model of aggregation with death is then derived and applied to the Isles of Scilly, UK patch-size data. The conclusion is that the power-law clustering observed in many vegetation ecosystems may simply be an aggregation effect and the exponential truncation observed when there is increased stress is due to an increase in the linear death rate of single units connected to a cluster.

8.2 Introduction to equations

My initial approach was to explicitly model the whole environment and species via a probabilistic cellular automata and use the subsequent patch-size distribution to make inferences about its relationship to the return rate. However, this does not model either the patch-size distribution or the return rate explicitly and does not give a theoretical understanding of how the power-law distribution emerges and under what conditions it should be observed.

The idea developed here is to model the patches themselves as opposed to individual sites. We denote $c_k(t)$ as the density of patches of size k at time t , where time is taken to be continuous i.e. $t \in [0, T] \subset \mathbb{R}$. A continuous model of patch-sizes can be studied, however for the present k shall take positive integer values only, $k \in \mathbb{N}/\{0\}$. A kernel of aggregation gives the rate at which patches of size i and j aggregate together to form a patch of size $i + j$, this kernel is denoted $K(i, j)$. Finally it is assumed there is a constant rate at which patches of size 1 or monomers enter the system. The governing master equation, also known as the Smoluchowski equation [Von Smoluchowski, 1916] is then

$$\frac{d}{dt}c_k = \frac{1}{2} \sum_{i+j=k} K(i, j)c_ic_j - \sum_{j \geq 1} K(j, k)c_jc_k + \delta_{k,1}. \quad (8.1)$$

Various properties are desirable for the kernel. Firstly commutativity, where the rate at which patches of size i and j aggregate does not depend on the ordering of the patches i.e. $K(i, j) = K(j, i)$. Secondly, scaling homogeneity, where the rate at which patches of a certain size aggregate scales by some factor $K(ai, aj) = a^\lambda K(i, j)$. The simplest kernel that satisfies these conditions is the constant kernel $K(i, j) = 1$. When this is the case the tail-solution (for large k) has the simple form [Hayakawa, 1987]

$$c_k \sim \frac{1}{\sqrt{4\pi}} \frac{1}{k^{3/2}}. \quad (8.2)$$

The tail of the patch-size distribution is a power law with exponent $-3/2$, where the power law nature of the solution is a consequence of the injection term (where births of patch size one enter the system) and the non-linear aggregation function. The equation can be solved analytically for more general kernels of the type

$$K(i, j) = i^{-a} + j^{-a}. \quad (8.3)$$

This generalisation needs an interpretation in the context of growing fixed vegetative clusters, which shall be discussed in section 8.3. This type of kernel also admits an analytic solution in the large patch-size limit [Krapivsky et al., 1999; Krapivsky et al., 2010] with a steady state distribution of the form $c_k \sim Ck^{-\tau}$ where

$$\tau = \frac{3-a}{2}, \quad (8.4)$$

$$C = \sqrt{\frac{1-a^2}{4\pi}} \cos\left(\frac{\pi a}{2}\right). \quad (8.5)$$

For a steady state to exist there must be the condition $-1 < a < 1$ and hence the scaling exponent can be found on the interval $\tau \in (1, 2)$. The dynamics of the equation can be probed by defining the cross-over time, which is the time taken for a density of patches of a certain size to reach its asymptotic value. The cross-over time for a patch of size k_* to the steady state solution c_{k_*} can also be calculated giving $t = (k_*)^z$ where $z = (1+a)/2$. The scaling of the cross-over time and the patch-size exponent can be related by the simple linear equation $\tau = 2 - z$. This gives a linear relationship between the static exponent at stationarity and its dynamic exponent, which may be associated with the return rate. Hence the conclusion is the power-law patch size and relationship between return rate and the resulting distribution can be explained by an aggregation process.

A real vegetation system is not purely defined by an aggregation process however. In particular in the previous example there is no death either of single vegetation units or patch clusters. Death may lead the fracturing of clusters and to changes in the exponent of the stationary distribution and so it is important to include in any model of vegetation clustering. In order to produce a model that is analytically tractable, it is assumed that the death rate is constant across all individual units and the death of individual units comes from the boundary of the cluster i.e. a death event in a cluster does not fragment the cluster into two or more clusters. A

modified Smoluchowski equation with a linear death term can then be produced as

$$\frac{d}{dt}c_k = \frac{1}{2} \sum_{i+j=k} K(i, j)c_i c_j - \sum_{j \geq 1} K(j, k)c_j c_k + \mu c_{k+1} - \mu c_k, \quad (8.6a)$$

$$\frac{d}{dt}c_1 = - \sum_{j \geq 1} K(j, 1)c_j c_1 + 1 + \mu c_2 - \mu c_1. \quad (8.6b)$$

8.3 An ecological interpretation of the aggregation kernel

The general additive aggregation kernel is of the form $K(i, j) = i^{-a} + j^{-a}$, where a represents the scaling parameter of the rate at which aggregates of a certain size join. If it is equally likely for a cluster of a certain size to aggregate with a cluster of any other size then the scale parameter $a = 0$. For a pure aggregation system with no fragmentation, this leads to a cluster scaling of $3/2$.

It is instructive to imagine a single unit or monomer coming into contact with a cluster and calculating the rate at which this occurs for larger as opposed to smaller clusters. If $a > 0$ then, assuming the size of the monomer is negligible, the monomer rate equation is $K(i) = i^{-a}$. This means smaller clusters are favoured and the growth rate reduces as clusters grow larger in size.

An ecological explanation of this could be due to the self-limitation through competition a larger cluster experiences with itself, thus reducing its potential for growth. Smaller clusters have more space and thus can grow at a quicker rate. When $a < 0$, larger clusters are favoured for growth compared with smaller clusters, this can be seen as a form of the Allee effect [Stephens and Sutherland, 1999]. Small clusters are more susceptible to environmental perturbation and as such, have a lower propensity for growth. Larger clusters of vegetation are able to regulate their environment more and thus have greater resources for growth (An example species where this holds is ribbed mussels [Bertness and Grosholz, 1985], where larger clusters provide protection and shelter for new mussels). A value for a then can give an indication of whether there is strong small cluster growth at the expense of large clusters forming or if the converse holds.

An alternative explanation of the aggregation exponent a is due to the edge effects of a cluster. A single individual vegetation unit aggregates to a cluster proportional to the edge of that cluster. If all cluster are non-fractal then it would be expected

that a vegetation unit aggregates at rate $i^{1/2}$, since the length of a non-fractal object scales as a square root with its area. For a general fractal cluster with boundary dimension d , it would be expected that an individual unit scales as $i^{1/d}$.

8.4 An analogue for the return rate

General aggregation-fragmentation systems do not have analytical solutions and only in a few special cases can exact solutions be found. Hence, some alternative way of describing the return rate of the aggregation system that is not dependent on the steady-state solution being known must be sought. One potential strategy to probe the dynamics is to consider the crossover time for a system. This can easily be computed if it is assumed the patch distribution has the asymptotic solution $c_k \approx Ck^{-\tau}$ (note τ is the Korcak exponent for this system). If the injection of new patches of size 1 occurs at rate 1 and the aggregation process preserves the total mass of the system, then total mass is the time t with the rate at which mass is added to the system 1. Applying this gives

$$t = \sum_{k=1}^{\infty} kc_k = \sum_{k=1}^{\infty} k^{1-\tau}. \quad (8.7)$$

When the system is evolving, there will necessarily be some maximum patch size k_{\times} , where there are no patches larger than this size (i.e. $c_k = 0$ for $k > k_{\times}$). This can then be used to approximately solve the above summation using the approximate integral rule to find the leading order exponent.

$$t = \sum_{k=1}^{k_{\times}} k^{1-\tau} \sim k_{\times}^{2-\tau}. \quad (8.8)$$

The cross-over time $t_{\times} = k_{\times}^{2-\tau}$ then gives the leading order time to which the density of patches of size k take to reach its limiting distribution. In order to acquire a rate from this the cross-over time assume that initially there are no patches of size k . The time taken for patches of size k to go from 0 to their asymptotic value is t_x . The asymptotic density of patches of size k is given by the tail solution $k^{1-\tau}$. Therefore the rate of change in the density of patches of size k , r_k can be found approximately as

$$r_k \approx \frac{\Delta c_k}{\Delta t} = \frac{k^{-\tau} - 0}{k^{2-\tau} - 0} = k^{-2} \quad (8.9)$$

The rate of change in the density of patches of size k therefore depends only upon this size of patches and not the power-law tail of the distribution. This analysis

applicable to a pure aggregation system with no loss of mass due to death indicates that there would be no relationship between the Korcak exponent and the return rate in the density. This therefore shows a limitation in studying a vegetation system as a pure aggregation process.

One of the issues with this approach is that there is no direct correspondence between the return rate in terms of the density of a population at or close to equilibrium and a rate of change in density for an aggregation process, which is far from equilibrium. For a system that has a heavier tale (i.e. the Korcak exponent τ is lower/ approaching one) the corresponding spatial pattern would have more large patches and fewer small patches and hence would look persistent from a spatial perspective. This is not the same as the type of persistence when only a single patch size is considered.

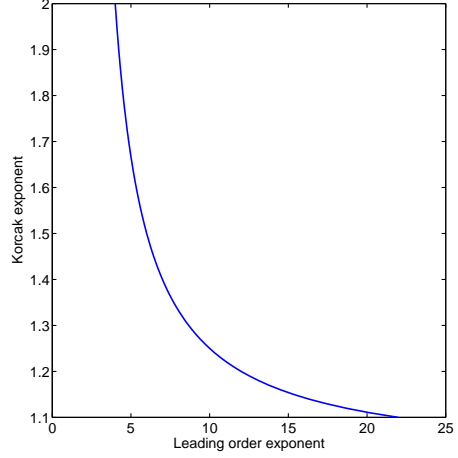


Figure 8.1: Plausible Korcak-return rate relationship.

A simpler system without a monomer injection term has been solved analytically [Von Smoluchowski, 1916]. This solution can then be used to look at how the asymptotic solution decays with respect to the new dynamics. The leading order exponent of these dynamics then sets the rate of the system. For initial data of the form $c_k(0) = Ck^{-\tau}$ and an constant kernel aggregation equation with no input of the form

$$\frac{d}{dt}c_k(t) = \sum_{i+j=k} c_i c_j - 2c_k \sum_{i=1}^{\infty} c_i. \quad (8.10)$$

Then the solution for large k is of the form

$$c_k \approx -\frac{1}{\Gamma(1-\alpha)} t^{-1} (Dt)^{-1/(\tau-1)} w^{-\tau}, \quad w = k/(Dt)^{1/(\tau-1)}. \quad (8.11)$$

The leading order term for t has exponent $2\tau/(\tau-1)$. Plotting this against the Korcak exponent τ gives a good close relationship to the one observed in the data (See Fig. (8.1)).

8.5 Constant aggregation with linear death

In this section we derive a new result for a more realistic model of vegetation aggregation with a death term that is linear in patch size. The strategy for deriving a solution is similar to the strategy in Krapivsky et al. [2010]. For a constant kernel $K(i, j) = 2$, Eq. (8.6) is rewritten as

$$\frac{d}{dt}c_k = \sum_{i+j=k} c_i c_j - 2c_k \sum_{j \geq 1} c_j + \mu c_{k+1} - \mu c_k, \quad (8.12a)$$

$$\frac{d}{dt}c_1 = -2c_1 \sum_{j \geq 1} c_j + 1 + \mu c_2. \quad (8.12b)$$

The asymptotic tail of the resulting patch-size distribution is then sought in order to gain an understanding of how the linear death rate affects the stationary distribution. A moment-generating function is used to find the steady state solution to this equation in a similar fashion to the one described in Krapivsky et al. [2010]. Firstly define the total number of all patches as $N = \sum_{k \geq 1} c_k$ and then sum Eq. (8.12) in order to obtain

$$\frac{dN}{dt} = \sum_{k \geq 1} \sum_{i+j=k} c_i c_j - 2 \sum_{k \geq 1} c_k \sum_{j \geq 1} c_j + \sum_{k \geq 1} \delta_{k,1} + \sum_{k \geq 1} \mu c_{k+1} - \sum_{k \geq 2} \mu c_k, \quad (8.13)$$

$$\frac{dN}{dt} = N^2 - 2N^2 + 1 - \mu c_1, \quad (8.14)$$

$$\frac{dN}{dt} = -N^2 + 1 - \mu c_1. \quad (8.15)$$

Dynamically, consider when N is at equilibrium. If $\mu = 0$ then the stationary solution is $N = 1$. If $\mu > 0$ then the equilibrium solution is necessarily bounded between one and zero as N and c_1 are always positive.

