School of Mathematics, Statistics, and Physics
Newcastle University
Newcastle upon Tyne
United Kingdom

Computational Modelling of the Spread of Tree Disease Through Forests

Sultan Abdullah Alyobi

Thesis submitted for the degree of Doctor of Philosophy

May 2021

Abstract

The impact of infectious tree diseases, such as Ash dieback, are wide-ranging: economic losses are estimated to be in the region of billions of pounds, for example with the cost of Ash dieback to the UK estimated to be £15 billion over the next century [1], risking thousands of jobs. Ecological losses can extend beyond the affected tree species to many dependent species, at a time when concern over the precariousness of many species is mounting; the loss of some tree species can also lead to desertification and permanent habitat loss. Palm trees, which are threatened by Bayoud disease, in particular can help to prevent desertification [2].

In order to control outbreaks of infections and mitigate this damage, it is desirable to develop reliable warning signals, which can indicate when a localised outbreak is likely to become an epidemic. Outbreak events can also be predicted or prevented by identifying characteristics, such as planting density or homogeneity, which may make a forest especially vulnerable or resilient. However, the possibility of identifying these through experiment is limited: tree diseases typically spread over long time scales, it is hard to control the relevant parameters, for example, maintaining a particular density of trees without continual human interference, and the areas of land required would be considerable. In this thesis, we will further develop existing mathematical models for disease spreading while making a broad examination of the effect of the parameters involved, in particular, the tree density, the infection rate and effective distance, and the uniformity of the forestscapes.

Our approach uses a lattice-based model, which combines a compartmental model of infection with a spatial component. Each lattice point represents either a tree or an empty space. We start with an infection with local interactions, such that an infected tree can only transmit the disease to its nearest neighbours. We consider two types of neighbourhoods: von Neumann and Moore, and quantify the observed dynamics by measuring the spreading velocity of the disease. As the forestscape density increases, we see a transition from local confinement to widespread outbreak at a critical density. The critical density is consistent with the percolation threshold, that is, the point at which the domain is spanned by a single connected cluster. We also further investigate an established framework of early warning signals such as the standard deviation, skewness and kurtosis of the velocity to predict the occurrence of the transition to the outbreak regime. Further measures, such as the mortality, also indicate a critical shift at the same critical density.

Previous work has focused on homogeneous forests, but this is not very representative of real forests, which frequently exhibit large scale features, such as man-made structures, or a natural clustering of trees. We develop an algorithm for generating forests with clustering, with a single parameter governing the degree of clustering. We characterise the observed distributions of trees in terms of Besag's function, which captures how far a distribution is from random at particular length scales. We also extract data on the distribution of trees in real forests, using the Mahalanobis distance, a measure of how close the properties of

an individual are to the properties of a set, to classify cells as either containing trees or not, and compare this to our synthetically generated data.

Using our synthetically generated clustered forestscapes in our local model, we observe that the dynamics strongly depend on the degree of clustering. We find that an increase in the level of tree aggregation suppresses the infection propagation. The critical density for the transition to the epidemic is consequently shifted to higher values for clustered forestscapes.

We then consider how the dynamics of an infection which spreads in a more realistic non-local way differ from an infection which only spreads locally. We observe that, in addition to the tree density and the degree of clustering, the dynamics are strongly influenced by the characteristic spreading distance, with some interesting interplay between the parameters. Measuring the spreading velocity, we again see that clustering suppresses the infection, but as the spreading distance increases, this effect is diminished. For the non-local infection, the transition to the epidemic regime is not clear when considering the velocity, but can be seen in the mortality. However, it occurs at much smaller tree densities than for the local infection. Overall, the dynamics of the infection can be remarkably altered by the switch from local to non-local interactions, with several observed effects that are relevant to understanding and predicting the propagation of tree diseases.

Acknowledgements

I would like to express my sincere gratitude to my supervisors Nick Parker and Andrew Baggaley for their guidance, patience, motivation and support throughout my journey in Newcastle. I thoroughly enjoyed working under their supervision and the knowledge they shared with me will be invaluable in my working environment.

I am extremely grateful to the following institutions for their financial support: my employer King Abdulaziz University who provided me with a scholarship while paying me 40% of my original salary; the government of Saudi Arabia for paying the tuition fees, annual tickets, medical insurance and generous monthly payment. Moreover, I am also thankful to the Saudi Embassy in London for their constant support during my journey.

Prior to coming to the UK, I knew that Newcastle University is recognised as a great institution. In all honesty, I found it extraordinary in many different areas. I want to thank Newcastle University for providing a first-class educational environment, 24/7 access to my office, valuable workshops and providing access to High-Performance Computing in an incredible location. Most importantly, I am grateful to the School of Mathematics, Statistics and Physics for their excellent organisation, talented IT experts and ensuring that everybody obtains a high-quality education.

I would like to thank my mother and father for the sacrifices they made and support. I would like to thank my wife and my kids, who came with me to the UK and helped me to get through the challenging times and who have also been great companions. Furthermore, I would like to take this opportunity to thank my brother and my sisters and the extended family for their continuous support and encouragement.

Finally, I want to thank Grange First School for being such an amazing and wonderful environment for my kids. They loved the school and had some fantastic times with their teachers and classmates. They will treasure the memories they have for the rest of their lives.

Contents

1 Why Model the Spread of Tree Diseases **1**

1.1 Tree Diseases 1

1.2 Economic & Ecological Impact 4

1.3 Different Applications to Model Tree Diseases 6

1.3.1 Individual-Based Modelling 6

1.3.2 Population Model (Compartment Model) 7

1.3.3 Continuum Model (Diffusion Model) 8

1.4 Lattice-Based Model 9

1.5 Motivation 10

1.6 The Aims and Objectives of This Study 11

2 Generating and Characterising Synthetic Forestscapes **13**

2.1 Introduction 13

2.2 Generation of Synthetic Forestscapes 13

2.2.1 Homogeneous Forestscapes 14

2.2.2 Heterogenous Forestscapes 15

2.3 Ripley's K Function 19

2.4 Besag's L Function 22

2.5 Periodic Boundary Conditions 22
2.6 Testing Besag's Function 24
2.6.1 Increasing the Disc Radius for a Fixed Box 25
2.6.2 Fixed Minimum Gap and Varying Radius 26
2.6.3 Fixed Radius and Varying Minimum Gap 27
2.6.4 Fixing the Radius and the Minimum Gap 27
2.6.5 Fixed Box Size and Varying Radius for 4 Discs 28
2.7 Conclusions 29

3 Percolation and the Properties of the Lattice 31

3.1 Introduction 31
3.2 What is Percolation Theory? 32
3.2.1 Clustering and Connectivity 32
3.2.2 Percolation on a Homogeneous Grid 33
3.3 Percolation Threshold 35
3.3.1 What is The Percolation Threshold? 35
3.3.2 Determining the Value of the Percolation Threshold 37
3.4 Generated Heterogeneous Data (Observed Tree Distribution) 38
3.4.1 Effect of Varying Clustering Parameter 39
3.4.2 Effect of Varying Seeds Rate 42
3.4.3 Effect of Varying Density 44
3.4.4 Effect of Varying Grid Size 46
3.5 Percolation Threshold for the Heterogeneous Data 47
3.5.1 Von Neumann Neighbourhood 48

3.5.2 Moore Neighborhood 49
3.6 Relating Von Neumann to Moore Neighbourhood 49
3.6.1 Statistical Approximations 50
3.7 Conclusions 54

4 Modelling the Infection Spreading in a Homogenous Forestscape 56

4.1 Introduction 56
4.2 The Model 57
4.2.1 Parameters 58
4.3 Examples of Infection Spreading 60
4.4 Mortality and Mortality Phase Diagram 61
4.4.1 The Phase Diagram for Mortality 64
4.5 Velocity and Velocity Phase Diagram 65
4.5.1 Velocity Definition 1 65
4.5.2 Velocity Definition 2 66
4.5.3 Velocity Definition 3 66
4.5.4 Time Series and Velocity distribution 66
4.5.5 Phase Diagram for Different Velocities 72
4.5.6 Catastrophic Shifts 76
4.6 Differences with Moore Neighbourhood 79
4.6.1 Infection Spreading Examples 79
4.6.2 The Phase Diagram for Mortality & Velocity 81
4.7 Conclusions 83

5 Modelling the Infection Spreading in a Heterogeneous Forestscape 85

5.1 Introduction . . . 85
5.2 Examples of Infection Spreading in a Clustered Field . . . 86
5.3 The Phase Diagram for Velocity . . . 87
5.4 The Phase Diagram for Mortality . . . 89
5.5 Conclusions . . . 90

6 Modelling the Infection Spreading Using Non-Local Interactions 92

6.1 Introduction . . . 92
6.2 Different Non-local Spreading Functions . . . 93
6.3 Verifying Numerics . . . 97
6.4 The Transient of the Infection Velocity . . . 98
6.5 Initial Infection Rate . . . 100
6.6 Examples of Infection Spreading . . . 104
6.7 Effect of Parameters on Infection Characteristics . . . 107
6.7.1 Characteristic Width of the Non-local Kernel . . . 108
6.7.2 Transmissibility of the Pathogen . . . 109
6.7.3 Forestscape Density . . . 112
6.7.4 Forestscape Homogeneity/Heterogeneity . . . 113
6.7.5 Universal Spreading . . . 118
6.7.6 Transition To Epidemic . . . 124
6.7.7 Further Investigation . . . 125
6.8 Distance to Nearest Infected . . . 131
6.9 Conclusions . . . 133

7 Comparing Synthetic and Real Forestscapes 136

7.1 Introduction . . . 136
7.2 Image Acquisition . . . 137
7.3 The Mahalanobis Distance . . . 137
7.4 Thresholding . . . 139
7.5 Computing Besag's Function on the Real Forestscapes . . . 141
7.6 The Effect of Changing Mahalanobis Sampling Box . . . 143
7.7 Characterising The Real Data . . . 145
7.7.1 Estimating Tree Density . . . 145
7.7.2 Estimating Gamma Through Least Square Difference . . . 145
7.7.3 Finland Data . . . 146
7.7.4 Canada Data . . . 149
7.8 Conclusions . . . 151

8 Conclusions **153**

A The Infection Spreading Algorithm **158**

B Forest Images From Finland **160**

C Forest Images from Canada **164**

Bibliography **168**

Glossary

ρ: tree density.

ρ_s: seeding rate.

β: infection rate (sometimes called the transmission probability or amplitude of the infection transmission).

T: infectious period (length of the infection).

N_n: the number of a selected cell neighbours.

N_c: the number of a selected cell possible connections.

γ: clustering parameter (a parameter that weights the probability according to the number of existing occupied neighbours).

γ_v: values for γ generated based on the von Neumann neighbourhood rules.

γ_m: values for γ generated based on the Moore neighbourhood rules.

$\mathcal{L}$: the grid size.

h: the strength of the clustering.

q: the characteristic lengthscale of the clustering.

S: number of susceptible trees.

I: number of infected trees.

R: number of dead trees

$\mathcal{M}$: the Mahalanobis distance.

T_c: the value of $\mathcal{M}$ at which we get the best representation of the forest trees.

$p(r)$: probability of a particular susceptible cell being infected by a particular infected cell at a distance r in a single timestep.

$p_1(r)$: probability that at least one tree at a distance r from a single infected becomes infected in a single timestep.

$p_2(r)$: probability that at least one tree at a distance greater than r becomes infected in a single timestep.

ΔI : number of new infections that a single infected tree is expected to directly generate from the initial condition of the infection to the first time step.

δ: minimum distance of a newly infected cell from any currently infected cell over the entire infection.

Chapter 1

Why Model the Spread of Tree Diseases

1.1 Tree Diseases

Infectious disease outbreaks, caused by non-native forest insects, pests, or other pathogens, can cause substantial damage to trees and agriculture crops. This can have a severe detrimental effect on both the wider ecosystem and the economy [3, 4, 5]. Many such outbreaks, with many distinct causes, have been observed in the 20^{th} and 21^{st} centuries, and with ever-increasing international trade, as well as climate change altering the existing ranges of species, further epidemics seem almost certain [6]. In order to reduce the impact of such diseases, it is necessary to develop an understanding of their behaviour; a key component of this is the development of robust models for disease transmission, which allow us to explore the effects of key parameters such as forest density and transmission rate. In this chapter, we will discuss some of the high profile tree diseases, the economic and ecological impacts of forest tree diseases, and different models for the propagation of pathogens through forests.

A number of studies have been undertaken to study tree diseases. According to data from the Food and Agriculture Organization of the United Nations (FAO), information concerning the effects of insect pests and diseases has been collected from nearly 75 nations around the world, with more than 300 distinct outbreaks [7]. It is notable that nearly 52 million hectares of forests are estimated to have been destroyed by disease in the period between 1980 and 2002.

There are several high profile tree diseases that have been documented. For instance, the emerald ash borer (EAB) (*Agrilus panipennis*) which is a common Asian species is

Figure 1.1: Aerial photograph of a forest in Ontario, Canada, showing the large impact of the emerald ash borer causing the death of ash trees [8].

known to cause significant damages to ash trees in the continent of North America outside its original home. This species was first detected in parts of North America in 2002 [9]. In Fig. 1.1, we have presented an example of the devastation caused by this beetle on ash trees. Although the species was accidentally introduced to the continent, these invasive organisms have caused great economic losses. In particular, within ten years it had killed millions of ash trees across the USA [10]. In Michigan state, it was reported that nearly 99 per cent of all ash trees whose stems are more than 2.5cm in diameter had been destroyed [11]. Since 2003, a minimum of 7 million ash trees have been killed in southeastern Michigan [9]. Unfortunately, all the different ash species in North America are susceptible to EAB [12]. It is most likely that EAB will functionally eliminate one of the most widely distributed tree species in the continent, causing severe ecological and economic ramifications [12]. In the period between 2004 and 2014 in the USA, this invasive species has caused the death of over 20 million forest and ornamental trees. Furthermore, some parts of Russia have also been affected by EAB, thus raising serious concerns inside mainland Europe. Although the pest has not yet been detected in any part of the United Kingdom, the government has been keen to ensure that appropriate preventative measures are implemented. In particular, the government is concerned about the dangers posed by the international movement of imported woods. It is worth remembering that the majority of the invasive pests are spread unintentionally via international trade.

Another serious disease for ash trees is ash dieback, which is caused by a native Asian fungus called *Hymenoscyphus fraxineus* [13]. This fungus usually causes fatality to the infected trees either directly or by weakening them. Ash trees across Europe are on the verge of being wiped out due to the deadly effects of both EAB and the ash dieback fungus. They are rampant across continental Europe, with up to 90% of ash trees infected in Denmark [14]. Ash dieback was confirmed in the UK in 2012, and is believed to have

arrived from Europe via imported plants and via wind. The mortality rate of this fungus in the UK is estimated to be between 90% and 99% [1].

Chestnut blight, another destructive disease, was unintentionally introduced into North America around 1900 as a result of the importation of infected stock from Japan. Since the introduction of the disease, nearly all of America's chestnut population has been destroyed [15]. Once a tree has been attacked by the blight, it results in a canker. Subsequently, the canker gradually destroys everything that is above the infection. However, the region below the infected area may survive, with the infected tree re-sprouting after a few years. According to Parker, [16], a blight is capable of surviving in the bark of dead trees, from where it then continues to infect any new sprouting tree.

Sudden oak death, caused by *Phytophthora ramorum*, is another destructive disease that has destroyed millions of tanoak and oak trees in California and southwest Oregon [17]. It was first discovered in 1995, but the origin of the disease still unclear. By 2030, it is estimated that the disease will have covered 14,000 km^2 in California which is nearly ten times the current area [18]. In what may be cause for concern, *P. ramorum* has also been found in nearly all woody plant species growing in mixed forests from central California to southern Oregon [19].

Dutch elm disease is a serious tree disease caused by a fungus called *Ophiostoma novo-ulmi*. This type of disease is known as a vascular wilt because it blocks the water transport systems of plants. Consequently, it causes the branches to wither and die. It spreads via elm bark beetles, and only presents in *Ulmus* and *Zelkova* genera in the *Ulmaceae* family. It is estimated that the disease has killed more than 60 million British elms during two epidemics, causing enormous economic effects [20]. Also, it is believed that the disease has caused the death of close to 97% of France's elms, and still continues to wreak havoc in other parts of the globe [21].

Date Palm (*Phoenix dactylifera*) is arguably one of the most well-known fruit trees in many parts of the middle east. It is a valuable cash crop for many farmers in Iran, Iraq, the Arabian Peninsula, and some parts of North Africa. However, this plant is also affected by several diseases with high-level of devastation such as Bayoud disease and inflorescence rot [22].

Bayoud disease, caused by the soil-borne fungus *Fusarium oxysporum* f. sp. *albedinis* [22], is arguably one of the most dangerous plant diseases known [23]. The pathogen is present in many parts of North Africa; it was first reported in Morocco more than a century ago [23], and subsequently spread to parts of Algeria [2]. It is estimated that it has destroyed nearly three millions trees in Algeria and twelve million in Morocco. It is imperative to note that the majority of commercial cultivators in North Africa are easily affected by Bayoud, as the disease is also known for its destructive habit of colonising the roots of other crops [22].

The inflorescence rot, *Mauginiella scaettae*, of data palm, known as Khamedj in some parts of North Africa, is another destructive disease to date palm trees. The disease was first reported in 1925 in Libya by Cavara [24]. A few years later, the disease had been reported in Iraq, Arabian Peninsula and some regions of North Africa. In the early 21st century the disease was also found in Spain [24]. The two outbreaks of this disease in the region of Basrah, in southeastern Iraq, during 1947-1948 and also 1977-1978, caused a decline in the annual harvest by 80% [25]. The region of Katif in Saudi Arabia faced a similar situation with a 70% decrease in the annual harvest [26]. This disease is considered to be the second most dangerous disease affecting date palms.

As discussed above, there are a great many high profile tree diseases that have caused devastation in many parts of the world. Clearly, these diseases have negative impacts on both the economic and ecological system. In the next section, we will discuss some of these ramifications.

1.2 Economic & Ecological Impact

Estimating the economic impact of diseases on trees is complicated, as although some damage is directly linked, for example in the tinder industry or costs to homeowners in removing diseased plants, other damage is less direct, such as reduced property prices and changes to stormwater runoff potentially leading to other property damage. Recent estimates have placed the overall annual ecomonic impact in the US to be billions of dollars [27].

The economic impact of a single disease alone can be devastating, as seen in many locations at many different times. The economic impact of sudden oak death in the state of California is expected to be in the region of tens of millions of dollars [28]. The chestnut population in the US has been almost eliminated since the introduction of chestnut blight; chestnut trees were known for their valuable timber with an industry worth $20 million in 1909 [29], equivalent to around $500 million today. In the UK, over the next 100 years, ash dieback is predicted to cost the British economy around £15 billion [1]. More than half of this cost is expected to occur in the next ten years. In New Zealand, managing the impacts of Dutch elm disease is estimated to cost the country over NZ$350 million [30].

The date palm is a critical source of income to many farmers in the Middle East and North Africa, and the appearance of palm tree diseases, such as Bayoud disease and the inflorescence rot, typically results in farmers being forced to flee their lands as a result of losing their primary source of income [22]. Date palm is of enormous economic importance in my home country Saudi Arabia. The total number of palm trees in Saudi is around 25 million, covering 170,000 hectares [31], which produced over one million tons of fruit in 2011 [32]. Moreover, the date fruit is of some religious significance, particularly during

periods of fasting [2].

Generally, it is notable that tree diseases have a direct bearing on a nation's economy, as the damage they cause can result in countries becoming consumers instead of being producers, as well as reducing the number of jobs in related sectors. Many countries recognise the impact of tree diseases and are increasing funding for research in this area. For example, the UK government spent some £37m on tree health research between 2012 and 2019 [33].

We now consider ecological impacts. The health of forests is very important for the ecosystem in several aspects. They help to preserve watersheds, which in turn conserves water and provides good habitats for plants and animals. Forests play a crucial role in gas regulation and carbon sequestration, storing carbon and releasing oxygen which is necessary for all complex life on earth [34]. Palm trees can help to prevent the process of desertification, which is extremely difficult to reverse and has severe and long-lasting negative impacts on ecosystems [2].

However, as we have discussed earlier in this chapter, non-native pests and pathogens can cause major tree decline and mortality in forest trees, which in turn can lead to adverse effects on the ecosystem. For example, a significant environmental impact is that ash trees are known to be the one and only food for more than 20 invertebrate species [14]. So, losing ash trees means losing these species, with the wider impacts on ecosystems difficult to predict but most likely significant, as trees are frequently keystone species [35, 36, 37]. This means that their existence has a disproportionate effect on the survival of other species and overall biodiversity. Another significant impact is that pests are known to interfere with the distribution of nutrient cycles, thus reducing seed production. The long-term destruction due to pests may lead to change in tree species that eventually alter productivity and biodiversity [27]. Furthermore, it is acknowledged that climate change is a considerable threat to forests, and the increase in temperature globally is responsible for the change in life circle and the outbreak frequency of a number of insects and pests specifically [38]. Changes to the climate can also modify the existing range of species, resulting in novel interactions, some of which are likely to be detrimental to tree health considering interactions when humans have introduced species previously.

In conclusion, the spread of diseases in a forest has economic, environmental and ecological impacts. Undoubtedly, different types of forest trees have different economic and ecological importance. Thus, it is imperative to understand how such pathogens propagate through a system of trees in an attempt to manage, control and minimise tree loss [39, 40]. We will try to better understand the propagation of tree diseases through mathematical modelling.

1.3 Different Applications to Model Tree Diseases

There are many models that can be utilised to study the spread of infection, such as the Individual-Based Model (IBM), the Population Model (Compartment Model) and the Continuum Model (Diffusion Model). In this section, we review and compare some of these models.

1.3.1 Individual-Based Modelling

Individual-based modelling (IBM) is a computational approach with spatial (continuous or discrete) and temporal dependence [41]. IBM provides a clear representation of the characteristics of the system by modelling the behaviour of its individual components [42]. These agents do not act randomly. Instead, they follow certain rules that govern their actions due to their environment and surrounding agents.

Often IBM has a stochastic component so that a meaningful result has to be obtained by repeating many stochastic realisations and averaging the behaviour. However, the advantage of IBM is that it is straightforward to incorporate different factors by way of adding different rules for the model. IBM will provide a starting point for the computational aspects of my thesis. Cellular automaton models, such as Conway's game of life, are a type of IBM with the restriction that agents move in discrete space; the cells migrate around a grid according to a set of rules [43].

In the early '90s people began using IBMs to model the spread of forest fires [44, 45, 46]. In this model, a forest is represented by a lattice. The lattice squares are either occupied by a tree or vacant. Each square can have one of three distinctive states: a healthy tree, a burning tree or a vacancy. Note that a vacancy represents either a burnt tree or no tree. In the simulation, the initial fire can be either set in a specific place in the lattice or placed in a random square, similar to the way lightning strikes a forest. Once a tree has a flame, it turns into a burning tree. Subsequently, a burning tree then spreads the fire to neighbouring trees, which spread the fire to their neighbours in turn. In this model, the fire can only permeate the forest if the tree density is large enough.

With the obvious parallels between the spread of a forest fire and the spread of disease, this model is a natural basis for a model of disease propagation in forests. Orozco-Fuentes *et al.* [47] extended this model by incorporating a stochastic component in the transmission, investigating the critical forest density and early-warning signals.

1.3.2 Population Model (Compartment Model)

Population modelling (PM) is an extremely important type of modelling that uses Ordinary Differential Equations (ODEs) to model and analyse a dynamical system. The population is split into different compartments, representing subsets of the population with different characteristics, which then interact through a set of ODEs [48]. Of the many types of PM, the most important with respect to this study is the SIR model. A population is split into three categories: susceptible (S), infected (I), and recovered (R). In the classical Kermack-McKendrick model [49], the populations evolve according to:

$$\frac{dS}{dt} = -\frac{\beta IS}{N}, \frac{dI}{dt} = \frac{\beta IS}{N} - \gamma I, \frac{dR}{dt} = \gamma I,$$

where N is the total population, β is the infection rate, and γ is the recovery rate. The model assumes that the population is fixed and completely homogeneous, so no births or deaths occur, and every member is equally affected, with no spatial structure or accounting for demographics such as age. It also assumes that there is no latent period, so an infected individual is immediately infectious, and the infectious period is the entire diseased period. An example of the population dynamics found from this model is shown in Fig. 1.2.

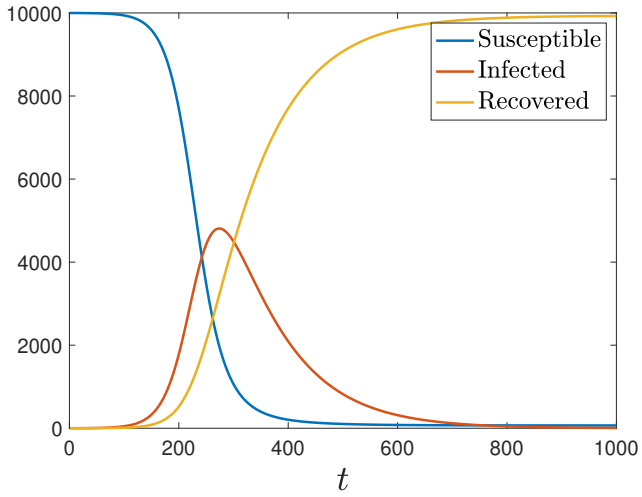

Figure 1.2: Population dynamics under the SIR model with $\beta = 0.03$, $\gamma = 0.01$, $N = 10000$, and $S = 9999$, $I = 1$, $R = 0$ initially.

The SIR model is extensively used for modelling human diseases, but can also be applied to modelling the spread of disease in forests as the underlying mathematics is rather generic. Many extensions to the SIR model have been proposed, such as the inclusion of background

birth and death rates, the introduction of additional compartments, for example, a deceased compartment and a latent compartment. There are many examples of alterations to the model for specific problems; for example, in [50], a variant is used to model Pine Wilt disease, which includes two levels of compartments to capture the spread of the disease by a nematode worm, which is in turn spread between the susceptible trees by beetles acting as carriers for the nematodes. Elsewhere, compartment models have been incorporated into bioeconomic models [39] which aim to model the interplay of biological disease propagation and economic choices.

We will make use of the SIR model in our simulations. However, the SIR model will not be our main approach, as the basic SIR model lacks any spatial dependence. More complex extensions to the SIR model have attempted to address this deficiency; however, they fall into the class of continuum models which we shall discuss next.

1.3.3 Continuum Model (Diffusion Model)

Population models can be extended to include spatial effects, with the governing equations then partial differential equations (PDEs) [51]. The advantage of this type of model is that it only needs to run a single time to obtain the average behaviour of a system, in contrast to the IBMs which need to be run many times to find the average behaviour. In the context of tree diseases, infections have been modelled using reaction-diffusion equations, identifying the possibility of travelling waves [52], and exploring the effects of periodic structure in the forestscape [53].

However, designing the model is generally far more difficult than for IBMs which include spatial dependence, as arbitrary rules that are simple to encode for individuals can not always be translated to the continuum representation in a straightforward manner. It is also worth noting that although the interactions are no longer stochastic, other factors may still result in the need to average over many ensembles. For example, if we want to understand how complicated, inhomogeneous spatial structures affect the model, we may not be able to say that a single field is representative of the type of structure we are interested in, and so may need to generate many structures from a distribution, and then simulate the continuum model on each of them. Although not a modelling issue as such, it can be very computationally intensive to use such models, as we need to update the status of every single point, rather than just the points of interest, such as points where an infection is active.

1.4 Lattice-Based Model

A lattice-based model can be viewed as a simple, locally-connected network, [54] discusses the implications of scale-free networks for disease transmission in plants; many of the tools developed in the context of animal diseases on networks could be usefully applied in this context". In modelling the spread of forest-fires, several authors have promoted lattice-based models, with individual trees or regions containing trees explicitly represented alongside their susceptibility and infectious status [44, 45, 46]. Such models are a variant of IBM models where the individuals are fixed on a regular grid, described in more detail in Chapter 2. The spatial distribution of trees is stochastic, and the transmission of the fire (or disease) from an infected tree to a susceptible tree occurs if and only if they are direct neighbours. In their recent paper 'Early warning signals in plant disease outbreaks' Orozco-Fuentes *et al.* modelled disease spreading in homogeneous forestscapes [47] following this method, but introduced an infection rate, i.e. a probability of infection being transmitted, to introduce stochasticity in time. This allowed them to identify two distinct phases of the disease dynamics (containment and widespread disease outbreak), criteria for being in either phase and the relation of this to percolation theory (discussed in Chapter 3), and statistical indicators of a potential outbreak. Our work here is based on this extended lattice model with an infection rate parameter.

The cells which may be occupied by trees are defined via lattice squares. Each cell can only be occupied by one tree; otherwise, the cell is empty. The population can be subdivided into three categories according to their epidemiological status: Susceptible (S), Infected (I) or Removed (R) [55]. Susceptible cells represent the presence of healthy trees which can be referred to as the S state. Infected cells which are represented by I signify trees that carry the infection. A tree will remain infectious until sufficient time passes, whereby it then becomes of the R state. Removed cells represent the population of trees destroyed by the infection. Once a tree dies from the infection, it is no longer infectious. Note that after a tree enters into the R state, it is unlikely that it will gain susceptibility again. This is due to the possibility that an invasion from another plant species will occur [56].

Since the spatial epidemic process is considered, it is usual to focus on the 2D case [55]. The fundamental unit of time used here is a time step, whereas the unit of length is a single lattice cell. Figure 1.3(a) represents the squared lattice before initialising the infection. Here, the lattice comprises only the occupied and unoccupied cells. The occupied cells are displayed in green, whereas the vacancies are displayed in white. Conversely, Fig. 1.3(b) represents the lattice after initialising the infection. Now some of the cells are infected or removed. Therefore, the infected are coloured in red, and the removed are coloured in brown.

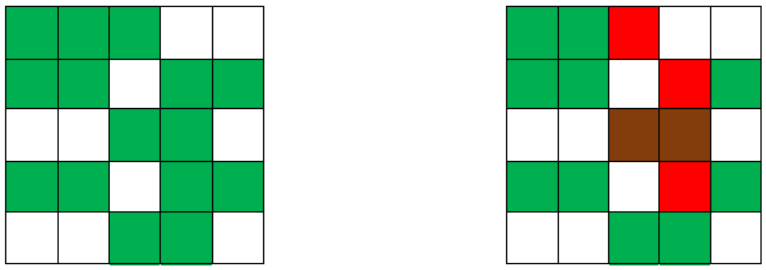

(a) (b)

Figure 1.3: Panel (a) represents a square lattice before initializing the infection while panel (b) represents the square lattice after initializing the infection. Susceptible, infected, removed, and empty cells are coloured in green, red, brown, and white respectively.

1.5 Motivation

The lattice-based models described above make two key assumptions: firstly, that the trees are distributed homogeneously in space, and secondly, that the infection is communicated by local, neighbour-to-neighbour interactions. Regarding the first of these assumptions, it is clear that trees, in reality, are typically distributed heterogeneously. This can be due to natural reasons, such as geographical landscape, species proliferation or competition with other species, or by human impact such as farming, planned forest management or construction. In Fig. 1.4, we show two different examples of how trees are distributed in real life. Regarding the naturally-occurring heterogeneity, several studies have examined this. In 2000, Manabe et al. concluded that clustering is a common pattern for most species, especially the rich tropical rainforest [57]. Also, a spatial distribution study of tree species in evergreen-deciduous broadleaf karst forests in southwest China and another study conducted in the Dinghushan Biosphere Reserve in China show that the clustered distribution is the dominant pattern in these forests. They also reveal that the percentage and intensity of clustering decreases as the spatial scale increases [58, 59]. We shall incorporate this into our model by developing a routine for generating synthetic heterogeneous forestscapes, with control of the degree of clustering through a single parameter.

(a) (b)

Figure 1.4: Two images of how trees are clustered in real life [60, 61].

Now turning to the second assumption, in reality, the infection spread is more complicated and depends on the pathogen. Timmermann et al. conclude that, of the possible mechanisms by which the fungal pathogen responsible for ash dieback is spread, the airborne transmission of spores is the most likely cause of primary infection [62]. A study by Chandelier et al. in 2014 investigated the distance at which the infectious agent was deposited, finding that while most remained within 50 m of the source, some proportion reached beyond this distance, potentially in concentrations sufficient to initiate new infections [63]. We shall model this by including a non-local spreading mechanism, with the infection rate now a function of the distance between the particular infected and susceptible trees, and spreading no longer restricted to direct neighbours. We also, for the neighbour-to-neighbour infection, use the Moore neighbourhood which is more realistic than von Neumann neighbourhood used by Orozco-Fuentes *et al.* [47], see Sec. 2.2.2 for definitions and discussions of these neighbourhoods'.

1.6 The Aims and Objectives of This Study

We aim to identify and develop mathematical models to understand how a disease propagates through forests, and to introduce the different factors affecting propagation, which reflect biological and geographical effects. If we can understand how pathogens spread, we could help guide intervention strategies to limit the loss of trees, providing ecological and economic benefits.

A number of challenges in modelling epidemics are discussed in [64, 65]. Although we do not directly address these in this work, our investigation of the effects of heterogeneous distributions, and the interplay of them with non-local interactions, is relevant to several of the topics raised, including "Capturing host spatial structure, even when data are limited", "Optimising dynamic controls in heterogeneous systems", and "How should we model contact structure in spatially heterogeneous populations?"

We will be basing our model on previous work using lattice-based models, as these capture both temporal and spatial aspects of the problem and allow us to easily incorporate specific behaviour, such as the exact way in which individual infection events occur.

In Chapter 2, we develop an algorithm for generating heterogeneous fields, controlled by a single parameter that induces clustering, which reduces smoothly to a homogeneous distribution as the parameter tends to zero. We quantify the level of clustering for different values of the clustering parameter using Besag's function.

In Chapter 3, we discuss the properties of our generated fields, both homogeneous and heterogeneous, in the context of percolation theory.

In Chapter 4, we simulate the infection spreading on homogeneous fields with nearest-neighbour propagation, probing the effects of the forestscape density and the infection transmission probability. We first verify Orozco-Fuentes *et al.* [47], and then extend this to Moore neighbourhoods, whilst exploring different options for quantifying disease propagation.

In Chapter 5, we study the disease propagation on our synthetic heterogeneous forestscapes, quantifying the spread and comparing this to the homogeneous case.

In Chapter 6, we introduce non-local spreading of the infection, to realistically model disease propagation. We consider the characteristics of some potential kernels and have selected the most appropriate, probe the effect of varying the effective distance at which spreading can occur, in conjunction with the heterogeneity, forest density, and infection rate.

In Chapter 7, we capture real forestscapes from publicly available aerial and satellite images, use an established method to extract information on their structure, and make a qualitative and quantitative comparison to our synthetic data.

Finally, we summarise our results in Chapter 8.

Chapter 2

Generating and Characterising Synthetic Forestscapes

2.1 Introduction

In this chapter, we will work towards generating heterogeneous forestscapes, which will then be used in simulations of disease propagation in later chapters. We propose an algorithm that generates homogeneous/heterogeneous forestscapes with a single parameter controlling the level of heterogeneity, and present different examples of homogeneous and heterogeneous forestscapes. Our algorithm uses the number of direct neighbours to populate points in the lattice, and we explore how different choices for what is considered a direct neighbour affects the resulting distribution. The generated distributions are quantified by using Ripley's K-function, which is a spatial statistical method that can describe how far a distribution is from random. We use Besag's function to normalise and standardise around zero. To help us to interpret Besag's function, we also test some simple distributions of points with clustering at well-defined scales.

2.2 Generation of Synthetic Forestscapes

We model a forestscape by a two-dimensional square lattice of $\mathcal{L} \times \mathcal{L}$ cells with populated cells represented by 1 and empty cells represented by 0. Each populated cell can represent either a single tree or a region populated with trees. The forestscape is flat and there is only one type of vegetation, with a spatial distribution that is either homogeneous or

heterogeneous. The degree of cell occupation is measured by the density, ρ, calculated as:

$$\rho = \frac{N_{\mathrm{T}}}{\mathcal{L}^2},$$

where N_{T} is the number of occupied cells.

2.2.1 Homogeneous Forestscapes

For the homogeneous field, each cell is populated independently of all other cells according to a Bernoulli trial with probability $p = \rho$, which is the desired overall density of the forestscape. Note that each cell is occupied with probability p and unoccupied with probability $1 - p$. Therefore, the level of occupancy is totally dependent on the value of ρ. It is apparent that the level of occupancy increases as ρ increases and reduces as ρ decreases. The total number of occupied cells then comes from a Binomial distribution: $N_T \sim B(\mathcal{L}^2, \rho)$. As we are sampling from a distribution, there is some deviation in the actual density obtained from the desired density. However, this deviation is negligible, with the standard deviation in the observed density reaching a maximum value of 10^{-3} when $\rho = 0.5$.

Figure 2.1 is an example of how points are distributed in a homogeneous field for a density $\rho = 0.62$. Thus, we will have approximately $500 \times 500 \times 0.62$ occupied cells. These cells are distributed uniformly and are not structured.

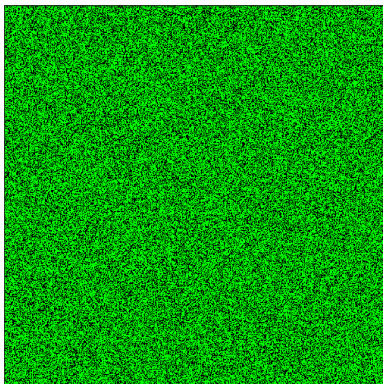

Figure 2.1: The point distribution for a homogeneous forestscape with density $\rho = 0.62$ and a lattice size $\mathcal{L} = 500$. Green cells represent the occupied sites, while the black cells represent the empty sites.

2.2.2 Heterogenous Forestscapes

We wish to generate heterogeneous distributions of occupied cells, as natural forests are not typically homogeneous. Doing so requires a different approach to the Bernoulli trial used in the previous section, as we need a mechanism which makes occupied cells more likely to be populated next to other occupied cells. Here we describe one possible approach to this.

We start with a $\mathcal{L} \times \mathcal{L}$ grid of unoccupied cells. We then seed the distribution by populating a proportion ρ_s of the cells, using the same process used in generating our homogeneous grids. The next step is the random selection of $\mathcal{L}^2(\rho - \rho_i)$ empty cells, where ρ is the prescribed density and ρ_i is the current density; i.e., the number of additional cells which we would need to populate to reach the desired density. We count the number of occupied neighbours N_n of each of the selected cells, and assign each of them a probability (to become occupied in the current iteration) according to:

$$p = \left(\frac{N_\mathrm{n} + 1}{N_\mathrm{c} + 1} \right)^\gamma ,$$

which we refer to as the clustering function, where N_c is the number of connected neighbours, and γ is a parameter that weights the probability according to the number of existing occupied neighbours. With these probabilities, we perform a Bernoulli trial for each selected cell, with success meaning the cell is to be populated. The occupation of cells is only updated when all trials within this step have been performed, so the number of neighbours and the probability associated with each cell is constant within each iteration. We repeat these steps, apart from the initial seeding step, until the desired density is reached.

In this work, we consider neighbouring cells based on a von Neumann neighbourhood ($N_\mathrm{c} = 4$) and a Moore neighbourhood ($N_\mathrm{c} = 8$), which we describe in more detail below. Since the method of connection is exceptionally important in relation to generating such a distribution, we are going to discuss the two methods of connections in the next two subsections, and we eventually proceed with one of them. We choose the arbitrary value $\rho_s = 0.01$ unless otherwise stated.

In Fig. 2.2, we present the calculated probability as the number of neighbours changes for various different values of γ. For $\gamma = 0$ all the selected cells will have a probability $p = \rho$ to be populated, which is the same as the homogenous forestscapes. However, as γ increases, selected cells with few neighbours become less likely to be populated than cells with many neighbours.

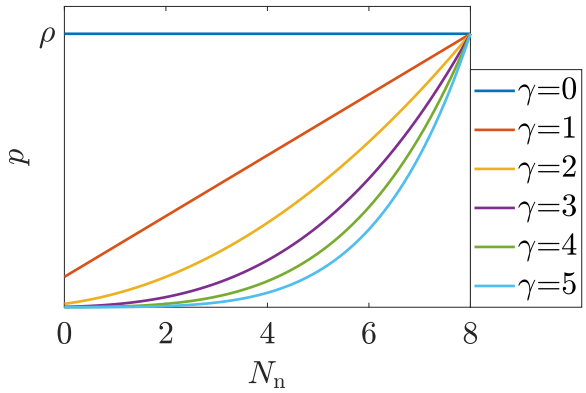

Figure 2.2: For Moore neighbourhood, values of p vs N_{n} for a range of values of γ.

Although we use integer values for γ, we are not actually restricted to them, and can use other real positive numbers. We also investigated whether using negative γ would allow the easy construction of distributions with anti-clustering. A normalisation is necessary to do this, as with the equation in its current form negative γ simply gives a constant probability of 1. An obvious choice of normalisation is to make the probability of a cell with no neighbours becoming occupied 1, which gives $p = (N_n + 1)^{\gamma}$, which we note is independent of the type of neighbourhood used. For $\gamma = -1$ this may lead to anti-clustering, however this is beyond the scope of this work.

Von Neumann Neighbourhood

This neighbourhood is defined in a 2D squared grid and consists of a central cell and its adjacent four neighbours [66]. It is named after John von Neumann, who used it to define the von Neumann cellular automaton and the von Neumann universal constructor within it [67]. Therefore for this neighbourhood, we have $N_{\mathrm{c}} = 4$ and $N_n \in [0, 4]$. Figure 2.3 illustrates examples of various levels of heterogeneous distribution generated based on the von Neumann neighbourhood. In this example, we vary the clustering parameter γ and fix the lattice size $\mathcal{L}$ and the density of occupied cells ρ. We see that as γ increases, occupied cells become more clustered together.

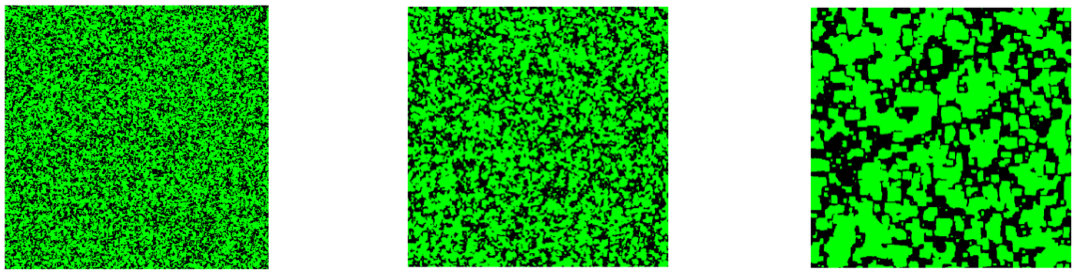

(a) $\gamma = 3$ (b) $\gamma = 5$ (c) $\gamma = 10$

Figure 2.3: The point distribution for heterogenous forestscapes created based on the von Neumann neighbourhood. Green cells represent the occupied cells, while the black cells signify the empty cells. We vary the clustering parameter γ for a lattice size $\mathcal{L} = 500$ and a density $\rho = 0.62$.

In this method, the probability that a cell which has no neighbours becomes occupied in one iteration of the generating algorithm can be calculated by:

$$p = (1/5)^{\gamma}.$$

Note that as γ increases, this probability decreases. Therefore, the points become more clustered together.

Moore Neighbourhood

This neighbourhood is defined in a 2D squared grid and consists of a central cell and the eight cells which surround it [68]. It is named after Edward Forrest Moore, the pioneer of cellular automata theory. For this neighbourhood, we have $N_{\mathrm{c}} = 8$ and $N_n \in [0, 8]$. Figure 2.4 displays examples of various levels of heterogeneity generated based on the Moore neighbourhood. In this example, we vary the clustering parameter γ and fix the lattice size $\mathcal{L}$ and the density of occupied cells ρ. We see that as γ increases, occupied cells become more clustered together.

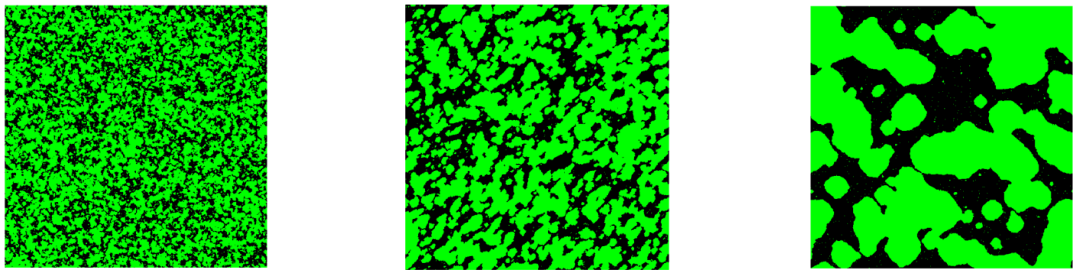

(a) $\gamma = 3$ (b) $\gamma = 5$ (c) $\gamma = 10$

Figure 2.4: The point distributions for the clustered data are created using the Moore neighbourhood. Green cells represent the occupied sites, while the black cells denote the empty sites. We vary the clustering parameter γ for a lattice size $\mathcal{L} = 500$ and a density $\rho = 0.62$.

In this method, the probability that a cell which has no neighbours becomes occupied in one iteration of the generating algorithm can be calculated by:

$$s = (1/9)^{\gamma}.$$

Just like for the von Neumann neighbourhood, this probability decreases as γ increases. However, the probability for Moore neighbourhood is much lower than von Neumann neighbourhood for a fixed value of γ.

Comparing the von Neumann and Moore neighbourhoods

We observe two key differences between the forestscapes generated using the von Neumann neighbourhood and the Moore neighbourhood. Firstly, the scale of clustering that arises for a particular value of γ is larger for the Moore neighbourhood. This is evident when comparing Fig. 2.3 and Fig. 2.4, where for each value of γ shown, the scale of the clusters is evidently larger when the Moore neighbourhood is used.

Secondly, for all values of γ, using the von Neumann neighbourhood forms squared clusters, with edges strongly aligned with the vertical and horizontal directions. On the other hand, using the Moore neighbourhood forms smooth rounded shapes, visually similar to clustering seen in nature. In Fig. 2.5, we fix $\gamma = 5$, and points are now only allowed to cluster around three specific seeds initialised at the start of the simulation. This allows us to observe that for large value of γ, the two methods can produce two different types of clustering. Principally, the von Neumann neighbourhood tends to produce squared blocks, whereas the Moore Neighbourhood tends to form circular shapes. If we increase γ, these square and circular shapes become more prominent. The clustering generated

based on the Moore neighbourhood is more realistic in capturing the curved distributions of natural forests. Thus, it is going to be our primary method for generating heterogeneous forestscapes.

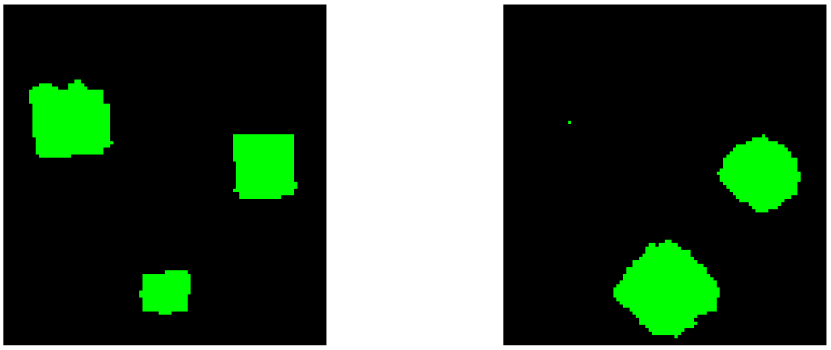

(a) (b)

Figure 2.5: Von Neumann neighbourhood (panel(a)) and Moore neighbourhood (panel(b)). All the parameters are fixed such that $\gamma = 5$, $\rho = 0.1$, $\mathcal{L} = 100$ and both grids are initialised with 3 seeds.

We note that not all cells selected as initial seeds ultimately nucleate clusters, as seen by the isolated cell in Fig. 2.5(b). This seems to be a relatively common occurrence, particularly for lower densities - briefly investigating, we found at least one isolated cell occur in 90% of distributions generated for the parameters used for this figure, dropping to 20% for $\rho = 0.7$ with all other parameters identical.

We note that the distributions generated using the Moore neighbourhood may also favour alignment with particular directions due to the underlying grid. It may be possible to avoid this by using a radially symmetric kernel to generate the clusters, e.g. a Gaussian of a fixed width, with a suitable cut-off distance to keep computational cost to a reasonable level.

2.3 Ripley's K Function

We see by eye that the clustering function leads to a visibly clustered distribution of occupied cells, but it is important to quantitatively characterise the clustering. This will allow us to precisely compare two distributions and to map out how the parameters involved change the clustering of occupied cells, e.g. the lengthscale and strength of the clustering. Hence, we have used a spatial analysis method termed Ripley's K function to describe how points are distributed in our grid (i.e. providing a description of the point pattern in our

generated grid). This function allows us to determine if the points in our grid are clustered, random or dispersed [69]. The importance of Ripley's K function comes from its ability to describe the point patterns at different scales so that we might have clustering at some scales and segregation at other scales. This method is generally used in plant ecology; it has been applied to describe herb distribution patterns [70], desert shrubs [71, 72], subtropical forest trees [58, 59] and tropical forest trees [73].

Consider an area A containing n points at locations $(x_i,\ y_i)$, with $i = 1, ..., n$, and distances between pairs of points $u_{i,j} = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}$. Ripley's K-function $K(r)$ can be calculated using the following formula:

$$K(r) = \left(\frac{A}{n^2}\right) \sum_{i=1}^{n} \sum_{j=1, j\neq i}^{n} I_r(u_{i,j}),$$

where

$$I_r(u_{i,j}) = \begin{cases} 0, & \text{if } u_{i,j} > r, \\ 1, & \text{if } u_{i,j} \leq r, \end{cases}$$

is an indicator function depending on how close two points are. So, this can be thought of as measuring the proportion of points which are within a particular distance to each other. Note that Ripley's K-function $K(r)$ has units of length squared and for homogeneous Poisson process $K(r) = \pi r^2$.

For points which are within r of the edge of the area under consideration, it is necessary to make a correction, as the number of points within r of them will be underrepresented. However, this can be avoided if the area considered is a subset of a larger area with the same spatial distribution. A typical approach, which we will use and which we describe in more detail in Sec. 2.5, is to produce a periodic region, the repeating part of which can then be tested without edge correction.

Figure 2.6 illustrates an example of a point pattern distribution and how to use Ripley's K function to study it. Ripley's K-function counts the number of neighbouring points found within a given distance (i.e. radius r) from each individual point, X_m. For each point on the grid, we consider circles of varying radius centred on the point of interest, ranging from a radius of 1 (the lattice spacing) up to the linear size of the grid. At that point, we compare the number of points within some distance to the expected number of points within this specific distance for uniform random data. If the number of points is greater than the expected number for a random distribution, the distribution is said to be clustered. Conversely, if it is smaller than for the random distribution, then the distribution is said to be dispersed.

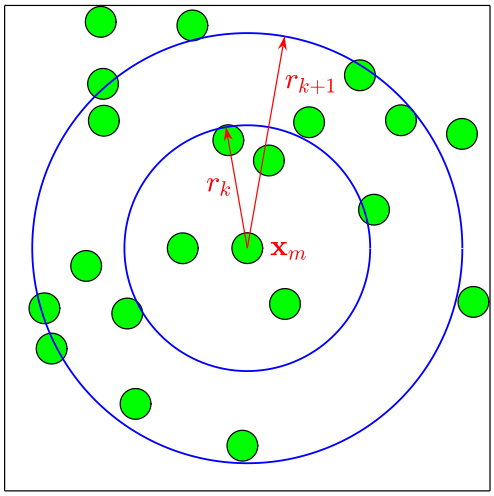

Figure 2.6: Around each point we draw circles. The radius of these circles, r, increases at a constant rate. Then we count the number of neighbours for each r and take the average.

To interpret the K function, we always conclude the following:

- When the observed value of K for a particular radius is smaller than its expected value for a random distribution, the lattice distribution is more dispersed than a random distribution.
- When its observed value is larger than its expected value, the lattice distribution is more clustered than random.

Figure 2.7 (a) gives an example of how to interpret Ripley's K-function. The expected number of points within a distance r is quadratic, as for a random distribution it is completely dependent on the area within a distance r.

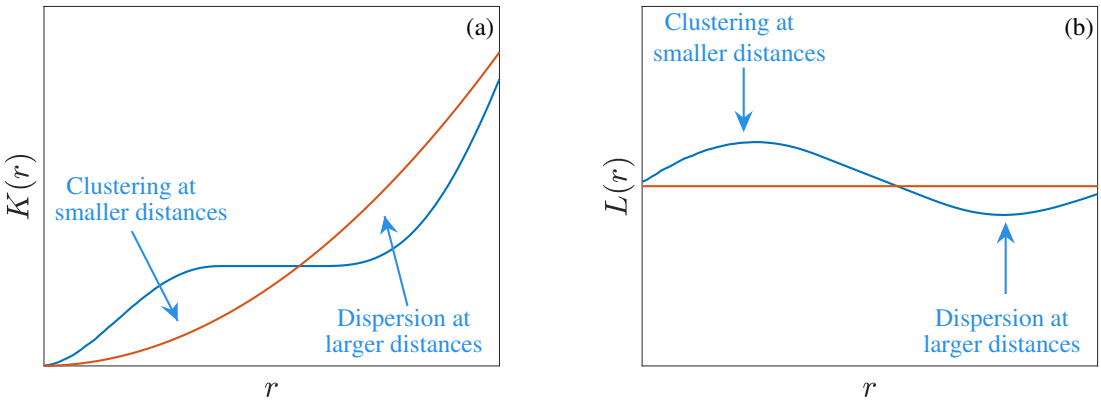

Figure 2.7: Panel (a) illustrates how to interpret Ripley's K-function, and (b) illustrates how to interpret Besag's function, for the same spatial distribution. The red curves show the expected value for a random distribution, while the blue curves show the calculated value for the particular distribution.

2.4 Besag's L Function

Given the dimensionality of $K(r)$, it is routinely transformed to Besag's L function. This allows $K(r)$ to be linearised and standardised around zero [74]. This transformation can be done using the following formula:

$$L(r) = \sqrt{\frac{K(r)}{\pi}} - r.$$

Now, the value of L can fall to one of three categories: zero, positive or negative values. Zero indicates that points are distributed randomly, which is similar to the completely random (Poisson) distribution. Positive values indicate that points tend to cluster together, while negative values indicate that the points are segregating. This is illustrated in Fig. 2.7 (b).

2.5 Periodic Boundary Conditions

It should be mentioned that if a calculated inter-point distance $u_{i,j}$ is greater than the distance between point $P_i(x_i, y_i)$ and the nearest plot boundary, part of the spatial neigh-

bourhood of point $P_i(x_i, y_i)$ lies outside the plot and hence, cannot be evaluated without a particular bias [75]. Therefore, to account for edge effects, periodic boundary conditions are used. Consider, when we were illustrating the point pattern in the previous section, we were interested in values of r that are up to the size of one dimension of the grid. However, this cannot be done if we only use the original grid. Thus, we used the periodic boundary conditions. To do this, we construct 8 copies of the original grid and place them around the original grid, as shown in Fig. 2.8. This allows us to calculate Ripley's function (and Besag's function) accurately up to the grid size $\mathcal{L}$ for points even at the edge of the original grid. Note that only points from the original grid are considered to be $P_i(x_i, y_i)$.

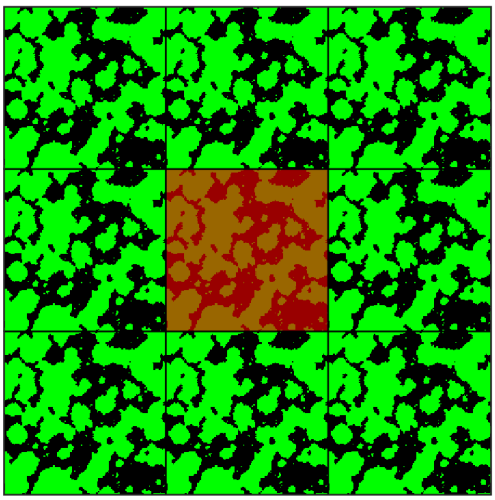

Figure 2.8: The original grid is shaded in red wrapped in by 8 copies (periodic boundary conditions).

Using periodic boundary conditions is not a unique approach to the problem. An example is the edge correction in the literature [75]. They include a discussion of five different edge corrections: a buffer zone, periodic boundaries conditions, and three variations of the weighted edge correction. We have tried the case when you have these extra regions surrounding the central region (i.e. a buffer zone), However, it is more computationally expensive to generate the synthetic domain over the larger grid, and so we proceed with the periodic boundary conditions.

2.6 Testing Besag's Function

To check the validity of the Besag's function and gain a basic understanding of how clustered distributions impact on Besag's function, we apply it to a distribution of points within a disc within the grid, as depicted in Fig. 2.9. In this figure, we have labelled two parameters of interest which are the radius of the disc r_p and the minimum gap between the border of the disc and the edges of the grid l_g. In panel (a), we have a single disc, while in panel (b), we display four discs. In the subsequent few sections, we are going to investigate the effects of changing r_p and l_g to gain an understanding of the relation between Besag's function and clustered distributions (who size is known a priori). This will help us in Chapter 3 when we use Besag's function to characterise the clustering produced by the clustering function (where the size of the clustering is not known a priori). We will start by considering one disc in the grid, but later consider multiple discs; we denote the number of discs by N_d.

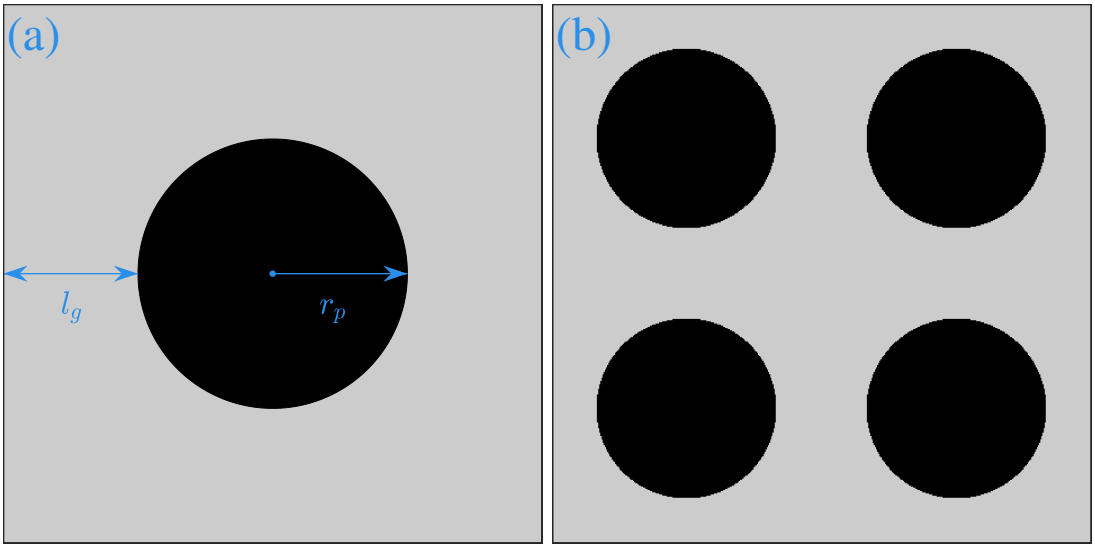

Figure 2.9: The grey region signifies the unoccupied cells, whereas the black region represents the occupied cells. The occupied cells are clustered together forming a disc shape with the radius r_p. The minimum gap between the border of the disc and the edges of the grid is named l_g.

When examining Besag's function, we are going to consider $h = \max(L(r))$, the height of the maximum value of $L(r)$ (i.e. the clustering amplitude), and $q = \text{argmax}(L(r))$, the position of the maximum value of $L(r)$ (i.e. the lengthscale of the clustering). Note that h is going to tell us the degree of the clustering or the strength of the clustering, and q is going to give us a measure of the characteristic lengthscale.

2.6.1 Increasing the Disc Radius for a Fixed Box

In Fig. 2.10(a), for a fixed grid size $\mathcal{L} = 100$, we increase r_p and consequently decrease l_g. We see a gradual decrease in the strength of the clustering h, and a gradual increase in the lengthscale of the clustering q. The decrease in the value of h as we increase r_p for a fixed grid size is because the grid is gradually filled with points (i.e. the increment of the density of the occupied cells). For these cases, as r_p increases the amplitude of the max h will be reduced, although the width of the peaks will keep increasing. We may conclude that as r_p increases for a fixed grid size, the $L(r)$ curves become flatter and flatter. This behaviour is expected because for a grid with $\rho = 1$, we expect $L(r)$ to be flat. In Fig. 2.10(b) and (c), we plot h and q respectively as functions of the disc radius r_p. We see that as r_p increases, the clustering amplitude h decreases, while the characteristic lengthscale q increases until $r_p = 24$, at which point it starts to take a constant value.

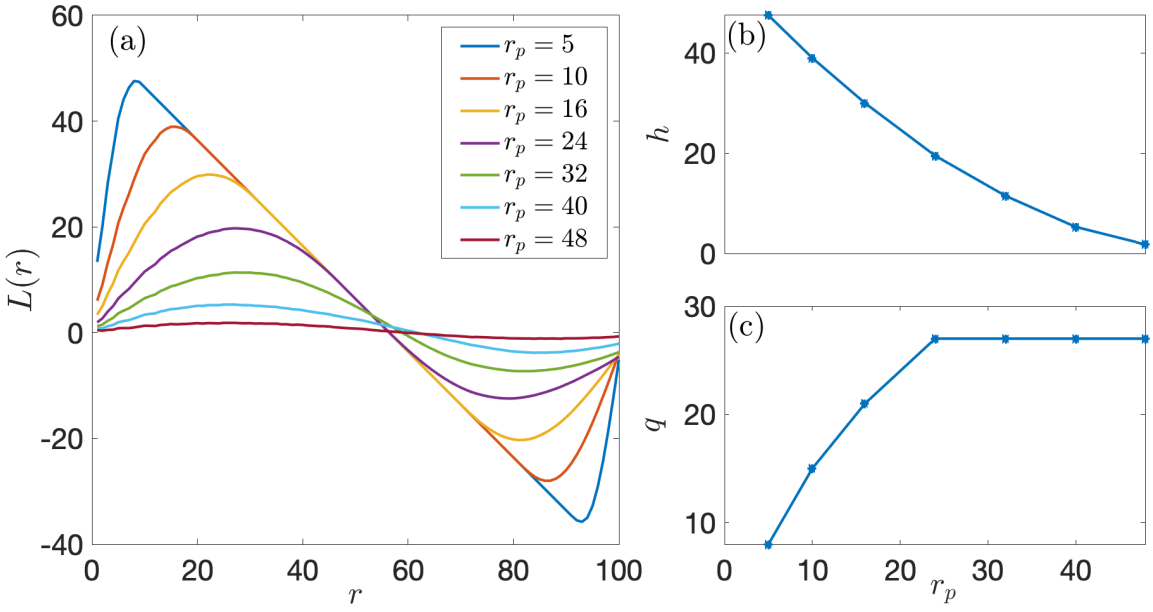

Figure 2.10: Panel (a), the calculated Besag's function for a domain of $\mathcal{L} = 100$ featuring a disc of radius r_p which is fully occupied with points (elsewhere the domain is empty of points). We gradually increase r_p and consequently decrease l_g. Panel (b), the amplitude of the clustering plotted as a function of the radius r_p. Panel (c), the lengthscale of the clustering plotted as a function of the radius r_p.

Considering the radius of the disc r_p, we notice two types of behaviour for the position of the maximum q. When $r_p \leq (0.25 \times \mathcal{L})$, as r_p increases the position of the maximum always increases to higher values of r. However, as r_p becomes higher than $(0.25 \times \mathcal{L})$, the

peak starts to flatten out, becoming very broad. Also, it becomes difficult to see by eye the exact location of the peak q. This is the effect of the normalisation, as the grid starts to fill with points. This can be visibly observed in Fig. 2.10.

For $r_p \leq (0.25 \times \mathcal{L})$, the peak will be between r_p and $2r_p$. As r_p increases beyond $(0.25 \times \mathcal{L})$, the peak will widen, and its maximum value will be attained at a value of r that is approximately the same as the position of the maximum when $r_p = (0.25 \times \mathcal{L})$.

2.6.2 Fixed Minimum Gap and Varying Radius

In Fig. 2.11 (a), we fix the minimum gap between the edge of the grid and the border of the disc $l_g = 40$. Note that $\mathcal{L} = r_p + 80$. We see that as r_p increases, the width of the peak also increases. Panel (b) shows that increasing r_p causes a slight decrease in the clustering amplitude h while panel (c) shows that it causes an increase in the characteristic lengthscale of the clustering q. This shows us that the value of q gives us a reasonable estimate of the radius, but it is not perfect.

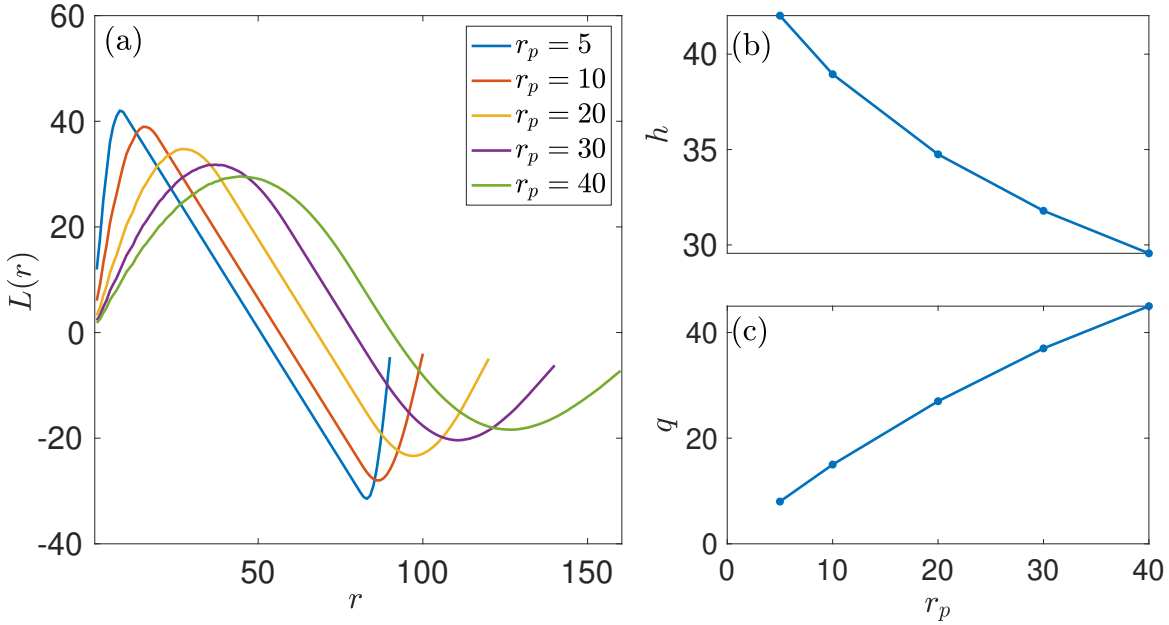

Figure 2.11: Panel (a): the calculated Besag's function for a domain which contains a disc of radius r_p which is fully occupied with points, and is empty everywhere else. We fix the minimum gap $l_g = 40$, and increase the radius r_p. Panel (b): the amplitude of the clustering plotted as a function of the radius r_p. Panel (c): the lengthscale of the clustering plotted as a function of the radius r_p.

2.6.3 Fixed Radius and Varying Minimum Gap

In Fig. 2.12, we increase the minimum gap l_g for a fixed disc radius $r_p = 20$. Note that $\mathcal{L} = 2(r_p + l_g)$. We observe that the amplitude of the clustering h gradually increases and the minimum of $L(r)$ gradually decreases. The lengthscale of the clustering q for the different curves are very close to each other; they are just after 20 and increasing in a small way. This is okay because we know that $r_p = 20$ is fixed and the grid size is increasing.

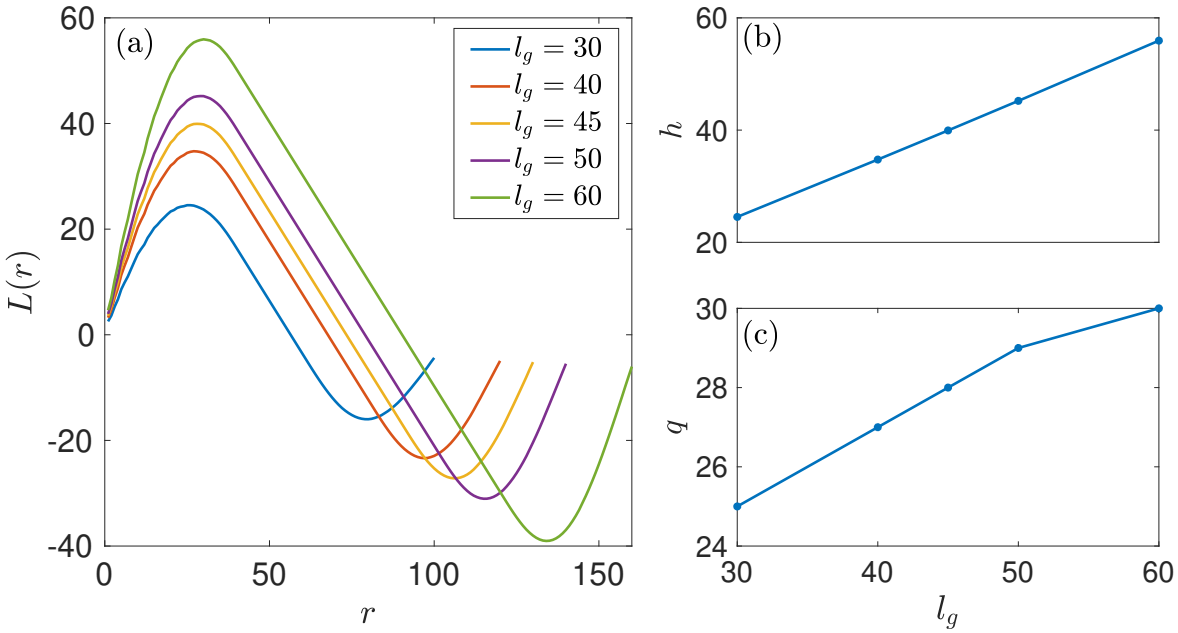

Figure 2.12: Panel (a): the calculated Besag's function for a domain which contains a disc of radius r_p which is fully occupied with points, and is empty everywhere else. We fix the radius $r_p = 20$, and increase the minimum gap l_g. Panel (b): the amplitude of the clustering plotted as a function of the minimum gap l_g. Panel (c): the lengthscale of the clustering plotted as a function of the minimum gap l_g.

2.6.4 Fixing the Radius and the Minimum Gap

In Fig. 2.13, we increase both l_g and r_p at a constant rate. Note that $\mathcal{L} = 2(r_p + l_g)$. Each time, we assign them the same value. We see that both the amplitude of the clustering h and the lengthscale of the clustering q increases.

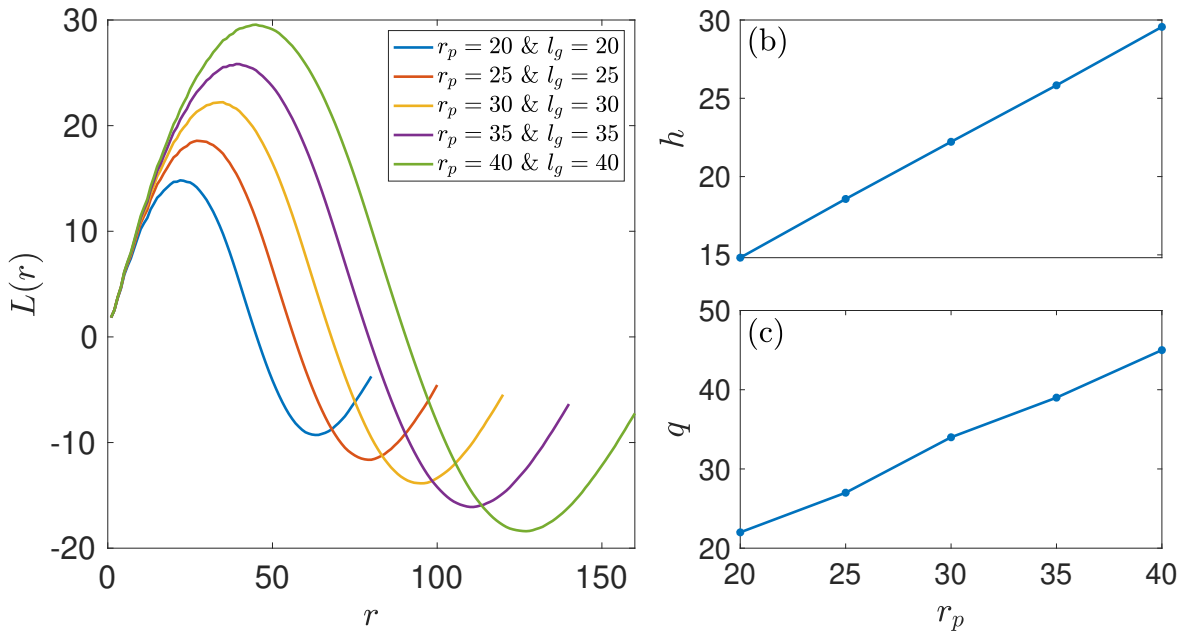

Figure 2.13: Panel (a): the calculated Besag's function for a domain which contains a disc of radius r_p which is fully occupied with points, and is empty everywhere else. We assign the radius r_p and the minimum gap l_g the same values while increasing them gradually. Panel (b): the amplitude of the clustering plotted as a function of the radius r_p. Panel (c): the lengthscale of the clustering plotted as a function of the radius r_p.

2.6.5 Fixed Box Size and Varying Radius for 4 Discs

Now, we are going to increase the number of discs N_d to see how that affects the values of $L(r)$. In Fig. 2.14, our $\mathcal{L} = 100$ domain has four discs ($N_d = 4$) of known sizes, and the discs are fully populated with occupied cells. In Fig. 2.14 (a), we calculate $L(r)$ at different scales r, and we see that increasing the number of the clusters causes an increase in the number of peaks. For very low densities, such as for $r_p = 5$, the peaks are very sharp; as r_p increases the "sharpness" of the peaks starts to diminish. Considering the largest peak, panel (b) displays the amplitude of the clustering h as function of the radius r_p. We see that increasing the disc radius causes a decrease in h. Panel (c) shows the lengthscale of the clustering q as a function of the disc radius. We observe that while $r_p \leq 15$ the lengthscale increases as r_p increases. As r_p becomes larger than 15, the lengthscale of the clustering starts to decrease.

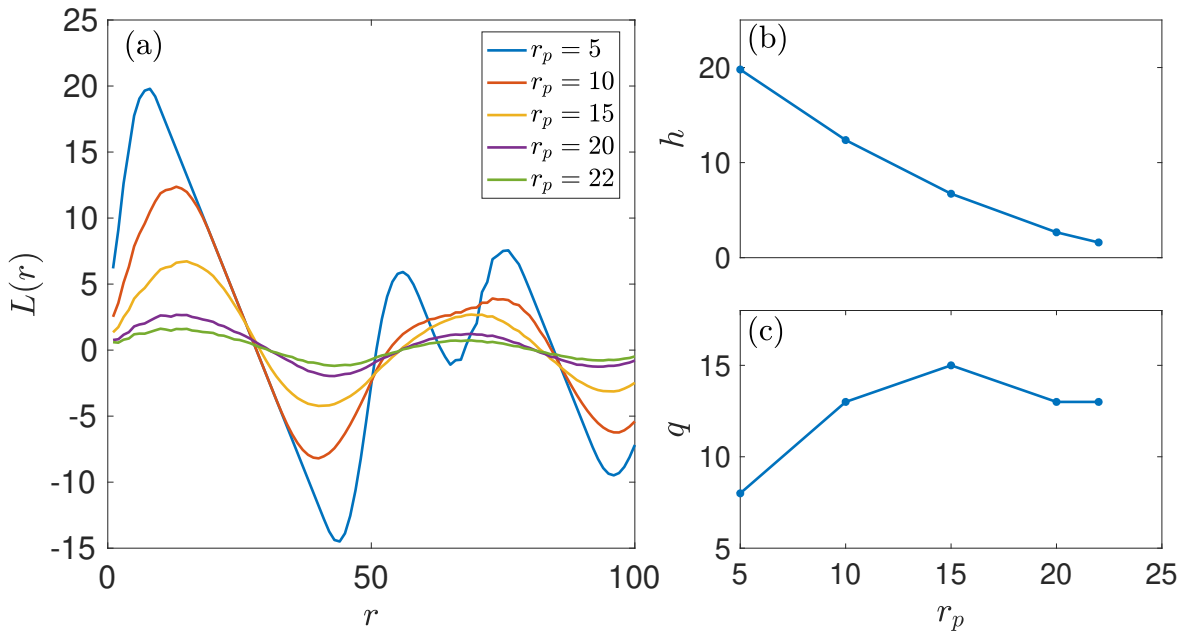

Figure 2.14: Panel (a): the calculated Besag's function for a domain containing four discs of radius r_p which are fully occupied with points, with the domain otherwise empty. We gradually increase r_p and consequently decrease l_g. Panel (b): the amplitude of the clustering plotted as a function of the radius r_p. Panel (c): the lengthscale of the clustering plotted as a function of the radius r_p.

2.7 Conclusions

We developed an algorithm that can generate homogeneous/heterogeneous forestscapes. Four parameters control our synthetic forestscapes: the lattice size $\mathcal{L}$, density ρ, seeds rate ρ_s, and a clustering parameter γ. For $\gamma = 0$, the forestscape is homogeneous. As γ increases the heterogeneity level increases.

Our synthetic forestscapes are generated based on two types of neighbourhoods: von Neumann and Moore. Our results showed that the clustering that results from them has two significant differences. First of all, using the same clustering parameter γ, the Moore neighbourhood tends to produce larger clustering than von Neumann. Also, the clusters generated based on the von Neumann neighbourhood tend to have squared shapes while clusters generated based on the Moore neighbourhood tend to have smooth and rounded (circular) shapes. The clustering based on the Moore neighbourhood is more similar to the curved distribution of natural forests, and so we will primarily use the Moore neighbourhood for generating heterogeneous forestscapes.

In order to compare distributions, we need to have a method to quantify the observed clustering. We use Ripley's K-function, which indicates whether a distribution is more clustered or more dispersed than a random distribution at particular lengthscales. This sensitivity to lengthscales allows us to compare both the strength of the clustering, and the lengthscales at which it occurs, for different distributions. This will allow us to precisely compare two distributions and to map out how the parameters involved change the clustering of the occupied cells. For ease of interpretation, we transform Ripley's K-function to the normalised, standardised Besag's function.

To test how Besag's function works, we generate a single disc that is fully populated with cells at the centre of the grid, and we use two parameters of interest: the radius of the disc r_p and the minimum gap between the edge of the disc and border of the grid l_g. We observe positive values of $L(r)$ at small scales which indicates clustering. This is reasonable because cells in the disc are clustered together at small scales.

For a fixed grid size, when r_p is very small, we see that the clustering amplitude $h = \max(L(r))$ has a high value. This value decrease as r_p increases. This is expected because the grid starts to become filled with points. Also, we observe that $L(r)$ has negative values at large scales for very high l_g and very low r_p. These negative values decrease as r_p increases. We observe that increasing the minimum gap between the edge of the disc and the border of the grid l_g causes an increase in the amplitude of the maximum h, while increasing the radius of the disc r_p causes a decrease in the amplitude of the maximum. The most important thing is the lengthscale of the clustering q (i.e. the position of the maximum). Clearly, as r_p increases until $\approx 0.25 \times \mathcal{L}$, the peak becomes wider, and the value of the location of the maximum increases. As r_p grows above $0.25 \times \mathcal{L}$, the peak becomes much broader, and the position of the maximum start to become stable. This is expected because the grid starts to get filled with points.

Then, we increase the number of discs to $N_d = 4$. Now we have multiple peaks. Considering the largest peak, we observe that increasing r_p causes a decline in the amplitude of the maximum h. Considering the lengthscale of the clustering q, we see that for $r_p \leq 15$, there is an increase in the lengthscale of the clustering q. As r_p increases beyond 15, the lengthscale of the clustering starts to become stable, or might even slightly decrease.

To test Besag's function, we have used initial controlled conditions where the points are clustered in clusters of known radius, and it confirms that the $L(r)$ function can give us a good estimate of the radius. It is not perfect. Nevertheless, it provides us with a measure of the radius.

Chapter 3

Percolation and the Properties of the Lattice

3.1 Introduction

To begin with, we shall forgo introducing dynamics into our modelling. We expect the dynamics to depend on the properties of the network, because we are relying on one tree passing the disease to its neighbouring trees. Thus, we are going to analyse the network first. Networks are widely observed in the natural world and have been applied in countless real-world applications, with perhaps the most important of these being electric power grids and aeroplane routes. Percolation theory can be used to model networks in terms of the connections between components [76].

Percolation theory was first described by Flory and Stockmayer during World War II, when they used the theory to describe the formation of macro-molecules from the chemical bonds between smaller molecules [77]. This process provides the foundation for gelation in which the solidification of jelly results from the formation of a chemical bond which spans across the whole system. A report called 'Percolation processes I. Crystals and mazes', written by Broadbent and Hammersley in 1957 [78] was the first mathematical application of percolation theory. In this paper, they studied the probability that the centre of a porous rock gets wet when it is immersed in a bucket of water. They found that it totally depends on the stone's porosity. Since then, percolation theory has been widely used across many areas such as physics, biology and geophysics through the use and development of fractal and scaling properties. Recently, the theory has been used to model forest fires, oil fields and fractals [76].

Clarifying the role of the clusters as the occupational probability changes is the main

idea of percolation theory [76]. Percolation is said to be present in a system if there is a spanning network of open sites, typically named the infinite cluster. For example, in a coffee percolator, hot water runs from the top, percolates through the gaps within the ground beans to provide the caffeinated drinks that many find irresistible. During this instance, the spaces between the ground coffee beans are the open sites, and percolation must have taken place; otherwise, no water would have filtered through. In essence, the objective is to model disease propagation in a forest by employing a forest grid consisting of sites that are allocated as open (occupied by trees), acting as the medium of transport for the infection. For infection propagation to occur, there must be a connected path from the origin of the disease outwards, i.e., a spanning network. Moreover, understanding the theory involved in this will help in the study of the simulated disease propagation.

3.2 What is Percolation Theory?

The origins and ideas of percolation theory, along with a variety of its real-world applications were explored in the introduction to this chapter. In the previous chapter, we created the point distribution for both homogeneous and heterogeneous data. Likewise, we provided various examples of how these distributions (i.e. grids of occupied sites) are represented. Now, we will introduce the formation of the percolating system and its key properties. First, we will illustrate percolation in a homogeneous field. Then, we will move gradually to discuss percolation for various heterogeneity levels. We will start our illustration by studying the connection between the occupied sites.

3.2.1 Clustering and Connectivity

A cluster is defined as a group of two or more occupied sites that are connected. The connectivity between occupied sites in the definition of the cluster varies from system to system. For instance, in the von Neumann neighbourhood (nearest neighbour) method of connection, only adjacent occupied sites are considered to be connected to an arbitrary site (i.e. a site is connected to sites that are to its north, south, west or east). However, for the Moore neighbourhood (i.e. next-nearest neighbour) method, the diagonal neighbours are also counted, so that a site can have eight neighbours at most. In Fig. 3.1, we display the two types of connectivity. The site under consideration is coloured in red, and the arrows show the location of the possible neighbours.

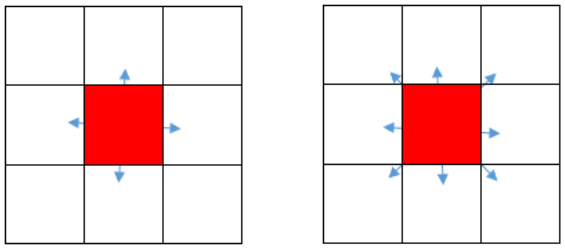

(a) (b)

Figure 3.1: Two types of connectivity are displayed in this figure. Panel (a) shows the von Neumann neighbourhood method of connection, while (b) illustrates the next-nearest neighbour (Moore neighbourhood) method. The site, which is coloured in red, represents an arbitrary site and the blue arrows indicate the location of its possible neighbours.

3.2.2 Percolation on a Homogeneous Grid

Now, we are going to consider a basic example for a homogeneous distribution of points. We use three different values of density ρ for 10×10 lattice grids. Note that, if ρ is set to zero, this will imply that all the sites are unoccupied, while if it is set to one, then this would imply that all the sites are occupied and that they form a single cluster. Identifying different clusters is the main concern regarding percolation theory, especially the biggest cluster which, if certain conditions are met, will dominate the entire system [76]. Thus, in this example, different clusters have been coloured with different colours. The aim is to develop an initial idea of the role of clustering as the density changes. Figure 3.2 displays three 10×10 homogeneous systems with different densities $0.35, 0.65$ and 0.75 in (a), (b) and (c) respectively. The method of connection assumed in this section is the von Neumann neighbourhood. For small values of ρ such as 0.35, the presence of clusters is rare and not guaranteed. When clusters appear in such a system, they are more likely to be very small and isolated by the unoccupied sites. As the density of the occupied sites of the lattice becomes higher, such as $\rho = 0.65$ and $\rho = 0.75$, the appearance of the biggest cluster is more obvious. In Fig. 3.2 (b), the biggest cluster spans the system from left to right or vice versa. In Fig. 3.2 (c), the biggest cluster spans the system from left to right and from bottom to top. This indicates that for the homogeneous data, as ρ increases, the biggest cluster becomes more dominant. When a single cluster of connected points spans the entire system, it is referred to as a spanning cluster, with the density of points at which such a cluster arises referred to as the critical density [76].

(a) (b) (c)

Figure 3.2: The coloured representation of a 2D percolating system represents each cluster with a different colour for increasing value of the occupational probability $\rho = 0.35$ (a), 0.65 (b) and 0.75 (c).

We now move to larger grids of 100×100 grid points. For $\rho = 0.59$, in Fig. 3.3 (a), we see that the biggest cluster starts to stretch across the system, taking on a dispersal structure. For $\rho = 0.60$, in Fig. 3.3 (b), a dramatic change has occurred to the biggest cluster. It now dominates and spans across the entire system with a sparse structure. We note that a small change in the density of a homogeneous grid can result in a significant difference.

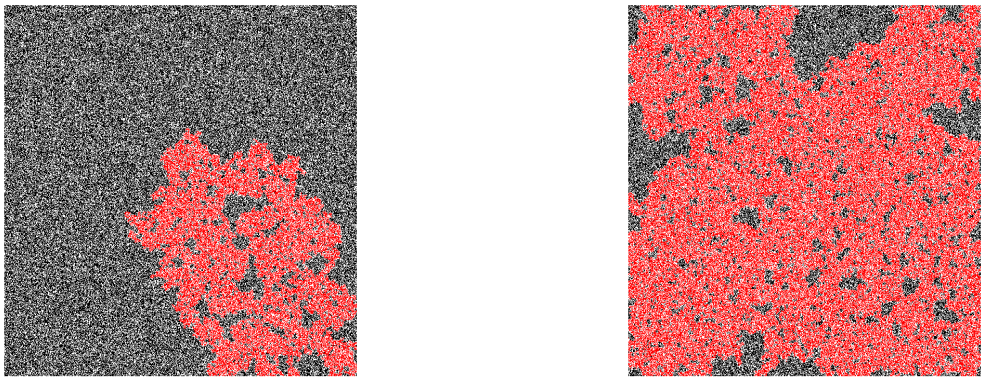

(a) (b)

Figure 3.3: Unoccupied sites coloured in white and occupied sites coloured in black with the biggest cluster highlighted in red. Panel (a) has $\rho = 0.59$ while panel (b) has $\rho = 0.60$.

Considering a 2D system, a spanning cluster arises when the cluster spans from left to right, top to bottom or both. The significance of the spanning cluster is not limited to extracting the principle information concerning the percolating system. Actually, it has equal importance in many scientific applications as regards percolation theory. It should

be mentioned that in physics, the spanning cluster is associated with phase transitions, which can be observed throughout nature. For instance, the solidification of jelly transpires after a spanning cluster first appears in the system; i.e. it is the first point at which the gelatine proteins begin to interlace, resulting in solidification. Other examples of this include the transition between insulators and conductors, and analogously modelling the spread of forest trees to produce prevention methods. The significance of the spanning cluster is further supported by the fractal properties it displays close to its formation. The characteristics of fractal behaviour are demonstrated by the occurrence of 'cracks' and 'holes' within the cluster; however, it is essential to demonstrate the existence of these in all length scales.

3.3 Percolation Threshold

First, we will define the percolation threshold ρ_c, followed by determining it for both the homogeneous and heterogeneous distributions.

3.3.1 What is The Percolation Threshold?

The percolation threshold refers to the emergence of long-range connectivity, i.e. a spanning cluster, in a network, at the critical density ρ_c (or more generally, occupation probability) of the sites in the network at which such a spanning cluster emerges [76]. The formal definition is for the emergence of this phenomenon in an infinite system, however, except for a very limited class of systems an exact value is not known, but can be estimated by numerical simulation [79].

Note that each type of system has a different critical density. These quantitative values can only be altered by two factors: how sites are connected (informally, the shape of the system), and the dimension of the system. The shape of the system has a significant impact on ρ_c. This is due to the connectivity of occupied sites being heavily dependent on ρ and the number of connections available. By changing the shape of the system, the connectivity also changes. For instance, using a triangular lattice, there are six possible neighbours to connect with. However, using von Neumann neighbourhood in the 2D square lattice, there are four possible neighbours to connect with. When the number of possible connections changes, the connectivity between the occupied sites alters. Hence the ρ_c value changes. The other factor is the dimension of the lattice. Changing this factor has an impact on the connectivity between sites, primarily through the possible number of sites to form a cluster with.

Table 3.1 displays different values of ρ_c for a square lattice. Here, the dimension, d,

is fixed and we change the configuration of neighbours which are considered connected. Figure 3.4 shows the basic connections (nearest neighbour N^2, next-nearest neighbour N^3, etc.) which are combined to make the neighbourhoods; for example, N^2 is the von Neumann neighbourhood, and $N^2 + N^3$ gives the Moore neighbourhood. We see that as the number of connections, and the distance at which points on the lattice are connected, increases, the critical density significantly decreases.

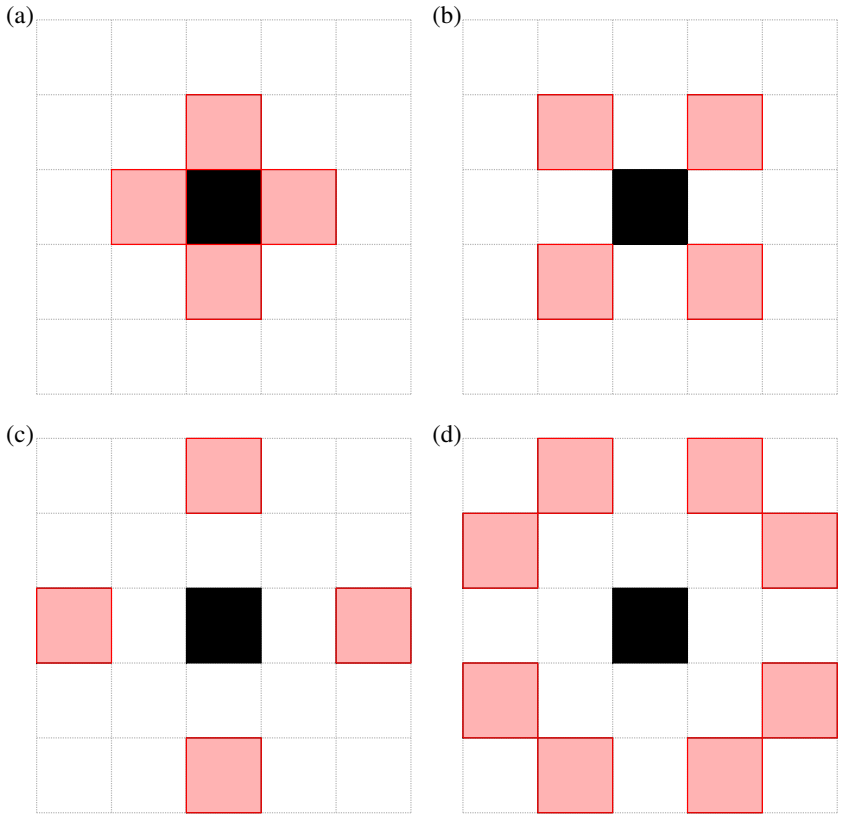

Figure 3.4: Black cell represnts the cell under consideration while red cells represents its neighbourhood. Different neighbourhood: (a) N^2, (b) N^3, (c) N^4 and (d) N^5.

Type of Lattice	neighbourhood	z	d	ρ_c
Square	N^2	4	2	0.593
Square	N^2 + N^3	8	2	0.407
Square	N^2 + N^4	8	2	0.337
Square	N^2 + N^3+ N^4	12	2	0.288
Square	N^2 + N^5	12	2	0.234

Table 3.1: The value of the percolation threshold, ρ_c, for various different lattices, highlighting their connectivity, z, and dimensions, d [80].

In the coming sections, we will estimate the ρ_c values for a range of heterogeneity levels. Very interesting values are expected for this complicated distribution because as γ increases, points tend to become more clumped together. For simplicity, we will start by determining this critical value for a homogenous field followed by approximating it for heterogeneous fields.

3.3.2 Determining the Value of the Percolation Threshold

To determine the value of the percolation threshold, we must define a new variable named $\mathcal{L}_{\max}$. This variable calculates the maximum linear size (height or width) of the biggest cluster in the system at any desired ρ. Then we normalise the maximum linear size of the biggest cluster by the lattice size $\mathcal{L}$, and we obtain $\Lambda = \mathcal{L}_{\max}/\mathcal{L}$. Note that Λ is calculated at each increasing increment of the value of ρ. In Fig. 3.5, the values of Λ are recorded for a single realisation at each density of a homogeneous grid, and averaged over 100 ensembles to decrease uncertainty, in (a) and (b) respectively. These calculations give a better understanding of how Λ changes.

In analysing Fig. 3.5(a), we see that for values of ρ that are much less than the percolation threshold, Λ increases only gradually with ρ, with very little noise. Around $\rho = 0.5$, however, Λ increases rapidly until it reaches unity, the maximum size of the system.

To see how much the behaviour alters from one ensemble to another, we calculate the standard deviation over the 100 ensembles shown in Fig. 3.5(b), with the results plotted in Fig. 3.6. We see a little deviation for small values of ρ; however, when ρ_c is approached from the left, a sudden increase is recorded. The sudden increase in the deviation reaches its maximum value just before the value of ρ_c is reached. This can be explained by the tenuous structure of the spanning cluster and the random process that generates the system on each iteration. These two properties are responsible for the production of the fluctuations in the data through the alteration of the size and the number of the produced clusters. As Λ

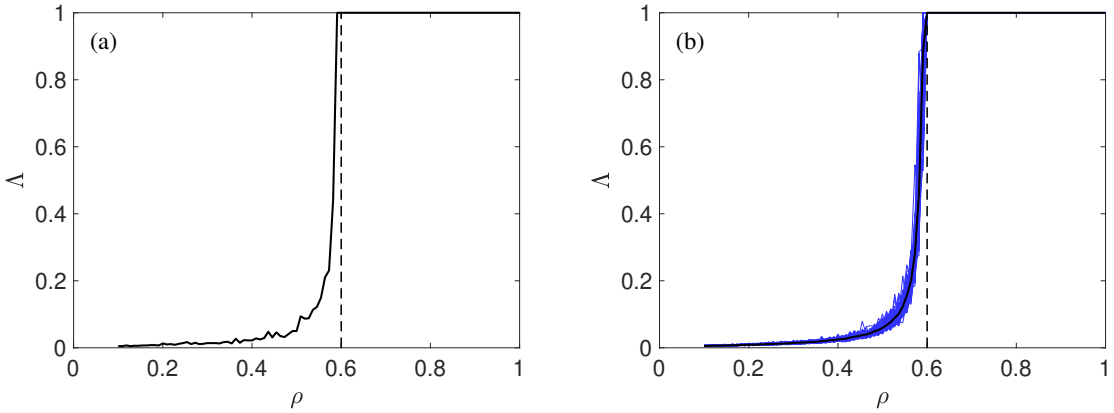

Figure 3.5: The normalised linear size of the largest cluster Λ plotted against ρ for one iteration in (a) and averaged over 100 ensembles in (b). The value of ρ_c is shown by a black dotted line. The system size here is 1000×1000.

plateaus, the deviation quickly diminishes to 0 because the spanning cluster has formed. Note that our calculated value of ρ_c is 0.60.

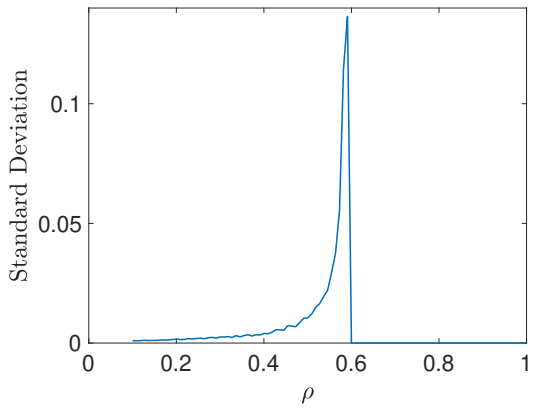

Figure 3.6: The standard deviation of Λ plotted against the value of the density ρ over 100 ensembles.

3.4 Generated Heterogeneous Data (Observed Tree Distribution)

The foundation of how we create the clustered distribution has been introduced in the previous chapter. Now we will illustrate the importance of the clustering parameter γ, the

density ρ, the seeding rate ρ_s and the lattice size $\mathcal{L}$. Previously, for the homogeneous data, we only discussed the importance of ρ and how it can change the connectivity between sites. Currently, it seems that most of these four parameters are critical. To see the effects of each single parameter, we will vary it and fix the other parameters. We will be examining how this affects the point distribution, and then apply Besag's function to show the associated clustering and segregation levels.

3.4.1 Effect of Varying Clustering Parameter

The clustering parameter is the way we vary heterogeneity. As γ increases, the level of heterogeneity is expected to increase. In Fig 3.7, we plot example point distributions for various γ while fixing the lattice size $\mathcal{L}$ to be 100, $\rho = 0.62$, and $\rho_s = 0.01$. For $\gamma = 0$, the generated point distribution is identical to the point patterns from the homogenous distribution. This is expected because the algorithm of generating the heterogeneous distribution with $\gamma = 0$ is equivalent to the algorithm of generating the homogeneous distribution. As γ increases, points become more clustered together. The gradual increase in the clustering level can be clearly seen as you go from panel (b) to panel (f) in Fig. 3.7. Similarly, we observe that the connectivity between occupied sites (points) for the heterogeneous field is stronger than the connectivity in the homogenous field. This is due to points favourably occupying sites with occupied neighbours in the generation algorithm.

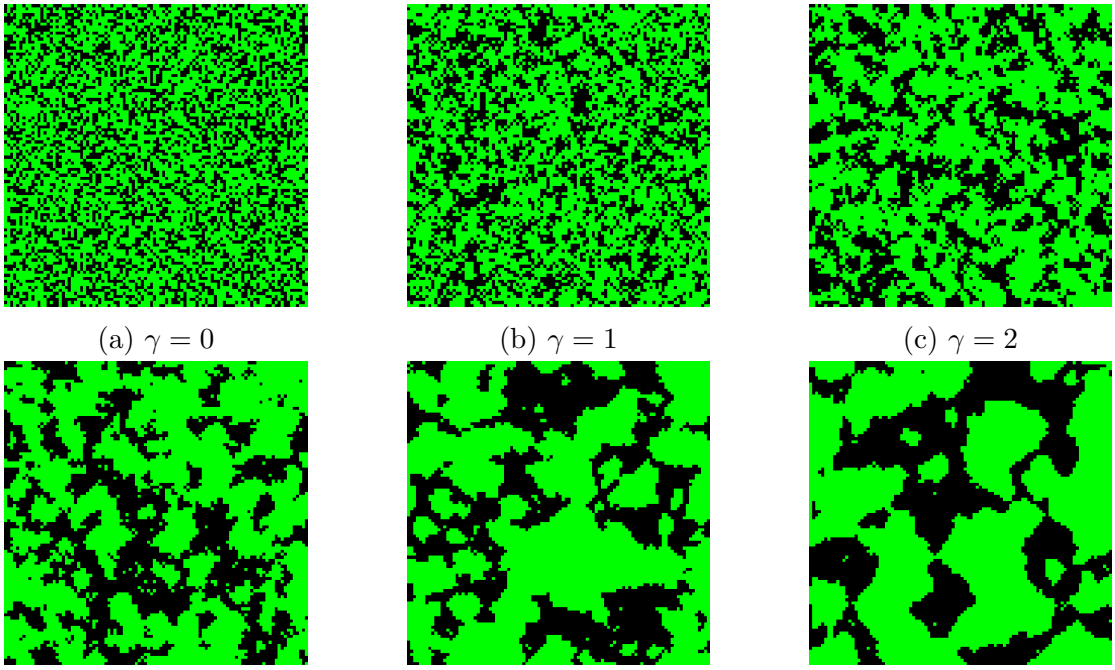

(a) $\gamma = 0$ (b) $\gamma = 1$ (c) $\gamma = 2$

(d) $\gamma = 3$ (e) $\gamma = 4$ (f) $\gamma = 5$

Figure 3.7: Generated clustered distributions for a range of values of the clustering parameter, from $\gamma = 0$ (homogeneous) to $\gamma = 5$ (producing large clusters), using the Moore neighbourhood method of connection. Green represents occupied sites, and black represents unoccupied sites. Here $\rho = 0.62$, $\mathcal{L} = 100$, and $\rho_s = 0.01$.

To quantitatively interpret the clustering, we will employ Besag's function $L(r)$, as described in Sec. 2.4. We will consider the clustering amplitude $h = \max(L(r))$, the height of the maximum value of $L(r)$, and the lengthscale of the clustering $q = \text{argmax}(L(r))$, the position of the maximum value of $L(r)$. In Fig. 3.8, we plot the mean of 30 ensembles of Besag's function $L(r)$ for a range of different γ. We note that when $\gamma = 0$, the dark blue curve slightly fluctuates around zero. This is expected because the point distribution for this particular γ is the same as that for the completely random (Poisson) distribution. As γ increases, the clustering amplitude increases, and the characteristic clustering lengthscale increases. To highlight these observations, h is plotted against γ in Fig. 3.8(b) and q is plotted against γ in Fig. 3.8(c).

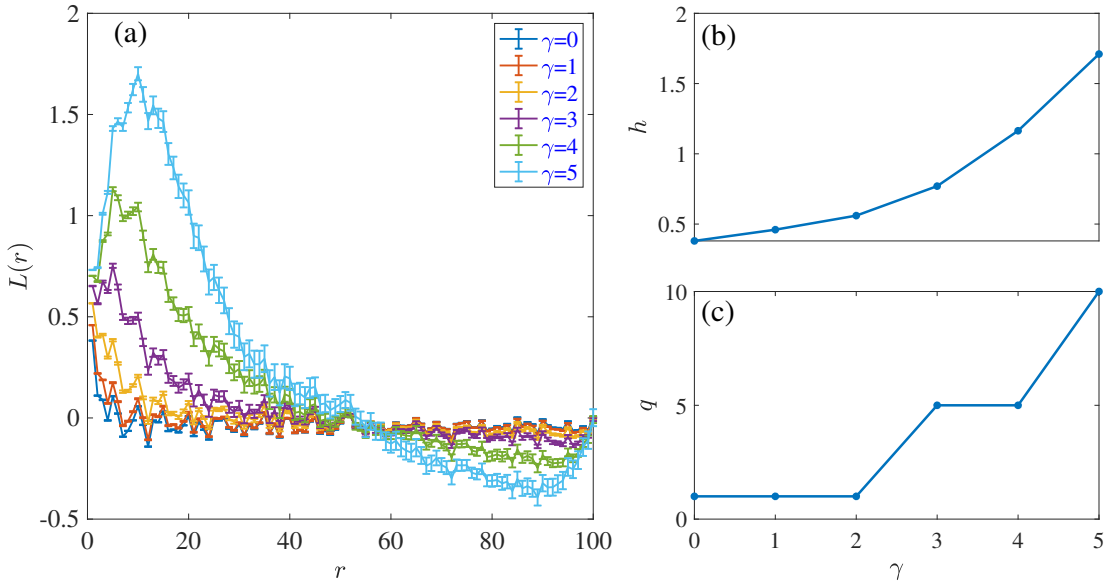

Figure 3.8: Panel (a): Besag's function for point distributions with different heterogeneity levels at different spatial scales r, for values of the clustering parameter γ between 0 and 5, averaged over 30 ensembles. The other parameters are fixed such that $\mathcal{L} = 100$ and $\rho = 0.5$. Panels (b) and (c) show the height and the position of the maximum of $L(r)$ respectively.

It is evident that these curves share similar features for different values of γ, in particular fluctuations across all scales and most obviously at small scales have the same structure. We believe that this is due to the discretisation of the grid, as the distances between points can only take particular values. To test this, we perturb these points with small amplitude noise. The idea is to add noise to the positions of the points, prior to calculating the $L(r)$ function. The separation between points on the grid is 1 unit. We added noise with amplitude 0.05 to the points first, and then we made a plot of $L(r)$. Then we repeated the procedure for amplitude 0.1 and then 0.2 - much smaller than the lengthscales we are interested in. By slightly randomising the positions away from the square grid, we expect to see that the structured noise arising from the discretisation gets blurred out. Studying Fig 3.9 we see that the curves no longer exhibit these structured fluctuations, while the general shapes of them are not affected. This is a clear indication that the discreteness of the grid causes the similar fluctuations.

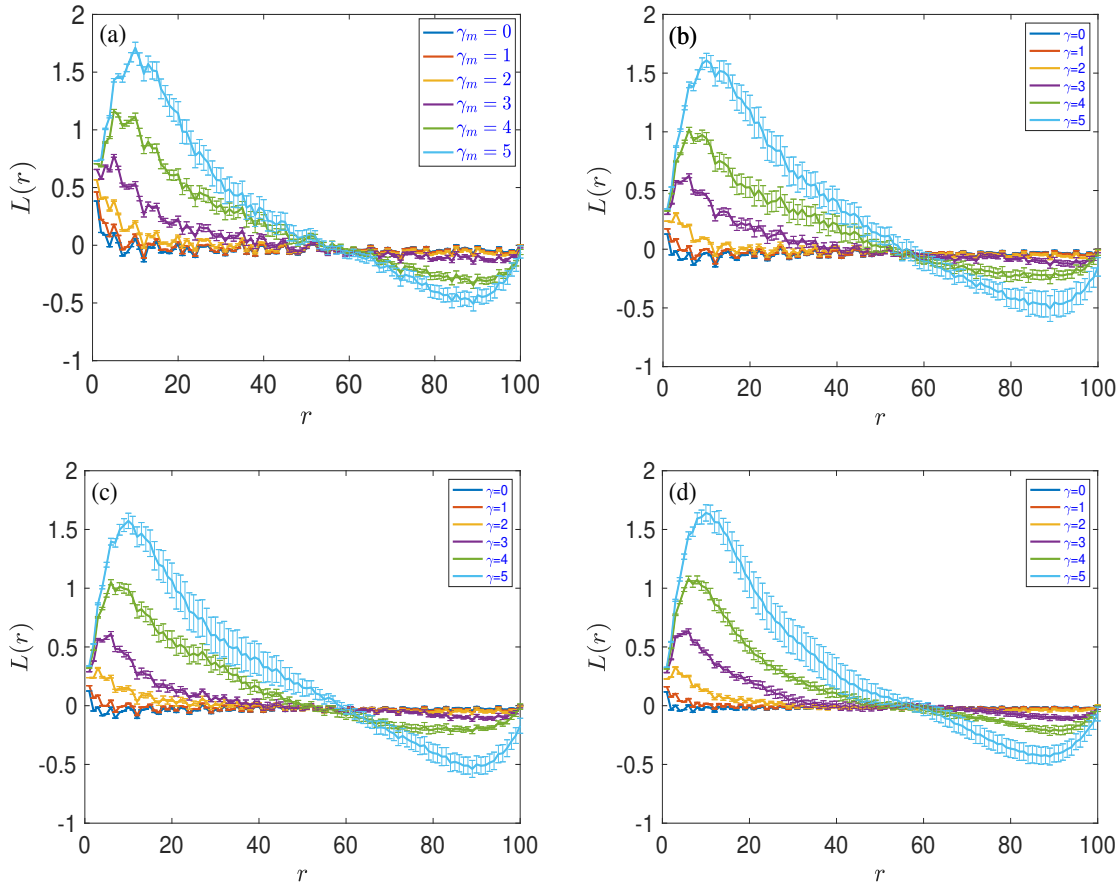

Figure 3.9: Panel (a): $L(r)$ calculated as for Fig. 3.8. Panels (b), (c), (d): $L(r)$ calculated after adding small amplitude noise to the positions of the points, with amplitudes 0.05, 0.1, and 0.2 respectively. The results are averaged over 30 ensembles.

3.4.2 Effect of Varying Seeds Rate

When we produce our clustered distributions by the clustering algorithm, we start with an empty grid. Then, we seed our grid with a percentage of the desired number of points. This percentage is called the seed rate and is donated by ρ_s. Our goal with the use of seed points is to decrease the computation time needed to generate these distributions, as when we start with an empty field it can take a long time, especially for large values of γ, for the first cells to become populated and act as nucleation sites for the clusters. However, in doing so, we do not want to strongly influence the characteristics of the resultant distributions, as introducing a second parameter controlling the clustering needlessly complicates later analysis. Here we will explore the effect of ρ_s in our heterogeneous field. We will see how an increase in ρ_s affects the overall clustering. In Fig 3.10, we plot example point distributions

for various seed rates ρ_s while fixing the lattice size $\mathcal{L} = 100$, $\rho = 0.62$, and $\gamma = 5$. We characterise the spatial distribution of these using Besag's L function, shown in Fig. 3.11.

Looking at the point distributions, we see that the seed rate can strongly affect the point distributions. We see that very small seed rates show a high level of clustering. As the seed rate increases the level of clustering decreases (points are less clustered together). This is reflected in the strength of the clustering, $h = \max(L(r))$, which decreases as ρ_s increases. This is due to the fact that the increment in the number of seeds implies that more points are scattered randomly (i.e. in the same way we generate our homogeneous distributions). We note that the characteristic length scale of the clustering, $q = \text{argmax}(L(r))$, also decreases as ρ_s increases. For $\rho_s \geq 0.1$, the number of seeds is relatively high, so it becomes less likely for the points to cluster together. For $0 \leq \rho_s \leq 0.01$, the difference in the resultant distribution is negligible, but the amount of time to generate the final distribution increases significantly as ρ_s decreases. We found that using $\rho_s = 0.01$ is the best choice for the seed rate. It produces a reasonable and controllable level of clustering for an affordable amount of computational time.