The moment-generating function $C(z, t) = \sum_{k=1}^{\infty} c_k z^k$ is now considered. Multiplying Eq. (8.12) by z^k and summing over all k gives the following

$$\begin{aligned} \frac{d}{dt}C &= C^2 - 2NC + z + \mu \sum_{k \geq 1} z^k c_{k+1} - \mu \sum_{k \geq 2} z^k c_k \\ &= C^2 - 2NC + z + \frac{\mu}{z}C - \mu C - \mu c_1. \end{aligned} \quad (8.16)$$

The new moment generating function defined as $A(z, t) = C(z, t) - N(z, t)$ is considered. The time derivative is calculated by combining Eq. 8.16 with Eq. 8.15

$$\begin{aligned}
\frac{d}{dt}A(z, t) &= \frac{d}{dt}C(z, t) + \frac{d}{dt}N(t) \\
&= C^2 - 2NC + \frac{\mu}{z}C + z - \mu C - \mu c_1 - 1 + N^2 + \mu c_1 \\
&= A^2 + \frac{\mu}{z}C - \mu C + z - 1 \\
&= A^2 + \mu \frac{1-z}{z}A + \mu \frac{1-z}{z}N + z - 1.
\end{aligned} \tag{8.17}$$

Note that the right-hand side is quadratic in terms of A . Setting the time-derivative to zero gives the steady-state solution of the moment-generating function as

$$A = \mu \frac{z-1}{z} + \sqrt{\mu^2 \frac{(1-z)^2}{z^2} - 4 \left(\mu \frac{1-z}{z}N + z - 1 \right)}. \tag{8.18}$$

In order to proceed it is assumed that the death rate μ is small and only the leading order term is kept. Hence

$$A \approx 2\sqrt{1-z - \mu \frac{1-z}{z}N}. \tag{8.19}$$

The strategy is to find A in terms of the power series $\sum_{k=1}^{\infty} c_k(z^k - 1)$. The expansion of $\sqrt{1-x}$ is used to obtain

$$A_{\text{approx}} = 2 \sum_{k=0}^{\infty} \frac{\Gamma(3/2)}{\Gamma(3/2-k)\Gamma(k+1)} (1 + \mu N)^{1/2-k} (-z - \mu N/z)^k. \tag{8.20}$$

Using the relationship $\Gamma(z)\Gamma(1-z) = \frac{\pi}{\Gamma(z)\sin(\pi z)}$, cancelling the $(-1)^k$ terms and absorbing all constants into a constant c term

$$A_{\text{approx}} = c \sum_{k=0}^{\infty} \frac{\Gamma(k-1/2)}{\Gamma(k+1)} (1 + \mu N)^{1/2-k} (z + \mu N/z)^k. \tag{8.21}$$

Using the binomial expansion, this becomes

$$A_{\text{approx}} = c \sum_{k=0}^{\infty} \sum_{i=0}^k \frac{\Gamma(k-1/2)}{\Gamma(k+1)} \frac{\Gamma(k+1)}{\Gamma(i+1)\Gamma(k-i+1)} (1 + \mu N)^{1/2-k} (\mu N)^{k-i} z^{2i-k}. \tag{8.22}$$

In order to find the k -th coefficient as $k \gg 1$ the leading order of the binomial is considered, hence only the terms where $i = k$ are kept in the expansion.

$$A_{\text{approx}} \approx c \sum_{k=0}^{\infty} \frac{\Gamma(k - 1/2)}{\Gamma(k + 1)} (1 + \mu N)^{1/2-k} z^k. \quad (8.23)$$

By using the asymptotic approximation $\Gamma(n+a)/\Gamma(n) \sim n^a$ and assuming k is large, the k -th coefficient in this expansion and hence the density of patches of size k is

$$c_k = k^{-3/2} \exp(-\Lambda k), \quad (8.24)$$

where $\Lambda = \log(1 + \mu N)$. The solution is therefore a power law with an exponential truncation with a factor Γ . When the death rate is 0, $\Lambda = 0$ and hence the patch-size distribution is a pure power law as is expected. A large death rate will lead to a solution that is dominated by an exponential decay term, hence the patch-size distribution is expected to have a smooth transition from a pure power law to an exponential distribution. A dimensionality argument of Eq. 8.12 [Connaughton et al., 2004] also leads to a power law exponent of the form $3/2$. We derived a more general version of Eq. 8.12 by considering a kernel of the form $K(i, j) = i^{-\alpha} + j^{-\alpha}$ and a general linear death rate of the form $\mu(k) = k^\beta$, which leads to the new general aggregation-fragmentation of the form

$$\frac{d}{dt} c_k = \frac{1}{2} \sum_{i+j=k} (i^{-\alpha} + j^{-\alpha}) c_i c_j - \sum_{j \geq 1} (j^{-\alpha} + k^{-\alpha}) c_j c_k + \delta_{k,1} + (k+1)^\beta c_{k+1} - k^\beta c_k. \quad (8.25)$$

A general solution of this equation cannot be found analytically, instead the rate equation is applied to the patch-size distribution data as described in Chapter 3. This analysis is discussed in Section 8.7, but for now there is a deviation into an alternative strategy to derive a solution to Eq. 8.12 when the continuous patch-size limit is taken.

8.6 Comparison to explicit spatial modelling of vegetation growth

In order to compare the model predictions of patch formation in an aggregation system with a constant death rate the prediction of the patch-size distribution obtained in Eq. 8.24 is compared to a simple probabilistic cellular automata model of vegetation growth. The cellular model is in a similar vein to the one discussed in Scanlon et al. [2007b], the model is defined on a toroidal lattice where each site

can exist in one of two states: alive (1) and dead (0). The alive state propagates through nearest neighbour growth at rate β , as well as through a background constant birth probability γ . The alive sites transition to a dead site with a constant death probability μ . Hence if n_x is the number of alive sites in the neighbourhood of site x , the transitions can be summarised as

$$P_x(0 \rightarrow 1) = \min\{1, \gamma + \beta n_x/4\}, \quad (8.26a)$$

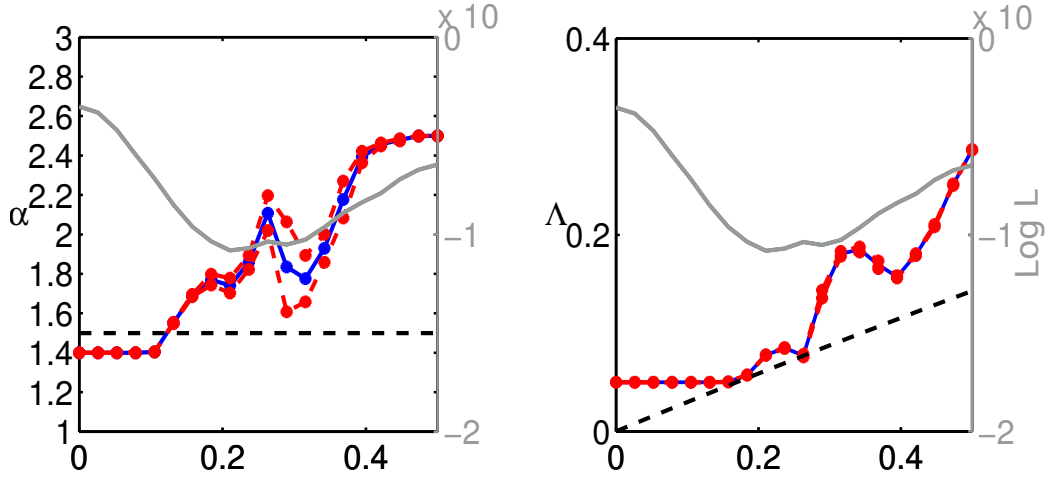
$$P_x(1 \rightarrow 0) = \mu. \quad (8.26b)$$

where the minimum function is used to guarantee the probability of transitioning to an alive state is one in the rare case when the sum of the two probabilities increases above one.

Simulations were conducted for constant aggregation and birth rate $\beta = 0.2, \gamma = 0.01$ and over a range of death rates. Simulations were ran for 600 time-steps and for lattice length $L = 500$. The final patch-size distribution was recorded for each simulation run and the following power-law with exponential truncation was fitted to the distribution using a maximum likelihood method

$$f(K = k) = Ck^{-\alpha} \exp(-\Lambda k), \quad (8.27)$$

for some normalising factor C . The resulting maximum likelihood estimators were found using a downhill simplex method implemented in Matlab R2014a [Lagarias et al., 1998]. The approximate solution to the aggregation equation predicts a constant power-law exponent α of $3/2$. This is close to the inferred value from simulation when the death rate is low (Fig. 8.2a), however for $\mu > 0.3$ the power-law exponent deviates from the theoretical value. The exponential factor Λ is non-zero when the death rate is zero (Fig. 8.2b), where the mean field solution predicts a zero exponential term. This deviation can be explained due to a finite-size effect, where the finite system size induces an exponential tail in the patch-size distribution. For increasing death rate, Λ does scale with the mean field prediction, although there are large deviations from this. Overall there is an increase in the exponential factor for increasing death rate as is predicted, however the functional form of the increase is not captured by the mean field approximation.



(a) comparison of power-law exponent from simulation to theoretical prediction (b) comparison of exponential factor from simulation to theoretical prediction

Figure 8.2: Exponents of patch-size distribution compared to simulations. The theoretical values for the power-law exponent α and the exponential factor Λ are derived in Eq. 8.24. As predicted for small values of the death rate the power-law component of the patch-size distribution is constant whilst there is an increase in the exponential component for increasing death rate. The likelihood of the fit (in grey) indicates that for intermediate values of μ there is a poor fit of the distribution to the simulation data

8.7 Empirical analysis of patch-size distribution using the aggregation-fragmentation equation

Comparing the Korcak exponent with empirical data is one way of relating empirical data to the general aggregation equation. However, in general the exact form of the aggregation equation is unknown. Furthermore, the type and effects of the fragmentation process can also be unknown and it is desirable to tease out these aspects from the data. Work has been done on the inverse problem, whereby an aggregation kernel is inferred from an empirical distribution. Recent work Jones et al. [2013], uses a least-squares method to fit a general kernel of the form $K(i, j) = \frac{q_0}{2} h(x)(i^\alpha j^\beta + j^\alpha i^\beta)$, where the exponents give the small and large asymptotic of the aggregation and $h(x)$ is a term of $O(1)$, where $x = i/j$. In order to preserve symmetry there is also the requirement that $h(x) = h(x^{-1})$.

Here I develop Bayesian approach in order to infer the parameters of the

aggregation-fragmentation equation

$$\frac{d}{dt}c_k = g \sum_{i+j=k} (i^\alpha j^\beta + j^\alpha i^\beta) c_i c_j - g \sum_{i=1}^N (i^\alpha k^\beta + k^\alpha i^\beta) c_i c_k + \mu(k+1)c_{k+1} - \mu(k)c_k + \tau \delta_{k,1}, \quad (8.28)$$

for the empirical patch-size distributions of the five sites from the Isles of Scilly, UK dataset. Model comparison for the different patches and different forms of the aggregation-fragmentation equation were then performed. This is done by comparing whether the fragmentation term has a linear death rate μ or one that is dependent on the current patch size $\mu(i)$. Three models of aggregation with a death rate linear in patch-size are considered. The first model assumes death is a constant for all patch sizes $\mu(i) = \mu$, the second assumes a square root form of the death rate $\mu(i) = \mu i^{1/2}$ and the third assumes a linear form of the death rate $\mu(i) = \mu i$. Finally, the inferred parameters are compared to the resulting dynamics of the system, in particular the relaxation time, as a way of understanding the persistence.

8.7.1 Developing the likelihood

The general patch-evolution equation for a patch-size distribution $C = (c_1, \dots, c_i, \dots, c_N)$ will evolve according to some functional form f_i representing aggregation and fragmentation (and also possibly birth and death) for each of the patch sizes. It is assumed there is some additive noise in the form of an i.i.d. random variable ϵ_t that is distributed as $\epsilon_t \sim N(0, \sigma^2)$

$$\frac{d}{dt}c_i = f_i(C) + \epsilon_t. \quad (8.29)$$

At stationarity the resulting likelihood has the form $L = \prod_{i=1}^N N(f_i(C), \sigma^2)$.

8.7.2 Results

The MCMC chain was run for all three models and all five sites. A burn-in time on 10^6 was found to be sufficient for convergence and the chain show good mixing properties. The converged chain was run for a further 10^6 time steps and thinning was performed on the resulting sample to produce uncorrelated samples from the posterior. These samples were then used to construct the posterior and infer the five parameters for each of the models.

As there are a total of fifteen separate posteriors constructed using the MCMC method the full results of each posterior are not shown. Instead the mean of the

marginal posteriors are given for each parameter as well as the 95% confidence intervals for each parameter are displayed in Table 8.2 and the relative log Likelihoods for each model is given in Table 8.1. The form of the aggregation parameters α and β are close to 0 across all sites and for all three models although the confidence intervals for each are large compared with the size of the support and hence is not being strongly selected for. The values of the aggregation exponent fit well with the estimated power-law exponent values for the seagrass meadows, where all values are close to 1.5 although there is variation between them. The marginal probabilities indicate that $\alpha, \beta < 0$ is more favoured than for $\alpha, \beta > 0$. Hypothesis testing for the three models was conducted using a Bayes factor for all of the sites where the relative log Bayes factor was calculated for site j from patch data D_j and model M_i as

$$B = \log \left(\frac{P(D_j|M_i)}{P(D_j|M_{\min})} \right), \quad (8.30)$$

where M_{\min} is the model with the lowest maximum a posteriori relative to the other models for that site. For all five meadows the constant death rate model $\mu(i) = \mu$ scores the highest relative log Bayes factor. Taking model one as the most likely, the aggregation exponents for all sites are within the same 95% confidence intervals, however site ogh has noticeably more negative aggregation exponents than the other sites. All other sites have similar parameters and overlapping confidence intervals.

	BLT	OGH	WBL	HTB	LA
Model 1	10.4181	9.1689	7.3466	10.424	7.4235
Model 2	7.89	0.83567	4.7453	7.3647	5.7792
Model 3	0	0	0	0	0

Table 8.1: Bayes factor comparison for all three models over the five sites. The relative Bayes factor log ratio was calculated by taking the maximum a posteriori (MAP) for each model and site respectively and dividing through by the lowest MAP. The Bayes factor indicates that the first model, where death is constant across all patches, has the strongest evidence.

8.8 Conclusion

Changing the focus away from explicit spatial modelling of vegetation patch formation and instead focusing on the dynamics of patch-size themselves gives a unique insight into the underlying aggregation-fragmentation processes. This chapter has primarily focused on solutions to equations where the aggregation kernel that gov-

erns that rate at which patches of two sizes will aggregate by either a constant or power law kernel. For a system where there is aggregation only the resulting patch-size distribution is that of a pure power law, with exponent that is dependent on the exponent of the power law aggregation kernel. The introduction of a linear death term, where an individual is lost from a patch at rate μ gives rise to a power law with exponential tail distribution of the form $c_k \sim k^{-\alpha} \exp(-\Lambda k)$. This solution holds generally when there is a linear death term and power-law aggregation kernel, even when the kernel is composed of a sum of two power-laws. Further, α is dependent on the specifics of the aggregation term alone and Λ is dependent on the death rate alone. This separation of the aggregation and fragmentation term implies, in principle, the ability to infer aggregation and death processes through observing the converged patch-size distribution alone, hence this is applicable to inferring process from a single spatial snapshot.

Scanlon et al. [2007a] predicts that there is a meltdown of a power-law distribution near a critical point in the system dynamics for a patch-size distribution of vegetation. The model used is a spatially-explicit one with a local growth term and a background death rate. The aggregation with constant death rate is one that is analogous as a mean field model to the one proposed. Hence through the derived solution in this chapter it is observed that there should always be an exponential tail to the distribution if the death rate is non-zero. Similar arguments have also been made recently [Pueyo, 2011], but notably none have explained the origin of a power law with exponential tail observed in vegetation systems. The derived model then, provides a theoretical origin to the observed spatial patterns in vegetation ecosystems that are under pressure that can be considered constant throughout space (rainfall, grazing etc.). This approach would be able to provide further insights into the nature of the patch-size distribution for other systems where disturbance may be spatially distributed.