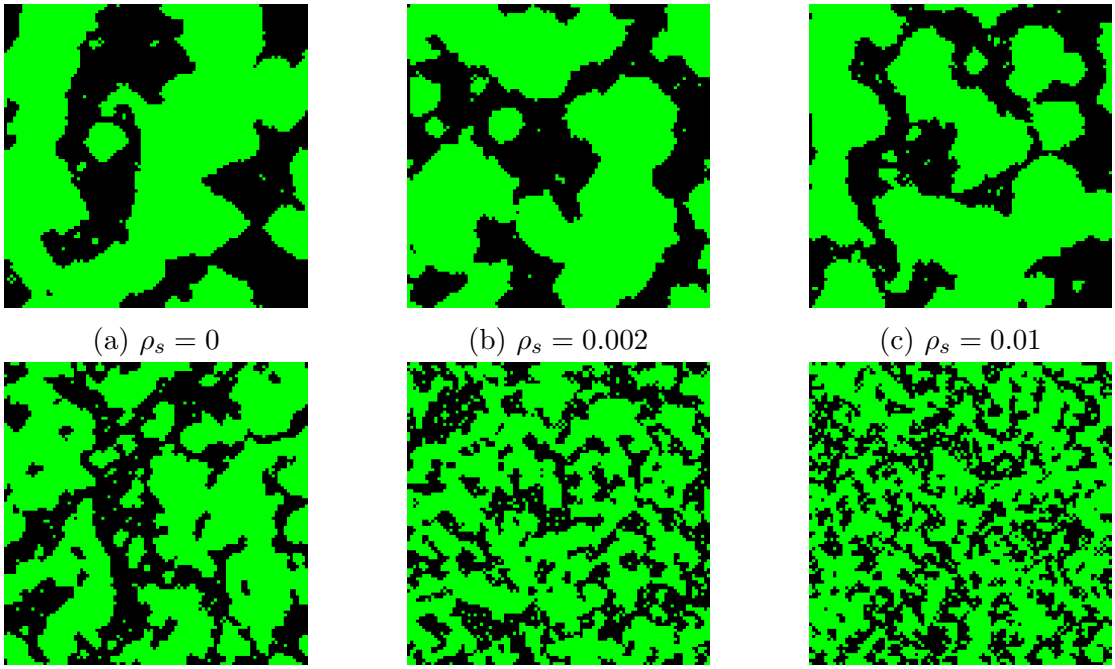

(a) $\rho_s = 0$ (b) $\rho_s = 0.002$ (c) $\rho_s = 0.01$

(d) $\rho_s = 0.1$ (e) $\rho_s = 0.3$ (f) $\rho_s = 0.5$

Figure 3.10: Heterogeneous distributions generated using our clustering algorithm, with increasing values of the seeds rate ρ_s, and other parameters fixed as $\rho = 0.62$, $\mathcal{L} = 100$ and $\gamma = 5$. Occupied sites are shown in green, with unoccupied sites in black.

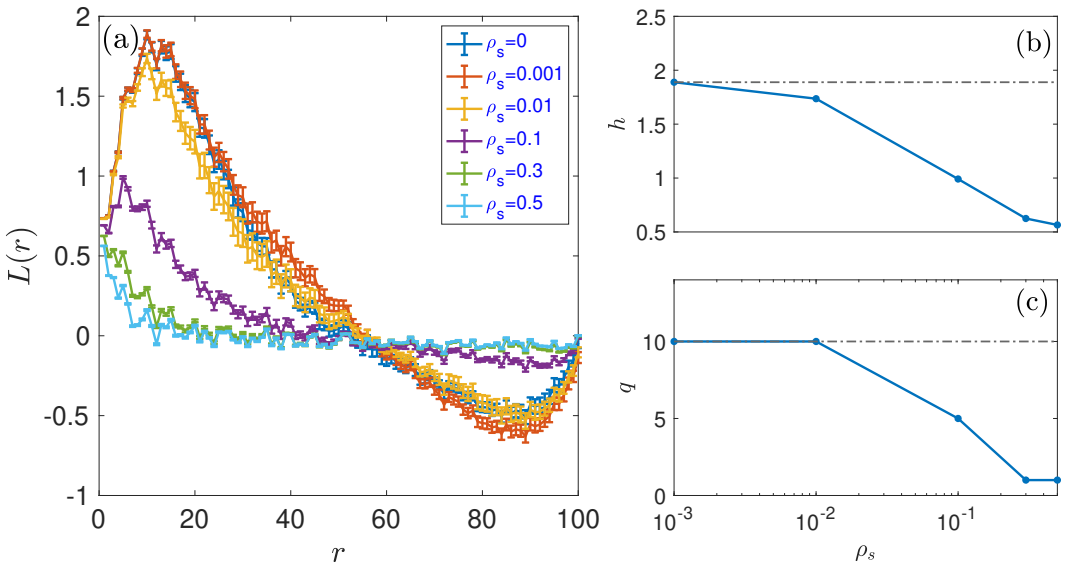

Figure 3.11: Panel (a), Besag's function for different seeds rate ρ_s, for $\mathcal{L}$=100, $\gamma = 5$ and $\rho = 0.5$, averaged over 30 ensembles. Panel (b), the maximum value of $L(r)$ plotted against ρ_s. Panel (c), the location of the maximum value of $L(r)$ plotted against ρ_s. The black dashed line represents $\rho_s = 0$.

3.4.3 Effect of Varying Density

In Fig 3.12, we plot example point distributions for various densities ρ while fixing the lattice size $\mathcal{L} = 100$, $\rho_s = 0.01$, and $\gamma = 5$. Looking at point distribution for each single density, we observe that an increase in ρ produces a larger population of points and decreases the number of isolated clusters. This is to be expected given that it is recognised that as density increases the connectivity between occupied sites also increases. Likewise, for this particular high clustering parameter, the connectivity is even stronger. These behaviours are going to be more obvious as you approach the critical density. Looking at these point distributions, the critical density for heterogeneous distributions will be different than homogeneous distributions. Studying the transition between $\rho = 0.7$ and $\rho = 0.85$, you can see that the number of isolated clusters for $\rho = 0.85$ is smaller than $\rho = 0.7$. This is due to the fact that as the density increase clusters start to merge together, forming a single large cluster (i.e. small cluster starts to become part of the spanning cluster).

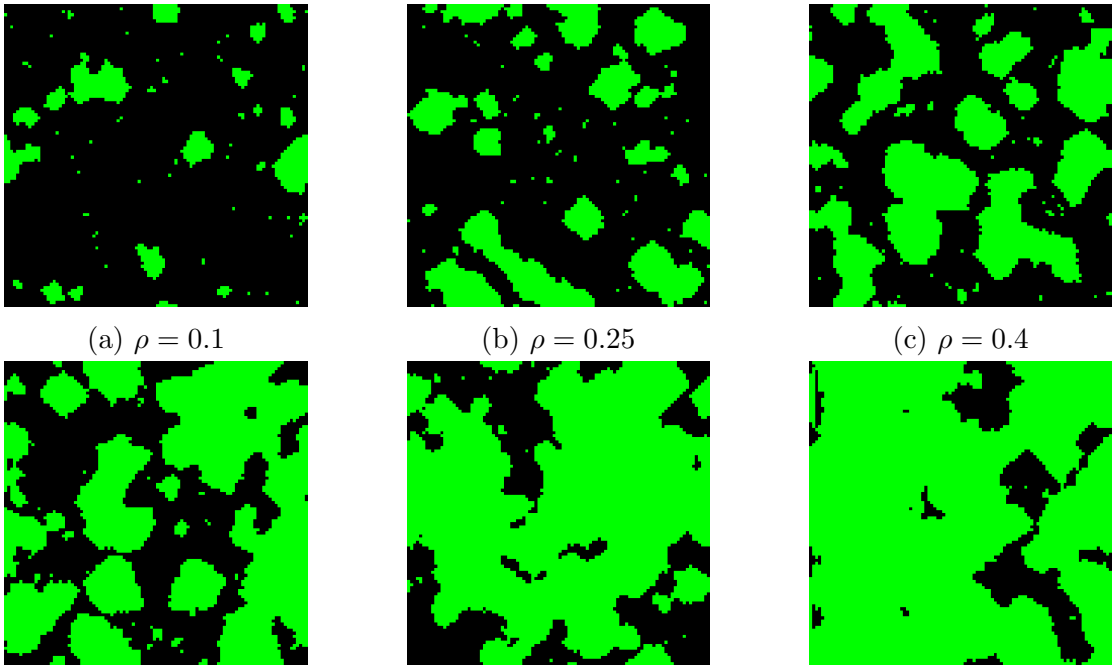

(a) $\rho = 0.1$ (b) $\rho = 0.25$ (c) $\rho = 0.4$

(d) $\rho = 0.55$ (e) $\rho = 0.7$ (f) $\rho = 0.85$

Figure 3.12: Heterogeneous distributions generated using our clustering algorithm, with $\mathcal{L} = 100$, $\gamma = 5$, and varying ρ. Occupied sites are shown in green, with unoccupied sites in black.

Figure 3.13(a) shows how the density of points affects the observed clustering amplitude and lengthscale. Here we vary ρ and fix everything else. We see that an increase in ρ causes a decrease in the strength of the clustering, which is displayed in Fig. 3.13(b). This is expected because as the density increases, the grid starts to get filled by points. On the other hand, the lengthscale of the clustering becomes constant, which is represented in Fig. 3.13(c).

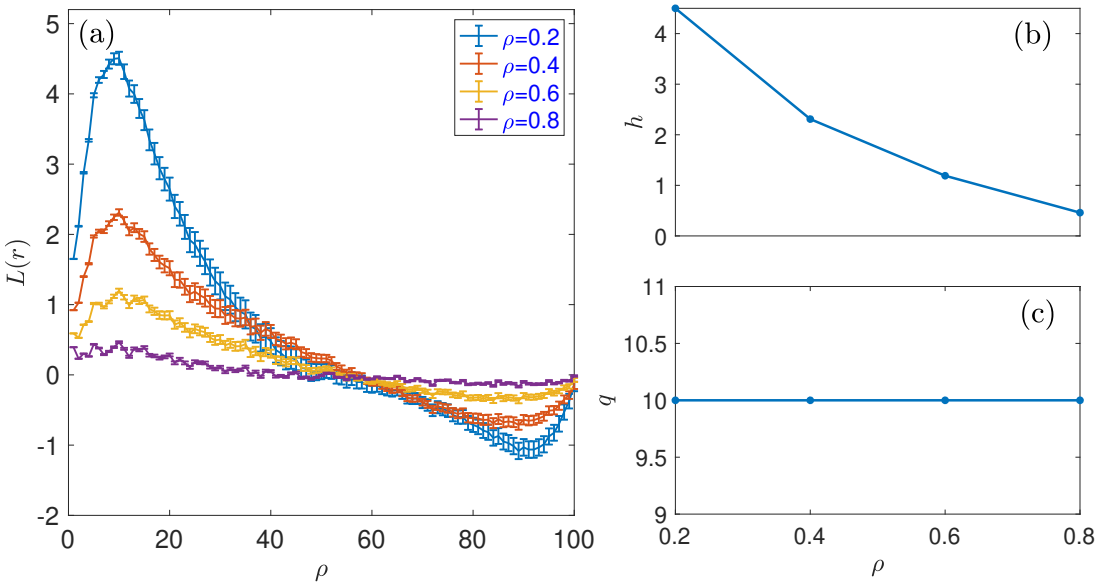

Figure 3.13: Panel (a): Besag's function for point distributions with different heterogeneity levels at different spatial scales r, for densities between 0.2 and 0.8, averaged over 30 ensembles. The other parameters are fixed such that $\mathcal{L} = 100$, $\gamma = 5$. Panels (b) and (c) show the height and the position of the maximum of $L(r)$ respectively.

3.4.4 Effect of Varying Grid Size

Now, we are interested in the effect of changing the grid size on the clustered distributions produced by the clustering algorithm. Figure 3.14 represents the point distributions for three different grid sizes $\mathcal{L}$=100, 200 and 400. We see that all the distributions have similar cluster size and strength. This shows the cluster size dictated purely by γ and not the grid size.

In Fig. 3.15, we plot Besag's function as a function of the lengthscale r for different lattice sizes $\mathcal{L}$. We observe that the lattice size has no effect in the strength of the clustering h =max($L(r)$) (i.e. the height of the maximum value for $L(r)$), and the lengthscale of the clustering q =argmax($L(r)$)(i.e. the position of the maximum).

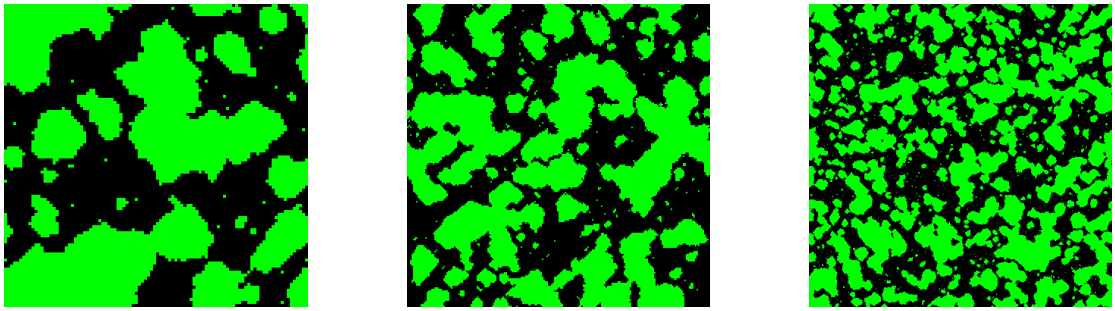

(a) $\mathcal{L} = 100$ (b) $\mathcal{L} = 200$ (c) $\mathcal{L} = 400$

Figure 3.14: Point distributions generated using our clustering algorithm, with $\rho = 0.5$, $\rho_s = 0.01$, $\gamma = 5$, and varying $\mathcal{L}$. Occupied sites are shown in green, with unoccupied sites in black.

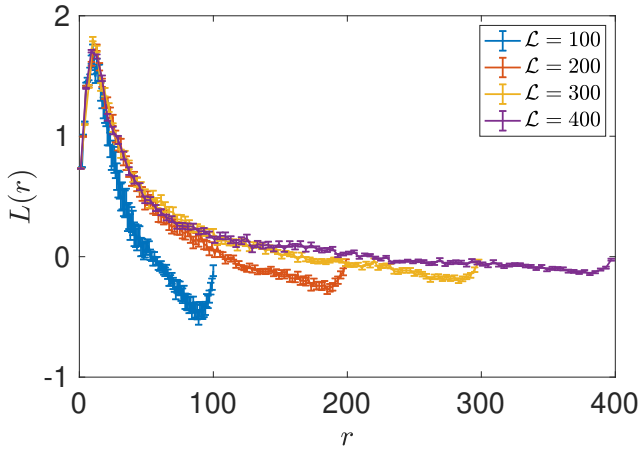

Figure 3.15: Besag's function for $\gamma = 5$, $\rho = 0.5$, $\rho_s = 0.01$ and different grid sizes. The results are averaged over 30 ensembles.

3.5 Percolation Threshold for the Heterogeneous Data

Note that only in this section and the rest of this chapter, values of γ generated based on the von Neumann neighbourhood rules are denoted by γ_v while values of γ based on the Moore neighbourhood rules are denoted by γ_m.

3.5.1 Von Neumann Neighbourhood

The known percolation threshold ρ_c for the von Neumann neighbourhood is 0.592746. We will now consider the normalised size of the largest cluster, Λ, as we did for the homogeneous distribution in Sec. 3.3.2. In Fig. 3.16, we display two examples for Λ calculated as a function of the density for two different γ_v. This shows how ρ_c is obtained for different γ_v. For $\gamma_v = 0$, the value of the percolation threshold is 0.6, which is approximately the same as the value of ρ_c from percolation theory. As γ_v increases, the value of the percolation threshold changes. We extend this to other values of γ_v.

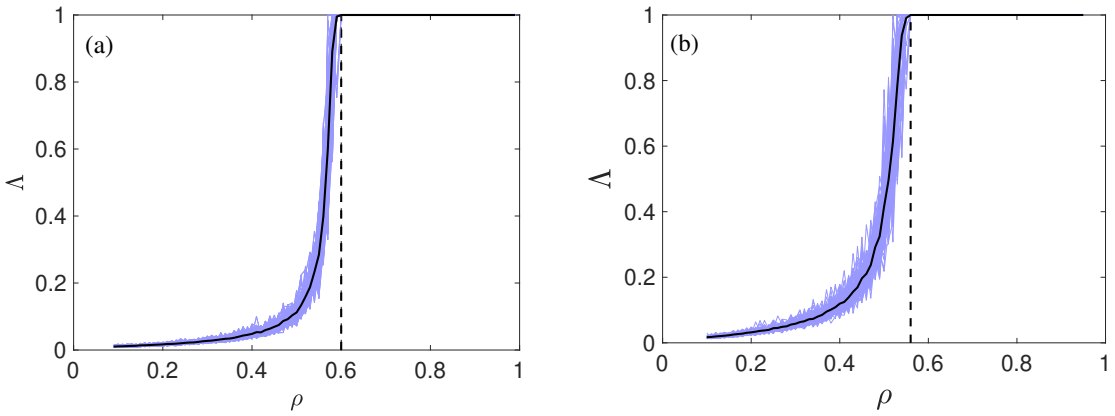

Figure 3.16: The normalised linear size of the largest cluster Λ as a function of ρ. Panel (a), for a homogeneous grid (i.e. $\gamma_v = 0$). Panel (b) for a grid with a clustered distribution of points with $\gamma_v = 2$.

In Table 3.2, we present ρ_c for the different heterogeneity levels. We use a 500×500 system. The clustering parameter has a considerable effect on the ρ_c values. This is because as γ_v increases, points tend to cluster together, causing the system to percolate either much faster or delay the percolation to higher densities. We observe that as γ_v increases from zero until three, clustering causes the spanning cluster to appear in the system at lower densities. As γ_v increases beyond three, the spanning cluster starts to occur at larger densities. We have done our investigations over 50 ensembles.

γ_v	0	1	2	3	4	5
ρ_c	0.60	0.57	0.56	0.56	0.57	0.61

Table 3.2: The calculated ρ_c values for the introduced clustering method for the von Neumann neighbourhood method.

3.5.2 Moore Neighborhood

For the Moore neighbourhood method of connectivity, the known value of ρ_c for entirely random data is 0.407 [80]. From Table 3.3, we observe that the calculated ρ_c values for different γ_m for this method of connection are higher than the known value of ρ_c for the Moore neighbourhood. Apparently, ρ_c values here are different than the calculated ρ_c values for the von Neumann neighbourhood. Evidently, both of these methods of connectivity have values of ρ_c that are changing quantitatively as the value of the clustering parameter changes.

γ_m	0	1	2	3	4	5
ρ_c	0.41	0.42	0.48	0.55	0.63	0.68

Table 3.3: The calculated ρ_c values for a range of values of γ_m. The method of connectivity used here is the next-nearest neighbour (i.e. the Moore neighbourhood) method.

3.6 Relating Von Neumann to Moore Neighbourhood

Although the critical densities are clearly different between systems with Moore and von Neumann connections, it is natural to ask whether they may produce similar levels of heterogeneity. To relate the von Neumann to the Moore neighbourhood, we need to compare $L(r)$ values for the two methods of connections using various values of γ. We need to find values of γ_v from the von Neumann neighbourhood that generate the same $L(r)$ curves as some other values of γ_m from the Moore neighbourhood. Then, we will generate the associated points distribution for these values of γ_v and γ_m to make sure that they produce the same cluster sizes.

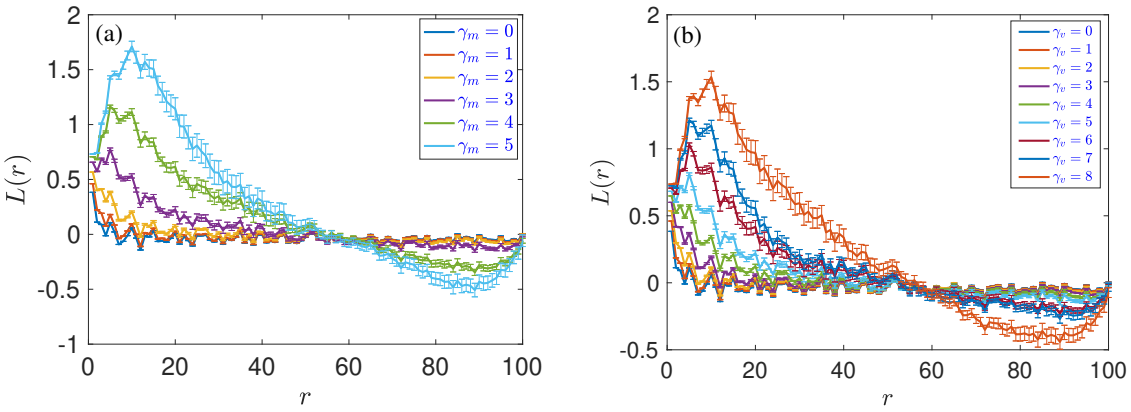

Figure 3.17: Besag's function $L(r)$ is calculated for the generated heterogeneous distribution with the lattice size $\mathcal{L} = 100$ and $\rho = 0.5$ as a function of the radius r (i.e. the spatial lengthscale) and the clustering parameter γ. The results are averaged over 30 ensembles and the error bars are attached to each curve. The method of connection used here is the Moore neighbourhood (panel(a)) and the von Neumann neighbourhood (panel (b)).

In Fig. 3.17, we plot the average mean of 30 ensembles of $L(r)$ values for a range of different γ for von Neumann(panel (a)) and Moore (panel (b)). Note that the $L(r)$ curve for $\gamma_v = 8$ form the von Neumann neighbourhood is the same as the $L(r)$ curve for $\gamma_m = 5$ form the Moore neighbourhood. This means that with more investigation, we might be able to relate both models.

3.6.1 Statistical Approximations

We are interested in approximating a value of γ_v from the von Neumann neighbourhood that is equivalent to a value of γ_m from the Moore neighbourhood. When we say equivalent, we mean that they have the same $L(r)$ values and produce similar point distributions.

Note that s_v and s_m refer to the probability of a site being occupied according to von Neumann and Moore neighbourhoods respectively, and N_v and N_m refer to the number of neighbours for the von Neumann and Moore neighbourhoods respectively. Then:

$$s_v = (N_v/5)^{\gamma_v},$$

where $N_v \in [1, 5]$,

$$s_m = (N_m/9)^{\gamma_m},$$

where $N_m \in [1, 9]$,

$$(N_v/5)^{\gamma_v} = (N_m/9)^{\gamma_m}$$
$$\log((N_v/5)^{\gamma_v}) = \log((N_m/9)^{\gamma_m})$$
$$\gamma_v \ \log(N_v/5) = \gamma_m \ \log(N_m/9)$$

Therefore;

$$\gamma_v = \left(\frac{\log(N_m) - \log(9)}{\log(N_v) - \log(5)}\right) \gamma_m,$$

where N_v does not equal 5.

Calculating γ_v for all the possible N_v and N_m values, and then taking the mean, we found out that $\gamma_v = 1.5790\gamma_m$. We will now test whether the relation $\gamma_v = 1.5790\gamma_m$ is sufficient to lead to similar point distributions. We test this visually in Fig 3.18. We used

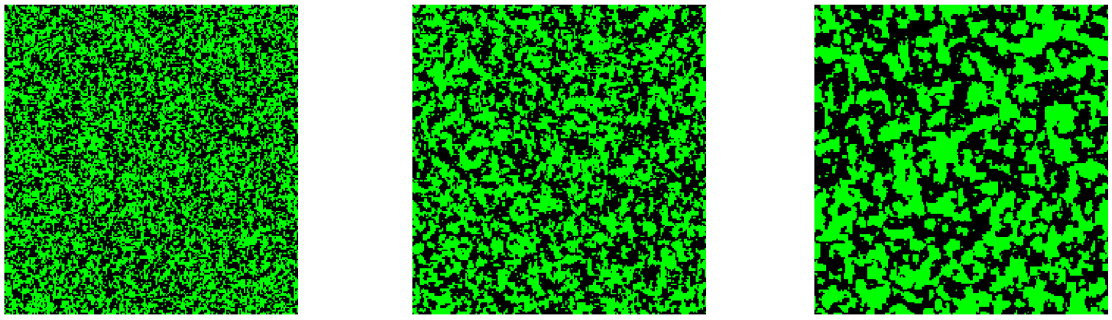

(a) $\gamma_v = 1.579$ (b) $\gamma_v = 3.158$ (c) $\gamma_v = 4.737$

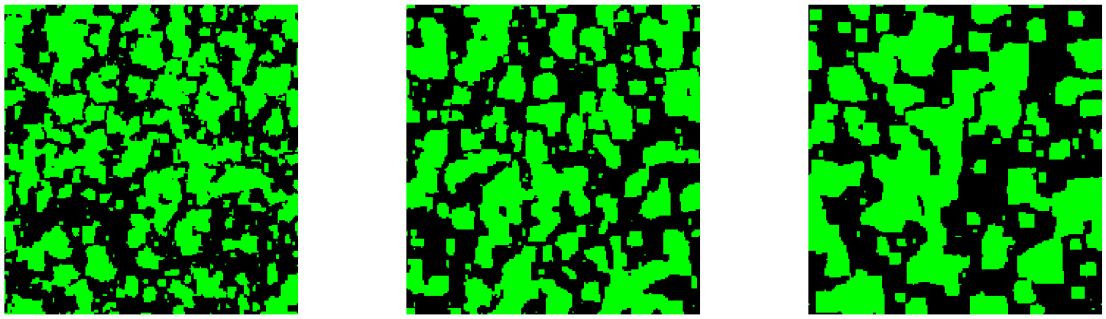

(d) $\gamma_v = 6.316$ (e) $\gamma_v = 7.895$ (f) $\gamma_v = 9.474$

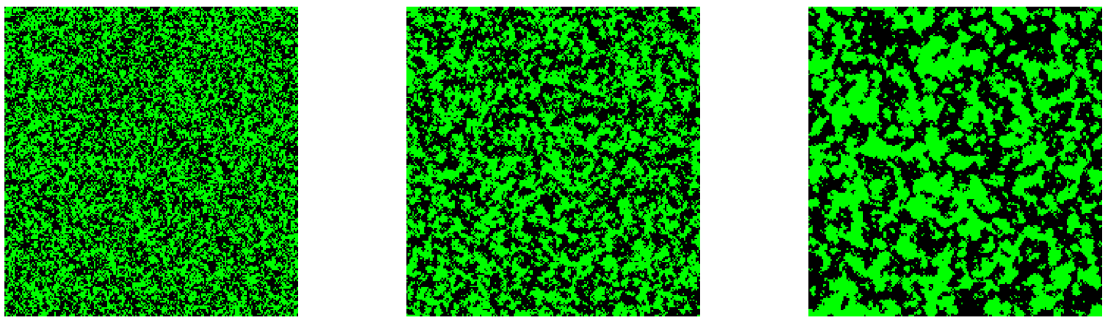

(g) $\gamma_m = 1$ (h) $\gamma_m = 2$ (i) $\gamma_m = 3$

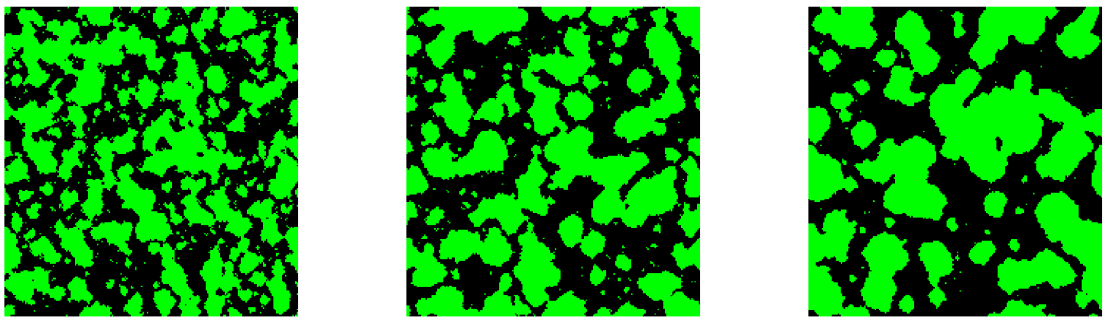

(j) $\gamma_m = 4$ (k) $\gamma_m = 5$ (l) $\gamma_m = 6$

Figure 3.18: Comparison of the clustered distribtions produced by our clustering algorithm when using von Neumann connectivity (top two rows) and Moore connectivity (bottom two rows), with $\mathcal{L} = 200$.

von Neumann in the top (panels(a-f)) and Moore in the bottom (panels (g-l)). The two methods have similar distributions. We see that these equivalent γ produces similar point distributions no matter what the shapes of the clusters (circles or squared). On the other hand, they have the same $L(r)$ values, see Fig 3.19. Since point distributions are similar, $L(r)$ values are expected to be similar.

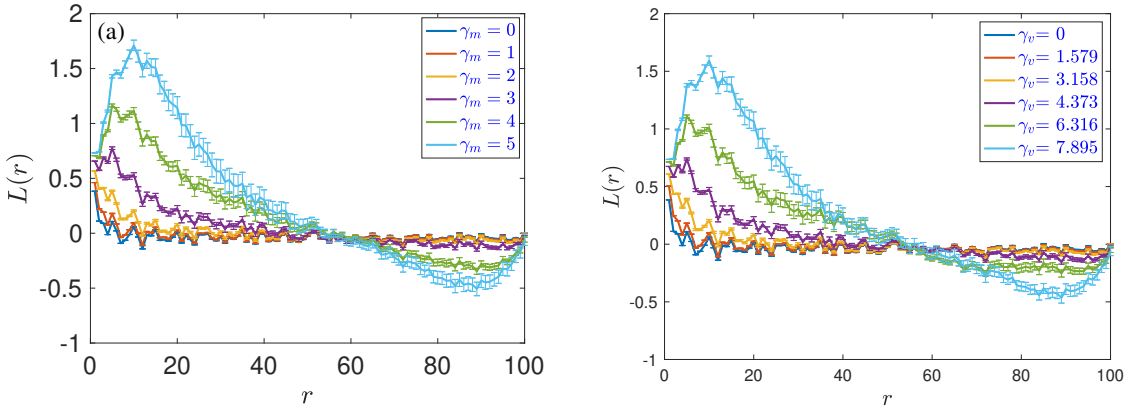

Figure 3.19: Besag's function $L(r)$ is calculated for the generated clustered distribution with the lattice size $\mathcal{L} = 100$ and density $\rho = 0.5$ as a function of the radius r (i.e. the spatial lengthscale) and the clustering parameter γ. The results are averaged over 30 ensembles, and the error bars are attached to each curve. The methods of connection used here are the Moore neighbourhood (a) and the von Neumann neighbourhood (b).

In Fig. 3.20, we repeated Fig. 3.19 for $\rho = 0.1$ and $\rho = 0.9$ for both von Neumann and Moore neighbourhood distributions to check if they still match. For $\rho = 0.9$ they still match. For $\rho = 0.1$, there is a slight difference in the highest value of γ. This might due to the density being low, and there are not enough points to produce clusters of the same size.

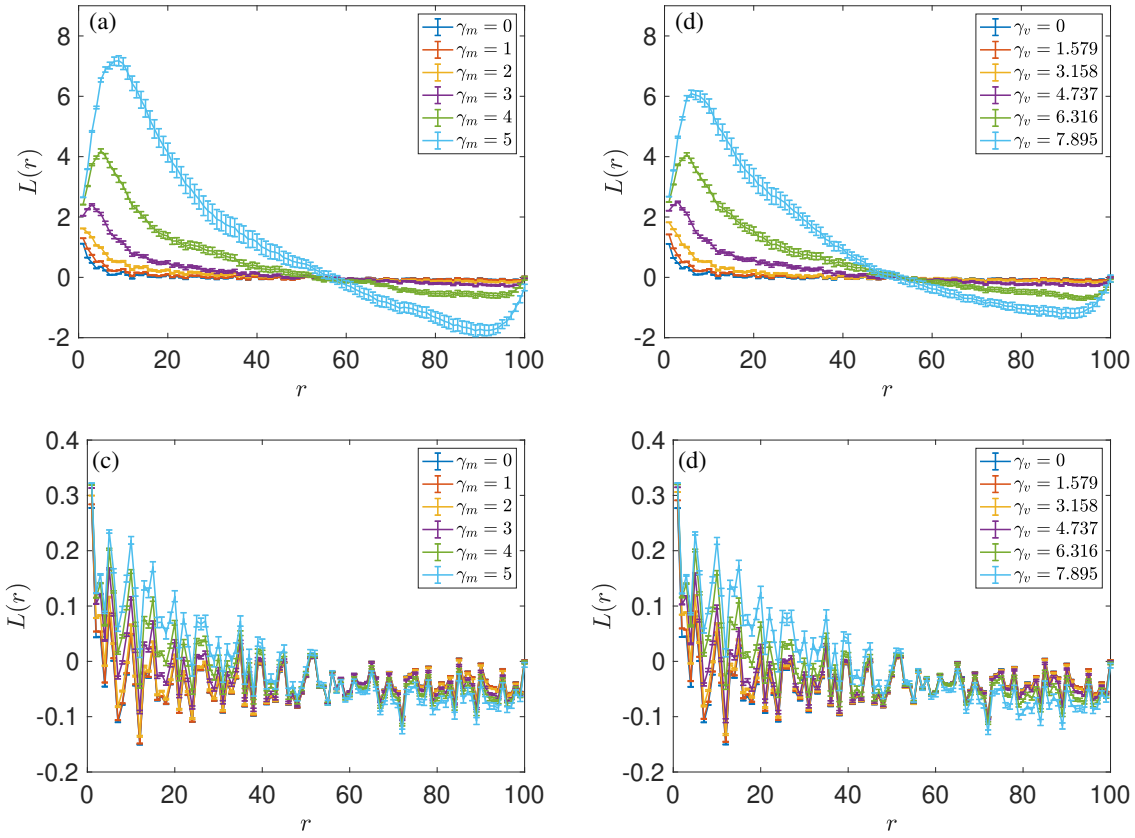

Figure 3.20: Besag's function $L(r)$ is calculated for the generated heterogeneous distribution with the lattice size $\mathcal{L} = 100$ as a function of the radius r (i.e. the spatial length scale) and the clustering parameter γ. The results are averaged over 30 ensembles, and the error bars are attached to each curve. The method of connection used here is the Moore neighbourhood (panel(a,c)) and the von Neumann neighbourhood (panel (b,d)). The density in the first row is 0.1 while in the second row is 0.9.

3.7 Conclusions

Prior to introducing dynamics into the clustered distribution produced by the clustering algorithm, we studied the network properties. We discussed the spanning cluster phenomenon and its first appearance in the system. We see that as heterogeneity increases, the occurrence of the spanning cluster occurs at different densities. Using the von Neumann neighbourhood as the method of connectivity between occupied sites, as we increase the clustering parameter γ_v from zero until three, the spanning cluster appears in the

system at lower densities. However, as γ_v increases beyond three, the spanning cluster starts to require a greater density. Using the Moore neighbourhood, we observed that as γ_m increases, the spanning cluster requires a larger density in order to first appear in the system. Note that in percolation threshold for Moore and von Neumann neighbourhood, even if we consider identical distributions, a distribution that contains a spanning cluster for the Moore neighbourhood may not do so for the von Neumann neighbourhood, if some connections are on diagonals.

Using the Moore neighbourhood to generate the clustered distribution, we examined the effect of the key parameters: γ_m, seed rate ρ_s, density ρ, and the grid size $\mathcal{L}$. First, we visualise the point distributions. Subsequently, through Besag's function, we characterised the observed spatial clustering and segregation of the generated point distributions. We established that as γ_m increases, the level of clustering increases. Conversely, as ρ_s increases, the level of clustering decreases. This is due to the fact that as ρ_s increases, more points are scattered randomly (i.e., in the same way, we generate our homogeneous distributions). We found that using $\rho_s = 0.01$ is the most appropriate choice because it produces a reasonable and controllable level of clustering for an affordable amount of computational time. The density ρ is another important parameter. As ρ increases, the population of points becomes larger. For all the different heterogeneity levels, as ρ increases, the spanning cluster starts to appear in the system. As γ_m increases, the appearance of the spanning cluster requires a larger ρ. Considering the grid size $\mathcal{L}$, it has no effect on the cluster size for the reason that the cluster size is dictated by γ_m.

The critical density, which is the density at which the spanning cluster first appears in the system, is different between systems with Moore and von Neumann connections. However, looking at the point distribution from the grids generated using the von Neumann neighbourhood and the Moore neighbourhood rules, although the shapes of the generated clusters are different, their sizes could be made the same. Comparing the produced point distributions and the associated Besag's function curves, we found that $\gamma_v = 1.5790\gamma_m$, where γ_v and γ_m are the clustering parameters to generate the point distributions from the von Neumann and the Moore neighbourhoods, respectively. However, in trying to link the two methods of connectivity in this way, we have only considered the cluster sizes. When we consider the percolation threshold for the two types of connection, we observe a monotonic increase with γ_m for the Moore neighbourhood, but not with the von Neumann neighbourhood; it may be that the differences between the two geometries are fundamentally irreconcilable by simply scaling the clustering parameter.

Chapter 4

Modelling the Infection Spreading in a Homogenous Forestscape

4.1 Introduction

In the beginning of this study, we have explored some of the high profile tree diseases and their ecological and economic impacts. We have also investigated different frameworks that can be utilised to model tree diseases. We decided to proceed with a computational approach called the lattice-based model, which has a stochastic spatial distribution and stochastic infection mechanism. Then, we outline the algorithmic process to generate our synthetic forestscapes. Note that in this chapter, we are only using homogeneous forestscapes. Moreover, we have understood the bulk properties of the lattice and the connections between the occupied sites. Now we are going to put the disease on top and then analyse the observed dynamics. Due to the complexity of the interaction between the individual components, it is difficult to predict and control the disease propagation in forests.

In this chapter, we will develop the model used by Orozco-Fuentes *et al.* [47], reproducing their results but also extending the analysis. We will use a lattice-based model incorporated with a stochastic process that mimics the interactions between the host and the pathogen. We will investigate disease spreading in a forest by measuring the spreading velocity of the pathogen on a lattice of susceptible trees. We will then use their definition of the spreading velocity, and consider two other possible definitions, making a comparison between them. To successfully predict the critical shifts between the two stable phases: local confinement and outspread outbreak, early-warning indicators for impending regime shifts are used. This framework is a well-known dynamical system approach. Even though it is difficult to control and predict the spread of plant diseases, this approach could be

useful for forecasting forest disease epidemics [47]. Another quantity to measure the impact of the disease spread is mortality. We will define it and produce its phase diagram as we vary the forestscape density and transmission probability.

4.2 The Model

An artificial forestscape will be represented by a 2D square lattice, of size $\mathcal{L} \times \mathcal{L}$, where $\mathcal{L} = 500$. The model has only one type of vegetation, and the forestscape is flat. Each site in our artificial forest can be in one of four states: Susceptible (S), Infected (I), Removed (R) or Empty (ϕ). Susceptible cells represent the population of healthy individuals, infected cells correspond to the sites that acquire the pathogen, removed cells correspond to the sites that are left by infected cells after the vegetation has died and empty cells represent sites where no vegetation can grow.

The distribution of trees within a single ensemble is fixed so that the only changes to the configuration come from the disease spreading; we do not model either growth in empty cells or regrowth in cells removed by the disease, and we do not consider trees being removed by other causes. This is justified by modelling the infection on time-scales much shorter than the growth of the susceptible species. Furthermore, in real forests, the destruction of one species at a site is often followed by a different species colonising the site, which has been observed in the wake of fires in savanna ecosystems [56]. Because of these factors, cells are unlikely to regain susceptibility, allowing us to neglect regeneration. We are also justified in not modelling removal by other causes, in particular, a constant background mortality rate, as the background mortality rate due to old age is much lower than the mortality rate during an epidemic. So it is reasonable to neglect it.

At the beginning of the simulation, the pathogen is initialised as a group of infected cells at the centre of our artificial forestscapes. This clump of infected cells is chosen to be a squared grid of size $l{=}\mathcal{L}/100$, which has been shown to be sufficient to avoid disease extinction at an early stage. Note that the nearest neighbour (von Neumann neighbourhood) method of infection is implemented. Therefore, the infection can spread to susceptible trees positioned to its north, south, east or west. Under these initial conditions the transient time, i.e., the time-lapse that contains remnants of the initial conditions, was found to be around 10 time steps, which is enough for the initially infected cells to die out. Since we are interested in the steady regime, we discard this transient from our final calculations. Note that the simulation runs until it either reaches one of the four boundaries, destroys all the susceptible trees, or dies out.

4.2.1 Parameters

Parameters have significant roles in this model. Essentially, they control the general behaviour of the system. Three parameters should be taken into account: tree density ρ, the transmissibility of the pathogen β and the infectious period T.

Tree Density

Tree density, ρ, has been thoroughly discussed in the methodology Chapter 2. In short, this parameter shows the average proportion of occupied sites.

For homogeneous forestscapes, individual cells are populated according to a Bernoulli trial with probability $p = \rho$, the desired overall density, where success means an occupied site and failure means an empty site. Taken as a whole, the density of the forestscape then comes from a Binomial $B(\mathcal{L}^2, \rho)$ distribution, with the expected number of successes (occupied sites) then $N_T = \mathcal{L}^2\rho$, to give an expected density of ρ. Although there is variation in the actual obtained value of ρ, this variation is small, with the standard deviation of ρ attaining a maximum value of 10^{-3} at $\rho = 0.5$. The occupancy level of cells is thus entirely dependent on ρ.

Infectious Period

The infectious period T, sometimes called the length of the infection, represents the time from which a cell acquires the infection to the time at which the cell dies and enters the R state. In this period, a cell can transmit the disease to its neighbouring cells. Following the work of Orozco-Fuentes *et al.* [47], we take $T = 10$.

When a susceptible cell acquires the infection, it becomes infectious for 10 time steps. We attach a variable τ to the newly infected cell and assign it the value $\tau = 1$. We then increment τ by 1 at each timestep, until $\tau = T$. In the subsequent timestep, the infected cell becomes a removed cell, and can no longer transmit the disease.

Figure 4.1 shows the configuration of a forestscape at the end of an example simulation, which here finished when the infection reached one of the boundaries. To visualise our simulation, infected cells are coloured with a range of colours starting from yellow at the beginning of the infection and ending up with red at the last time step concerning the infection. Likewise, susceptible cells are coloured in green, removed cells are coloured in brown and empty cells are in black. From the colour bar in Fig. 4.1, we can observe the corresponding colours for infected trees in each time step.

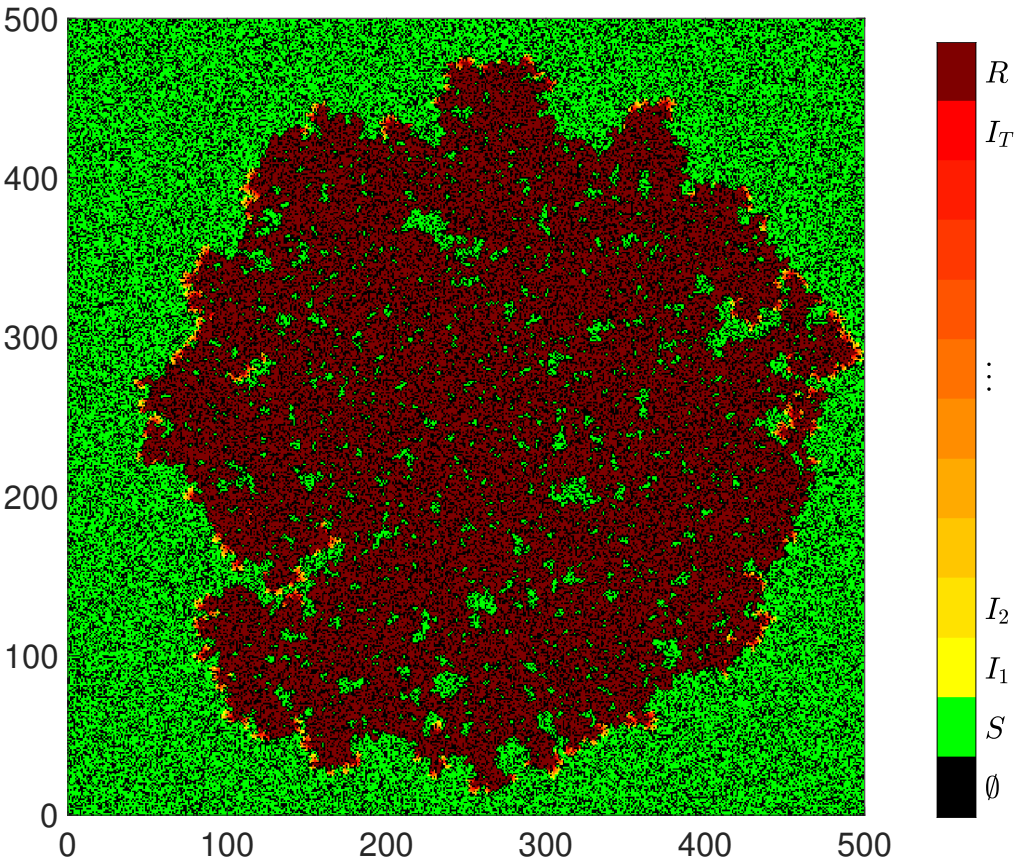

Figure 4.1: The final configuration of the disease propagation obtained for a lattice of size $\mathcal{L} = 500$, tree density $\rho = 0.62$, transmissibility $\beta = 0.5$ and infectious period $T = 10$.

Transmissibility of the pathogen

Transmissibility of the pathogen β represents the probability that a disease is transmitted to a neighbouring susceptible tree at each time step from an infected tree during its infectious period. Transmission is considered independently for all susceptible cells. The probability of a susceptible tree S_i becoming infected in a single time step is given by:

$$P(S_i \to I) = 1 - (1 - \beta)^{n_i},$$

where n_i is the number of infected neighbouring cells. The type of connectivity used in this chapter is the von Neumann neighbourhood, so $n \in [0, 4]$. First, we will use a range

of transmission probabilities when calculating the average velocity or mortality. Then this parameter will be fixed while investigating the effect of the other parameters.

Here we consider the basic case where β is constant throughout the infection, and identical for all trees, but the model could easily be extended. Firstly, variation in the level of disease tolerance of individual cells could be modelled by sampling β from an appropriate distribution for all cells. Secondly, we could model variation with time, for example by making β depend on τ to model the transmissibility varying depending on the stage of the disease, or by making β depend on the time step in some periodic manner to model seasonal effects.

4.3 Examples of Infection Spreading

We observe three clear behaviours of the infection at different densities. At lower densities, we see the infection dying out without spreading, and at higher densities, we see the infection spreading uniformly in all directions and infecting all cells. Then, at intermediate densities, we see the infection spreading throughout the forestscape, but in a non-uniform way which does not infect every cell.

Figure 4.2 (a) shows an example of lower tree density $\rho = 0.35$. We can clearly note that at this relatively small density, the infection could not propagate through the forest due to the limited number of susceptible cells. Moreover, the host distribution is sparse, and so the domain is not well connected. Figure 4.2 (c) shows an example for $\rho = 0.90$ which has a very high density. We see that the disease annihilates most of the cells. It is obvious that the infection kills every cell that comes in its way up to the boundaries. Note that the number of infectious cells at the last time step is exceedingly high. This is due to the fact that if we did not stop the model at this stage, the infection would destroy almost every cell. However, for some intermediate values of ρ, there is an interesting behaviour for which the disease spreads through the lattice, and there is a cluster of connected cells from the centre of the domain to the boundaries, but without total annihilation of the host population. An example of this is given in Fig. 4.2(b). In the rest of this chapter, we will quantify these different behaviours in terms of mortality, the spreading velocity and early-warning indicators.

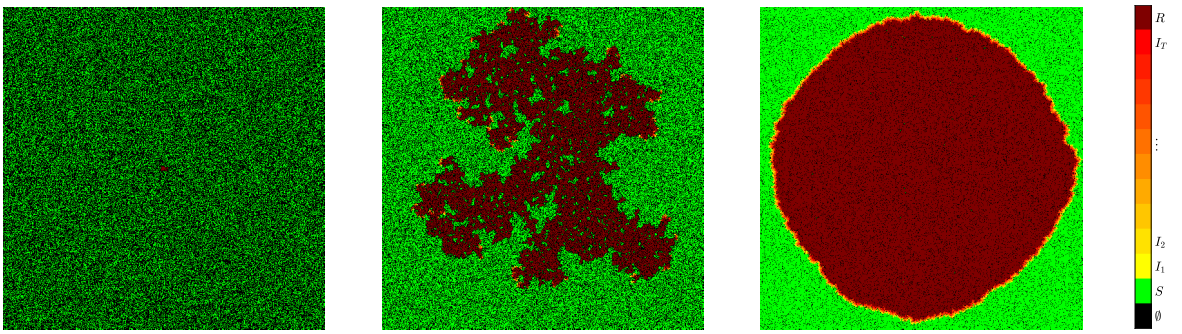

(a) (b) (c)

Figure 4.2: The final configuration of the disease propagation obtained for a lattice of size $\mathcal{L} = 500$, transmissibility $\beta = 0.50$, infectious period $T = 10$ and tree density $\rho = 0.35$ (a), 0.60 (b) and 0.90 (c).

4.4 Mortality and Mortality Phase Diagram

The mortality rate μ, also known as the death rate, measures the number of deaths in a population as a proportion of the population. Here the population is equivalent to the area considered. There are many ways to calculate mortality, all of which are dependent on how we define the area used in the calculation. Therefore, we have defined it according to five different areas and compared the outcomes. Mathematically, the mortality rate can be calculated using the following formula:

$$\mu = \frac{\mathcal{D}}{N_0}$$

where $\mathcal{D}$ is the sum of the number of currently infected and removed cells and N_0, is the number of occupied cells in the considered area at the end of the simulations. We have considered the number of currently infected in the calculation because, in our simulations, the currently infected will eventually die after a specified amount of time has passed. Note that mortalities are calculated at the end of the simulations.

We consider five different ways of evaluating the mortality from our simulations, denoted μ_1, μ_2, μ_3, μ_4 and μ_5. These are defined as follows:

- For Mortality 1 μ_1, we include the entire region.
- For Mortality 2 μ_2, we include trees within the mean radius of the currently infected trees, illustrated in Fig. 4.3 (a).

- Mortality 3 μ_3 (Fig. 4.3 (b)) is the same as μ_2, but uses the minimum radius of currently infected trees instead of the mean.
- For mortality 4 μ_4 (Fig. 4.3 (c)), we take a creative approach. We select all trees within radius 2 of any tree that has died to define the region.
- Mortality 5 μ_5 (Fig. 4.3 (d)) is the same as μ_4 but attempts to include isolated regions by testing whether they are connected to the exterior uninfected region, and including them if they are not.

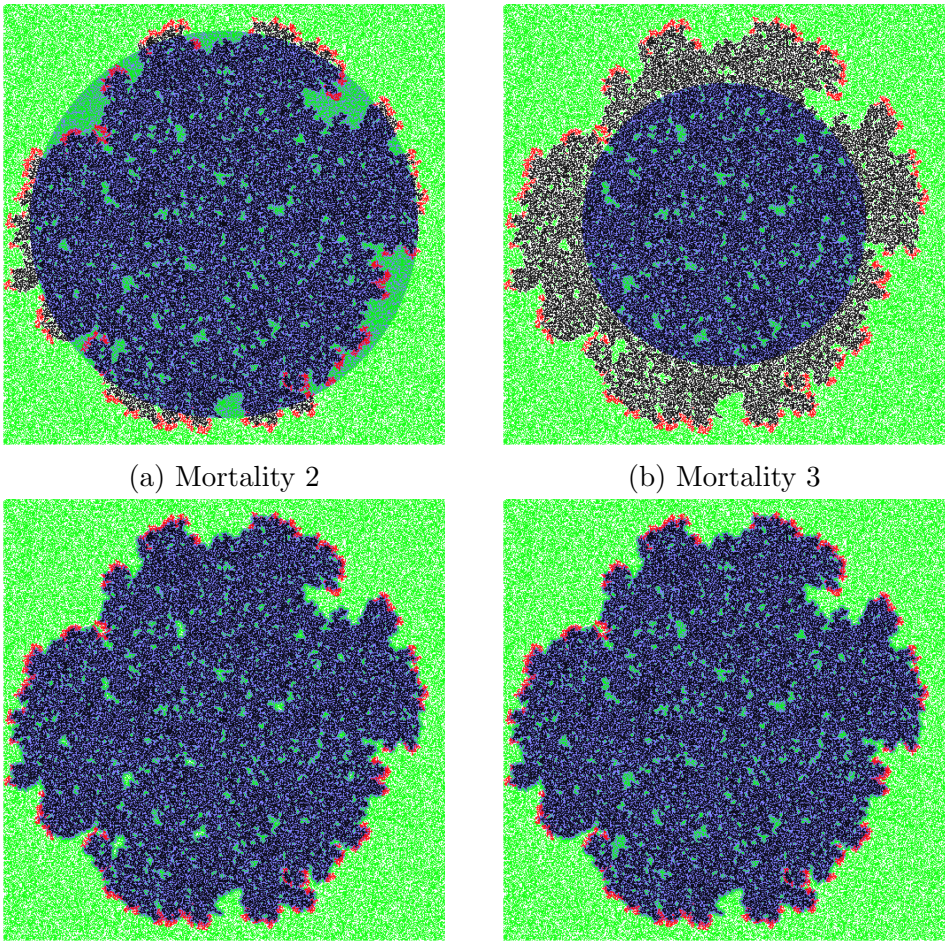

(a) Mortality 2 (b) Mortality 3

(c) Mortality 4 (d) Mortality 5

Figure 4.3: Illustration of the mortality measures discussed in the main text for an infection which has just reached the boundaries. The simulation shown is on a homogeneous forestscape, with $\mathcal{L} = 500$, $\rho = 0.63$, $\beta = 0.90$, and using the von Neumann neighbourhood. Panel (a) shows μ_2, (b) μ_3, (c) μ_4, and (d) μ_5. The colour-scheme used here shows susceptible cells in green, infected cells in red, removed cells in black, and empty cells in white, with the region considered for the mortality calculation shown in blue.

The calculation for μ_2 includes regions which the infection has not yet reached. It is probably a reasonable estimate of the mortality when the infection spreads isotropically (observed to happen with a high infection rate and high tree density here, but potentially affected by the heterogeneity discussed in Chapter 5), otherwise poor. Furthermore, if the infection dies out, this measure will be undefined. The calculation for μ_3 is probably more accurate than μ_2, as it doesn't include regions beyond the infection, but it may lose a lot of data for anisotropic cases. As with μ_2, it will be undefined if the infection dies out. The calculation for μ_4 may give a higher mortality rate than expected for low densities/clustered forestscapes, as isolated surviving regions are not included.

In Fig. 4.4, we plot the ensemble averages of the five different types of mortality as a function of tree density. Note that, for μ_2 and μ_3, the ensemble average excludes infections which die out, as the mortality is undefined for these. This leads to a paucity of data for the lowest densities, with no available data below $\rho = 0.57$, and large errorbars near the critical threshold (with the point at $\rho = 0.57$ actually coming from a single realisation which did not die out). If we were to exclude the infection when it died out without reaching the boundary, all the five types would be zero under the critical transition. Above the critical transition, all the five types of mortality have high values and increase. In terms of epidemic and confinement, all the different types of mortality tell the same story. Thus, we are going to proceed with the first type of mortality, as it the simplest of these measures.

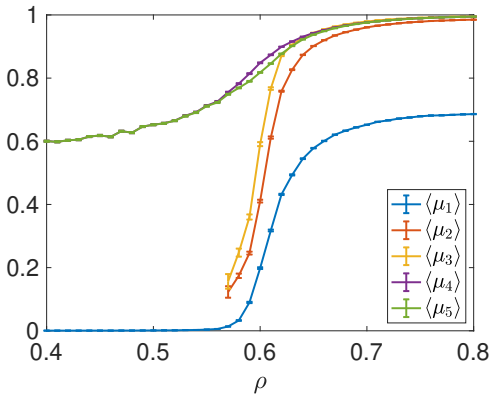

Figure 4.4: The calculated mortality at the end of the simulations for the different methods, for 1000 realisations at each density shown, with $\mathcal{L} = 500$ and $\beta = 0.5$.

We note that $\langle \mu_1 \rangle$ reaches a maximum of 0.65, while that others attain a maximum of 1. This is simply due to the first measure including the corners of the field, which are not infected at the end of the simulation when the infection meets the boundary elsewhere, and which are excluded in the other measures. However, as the transition to epidemic is our principal concern, we have chosen the simplest measure.

4.4.1 The Phase Diagram for Mortality

In Fig. 4.5, the ensemble-averaged mortality $\langle \mu_1 \rangle$ is measured as a function of tree density for different transmission probabilities β. From these results, we observe the occurrence of a critical transition that separates two stable phases: local confinement and disease outbreak. The existence of such a critical value indicates the presence of a phenomenon termed the spanning cluster, which is a spatially connected cluster of infected and removed cells, occurring at a critical density ρ_c. We observe a sudden shift around $\rho = 0.60$, where the mortality shoots up abruptly, which is around the percolation threshold of the von Neumann neighbourhood (i.e. $\rho_c = 0.592746$). Below the critical density, the mortality is almost zero. Above the criticality, the mortality is high and increases.