The model also gives insight into how there can be a continuous array of power-law exponent observed in nature. The aggregation with no death model predicts that power-laws exponents in the range $(1, 2)$ are physically possible, which is what has been observed in the Isles of Scilly data as well as in semi-arid systems and mussel beds. The model therefore predicts that a change in the exponent of a patch-size distribution is related to a change in the structure of the aggregation kernel. A simple dimensionality argument can be used to show that in the aggregation and death model with a kernel that has a general power law scaling, the resulting stationary

distribution will have the same exponent as that in the model with no death. The drawback to this approach in our context is that there is no explicit way of defining the return rate of the system, although the relaxation time for a characteristic patch-size provides an ad-hoc way of assessing how the system would evolve following a perturbation. The conclusion of how to relate the patch-size distribution to the system dynamics is that both the power-law exponent and the presence of an exponential cut-off does give an indication of the underlying dynamics. More complex fragmentation processes than the one discussed would alter these conclusions however, as a non-linear fragmentation process will also lead to self-similar solutions and thus the two processes are confounded when only the stationary state is observed [Ernst and Van Dongen, 1987], such processes include storms and other strong weather events that could split a single cluster of vegetation into multiple clusters.

	BLT	OGH	WBL	HTB	LA
Model 1	α	(-0.91,-0.06,0.85)	(-0.93,-0.21,0.66)	(-0.91,-0.05,0.85)	(-0.91,-0.08,0.86)
	β	(-0.91,-0.05,0.87)	(-0.92,-0.15,0.77)	(-0.9,-0.05,0.86)	(-0.9,-0.03,0.88)
	μ	(0.33,0.8,1.22)	(0.36,0.74,1.11)	(0.32,0.79,1.22)	(0.41,0.83,1.21)
	τ	(23.43,36.61,51.62)	(33.24,43.15,55.44)	(23.61,35.46,78)	(28.95,44.76,62.47)
	g	(0.01,0.07,0.18)	(0.01,0.07,0.16)	(0.01,0.07,0.17)	(0.01,0.08,0.18)
Model 2	α	(-0.92,-0.11,0.82)	(-0.93,-0.28,0.51)	(-0.91,-0.1,0.81)	(-0.92,-0.13,0.81)
	β	(-0.92,-0.09,0.85)	(-0.92,-0.17,0.7)	(-0.91,-0.07,0.86)	(-0.92,-0.07,0.86)
	μ	(0.34,0.71,1.05)	(0.01,0.12,0.25)	(0.12,0.57,1.02)	(0.33,0.66,0.99)
	τ	(38.07,46.09,55.17)	(18.59,27.53,37.38)	(15.2,26.27,40.33)	(39.33,49.62,60.72)
	g	(0.01,0.07,0.17)	(0.03,0.14,0.26)	(0.01,0.1,0.23)	(0.01,0.09,0.21)
Model 3	α	(-0.92,-0.13,0.76)	(-0.94,-0.34,0.46)	(-0.91,-0.11,0.79)	(-0.92,-0.17,0.74)
	β	(-0.91,-0.09,0.83)	(-0.93,-0.17,0.68)	(-0.91,-0.11,0.82)	(-0.92,-0.08,0.84)
	μ	(0.01,0.08,0.18)	(0.01,0.03)	(0.01,0.08,0.17)	(0.01,0.11,0.23)
	τ	(26.99,38.09,49.17)	(14.93,21.08,28.55)	(15.05,25.69,34.9)	(23.59,33.45,43.71)
	g	(0.05,0.15,0.25)	(0.03,0.15,0.28)	(0.04,0.15,0.27)	(0.05,0.15,0.26)

Table 8.2: Inferred parameters and 95% confidence intervals for all three models and five sites, parameter estimates are taken as the mean of the posterior marginalised for each of the parameters. α, β are the exponents of the aggregation kernel, μ is the death rate, τ is the rate of births of size one and g is the rate of aggregation.

Bibliography

- David J Aldous. Deterministic and stochastic models for coalescence (aggregation and coagulation): a review of the mean-field theory for probabilists. *Bernoulli*, pages 3–48, 1999.
- Mark D Bertness and Edwin Grosholz. Population dynamics of the ribbed mussel, *geukensia demissa*: the costs and benefits of an aggregated distribution. *Oecologia*, 67(2):192–204, 1985.
- Colm Connaughton, R Rajesh, and Oleg Zaboronski. Stationary kolmogorov solutions of the smoluchowski aggregation equation with a source term. *Physical Review E*, 69(6):061114, 2004.
- MH Ernst and PGJ Van Dongen. Scaling laws in aggregation: Fragmentation models with detailed balance. *Physical Review A*, 36(1):435, 1987.
- Frédéric Guichard, Patti M Halpin, Gary W Allison, Jane Lubchenco, and Bruce A Menge. Mussel disturbance dynamics: signatures of oceanographic forcing from local interactions. *The American Naturalist*, 161(6):889–904, 2003.
- Hisao Hayakawa. Irreversible kinetic coagulations in the presence of a source. *Journal of Physics A: Mathematical and General*, 20(12):L801, 1987.
- George A Jackson. A model of the formation of marine algal flocs by physical coagulation processes. *Deep Sea Research Part A. Oceanographic Research Papers*, 37(8):1197–1211, 1990.
- Peter P Jones, Robin C Ball, and Colm Connaughton. Nonlinear least-squares method for the inverse droplet coagulation problem. *Physical Review E*, 88(1):012138, 2013.
- Sonia Kéfi, Max Rietkerk, Concepción L Alados, Yolanda Pueyo, Vasilios P Papanastasis, Ahmed ElAich, and Peter C De Ruiter. Spatial vegetation patterns and

imminent desertification in mediterranean arid ecosystems. *Nature*, 449(7159): 213–217, 2007.

P. L. Krapivsky, J. F. F. Mendes, and S. Redner. Influence of island diffusion on submonolayer epitaxial growth. *physical review B*, 59:15950–15958, June 1999.

Pavel L Krapivsky, Sidney Redner, and Eli Ben-Naim. *A kinetic view of statistical physics*. Cambridge University Press, 2010.

Jeffrey C Lagarias, James A Reeds, Margaret H Wright, and Paul E Wright. Convergence properties of the nelder–mead simplex method in low dimensions. *SIAM Journal on optimization*, 9(1):112–147, 1998.

Hiro-Sato Niwa. School size statistics of fish. *Journal of theoretical biology*, 195(3): 351–361, 1998.

Salvador Pueyo. Desertification and power laws. *Landscape ecology*, 26(3):305–309, 2011.

Todd M Scanlon, Kelly K Caylor, Simon A Levin, and Ignacio Rodriguez-Iturbe. Positive feedbacks promote power-law clustering of kalahari vegetation. *Nature*, 449(7159):209–212, 2007a.

Todd M Scanlon, Kelly K Caylor, Simon A Levin, and Ignacio Rodriguez-Iturbe. Positive feedbacks promote power-law clustering of kalahari vegetation. *Nature*, 449(7159):209–212, 2007b.

Philip A Stephens and William J Sutherland. Consequences of the allee effect for behaviour, ecology and conservation. *Trends in Ecology & Evolution*, 14(10): 401–405, 1999.

Marian Von Smoluchowski. Drei vortrage uber diffusion. brownische bewegung und koagulation von kolloidteilchen. *Z. Phys.*, 17:557–585, 1916.

EJ Weerman, J Van Belzen, M Rietkerk, S Temmerman, S Kéfi, PMJ Herman, and J Van de Koppel. Changes in diatom patch-size distribution and degradation in a spatially self-organized intertidal mudflat ecosystem. *Ecology*, 93(3):608–618, 2012.

Chapter 9

Vegetation disease on spatial pattern formation

The sun comes up just about as often as it goes down, in the long run,
but this doesn't make its motion random

(Donald Knuth)

9.1 Introduction

The work outlined in the previous chapters has concentrated on the formation of spatial order in sessile communities due to the short and long-range interactions between the organism and its environment. This has included how spatial pattern can be informative about the underlying dynamics and how the distribution of patch sizes can arise due to an underlying aggregation process. This chapter concerns the effects of processes on these resulting patterns. Spatial pattern is an important factor in both species and disease spread [Fridley et al., 2007; Real and McElhany, 1996]. On short time-scales, biological invasion is dependent on the connectivity of appropriate habitat, whereas on longer time-scales there is an interaction between the invading species and its host, leading to dynamics which are not separable. There are therefore two regimes of study, where biological invasion is fast on the time-scale of the host dynamics and where the time-scales are not separated and host-species dynamics must also be explicitly taken into account.

Further consideration can be made in the context of a host-parasite system, where either the host or parasite is evolving. An example is in how a disease responds to the fragmentation of its host, whereby the dominant trait of the disease varies depending on the spatial properties of the host. Prudence is a term used to describe a parasite or disease that evolves in a way that it limits its own spread, the consequence of which is that the host population has time to recover and the parasite remains endemic in the population [Lion and Boots, 2010]. Local clustering can lead to parasites or disease with low virulence, while a co-evolution of both parasite and host leads to hosts with high resistance as well as parasites with low virulence [Best et al., 2011].

Once again the case study of seagrass is used as an example of a vegetation system where there are spatial processes occurring on the system. Seagrass is primarily a clonal spreading vegetation where its entire life-cycle is sub-tidal. A wasting disease, caused by a slime-mould protist *L. zosterae* [Muehlstein et al., 1991], led to a large Atlantic-wide epidemic that severely reduced its numbers in the 1930s [Muehlstein, 1989; Orth et al., 2006]. Currently wasting disease is endemic in most Atlantic seagrass communities and there are several open questions related to the epidemic dynamics [Bull et al., 2012]. Primarily, it is not clear how the same disease could cause such a catastrophic decline in seagrass numbers previously, yet remains at endemic levels more recently. This question provides the motivation for understanding how the geometry of a vegetation distribution affects its ability to regulate disease

spread.

The chapter begins with an investigation of how a biological species spreads when the environment (i.e. the vegetation population) can be considered static. In particular the role of heterogeneity and a fractal structure in the distribution of vegetation shall be investigated to elucidate the relationship between environmental heterogeneity and species diffusion. The second half of the chapter shall be concerned with the impact of spatial pattern on the dynamics of a pathogen, where the rate of the dynamics of the pathogen and the rate of the dynamics of the vegetation are comparable. In particular, the focus is on the role of banding in vegetation in limiting the spread and impact of a virulent disease. The final section of the chapter focuses on the origins of banding in the presence of a pathogen from an evolutionary viewpoint. The main hypothesis of this section is whether banding in vegetation can be viewed as an evolved trait in the sessile host species in the presence of a disease. In which case, under what conditions would such an evolved trait be expected to rise and how generally does banding impact the spread and distribution of the disease.

9.1.1 Modelling the spread of a species in a disordered environment

Normal diffusion is characterised by Fick's law where the mean squared displacement of a diffusing species scales as the square of the time i.e. if $X(t)$ is a random variable representing the position of a diffusing particle then Fick's law states that

$$E[X^2(t)] \sim t^2. \quad (9.1)$$

Anomalous diffusion in contrast is where correlation functions disobey Fick's law and has been well studied in the past 30 years [Andow et al., 1990; Bouchaud and Georges, 1990; Cohen and Murray, 1981]. To begin with, how the probability distribution of a concentration evolves in time in a disordered media was studied, where there is some intrinsic noise that is static on the time-scale of the diffusion process. In order to understand the diffusion properties and how the rate of growth is affected by a disordered environment a model of long-range diffusion with logistic growth was developed. This is achieved using an integro- reaction-diffusion equation, where the diffusion term is replaced with a convolution of a Gaussian. For a concentration of some species u , the evolution equation is

$$\frac{\partial u}{\partial t} = r(\mathbf{x})u(1 - u) + D \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} u(x - p, y - q)k(p, q) - u(x, y)k(p, q)dqdp. \quad (9.2)$$

It is assumed that the spread of a biological species is Gaussian with some constant variance σ , hence $k(x, y) \propto \exp(-\frac{1}{2\sigma^2}(x^2 + y^2))$. Note that if $\sigma \rightarrow 0$ and r is a constant then Eq. 9.2 reduces to the Fisher's equation i.e.

$$\frac{\partial u}{\partial t} = ru(1 - u) + D\nabla^2 u, \quad (9.3)$$

of which travelling wave fronts $u(z)$ are a solution for wave speed $c > 2$ with $z = x - ct$. It is clear that if $r(\mathbf{x})$ is not homogeneous across space then this will alter the wave properties of the equation. It would be expected that in the case where space is discretised and $r(\mathbf{x})$ is an i.i.d. Bernoulli random trial across space then a retardation of the growth velocity is expected. In one spatial dimension, the probability of having a gap of size k where the growth rate is zero would be $\mathbb{P}(\text{gap } k) = (1 - p)^k p^2$. In this gap r is zero and hence the governing evolution is a pure diffusion process. Smaller p implies that the probability of larger gaps exist, until in the limit as $p \rightarrow 0$, where there is a single gap of infinite size and the evolution is governed by a pure diffusion process.

9.1.2 The Multiplicative cascade model

The focus of the generalised Fisher model is to understand how different environmental properties, such as environmental heterogeneity affects the diffusion and general fecundity of a biological species being introduced to an environment. In order to explore the consequences of environmental heterogeneity we desire a neutral model of the environment where properties such as density can be controlled and others such as heterogeneity can be varied.

The Multiplicative cascade model [Meakin, 1987; Meneveau and Sreenivasan, 1987] has been used to study a wide-range of spatial heterogeneities in physical sys-

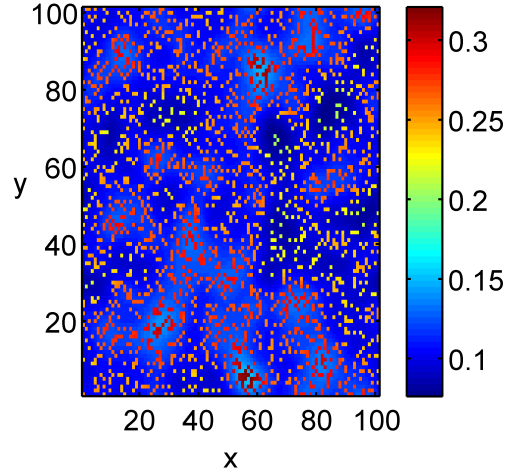


Figure 9.1: An equilibrium distribution for a disordered $r(\mathbf{x})$ drawn from a Poisson point process with probability $\rho = 0.3$. The pattern has long-range structure induced by areas where the density of occupied lattices is higher than the average, this leads to a diffusion pattern that is highly heterogeneous even when $r(\mathbf{x})$ is drawn from a homogeneous distribution.

tems such as rainfall [Schertzer and Lovejoy, 1987] and in the distribution of galaxies [Martínez et al., 1993]. Its construction is relatively straightforward: begin with a discrete lattice of a certain size $(N \times N)$ and four parameters $q_1, q_2, q_3, q_4 \in [0, 1]$, where each q is a general quantity used to as a measure for how the mass of a system is distributed through the spatial hierarchies. The lattice is split into four equal quadrants and a q is randomly assigned to each without replacement. Each quadrant is then split into four further quadrants and the the process continues down to the m -th level, where $m = \log_2 N$. A probability is assigned to a cell i by multiplying together all the preceding probabilities in the cell's hierarchy. Hence, the probability at cell i is

$$p_i = \frac{\prod_{j=1}^{j=m} q_{i_j}}{(q_1 + q_2 + q_3 + q_4)^m}, \quad (9.4)$$

where i_j indexes the sequence of q that were chosen at each hierarchy that i is contained in. If the difference between the q_i is large then this leads to a highly heterogeneous measure over the lattice. Correspondingly, if $q_1 = q_2 = q_3 = q_4$ then each p_i has the same value and hence the measure is completely homogeneous.