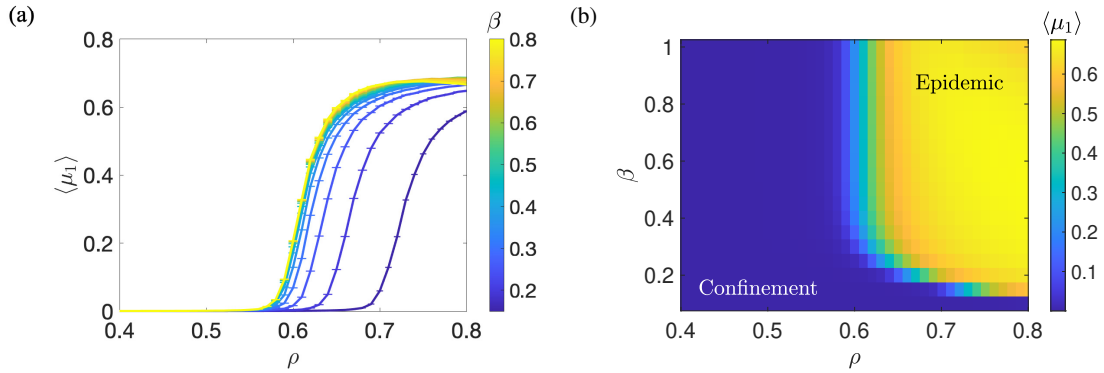

Figure 4.5: Panel (a): for homogeneous forestscapes, $\mathcal{L} = 500$, 1000 ensembles and von Neumann neighbourhood, the averaged mortality μ_1 is plotted as a function of tree density and for various infection rates. Panel (b): the phase diagram of the averaged mortality μ_1, with regions of local confinement and disease outbreaks indicated.

The phase diagram in Fig. 4.5 shows the effects of both tree density ρ and the transmissibility of the pathogen β. The density of the susceptible individuals is significant because the occurrence of the critical transition between both stable phases is dependent on it. For low tree density, the infection dies out quickly due to the limited host population, while for high density the epidemic spreads all over the region. For intermediate values of ρ, the system displays a critical transition between local confinement and disease outbreak. Around this critical value, the disease does not annihilate the whole population of trees, but rather spreads through the artificial forest as an active cluster of sick trees interspersed with healthy susceptible individuals. We also note that for these values around the critical density, the system reveals a fractal behaviour. Figure 4.2 (b) shows an example for $\rho = 0.60$.

The single effect of the transmissibility of the pathogen β is a displacement of the critical value of ρ. For lower β, the density surpassing the theoretical critical density is not sufficient to induce a transmission to epidemic. Although the forestscape is connected, the infection is suppressed by the low infection rate. Conversely, as we increase β, we find that the density at which a transition to the epidemic mode occurs tends towards the theoretical value of ρ_c, observing little change in the observed critical density for $\beta > 0.4$. In Fig. 4.5 (a), we see that for $\beta = 0.15$, the transition occurs at $\rho \approx 0.7$, but for $\beta = 0.5$, it occurs at $\rho \approx 0.6$.

4.5 Velocity and Velocity Phase Diagram

The spreading of the disease through a forest has the effect of separating two domains, which are the susceptible cells S and the removed cells R, by a transient interface of the infected cells I. To quantify the observed dynamic, we calculate the average spreading velocity of the pathogen. This calculation can generate considerable information concerning the behaviour and the potential impact of the infection. There are many different ways to calculate velocity. Here, we are going to present three different types and then make a comparison between them. Based on this, we will establish the preferred definition of the spreading velocity to use throughout the rest of this thesis.

4.5.1 Velocity Definition 1

The spreading velocity v is defined in terms of the number of susceptible individuals S at each time step (i.e. we measured the velocity in terms of the green areas). This velocity is defined as:

$$v(t) = -\frac{\Delta\sqrt{S(t)}}{\Delta t}, \tag{4.1}$$

where S is the total number of susceptible trees available for the infection. This type of velocity was used by Orozco-Fuentes *et al.* [47]. Note that here we consider the change over two timesteps, i.e.

$$v(t) = -(\sqrt{S(t+1)} - \sqrt{S(t-1)})/2.$$

4.5.2 Velocity Definition 2

In the second type the spreading velocity u is defined in terms of the total number of affected trees M available at each time step (i.e. the sum of I and R), which is equivalent to an area. Taking the square root of this gives a lengthscale, and so we construct a velocity as:

$$u(t) = \frac{\Delta\sqrt{M(t)}}{\Delta t}. \tag{4.2}$$

Note that here we consider the change over two timesteps, i.e.

$$u(t) = (\sqrt{M(t+1)} - \sqrt{M(t-1)})/2.$$

4.5.3 Velocity Definition 3

The third type is calculated by the following formula:

$$w(t) = \frac{\Delta\bar{R}_i(t)}{\Delta t}, \tag{4.3}$$

where R_i is the distance for each infected tree from the centre of the domain and $\bar{R}_i$ the mean calculated at each time step. Note that here we consider the change over one timestep. However, it could also be defined as a time-centred difference.

$$w(t) = \bar{R}_i(t) - \bar{R}_i(t-1).$$

4.5.4 Time Series and Velocity distribution

In this section, we examine the time series and distribution of the measured velocities, for simulations with parameters that lead to behaviours of particular interest, and for parameters that give limiting cases. This will help us to understand the differences between the different velocity measures and how to interpret them.

For each single ensemble, at each tree density, we initialise the infection at the centre of the domain. The spreading velocity is then calculated using each of the three types v, u and w. The spreading velocity is extracted from the time series. Therefore, in this section, we are going to show time series examples for the three different types of velocity. Also, we will show some results of the PDF of the velocity and the time-averaged velocity.

First we will consider the trivial case where $\rho = 1$ and $\beta = 1$. Figure 4.6 (a) shows the number of susceptible S, infected I and removed R trees at each time step while panel (b) shows the spreading velocity. Note that this is the limiting case, where the maximum possible values of ρ and β result in the maximum possible velocity. Using the first definition of velocity, v increases throughout the simulation. For our second type, the velocity goes through a brief transient period, settling at $u = \sqrt{2}$. This is consistent with what we expect geometrically - the disease spreads out as a diamond shape (see Fig. 4.6(c)). Each side of the diamond scales with time like $\sqrt{2}t$, and then using the definition for u, we find $u{=}\sqrt{2}$. Note that if ρ and β are less than one, then the spread becomes more circular. For our third definition, we again see some transient behaviour, with w then settling to a constant value of ≈ 0.8. This simply corresponds to the average speed of the edge of the diamond.[1] The transient behaviour is due to the effect of the initial conditions because we initialise the infection as a 5×5 grid in the centre of the domain. We note that the first definition of velocity v goes like $2t/(500^2{-}2t^2)^{1/2}$. This formula is ploted as a black dashed line in Fig. 4.6(b).[2]

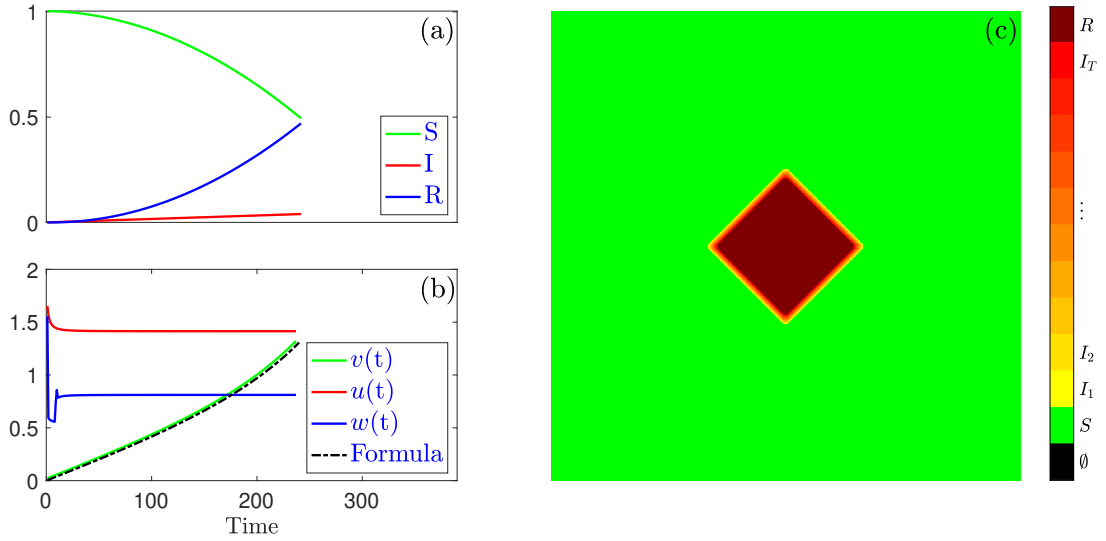

Figure 4.6: Panel (a) SIR compartment occupancy, (b) time series, (c) forestscape configuration for the trivial case $\beta = 1$ and $\rho = 1$, with $\mathcal{L} = 500$ and using the von Neumann neighbourhood.

Figure 4.7 shows the time series of the velocities v, u and w for lower densities where the disease is confined, a moderate density close to the critical density, and a higher density

[1] After N steps, the average distance to the edge of the diamond will be $\frac{1}{N^2}\sum_{i=0}^{N-1}\sqrt{N^2-2iN+2i^2}$. This converges to around 0.811.

[2] If we move to Moore neighbourhood, $(4t/(500^2{-}4t^2)^{1/2})$ can be verified in the same manner as Fig. 4.6(b).

where we have disease outbreak. We can make the following observations from Fig. 4.7 about the behaviour of the velocity in time:

- In the confinement phase such as Fig. 4.7 (a) and (b), we see that, for $\rho \ll \rho_c$, the infection dies out at an early time. In particular for $\rho = 0.52$ which is less than ρ_c, the infection dies out after only 50 time steps which is an extremely small number. Also, we note that the green velocity in Fig 4.7 (a) and (b), appears to be a flat line, but it is not. It is actually fluctuating around some small value (i.e. a little value such as 0.01), evident in the associated histogram for density $\rho = 0.52$ - Fig. 4.8 (a).

- In the transition region, as ρ approaches ρ_c from both sides (i.e. $\rho \approx \rho_c$), the infection can take on a fractal behaviour and thus require more time to reach the edge of the forest. An example for $\rho = 0.62$ is shown in Fig. 4.7(c). We see that the number of time steps is very high, around 600.

- In the epidemic phase such as Fig 4.7 (d) and (e), we observe that for high ρ, when the infection begins to spread in the epidemic regime, the velocity v starts increasing with time. In contrast, the other two velocities are reasonably constant. Considering the number of time steps required to reach the edge of the forestscapes, we see that as ρ increases, the infection needs less time to reach the edges. For $\rho = 1$, which is much higher than ρ_c, the infection only needs 370 timesteps to reach the boundaries.

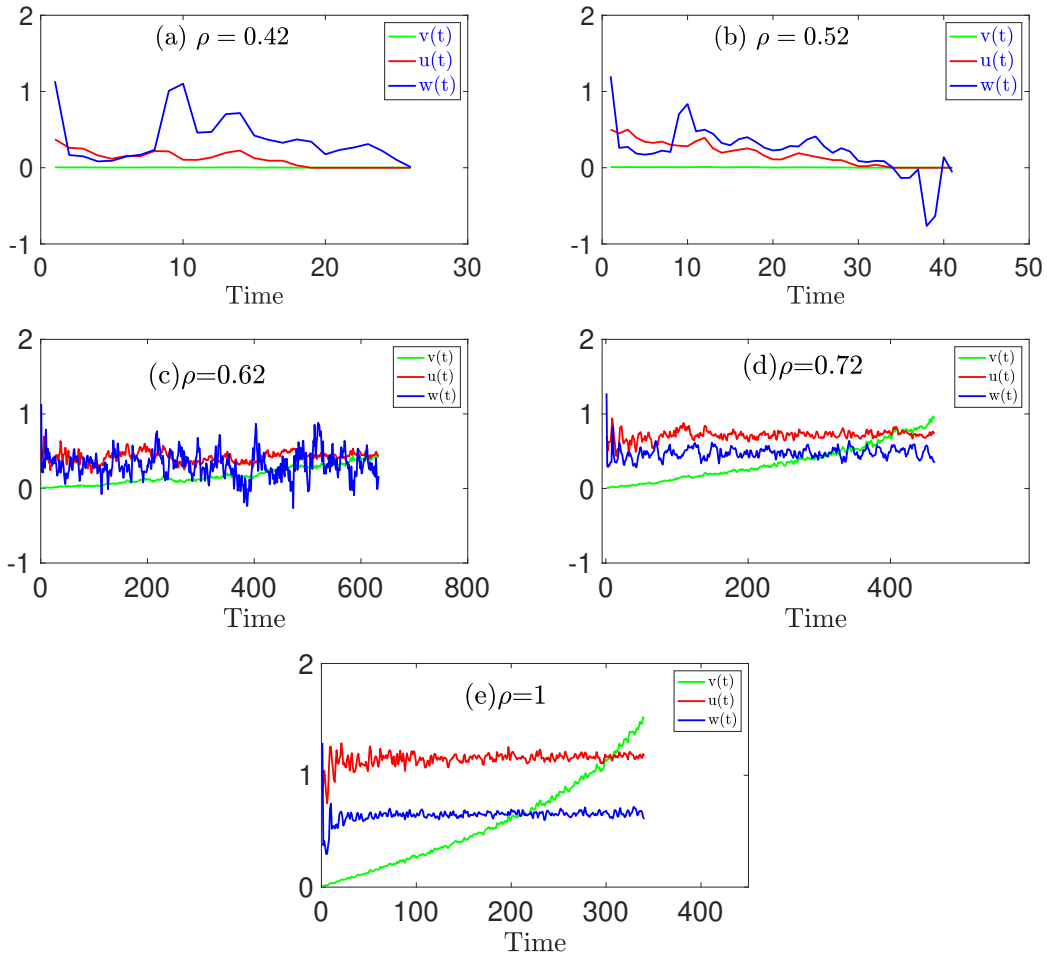

Figure 4.7: Time series for the propagation velocity calculated with the three different definitions v, u and w, and represented with different colours for $\beta = 0.5$, $\rho = 0.42$ (a), $\rho = 0.52$ (b), $\rho = 0.62$ (c), $\rho = 0.72$ (d) and $\rho = 1$ (e).

Shifting our attention to the velocity distribution, we combine all the data from all the ensembles, and produce a PDF of the velocity using the different measures and at densities below, around, and above the critical density in Fig. 4.8. We also take the time-averaged velocity for each simulation and produce a PDF of this in Fig. 4.9. Note that for all the produced PDFs, the initial transient is discarded.

In the PDFs using the non-time-averaged data, we see that as ρ increases past the critical threshold, v becomes more spread out, while u and w become less spread out. This reflects our observation that in the time series (Fig. 4.7) v increases with time in the epidemic regime, while u and w are fairly constant. The PDF for v principally represents the systematic variation in the time series, rather than the random variation in the time

series for u and w. We also see that w is much noisier below the critical threshold, and can also take negative values. This happens when the infection is in the fractal regime, as it can spread back towards the origin. This effect is observable only in w, as this measure is defined based on the distance of infected points from the origin, while the other measures are defined based on the area that is susceptible or affected, which is not sensitive to the direction of the spread. We also note that for densities much greater than the critical density, we do not observe negative velocities for w.

With reference to the time-averaged data, we see that for all of the measures, the position of the peak of the distribution shifts towards higher values as the density is increased. We note that the peak of the PDF for all data coincides with the peak of the PDF for the time-averaged data for u and w, but not for v. Note that if we use the full-time series, including the transient period, the results are not significantly affected.

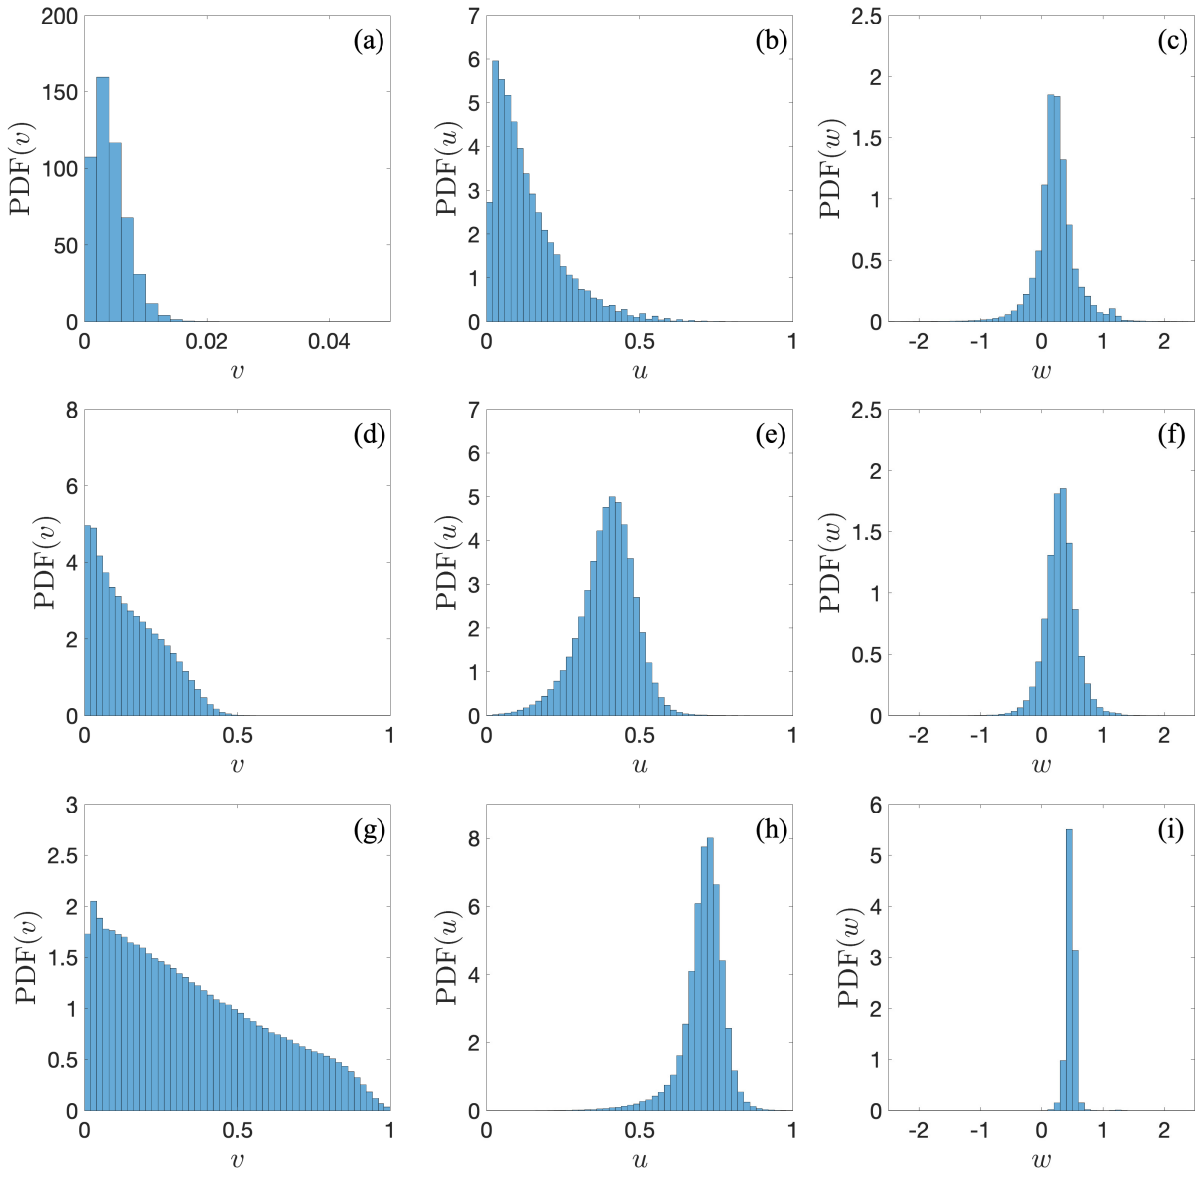

Figure 4.8: PDF of the velocity over all time steps, having discarded the first 10 timesteps of each realisation, for 10^4 realisations, using v (first column), u (second column), and w (third column), for $\rho = 0.52$ (first row), $\rho = 0.62$ (second row), and $\rho = 0.72$ (third row), and with $\beta = 0.5$ for all. Note the different scales on the different axes.

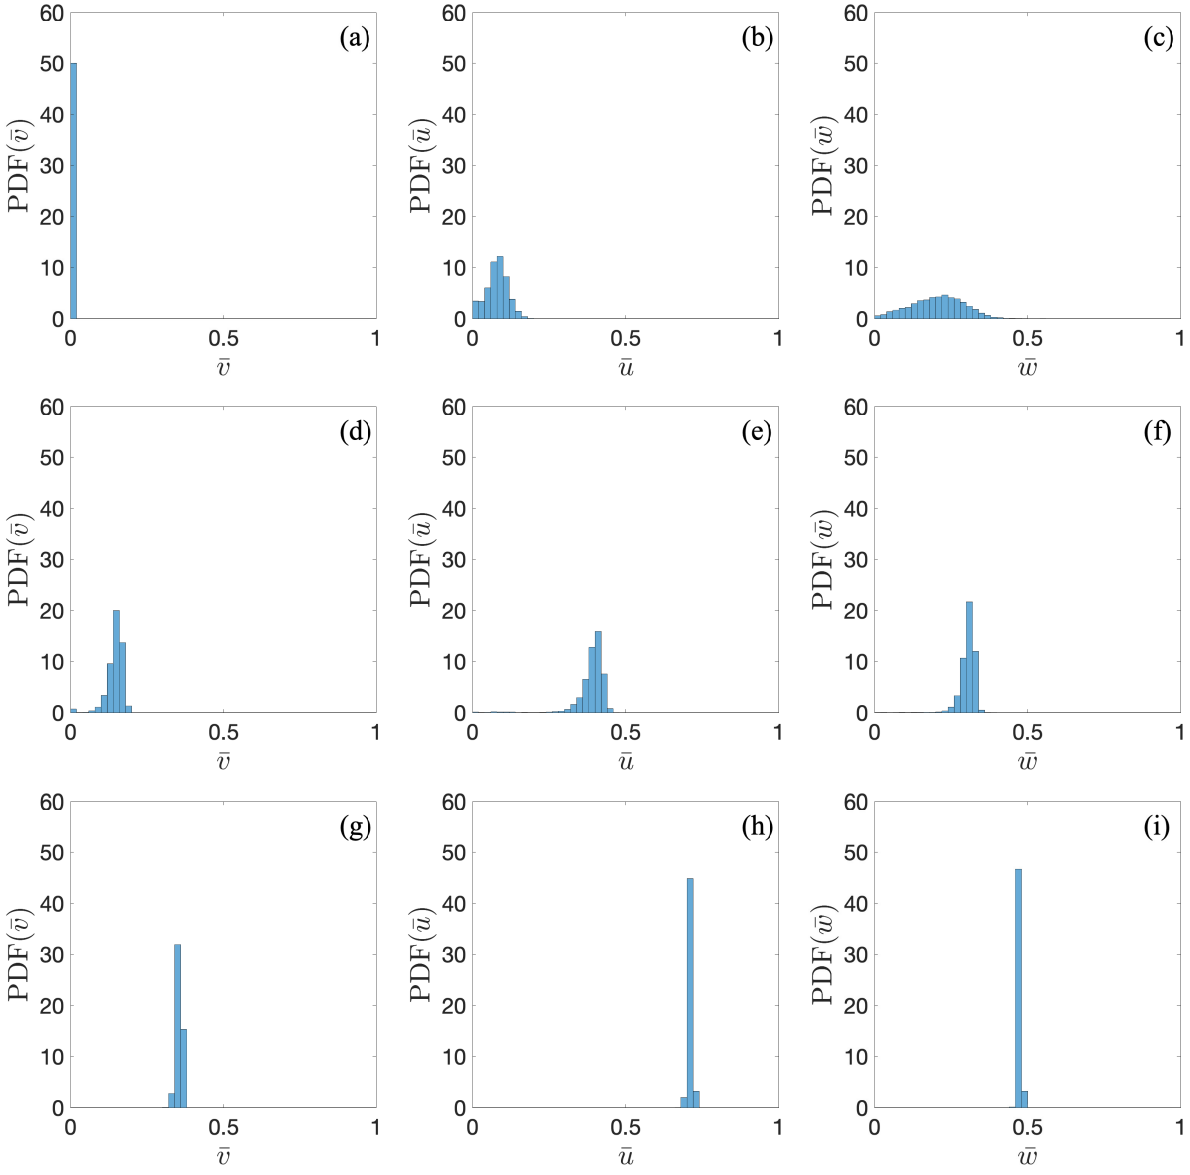

Figure 4.9: PDF of the time-averaged velocity, for 10^4 realisations, using v (first column), u (second column), and w (third column), for $\rho = 0.52$ (first row), $\rho = 0.62$ (second row), and $\rho = 0.72$ (third row), and with $\beta = 0.5$ for all. Note the different scales on the different axes.

4.5.5 Phase Diagram for Different Velocities

The mean spreading velocity ($\langle\bar{v}\rangle$, $\langle\bar{u}\rangle$ and $\langle\bar{w}\rangle$) is calculated by taking the time and ensemble averages of the spreading velocity. Note that we remove the first 10 time steps

of each ensemble from our analysis, to remove the transient period.

In Fig. 4.10, $\langle \bar{v} \rangle$, $\langle \bar{u} \rangle$ and $\langle \bar{w} \rangle$ are obtained as a function of the tree density ρ for various infection rates β. Note that the mean spreading velocity is calculated over the time until the infection dies or it reaches the edge. In all the different cases, we record a non-zero velocity. The reason this is non-zero is because we are only defining the velocity when the infection is actually spreading. Figure 4.11 is the same as Fig. 4.10, but if the infection dies out, the velocity is set to zero; in effect, the averaging is taken in the limit of infinite time and averages to zero in this case. We see that the velocity is zero below the critical density, while the value of the velocity does not change at or above the critical density.

In practice, when people observe an infection spreading, they would like to know whether they are in the confinement or outbreak phase. If somebody is interested in understanding if the real infection is spreading through the forest, what would you tell them if they were able to measure the velocity of the infection spreading? At that point, they can determine whether they are in the confinement or the outbreak phase. However, in Fig. 4.11, we are changing the results because we know it does not spread to the edge; that is, we are changing the results when we know the final results. However, that does not really help to tell somebody who is looking at the infection spreading what is actually going on. Therefore, although the plots in Fig. 4.11 look better and consistent, they provide less information because we have focussed on the epidemic phase here. In practice, we use the results in Fig. 4.10, where we remove the transient and include the cases where the infection die out.

Considering the transmission probability β: as β increases, the critical threshold tends toward the critical threshold of the von Neumann neighbourhood. For $\beta > 0.25$, the transition is around 0.60. However, lower values of β shift the occurrence of the epidemic towards larger values of tree density. This is identical to the behaviour we see in the mortality phase diagram.

As we have seen in mortality, there is a critical value for the density ρ around 0.60 at which we observe a transition to epidemic. We note that this value is close to the percolation threshold $\rho_c = 0.592746$ found in percolation theory when using a von Neumann neighbourhood [76, 81]. However, the introduction of stochasticity by the infection rate can result in significant deviations from formal percolation theory. As the connection between two cells is not sufficient to guarantee transport between them, we observe that as the probability of transmission reduces, exceeding the critical threshold for percolation is not sufficient to guarantee percolation through the system.

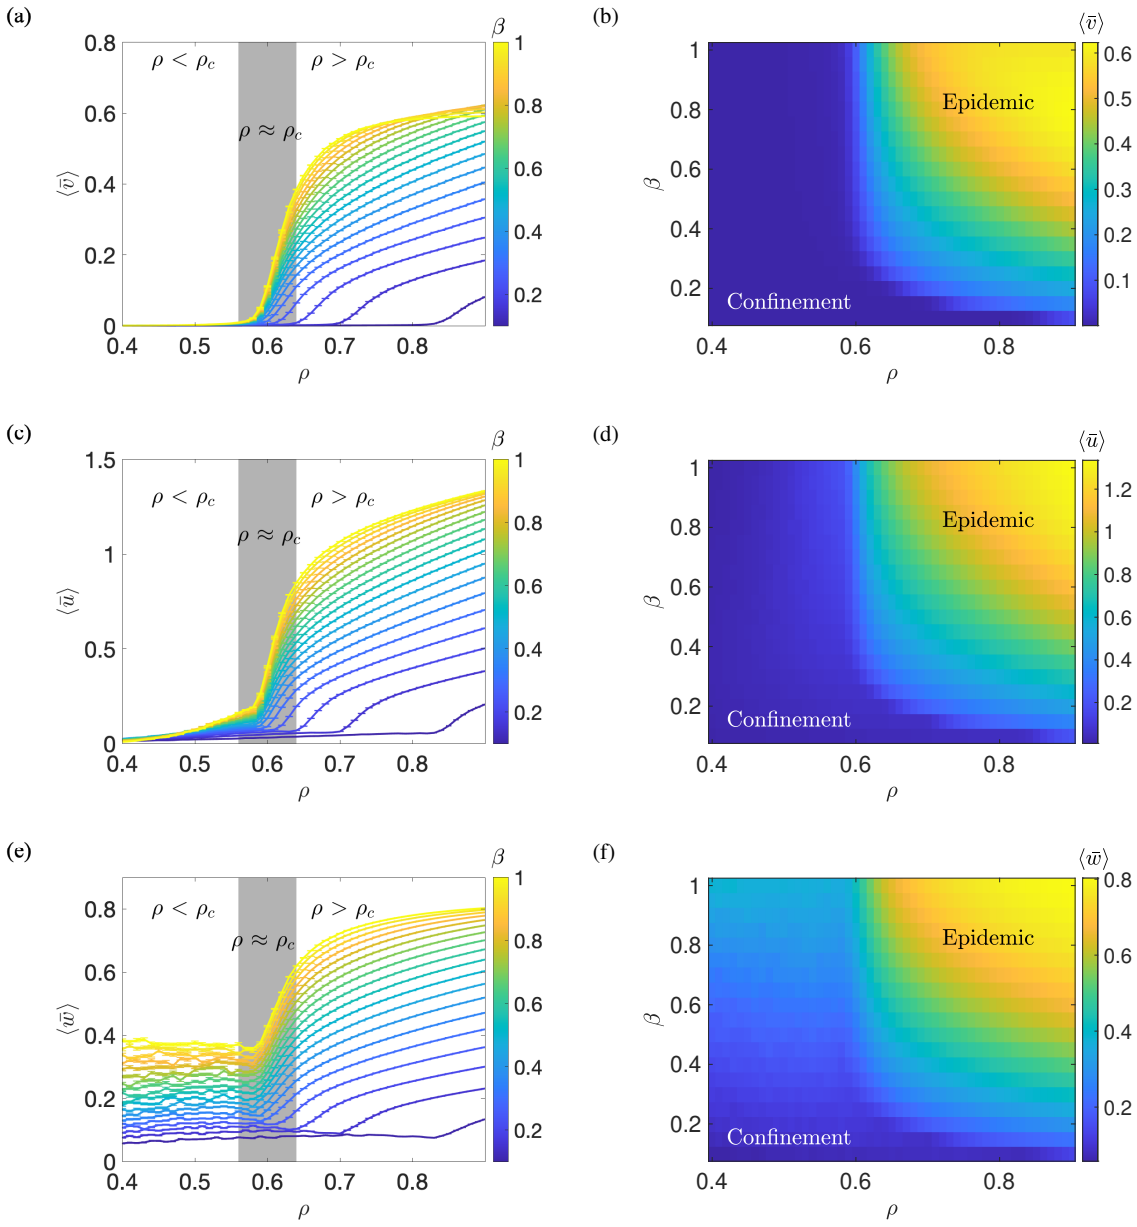

Figure 4.10: Panels (a), (c), and (e): the time and ensemble averaged spreading velocity for the different velocity measures are plotted against tree density ρ for different infection rates β for an $\mathcal{L} = 500$ forestscape. The results have been averaged over 1000 ensembles, with the first 10 time steps from each simulation discarded to account for the transient time caused by the initial conditions. Panels (b), (d), and (f) show the phase diagram of the disease propagation in the artificial forest, with regions of local confinement and disease outbreaks indicated.

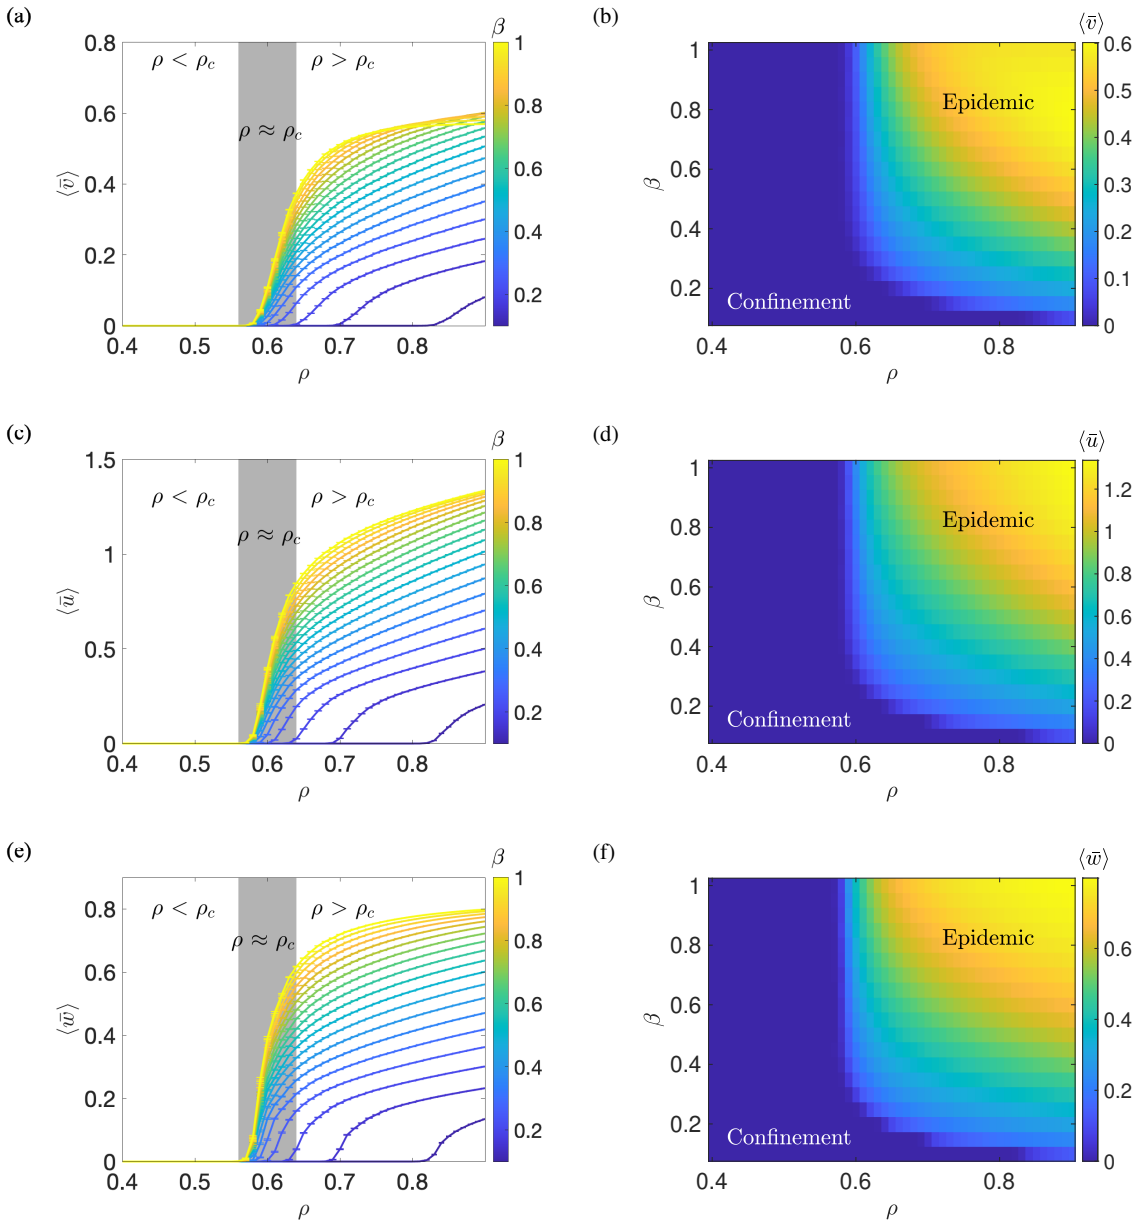

Figure 4.11: Panels (a), (c), and (e): the time and ensemble averaged spreading velocity for the different velocity measures are plotted against tree density ρ for different infection rates β for an $\mathcal{L} = 500$ forestscape. The results have been averaged over 1000 ensembles, with realisation where the infection dies without reaching the boundaries discarded. Panels (b), (d), and (f) show the phase diagram of the disease propagation in the artificial forest, with regions of local confinement and disease outbreaks indicated.

4.5.6 Catastrophic Shifts

In recent epoch years, there has been increasing interest in the use of early-warning signals for abrupt transitions in the dynamics of complicated systems. Although the specifics of systems can obviously be completely unrelated, it transpires that the dynamics of systems close to a critical point share generic properties which can be exploited to predict the onset of critical transitions, also referred to as a catastrophic shift or tipping point, between alternative stable phases [82, 83]. The early-warning indicators themselves are then statistical properties that depend on the system parameters, which change in distinct ways before a critical transition. Because of the generic nature of these indicators, they are suitable for predicting such shifts in systems even when the underlying system dynamics are not well-understood [84, 85]. This technique has been applied to predict catastrophic regime changes in a wide variety of biological and ecological systems, including the simulation of dieback of the Amazon rainforest in climate models [86], desertification processes [87], and as a model for monitoring tree mortality [88]. Here, we will examine how early-warning signals can be used to predict the transition from a confined outbreak of tree disease to an epidemic outbreak. We will consider how the standard deviation, skewness, and kurtosis of the velocity distribution depends on the forestscape density, which we have already seen can drive a shift between these regimes.

The standard deviation describes how spread out values are from the mean. Below the critical transition, infections consistently die out, while above the critical transition, they spread in a consistent fashion. However, around the critical transition, the infection can enter a fractal regime, with a lot of variation in the velocity. Hence, the standard deviation will be small far above and far below the critical transition and will reach a maximum value near the critical transition. Removing the transient for v and u (i.e. Fig. 4.12 (a) and Fig. 4.13 (a) respectively), we note a small but consistent increase in the standard deviation below the critical density, an abrupt increase as we reach the critical point, and far above the critical threshold a reduced, almost stable standard deviation. For v, the standard deviation is higher above the critical threshold, while for u it is lower. On the other hand, the standard deviation of w, which is displayed in Fig. 4.14(a), has a different behaviour below the critical density. It has a high standard deviation below the critical density, which is because of the noise for this definition of velocity below the critical density, as seen in Fig. 4.7, while above the critical density it has a much lower standard deviation.

Someone who observes the infection spreading in various satellite images is able to calculate the standard deviation of the velocity. There are big differences in the standard deviation, depending on whether the disease is confined or in the outbreak state. For example, using w, if the standard deviation is high, the infection is expected to be contained, while if the standard deviation is low, observers should be concerned that the disease will become epidemic.

The third moment, the skewness, describes the asymmetry of the distribution. Positive skewness means the distribution will have a tail to the right, and the bulk of the data will be to the left, and vice-versa for negative skewness. For the skewness of v, which is represented in Fig. 4.12(b), we observe positive values below the critical density with a slight increase as we approach the critical density. Around the transition, the skewness decreases abruptly and becomes negative. Above the critical density, skewness becomes less negative, taking a stable value around -0.5. For the skewness of u and w which are represented in Fig. 4.13(b) and Fig. 4.14(b) respectively, once we are above or below the transition, the statistics are close to Gaussian. Clearly, the velocities are from a symmetric distribution because the skewness is zero. Then at the transition, we have negative skewness which means long negative tail. To check how close we are to a normal distribution, in Fig 4.15, we display the QQ Plot of the sample data versus the standard normal for three different densities which are below, around and above the critical density (0.52, 0.62 and 0.72 respectively). We see that for density below the critical density very small but almost normal. Around the critical density, the negative tail expands, and it starts to produce high velocities. Above the critical density, we have a normal distribution.

The fourth moment, the kurtosis, describes how heavily tailed the distribution is. The standard normal distribution has a reference value of 3, with values above this indicating heavier tails, and so more extreme values, and values below this indicating less heavy tails, and so less extreme values. Close to a transition, strong perturbations in the dynamics can lead to extreme values, which manifests as a larger kurtosis value. Indeed, in panel (c) of Figs. 4.12 - 4.14, the kurtosis begins to rise above the reference value within the transition region, reaching a sharp maximum just inside the epidemic region. For the first velocity type, a small peak is also evident at around $\rho = 0.57$, at the start of the transition.

To sum up, the standard deviation results are much more sensitive below the critical point, which is because the simulation itself is not that big (i.e. it dies out at an early age). The other effect comes from the way the velocity is defined. Above the transition, there is no significant change because you have more data which makes sense. Considering the skewness or kurtosis, we do not see any significant difference when we change the velocity definition.

We also examined whether the inclusion or exclusion of the transient period at the beginning of the infection has a significant influence on the results. We found that the exclusion does not have a considerable effect on the results.

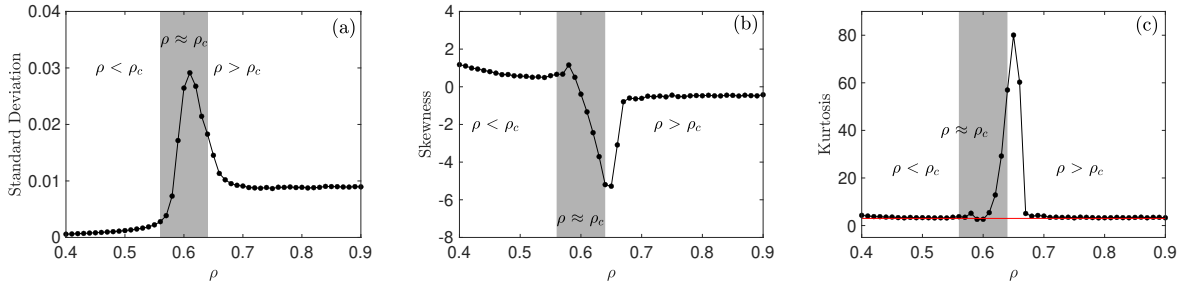

Figure 4.12: These tests consider the velocity v. First 10 time step is discarded. The test was done for 10^4 ensembles.

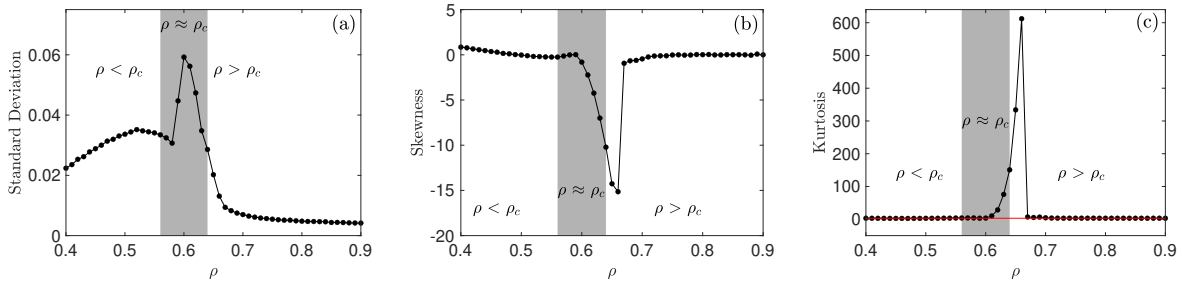

Figure 4.13: These tests consider the velocity u. First 10 time step is discarded. The test was done for 10^4 ensembles.

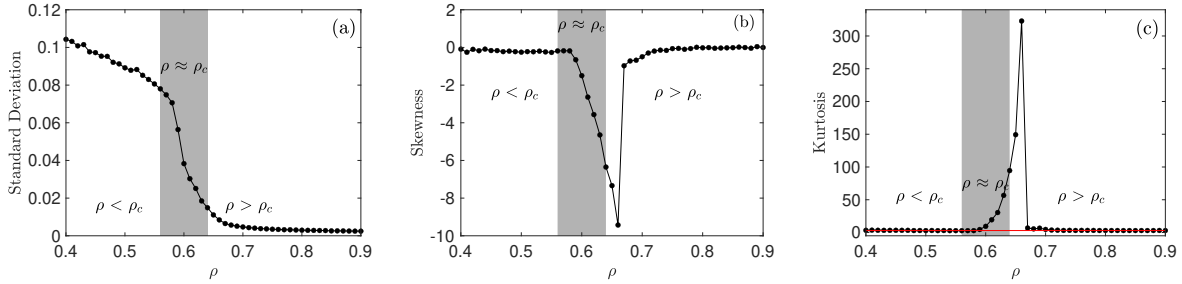

Figure 4.14: These tests consider the velocity w. First 10 time step is discarded. The test was done for 10^4 ensembles.

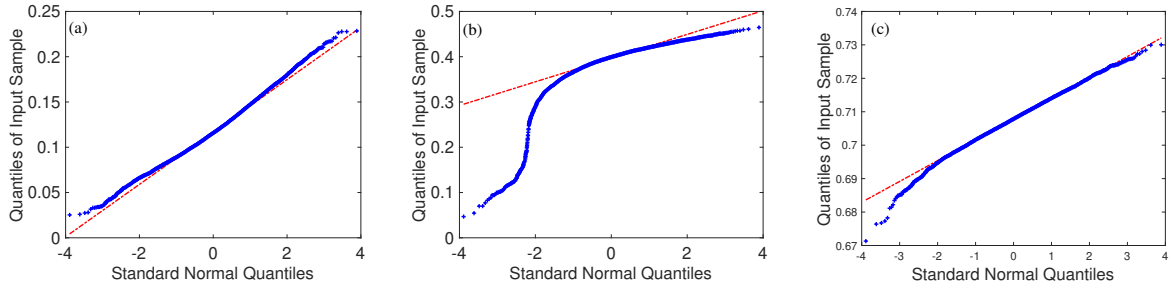

Figure 4.15: For u, we display the QQ Plot of the sample data versus the standard normal for three different densities which are below, around and above the critical density (0.52, 0.62 and 0.72 respectively).

4.6 Differences with Moore Neighbourhood

In the previous sections, we investigated the infection propagation in a lattice-based model using the von Neumann neighbourhood, where an infected tree can only transmit the disease to direct neighbours that are to its north, east, south and west. For quantifying the observed dynamics, we have showed that the velocity u (eq. 4.2) is reasonable and gives consistent results. Therefore, we are going to use it in the following chapters. We have also used another statistical measure which is the mortality rate. For infection spreading following the von Neumann neighbourhood rules, we observed a critical transition at which a confined disease transforms into an outspread epidemic. The abrupt shift in the mortality and the velocity depends on the tree density.

In this section, we are going to use the Moore neighbourhood rules, where an infected tree is allowed to spread to the diagonals as well. The Moore neighbourhood is expected to tell a similar story to the von Neumann neighbourhood. However, the Moore neighbourhood is closer to a realistic spatial connectedness, where we do not expect dynamics to be confined to cardinal directions. Using the Moore neighbourhood, the transition to the epidemic is expected to occur at much lower densities. Also, the established framework of early warning signals of the velocity will be calculated. We expect an abrupt change in the standard deviation, skewness and kurtosis of the velocity around the critical density.

4.6.1 Infection Spreading Examples

In this section, we are going to display a few examples of the infection spreading using the Moore neighbourhood. In Fig. 4.16 (c), we show the infection spreading for the trivial case (i.e. tree density $\rho = 1$ and infection rate $\beta = 1$). We see that the infection spreads

like a square. Geometrically, each side of the square grows in time like $2t$. Thus, by the definition of the velocity in eq. (4.2) for $\rho = 1$ and infection rate $\beta = 1$, the velocity of the infection for the Moore neighbourhood is 2. In Fig. 4.16 (b), we plot the velocity as a function of time, we see that the velocity is a straight line at $u = 2$. In panel (a), we display the SIR compartment occupancy. Note that for lower density ρ and infection rate β, the infection spreads in a circular way.

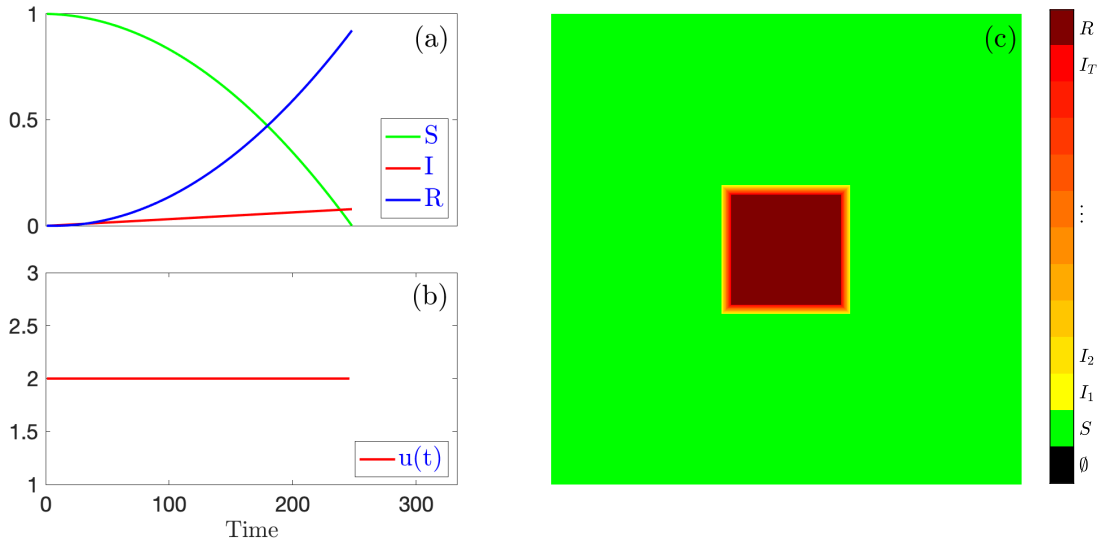

Figure 4.16: Panel (a) SIR compartment occupancy, (b) time series and (c) infection spreading using the Moore neighbourhood, for the trivial case $\beta = 1$ and $\rho = 1$, with $\mathcal{L} = 500$.

In Fig. 4.17, we show three different examples of the infection spreading where the infection is allowed to spread using Moore neighbourhood rules until the infection either dies out or reaches the boundaries. We fix the infection rate $\beta = 0.5$, and we vary the density. For the lower density, $\rho = 0.32$, the infection dies out almost immediately, while for tree density $\rho = 0.41$ the infection takes on a fractal behaviour. This density is around the Moore neighbourhood's critical density, which is 0.407. This tells us that we might experience a transition to epidemic around this density. For the higher density, $\rho = 0.62$, the infection propagates evenly in all directions. From Fig. 4.17, we can clearly see that the critical density at which we observe a transition to epidemic for the Moore neighbourhood is going to be much lower than the critical density for the von Neumann neighbourhood.

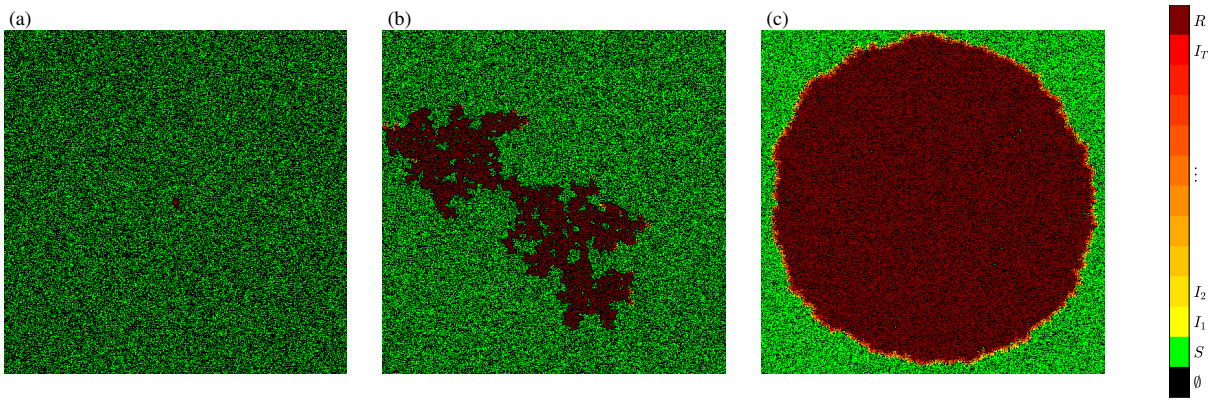

Figure 4.17: Examples of the final configuration of the infection propagation using the Moore neighbourhood for infection rate $\beta = 0.5$, $\mathcal{L} = 500$, and for different tree densities $\rho = 0.32$ (a), $\rho = 0.41$ (b), and $\rho = 0.62$ (c)

4.6.2 The Phase Diagram for Mortality & Velocity

Now, we are using Moore neighbourhood, so each infected tree during its infectious period T has a probability β to infect all its 8 direct neighbours. In Fig. 4.18, the ensemble-averaged mortality $\langle \mu_1 \rangle$ is measured as a function of tree density for different transmission probabilities β. We observe a sudden shift around $\rho = 0.4$, where the mortality shoots up abruptly, which is around the percolation threshold of the Moore neighbourhood. Below the critical density, the mortality is almost zero. Above the criticality, the mortality is high and increases. Obviously, the critical transition depends on the density of the forestscapes and is around the Moore neighbourhood critical density which is 0.407.

Considering the transmission probability β, as β increases, the critical threshold moves towards the critical threshold of the Moore neighbourhood. Apparently, for β>0.25, the transition is around 0.40. However, lower values of β shift the occurrence of the epidemic towards larger values of tree density.

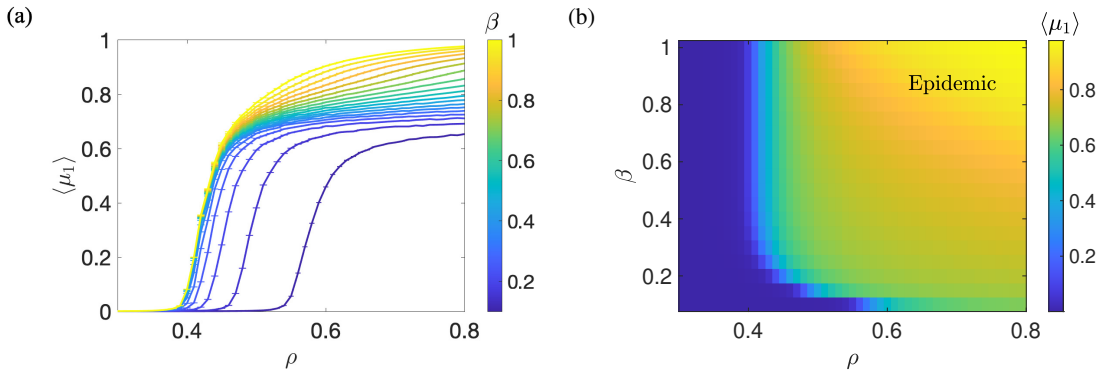

Figure 4.18: Panel (a), for homogeneous forestscapes, $\mathcal{L} = 500$, 1000 ensembles and Moore neighbourhood, the averaged mortality $\langle \mu_1 \rangle$ is plotted as function of tree density and for various infection rates. Panel (b), the associated phase diagram of the mortality, with regions of local confinement and disease outbreak indicated.

In Fig. 4.19, we measure the time-ensemble averaged velocity $\langle \bar{u} \rangle$ as a function of tree density for various transmission probabilities β. Obviously, the velocity shows the same critical transition to epidemic as the mortality. The transition to epidemic occurs at values of densities that are much lower than for the von Neumann neighbourhood. This is around the Moore neighbourhood critical density discussed in Chapter 3.