A measure can be defined on the lattice by taking a set C that is a sub-set of the lattice with c_i lattice sites and defining the measure as $\mu(C) = \prod_{i=0}^n p_{c_i}$. This gives the total probability of an arbitrary point landing in C . Notice this measure is a complete description of the model and the model could have easily been defined with μ as a starting point. The properties of μ can be analysed by considering the generalised moments

$$M_r(q) = \sum_{\mathcal{M}_r} \mu(C)^q, \quad (9.5)$$

where \mathcal{M}_r is the set of r -mesh squares C for which $\mu(C) > 0$. As the model is defined constructively by splitting the lattice into four quadrants at each point, we can consider how the generalised moments change as the r -mesh squares halve, i.e. take r to be equal to $N2^{-n}$ for increasing values of n up to m . For the n -th level, denote l_i as the number of times the parameter p_i has been picked. This leads to the relationship $l_1 + l_2 + l_3 + l_4 = n$; furthermore, the number of ways this sequence can be chosen is $\frac{n!}{l_1!l_2!l_3!l_4!}$. As such we may now write down an explicit formula for the moments

$$M_{N2^{-n}}(q) = \sum_{l_1+l_2+l_3+l_4=n} \frac{n!}{l_1!l_2!l_3!l_4!} p_1^{q l_1} p_2^{q l_2} p_3^{q l_3} p_4^{q l_4} = (p_1^q + p_2^q + p_3^q + p_4^q)^n. \quad (9.6)$$

It is clear from this equation that the moments scale as a power law for decreasing r . Define the exponent $\beta(q)$ as the following limit

$$\beta(q) = \lim_{r \rightarrow 0} \frac{\log M_r(q)}{\log r}. \quad (9.7)$$

Hence, for the Multiplicative cascade model $\beta(q) = \log(p_1^q + p_2^q + p_3^q + p_4^q)/\log(2)$. $\beta(q)$ contains an entire description of the multifractality of the measure μ . The Rényi dimensions, which are also used to characterise the multifractal nature of a measure through box-counting can be related by the fomula $D_q = \beta(q)/(1 - q)$ [Falconer, 2013]. D_0 is the standard box-counting dimension, when $p_1, p_2, p_3, p_4 > 0$ $D_0 = \log(4)/\log(2) = 2$. Hence the box-counting dimension of the multiplicative cascade model as has been defined here will almost always be 2. This may seem surprising considering these can include some quite heterogeneous situations where parts of the lattice are very sparse, however the measure will always be strictly positive ($\mu > 0$) everywhere and as such the cascade will always have non-empty boxes of arbitrary size assuming an infinite lattice. For a finite lattice finite size effects come into play, which would lower the box-counting dimension to a value less than 2. If $p_i = 0$ for a single $i \in \{1, 2, 3, 4\}$ then the box-counting dimension $D_0 = \log(3)/\log(2) \approx 1.59$. This is due to parts of the measure being 0 at every box size r . Similarly if $p_i = 0$ for two $i \in \{1, 2, 3, 4\}$ then $D_0 = \log(2)/\log(2) = 1$. The measure's support is therefore geometrically similar to a line.

For purposes here it is more elucidating to consider the multifractal spectrum, which gives a measure of the heterogeneity of the measure [Stanley and Meakin, 1988]. We can relate this quantity to the coarse-grained multifractal spectrum $f(\alpha)$ via a Legendre transform

$$f(\alpha) = \inf_{-\infty < q < \infty} \{\beta(q) + \alpha q\}. \quad (9.8)$$

As $\beta(q)$ is continuously differentiable we can calculate $f(\alpha)$ by differentiating the term inside the infimum of the Legendre transform and setting to zero. This gives the form of α as

$$\alpha = -\frac{p_1^q \log p_1 + p_2^q \log p_2 + p_3^q \log p_3 + p_4^q \log p_4}{(p_1^q + p_2^q + p_3^q + p_4^q) \log 2}. \quad (9.9)$$

Inputting this into the definition of $f(\alpha)$ we derive the parametric form of the multifractal spectrum

$$f(\alpha) = \frac{\log(p_1^q + p_2^q + p_3^q + p_4^q)}{\log 2} - \frac{q(p_1^q \log p_1 + p_2^q \log p_2 + p_3^q \log p_3 + p_4^q \log p_4)}{(p_1^q + p_2^q + p_3^q + p_4^q) \log 2}. \quad (9.10)$$

Observe that if $p_1 = p_2 = p_3 = p_4$ then $f(\alpha) = c$ and $\alpha = c$ for all $-\infty < q < \infty$. As the p_i deviate from each other α occupys are larger range implying an increasing heterogeneity of the measure μ . We can characterise this heterogeneity by calculating $-1/|f''(\alpha)|$ around its maximal point (when $q = 0$).

$$\begin{aligned} f'(\alpha) &= \frac{df(\alpha(q))}{dq} \frac{dq}{d\alpha(q)} \\ &= \frac{-\alpha(q) - q\alpha'(q) + \alpha(q)}{\alpha'(q)} = -q, \\ \implies f''(\alpha) &= \frac{df'(\alpha(q))}{dq} \frac{dq}{d\alpha(q)} \\ &= -\frac{1}{\alpha'(q)} \\ &= \frac{(\sum_i p_i^q)^2 \log 2}{(\sum_i p_i^q)(\sum_i p_i^q (\log p_i)^2) - (\sum_i p_i^q \log p_i)^2}. \end{aligned}$$

The heterogeneity index is thus

$$-1/|f''(\alpha(0))| = \frac{4 \sum_i (\log p_i)^2 - (\sum_i \log p_i)^2}{16 \log 2}. \quad (9.11)$$

Note that when $p_1 = p_2 = p_3 = p_4$ we should have $-1/|f''(\alpha(0))| = 0$ as the measure is entirely monofractal. As an example take $p_1 = p_2 = p$ and $p_3 = p_4 = q$. The heterogeneity index in this case is

$$-1/|f''(\alpha(0))| = \frac{(\log p - \log q)^2}{\log 16}, \quad (9.12)$$

we can see that in the case where $p = q$ then $-1/|f''(\alpha(0))| = 0$, as we would expect since the measure is entirely monofractal and the multifractal spectrum collapses on a single point with width 0. As the probabilities are normalised $p + q = \frac{1}{2}$ and as such as $p \rightarrow 0$ then $q \rightarrow \frac{1}{2}$. As $p \rightarrow 0$, $-1/|f''(\alpha(0))| \rightarrow \infty$, hence the larger the difference between the probabilities p and q the more heterogeneous the measure and the corresponding realisations of the model will be.

9.1.3 Methods

The simulations were initialised by drawing $r(\mathbf{x})$ from a multiplicative cascade model with parameters $(q_1, q_2, q_3, q_4) = (1, 1, p, q)$ where p and q are in the range $0 < p, q \leq 1$. The first step is to normalise the parameters, $p_i \rightarrow q_i / \sum q_i$. $r(\mathbf{x})$ is then generated by producing the matrix P where each row is a random permutation of (p_1, p_2, p_3, p_4) . A fixed number of points is selected and for each a row vector is produced with indexes randomly drawn based on the probabilities of each column. For example if the number of hierarchies is 3 then a randomly drawn matrix P would have the form

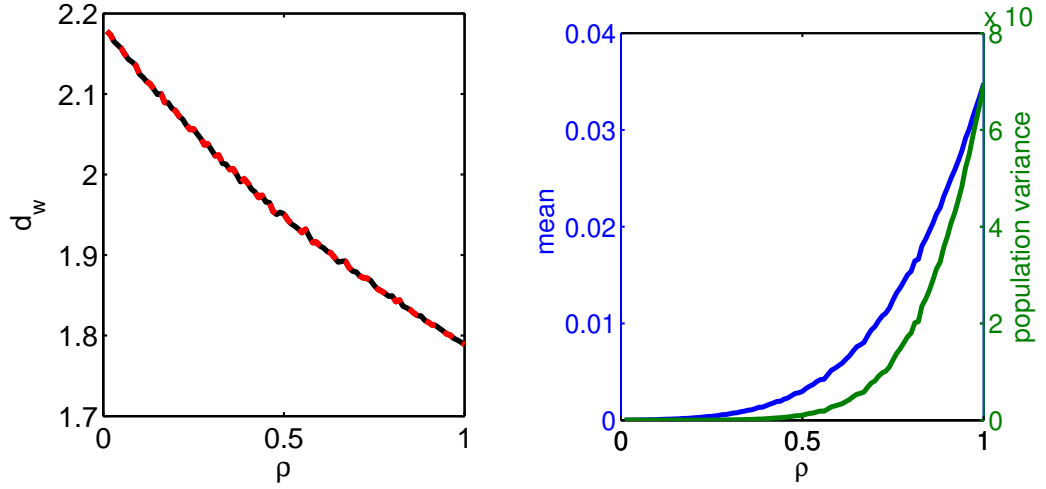
$$\begin{pmatrix} p_2 & p_4 & p_4 \\ p_1 & p_2 & p_3 \\ p_3 & p_1 & p_2 \\ p_4 & p_3 & p_1 \end{pmatrix}.$$

A random vector based on this matrix could then be $(1, 2, 2)$ (based on the relative probabilities). This vector defines the co-ordinates for this point. In this case the point is in the second quadrant of the second quadrant of the first quadrant of the lattice. There is the possibility that due to the hierarchy being terminated for some finite m that some points may overlap.

Once $r(\mathbf{x})$ has been drawn, the integro-differential equation described in Eq. 9.2 was solved numerically using a Runge-Kutta 4th order method in the interval $[0, T]$ with an initial condition u_0 , where $u_0(x) = 1$ when x is the central co-ordinate and 0 everywhere else.

9.1.4 Observables

Certain observables are taken from the numerical solution to Eq. 9.2 to ascertain how the multifractal features of $r(\mathbf{x})$ affect the dynamic properties of the species spread. The mean and variance of the population over space was recorded for the final time point T . Additionally, as we are interested in how $r(\mathbf{x})$ alters the diffusion properties we record the exponent of the average displacement for early time $\langle (x(t) - x_0)^2 \rangle \sim t^{2/d_w}$. This was done by weighting the squared distances from the centre of the lattice by the density and then averaging at each time point. An exponent was then fitted to these quantities on a log-log plot using linear regression. This produces an estimate for the dynamic exponent d_w , which measures the deviation from Fickian diffusion.



(a) d_w varying in an approximately linear relationship with ρ . (b) mean and variance of species density at time T .

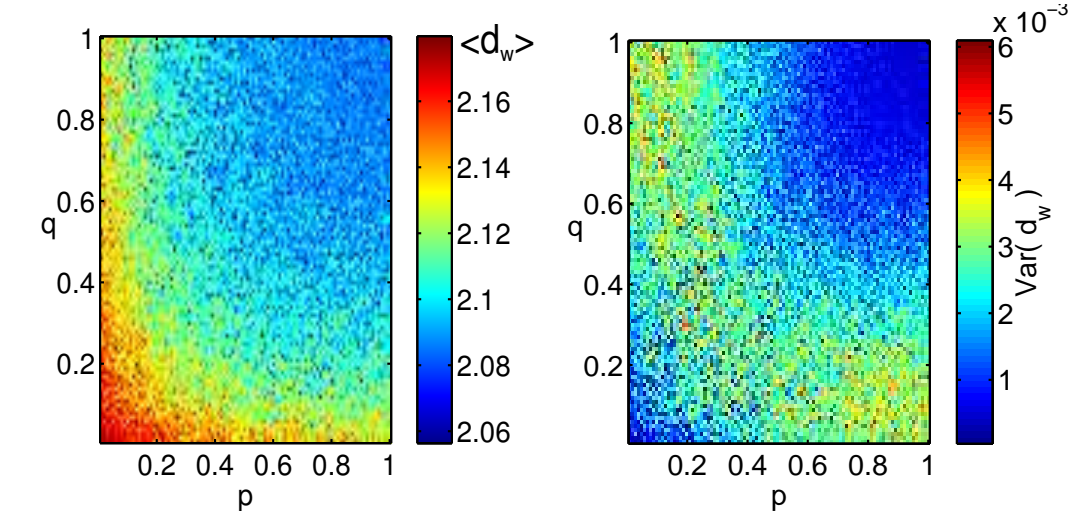
Figure 9.2: Summary of $r(\mathbf{x})$ for a Poisson point process with density ρ . Fig. 9.2b gives the relationship between d_w as the probability of lattice site occupation parameter ρ increases. There is an approximately linear relationship as diffusion moves from the sub-diffusive regime ($\rho < 0.4$) into the super-diffusive regime ($\rho > 0.5$). There is a non-linear relationship between the mean and variance of the population against the density parameter ρ .

9.1.5 Results

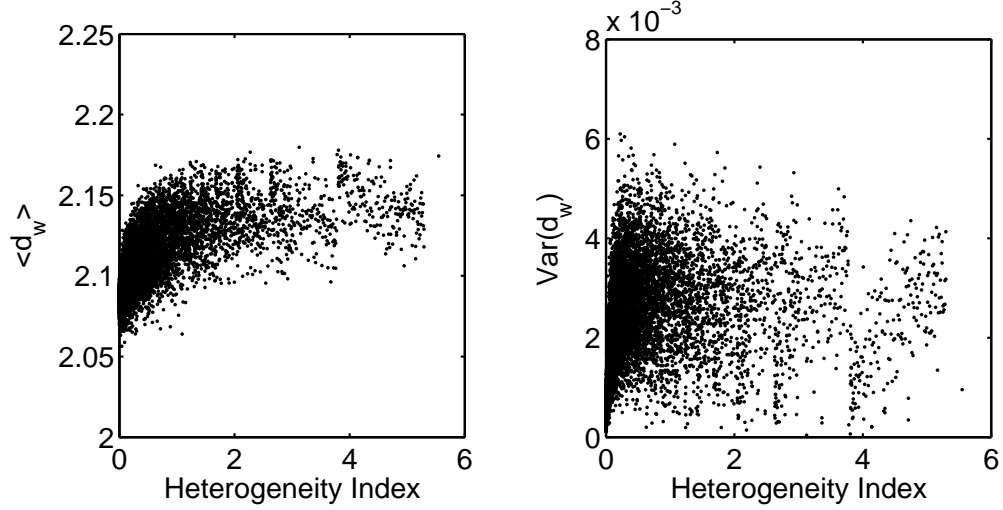
As a model of comparison, Eq. 9.2 was numerically solved for a $r(\mathbf{x})$ drawn from a Poisson point process with a cut-off at the smallest lattice size. This produces a homogeneous, but patchy landscape that can be controlled using the density parameter ρ that is the probability of a randomly chosen lattice site being occupied. d_w has an approximately linear relationship with ρ (Fig. 9.2a) i.e. as occupancy increases the rate of diffusion also increases.

Fig. 9.2b shows the mean population after some time T . The equilibrium density would be expected to scale as a linear function of ρ . The relationship is in fact super-linear due to the population at time T still being in a transient state and hence the super-linear growth is due to the propagation of the species wave-front. This is also reflected in the relationship between population variance and probability of occupancy ρ . The increase in variance is due to the super-diffusive transient dynamics, where there is a high degree of variability along the wave front.

A summary of when the growth rate $r(\mathbf{x})$ is drawn from the multiplicative cascade model with parameters $(1, 1, p, q)$ is given in Fig. 9.3. Each point (p, q) is averaged over 100 realisations of the multiplicative cascade model for those parameters. d_w is in the sub-diffusive regime when p and q are low, which corresponds to a high heterogeneity index (as indicated in Fig. 9.3c). As $(p, q) \rightarrow (1, 1)$ the species spread becomes super-diffusive and $r(\mathbf{x})$ is drawn from an increasingly homogeneous distribution. The variance of d_w (Fig. 9.3c) is maximised at $(1, 0)$ and $(0, 1)$. In this region the box-counting dimension is $\log(3)/\log(2)$, which is less than the other regions in parameter space excluding $(0, 0)$. Although the regions excluding $(0, 0)$, $(1, 0)$, $(0, 1)$ have box-counting dimension 2 for an infinite lattice size, finite size effects reduce the box-counting dimension and cause it to vary continuously between 2, $\log(3)/\log(2)$ and 1.



(a) Empirical mean for d_w for p, q parameter values. (b) Empirical variance for d_w for p, q parameter values.



(c) Relationship between mean of d_w and heterogeneity index. (d) Relationship between variance of d_w and heterogeneity index.

Figure 9.3: Summary of when growth rate $r(\mathbf{x})$ drawn from a multiplicative cascade model with parameters $(1, 1, p, q)$ where (p, q) are taken from the unit interval. Each point (p, q) is averaged over 100 realisations of the Multiplicative Cascade model for those parameters. d_w is in the sub-diffusive regime when p and q are low, which corresponds to a high Heterogeneity index (as indicated in Fig. 9.3c). As $(p, q) \rightarrow (1, 1)$ the species spread becomes super-diffusive and $r(\mathbf{x})$ is drawn from an increasingly homogeneous distribution. The variance of d_w (Fig. 9.3c) is maximised at $(1, 0)$ and $(0, 1)$, in this region the box-counting dimension is $\log(3)/\log(2)$, which is less than the other regions in parameter space excluding $(0, 0)$. Although the regions excluding $(0, 0)$, $(1, 0)$, $(0, 1)$ have box-counting dimension 2 for an infinite lattice size, finite size effects reduce the box-counting dimension and cause it to vary continuously between $2, \log(3)/\log(2)$ and 1.

9.1.6 Discussion

Although both the monofractal form of $r(\mathbf{x})$ and the multifractal form of $r(\mathbf{x})$ both have a box-counting dimension in the large lattice limit of 2, they both exhibit different dynamic scaling properties. The variance of the d_w exponent is narrower than in the case where $r(\mathbf{x})$ is a homogeneous Poisson point process. Comparison may be difficult however as when $r(\mathbf{x})$ is varied in the multifractal case, the density of points is kept at a constant, whereas it varies continuously between 0 and the N lattice points.

For the density chosen in the multifractal model ($\rho = 0.3$) the regime is entirely sub-diffusive ($d_w > 2$). The heterogeneity index gives a measure of the diffusive exponent d_w in this regime for a fixed density. The variance of d_w also depends on the heterogeneity index, but rather has a non-linear relationship with a peak followed by a linear decay. The variation in the diffusive exponent is thus maximised when the heterogeneity index is low. This is perhaps a surprising result as an increase in the heterogeneity index increases the variation of local scaling within r , which in turn would lead to an increasing variance in the scaling of the diffusion rate. d_w is however a *global* property of the lattice and hence the effect of increasing heterogeneity may be masked by taking the average over the lattice sites. It would be interesting to define a local scaling parameter for the species diffusion process and hence be able to ascertain how the heterogeneity affects the diffusion across the space.

9.2 Disease diffusion in a changing environment

The previous section examined the relationship between scaling and disease dynamics on a static environment where the rate of the disease dynamics is at a much faster rate than the rate of the vegetation process. Often the rate at which the disease propagates is on a similar length-scale to the rate of the vegetation process. We therefore consider a disease process that has a spatial scale l_D and a rate of infection β proportional to the rate of growth for a vegetation process with competition as described in Chapter 5. The purpose of the next section shall be to ascertain the effect of a disease process on the pattern formation and dynamical properties of the competition offset model. In order to proceed the competition offset model is modified such that each site can be in a third diseased state.

9.2.1 Disease model

The model is as follows, the system is a $N \times N$ lattice Ω where each $\omega \in \Omega$ can be in one of three states: E, O, D . E stands for an empty site; O is a site occupied with healthy vegetation and D is a site occupied with vegetation in a diseased state. For convenience, we define an indicator function I_X , which is defined as follows

$$I_X(\omega) = \begin{cases} 1 & \text{if in state } X, \\ 0 & \text{else.} \end{cases} \quad (9.13)$$

Similarly $I_X(\Omega)$ is the indicator matrix that contains $I_X(\omega_{ij})$ as its elements. The disease propagates with a kernel k_d which is Gaussian with mean 0 and variance l_d . The dynamics can then be written according to the probability of a site transitioning from state X to state Y in a time-step as (Fig. 9.4)

$$P(E \rightarrow O) = k_r * I_O(\Omega), \quad (9.14a)$$

$$P(O \rightarrow D) = \beta(k_d * I_D(\Omega)), \quad (9.14b)$$

$$P(O \rightarrow E) = (k_c * I_O(\Omega))^k, \quad (9.14c)$$

$$P(D \rightarrow E) = \gamma, \quad (9.14d)$$

where $*$ is the standard two-dimensional convolution operator. As a brief reminder of the previously discussed dynamics, the kernel k_r represents reproduction due to local clonal shooting and long-range sexual reproduction and takes the form of a Gaussian with zero mean and variance l_r . The transition from occupied site to unoccupied is due to death from competition and is mediated via the competition kernel k_c with

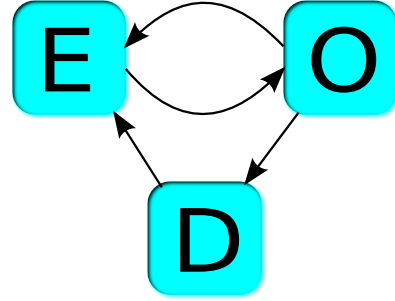
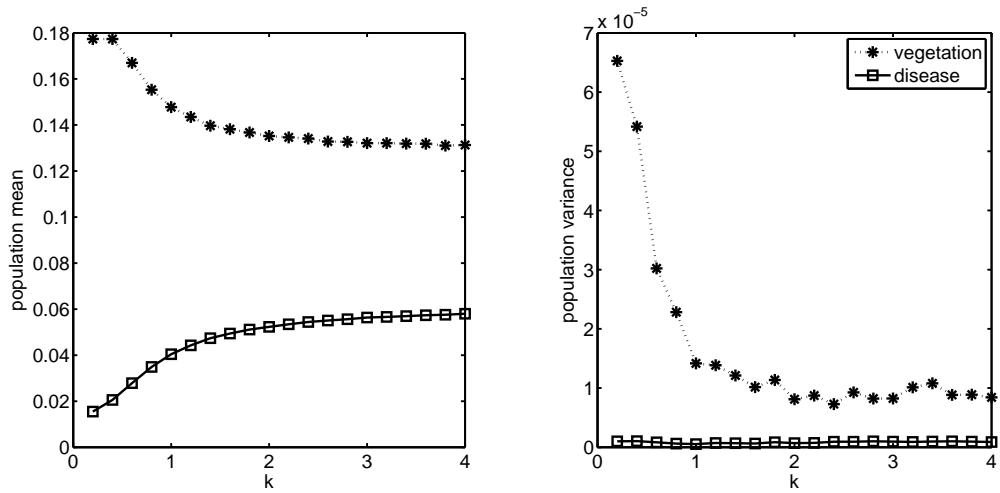


Figure 9.4: Schematic diagram of the disease model (Eq. 9.14)

an offset that is mediated by a mean with radius r and angle θ and a variance l_c . This represents death due to competition factors such as hydrological scouring and resource depletion. The offset comes due to environmental factors such as gradient or prevailing current. There is also a dimensionless parameter k , which represents the strength of the competition and controls the kurtosis of the competition kernel. For low k the effect of the competition is more uniform and hence represents



(a) Population mean for a range of competition values. (b) Population variance in time for range of competition

Figure 9.5: mean and variance for the population of vegetation and diseased-vegetation over a range of k , leaving other demographic parameters fixed at $\sigma_1 = 0.5, \sigma_2 = \sigma_3 = 1, r = 10, \theta = \pi/4, \gamma = 0.2, \beta = 2$. Decreasing competition (increasing k) leads to lower density of healthy vegetation and higher disease prevalence. This in turn, also leads to a lower variance in the vegetation population. Vegetation dynamics with strong competition was found to go on longer excursions than in the low competition regime where clusters quickly grow and are then invaded and quickly eradicated by disease.

stronger competition. For larger k the competition is weaker going to 0 as $k \rightarrow \infty$.

9.2.2 Competition in regulating disease spread

As a first investigation of how spatially-distributed competition controls disease spread we can compare the disease dynamics for a system which has no competition ($k \gg 0$), to a system where there is strong competition with offset such that banding is exhibited. Here we introduce a constant reservoir of infection by allowing all susceptible sites to become infected at each time-step with a small probability (10^{-4}). For a system that has no competition there are regular epidemics that spread throughout the population leading to a high degree of variability in the vegetation population and a high level of disease (Fig. 9.6). In contrast where competition is strong, but all other parameters are the same as the previous example, the vegetation forms into a banded structure. In this case the diseased state remains endemic at levels far lower than the healthy vegetation state. This is achieved as diseased patches are contained due to competition effects, meaning the disease cannot continue to propagate attacking other healthy patches, thus giving previously diseased patches time to recover (Fig. 9.7).

In order to determine how the strength of competition affects the prevalence of disease a number of simulations were performed over a range of k . For each $k \in \{0.2, 0.4 \dots, 4\}$, 100 replicate simulations were carried out for 10^4 time-steps, with other parameters held constant such that the disease process occurred on the same spatial scale as the competition and vegetative growth was primarily local ($\sigma_1 = 0.5, \sigma_2 = \sigma_3 = 1, r = 10, \theta = \pi/4, \gamma = 0.2, \beta = 2$). There was also a background disease rate of 10^{-5} , this was to ensure that the disease could never be completely eradicated, thus clusters of disease emerge spontaneously throughout the lattice. This rate is biologically reasonable as some long-range infection events may occur that are not captured by the local spreading term. For example, in seagrass, diseased shoots can become detached and float on currents where they can come into contact with susceptible leaves [Moore and Short, 2006]. The vegetation process has no background birth rate and as such there is a probability of complete extinction of the vegetation, although this was not observed for the parameter values and the simulations that were studied. In the presence of low competition ($k > 1$), disease is prevalent and higher than where competition is strong ($k < 1$), competition therefore helps to regulate the incidence of endemic disease for this particular model (Fig. 9.5a). The variances (Fig. 9.5b) of the vegetation population are also strongly affected by the interaction between competition and disease. For high competition values, bands of vegetation form that are susceptible to infection and death due to disease, this leads to the population performing large excursions away from the mean as large bands are infected leaving gaps that recover slowly due to the lower fecundity of the vegetation compared to the in the case of low competition.

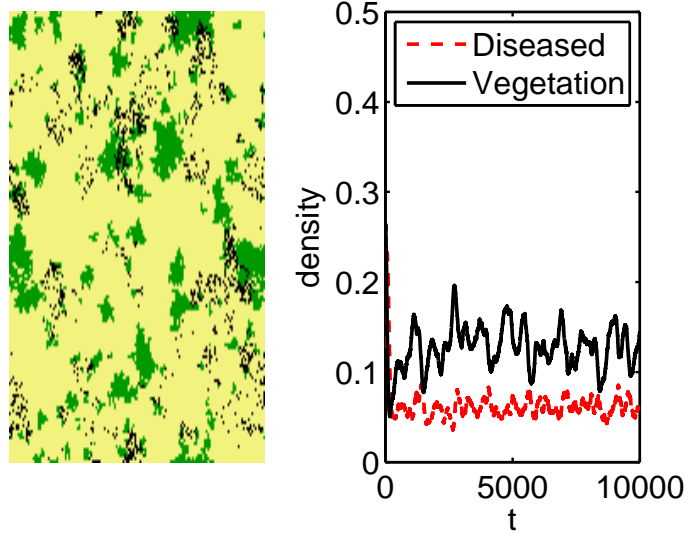


Figure 9.6: realisation of dynamics for the model described by Eq. 9.14 with parameters $l_r = 0.5, l_d = 1, \gamma = 0.2, \beta = 2$ and where there is no competition i.e. $k \rightarrow \infty$. The figure on the left shows a typical snapshot of the spatial distribution of healthy vegetation (in green), the diseased state (in black) and the empty site (in yellow). The right-hand side shows the time-series for vegetation and disease.