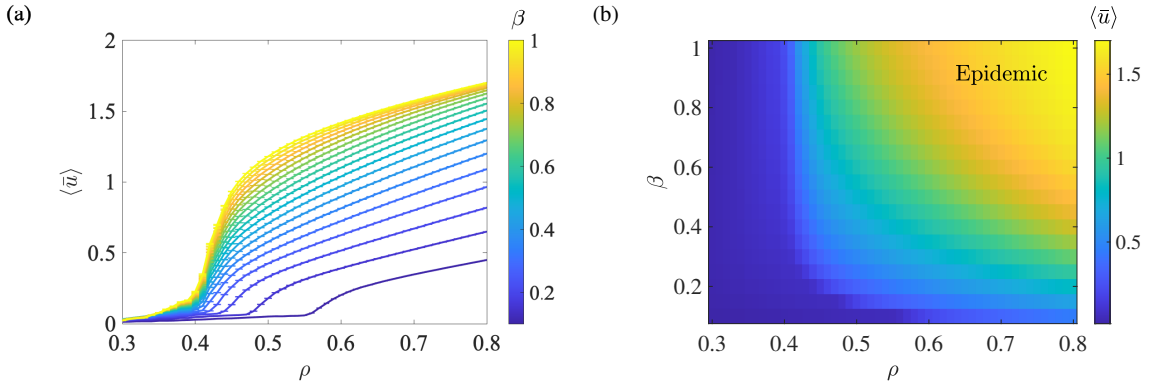

Figure 4.19: Panel (a), the averaged spreading velocity $\langle \bar{u} \rangle$ for homogeneous forestscapes is plotted against tree density ρ for different infection rates β. The infection follows Moore neighbourhood rules of infection on an $\mathcal{L} = 500$ grid, with results averaged over 1000 ensembles. Panel (b) is the associated phase diagram of the disease dispersal, which indicates regions of local confinement and disease outbreaks.

In Fig. 4.20, we measure the standard deviation, skewness and kurtosis of the velocity u. We see that an abrupt change in all the established early warning signals occurs around the critical density.

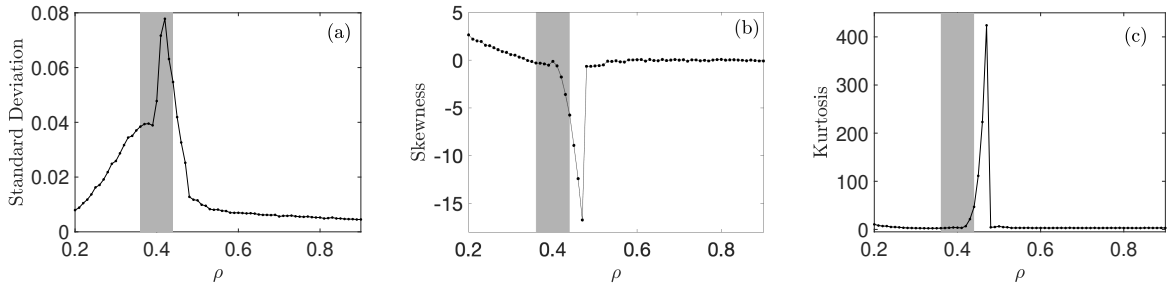

Figure 4.20: These tests consider the velocity u. The first 10 time steps are discarded. The test was done for 10^4 ensembles.

4.7 Conclusions

To model the propagation of a disease on a homogeneous forestscape, we have used a lattice-based model incorporated with a stochastic process that mimics the interactions between the host and the pathogen. We initialized the infection by a patch of infected cells in the centre of the forestscape. Once a healthy susceptible cell becomes infected, it remains infectious for T time steps before it becomes of a removed cell. T was chosen empirically and based on a choice made by Orozco-Fuentes *et al.* [47]. Note that the infection dies out when it reaches the boundaries or when it kills all the susceptible cells. We have investigated a disease propagation in a forest through measuring the spreading velocity of the pathogen on a lattice of susceptible trees. We observe an initial transient at the start of the simulation. This transient occurs at the first few timesteps around the first ten timesteps. This is consistent with the fact that by this time, the initially infected cells already died out. Therefore, we exclude this transient from our calculation of the velocity. We have used three different types of velocities and made a comparison between them. Our findings illustrate that there exists a critical transition that separates two stable phases: local confinement and disease outbreak, that depend on the density of the susceptible cells. We further investigate an established framework of early warning indicators such as standard deviation, skewness and kurtosis of the velocity to predicate the abrupt change between the two stable regimes. Our work indicates that they could be useful to predicate forest epidemics. Even though it is difficult to control and predict the spread of plant disease, our approach could be helpful in forecasting forest disease epidemics. The results from our study suggested suitable densities when considering plantation industry which are densities that are below the critical transition. They probably slow down the propagation

making forests more resilient. Consequently, it allows industry more time for management in case of the occurrence of a disease. Also, for these densities, the infection might die out without intervention.

Basically, what we do in this chapter is that we independently reproduce the paper published by Orozco-Fuentes *et al.* [47]. In addition to their definition of velocity, we have used two more definitions. We found that they all lead to the same conclusion. We made a comparison between these definitions and proceeded with the more sophisticated one which is u. It is interesting to note that, despite the issues identified with some of the velocity measures, all three can act as early warning indicators - that even a poor estimate of the velocity can work in this framework is perhaps indicative of its robustness. In addition, we introduce another statistic, the mortality rate. We find that the mortality tells the same story as velocity, that there is a critical transition separating two stable phases: confinement and epidemic. The value of the critical density is consistent between these two measures.

The method of infection used by Orozco-Fuentes *et al.* [47], which is the von Neumann neighbourhood can be made more realistic by including the diagonals (i.e. using Moore neighbourhood). This motivates us to introduce the Moore neighbourhood. Using the Moore neighbourhood to allow the infection to spread from an infected tree to its neighbouring trees, we measured the mortality and the spreading velocity. We found that the Moore neighbourhood tells an almost identical story as the von Neumann neighbourhood. The only difference is that the critical density for the Moore neighbourhood is around 0.40, while the threshold for the von Neumann neighbourhood is approximately 0.60, again coinciding with the percolation threshold and the emergence of a spanning cluster. This reduction in critical density for the Moore neighbourhood can be understood as the network being more connected than the von Neumann neighbourhood.

Chapter 5

Modelling the Infection Spreading in a Heterogeneous Forestscape

5.1 Introduction

In Chapter 2, we generated heterogeneous distributions for forestscapes using two types of neighbourhood: von Neumann and Moore. We concluded that the Moore neighbourhood produces more realistic heterogeneous forestscapes, and so decided to proceed with it. All the synthetically generated forestscapes used in the rest of this thesis are produced using the Moore neighbourhood.

Having studied the infection spreading on homogeneous forestscapes, we now consider heterogeneous forestscapes, which are more realistic. Previously, when we studied the homogeneous fields, we determined that quantitative results show the occurrence of a critical transition between two stable phases, specifically, local confinement and an extended epiphytotic outbreak. We have investigated the effect of both the transmission probability β and the forestscape density ρ. Both parameters regulate the evolution of the disease, with the most effective one being the tree density, while the transmission probability, as it decreases below 0.3, only causes a displacement of the critical value towards higher values of the tree density.

Considering the heterogeneous forestscapes, in addition to the tree density ρ and the infection rate β, the clustering parameter is of a particular interest. In this chapter, we will put the disease on top of the synthetically generated heterogeneous forestscapes. Then, we will investigate the disease propagation in these forestscapes by means of measuring the spreading velocity of the pathogen on susceptible trees. We will do that for different heterogeneity levels, ranging from homogeneous forestscapes (i.e. $\gamma = 0$) to highly clustered

ones (i.e. $\gamma = 5$), for different tree densities and infection rates. We are interested in how the heterogeneity affects the disease propagation.

5.2 Examples of Infection Spreading in a Clustered Field

In this section, we will show examples of how the infection spreads in heterogeneous forestscapes, comparing them to examples for the infection spreading in homogeneous forestscapes.

Figure. 5.1 shows an example of the propagation in the epidemic regime with $\rho = 0.7$, $\beta = 0.5$ and $\gamma = 5$. Although we are clearly in the epidemic regime, there is some suppression of the infection in some directions, because the infection is not able to spread over the gaps introduced by the clustering.

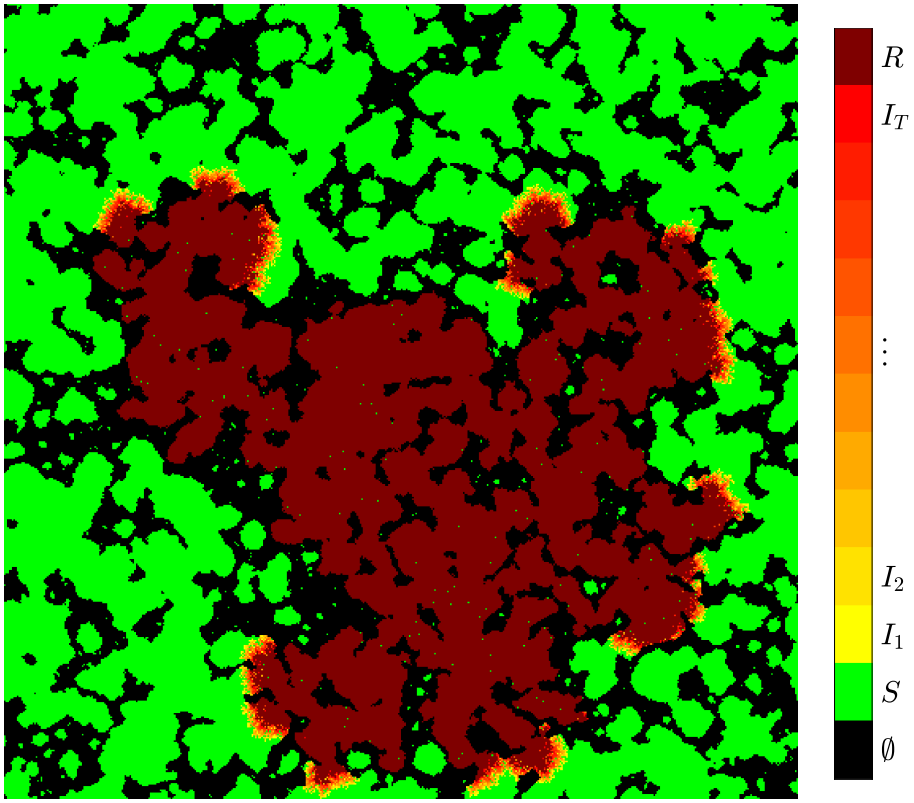

Figure 5.1: An example of the local infection in a synthetic heterogeneous forestscape. Here $\gamma = 5$, $\mathcal{L} = 500$, $\beta = 0.3$, $\rho = 0.7$, and we use the Moore neighbourhood as method of infection.

In Fig. 5.2, we show an example of the infection propagation for the heterogeneous forestscapes with $\gamma = 5$ and compare it to homogeneous forestscapes with all other parameters fixed. We observe that, for the heterogeneous forestscapes, the disease could not spread to the boundaries and dies out at an early stage (i.e. in the confinement regime), while for the homogeneous forestscapes we are definitely well into the epidemic regime. This is a clear indication that clustering can cause suppression of the infection propagation.

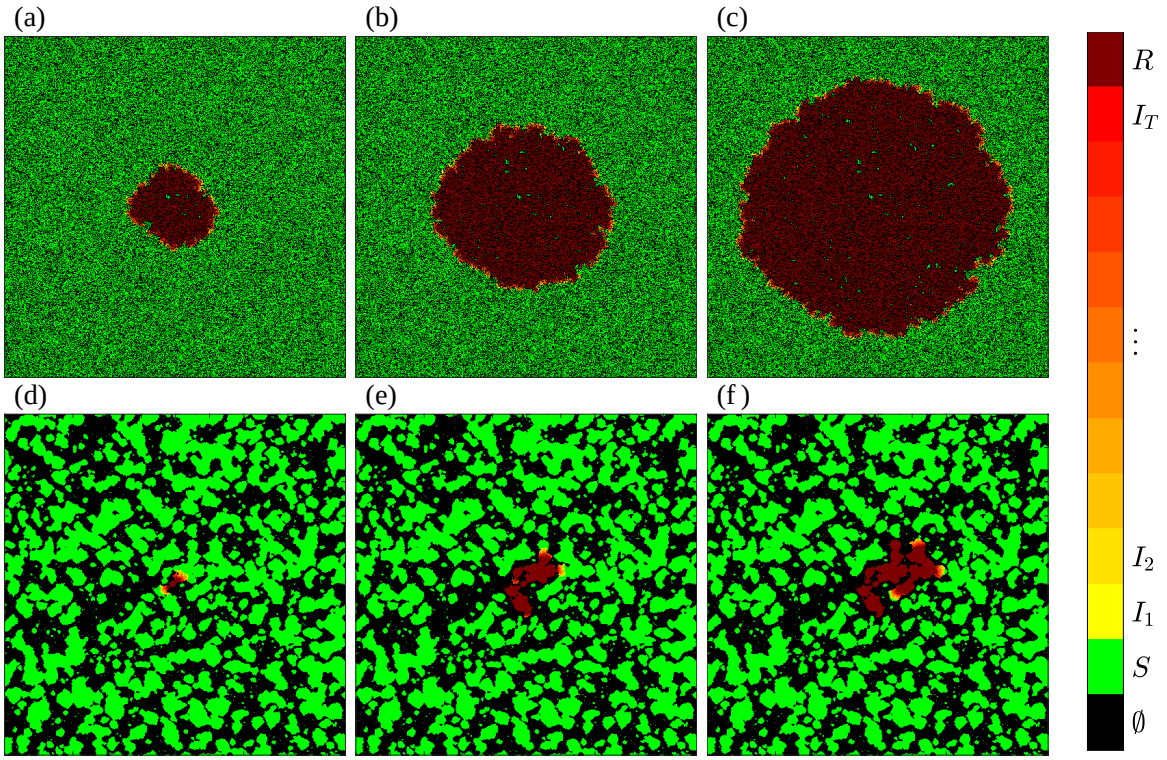

Figure 5.2: Comparison of how the local infection propagates for homogeneous and heterogeneous forestscapes. The top row shows the progression of an infection for a homogeneous forestscape, with snapshots at times $t = 100$ (a), $t = 200$ (b), and $t = 300$ (c). The bottom row shows the progression of an infection for a heterogeneous forestscape, with $\gamma = 5$, at times $t = 20$ (d), $t = 60$ (e), and $t = 100$ (f). Both simulations use the Moore neighbourhood to propagate the infection, and have $\mathcal{L} = 500$, $\beta = 0.5$, and $\rho = 0.5$.

5.3 The Phase Diagram for Velocity

To see the effect of the different parameters, we are going to show phase diagrams for the different heterogeneity levels and the different infection rates. We will do that for a range

of densities between 0.3 and 0.8. In Fig. 5.3, we plot the averaged spreading velocity as a function of tree density for various level of heterogeneity. We see that an increase in the heterogeneity level suppresses the velocity of the infection, and causes the transition to epidemic to require a larger density.

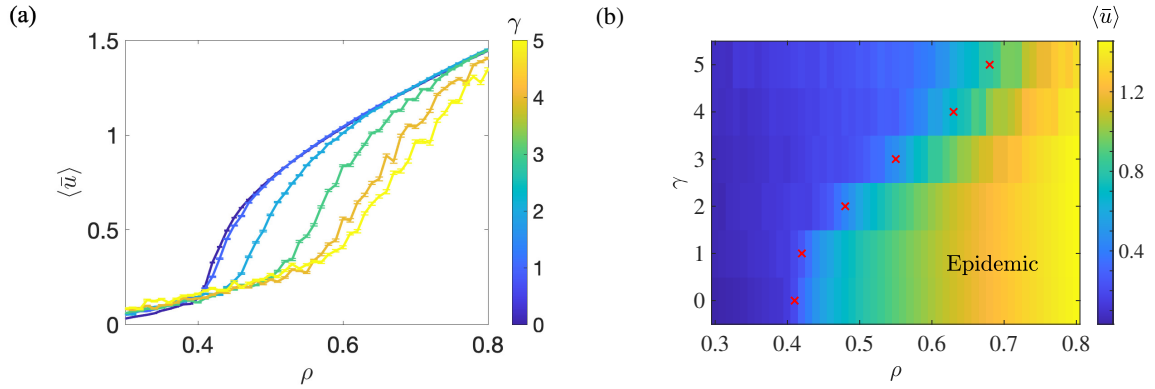

Figure 5.3: Panel (a), the averaged spreading velocity $\langle\bar{u}\rangle$ plotted against tree density ρ for different heterogeneity levels γ. The other parameters are $\mathcal{L} = 500$, $\beta = 0.5$, and the results have been averaged over 1000 ensembles. Panel (b) shows the associated phase diagram of the disease dispersal which indicates regions of local confinement and epidemic, with the percolation threshold values ρ_c previously identified (Table 3.3) superimposed as red crosses.

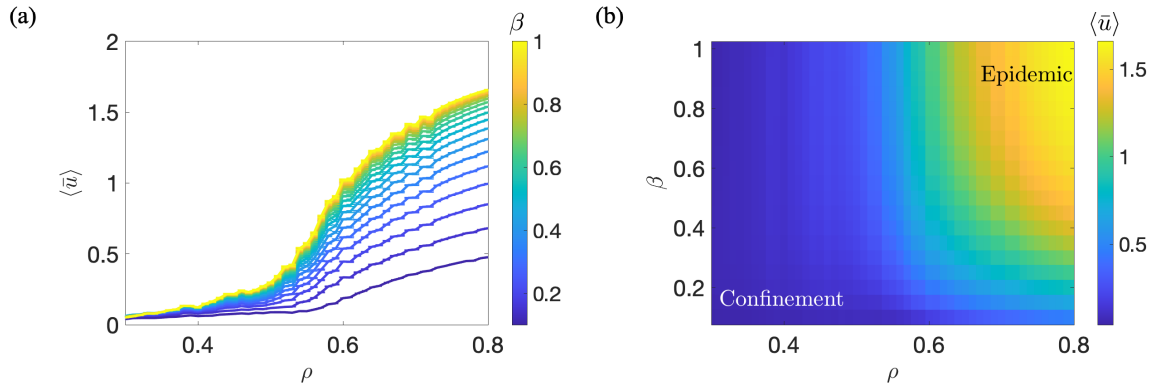

Figure 5.4: Panel (a), the averaged spreading velocity $\langle\bar{u}\rangle$ for heterogeneous forestscapes with $\gamma = 3$, plotted against tree density ρ for different infection rates β. Here $\mathcal{L} = 500$, and the results have been averaged over 1000 ensembles. Panel (b) shows the associated phase diagram of the disease dispersal, with regions of local confinement and epidemic.

The increase in the critical density is consistent with our investigation into the percolation threshold for heterogeneous forestscapes in Chapter 3. There we found that, for the Moore neighbourhood, increasing the heterogeneity parameter increased the percolation threshold, the point at which a spanning cluster emerges. In Fig. 5.3 (b), we show the phase diagram for the propagation velocity as we vary ρ and γ, with the values of ρ_c previously identified superimposed. We see that the critical density for transition to the epidemic is in good agreement with these values. We speculate that this may be related to the actual form of the spanning cluster for highly clustered distributions.

In Fig. 5.4, we plot the time- and ensemble-averaged velocity against tree density for various infection rates β, for heterogeneous forestscapes with $\gamma = 3$. Reducing the infection rate results in the reduction of the propagation velocity, and the shifting of the critical density to higher values.

5.4 The Phase Diagram for Mortality

Having studied the phase diagram for the velocity, we are now going to study the phase diagram for the mortality. We will investigate the ensemble-averaged mortality for the various heterogeneity levels. We will also investigate the mortality as we change the infection rate for heterogeneous forestscapes with $\gamma = 3$. In Fig. 5.5, we plot the mortality against tree density for a range of heterogeneity levels. We observe that, as with the velocity, the

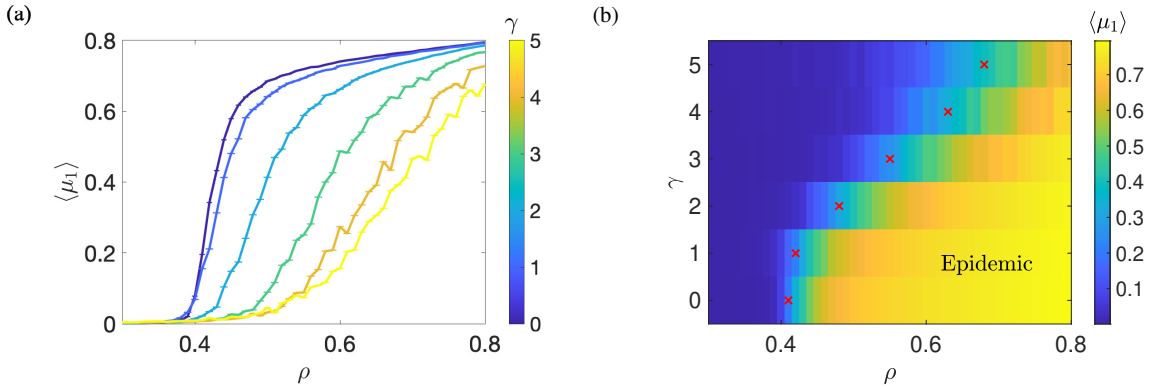

Figure 5.5: Panel (a), the ensemble-averaged mortality $\langle \mu_1 \rangle$ is plotted against tree density ρ for different heterogeneity levels γ. The other parameters are $\mathcal{L} = 500$, $\beta = 0.5$, and the results have been averaged over 1000 ensembles. Panel (b) shows the associated phase diagram of the mortality rate, which indicates regions of local confinement and disease outbreaks, with the percolation threshold values ρ_c previously identified (Table 3.3) superimposed as red crosses.

mortality is suppressed by the increased heterogeneity, and the critical density is shifted to higher densities for more heterogeneous forestscapes.

In Fig. 5.6, we plot the mortality for heterogeneous forestscapes with $\gamma = 3$ against tree density for various values of infection rate. We observe that the mortality is again reduced as the transmission rate is reduced, with the critical density again shifted to higher values, although the effect is not as clear as for the velocity phase diagram.

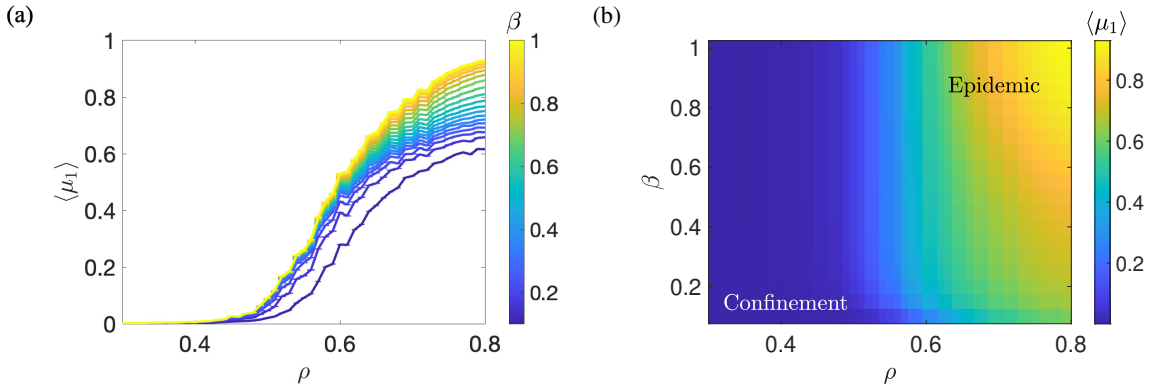

Figure 5.6: Panel (a), the ensemble-averaged mortality $\langle \mu_1 \rangle$ for heterogeneous forestscapes with $\gamma = 3$, plotted against tree density ρ for different infection rates β. Here $\mathcal{L} = 500$, and the results have been averaged over 1000 ensembles. Panel (b) shows the associated phase diagram of the mortality, with regions of local confinement and disease outbreak indicated.

5.5 Conclusions

Using our local model with Moore neighbourhood propagation on heterogeneous forestscapes, we observe that the dynamics are strongly changed by the degree of clustering. We see that increasing the level of tree aggregation reduces the propagation velocity, which is accompanied by a shift of the critical density, which separates confinement and epidemic behaviour, to higher values.

As with homogeneous forestscapes, we have quantified the infection spreading by measuring the spreading velocity of the pathogen, and the mortality rate. Here, we have investigated how the transmission probability and the forestscape density affect these measures, and the transition to epidemic, in conjunction with the heterogeneity parameter.

The effect of the transmission probability and the forestscape density is the same as for the homogeneous forestscapes; a critical transition occurs as the density increases, with the

density at which this occurs being shifted to higher values as the transmission probability decreases. With the inclusion of the heterogeneity parameter, we observe the suppression of the propagation velocity and mortality as the heterogeneity increases, again accompanied by an increase in the critical density. The increase in the critical density is consistent with our investigation into the percolation threshold for heterogeneous forestscapes in Chapter 3. For future work, we could quantify the effects of the isolated sub-forests that arise at large values of the clustering parameter γ.

Note that we attempted to use the statistics from Chapter 4 which are standard deviation, skewness and kurtosis of the velocity as indicators of catastrophic events. However, this was not useful, and we could not predict the transition to epidemic when we initialize the disease in our heterogeneous forestscapes. This is because these statistical tests are very sensitive, and the curves of these statistical tests become noisy.

Chapter 6

Modelling the Infection Spreading Using Non-Local Interactions

6.1 Introduction

So far, all of the interactions considered have been local, with the infection only able to spread to direct neighbours of infected trees. In reality, the spread of infection is not local, but non-local (i.e. has a non-local kernel) [62]. A study by Chandelier *et al.* in 2014 investigated the distance at which the infectious agent (in the case of that study, fungal spores) was deposited, finding that while most remained within 50m of the source, some proportion reached beyond this distance, potentially in concentrations sufficient to initiate new infections [63]. Relating this to our forestscapes, 50m would roughly translate to 5 pixels in our simulations, but there are some difficulties in translating this directly into a distribution for the spread of infection. In real life, the spread of an infectious agent could be affected by numerous factors, including the topology of the surrounding area (in particular the altitude of infected sites), weather, and airflow patterns. Furthermore, infection is not simply a boolean parameter but relates to the concentration of the infectious agent. This may be further complicated by variations in the concentration of the infectious agent emitted over the course of an infection, in different seasons, or the susceptibility of individuals to the infection [62, 89, 90].

Given the fairly limited data available (see Chandelier *et al.* [63], table 3) and the myriad of factors that could be involved, separating this all out to arrive at a model of the spread that captures all of this is clearly not a viable option. Moreover, the means by which different tree diseases propagate are varied, and developing a hyper-specific model for one will likely be of no use for another. We instead develop a generic model for the infection rate as a spatial distribution, which has three components: the form of the

distribution (e.g. linear, power-law, Gaussian), and two parameters; the transmissibility and a characteristic distance. Such a model allows us to probe non-local effects in a generic manner which could be of use in many different scenarios; we also note that this model could be straightforwardly extended to include many of the aforementioned effects if it became apparent that some were crucial in describing the spread of a particular disease.

In this chapter, with the infection rate now modelled as a function of the distance between the particular infected and susceptible trees, rather than being restricted to direct neighbours, we investigate the effect of non-local interactions on the spread of the infection. This is done in conjunction with the previously probed parameters: density, the transmissibility of the pathogen, and the degree of heterogeneity of the forestscape. We again quantify the behaviour of the disease in terms of the velocity, mortality, and transition to epidemic, using simulated homogeneous and heterogeneous fields. To better understand the spread, we also consider other statistics, such as the number of infected trees at each time step and the probability density function of the minimum distance of a newly infected cell from any currently infected cell.

6.2 Different Non-local Spreading Functions

In this section, we will discuss how a non-local component is included in our model and compare some possible choices for the kernel function. To move from infections being spread by direct contact between neighbouring cells, to infections being spread by non-local interactions, we introduce a spatial component to the infection rate. Rather than a constant probability for transmission between direct neighbours, we base our calculations on the probability $p(r)$ of a particular susceptible cell being infected by a particular infected cell at a distance r in a single timestep. We expect that $p(r)$ will have a maximum value at the origin, and will decrease to 0 far away from the infected cell, either at infinity or by some finite distance. There are several obvious potential choices satisfying this, which we compare in this section. For each choice below, the amplitude of the infection transmissibility is denoted by β, and the characteristic width of the distribution by σ. We consider four kernels for non-local interactions:

- Linear:

$$p(r) = \begin{cases} \beta\left(1 - \dfrac{r}{\sigma}\right), & r \leq \sigma, \\ 0, & r > \sigma. \end{cases} \tag{6.1}$$

- Circle:

$$p(r) = \begin{cases} \beta\sqrt{\dfrac{\sigma^2 - r^2}{\sigma^2}}, & r \leq \sigma, \\ 0, & r > \sigma. \end{cases} \tag{6.2}$$

- Power Law:

$$p(r) = \begin{cases} \beta, & r \leq \sigma, \\ \dfrac{\beta}{(r/\sigma)^\alpha}, & r > \sigma, \end{cases} \tag{6.3}$$

  where α is a positive constant which gives the degree of the power law. In this analysis we use $\alpha = 2$, which we will discuss subsequently.

- Gaussian:

$$p(r) = \beta \exp\left\{\frac{-r^2}{2\sigma^2}\right\}. \tag{6.4}$$

In Fig. 6.1 we show the different kernels with the same σ on a single plot to directly compare their form, and also show them on a logarithmic scale to highlight the differences in the tails of the distributions.

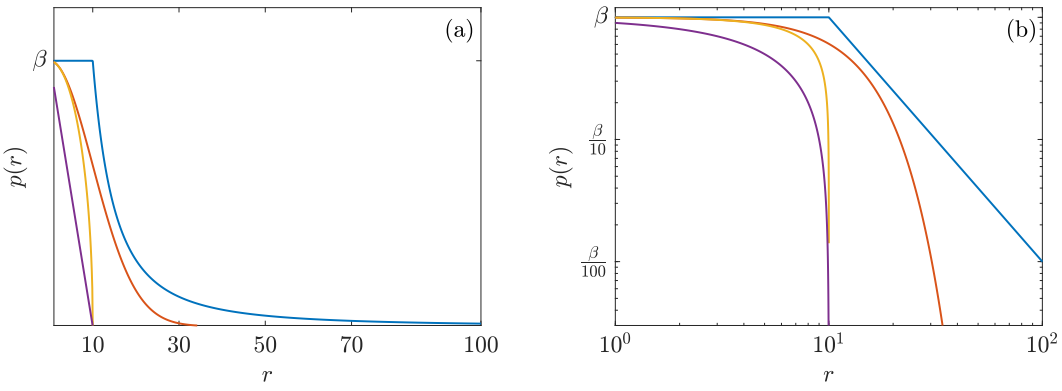

Figure 6.1: Comparison of the different kernels: linear (purple), circle (yellow), Gaussian (red), and power law with exponent $\alpha = 2$ (blue), with $\sigma = 10$ for all. Panel (a) shows the different kernels on a linear scale, while panel (b) shows the same plot on a logarithmic scale.

Let us now compare the properties of these distributions, and how we might anticipate an infection using them as kernels to behave.

Firstly, their overall forms. Both the linear and circle kernel have a fixed distance, given by σ, beyond which the probability of the infection spreading is zero. In contrast, the Gaussian and power-law kernels have tails beyond a distance σ, with the tail of the power-law being far more pronounced at large distances (see Fig. 6.1 (b)). A notable difference between the power law and the other kernels is the behaviour near the origin: within σ of the origin, the power law has a constant rate of infection, while the other kernels have increasing rate of infection the closer to the origin we are.

We now consider the probability $p_1(r)$ that at least one tree at a distance r from a single infected becomes infected in a single timestep, and the probability $p_2(r)$ that at least one tree at a distance greater than r becomes infected in a single timestep, for a forestscape completely populated ($\rho = 1$) with susceptible trees. Rather than integrating the kernels over a band of width 1 for the former, or from r to infinity for the latter, we calculate this using:

$$p_1(r) = 1 - \left[\prod_{i \text{ s.t. } r \leq r_i < r+1} (1 - p(r_i)) \right], \tag{6.5}$$

where the product is taken over all cells at a distance between r and $r+1$ from the infected cell, and r_i is the distance to the i^{th} cell, and

$$p_2(r) = 1 - \left[\prod_{i \text{ s.t. } r_i \geq r} (1 - p(r_i)) \right], \tag{6.6}$$

where the product is taken over all cells at a distance $\geq r$. This is done to account for the fact that although the kernel is defined on a continuous domain, the infections are simulated on a discrete grid. The results of this calculation are shown in Fig. 6.2. It is apparent that, depending on the kernel used and the transmission probability, the infection can travel far further in a single timestep than the characteristic width of the kernel, particularly in the case of the power law kernel.

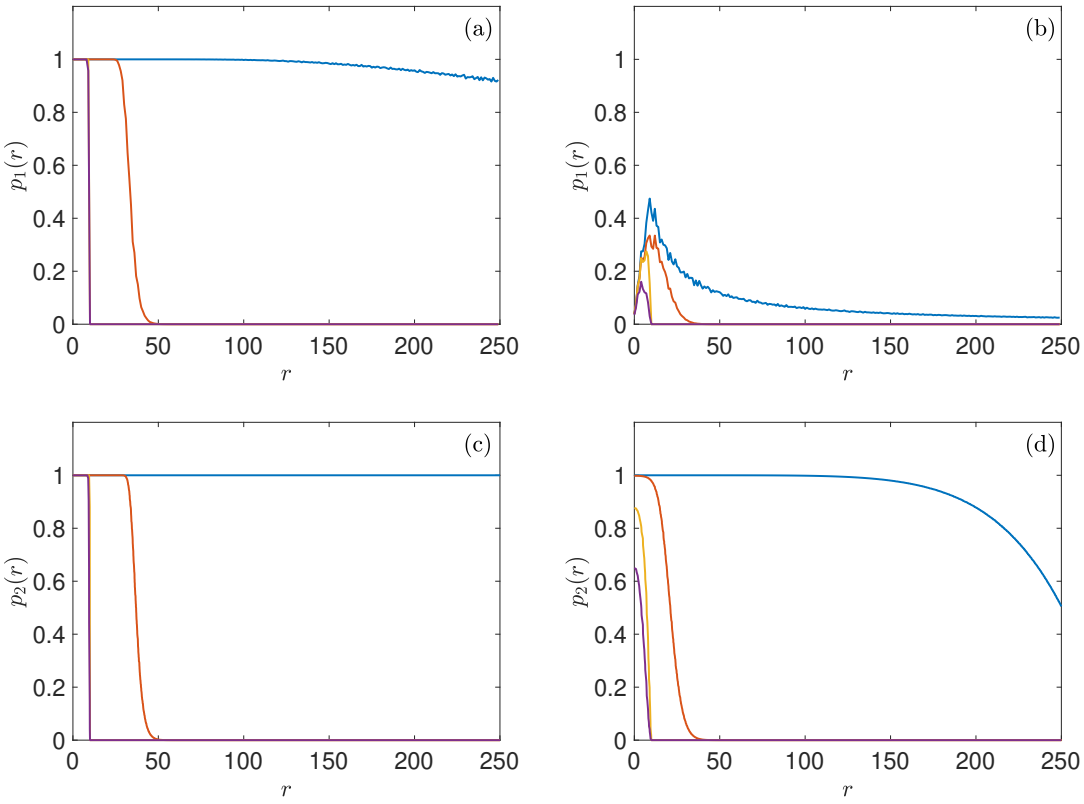

Figure 6.2: The probability of at least one tree at a distance r from a single infected tree becoming infected in a single timestep, calculated using Eq. (6.5)(panels (a), (b)), and the probability of at least one tree at a distance r or greater becoming infected in a single timestep, calculated using Eq. (6.6) (panels (c), (d)). Panels (a) and (c) correspond to $\beta = 1$, panels (b) and (d) correspond to $\beta = 0.01$, with $\sigma = 10$ for every distribution here. The lines show different kernels: linear (purple), circle (yellow), Gaussian (red), power law (blue). On panels (a) and (c) the linear and circle distributions coincide. The apparent noise on some curves is simply due to the discretisation of the domain.

Considering the maximum distance at which we can expect at least one new tree to become infected (Fig. 6.2 (c,d)), for both the linear and the circle kernels with $\beta = 1$ this is essentially σ, while for the Gaussian and the power-law it is far beyond σ. For $\beta = 0.01$, this becomes comparable to σ for the Gaussian, and much lower for the linear and circles kernels, while remaining far beyond σ for the power law.

In order to extract measures of the infection spread, we need a reasonably long time series to work with, as we expect our simulations to start with a transient period set by the initial conditions. Then, the longer the time series is, the better we can expect our statistics

to be. As evident from Fig. 6.2, we can expect the infection to travel much further from an infected cell in a single timestep if we use a power-law kernel than if we use one of the others, while keeping the value of σ constant between them, and it is reasonable to expect that the infection will then reach the edge of the forestscape sooner. Increasing the size of the grid is too computationally expensive to consider doing for more than a handful of runs, at which point we lose the advantage of taking statistics over many ensembles. On the other hand, we could try and mitigate having short time series by increasing the number of realisations to compensate. In practical terms, it would be best to make sure we have reasonably long time series, on the order of 100 timesteps.

We will focus our investigation on the Gaussian kernel. It has long tails, unlike the linear and circle kernels, which is appropriate for investigating non-local behaviour. Unlike the power-law there are no obvious complications (the effect of having a constant infection rate close to the origin in order to avoid a singularity; the potential difficulty in producing good quality data from long time series) - and there is only a single parameter to control, namely σ, which has a clear interpretation, while the power law has the exponent to choose. In this preliminary investigation, we have used $\alpha = 2$, but this is a somewhat arbitrary choice, and there is no obvious physical justification for making a particular choice for α, particularly without good experimental data on how these distributions are in reality. Although other choices of α might allow us to avoid the difficulty with $\alpha = 2$ of reaching boundaries almost instantaneously, the choice seems largely arbitrary. Furthermore, the choice of the Gaussian is a conventional choice for the investigation of spreading behaviour [91].

6.3 Verifying Numerics

We want to verify that our numerical simulations are propagating the infection correctly. To do this, we will initialise an infection, with parameters $\beta = 1$, $\sigma = 40$, with 3 infected cells at $x = 200$, $x = 400$, and $x = 600$ on an 800×800 forestscape with $\rho = 1$. We use multiple cells because we want to check not only that the distribution of infections from a single cell is as expected, but that the distributions of infections from multiple infected cells are interacting correctly. We then simulate the first timestep 10^4 times and count the frequency with which each cell along the line passing through the initially infected cells becomes infected, using this to estimate the probability of them becoming infected. In Fig. 6.3 (a) we illustrate the positions of the initially infected cells, and the typical distribution of cells infected in a single timestep. We can also calculate the theoretical probability of the cells on this line becoming infected:

$$P(x) = P_1(x) + P_2(x)\bar{P}_1(x) + P_3(x)\bar{P}_2(x)\bar{P}_1(x),$$

where $P_i(x)$ is the probability of the cell at position x being infected by the i^{th} initially infected cell in a single timestep, and $\bar{P}_i(x) = 1 - P_i(x)$ is the probability of it not being infected. In Fig. 6.3 (b) we plot the theoretical probability of the cells becoming infected at position x (dashed-black), and the frequency with which they become infected in our simulations (blue). We see that the outcomes from the averaged 10^4 simulations match the expected analytical results, and we are confident that we are propagating the infection correctly.

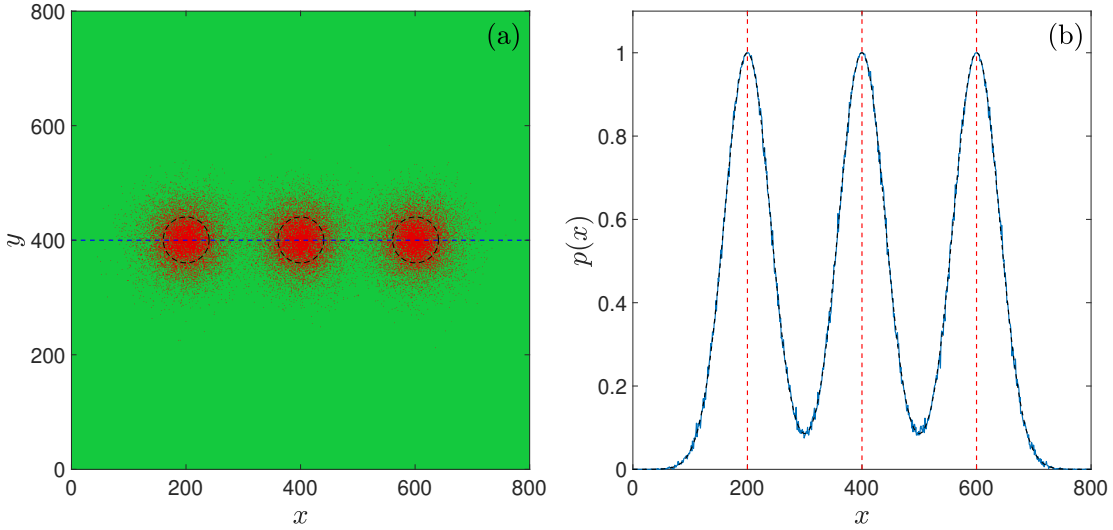

Figure 6.3: Panel (a): the forestscape that we test out numerics on, with three initially infected trees at $x = 200$, $x = 400$, and $x = 600$, all with $y = 400$, with susceptible cells shown in green and infected cells in red. The distance σ from each of them is indicated by a black dashed line, with a typical realisation of the cells infected after a single timestep shown. The line along which we calculate the analytic and simulated probablity is indicated by a blue dashed line. In panel (b) we show the analytical (black dashed line) vs. simulated (blue line) probabilities of the trees becoming infected at position x along the blue dashed line indicated in (a).

6.4 The Transient of the Infection Velocity

When analysing the velocity of the infection using the same techniques as for the local interactions, we observe a transient period during which the velocity increases, followed by a period in which the velocity is roughly constant, with some random noise, and finally, a period in which boundary effects become significant and the velocity decreases. We remove

both the initial transient period and the final period algorithmically in order to focus on the steady-state.

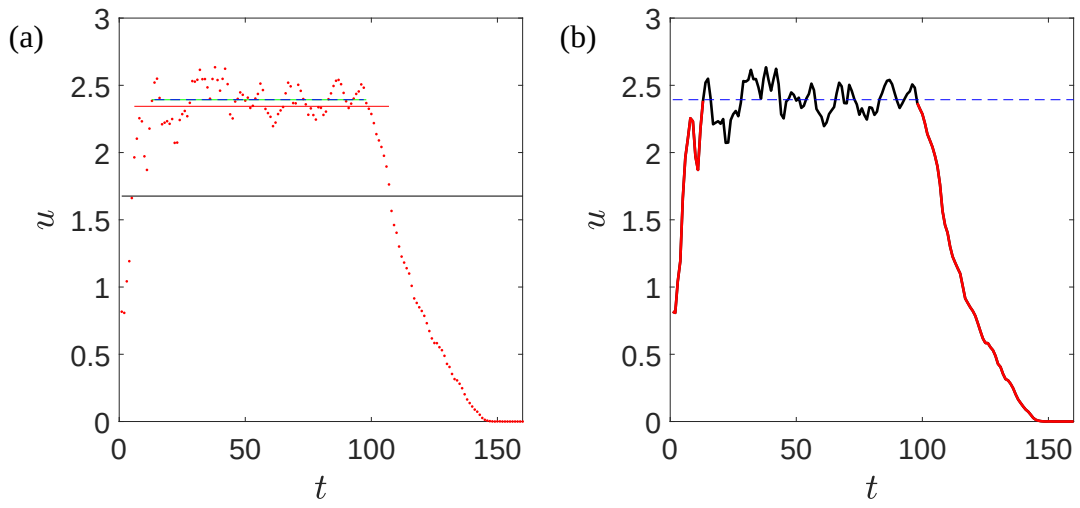

Figure 6.4: Illustration of the process used to remove the transient sections of the velocity time-series for one simulation on a homogeneous forestscape, with parameters $\sigma = 3$, $\rho = 0.3$, $\beta = 0.5$. Panel (a) shows the velocity vs. time (red points), the mean velocity (black), and the subsequent means found by discarding data below the previous means as described in the main text (red, blue, green lines), which converge. Panel (b) shows the time series again, with the sections discarded in red, the section used to calculate the mean in black, and the mean as a blue dashed line.

The algorithm used to remove the initial and final transient periods, illustrated in Fig. 6.4 is as follows: we calculate the mean velocity over the entire time series. We then identify the first and last times at which the velocity exceeds the mean, and discard the data outside this interval. We then calculate the mean velocity over this interval and again identify the first and last times at which the velocity is above this new mean, discarding the data outside this interval. This is repeated until the mean value found converges, and no more data is discarded in a single iteration.

We note that this approach is not guaranteed to converge: if the time series only has a very small section of non-transient data, or does not have a non-transient section, it will not return a value. However, when we inspect the regions of our data identified, they are consistent with what we would identify as non-transient by eye.

6.5 Initial Infection Rate

In order to explore the effect of non-local interactions in contrast to local interactions, we need to use consistent parameters. However, when we compare simulations with the same infection rate, it is apparent that far more trees become infected in a single timestep with non-local interactions. To make direct comparisons we work to make the initial infection rate ΔI, defined subsequently, the same between the two types of simulations so that we are not just seeing the effect of far more infections being created in a single timestep, but the effect of the non-local interactions. The initial infection rate ΔI is defined as the number of new infections that a single infected tree is expected to directly generate from the initial condition of the infection to the first time step. At the same time, we also want to make ΔI independent of the width σ of the Gaussian. The area under the infection probability curve $p(r)$, multiplied by the tree density ρ, gives us ΔI for the non-local infection. Introducing a correcting factor, which will depend on σ, will allow us to normalise the non-local infection as desired.

For the local case, ΔI is simply:

$$\Delta I = \rho\beta N, \tag{6.7}$$

where N is the number of neighbours considered by the algorithm. For the non-local case with a Gaussian kernel of width σ, we have:

$$\begin{aligned}
\Delta I &= \rho \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} p(x,y)dxdy \\
&= \rho \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \beta \exp\left(\frac{-(x^2+y^2)}{2\sigma^2}\right)dxdy \\
&= \rho\beta \int_{-\infty}^{\infty} \exp\left(\frac{-x^2}{2\sigma^2}\right)dx \int_{-\infty}^{\infty} \exp\left(\frac{-y^2}{2\sigma^2}\right)dy \\
&= \rho\beta(\sqrt{2\sigma^2\pi}) \cdot (\sqrt{2\sigma^2\pi}) \\
&= \rho\beta 2\sigma^2\pi.
\end{aligned} \tag{6.8}$$

To make ΔI for the non-local infection independent of σ, and equal to ΔI for the local infection using the Moore neigbourhood, with the same β, we redefine $p(r)$, to give our normalised non-local kernel:

$$p(r) = \beta \left(\frac{4}{\pi \sigma^2} \right) \exp \left(\frac{-r^2}{2\sigma^2} \right). \tag{6.9}$$

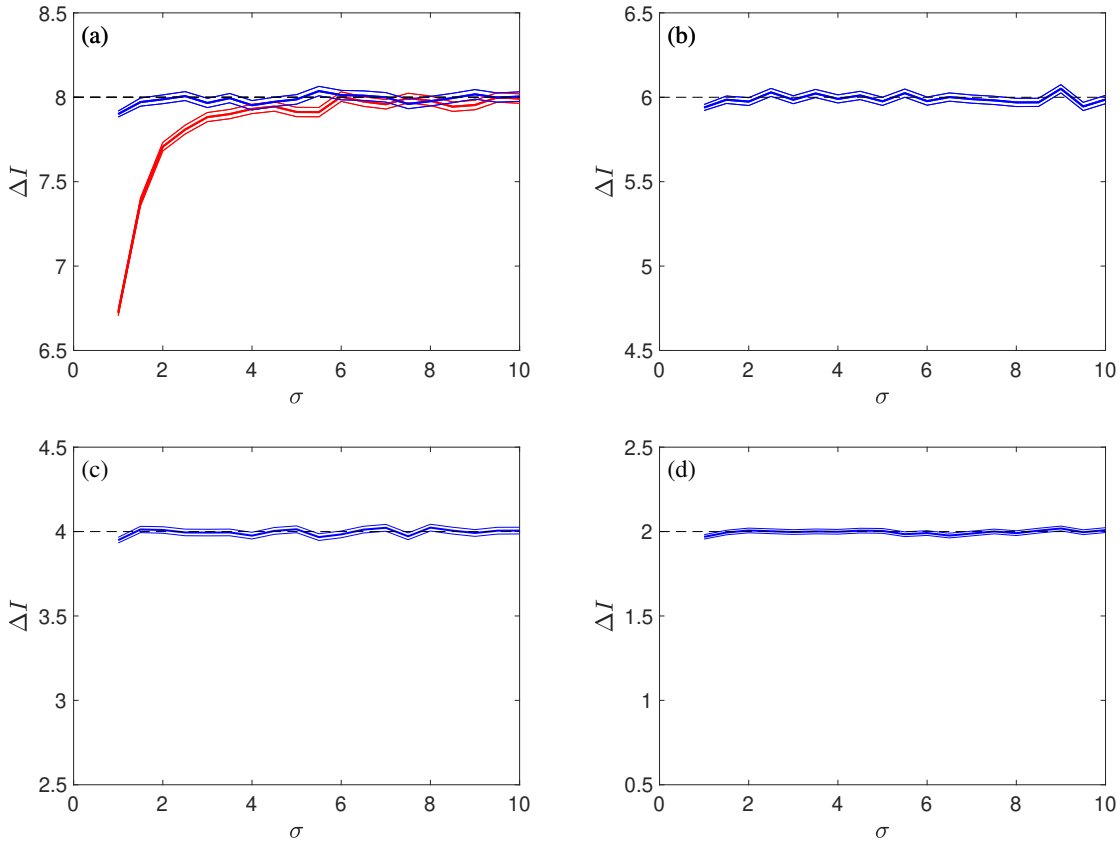

Figure 6.5: Initial infection rate ΔI estimated from 1000 ensembles using a non-local kernel for $\rho = 1$ (a), $\rho = 0.75$ (b), $\rho = 0.5$ (c), $\rho = 0.25$ (d), with $\beta = 1$ for all panels, for the uncorrected normalised kernel (red) and the corrected normalised kernel (blue), with the expected value for local simulations shown as a black dashed line.

In Fig. 6.5 (a) we show the initial infection rate as a function of σ, calculated as the number of trees infected by a single infected tree at the centre of the forestscape in a single timestep. We use $\beta = 1$ and $\rho = 1$, and so the corresponding ΔI for the 8-neighbourhood simulations is 8. Considering the red curve, which is derived from simulations using the normalised kernel, eq. (6.9), we observe that, for larger values of σ, our ΔI is in good agreement with the 8-neighbourhood simulations, indicating that the normalisation works.

However, the value for the normalised kernel falls increasingly short of the expected values as we move to smaller σ. It turns out that this is because of the discretisation of the forestscape. When we calculated the area under the curve above eq. (6.8) we treated the region as continuous, while it is actually on a discrete grid. Crucially, the region occupied by the infected tree should be excluded from the calculation to find the expected number of new infections accurately. Our corrected ΔI is then:

$$\Delta I = \rho\beta(2\sigma^2\pi - A_{\text{exc}}), \tag{6.10}$$

where the area excised is given by:

$$A_{\text{exc}} = \int_{-1/2}^{1/2}\int_{-1/2}^{1/2} \exp\left(\frac{-(x^2+y^2)}{2\sigma^2}\right)dxdy,$$

corresponding to the 1×1 [pixel]^2 region centered on the origin, which we evaluate numerically. Again matching this with the ΔI for local simulations, our corrected non-local kernel is:

$$p(r) = \beta\left(\frac{8}{2\pi\sigma^2 - A_{\text{exc}}}\right)\exp\left(\frac{-r^2}{2\sigma^2}\right). \tag{6.11}$$

We compare the initial infection rate with this correction against the uncorrected rate in Fig. 6.5 (a) where the uncorrected kernel is displayed in red and the corrected kernel displayed in blue, and see that the corrected version is much closer to the local simulations at small σ. We will use the corrected, normalised kernel in all subsequent calculations, referring to it as simply the normalised kernel. For Fig. 6.5 (b-d), we show ΔI as σ changes for different values of tree density, noting that the observed ΔI is in good agreement with the expected value for all densities.

In Fig. 6.6 we show the proportion of the total area which is excised. Below $\sigma = 1$, this proportion increases rapidly; for these values, only the tails of the distribution lie outside the already infected site, and so values of $\sigma < 1$ are not appropriate. In Fig. 6.7 we compare the velocity, calculated using eq. (4.2), for the unnormalised transmission probability, eq. (6.4), with the velocity for the normalised transmission probability, eq. (6.11). We observe that the velocity is considerably lower (by a factor of about 2) for the normalised transmission probability; however, the velocity still increases as σ increases, suggesting that this is also dependent on the non-local features of the transmission mechanism, and not just the initial infection rate. Comparing the averaged velocity $\langle\bar{u}\rangle$ to the local model, we observe that the values of $\langle\bar{u}\rangle$ for the local model increase as tree density increases.

For the non-local model, in addition to the tree density, the new parameter, which is the characteristic distance of the infection σ, plays a critical role. We observe that as σ increases, the averaged velocity increases.

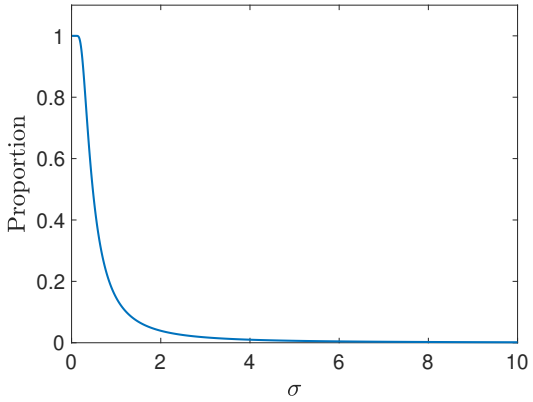

Figure 6.6: Proportion of the area under the Gaussian kernel which is removed by the correction.

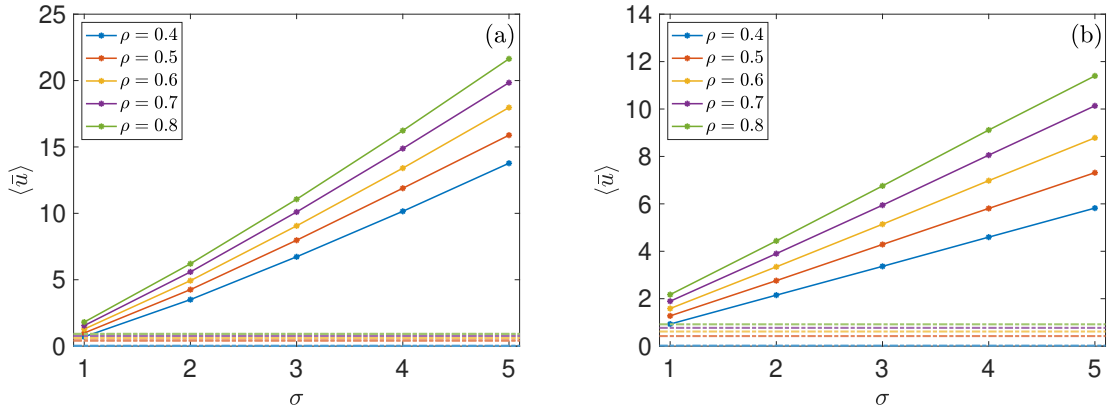

Figure 6.7: Infection velocity for unnormalised (a) and normalised (b) transmission probabilities, for 500×500 homogeneous forestscapes, with $\beta = 0.5$, for a range of ρ and σ. Note the different y-axis scales between the two panels. The dashed lines with the same color coding for the different densities represent the averaged velocity $\langle \bar{u} \rangle$ for the local model using the same parameters.

6.6 Examples of Infection Spreading

Later in this chapter, we will systematically compare our non-local simulations with the previous local simulations, again probing the effect of the key parameters. Before doing so, we show a number of individual simulations which highlight some interesting behaviour.

In Fig. 6.8, we show snapshots of two simulations, one with $\gamma = 0$, corresponding to a homogeneous forestscape, and one with $\gamma = 5$, corresponding to a heterogeneous forestscape, with all other parameters the same. We note that, although the mean distance of infected trees from the origin at a particular time appears to be roughly equivalent, there is considerably more anisotropy for larger values of γ.