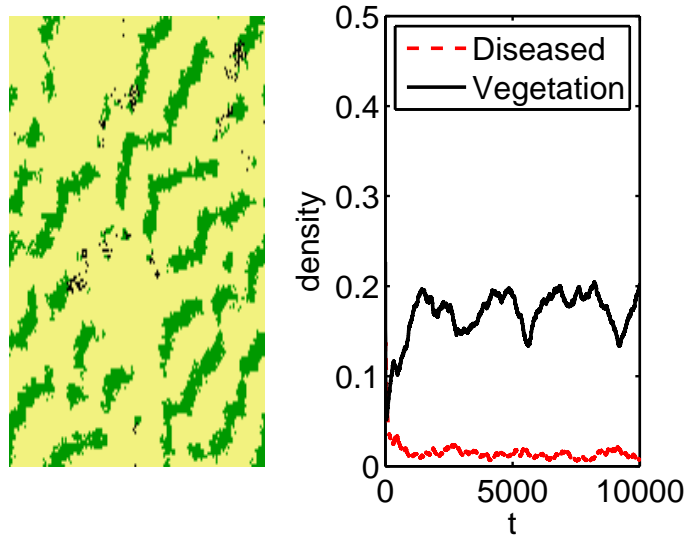


Figure 9.7: realisation of dynamics for the model described by Eq. 9.14 with parameters $l_r = 0.5, l_d = 1, \gamma = 0.2, \beta = 2$. Here there is strong competition present given by $l_c = 1$ and $k = 0.1$.

9.3 Evolutionary model for pattern formation in reaction to disease

Models where disease evolve into a critical state where cluster sizes are scale-free have been previously considered [Socolar et al., 2001]. Regular spatial vegetation patterns in vegetation have often been associated with environmental interaction, such as the presence or absence of nutrients or ground water or the presence of biotic interaction [Klausmeier, 1999; Rietkerk and Van de Koppel, 2008]. These spatial patterns have been shown to provide global benefit to the species, by allowing species to exist in environments that would otherwise be unfavourable and not permit their existence. In section 9.2 the effect of disease on spatial pattern was explored. It was shown that the presence of disease has a large impact on the variability of the vegetation dynamics, thus making it more vulnerable to extinction. Regular pattern formation in the form of banding was found to have a large impact on the disease dynamics. As the amount of spatial competition, and thus banding increases, the prevalence of disease decreases. This would naturally lead to areas of vegetation with banding surviving, whilst areas of vegetation with no banding going extinct. This represents group-level selection, where one group is able to proliferate at the expense of another group and hence selection is occurring at the large population level. For a global pattern to form, selection should also act on the level of individuals. The hypothesis then is that the amount of spatial competition an individual feels is a heritable trait and as such is selected for in the presence of disease. This leads to dynamics where, locally if there is no disease a plant is better to switch to have low mortality due to spatial competition as this will allow its offspring to proliferate at a higher rate. Where locally there is a strong disease presence, the spatial competition provides a way of isolating offspring from other patches that are in a diseased state, thus increasing their reproductive success. There is also a free-loader effect, where in a patch where all individuals feel strong spatial competition, it is better to switch to the strategy where spatial competition is low and thus offspring can proliferate at the expense of neighbouring vegetation.

These ideas will be explored using a modified model of the version outlined in section 9.2. The changes are with the competition parameter k , which will now be a spatially explicit variable $k(\mathbf{x})$. This is now modelled as a heritable trait. Here asexual reproduction is assumed, when there is a birth event at an empty site $\omega \in \Omega$, a parent is randomly selected according to the probabilities of a birth event at the lattice site ω for each individual. We also assume that k can either take high or low

values ($k_{\{H,L\}} = \{0.01, 100\}$) and at a birth event there is some small probability λ that the trait mutates from high to low/low to high value. Thus each lattice site can be in one of four states: E -empty site; L -occupied with low competition trait; H - occupied with high competition trait; D -diseased. The transition probabilities are therefore

$$P(E \rightarrow L) = k_r * I_L(\Omega) + \lambda k_r * I_H(\Omega), \quad (9.15a)$$

$$P(E \rightarrow H) = k_r * I_H(\Omega) + \lambda k_r * I_L(\Omega), \quad (9.15b)$$

$$P(\{H, L\} \rightarrow D) = \beta(k_d * I_D(\Omega)), \quad (9.15c)$$

$$P(\{H, L\} \rightarrow E) = [k_c * (I_L(\Omega) + I_H(\Omega))]^{k_{\{H,L\}}}, \quad (9.15d)$$

$$P(D \rightarrow E) = \gamma. \quad (9.15e)$$

We can explore the model dynamics when disease is not present i.e. when $\beta = 0$. In this region of parameter space L can always out compete H due to the fact that H feels increased competition than L . This can be studied using a phenomenological non-spatial mean field model. The density of L is denoted x_L and the density of H is x_H . It is assumed that in the small population limit growth of both populations is exponential. Competition is a second order process that is asymmetric in the population. Therefore the model dynamics may be written as

$$\dot{x}_L = r_L x_L (1 - a_L(x_L + x_H)), \quad (9.16a)$$

$$\dot{x}_H = r_H x_H (1 - a_H(x_L + x_H)). \quad (9.16b)$$

These are the competitive Lotka-Volterra equations in two-dimensions. By performing non-dimensionalisation we may reduce the number of parameters in the system. Using the substitutions $\rho = r_H/r_L, \tau = r_L t, u_L = a_L x_L, u_H = a_H x_H, b_L = a_L/a_H, b_H = a_H/a_L$ we arrive at the non-dimensional form of the equation

$$\dot{u}_L = u_L (1 - u_L - b_L u_H), \quad (9.17a)$$

$$\dot{u}_H = \rho u_H (1 - u_H - b_H u_L), \quad (9.17b)$$

by setting the derivatives to zero, three fixed points of the system can be determined $(0,0), (0,1)$ and $(1,0)$. The stability properties of these fixed points can then be calculated via the Jacobian

$$\begin{pmatrix} 1 - 2u_L - b_L u_H & -b_L u_L \\ -\rho b_H u_H & \rho(1 - 2u_H - b_H u_L) \end{pmatrix}. \quad (9.18)$$

By substituting in the values of the fixed points and calculating the corresponding eigenvalues, it can be shown that the origin $(0, 0)$ is always unstable. The other two fixed points are either stable or unstable depending on the values of b_L and b_H . If we assume that $a_H > a_L$ then the fixed point $(1, 0)$ is globally attractive and $(0, 1)$ is unstable. A number of realisations of the stochastic process were compared to the Lotka-Volterra dynamics (Fig. 9.8). Qualitatively, the realisations of the process conform to the mean field dynamics indicating that in the absence of disease, vegetation with low competition out competes vegetation that experience high competition.

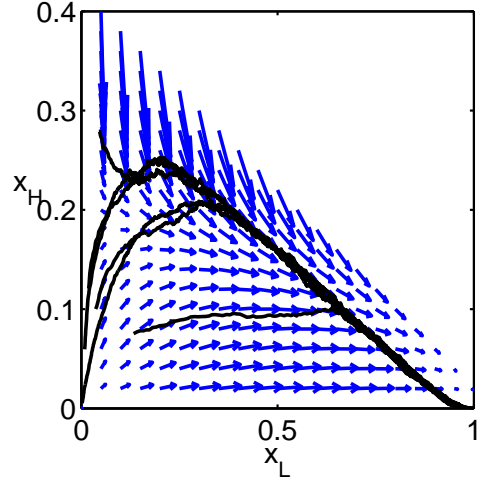


Figure 9.8: Realisations of altruistic competition model with no mutation and disease.

9.3.1 Evolutionary model for competition that is transmitted

When considering the competition parameter as a property of an individual vegetation, there is now an asymmetry between competition transmitted and competition felt by surrounding vegetation. The model outlined in Eq. 9.15 is considered for when competition is felt by surrounding vegetation as opposed to being transmitted. In this sense the strategy of having high competition is altruistic; when local density is high vegetation with high competition will have a higher probability of mortality, thus allowing vegetation with low competition to proliferate.

We can also consider the opposite case, when competition is transmitted rather than received. Consider if each individual has a variable c that is in the unit interval $[0, 1]$ and represents the strength of the competition that the individual transmits. The competition felt by an individual is then the weighted sum of all c weighted by the competition kernel. During a birth event a site becomes occupied and randomly picks an occupied site in its neighbourhood to be a parent. This site then receives the competition c from its parent. There is also a probability λ that the competition parameter mutates and becomes u , where u is a random variable drawn from the uniform distribution on the unit interval. The model space therefore has two processes operating on it: the finite-state Markov chain that characterises the disease and growth process $E \leftrightarrow O \rightarrow D \rightarrow E$ and the continuous-state Markov chain

that characterises the heritable trait process c . The state of the vegetation-disease process is again given by Ω . The state of the heritable trait c is given by C , where each site has the value c_{ij} . The dynamics may be summarised as

$$P(E \rightarrow O) = (1 - r_c c_{ij}) k_r * I_O(\Omega), \quad (9.19a)$$

$$P(O \rightarrow D) = \beta(k_d * I_D(\Omega)), \quad (9.19b)$$

$$P(O \rightarrow E) = k_c * [I_O(\Omega) \circ C], \quad (9.19c)$$

$$P(D \rightarrow E) = \gamma, \quad (9.19d)$$

$$f(c_{ij}|E \rightarrow O) = \lambda + \frac{(1 - \lambda)}{n_{ij}} \sum_{k=0}^{n_{ij}} \delta(c_{n_{ij};k} - c_{ij}), \quad (9.19e)$$

where \circ is the Hadamard-Schur product defined by $(A \circ B)_{ij} := A_{ij} B_{ij}$; n_{ij} is the number of sites in state O in the neighbourhood of site ij ; $n_{ij;k}$ is the k th neighbour of site ij and δ is the kronecker-delta function. The evolution of c is characterised by a probability density function f due to the fact that c is a continuous variable. Only sites where there is a birth event can update the competition parameter c , hence the probability density function is conditioned on there being a birth event $E \rightarrow O$. It is also assumed that there is some cost associated with the transmission of competition leading to death. This cost would be due allelopathic interaction such as through the production of toxins that lead to the death of surrounding vegetation [Gopal and Goel, 1993]. The energy cost is assumed to be linearly proportional to the amount of competition transmitted, hence this reduction in reproductive potential is given by the probability $1 - r_c c_{ij}$.

We may calculate what the expected value of C_i is given that there is no selection pressure for c . The expectation of the random variable C_{ij} given a birth event is given by the following

$$\begin{aligned} \mathbb{E}[C_{ij}|E \rightarrow O] &= \int_0^1 x f(x|E \rightarrow O) dx \\ &= \int_0^1 \left[\lambda x + \frac{(1 - \lambda)x}{n_{ij}} \sum_{k=0}^{n_{ij}} \delta(c_{n_{ij};k} - x) \right] dx \\ &= \frac{\lambda}{2} + \int_0^1 \frac{(1 - \lambda)x}{n_{ij}} \sum_{k=0}^{n_{ij}} \delta(c_{n_{ij};k} - x) dx. \end{aligned}$$

In order to evaluate this integral, assume that c_{ij} are spatially uncorrelated, hence the summation term disappears. The $\delta(c_{n_{ij};k} - x)$ term can be approximated by

considering the expectation over all sites. This implies that the δ function can be approximated by the probability density function g representing the distribution of c over all alive sites. Since c is not selected for, we may assume that this is uniform i.e. $g(x) = 1$. Hence

$$\begin{aligned}\mathbb{E}[C|E \rightarrow O] &= \frac{\lambda}{2} + \int_0^1 (1 - \lambda)xg(x) \, dx \\ &= \frac{1}{2}.\end{aligned}$$

So if the expected value of C deviates strongly from a $1/2$ we may reject the hypothesis that c is evolving under no selection pressure. In order to calculate the variance of c under no selection pressure, we assume that each C_i is independently drawn from a uniform distribution on the interval $[0, 1]$. We may then calculate the variance of the sum of these i.i.d random variables using a moment-generating function. Define the moment-generating function as

$$m_C(t) := \mathbb{E}[e^{tC}]. \quad (9.20)$$

Also define the random variable Y to be the sum of all C_i in the lattice that are in the occupied state i.e. if the total number of occupied sites on the lattice is n , then the moment-generating function for Y is defined as

$$m_Y(t) := \mathbb{E}[e^{t \sum_{i=0}^{n-1} C_i}] = (\mathbb{E}[e^{tC}])^n = m_C(t)^n. \quad (9.21)$$

The moment-generating function for C can be easily calculated from the definition of the expectation for a probability density function.

$$m_C(t) = \frac{e^t - 1}{t}, \quad (9.22)$$

hence the moment-generating function for the random variable Y is

$$m_Y(t) = \left(\frac{e^t - 1}{t} \right)^n. \quad (9.23)$$

To calculate the variance observe the relationship between the moments of Y and the derivatives of the moment-generating function

$$\mathbb{E}[Y^l] = m_Y^{(l)}(0). \quad (9.24)$$

We may therefore calculate the variance by first calculating the first derivative of m_Y

$$m'_y(t) = n \frac{te^t - e^t + 1}{t^2} \left(\frac{e^t - 1}{t} \right)^{n-1},$$

$$\lim_{t \rightarrow 0} m'_y(t) = \frac{n}{2},$$

where we have used L'Hôpital's rule for the indeterminate fractions. The second derivative at zero may be calculated in a similar fashion

$$m''_y(t) = n \left(\frac{e^t - 1}{t} \right)^{n-1} \left(\frac{t^3 e^t - 2t^2 e^t + 2te^t - 1}{t^4} \right)^{n-1}$$

$$+ n(n-1) \left(\frac{te^t - e^t + 1}{t^2} \right)^2 \left(\frac{e^t - 1}{t} \right)^{n-2},$$

$$\lim_{t \rightarrow 0} m''_y(t) = \frac{n}{3} + \frac{n(n-1)}{4}.$$

Hence the variance of Y may be calculated using the definition

$$\begin{aligned} \text{Var}(Y) &= \mathbb{E}[Y^2] - \mathbb{E}[Y]^2 \\ &= \frac{n}{3} + \frac{n(n-1)}{4} - \left(\frac{n}{2} \right)^2 \\ &= \frac{n}{12}. \end{aligned}$$

The average of C for a given configuration is given by $\bar{c} = \frac{1}{n} \sum_{i=0}^{n-1} c_i$. In order to find the variance of \bar{c} divide the random variable Y by n . Using the standard properties of the variance when multiplying by scalars we have that

$$\text{Var}(\bar{c}) = \frac{1}{12n}. \quad (9.25)$$

We may then compare this value to the value of \bar{c} in simulations. If the simulated \bar{c} is not in the range $(\frac{1}{2} - \frac{1}{12n}, \frac{1}{2} + \frac{1}{12n})$, then we may conclude that c is being selected for under the vegetation-disease dynamics.