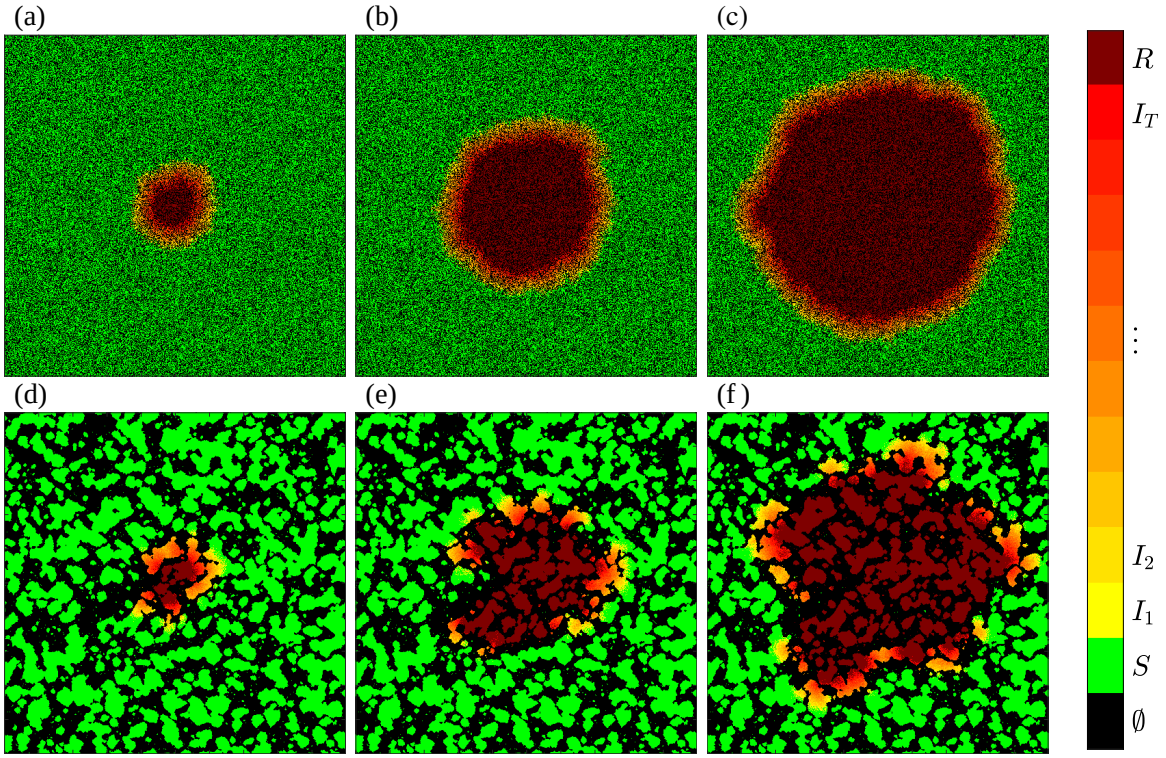

Figure 6.8: An example from our model of how the non-local infection propagates in our synthetic forestscapes. Using a non-local Gaussian kernel as the model of infection, we fix $\mathcal{L} = 500$, $\beta = 0.5$, $\sigma = 3$, $\rho = 0.5$, and we use two different forestscapes, one homogeneous with $\gamma = 0$ (a-c), and one hetereogeneous with $\gamma = 5$ (d-f). Snapshots are shown (from left to right) at times $t = 20$, 40, and 60.

In Fig. 6.9 we show the progression of three simulations in the same forestscape, with different values of σ. We note that the infection spreads faster for higher values of σ, even though they have the same initial transmission rate according to our normalisation of β

(see Sec. 6.5). We will discuss this effect in more detail in Sec. 6.7.1, and specifically for heterogeneous forestscapes in Sec. 6.7.4.

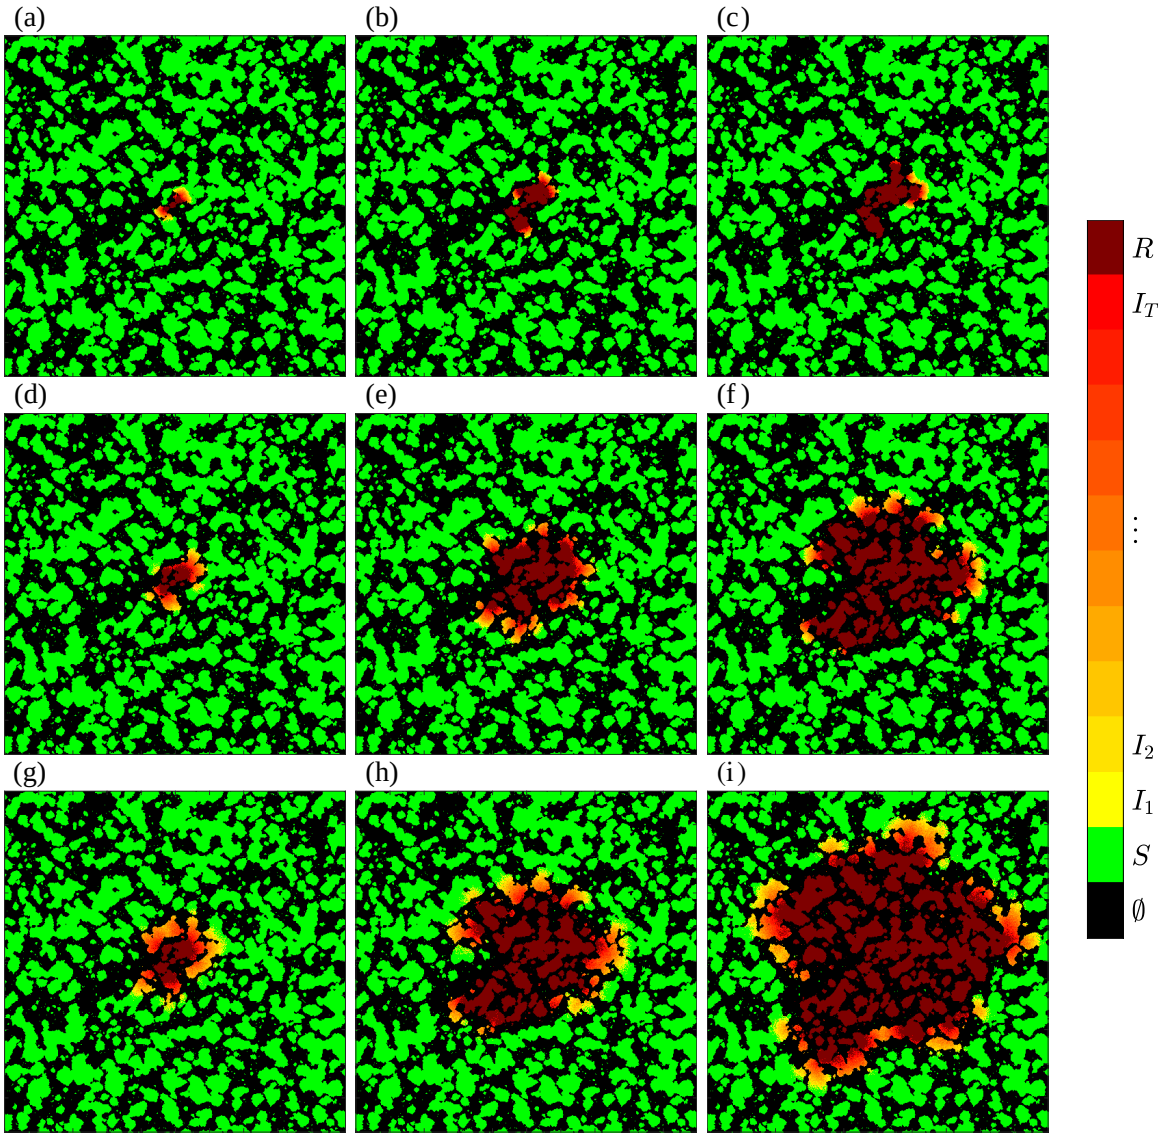

Figure 6.9: An example from our model of how the non-local infection propagates in our synthetic forestscapes for $\sigma = 1$ (a-c), $\sigma = 2$ (d-f) and $\sigma = 3$ (g-i) and everything else is fixed.

In Fig. 6.10 we examine three types of infection on a heterogeneous forestscape: one in which both σ and β have small values, one in which σ is large, but β is small, and one in which they are both large. Unsurprisingly, when they are both low, the infection dies out at an early stage, and when they are both high, the infection is rapid and has a high mortality rate. However, when σ is large, and β is small, we observe an interesting

phenomenon: although the infection manages to spread throughout the forestscape, with the exception of a few small regions, in all of the regions through which the infection passes some proportion of the trees survive, with trees on the edge of clusters more likely to survive than those in the centre of clusters.

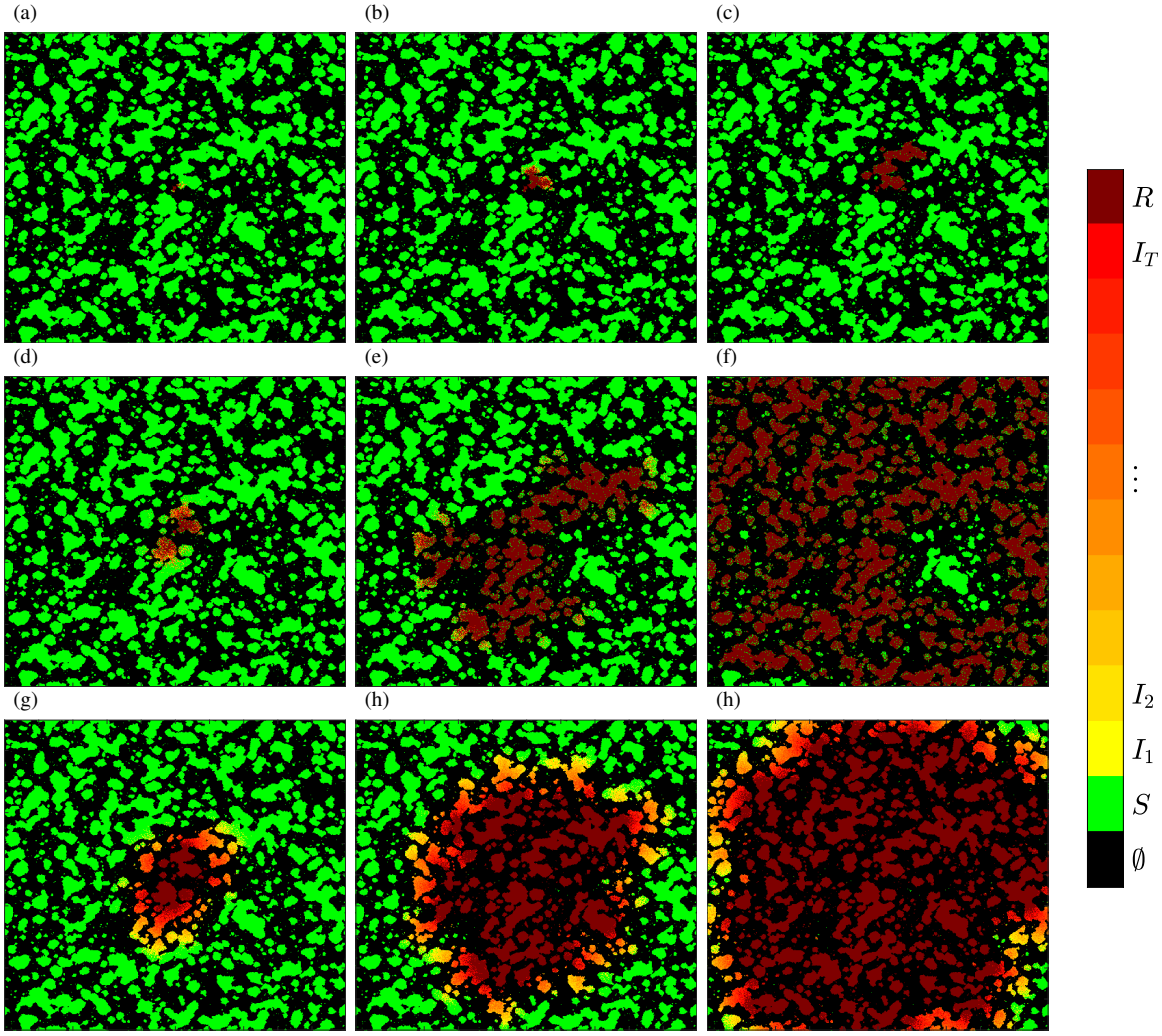

Figure 6.10: An example from our model of how the non-local infection propagates in our synthetic heterogenous forestscapes with $\gamma = 5$ as we change the transmission probability β and the Gaussian width σ . We fix $\mathcal{L} = 500$ and $\rho = 0.41$ (around the well-known Moore neighbourhood critical threshold). For (a-c), we use $\sigma = 2$ and $\beta = 0.05$, and show snapshots at times t=20, 60 and 200. While for (d-f), we use $\sigma = 5$ and $\beta = 0.05$ (i.e. we only increase σ), with snapshots at $t = 20$, 260 and 560. For (g-i), we use $\sigma = 5$ and $\beta = 0.5$, with snapshots at $t = 20$, 40 and 60.

We quantify the observed phenomenon when σ is large and β is small in Fig. 6.11 (a). For each populated cell (i.e. susceptible or dead at last time step) we calculate the

proportion of the surrounding region which is populated, simply by counting the number of occupied cells within a chosen radius. We then bin the trees based on this proportion and count the proportion of trees in each bin which are left alive after the infection has died out. We observe that less than 20% of tree with $> 90\%$ of the surrounding region populated survive, while in contrast, more than 60% of trees with $< 20\%$ of the surrounding region populated survive. In panels (b) and (c), we use the same forestscape as panel (a), and we only increase β. The results of Fig. 6.11 are consistent with what we see in Fig. 6.10.

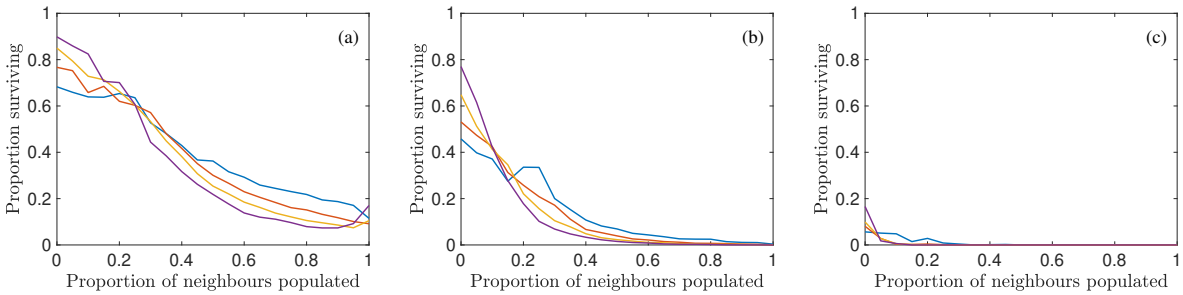

Figure 6.11: Proportion of trees which survive, as a function of the proportion of the region around them which is populated, calculated as described in the main text. We use different radius (3, 5, 7, 9 corresponding to blue, red, yellow, and purple lines respectively) to calculate the proportion of the surrounding trees which are populated. We note that the broad features are consistent whichever radius we use to calculate the populated proportion of the surrounding region. We have used $\beta = 0.05$, 0.1 and 0.5 for panel (a), (b) and (c) respectively.

6.7 Effect of Parameters on Infection Characteristics

We characterise the observed dynamics in our forestscapes through measuring the spreading velocity of the pathogen. We also consider other statistics such as the number of infected trees at each time step, mortality, and the PDF of the distance to the nearest infected tree for the newly infected trees.

Using these measures, many parameters need to be investigated: the transmissibility of the pathogen β, the tree density ρ, the level of heterogeneity γ, and the characteristic width of the non-local kernel σ. We shall also examine the combined effect of these different parameters. Here we are mostly focusing on parameter values in the epidemic regime, while later on, we will explore the behaviour, including confined regimes, over a larger parameter space.

6.7.1 Characteristic Width of the Non-local Kernel

The characteristic width σ is a new parameter that we use to model the non-local infection, governing how far the infection is able to spread in a single timestep. We show the averaged spreading velocity, calculated using eq. (4.2), vs σ for a range of transmissibilities β in Fig. 6.12. Panel (a) shows the curves on a linear scale, and we note that the velocity appears to scale linearly with σ: this is confirmed in panel (b), which shows the same curves on a logarithmic scale, with the scaling holding for all values of β, with the exception of the lowest values ($\beta \simeq 0.2$, $\beta \simeq 0.3$) where the scaling seems to be failing for the lowest values of σ.

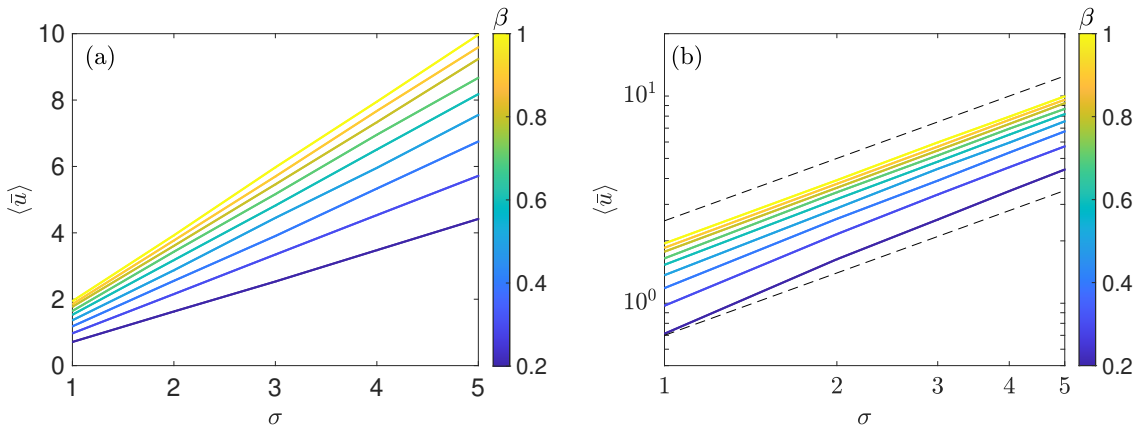

Figure 6.12: Effect of the characteristic width σ of the non-local kernel on the infection velocity, on a linear scale (a) and on a logarithmic scale (b), with dashed black lines indicating linear scaling. This shows the relation for a range of values of β, with tree density $\rho = 0.5$, and the heterogeneity parameter $\gamma = 0$, i.e. homogeneous forestscapes.

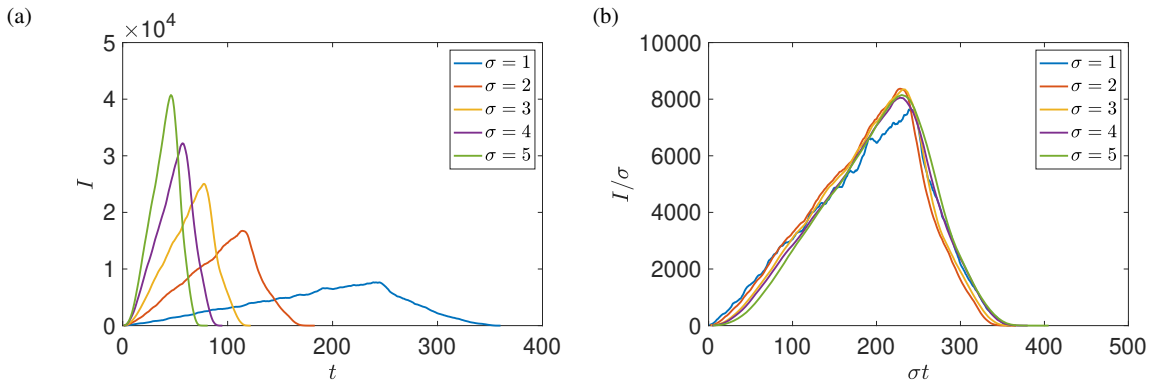

Figure 6.13: The number of infected cells vs. time for a simulation with $\beta = 0.5$, $\rho = 0.5$, and $\gamma = 0$. Panel (a) shows the raw values, while panel (b) shows the collapse onto a single curve under the scaling $I \to I/\sigma$, $t \to \sigma t$.

In Fig. 6.13 (a) the number of infected trees is plotted as a function of time for a range of values of σ. We see that as σ decreases, the rate of infection is lower, and the maximum value of the curve is considerably smaller. Note that the decrease of the curves at late times is due to the infection reaching the boundaries. In panel (b) we show the collapse of the curves onto a single curve when we scale $t \to \sigma t$ and $I \to I/\sigma$, which we find empirically.

We note that the transmission probability is normalised, such that the initial infection rate is independent of σ. However, despite this, we still see strong scaling with σ; there is something intrinsic about increasing long-range effects that increases the rate at which the infection spreads through the domain.

6.7.2 Transmissibility of the Pathogen

For the local infection, the single effect of the transmissibility of the disease β is a displacement of the critical tree density for the transition to epidemic, towards higher values of tree density. In this section, we would like to investigate how the move to non-local interactions influences the effect of β.

In Fig. 6.14 (a), we plot the (time- and ensemble-averaged) velocity as a function of β for various σ. We see that, for all values of σ, as β increases the velocity of the infection increases, most rapidly for the smallest values of β. Using the scaling $\langle \bar{u} \rangle \propto \sigma$ identified in the previous section, we show $\langle \bar{u} \rangle / \sigma$ in panel (b) and observe all of the curves collapsing onto a single line. We thus conclude that the effect of the transmissibility on the infection velocity is independent of the characteristic scale σ of the infection kernel. Note that the

transmissibility rate is normalised by σ as discussed in Sec. 6.5; however, the infection velocity is still dependent on σ: this is purely a result of the non-local interactions. We considered the curves on logarithmic axes and saw no obvious simple form for the curves.

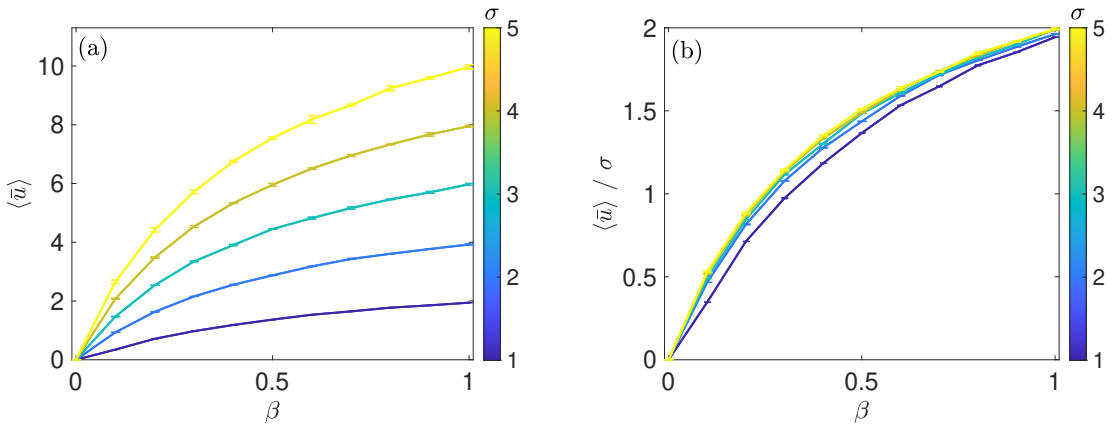

Figure 6.14: Effect of the transmissibility of the pathogen on the infection velocity. Panel (a) shows this relation for a range of values of σ, with the tree density $\rho = 0.5$, and the heterogeneity parameter $\gamma = 0$, i.e. homogeneous forestscapes. Panel (b) shows the collapse of the curves in panel (a) onto a single curve when we scale by the characteristic width of the infection kernel σ.

In Fig. 6.15 (a), we plot the number of infected as a function of time. As β gets larger, we have a much larger gradient and a much more rapid rate of infection, but it hits the boundaries sooner because it is faster. Then, in panel (b) we show how the number of infected cells scales with the transmissibility of the pathogen. For all the different values of β with all other parameters being constant, we empirically find that scaling $t \to \sqrt{\beta}t$ and $I \to I/\sqrt{\beta}$ collapses all the curves onto a single line. We note that for the lowest value of β, the curve falls somewhat lower than the other scaled curves. We investigate some possible causes of this in Sec. 6.7.5.

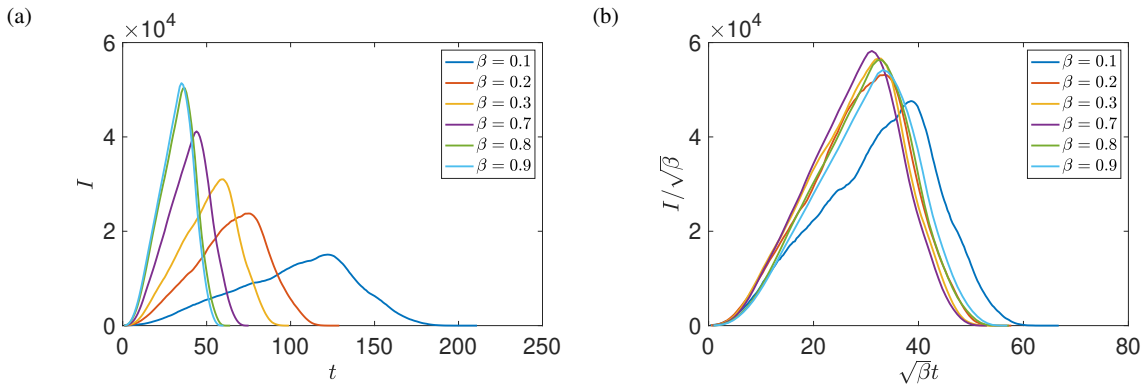

Figure 6.15: The number of infected trees as a function of time for $\sigma = 4$, $\rho = 0.5$, $\gamma = 0$ (a), and the same data scaled with $t \to \sqrt{\beta}t$ and $I \to I/\sqrt{\beta}$ (b).

In Fig. 6.16, we show two examples of a forestscape at late stages of an epidemic. In panel (a) $\beta = 0.1$, and in panel (b) $\beta = 0.2$. Under the lower value of β, many surviving clusters of cells are evident, whereas for the higher value of β, although there are isolated surviving cells, the proportion is less significant. Looking at the interface between the susceptible and removed cells, it is clear that the lower value of β allows a less uniform interface to develop, resulting in isolated clusters, which then have some protection from the epidemic due to this isolation, combined with the already lower value of β.

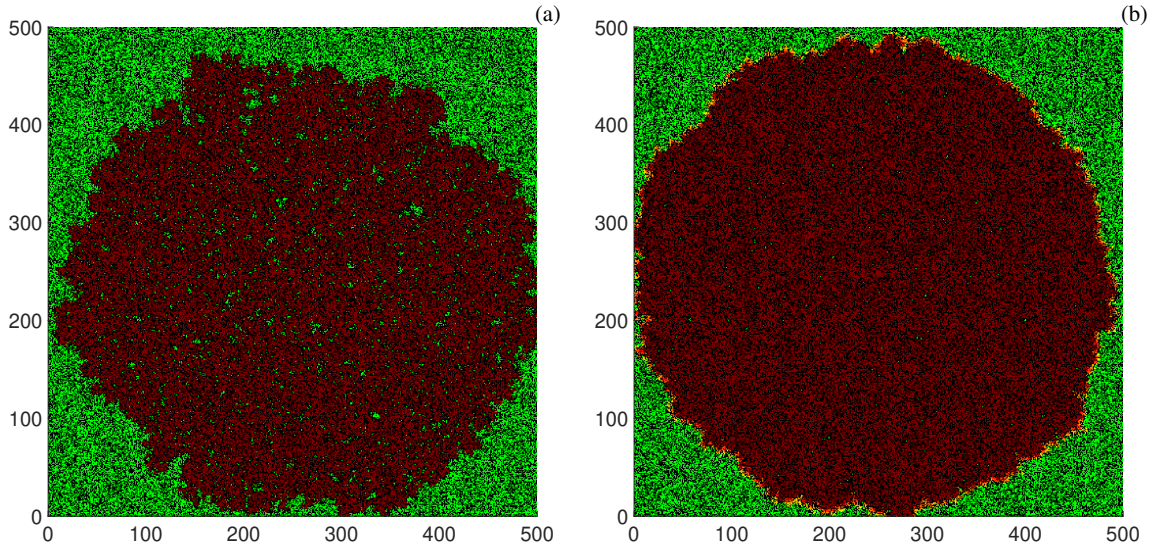

Figure 6.16: Status of cells at a late stage of an epidemic for $\beta = 0.1$ (a) and $\beta = 0.2$ (b), with $\sigma = 1$, $\rho = 0.5$, and $\gamma = 0$.

6.7.3 Forestscape Density

For the local infection discussed in Chapter 4, the most important parameter is tree density. A disease transforms from a confinement state to an outbreak state if and only if tree density is sufficiently large. Considering the homogeneous forestscapes, the critical density for the transition to an epidemic also depends on the method of infection. For example, if we use the von Neumann neighbourhood, the critical density is ≈ 0.60 while if we use the Moore neighbourhood the transition is ≈ 0.41. For heterogeneous forestscapes discussed in Chapter 5, this transition also depends on the heterogeneity level.

In this section, we examine how the dynamics of the infection depend on the forestscape density for non-local infections. In this section, we are only interested in homogeneous forestscapes. Heterogeneous forestscapes will be discussed in more detail in the next section. Then, we will discuss the transition to epidemic in Sec. 6.7.6, as it is strongly influenced by some of the other parameters in our model.

In Fig. 6.17, we show the infection velocity as a function of ρ for a range of values of the characteristic width of the infection. As the tree density increases the velocity increases in a manner that appears close to linear; however, we note that the gradient is slightly reduced at the highest values of ρ. This is more pronounced at the lowest values of σ, with the curve for $\sigma = 1$ also exhibiting a kink at around $\rho = 0.2$. Attempting to find a scaling for $\langle \bar{u} \rangle$ with ρ empirically yielded no clear simple relation. We do not see a phase transition for u, can we see a phase transition for the mortality μ_1? We will investigate this in more detail in Sec. 6.7.6.

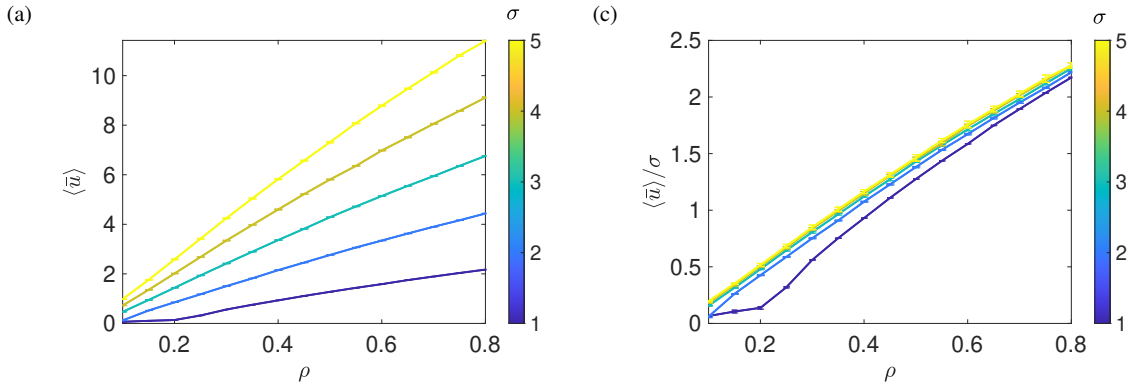

Figure 6.17: Effect of the forestscape density on the infection velocity. Panel (a) shows this relation for a range of values of σ, with the transmissibility rate $\beta = 0.5$, and the heterogeneity parameter $\gamma = 0$, i.e. homogeneous forestscapes. Panel (b) shows the same curves as panel (a) with the averaged velocity divided by σ.

In Fig. 6.18, we illustrate how the number of infected cells scales with the tree density. We empirically find that, for different values of ρ and constant σ, scaling $I \to I/(\rho^{3/2})$ and $t \to \sqrt{\rho}t$ collapses all the curves onto a single line.

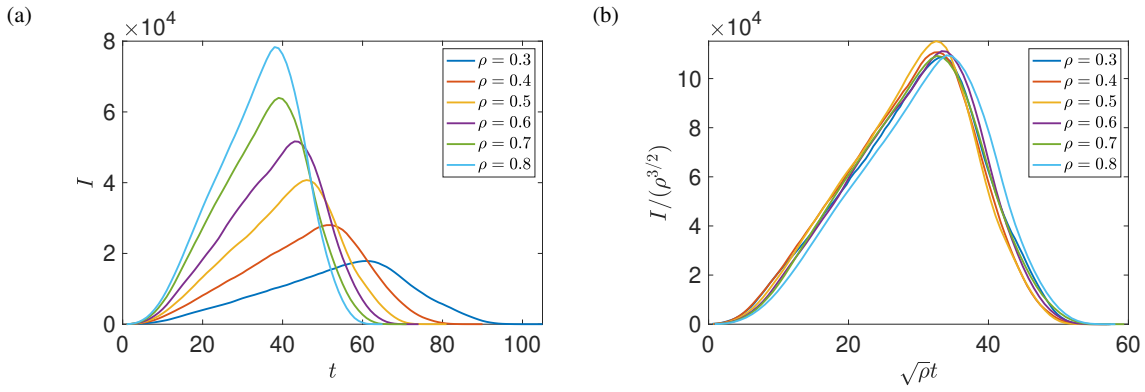

Figure 6.18: Panel (a): number of infected trees vs. time for homogeneous forestscapes with range of tree densities and $\sigma = 5$. Panel (b): curves scaled according to $I \to I/(\rho^{3/2})$, $t \to \sqrt{\rho}t$.

6.7.4 Forestscape Homogeneity/Heterogeneity

In Chapter 5, when we simulate the spread of disease on heterogeneous forestscapes using local interactions, we found that heterogeneity suppresses the disease propagation. We now examine how heterogeneity affects the dynamics of the infection when we have non-local interactions.

In Fig. 6.19, we show the effect of changing β when $\gamma = 5$ (i.e. for heterogeneous landscapes). We see that the velocities have similar behaviours as the homogenous forestscapes (cf. Fig, 6.14), but the scaling for σ only works for the higher values of σ. We interpret this as the non-local interactions enabling the infection to "jump" over the gaps in the clustered distribution for the higher values of σ, but being unable to do so for the lower values of σ, as the gaps become larger than the characteristic spreading distance.

We highlight this behaviour in Fig. 6.20, in which we plot the number of infected trees against time for heterogeneous forestscapes for $\sigma = 1$ (panel (a)) and $\sigma = 5$ (panel (b)). Firstly, we note that the peak number of infected cells is much greater for the $\sigma = 5$ case, and is reached much sooner. In Fig. 6.20, we show the number of infected trees as a function of time for different β. For Fig. 6.20, we attempted the scaling identified in Fig. 6.15, we note that the curves follow the same scaling for $\sigma = 5$, but they fail for $\sigma = 1$. Far more fluctuations are evident for the lower σ infections, which is again attributable

to the heterogeneity. From the perspective of the infection, with small σ the forestscape will appear extremely granular, preventing consistent transmission, while for the large σ infections the spatial fluctuations below the characteristic spreading distance will appear more like a random field, allowing smooth transmission.

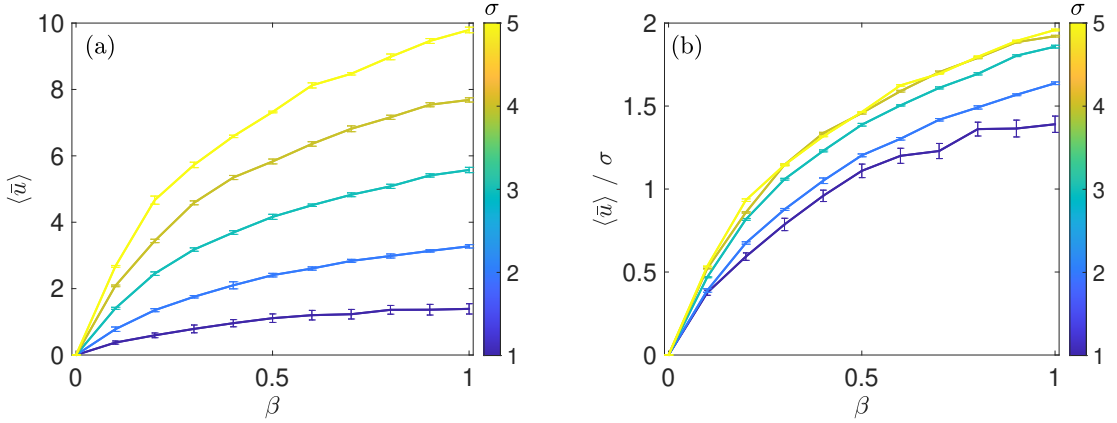

Figure 6.19: Effect of the transmissibility of the pathogen on the infection velocity. Panel (a) shows this relation for a range of values of σ, with the tree density $\rho = 0.5$, and the heterogeneity parameter $\gamma = 5$, i.e. heterogeneous forestscapes. Panel (b) shows that the collapse of the curves in panel (a) onto a single curve, when scaled by the characteristic width of the infection kernel σ, happens only for larger σ, in contrast to the homogeneous forestscapes (cf. Fig. 6.14 (b)).

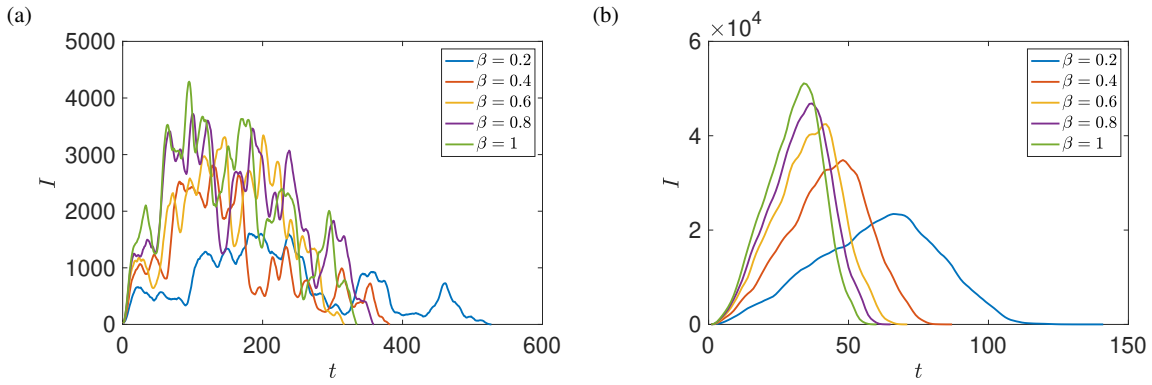

Figure 6.20: The number of currently infected trees plotted as a function of time for $\sigma = 1$ (a), and $\sigma = 5$ (b), with $\rho = 0.5$, and $\gamma = 5$ (hetereogeneous forestscapes).

We examine this from another perspective in Fig 6.21, in which we show how the

number of infected trees vs time is affected by changing the heterogeneity parameter, for a $\sigma = 2$ (panel (a)) and $\sigma = 5$ (panel (b)). For $\sigma = 5$ there is no significant difference evident between the different levels of heterogeneity, although by $\gamma = 5$ the curve has begun to show some fluctuation. However, for $\sigma = 2$, fluctuations onset for much lower values of γ, and for the highest values, the infection is suppressed, with a much lower peak (although the ultimate mortality rate is still unchanged here).

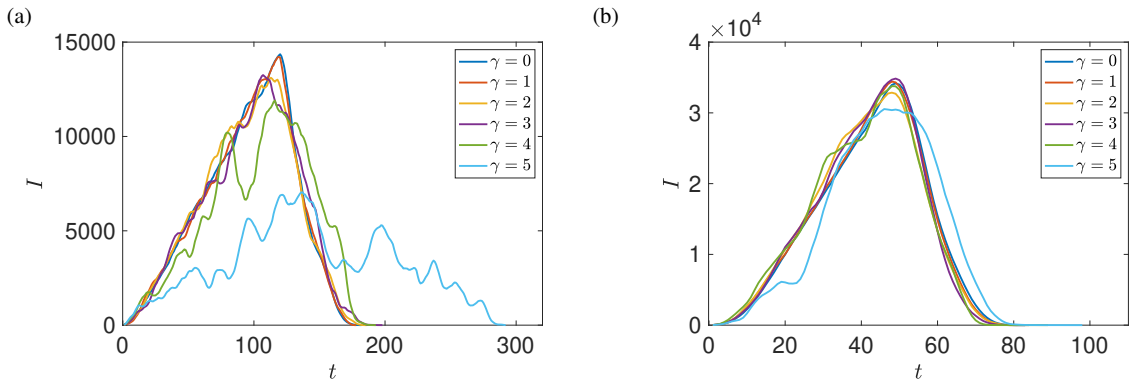

Figure 6.21: Number of currently infected trees as a function of time for various values of γ, with $\rho = 0.44$ and $\sigma = 2$ (a), $\sigma = 5$ (b). We see large fluctuations in the number of infected cells in panel (a), and note that increasing σ reduces the fluctuations in panel (b).

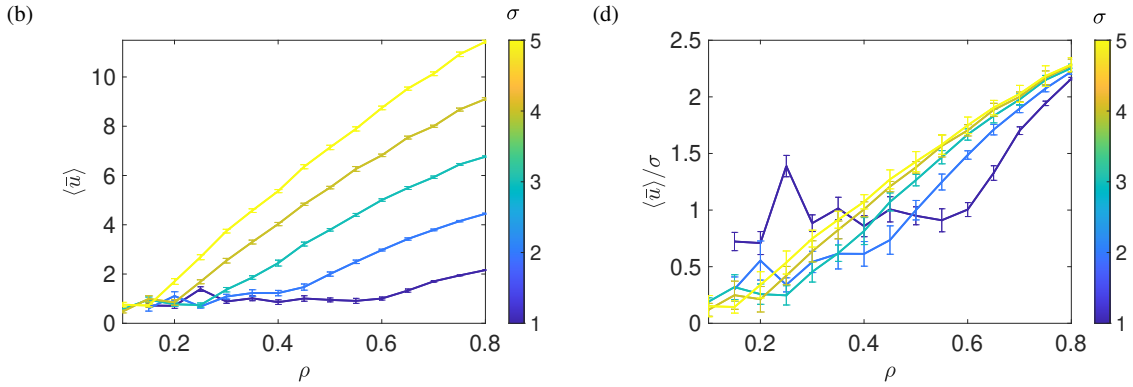

Figure 6.22: Effect of the forestscape density on the infection velocity. In panel (a) we show how the propagation velocity changes with the density for $\gamma = 5$ (heterogeneous), for different σ. In panel (b) we show the same curves scaled by σ.

In Fig. 6.22 we show how the infection velocity is affected by the forestscape density for

heterogeneous forestscapes, for comparison with Fig. 6.17 which gives this information for homogeneous forestscapes. We observe that the velocity is suppressed at lower densities, with the density at which this suppression is overcome being highest for the lowest values of σ.

We summarise this in Fig. 6.23, which shows how the tree density ρ and non-local infection width σ impact the propagation velocity.

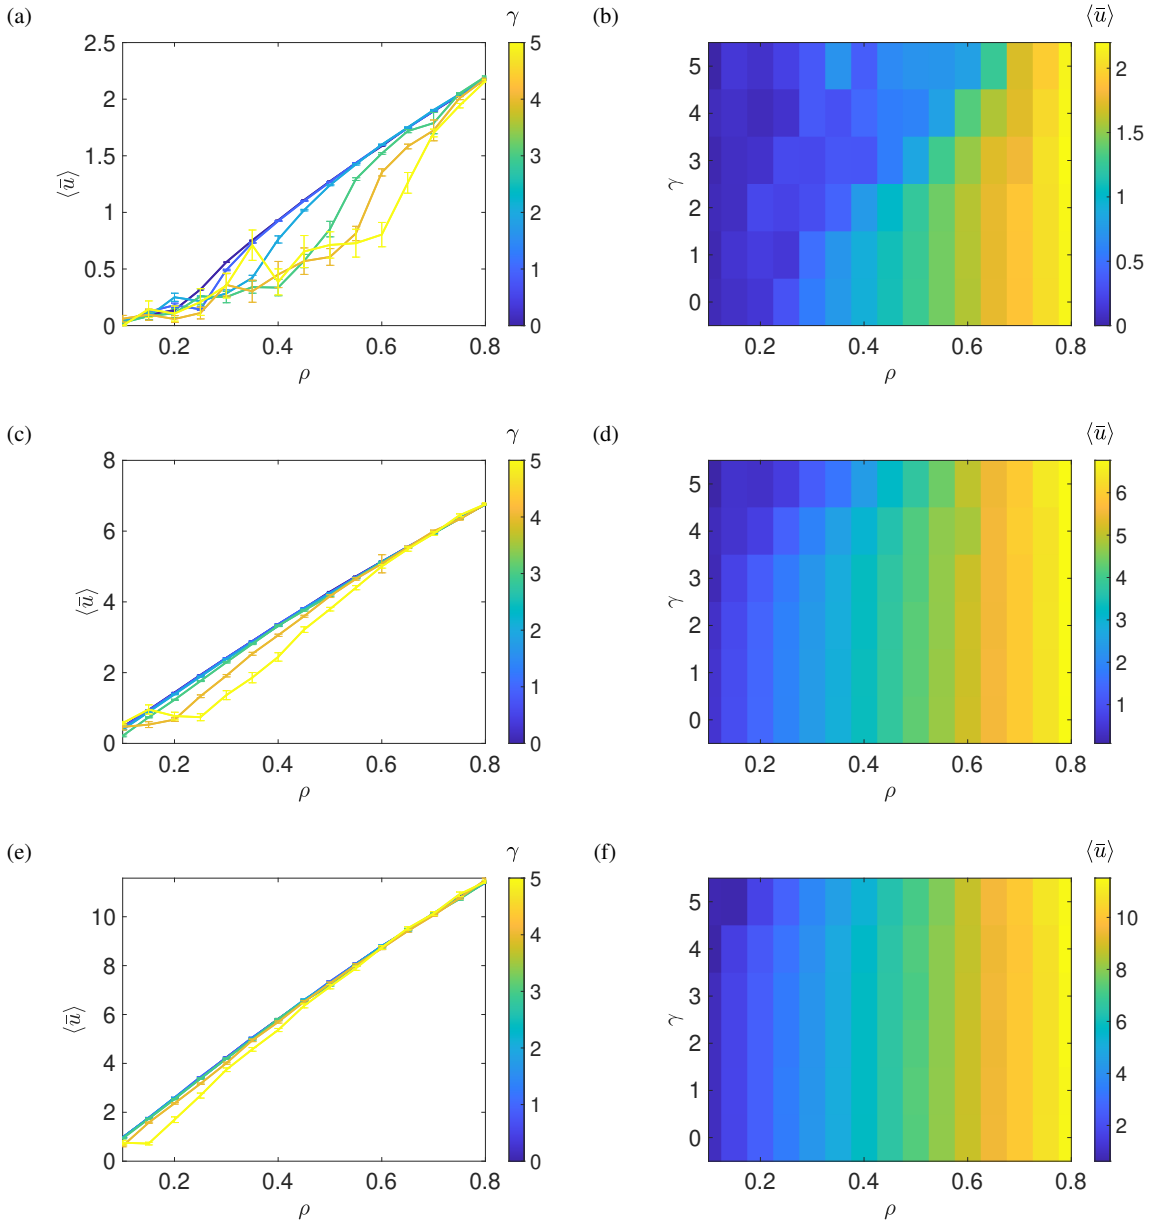

Figure 6.23: Panels (a), (c), and (e): the averaged spreading velocity of the infection $\langle \bar{u} \rangle$ for different heterogeneity levels of forestscapes is plotted against tree density ρ for transmission probability $\beta = 0.5$ and non-local infection width $\sigma = 1$, 3 and 5 ((a), (c) and (e) respectively). Panels (b), (d) and (f) are the associated phase diagrams of the infection velocity.

The general trend is that higher values of γ suppress the infection velocity. However, we observe two regions where this is disrupted. Firstly, this trend is less noticeable as we move to higher values of σ and is barely noticeable for $\sigma = 5$. A physical explanation for the

reduced velocity for more heterogeneous forestscapes is that, as well as cells being clustered, larger gaps emerge, which the infection is unable to spread over. The diminished effect is easily understood here as the consequence of the characteristic scale of the non-local kernel becoming comparable to or larger than the scale of the gaps.

We also note that the velocity is not suppressed for the highest densities, even for the smallest values of σ. Again, this can be understood if the suppression is linked to gaps in the forestscape: if the density is high enough, large gaps are less frequent or no longer present at all, and no longer significantly suppress the infection velocity.

6.7.5 Universal Spreading

To find out the difference between simulations which follow the universal curve and those which don't, it is instructive to visualise them. Do they look different? If so, is this difference true for other simulations that do/don't fit the universal curve? Therefore, in the following sections, we show three examples that investigate the different parameters: σ, ρ and β for different γ.

Different Characteristic Width

In Fig. 6.24, we display some examples that compare the dynamics for a simulation that follow the universal curve and ones that don't. The first row is for homogenous fields (i.e. $\gamma = 0$), and the second row is for heterogeneous fields with $\gamma = 5$. Here we increase the characteristic width σ as we go from the right to the left. We have $\sigma = 1$, $\sigma = 3$ and $\sigma = 5$ for the first, second and third columns respectively. We note that as σ increases the width of the front of the infection increases. For the homogeneous forestscape, the infection spreads evenly in all directions. However, for the heterogeneous forestscapes and with the lowest infection width $\sigma = 1$, the infection dies out at an early stage, and we have a confined disease. For $\sigma = 3$, we are in the epidemic regime, but the infection propagation becomes less circular and tends to have some surviving clusters. Again, this shows that heterogeneity suppresses the infection propagation, but as σ increases, this becomes less evident.

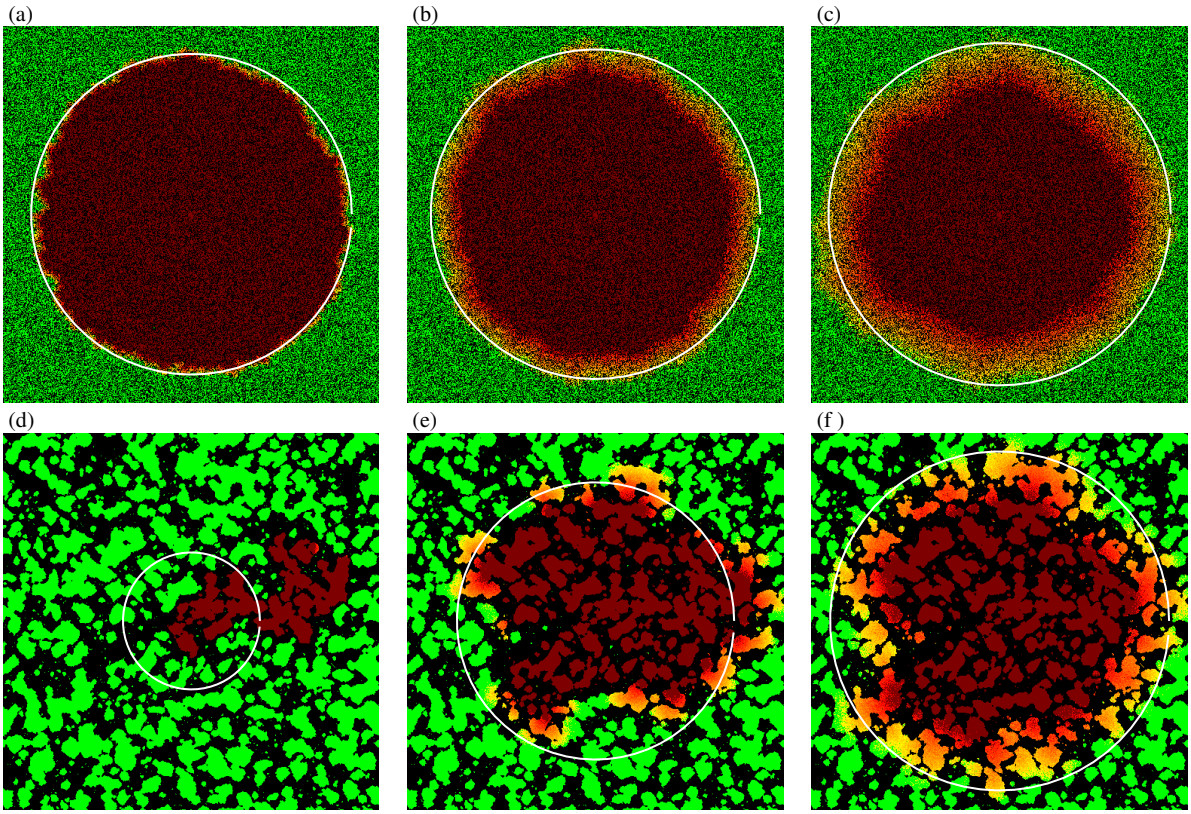

Figure 6.24: Late-stage infection configurations. We fix $\beta = 0.5$ and $\rho = 0.5$. The top row has $\gamma = 0$ (i.e. homogeneous forestscape), the bottom row has $\gamma = 5$ (i.e a heterogeneous forestscape). We vary the characteristic width σ to be 1, 3 and 5 in the first, second and third columns respectively. The white circles have the same area as the infected/removed cells (factoring for density), to indicate how far from circular the infection configurations are.

In Fig. 6.25, we show the number of infected as a function of time for the simulations in Fig. 6.24. We scale with respect to σ. This will allow us to see the difference between dynamics following the universal scaling and the ones that do not. We see that all the curves for the homogeneous forestscape, which spread evenly in all directions, follow the universal scaling. For the heterogeneous forestscape, the situation is different. For example, considering the dynamic in Fig. 6.24 (d), the infection dies out at an early stage. Also, the associated curve in Fig. 6.25 does not follow the universal scaling. On the other hand, some of the curves for the heterogeneous forestscapes such as $\sigma = 3$, we know that we are in the epidemic region, but the curve does not scale very well with other curves, with large deviations evident. This is associated with what we see in the dynamic in Fig. 6.24 (e) where the gaps cause suppression to the dynamics in some of the directions. While for $\sigma = 5$, the associated curve follows the universal scaling more closely because the infection

can spread over larger distance at each time step and can transmit over the gaps.

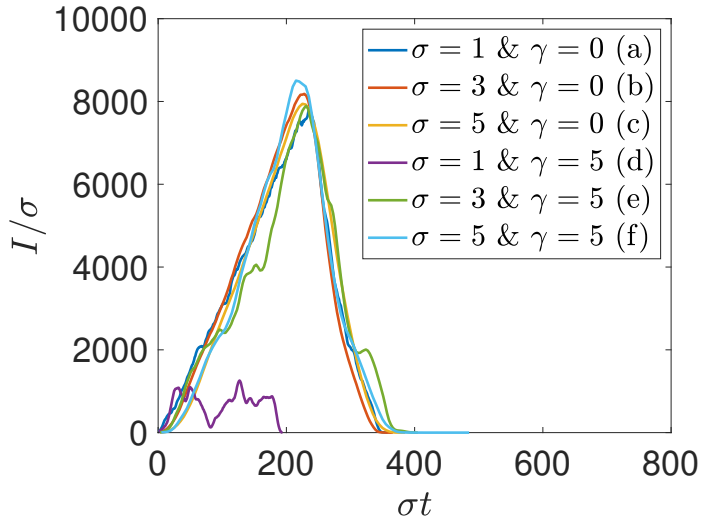

Figure 6.25: The number of infected as a function of time for the simulations shown in Fig. 6.24, normalised for σ.

We may conclude that infections which follow the universal scaling spread evenly in all directions; diseases which don't follow the universal scaling do not spread evenly, and tend to have surviving clusters at the end.

Different Tree Densities

In Fig. 6.26, we show some examples that compare the dynamics for some simulations that follow the universal curve and some that don't. The first row is for homogeneous fields (i.e. $\gamma = 0$), and the second row is for heterogeneous fields with $\gamma = 5$. We decrease the tree density as we go from left to right. We note that as ρ decreases, the infection propagation becomes less circular. This is more evident for the heterogeneous forestscape with $\gamma = 5$. For the heterogeneous forestscape and the lowest tree density, shown in Fig. 6.26 (f), the infection dies out at an early stage. This suppression can be overcome by increasing σ or ρ.