For the altruistic competition model the calculation of the fluctuations in the null model are simpler. Take c_i to be the indicator for occupied site i as to whether the site is in state L . If we assume there is no selection pressure on c_i and that the c_i are independent, then we may assume each c_i is drawn from a Bernoulli random trial with probability $1/2$. The random variable Y representing the sum of the c_i

would therefore be binomially distributed with probability mass function

$$P(Y = k) = \binom{n}{k} \left(\frac{1}{2}\right)^n. \quad (9.26)$$

The mean and variance are then calculated in a straightforward manner to be

$$\mathbb{E}[Y] = \frac{n}{2} \quad \text{Var}[Y] = \frac{n}{4}. \quad (9.27)$$

The proportion in state L is equal to Y/n and hence the expectation and variance are

$$\mathbb{E}[L] = \frac{1}{2} \quad \text{Var}[L] = \frac{1}{4n}. \quad (9.28)$$

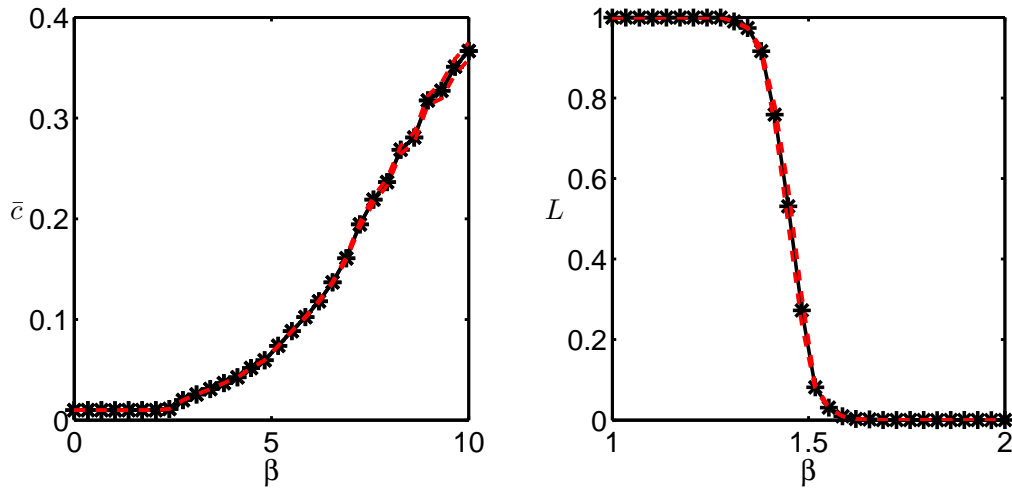
hence for simulations where the proportion of occupied sites in state L that are not in the range $(\frac{1}{2} - \frac{1}{4n}, \frac{1}{2} + \frac{1}{4n})$ are undergoing strong selection.

9.4 Results

9.4.1 Altruistic competition

Simulations were performed on a 150×150 lattice for 10^4 time-steps. For each region in parameter space simulations were repeated 50 times. For the altruistic model, the lattice was split in two: in the first half a population of vegetation in state L were placed randomly with probability 0.2. For the right half a population of vegetation in state H were placed randomly with probability 0.2. The simulations were allowed to evolve under the dynamics described in Eq. 9.15.

For the altruistic model there is a strong relationship between the proportion of L in the population and the ambient infection probability. When the ambient infection probability is large (10^{-3}) the final proportion of H in the population is low (less than 0.1). However, when the ambient infection probability is low (10^{-5}) the proportion of L begins to reduce. For increasing infection rate (Fig. 9.9b) the proportion of L is lowered significantly in the population to around 0.5 for high infection rate. The other parameters are kept fixed for $\gamma = 0.1, l_1 = 0.5, l_2 = 2, l_3 = 1, \mu = 0.01, r = 2$. A partial explanation for the large difference in dynamics when the ambient infection probability is small or large is due to the connectedness of the population in the L state. If there are many infections occurring then patches in the L state are disconnected, hence H patches do no better than L patches and L can invade H . Alternatively, when the ambient infection probability is small, the L patches have time to recover between outbreaks. This allows H to become a



(a) Increasing proportion of vegetation exhibiting spatial competition in the presence of disease. The population of vegetation with no imposed spatial competition for values of $\beta < 2$ are at the background mutation probability 0.05. With increasing force of the infectious agent, the proportion of vegetation with spatial competition increases as a fragmented landscape becomes the more dominant strategy. Parameters are $\gamma = 1, l_1 = 0.5, l_2 = 2, l_3 = 1, \mu = 0.01, r = 0, r_c = 0.25$.

(b) Evolving competition in an altruistic setting. Parameters are $\gamma = 0.1, l_1 = 0.5, l_2 = 2, l_3 = 1, \mu = 10^{-7}, r = 2$ and ambient probability of infection 10^{-7} . Note that this relationship does not hold when 10^{-3} where L is always the more successful strategy.

Figure 9.9: The effect of changing β on competition in the selfish and altruistic setting

single connected component, which is then wiped out when an outbreak does occur, thus reducing the survival of H patches. Although the proportion of L to H does approach 0.5 it is still significantly in the selection regime as it is outside of the range $(1/2 - 1/(4n), 1/2 + 1/(4n))$.

There is a strong interaction between the offset parameter and the persistence of vegetation in the L state. When the offset parameter $r = 10$, indicating that the effects of competition are felt from the existence of vegetation 10 lattice sites away, vegetation in state L can always out compete vegetation in state H regardless of the force of infection on the disease process. When the offset to the competition is small, $r = 2$, the feedback between presence of vegetation and competition is now more local. Sites in state H are patchy and negatively correlated with surrounding vegetation sites, this reduces the local density of vegetation below the percolation threshold, thus preventing the disease from spreading. For larger offsets, bands of continuous vegetation form, although these bands can reduce the impact of disease, they cannot out compete vegetation in state L due to the feedback between competition and disease being spatially separated.

Varying the infection rate β , the model displays properties analogous to a phase transition. For values of the force of infection that approach a critical value β_c the proportion of vegetation in state L on the lattice shows a sharp discontinuity. Below the value β_c , vegetation with the L strategy dominate and out competes the H strategy. For $\beta > \beta_c$, H out competes L , leading to a population of vegetation purely in state H (ignoring the transitory appearance of L vegetation due to mutation). Phenomena analogous to phase transitions have been observed in other cellular automata models of epidemics

[Fuentes and Kuperman, 1999], however the disease process driving a transition between two competing strategies of spatial competition is to our knowledge unique.

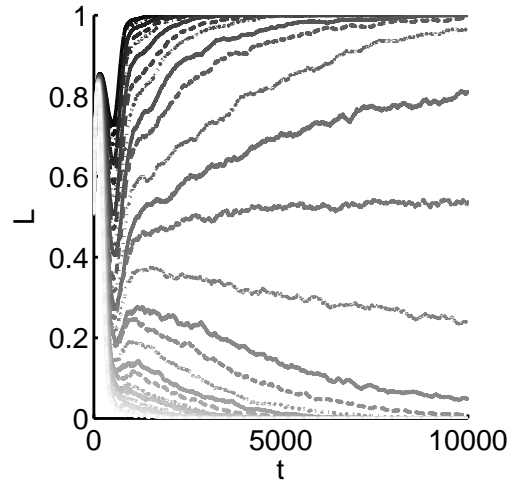


Figure 9.10: Realisations of altruistic competition model showing critical slowing down of relaxation time near critical force of infection β_c .

As β approaches the critical point β_c , a slowing down in the relaxation time is observed (Fig. 9.10). For values of β far away from the critical point, one strategy quickly dominates thus the relaxation time is short. Near the critical point, there is a transient co-existence of the two vegetation types, where periods of transient co-existence increase as β approaches the critical value. Theoretically, at the critical point both forms of vegetation should exist for all time, however as the transition is sharp, the value of β would have to be finely-tuned, which in simulations would not be possible due to finite precision. Fluctuations of the two vegetation types also increase sharply around the critical value. Fluctuations in the diseased state are high when β is lower than the critical point, increases sharply at the critical point and then reduces to a lower background level when $\beta > \beta_c$. Fluctuations in the density of vegetation also peak around the critical point.

9.4.2 Selfish competition

Simulations were performed on a 150×150 lattice for 10^4 time-steps. For each region in parameter space simulations were repeated 50 times. A spatially independent random starting configuration was chosen with density around 0.2. Each occupied site had a competition parameter c randomly drawn from the unit interval. The simulations were allowed to evolve under the dynamics described in Eq. 9.19 with model parameters $\gamma = 1, l_1 = 0.5, l_2 = 2, l_3 = 1, \mu = 0.01, r_c = 0.25$. The average competition parameter was measured for repeated simulations over various values of β (Fig. 9.9a). When disease is not present ($\beta = 0$) vegetation with no spatial competition dominates. For increasing β the proportion of vegetation with a non-zero competition parameter increases. The distribution of c is exponential with mean \bar{c} that increases for increasing β . For larger β , the disease overwhelms the vegetation leading to extinction (hence these extreme values of β are not shown in Fig. 9.9a).

9.5 Conclusion

This chapter has explored processes that co-occur and interact with vegetation. The focus has been in two parts: processes that can be considered fast on the time-scale of vegetation growth and processes whose rates are comparable to the rates of vegetation dynamics. In the first case, where the vegetation pattern can be considered static, the vegetation patches were modelled using a multifractal measure to characterise its scale-invariance and heterogeneity. This links well with the spatial analysis work performed in Chapter 4, where the local scaling properties of seagrass mead-

ows were analysed and found to have varying local scaling throughout the meadows. How this variation in local scaling affects species distributions is of interest to applied ecology research [Tokeshi and Arakaki, 2012]. For example, in a distribution study of Atlantic cod in an eelgrass meadow it was found that the maximum density of cod occurred where the local scaling of the eelgrass was intermediate between 1 and 2 [Thistle et al., 2010]. It was hypothesised that this density distribution can be explained due to an increased biodiversity where the fractal measure is on intermediate scales due to the heterogeneity, which provides more available niches for species to occupy. The section focused on how a measure of heterogeneity, calculated via the multifractal spectrum, can be used to assess the ability for species to colonise the vegetation distribution. If such a measure can be identified then it could be used to rapidly assess the quality of a vegetation ecosystem by determining how a species with a known dispersal would colonise. The simulated fragmented landscapes was found to have a diffusion exponent in the sub-diffusive regime, where the dispersal process became more sub-diffuse for a more heterogeneous landscape. There are however issues with using the multiplicative cascade model as a neutral model of vegetation heterogeneity. Although in the limit of an infinite lattice the fractal dimension will always be 2 and the density of points can always be guaranteed to be some value, this is not necessarily true of finite lattices. This is due to as the multiplicative measure becomes more heterogeneous there are regions where the probability of occupancy are very low. For an infinite lattice these regions would still contain points, but for a finite lattice this does not necessarily hold. There are also issues where the measure is close to 1, in these regions multiple points might occupy the same lattice site, thus reducing the overall density. The model therefore gives an indication of how heterogeneity decreases the diffusivity of a colonising species, but does not separate heterogeneity entirely from density, which would be required in order to establish if the effect is genuinely due to an increase in heterogeneity and not just due to a decrease of density or the box-counting dimension.

The interaction between disease dynamics and vegetation dynamics was further explored where the rate of dynamics of the disease and vegetation are comparable. The importance of the vegetation competition in regulating the spread of disease was assessed by comparing the prevalence of disease against the strength competition between vegetation that induces banding. Increased competition was found to limit the presence of the diseased state, but increase the variance of the population due to bands of susceptible vegetation becoming infected leading to a collapse of a significant proportion of the total population. In the limit where there is no compe-

tition, the model is akin to the forest fire model [Bak et al., 1990] where the system naturally evolves into a critical state, where the distribution of outbreaks of disease follows a power law distribution.

The evolutionary underpinning of banding in the presence of disease was explored by considering a vegetation population with two traits: high and low competition that is felt by surrounding vegetation. For small infection rates the low competition trait dominates and the dynamics are akin to Lotka-Volterra competition. There is a critical infection rate, however, after which the high competition trait dominates and can be sustained even for relatively large banding.

A selfish competition model was also explored where competition is transmitted rather than felt. Higher rates of infection lead to an increase in the average competition in the population. One possible mechanism of this could be where plants are able to, through toxins or other means, decrease the reproductive success of vegetation in the surrounding area. An example of auto-allelopathy can be found in white clover, where its presence has been shown to decrease the density of surrounding vegetation including itself [Macfarlane et al., 1982] as well as alfalfa [Jennings and Nelson, 2002]. There has been to date no evidence for the auto-allelopathy effect in seagrasses and hence the selfish competition model, where competition is transmitted, but not felt may not be appropriate to apply to the ecosystem. There is however evidence that long-range competition does effect seagrass density and this competition is dependent on the strength of local currents and wave-action [Van Der Heide et al., 2010]. There is therefore a hypothesis that this competition may not necessarily arise only due to the vegetation-environment interaction, but due to some evolutionary mechanism that arose in the presence of wasting disease.

Bibliography

- DA Andow, P M Kareiva, Simon A Levin, and Akira Okubo. Spread of invading organisms. *Landscape Ecology*, 4(2):177–188, 1990.
- Per Bak, Kan Chen, and Chao Tang. A forest-fire model and some thoughts on turbulence. *Physics letters A*, 147(5):297–300, 1990.
- Alex Best, Steve Webb, Andy White, and Mike Boots. Host resistance and coevolution in spatially structured populations. *Proceedings of the Royal Society B: Biological Sciences*, 278(1715):2216–2222, 2011.
- Jean-Philippe Bouchaud and Antoine Georges. Anomalous diffusion in disordered media: statistical mechanisms, models and physical applications. *Physics reports*, 195(4):127–293, 1990.
- James C Bull, Emma J Kenyon, and Kevan J Cook. Wasting disease regulates long-term population dynamics in a threatened seagrass. *Oecologia*, 169(1):135–142, 2012.
- Donald S Cohen and James D Murray. A generalized diffusion model for growth and dispersal in a population. *Journal of Mathematical Biology*, 12(2):237–249, 1981.
- Kenneth Falconer. *Fractal geometry: mathematical foundations and applications*. John Wiley & Sons, 2013.
- JD Fridley, JJ Stachowicz, S Naeem, DF Sax, EW Seabloom, MD Smith, TJ Stohlgren, D Tilman, and B Von Holle. The invasion paradox: reconciling pattern and process in species invasions. *Ecology*, 88(1):3–17, 2007.
- MA Fuentes and MN Kuperman. Cellular automata and epidemiological models with spatial dependence. *Physica A: Statistical Mechanics and its Applications*, 267(3):471–486, 1999.

- Brij Gopal and Usha Goel. Competition and allelopathy in aquatic plant communities. *The Botanical Review*, 59(3):155–210, 1993.
- John A Jennings and C Jerry Nelson. Zone of autotoxic influence around established alfalfa plants. *Agronomy Journal*, 94(5):1104–1111, 2002.
- Christopher A Klausmeier. Regular and irregular patterns in semiarid vegetation. *Science*, 284(5421):1826–1828, 1999.
- Sébastien Lion and Mike Boots. Are parasites “prudent” in space? *Ecology letters*, 13(10):1245–1255, 2010.
- MJ Macfarlane, D Scott, and P Jarvis. Allelopathic effects of white clover 2. field investigations in tussock grasslands. *New Zealand journal of agricultural research*, 25(4):511–518, 1982.
- Vicent J Martínez, Silvestre Paredes, and Enn Saar. Wavelet analysis of the multifractal character of the galaxy distribution. *Monthly Notices of the Royal Astronomical Society*, 260(2):365–375, 1993.
- Paul Meakin. Diffusion-limited aggregation on multifractal lattices: A model for fluid-fluid displacement in porous media. *Physical Review A*, 36(6):2833, 1987.
- C Meneveau and KR Sreenivasan. Simple multifractal cascade model for fully developed turbulence. *Physical Review Letters*, 59(13):1424, 1987.
- Kenneth A Moore and Frederick T Short. *Zostera*: biology, ecology, and management. In *Seagrasses: biology, ecology and conservation*, pages 361–386. Springer, 2006.
- Lisa K Muehlstein, David Porter, and Frederick T Short. *Labyrinthula zosterae* sp. nov., the causative agent of wasting disease of eelgrass, *zostera marina*. *Mycologia*, pages 180–191, 1991.
- LK Muehlstein. Perspectives on the wasting disease of eelgrass *zostera marina*. *Diseases of aquatic organisms*, 7(3):211–221, 1989.
- Robert J Orth, Tim JB Carruthers, William C Dennison, Carlos M Duarte, James W Fourqurean, Kenneth L Heck, A Randall Hughes, Gary A Kendrick, W Judson Kenworthy, Suzanne Olyarnik, et al. A global crisis for seagrass ecosystems. *Bioscience*, 56(12):987–996, 2006.

- Leslie A Real and Paul McElhany. Spatial pattern and process in plant-pathogen interactions. *Ecology*, pages 1011–1025, 1996.
- Max Rietkerk and Johan Van de Koppel. Regular pattern formation in real ecosystems. *Trends in Ecology & Evolution*, 23(3):169–175, 2008.
- Daniel Schertzer and Shaun Lovejoy. Physical modeling and analysis of rain and clouds by anisotropic scaling multiplicative processes. *Journal of Geophysical Research: Atmospheres (1984–2012)*, 92(D8):9693–9714, 1987.
- Joshua ES Socolar, Shane Richards, and William G Wilson. Evolution in a spatially structured population subject to rare epidemics. *Physical Review E*, 63(4):041908, 2001.
- H Eugene Stanley and Paul Meakin. Multifractal phenomena in physics and chemistry. *Nature*, 335(6189):405–409, 1988.
- Maria E Thistle, David C Schneider, Robert S Gregory, and Nadine J Wells. Fractal measures of habitat structure: maximum densities of juvenile cod occur at intermediate eelgrass complexity. *Marine Ecology Progress Series*, 405:39–56, 2010.
- M Tokeshi and S Arakaki. Habitat complexity in aquatic systems: fractals and beyond. *Hydrobiologia*, 685(1):27–47, 2012.
- Tjisse Van Der Heide, Tjeerd J Bouma, Egbert H Van Nes, Johan Van De Koppel, Marten Scheffer, Jan GM Roelofs, Marieke M Van Katwijk, and Alfons JP Smolders. Spatial self-organized patterning in seagrasses along a depth gradient of an intertidal ecosystem. *Ecology*, 91(2):362–369, 2010.

Chapter 10

Conclusion

The truth is not always beautiful, nor beautiful words the truth.

(Lao Tzu - *Tao te Ching*)

The main thesis explored throughout has been what information on the underlying dynamics can be garnered from a single spatial snapshot of a vegetation ecosystem. In recent years there has been an explosion in the technology and accessibility of this type of data that can be used in the ecological sciences. There are however, gaps in how this data can be used to supplement or even replace some aspects of traditional ecological survey techniques. This thesis has tried to close the gap between these two forms of data by assessing how various spatial statistics can be used and developed in order to assess certain features of the dynamics.

As a way of testing these hypotheses the developed methods have been applied to the example ecosystem seagrass. Seagrasses are a primarily clonal form of marine vegetation that form a number of characteristic spatial patterns including regular and scale-free. They are thus an ideal testing ground for assessing theories on the interaction between vegetation and environment and how this impacts the dynamics, hence throughout there has been a focus on this system although other example systems such as semi-arid ecosystems and mussel beds have been explored. The conclusions drawn should be more generally applicable to ecosystems where a single vegetation or single functioning group dominates.

The exploration of pattern and process has taken a number of avenues of investigation. Initially, based on previous studies the idea was to use the patch-size distribution in order to measure the persistence of the system. This took the form of exploring a number of plausible models of vegetation growth in the presence of intra-specific competition. The modelling has focused on the use of probabilistic cellular automata (PCA), which is in the class of Markov chain models and has been widely used in spatial ecology. There are a number of advantages to this approach: the rules that govern the dynamics focus on individual sites and hence plausible mechanisms of plant dynamics can be developed into the model, rather than aggregating these terms as would be done in a partial differential equation approach. PCA models are capable of producing both regular spatial patterns as well as scale-free and fractal patterns. Occupancy data, where an image is divided up into a regular grid, where each site either indicates it is occupied or not is easily compared to the resulting distributions from a PCA model and, as the basis of the model is a Markov chain, the full machinery of Markov chain theory can be used in the analysis of the resulting model.

PCA modelling was used to gain traction on the use of spatial heuristics as indicators of the underlying dynamic process. Two main fractal measures were identified along with several measures of persistence including the return rate and the time to equilibrium following a disturbance. The boundary dimension was found to be a good predictor of the growth rate for a growing cluster, however this became a poor indicator once the cluster had been established. The patch-size distribution was able to capture the persistence of the system as measured by the return rate under some circumstances when the demographic parameters were constant. However, if other parameters were varied then it did not reveal a strong relationship to the return rate. This provides insight into when a single spatial measure is a good indicator of persistence or growth and when they are not applicable. Indeed, any heuristic will necessarily only give some insight into the dynamics of the system, dependent upon a number of factors. The main conclusion is that the power-law or Korcak exponent of the patch-size distribution is an applicable measure for comparison between sites only when the environmental parameters differ between simulations. This was compared to the seagrass data obtained from the Isles of Scilly, UK where a similar relationship was found to the one observed in the simulations. Although there is this striking correlation between the simulations and the data, there is no explicit theoretical relationship established between the two measures beyond an ad-hoc argument that a more patchy environment would be expected to be less persistent.

The theoretical relationship between the Korcak exponent and the persistence was explored more by considering the whole vegetation system as an aggregation process. This method involved producing a mean field model of aggregation, defined using the kernel $k(i, j)$ which is the rate of aggregation between patches of size i and patches of size j . The births were modelled as a monomer injection term where individual patches entered the system at a constant rate. Death of the individual sites was also modelled at a constant rate. The resulting patch-size distributions conform well to the patch-size distributions in the seagrass data as well as other example systems. The aggregation process produces a power-law distribution of the patch-sizes whilst the death rate induces an exponential tail to the distribution. The model also predicts that the exponential tail increases for increasing values of the death rate, which has been observed in real systems such as semi-arid ecosystems, however until now only comparisons between simulations and data have been used to gain this insight. The aggregation model provides theoretical justification towards the relationship between the exponential tail and the death rate. A power-law aggregation kernel leads to a power-law patch-size distribution with varying exponent

that is linearly dependent on the exponent of the aggregation kernel. This provides a theoretical link between the dynamics of the system and the resulting exponent of the patch-size distribution or Korcak exponent. No direct comparison between the return rate and the exponent can be made for this model as there is no clear way to define a return rate for the system. The exponent can, however, be related to the relaxation time, where an increasing kernel exponent leads to a decreasing patch exponent and an increasing relaxation time.

The probabilistic cellular automata model of vegetation growth with spatial competition was able to reproduce broad spatial patterns found in a variety of spatial patterns including scale-free and regular spatial patterns. In real-systems both are present, where on smaller scales regular patterns may dominate, whereas on larger scales, more scale-free patterns emerge. This discrepancy is difficult to characterise using a single model of vegetation growth where the variance of the vegetation growth and the competition are pre-defined. As such, for cases where banding phenomena is being studied there is an offset to the competition term and whereas when scale-free phenomena is being studied an offset was not considered. In a real system, the variance of competition would be dependent on a number of factors such as wave action and current strength, these would naturally vary in space and time and hence lead to a range of regular and scale-free patterning. In order for the model to remain parsimonious the parameters of the system were kept constant in order to probe the various dynamics.

PCA model fitting was considered for vegetation occupancy data. In order to extract the relevant details of the dynamics, it was assumed that the spatial pattern had a regular structure in the form of vegetation bands. These spatial patterns produces strong spatial correlations that can be used in order to conduct inference on the PCA model. The main idea is to calculate the rate of change of the spatial correlation structure given a certain spatial snapshot. The novel method was to use this rate to construct a synthetic likelihood that can be used to perform inference. The method was shown to be valid when using synthetic data and some cases of real data, however was not able to capture the necessary parameters for other data where banding was present.

The final section dealt with the effect of spatial pattern on processes in a vegetation community. This took the form of noting certain fractal heuristics that are related to spatial heterogeneity and using various spatial patterns where this hetero-

geneity was varied in order to study the effect of the changing diffusion properties of an invading species or disease. It was hypothesised that a higher heterogeneity of the spatial pattern leads to a slowing of the diffusion. Although this was partially observed in the model, as density of the lattice could not be controlled for, the evidence for the hypothesis remains inconclusive. How disease spreads in a vegetation system where there are strong spatial competition effects between the vegetation was also explored. Banding was found to reduce the presence of disease by reducing the average cluster size of susceptible vegetation.

Finally an evolutionary explanation of banding was proposed by assessing the dynamics of vegetation with a distribution of spatial competition traits. There was a sharp transition dependent on the virulence of the disease between where the low competition trait dominated to where the high competition trait dominated. The ability for a vegetation species to evolve a strategy such as this as opposed to the apparent competition arising due to the interaction between vegetation and environment is not established. However, the model does show the plausibility of such a mechanism in explaining the rise of strong regular spatial structures. Seagrass is affecting by a wasting disease that lead to a large-scale epidemic in the 1930s where a large percentage of the North Atlantic seagrass was decimated. Various hypotheses have been proposed as to how this epidemic came about. There is another question as to why there have been no new epidemics since. Spatial regulation of seagrass may therefore go some way to explaining the current endemic state of the disease although for the evolution model, the force of infection needed for high competition to dominate is large. Therefore other factors would have to be considered for this discrepancy to be understood.

10.1 Future work

For regular spatial patterns the natural spatial statistic to use to investigate patterning is the spatial correlation function. For scale-free vegetation patterns, spatial statistics such as the boundary dimension and the patch-size distribution were used to investigate these aspects of the spatial pattern. A true vegetation system will naturally have a combination of both regular and scale-free pattern. The focus has been on isolating one of these types of spatial pattern and essentially ignoring the other aspects. This raises the question of being able to use a spatial statistic that is able to capture the full features of a spatial pattern that includes both regular and scale-free. One possible way to achieve this is to use the synthetic likelihood in-

ference method developed by Wood [2010]. The method requires various summary statistics that capture the full features of the dynamics (in this case the spatial dynamics) and then producing a likelihood based upon these measurements. The likelihood can then be used in parameter estimation for spatial data over a range of scales. There would be certain drawbacks to this method, chiefly among which is that a single tractable model may not be able to capture the full aspects of both the regular and scale-free patterning. At a certain range of length-scales and spatial locations only single factors may dominate, such as current or clonal growth. This allows certain parameters in the model such as strength of competition to be considered constant over the spatial and temporal range being considered. For larger scales, the parameters would vary throughout space as current strength and other aspects such as aspect vary. This would therefore require either modelling the functional form of this variation in space or else using covariates other than those taken from a spatial snapshot to inform the model. Another possible method that could be used to scale-up the inference analysis would be to assume the model parameters are constant over a certain spatial range. The inference method can then be performed on each of these spatial snapshots and then the parameters can be stitched together, to give a description of how spatial competition and reproduction varies through space. The drawback of this method would be that the surrounding sites do not inform the site under which inference is being performed. It also assumes that there can be step-changes in the reproduction and competition parameters in space, which may not be realistic.

Chapter 9 considers how processes occurring with a spatially-distributed vegetation ecosystem impact the spatial pattern and how the spatial pattern impacts the process. These ideas were explored by considering how diffusive processes such as disease and species invasion are impacted by the spatial pattern. This considered two forms of pattern: scale-free in the form of a multifractal multiplicative cascade model and regular in the form of banding. The scale-free distribution was considered static, this then relates to a class of diffusion processes known as quenched disorder models where the spectral dimension (a dimension that quantifies the diffusion process) can be directly related to the box-counting dimension of the fractal [Bouchaud and Georges, 1990]. For more general models the characteristics of the diffusion process do not have a known analytical solution and hence simulation was used instead. As these solutions do exist for simpler systems, there may exist analytic solutions for this system. Finally, the evolution of traits such as competition were used to investigate how evolution may play a role in the formation of pattern

in the presence of disease. This work is speculative and thus there are many questions raised that would need either field work, modelling or a combination of both in order to investigate. For the evolution modelling only competition strength was considered a heritable trait and the other individual traits such as the reproduction kernel were assumed to be homogeneous throughout all individuals. Other aspects of vegetation colonisation and growth may also be heritable and hence a more general evolutionary model would need to be constructed in order to explore this type of phenomena.

10.2 Conclusion

Pattern formation and persistence in vegetation ecosystems has received a large amount of attention in recent years. Research has focused on exploring plausible mechanisms that can induce pattern formation as well as other variables associated with the strength of patterning. The contribution laid out here has focused less on the plausible mechanisms that lead to pattern formation, but more on what information the pattern gives to the persistence of the underlying dynamics as well as what the pattern can say about the parameters that generate it. The original hypothesis is that certain spatial features of a vegetation system can be used as heuristics to estimate the dynamical persistence and to ascertain what other information the spatial pattern gives about the general dynamics of the system. The work here has clarified the relationship between pattern and persistence, provided theoretical insight into the origin of the relationship and developed methods of inference on spatial pattern. This work shows the importance of understanding the role of spatial pattern in the rapid assessment of ecosystem dynamics and points towards future work that can assess dynamic properties in a changing landscape.

Bibliography

Jean-Philippe Bouchaud and Antoine Georges. Anomalous diffusion in disordered media: statistical mechanisms, models and physical applications. *Physics reports*, 195(4):127–293, 1990.

Simon N Wood. Statistical inference for noisy nonlinear ecological dynamic systems. *Nature*, 466(7310):1102–1104, 2010.