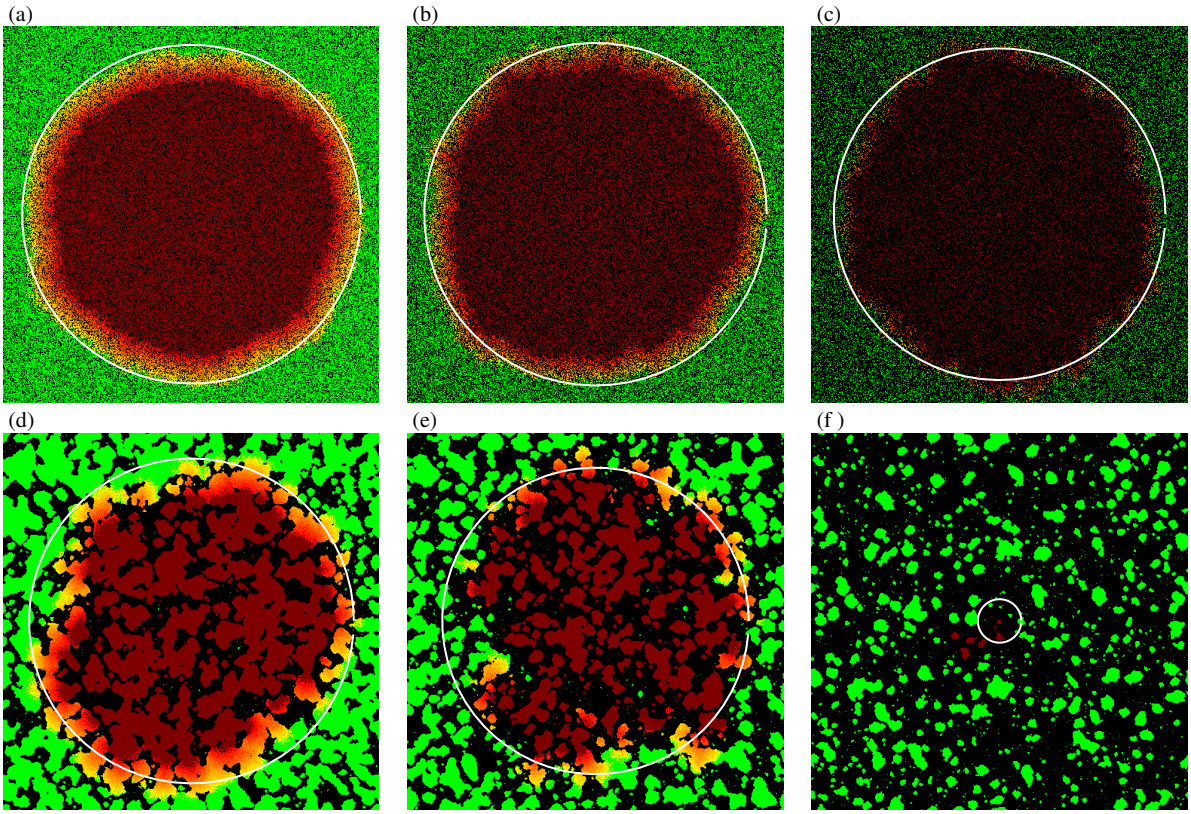

Figure 6.26: Late-stage infection configurations. We fix $\beta = 0.5$ and $\sigma = 3$. The top row has $\gamma = 0$ (i.e. homogeneous forestscape), the bottom row has $\gamma = 5$ (i.e a heterogeneous forestscape). We vary the density ρ to be 0.65, 0.45 and 0.20 for the first, second and third columns respectively. The white circles have the same area as the infected/removed cells (factoring for density), to indicate how far from circular the infection configurations are.

In Fig. 6.27, we display the number of infected trees as a function of time for all the simulations shown in Fig. 6.26. We note that for the lowest tree density and homogeneous forestscape, the associated curve for the number of infected does not follow the universal scaling. For the heterogeneous forestscape, the infection dies out at an early stage and does not spread to the boundaries; consequently, the number of infected does not follow the universal scaling. For high tree density such as $\rho = 0.65$, both homogeneous and heterogeneous forestscapes follow the universal scaling. For medium density such as $\rho = 0.45$, we note that the curve for the homogeneous forestscape follow the universal scaling while the curve for the heterogeneous forestscape does not follow it very well.

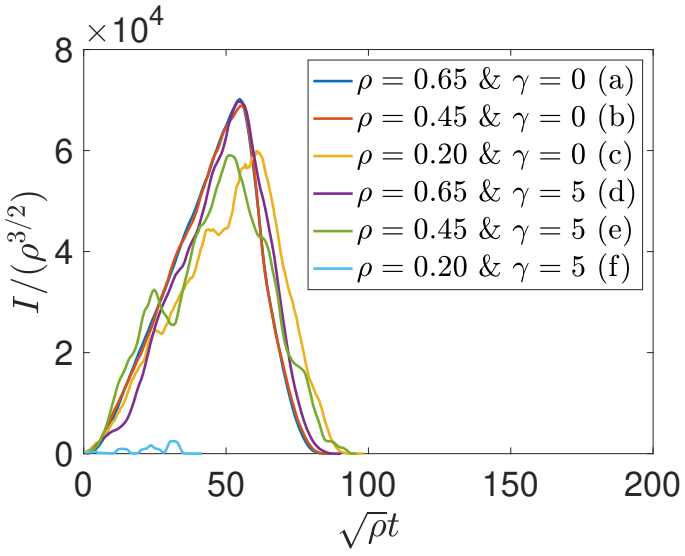

Figure 6.27: The number of infected as a function of time for the simulations shown in Fig. 6.26, normalised for density.

Different Transmissibility

Recall that in Sec. 6.7.2, we have shown that the scaling for β works only for $\beta > 0.2$. In Fig. 6.28 we increase the transmissibility β as we go from left to right. The first row is for homogeneous fields (i.e. $\gamma = 0$), and the second row is for heterogeneous fields with $\gamma = 5$. We observe that there is a small difference between $\beta = 0.1$ and $\beta = 0.5$, while there is a non-negligible difference between them and $\beta = 0.05$. The general trend is that as we increase β, the infection can spread more evenly and slightly faster. Also, we can clearly see that an increase in the transmissibility of the pathogen can decrease the number of surviving trees. As β decreases from 0.5 to 0.1, we notice some surviving clusters. However, for the lowest value of β in this example (i.e. $\beta \leq 0.05$), trees at the edge of the clusters are more likely to survive, which means lower mortality rate for the disease. The most important thing here is that as β decreases, more trees are going to survive. Probably, this will cause the universal scaling to fail.

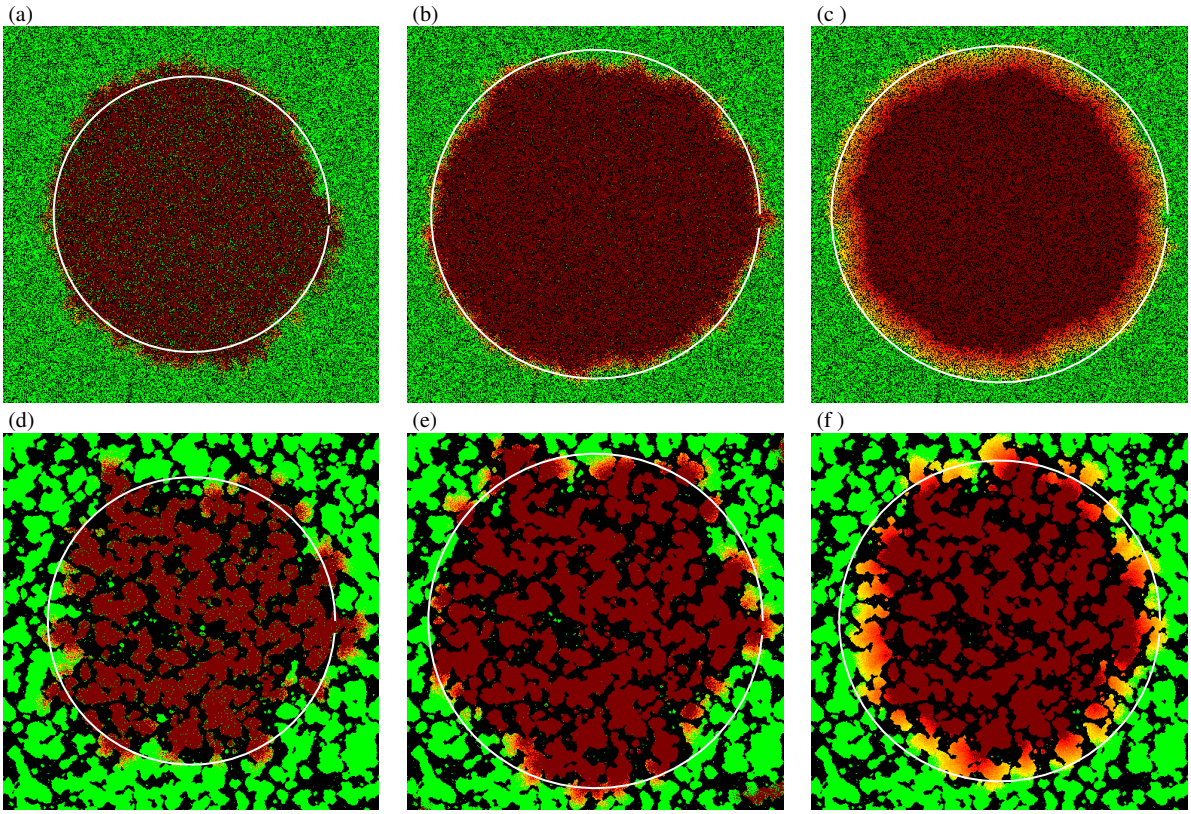

Figure 6.28: Late-stage infection configurations. We fix $\rho = 0.6$ and $\sigma = 3$. The top row has $\gamma = 0$ (i.e. homogeneous forestscape), the bottom row has $\gamma = 5$ (i.e a heterogeneous forestscape). We vary the trnasmssiblity β to be 0.05, 0.1 and 0.5 for the first, second and third columns respectively. The white circles have the same area as the infected/removed cells (factoring for density), to indicate how far from circular the infection configurations are.

In Fig. 6.29, we plot the number of infected trees as a function of time for the dynamics in Fig. 6.28. We note that the universal scaling works for the largest value of β for both the homogeneous and the heterogeneous forestscape. As β decreases, the scaling starts to fail, however, the curves for the homogeneous and the heterogeneous are roughly comparable.

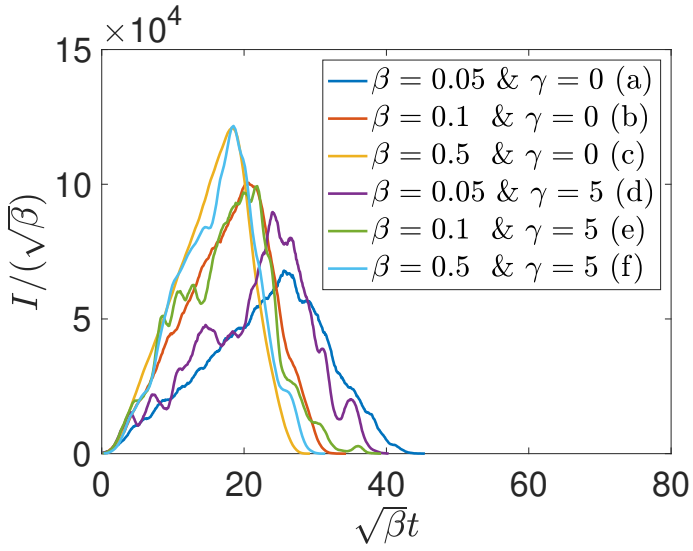

Figure 6.29: The number of infected as a function of time for the simulations shown in Fig. 6.28, normalised for transmissibility rate.

6.7.6 Transition To Epidemic

For the local infection when we calculate the spreading velocity of the disease, our results illustrate the existence of a critical transition separating two stable phases: local confinement and disease outbreak, that depends on the density of the susceptible individuals. Considering the spreading velocity for the non-local infection in Fig. 6.23, this transition is not clear anymore. The velocity increases steadily as the density increases, but there is not a clear point at which the behaviour of the velocity rapidly changes with the density. We will now examine how the mortality rate changes with the various parameters, and whether a clear transition to epidemic is discernible.

In Fig. 6.30, we calculate the mortality as a function of tree density, for a range of values of the infection width σ, for both homogeneous and heterogeneous forestscapes. For $\sigma = 1$, we see a clear transition to the epidemic as the density increases, with the density at which this transition occurs increasing as γ increases. This is consistent with the infection being suppressed by gaps in the forestscape, which have less effect as the density increases and the gaps shrink. For higher values of σ, we see that the density at which the transition occurs is shifted to lower values; this is again consistent with the gaps becoming less significant as the characteristic scale of the non-local kernel becomes comparable to the characteristic scale of the gaps. Note that the transition to epidemic is not clear from the velocity alone, unlike the local infection, where the transition was clear from both the mortality and the velocity separately. This highlights the importance of considering both of these measures

when modelling spatial infections.

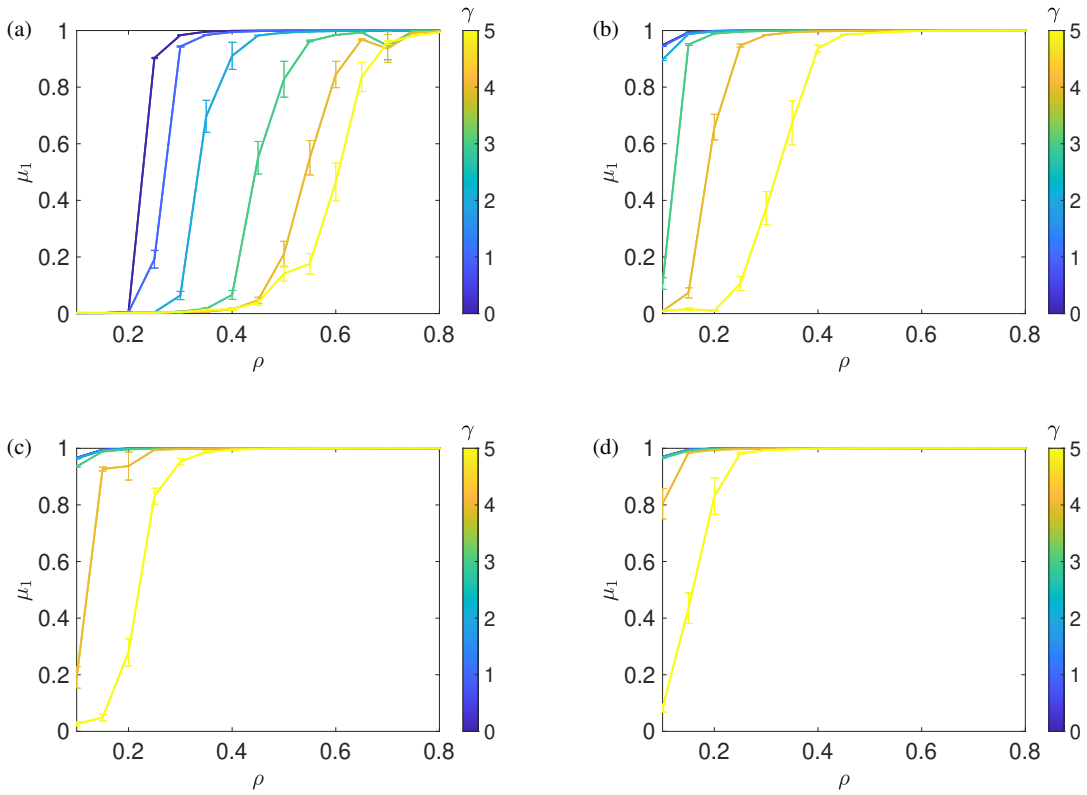

Figure 6.30: Effect of the heterogeneity level on the mortality for $\beta = 0.5$, a range of values of γ, and a $\sigma = 1, 3, 4, 5$ (panels (a-d) respectively).

6.7.7 Further Investigation

So far, we have performed our analysis across different parameters. Some interesting behaviour is noticed around the known threshold for connectedness with the Moore neighbourhood, which is 0.407 [80]. The non-local effect is evident around this density. At the start of this section, we consider densities between 0.38 and 0.44. We are interested in investigating what influences the shape of the averaged velocity. Then we will apply the analysis for a wide range of densities and include analysis of mortality.

In Fig. 6.31 (a), for $\rho = 0.44$ and $\beta = 0.5$ the time- and ensemble-averaged velocity is calculated at different values of σ for various levels of heterogeneity. For lower σ, the averaged velocity of the infection decreases as γ increases. However, this decrease becomes less strong as σ increases. To highlight this, in Fig. 6.31 (b-d), we fix the width σ to

be 2, 3 and 5 respectively, and the averaged velocity is calculated at different values of γ and ρ. We see that the averaged velocity is lower as γ increases, but for the largest values of σ, this effect is less strong. That is, the drop off in the spreading velocity is only noticeable when σ is small and goes away as σ increases. Simply, when σ becomes larger than the characteristic scale of the clustering, the effect of the clustering on the characteristic velocity of the infections becomes negligible. For smaller γ and high σ, you can see the effects of the clustering disappears because the transmission can now jump between clusters. Note that these densities are around the so-called Moore threshold ($\rho = 0.407$), but if we increase the tree density, this drop off will gradually become less apparent. This is reasonable because the separation between clusters become smaller.

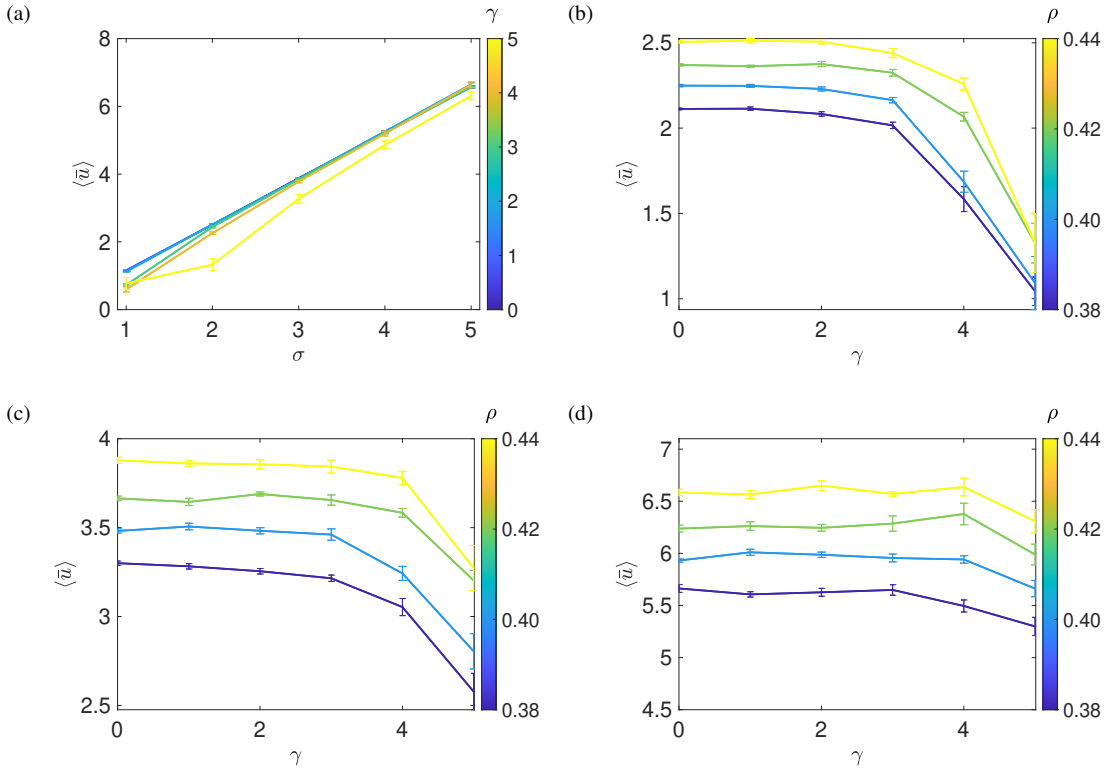

Figure 6.31: Panel (a), for $\rho = 0.44$ and $\beta = 0.5$ the averaged velocity is calculated at different values of σ for various levels of heterogeneity. Panel (b - d), we fix the width σ to be 2, 3 and 5 respectively, and calculate the averaged velocity at different values of γ and ρ.

In Fig. 6.32 (a), for a homogenous forestscape with $\rho = 0.44$, we empirically find a transmission probability β such that a pathogen with $\sigma = 5$ has the same velocity as a pathogen with $\sigma = 2$ and $\beta = 0.5$. That is, we want to reduce β for $\gamma = 0$ in the yellow curve in Fig. 6.31 (d) so it has the same velocity as $\gamma = 0$ in the yellow curve in Fig. 6.31 (b). The desired β was found to be 0.12. This was not straightforward and required several

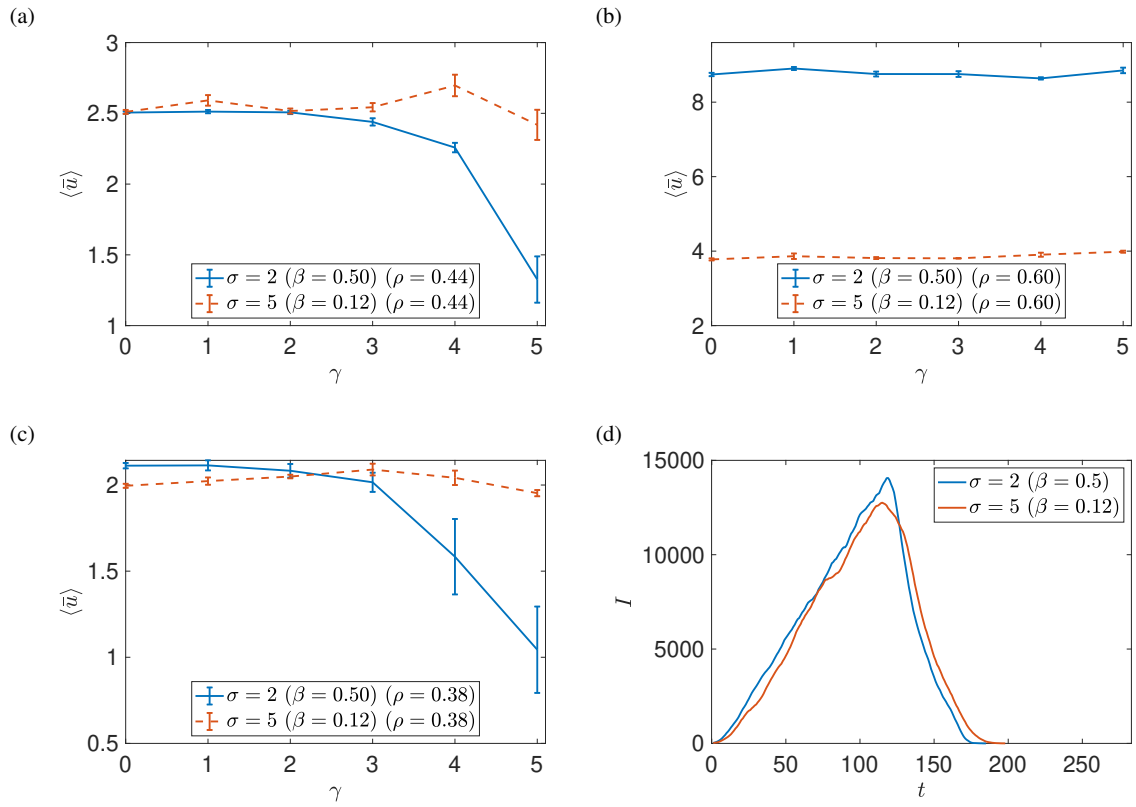

Figure 6.32: Panel(a), we used the top curve from Fig. 6.31 (b) (blue curve) and reduced the transmission probability β of Fig. 6.31 (d)(dashed red curve). We empirically found the best β to make the two homogeneous fields ($\gamma = 0$) have the same velocity, which happens to be 0.12, then we apply this to the different heterogeneity levels. Panel(b), we apply the same β that we found empirically to a higher tree density $\rho = 0.6$. Panel(c), we do the same as panel(b), but for a smaller tree density $\rho = 0.38$. Panel (d), we plot the number of infected trees for $\gamma = 0$ and $\rho = 0.44$.

attempts and human feedback to be successfully achieved. Then we are interested to see what happens as the heterogeneity level increases. We wanted to see if the red curve in Fig. 6.32 collapses into the blue curve as β decreases, or retains its shape; we find that it retains its shape as set by σ (i.e. you cannot map from one σ to another by only changing β). Reducing the transmission probability seems to have a single effect, which is to reduce the averaged velocity, but it does not affect the behaviour of the averaged velocity as γ changes. For Fig. 6.32 (b) & (c), we wanted to check what will happen if we change ρ. Can we use the same β which we find empirically for the density in panel (a)? The answer is, unfortunately, no. So, when we change density, we need to find a β empirically, so the two different σ have the same velocity (requiring further human feedback). In panel (d) we

compare the number of infected trees at each timestep for the homogeneous forestscapes in panel (a).

In Fig. 6.33 (a) and (b), we plot the velocity as a function of tree density for $\beta = 0.5$, $\sigma = 2$ and different heterogeneity levels. In Fig.6.33 (c) and (d), we reduced β for calculating the time and ensembles averaged velocity which is presented in Fig. 6.23 (e) and (f), we see that reducing β reduces the averaged velocity, but it does not affect the shape of the curves.

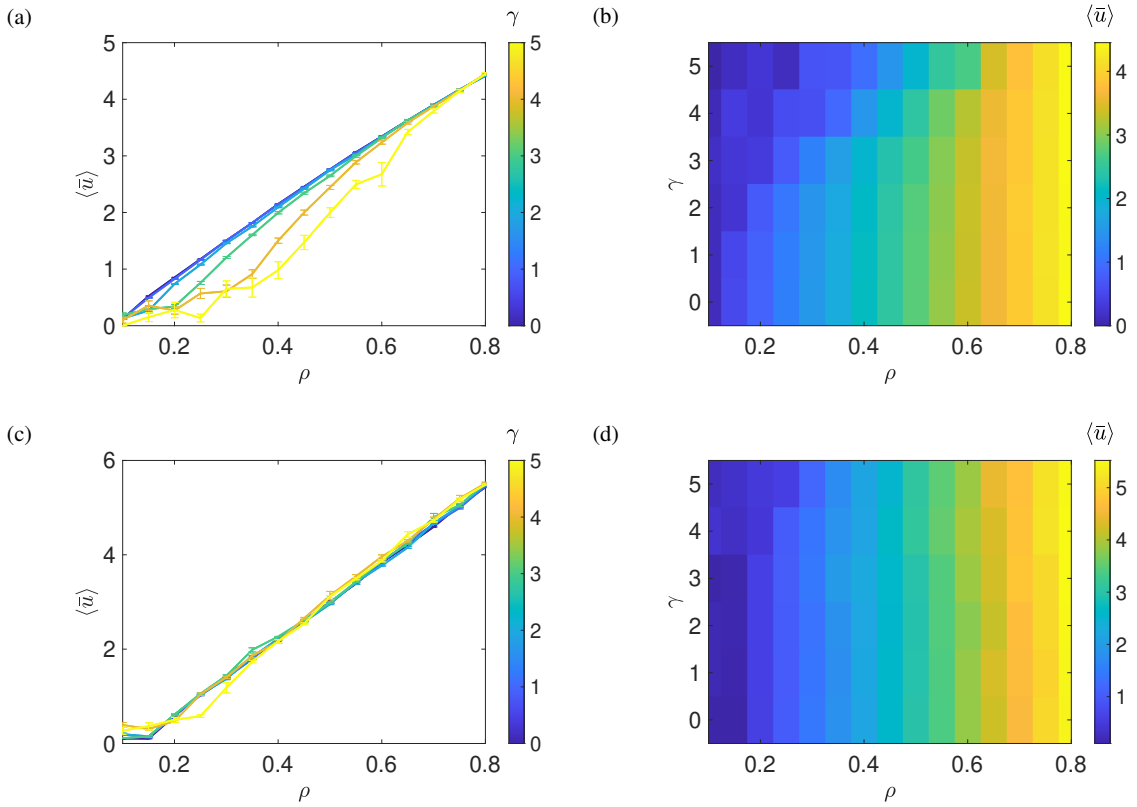

Figure 6.33: Panels (a) and (c): the averaged spreading velocity of the infection $\langle\bar{u}\rangle$ for different heterogeneity levels of forestscapes is plotted against tree density ρ ,while (b) and (d) are the associated phase diagrams. In the first row, we used $\beta = 0.5$ and $\sigma = 2$ while in the second row we used $\beta = 0.12$ and $\sigma = 5$ (i.e a reduced β for a high σ).

In Fig.6.34(b), we reduced β for mortality which is presented in panel (a), we see that reducing β significantly affects lower densities. The transition to epidemic now occurs at larger densities.

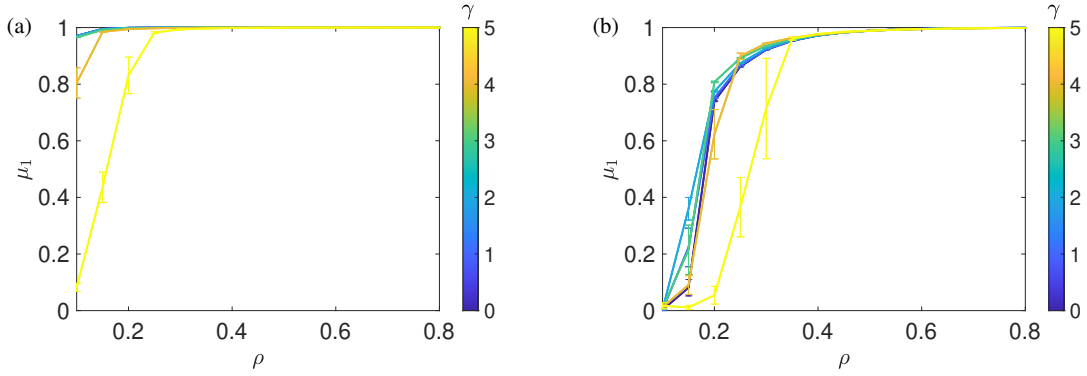

Figure 6.34: Panel (a) for $\beta = 0.5$ and $\sigma = 5$, the mortality is calculate at different densities for a range of heterogenity levels γ while panel (b) is the same as panel (a), but for $\beta = 0.12$.

In Fig. 6.35, for $\rho = 0.44$ and $\beta = 0.5$, the mortality μ_1 is calculated at a range of heterogeneity levels for various σ. We see that as γ increases, the mortality decreases, but for a relatively high value of σ, these effects become less strong, and the disease kills all the trees in its way.

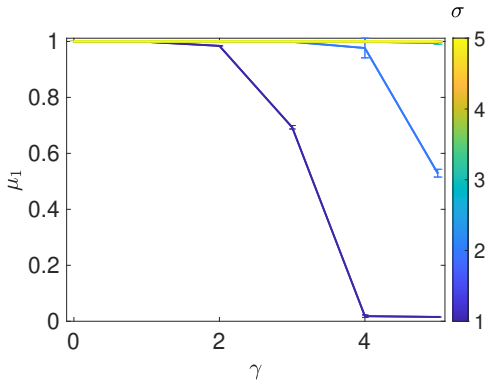

Figure 6.35: For $\rho = 0.44$ and $\beta = 0.5$, the mortality μ_1 is calculated for a range of heterogeinity levels.

In Fig. 6.36, the ensemble-averaged mortality $\langle \mu_1 \rangle$ is measured as a function of transmission probabilities β for different values of the characteristic width of the non-local kernel σ . From these results, we observe the occurrence of a critical transition that separates two stable phases: local confinement and disease outbreak. In panel (a), for the lowest value of σ and homogeneous forestscapes, the transition to epidemic is observed around $\beta = 0.05$.

As σ increases, the transition is observed around $\beta = 0.02$. Moving to heterogeneous forestscapes with $\gamma = 5$ in panel (b), for higher values of σ, we observe the same critical transition which occurs around the same value as the homogeneous forestscapes. This is expected. For these high values of σ and tree density $\rho = 0.60$, the effect of the clustering is diminished here because the disease can jump over gaps between the clusters. However, for heterogeneous forestscapes with the lowest value of σ, the infection is suppressed by the gaps between the clusters of trees which the infection could not jump over. Consequently, the lowest value of σ for the heterogeneous forestscape has lower mortality than higher σ.

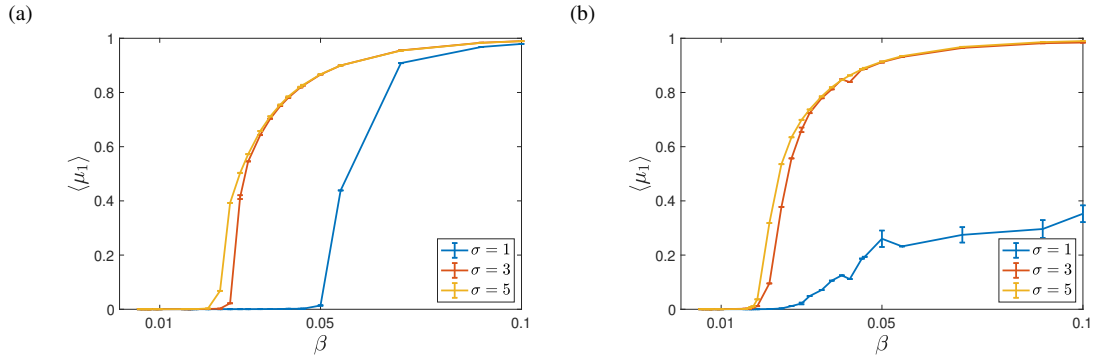

Figure 6.36: Panel (a), for homogeneous forestscapes, $\mathcal{L} = 500$, tree density $\rho = 0.60$ and 10 ensembles, the averaged mortality μ_1 is plotted as a function of transmissiblity β and for a range of the characteristic width of the non-local kernel σ. Panel (b), we did the same thing as panel (a), but for heterogeneous forestscapes with $\gamma = 5$.

In Fig. 6.37, for a fixed homogeneous forestscape with density $\rho = 0.60$, we display snapshots for different combination of σ and β which have the same mortality ($\mu_1 = 0.57$). In panels (a-c), we the display snapshots at time $t = 300$ while in panels (d-f) the snapshots are taken at the end of the simulations. Picking combinations of parameters that produces the same mortality allows us to see the differences in the way the disease propagates.

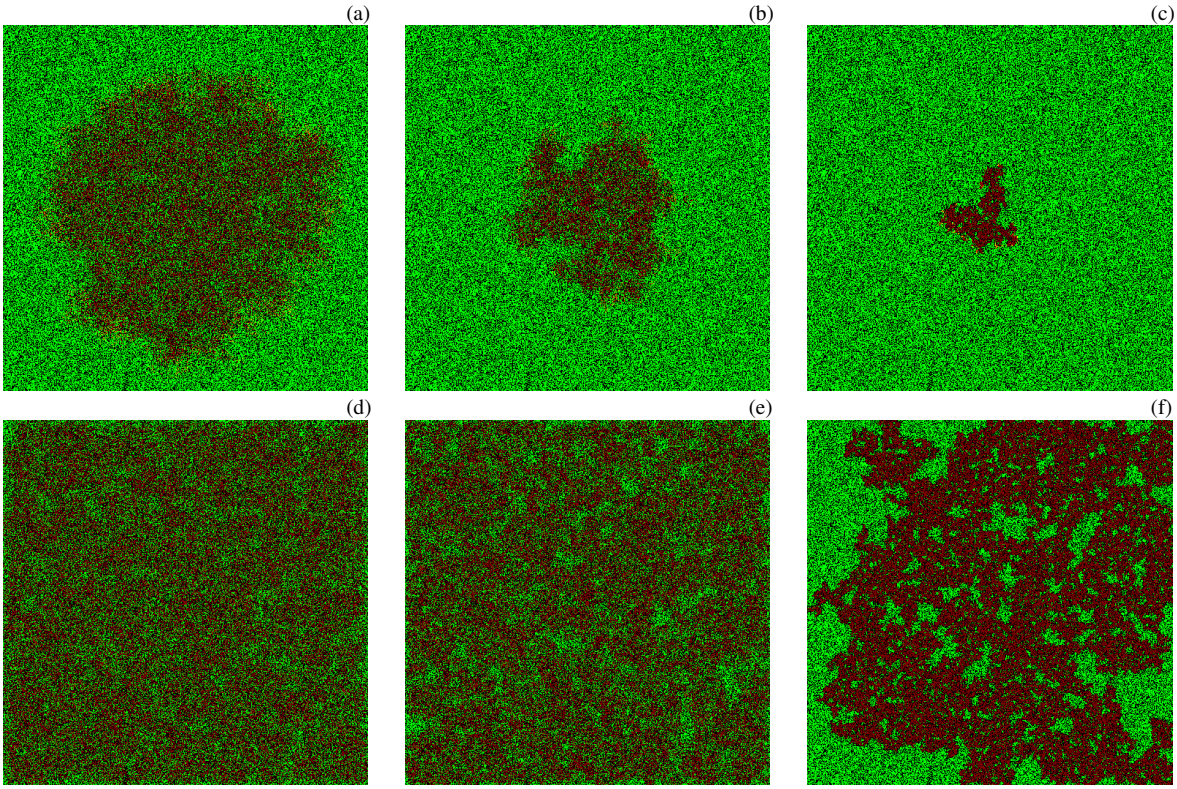

Figure 6.37: For homogeneous forestscapes, $\mathcal{L} = 500$ and tree density $\rho = 0.60$, the first column ($\sigma = 5$ & $\beta = 0.0320$), the second column ($\sigma = 3$ & $\beta = 0.0328$), and the third column ($\sigma = 1$ & $\beta = 0.0550$). In panels (a-c), we the display snapshots at time $t = 300$ while in panels (d-f) the snapshots are taken at the end of the simualtions. Note that at the end of the simualtions all these have the same mortality $\mu_1 = 0.57$.

6.8 Distance to Nearest Infected

With the introduction of non-local interactions, it is possible for infections to spread further than to their nearest neighbours. In this section, we are interested in how far newly infected trees are from the previously infected trees.

We calculate the probability density function (PDF) of the minimum distance δ of a newly infected cell from any currently infected cell over the entire infection. In Fig. 6.38 we use $\rho = 0.6$ and vary σ and γ. We observe that changing γ has little effect on the shape of the PDFs. In all cases, the PDF has a strong peak at $\delta = 1$, indicating that most new infected cells have a direct neighbour that is infected. Unsurprisingly, as we increase σ,

the tails of the distribution become more pronounced. This is most dramatic as we change from $\sigma = 1$ to $\sigma = 3$: note that no values of δ beyond around four are seen for $\sigma = 1$ - though of course, such values are possible but likely too rare for us to observe without a larger number of realisations.

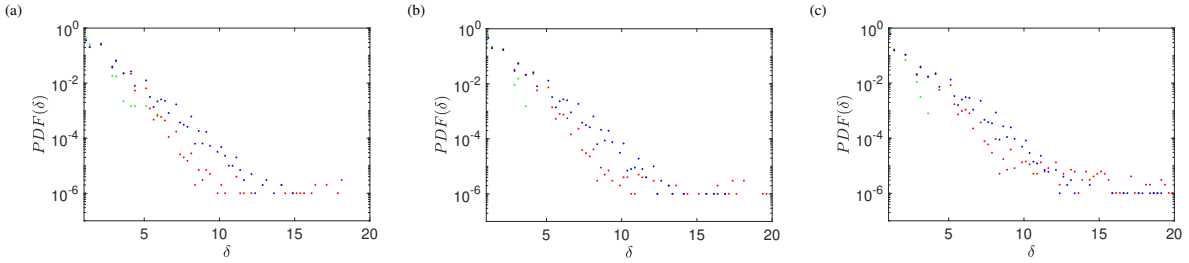

Figure 6.38: For $\rho = 0.60$, the probability density function of the minimum distance of a new infected tree from any infected tree PDF(δ), calculated for $\sigma = 1, 3, 5$ (green, red and blue respectively). Panel (a) is for a homogeneous forestscape, while panels (b) and (c) are for $\gamma = 3$ and $\gamma = 5$, i.e. heterogeneous forestscapes

In Fig. 6.39 we use $\gamma = 0$ and vary σ and ρ. We observe that varying the density has little effect on the overall shape of the PDFs while increasing σ both broadens the distribution and changes the shape of it, with a more pronounced kink in the shape of the distribution becoming evident with increasing σ. For the lower values of σ long tails become evident for the highest values of ρ. On the other hand, for the higher values of σ, the tails are longer for the lowest densities.

In Fig. 6.40 we use $\rho = 0.45$ and vary σ and γ. For the lowest value of σ, we again see a sudden cut-off in the distribution beyond around $\delta = 4$. Varying γ appears to have little effect on the shape of the distribution.

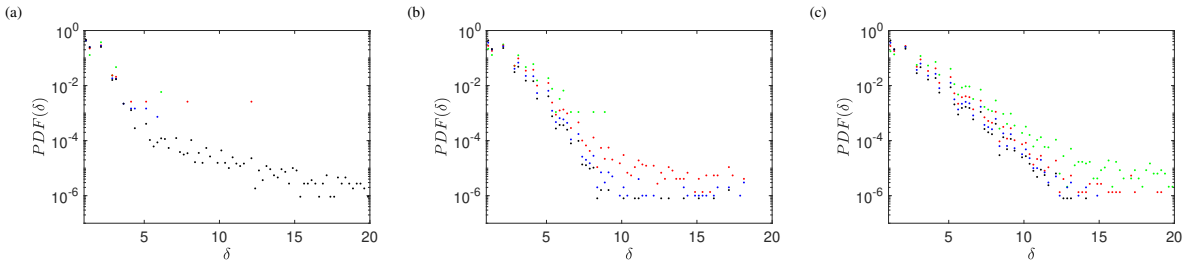

Figure 6.39: For $\gamma = 0$, the probability density function of the minimum distance of a newly infected tree from any infected tree PDF(δ) calculated for $\rho = 0.3, 0.45, 0.6, 0.8$ (green, red, blue and black respectively). We do this for $\sigma = 1, 3, 5$ which are represented in panels (a-c) respectively.

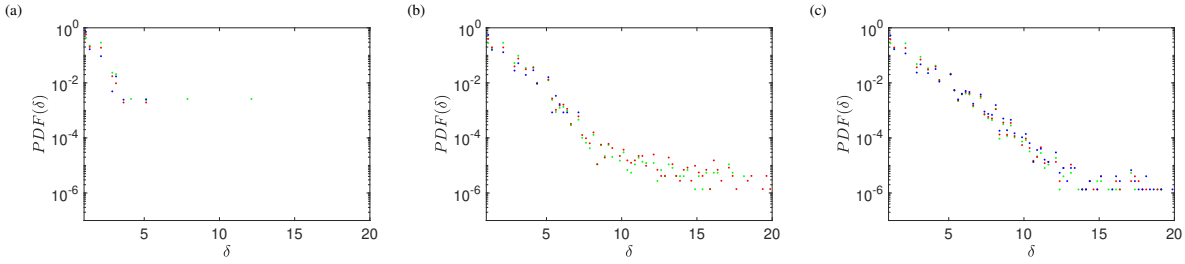

Figure 6.40: For $\rho = 0.45$, the probability density function of the minimum distance of a newly infected tree from any infected tree PDF(δ) calculated for $\gamma = 0, 3, 5$ (green, red and blue respectively). We do this for $\sigma = 1, 3, 5$ which are represented in panels (a-c) respectively.

6.9 Conclusions

For the non-local infection, we model the transmission probability as a function of the distance between the infected and susceptible cells, rather than taking some constant value for direct neighbours and zero elsewhere. We develop a generic model for spatial dependence of the transmission probability, using a Gaussian kernel, which has two parameters: the transmissibility β and a characteristic distance σ. In order to make direct comparisons to the local infection simulations, we make a correction to β such that the initial infection rate ΔI is kept the same as σ is changed.

Using the same tools as for the local infections, we have probed the effect of the characteristic distance of infection transmission, forestscape density, the transmissibility of the

pathogen, and the degree of heterogeneity of the forestscape, as well as the combined effects of these parameters.

Firstly, we note that the behaviour of the infections is strongly dependent on the value of σ, even after we have normalised the initial infection rate with respect to it. The spreading velocity is still directly proportional to σ, and having made this normalisation we conclude that this is an effect purely due to the non-local nature. The value of σ is also important in the suppression of the infection velocity by heterogeneity, and the critical value at which transition to epidemic occurs.

We find that the heterogeneity parameter can suppress the infection velocity, which we interpret as the effect of larger gaps emerging in the forestscape, which the infection is unable to transmit over. This is consistent with the two observed parameter regions in which this effect is diminished. Firstly, for the largest values of σ, the suppression is negligible: with larger values of σ the infection is able to transmit over these gaps unhindered. Then, we note that the suppression is also diminished at higher densities: if the forestscape is more dense, clearly the number or size (or some combination of these) of gaps must be lower, so there are fewer obstacles to the infection.

The local infection has two stable phases: local confinement and disease outbreak (epidemic), with a clear critical transition between them which is dependent on the forestscape density. Underlying this is the connectedness of the forestscape - the critical transition occurs at the density at which percolation theory tells us that the regions are connected, for both the von Neumann and Moore neighbourhood models. However, for non-local infections, it is less clear whether a forestscape can be considered to be connected, as the neighbourhood determining the connectedness is no longer clearly defined. When we examine the velocity as a function of density, there is no clear transition.

We instead examine the mortality rate, finding that, as for the local infection, a critical transition occurs above a certain density. However, for almost all combinations of parameters, the critical density is considerably lower than for the local infections. We observe two trends: higher values of heterogeneity parameter lead to higher critical densities, and higher values of σ lead to lower critical densities. We note that for some cases the mortality is already almost 1 at the lowest value of density which we use, $\rho = 0.1$.

For the transmissibility of the pathogen, for all values of σ, as β increases the velocity of the infection increases, most rapidly for the smallest values of β.

We empirically identify how the number of infected trees vs time scales with the various parameters. Although the scalings apply to a broad subset of the simulations, they are not universal, typically failing for combinations of parameters which are close to the boundary between contained and epidemic infections; the critical region. Examining cases where the scaling fails, we identify two behaviours. In the first, the infection spreads in an anisotropic manner, while tending to spread isotropically when the scaling does hold. In

the other, we observe a proportion of the cells in the affected regions surviving - we also identify an interesting phenomena, where cells that are close to the edge of clusters have a greater probability of surviving. We also attempted to combine the individual scaling curves found into a single master curve, finding: $I \to I/(\rho^{3/2}\sigma\sqrt{\beta})$ and $t \to \sigma\sqrt{\rho\beta}$. This seems to describe a good amount of the behaviour observed; however there are regions where the individual scaling laws, here combined to produce this master curve, fail and so the master curve naturally also fails in these regions.

Clearly, there are many rich aspects of the infection dynamics which stem from the inclusion of a non-local kernel, and the combined effects of this with other parameters. Even with the smallest values of σ, and with the infection probability normalised to match the local infection, the behaviour of the non-local infection can be drastically different from the behaviour of the local infection with all other parameters the same. Many of the observed effects may be extremely relevant to understanding and managing the propagation of tree diseases, and could not emerge from a model containing only local infection spreading. For future studies, it would be interested to quantify and analyse the "front width" of the infection dynamics. We expect that it is proportional to σ. But, it is potentially something interested for the following studies.

Chapter 7

Comparing Synthetic and Real Forestscapes

7.1 Introduction

In the previous chapter, we simulated the infection spreading on synthetic forestscapes, which we have generated as either homogeneous fields, or heterogeneous fields in which the distribution is clustered. It is not feasible to conduct experiments in real forests to study the spread of tree disease: the process may take a very long time (probably decades), requires an extensive area of woods, and most importantly, it is not ethical to initialise tree diseases in real forests. Hence, being able to model disease outbreak in a realistic setting is crucial, with one aspect of this being the synthesis of realistic forestscapes.

In this chapter, we address the question of how well suited our synthetically generated forestscapes are for modelling real forests. In order to do this, we will extract publicly accessible real forest images provided by Google maps. We will gather images that are approximately 10km×10km from natural forests from two different countries: Finland and Canada, chosen because they belong to the same Köppen climate group. Also, they were chosen because of the availability of data. From Finland and Canada, we can get good quality data that we can analyze. Then, using the Mahalanobis distance based on colour, we will classify each pixel as being either wooded or not.

We will then compare our synthetically generated forestscapes to the real data in terms of two parameters: the tree density, and the heterogeneity level. Comparing the tree density is straightforward. For the heterogeneity level, we will again use Besag's function to characterise the spatial clustering and segregation in the real forestscapes. This will in turn be used to estimate the value of γ which produces synthetic forestscapes most similar

to the particular real forestscape.

7.2 Image Acquisition

The forest trees images are extracted from Google maps. Then within MATLAB, we cut out regions that are 10km×10km, based on the scale provided. The resolution depends on the source, and we do not have precise control over the zoom level, so these images range from $650 \times 650\text{px}^2$ to $1040 \times 1040\text{px}^2$, with a single pixel then covering a region from approximately 10^2m^2 to 15^2m^2, which is on the order of magnitude of the area occupied by a mature tree. We acquired fifteen different images from Finland and another fifteen from Canada. We have tried to select completely natural forests with as few roads as possible. We use the Portable Network Graphic (PNG) format to save these images because it is a lossless format [92]. A typical image is shown in Fig.7.1(a).

Google maps gather their images from various sources such as aerial photography and multiple satellites. The source of my images is Imagery@ Landsat / Copernicus, Maxar Technologies, Imagery@2019 TerraMetrics, Map data @2019. The highest resolution for Landsat is 15m/pixel, and Copernicus is 10m/pixel.

7.3 The Mahalanobis Distance

To distinguish between what is a tree and what is not a tree, we have used the Mahalanobis distance in colour space. Mahalanobis distance is a useful measure of the distance between a point and a distribution, standardised by the variation in the distribution. This method was introduced by an Indian-Bengali statistician called Prasanta Chandra Mahalanobis in a context of a racial likeness study. It has since been widely used in the field of cluster analysis and classification [93, 94, 95]. Since our data are satellite images and aerial photographs for forest trees, and we are only interested in classifying which regions are trees, then the Mahalanobis distance is a suitable method.

We first select a region which can be easily visually identified as totally populated with trees. The Mahalanobis distance $\mathcal{M}$ from this subset, for a particular pixel, can be calculated using the following formula:

$$\mathcal{M} = \sqrt{(\Psi - \mu)^T S^{-1} (\Psi - \mu)}$$

where Ψ is a 3-vector of the red, green and blue (RGB) colour channels for the pixel under consideration in the original image, and S and μ are the covariance matrix and mean of

the RGB colour channels for the selected populated region. It is desirable to select as large a region as possible to initially identify as trees, as this will give better statistics for this class, which makes the Mahalanobis distance more reliable. Calculating the covariance for a box of trees allows us to see how the three colour channels influence each other. The Mahalanobis distance then allows us to see how far a value is from the mean of a distribution, in a way that standardises the variance and covariance.

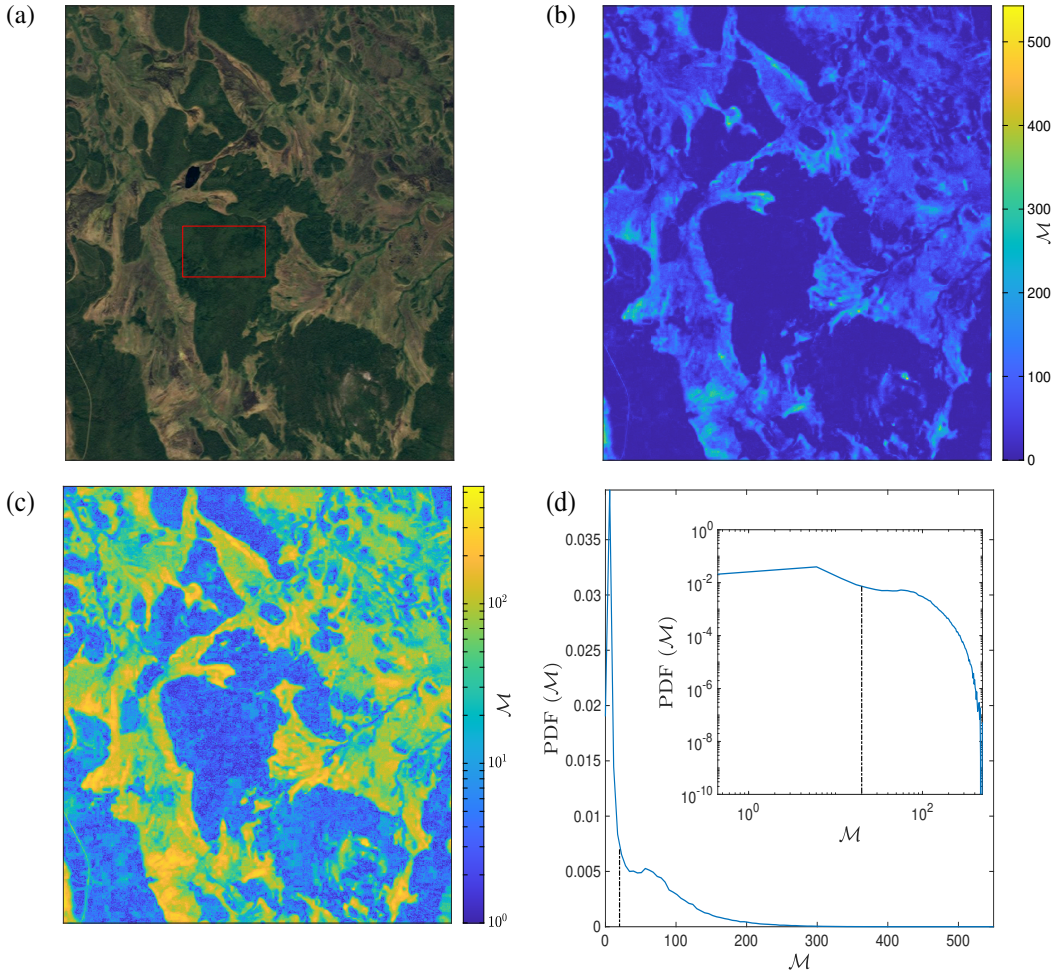

Figure 7.1: Illustration of the process used to transform our aerial photograph to a binary map of the forest area. Panel (a) shows the original RGB image from Finland, with the sampling box highlighted as a red box. Panel (b) shows the calculated Mahalanobis distance $\mathcal{M}$ for each pixel from the sampling box, shown again in panel (c) on a logarithmic scale. Panel (d) shows the PDF of $\mathcal{M}$, with our chosen T_c shown as a black dashed line, with this shown on a logarithmic scale in the inset.

In Fig. 7.1 we illustrate the process of classification. Panel (a) shows the $10 \times 10\text{km}^2$ image extracted from Google maps, with the region selected as our subset entirely populated by trees highlighted with a red box. Panel (b) shows the Mahalanobis distance of each pixel from the selected region, with panel (c) showing the same information on a logarithmic scale to highlight the distinction between trees and not-trees.

In Fig. 7.2, we show the distribution of colours in each of the colour channels, with the subset selected to calculate the Mahalanobis distance indicated. We note that the values are fairly tightly clustered, which is a good sign and shows that we have picked a reasonable characterisation box for our distribution of trees.

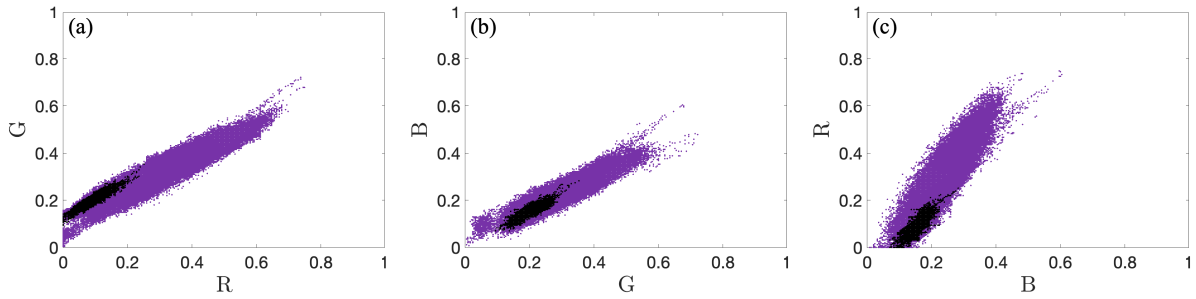

Figure 7.2: For the example image shown in Fig. 7.1 (a) we plot the three different colour channels, RGB, against each other in purple. Each combination of two colour channels from the original image and from the characteristic box of trees are plotted against each other. The black dots show the selected characteristic tree box.

7.4 Thresholding

We now seek to convert the images to a binary map of trees and not-trees using the Mahalanobis distance. We empirically threshold these distances until we get a good binary representation of the forest. That is, we are looking for a threshold of Mahalanobis distances, T_c, such that every value less than this value is considered to be a tree. We start from zero and then increase this value at a constant rate until we get the desired forest representation.

The main drawback of this process is how much time and supervision is needed. However, if we examine the probability density function (PDF) of the Mahalanobis distance, we note that we have a large peak close to the origin, which contains our trees, and a broader distribution of points which are not trees further from the origin. By looking at where this peak approximately ends, we can get a sensible initial value of T_c to test, which speeds the process up. In Fig. 7.1 (d), we show where our largest peak approximately ends

(i.e. roughly around 20). For this particular case, our desired threshold is $T_c = 20$.

To illustrate the importance of picking an appropriate value for T_c, we now introduce a new quantity, ϵ, which denotes the difference between the threshold value used, and the actual value of T_c identified. In Fig. 7.3, we illustrate how the classification can change for positive and negative values of ϵ. In panel (a) we show the original image. In panel (b),

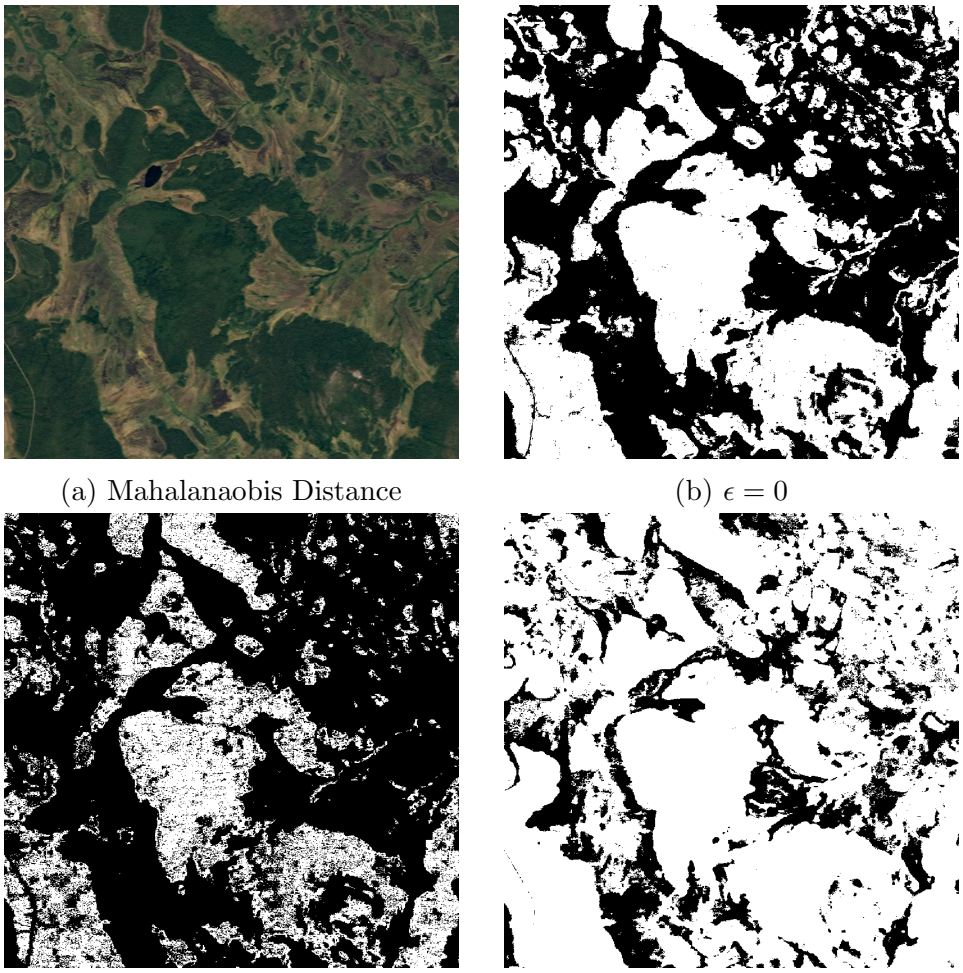

(a) Mahalanaobis Distance (b) $\epsilon = 0$

(c) $\epsilon = -15$ (d) $\epsilon = 50$

Figure 7.3: Illustration of the importance of picking the right T_c. Panel (a) shows the original image of the forestscape. Panel (b) shows the classification, with trees in white and not-trees in black, when we use $T_c = 20$, which we use as our baseline ($\epsilon = 0$). In panel (c) we use a threshold value which is too small, while in (d) we use a threshold value which is too large.

we show the binary map of trees produced using $T_c = 20$ as recommended by the edge of the large peak in the PDF. Comparing this to the original image by eye, this seems to

be an accurate classification. There are no obvious regions which have been incorrectly classified, the boundaries between classifications are generally smooth, and features, such as the road in the lower left, have been captured accurately. Then, in panel (c), we show how using a negative ϵ, i.e. a too small threshold value, affects the classification. There are now many regions which are clearly populated with trees in the original image which have not been classified as trees. The boundaries between regions are no longer as smooth and clear-cut, and regions which are mostly identified as populated here are very noisy. The road feature is also over-represented here. In contrast, in panel (d) we show how a positive ϵ, i.e. a too large threshold, can affect the classification. We now see many regions which are clearly not populated with trees have been classified as trees. Again, the boundaries between regions are not as smooth and clear-cut, and there is a lot of noise in regions which were previously not identified as trees. In this case, the road feature has mostly been lost.

7.5 Computing Besag's Function on the Real Forestscapes

We have now reduced our original images to binary matrices of regions populated by trees, which appear by eye to be good representations of the actual forestscapes. We will now move on to characterising the spatial distribution of trees, using Besag's function. However, it is computationally expensive to use all the points from the data set as test points for Besag's function. For higher densities, the calculation time can take up to 5 days on a desktop computer for a single run for a 400×400 image, so calculating this for all 30 images at their full resolution is not feasible for us.

To make progress, we reduce the computation time in this section through downsampling the image, then using all of the points in the downsampled image to calculate Besag's function. This changes the resolution of the images, so the smallest scale which we have precise information about will increase. This will affect the result of Besag's function for the smallest features, but should still be consistent for the larger features.

Rather than resizing our images after the classification stage, we instead resize the matrix of Mahalanobis distances. The aim of doing this is to interpolate the Mahalanobis values at the downsampled points, which can then be used to classify, rather than downsampling after the classification step, when we only have binary values to try and interpolate. Figure 7.4 shows the binary image for the original image without any downsampling, and the binary matrices produced from the resized Mahalanobis distance matrices. By eye, the characteristics of the image do not appear to be changed until we resize by a factor of 10, although some of the smallest features are naturally lost before this point.

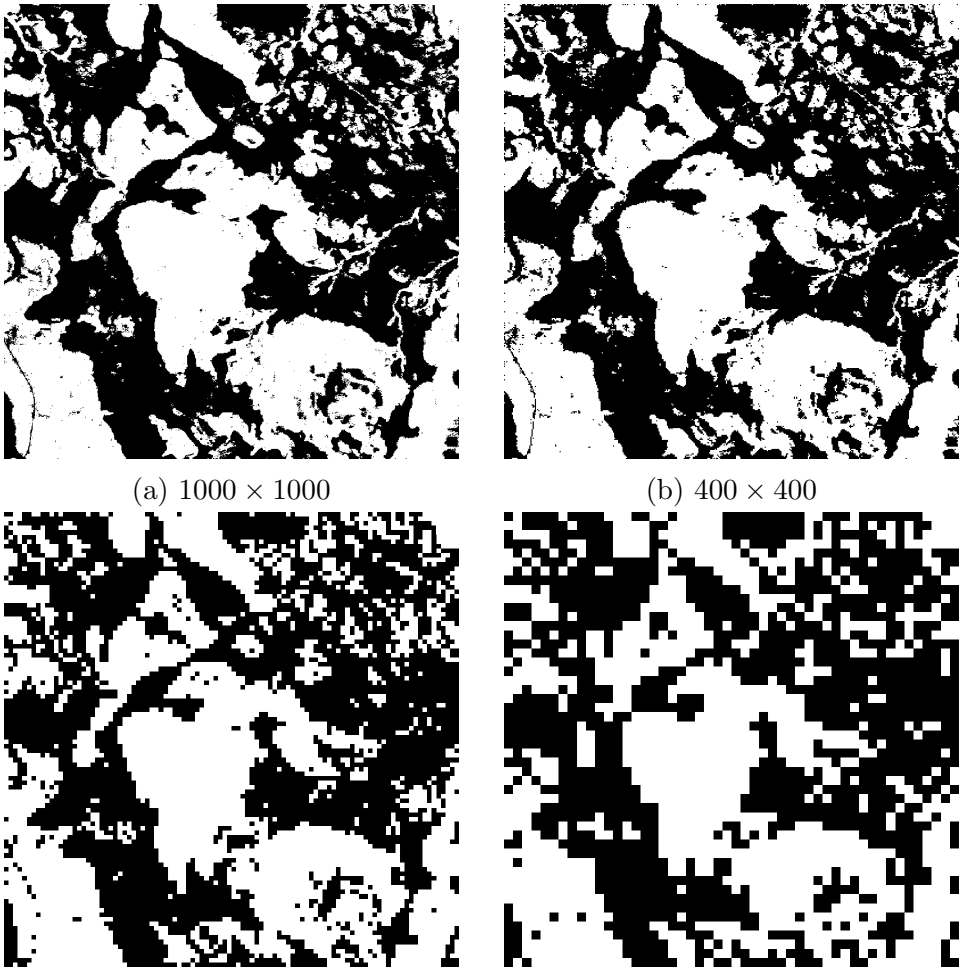

(a) 1000×1000 (b) 400×400

(c) 100x100 (d) 50x50

Figure 7.4: Panel (a) shows the original binary map corresponding to Fig. 7.1(a). Panels (b) - (d) show the downsampled maps as various resolutions.

We are now using real data with different size grids that we can not control in the same way as our model, so it is natural to introduce a normalisation to Besag's function. We normalise r as $(D/\mathcal{L})r$, $L(r)$ as $(D/\mathcal{L})L(r)$, where D is the length of the box under consideration, so for our real forests $D = 10$km. In Fig. 7.5, we plot Besag's function for a range of image sizes. We see that all the curves collapse onto each other, which implies changing the resolution of the image does not affect the curve of Besag's function, and confirms our normalisation. Note that periodic boundaries are used to be comparable to numerical approach, but this maybe an issue for real, less homogeneous data.

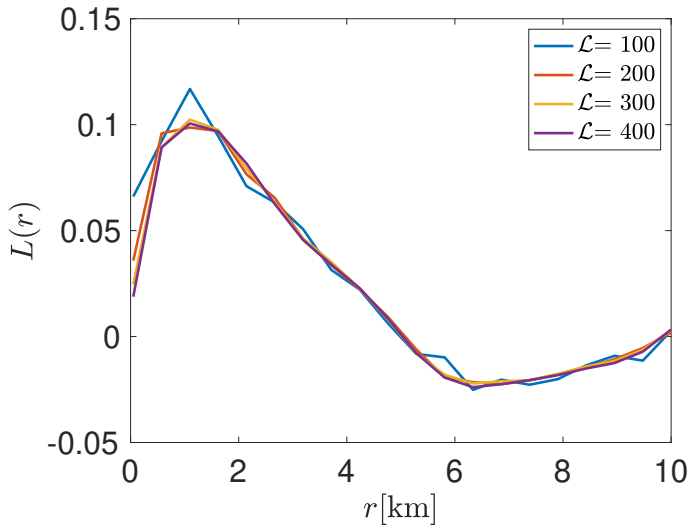

Figure 7.5: Besag's function at different scales measured in kilometres for the same image but downsampled. We see that all the curves collapse onto one another, which implies that changing the size of the image does not affect the overall Besag's values.

Note that if we decrease the resolution to be 100×100 or less, the $L(r)$ becomes noisier especially for small scales which is reasonable and expected, and we should avoid these lower resolutions. Moreover, at this resolution, we no longer capture the small-scale features such as roads. Therefore, we have decided to go with grid sizes that are 200×200 because they are both computationally affordable, and they do not lose too much information.

7.6 The Effect of Changing Mahalanobis Sampling Box

To transform our satellite images to binary matrices through Mahalanobis distances $\mathcal{M}$, we always pick a sampling box which is a subset from the original image that is fully populated with trees. When we pick this sampling box, we try as much as we can to have subsets that only contains trees.

It is important to check the effect of changing the sampling box. Therefore, in Fig. 7.6(a), we have picked two different sampling boxes to transform our satellite images to binary matrices through Mahalanobis distances $\mathcal{M}$. In the first case, highlighted in red, we have selected a large region which we are confident contains no, or very few, non-trees. In the

second case, highlighted in cyan, we have selected a relatively small region, which we are again confident contains no, or very few, non-trees.

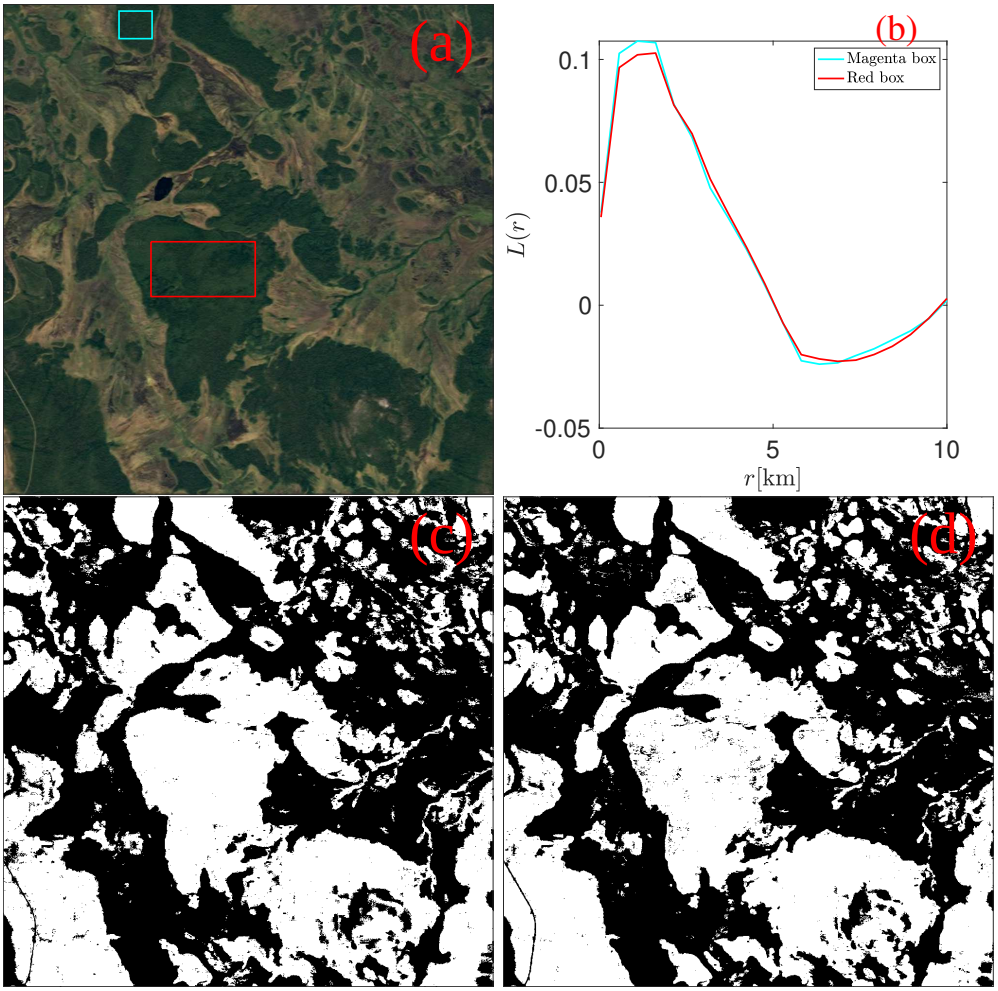

Figure 7.6: Panel (a), we have picked two different sampling boxes to transform our satellite image to a binary matrix through Mahalanobis distances $\mathcal{M}$. Panel (b), we calculated Besag's function for the produced binary matrices . The binary matrix produced from the red box represented in panel (c) while the binary matrix from the cyan box displayed in panel (d).

Then, in Fig. 7.6(b), we calculated Besag's function for the produced binary matrices. The binary matrix produced based on the red box has a $\rho = 0.4678$ and represented in in Fig. 7.6(c) while the binary matrix from the cyan box has $\rho = 0.4671$ and displayed in in Fig. 7.6(d). Therefore the relative difference is small (0.0007/0.4678=0.00150), which is

good. We can also consider what proportion of pixels are given the same classification by using each sampling box, and find that 94.6% are in agreement. This is expected because we are interested in the typical colour of forest trees. This typical colour can be produced by different tree box. Note that this characteristic box of trees needs to be as large as possible, and needs to avoid any empty spaces inside it. The second condition is much stronger than the first one. This is because picking a box with some empty spaces can result in a poor representation of the forest trees. However, picking a small box does not always lead to unusable results.

7.7 Characterising The Real Data

We can now calculate Besag's function with confidence that our downsampling does not affect our accuracy, and so we are in a position to start comparing the extracted data to our synthetic data. We will do this by finding the heterogeneity parameter, which produces the synthetic data with the most similar Besag's function to the real data.

7.7.1 Estimating Tree Density

Tree density can be obtained through a straight forward calculation from the binary map of the forestscapes. It is calculated using the following formula:

$$\rho = \frac{N_{\mathrm{T}}}{\mathcal{L}^2},$$

where N_{T} is the number of trees. As well as being an important parameter in its own right, we also need to know the tree density to be able to estimate the heterogeneity levels, as we have seen in Sec. 3.4 that the clustering depends on the density as well as the heterogeneity parameter γ.

7.7.2 Estimating Gamma Through Least Square Difference

Knowing the density for a forestscape, we can now estimate the heterogeneity parameter. The least-squares method allows us to estimate a parameter which best describes a set of data, by minimising the deviation between the data predicted by the parameter and the data observed. To estimate γ for the natural forestscapes, we wish to find the value of γ that minimises:

$$SS = \sum_i [\langle L_\gamma(r_i) \rangle - L'(r_i)]^2,$$

where $\langle L_\gamma(r_i) \rangle$ is the ensemble-averaged Besag's function for the synthetic heterogeneous data with clustering parameter γ, and $L'(r_i)$ is the Besag's function for the natural forestscape under consideration, summed over all radii at which we calculate Besag's function. We believe that the minimum SS value can tell us the best possible fit for the real data (i.e. the optimum γ for that best possible fit from the synthetic heterogeneous data).

Figure. 7.7 shows an example of how to estimate the optimum γ for each real forest image by using the least-squares method. We calculate SS over a range of values of γ, and then take the value at which the minimum value of SS occurs to be the optimum choice for γ, and so in the example shown we find $\gamma = 5$. Note that we repeat this process for all the real forest images.

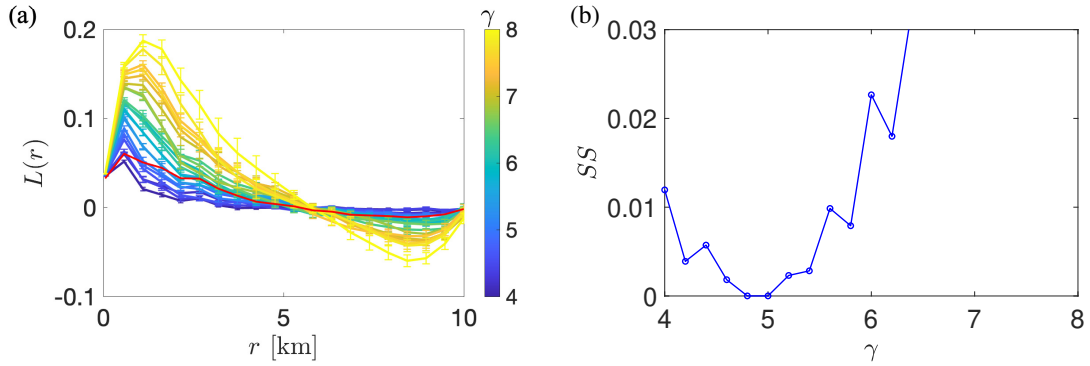

Figure 7.7: Panel (a), the red curve represents the calculated Besag's function for a real image from Finland which has a certain density, while the other curves represent the calculated Besag's function for the same density for different values of γ averaged over 30 ensembles. Panel (b), the least square difference SS for the real image from Finland.

7.7.3 Finland Data

In this section, we present the density and optimum γ estimated for images from Finland. The acquired images, the choice of sampling boxes, and the resulting binary representations are given in Appx. B. We summarise the densities and optimum γ for the forestscapes in table 7.1. The mean density is 0.5345 and the mean optimum γ is 5.680. Using the total sum of square errors over all forestscapes, the overall best-estimate is $\gamma = 6.0$.

Plot ref.	Density	Optimum γ
Fig. 7.1(a)	0.4678	6.6
Fig. B.1(a)	0.5165	6.6
Fig. B.1(c)	0.5562	5.8
Fig. B.1(e)	0.4893	5.0
Fig. B.1(g)	0.5012	5.0
Fig. B.1(i)	0.6178	5.4
Fig. B.1(k)	0.6937	4.8
Fig. B.1(m)	0.3470	7.4
Fig. B.1(o)	0.3897	6.4
Fig. B.2(a)	0.6048	5.2
Fig. B.2(c)	0.6329	5.6
Fig. B.2(e)	0.5649	5.4
Fig. B.2(g)	0.5783	4.6
Fig. B.2(i)	0.6107	5.6
Fig. B.2(k)	0.4471	6.0

Table 7.1: The densities and optimum heterogeneity parameter for the 15 different 10km $\times$ 10km satellite images for forests from Finland.

In Fig. 7.8, we plot the optimum γ against tree density. Generally, we observe that forestscapes with lowest tree density has high levels of heterogeneity, while forestscapes with the highest density have lower heterogeneity levels. We fit a line through our data, which shows a negative slope, so as tree density increases, the level of heterogeneity decreases.

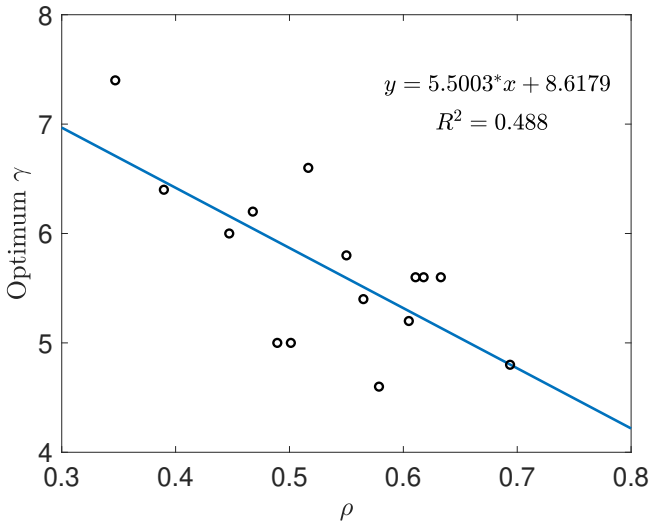

Figure 7.8: Optimum heterogeneity parameter vs. tree density, with linear regression.

In Fig. 7.9, we plot the individual Besag's functions from the 15 images from Finland. Then, we calculate the averaged Besag's function for the average density from Finland for different values of γ and compare this with the ensemble-averaged Besag's function from the real data. As observed previously, the peak is sharp for the generated data and more rounded for the real data.

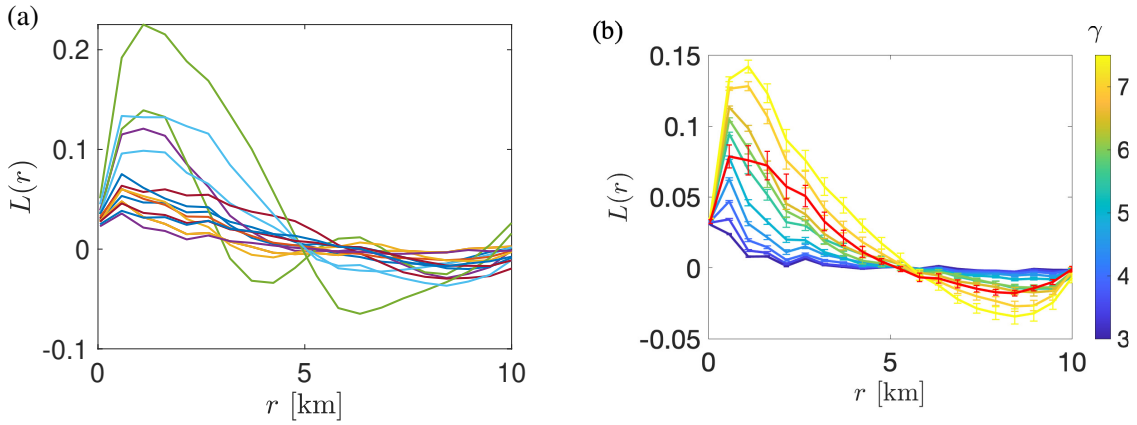

Figure 7.9: Panel (a), the individual Besag's functions for the real data from Finland. Panel (b), Besag's function for the generated data for different γ. The red curve is the ensemble-average of the real data.

7.7.4 Canada Data

In this section, we present the density and optimum γ estimated for images from Canada. The acquired images, the choice of sampling boxes, and the resulting binary representations are given in Appx. C. We summarise the densities and optimum γ for the forestscapes in table 7.2. The mean density is 0.6116 and the mean optimum γ is 5.880. Using the total sum of square errors over all forestscapes, the overall best-estimate is $\gamma = 5.8$.

Plot ref.	Density	Optimum γ
Fig. C.1(a)	0.4170	4.6
Fig. C.1(c)	0.6971	6.4
Fig. C.1(e)	0.7288	6.4
Fig. C.1(g)	0.5753	5.6
Fig. C.1(i)	0.6522	6.4
Fig. C.1(k)	0.6458	5.6
Fig. C.1(m)	0.6561	6.2
Fig. C.1(o)	0.5342	5.4
Fig. C.2(a)	0.6338	7.8
Fig. C.2(c)	0.6572	4.2
Fig. C.2(e)	0.6478	5.6
Fig. C.2(g)	0.5192	7.8
Fig. C.2(i)	0.6866	5.4
Fig. C.2(k)	0.5405	5.6
Fig. C.2(m)	0.5838	5.2

Table 7.2: The density and optimum heterogeneity parameters for 15 different 10km x 10km satellite images of forests from Canada.

In Fig. 7.10 we plot the optimum heterogeneity parameter against density. Discarding outliers and fitting a line to the data, we find that, in contrast to the data from Finland, the optimum heterogeneity parameter increases as the tree density increases.

Although the climate is comparable between Canada and Finland, we observe opposite trends in the relation between density and the optimum clustering parameter. It seems that there are factors other than the climate affecting this relation, but speculation on what they may be is beyond the scope of this work.

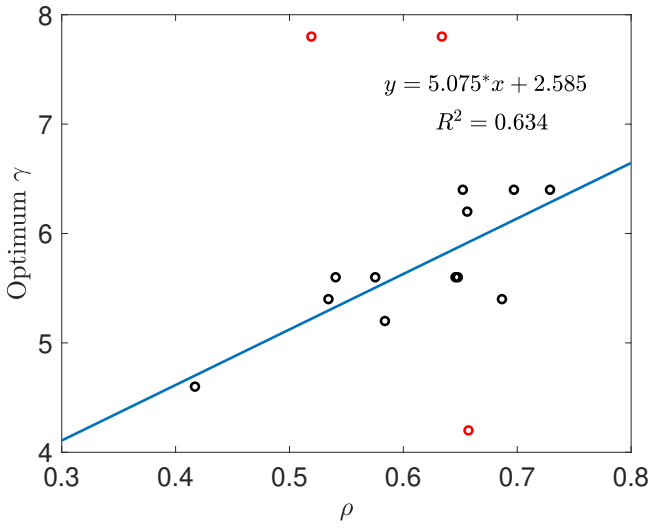

Figure 7.10: Optimum heterogeneity parameter vs. tree density, with linear regression. Outliers are shown as red discs, with the remaining data shown as black discs.

The outliers correspond to the forestscapes shown in Fig. C.2 (a), (c), and (g). On visual inspection, there is nothing obvious that marks these particular forestscapes as distinct from the rest of the dataset. Note that the outliers have removed from the line of the best fit, and no further analysis for this data set was performed.

In Fig. 7.11, we plot the individual Besag's functions from the 15 images from Canada. Then, we calculate the averaged Besag's function for the average density from Canada for different values of γ, and compare this with the ensemble-averaged Besag's function from the real data.

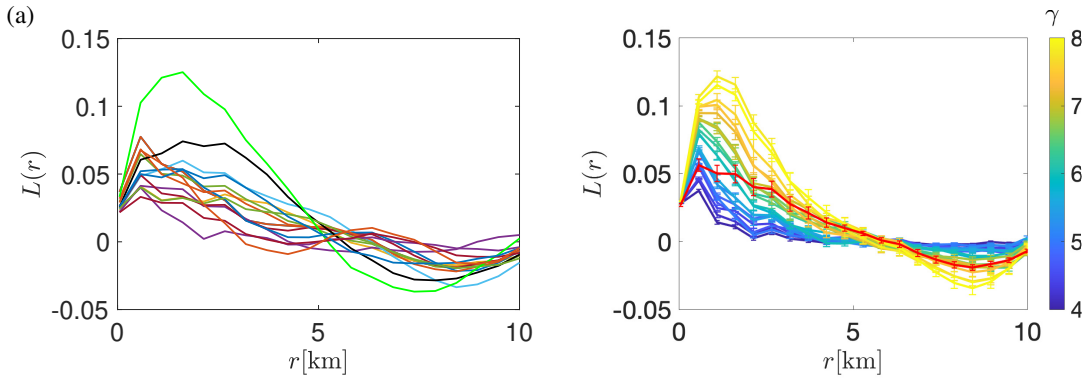

Figure 7.11: Panel (a): the individual Besag's functions for the real data from Canada. Two curves in particular stand out as different than the others: the green ($\rho = 0.51$, Fig. C.2 (h)) and the black ($\rho = 0.63$, Fig. C.2 (b)). Panel (b): the ensemble-averaged Besag's function for this data.

7.8 Conclusions

We used forests images provided by Google Maps for natural forests from two different countries: Finland and Canada. Then, using MATLAB, we cut out regions that are $\simeq (10\text{km})^2$. In total, we extracted 15 distinct natural forest images from each country.

To give a clear characteristic of the trees, we use a strategy based on the Mahalanobis distance. Because this is based on sampling a subset of each image, rather than trying to classify trees in all images simultaneously, we do not have to worry about differences between individual images, such as the colour balance or light levels being different. We can see by eye that the Mahalanobis distance gives a clear division between region with trees and the regions without trees.

We used the Mahalanobis distance to convert the forest images into binary images, for trees and non-trees. Having binary images allows us to estimate the tree density and the heterogeneity level for each of the natural forest images. Calculating tree density was straightforward, but the heterogeneity level for each image was estimated through least square difference.

We observed that, for the forestscapes captured from Finland, the heterogeneity parameter tends to decrease as density increases, while for the forestscapes captured from Canada, the trend is reversed, and we see the heterogeneity parameter tending to increase as the density is increased, even though they were chosen from natural forests in the same Köppen climate group. Reviewing the data by eye, it is clear that the synthetic forestscapes

are not a perfect match for the real data, as the real data can contain a wide variety of more complicated structures and systematic properties that are not included in our generation of synthetic data. However, in terms of the degree of clustering and segregation, we have produced reasonable representative data to work with.

We conclude that our method of generating clustered forestscapes is appropriate for generating data with comparable clustering to real data, but would need to be further developed if the effects of more complicated structures on disease spread were to be investigated.

Chapter 8

Conclusions

The goal of this thesis was to model disease propagation in natural forestscapes. We started with existing models and then developed them to have a better understanding of the problem. First, we independently reproduced the results from a simplified model used by Orozco-Fuentes *et al.* [47] where they used a temporal and spatial lattice-based model incorporated with a stochastic process to mimic the interactions between the susceptible population and the pathogen. We used homogeneous forestscapes, and allowed the diseases to spread via the von Neumann neighbourhood rules, where an infected cell can only infect cells to its right, left, above and below with an infection rate β. We characterised the observed dynamics by measuring the spreading velocity. We found that there is a critical transition, depending on the density of the forestscapes, that separate two stable phases: local containment and widespread outbreak. Also, we used the established framework of early warning signals for the velocity to predict the occurrence of the transition to epidemic. We defined the velocity in three distinct ways and implemented the early warning signals of the velocity for the three definitions. We showed that they all lead to the same transition to epidemic, which is dependent on the tree density. Regarding the infection rate β, we found that as β increases, the transition to epidemic tends to the critical threshold of percolation theory ρ_c. We also introduced the mortality rate, which is another statistical measure to quantify the observed dynamics. The mortality rate showed an identical story for the transition to the epidemic, and the effect of the infection rate, as the velocity.

The von Neumann neighbourhood is not the best method to model neighbour-to-neighbour infection propagation. To make the model more realistic, we studied the infection propagation using the Moore neighbourhood, which allows the disease to spread to the four cardinal directions and the diagonals as well. Measuring the spreading velocity, we noticed that the Moore neighbourhood rules tell the same qualitative story as the von Neumann neighbourhood rules. However, for the Moore neighbourhood, the transition to epidemic occurs at much lower tree density. This is associated with the appearance of the

spanning cluster in the percolating system. For the von Neumann rules, its first appearance in the system is ≈0.60 while for the Moore neighbourhood is ≈0.40.

The distribution of natural forestscapes is actually heterogeneous with different degrees of heterogeneity. The heterogeneity of the forestscapes can be natural or anthropogenic. Therefore, we synthetically generated forestscapes with different heterogeneity levels, with a single parameter, γ, governing the degree of heterogeneity. For $\gamma = 0$, the forestscape is homogeneous, and as γ increases, the heterogeneity level increases. Before introducing the dynamics into our synthetically generated forestscapes, we studied the properties of the network. We discussed the spanning cluster phenomenon and its first appearance in the system. We see that as heterogeneity increases, the occurrence of the spanning cluster happens at different densities. Using the Moore neighbourhood as a method of connectivity between occupied cells, we found that as γ increases, the spanning cluster requires a larger density to first appear in the system. Then, through Besag's function, we characterised the observed spatial clustering and segregation of the generated forestscapes.

Initialising the disease in these heterogeneous forestscapes using Moore neighbourhood rules, we observed that heterogeneity suppressed the velocity of the infection spreading, causing the epidemic to occur at much larger densities. That is, as γ increases, the abrupt shift from confinement to epidemic occurs at much larger density. The mortality rate showed the same behaviour as the velocity. Again, this is associated with the delay caused by the clustering to the first appearance of the spanning cluster as γ increases. On the other hand, the infection rate has a sole effect which is that as β decreases, it causes the transition to epidemic to occur at larger ρ. This is an identical effect of the infection rate as we have seen in the homogeneous forestscapes.

The spread of tree disease, pests or other pathogens is actually non-local. Therefore, we have used non-local kernels to model the infection spreading. We generalise the method of infection from a constant probability of transmission between direct neighbours to a transmission probability that depends on the separation between cells. After investigating several possible distributions, we decide to use a Gaussian kernel, controlled by the standard deviation of the Gaussian σ, which can be thought of as the characteristic spreading distance. We then normalise this distribution such that the number of new infected expected in a single timestep at the start of an infection is independent of σ, allowing us to focus on how the non-local spreading affects the dynamics and make direct comparisons with the local infection.

We then explore how the established parameters, tree density, transmissibility, and heterogeneity, affect the infection dynamics in this context, and how the new parameter σ interacts with them. The connectedness of forestscapes was the critical component for predicting the transition to epidemic in local infections. However, the connectedness of a forestscape is no longer clearly defined when we introduce spatial dependence of the

transmission. We also find that the velocity is no longer a clear indicator of whether an infection is widespread or contained. Even so, the mortality is still a good indicator of the state the infection is in. We again observe a critical transition at a particular density and find that for almost all combinations of parameters, the critical density is lower than for the corresponding local infection. The trends in the critical density are broadly similar to those for the local infections, with decreasing β shifting the critical density to higher values, and increasing γ shifting the critical density to higher values. However, we also observe that increasing σ can dramatically alter the critical density, even for very low β and very high γ.

The interaction between γ and σ is most apparent in the suppression of the infection velocity. The suppression of the infection velocity by γ observed for local infections is again seen for non-local infections, but is diminished or even completely overcome by increasing σ. This is due to the non-local spreading allowing the infection to "jump" over gaps in the forestscape, with competition between increased heterogeneity producing increasingly "gappy" forestscapes, and increased σ making it more likely that an infection will manage to "jump" the gaps.

We then investigate the number of infected trees I as a function of time, and empirically identify how this scales with the various parameters. For fixed ρ and β, and varying σ and γ we find that rescaling curves $I \to I/(\sigma)$ and $t \to \sigma t$ causes the majority of curves to collapse. Likewise, for fixed σ and β, and varying ρ and γ, we find scaling $I \to I/(\rho^{3/2})$ and $t \to \sqrt{\rho}t$, and for fixed σ and ρ, and varying β and γ, we find scaling $I \to I/(\sqrt{\beta})$ and $t \to \sqrt{\beta}t$.

The scaling does not hold for all parameter combinations, particularly tending to fail for large γ, low β, and low σ. We identify two easily observable patterns in the dynamics which accompany deviation from the scaling. Firstly, in many of these cases, the infection does not spread evenly, but spreads further in some directions than others; this can be fairly subtle, with the infection wavefront showing moderate fluctuations, or it can be very pronounced, with the infection spreading entirely in one direction. Then, for some cases where the transmissibility is low, and the infection width is low, we see a systematic increase in surviving cells, with increasingly large clusters entirely escaping the disease, and with many cells close to the edge of clusters not being infected.

There are several complex behaviours observable from this simple implementation of non-local interactions which are not present in local infection simulations. Even using the smallest values of σ, and with the transmissibility normalised to match the local case, the dynamics observed can be radically different. This is relevant to any attempts to establish forest configurations which are robust or vulnerable to disease outbreaks; for example, investigating the optimal size of breaks to disrupt disease propagation could not be done in a local context.

Finally, we extract real forestscapes from satellite images of Finland and Canada, classifying the points as either trees or not trees using the Mahalanobis distance. We estimate the heterogeneity parameter of each forestscape, and compare them to our synthetically generated data. We conclude that our synthetically generated forestscapes are a reasonable model for the clustering and segregation seen in real forests, but do not include other structure present in real forests, such as large scale anisotropy and large scale structures, such as rivers and roads.

There any many interesting directions that we could take in future with this work. A simple and potentially useful investigation, using the same model with non-local infections, could be done to investigate the size of breaks in a forestscape that are necessary to prevent the spread of infections, depending on the characteristic infectious distance and transmissibility.

The model could also be extended to include more realistic scenarios. For example, populations often exhibit a range of susceptibility to disease. This could be modelled by assigning individual transmissibility probabilities, drawn from a known distribution, to individual trees. A simplified version could simply consider two populations, one with high susceptibility and one with low susceptibility.

Another option for extending the model is to move away from the grid - in the real world, the positions of trees are not confined to a grid, and as the probabilities are based on separation in the non-local model there is no need to confine ourselves. Although it was useful here to allow a comparison between the local and non-local models, moving away from a grid could allow a more realistic examination of how tree density affects infections. An alternative approach which has been demonstrated in literature, e.g. in the context of disease outbreaks in wild animal populations [96] and infection transmission in territorial populations [97] is to use points which aren't fixed in a lattice. This may allow us to map to reality better. The framework we put forward in Chapter 6, where we have used Gaussian kernel interactions, is not confined to a grid, so further studies would be able to use this without additional modification.

Finally, one could do further analysis than our analysis and look at how properties such as q related to the tree density or the overall dynamics. They could look at how these affect velocity or mortality.

Appendices

Appendix A

The Infection Spreading Algorithm

This algorithm describes the full simulation, from the initialisation with a chosen input forestscape, determining which cells change from susceptible to infected, and from infected to removed, in each timestep, and the termination of the simulation, for both local and non-local infection spreading mechanisms.

- input matrix describing location of (susceptible) trees
- create matrices describing location of infected trees and dead trees
- create matrix describing how long trees have been infected for
- while (there is at least 1 infected tree & the infection has not reached the boundary)
- for (each infected tree)
- for (each susceptible cell if non-local, each occupied neighbour if local)
- get probability P_{inf} that infection spreads from that infected tree to that suspectible
| tree based on distance between them
- if ($P_{\text{inf}} > U(0,1)$)
- remove susceptible from matrix of susceptibles, store in new infected
- end if
- end for
- end for
- add all new infected trees to matrix of infected trees
- increment matrix of length of infection by 1 for all infected trees
- remove trees that have been infected for 10 time steps from matrix of infected and add
| matrix of dead
- end while

where P_{inf} is a function of r defined according to eq. (6.4) for the case where we have non-local interactions, and has the form:

$$P_{\text{inf}} = \begin{cases} \beta, & \text{if direct neighbours,} \\ 0, & \text{otherwise.} \end{cases}$$

for the 4- and 8-neighbourhood interactions.

Rather than finding the analytic probability of a susceptible tree being infected given the configuration of already infected trees, we find the probability of it being infected by each infected tree separately, and take the infection to have spread if it has been infected by any of these infected trees - so this is not a sum of probabilities, but whether it has been infected by the first tree, or the second, or the n^{th}.

Appendix B

Forest Images From Finland

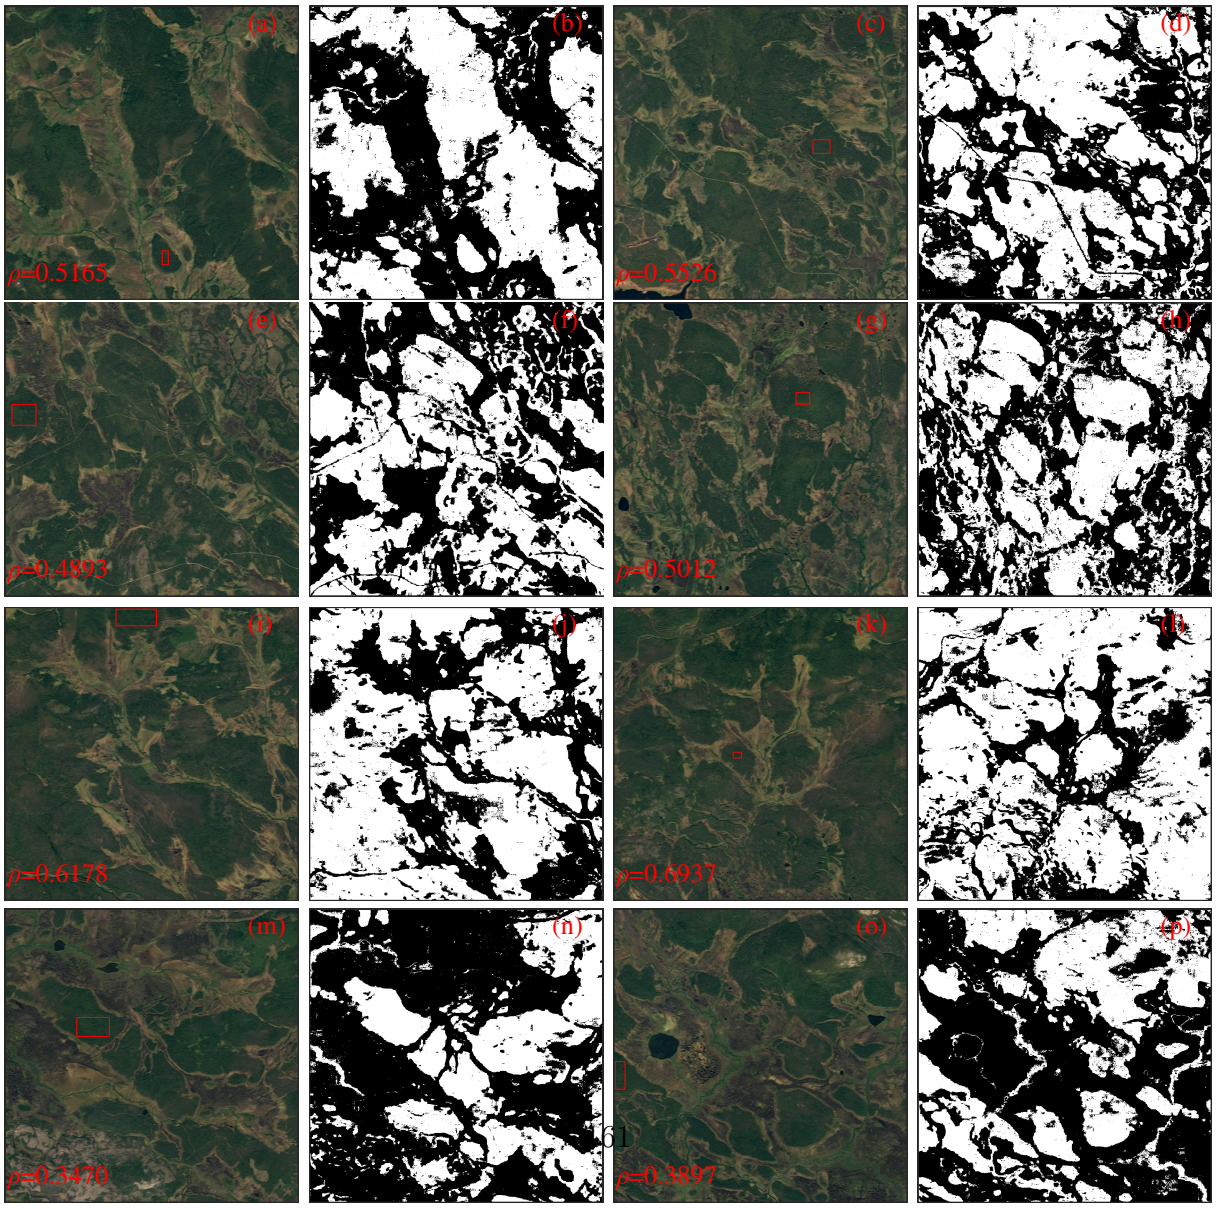

Figure B.1: Real images with the sampling boxes highlighted in red, and their binary matrices.

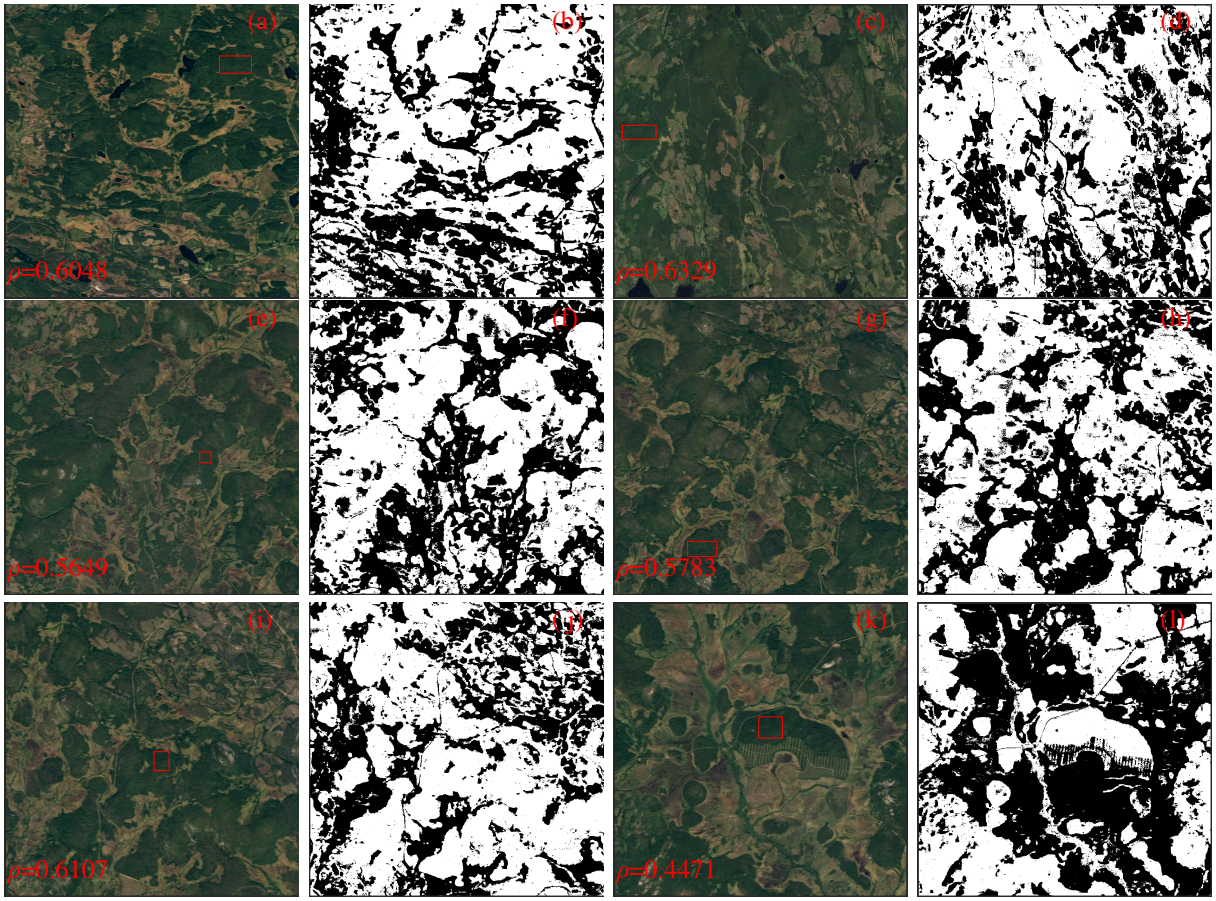

Figure B.2: Real images with the sampling boxes highlighted in red, and their binary matrices.

In Fig. B.1 and Fig. B.2, we display our 14 images from Finland alongside the sampling box, coloured in red, tree density ρ, and their binary matrices. Note that for our analysis of the images from Finland, we used 15 images. The first image was displayed and analysed in more details through Chapter 7 while the other 14 are shown here. In Fig. B.3, we plot Besag's function for each image from Finland in red. Then, using the density of the real forestscape, we calculated Besag's function for a range of different γ between 4 and 8. Calculating Besag's function for this range of γ allowed us to estimate the optimum γ for each particular image.

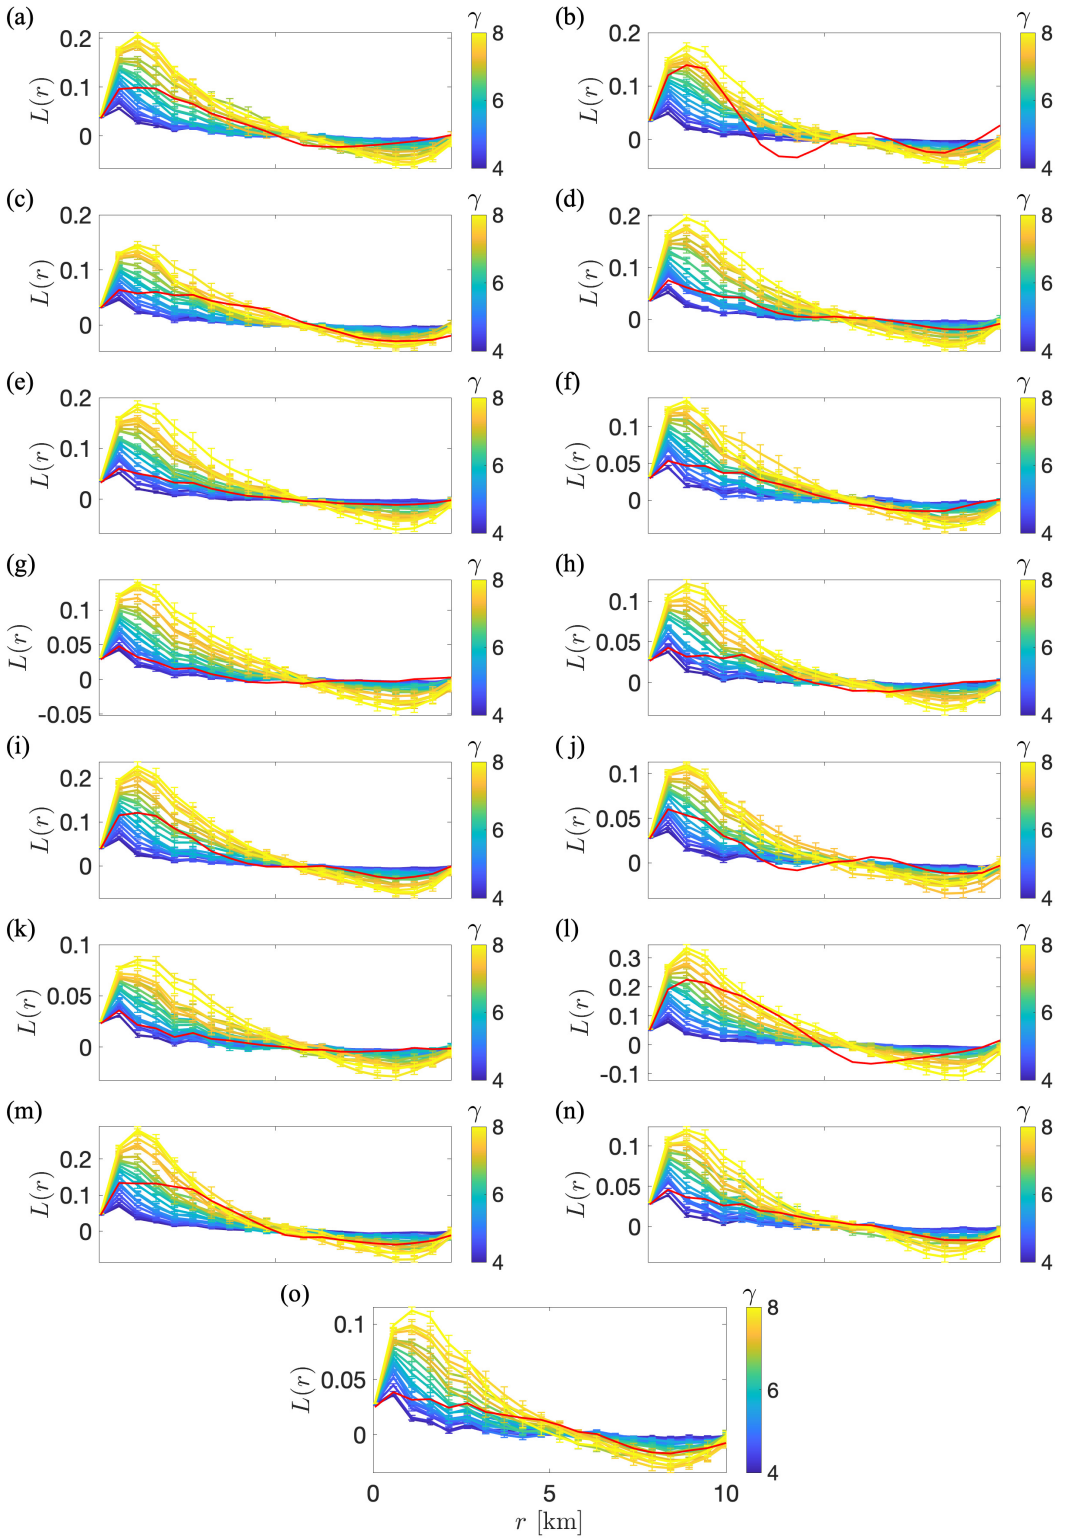

Figure B.3: The red curve represents the calculated Besag's Function for a real image from Finland which has a certain density, while the other curves represent the calculated Besag's function for the same density for different values of γ averaged over 30 ensembles.

Appendix C

Forest Images from Canada

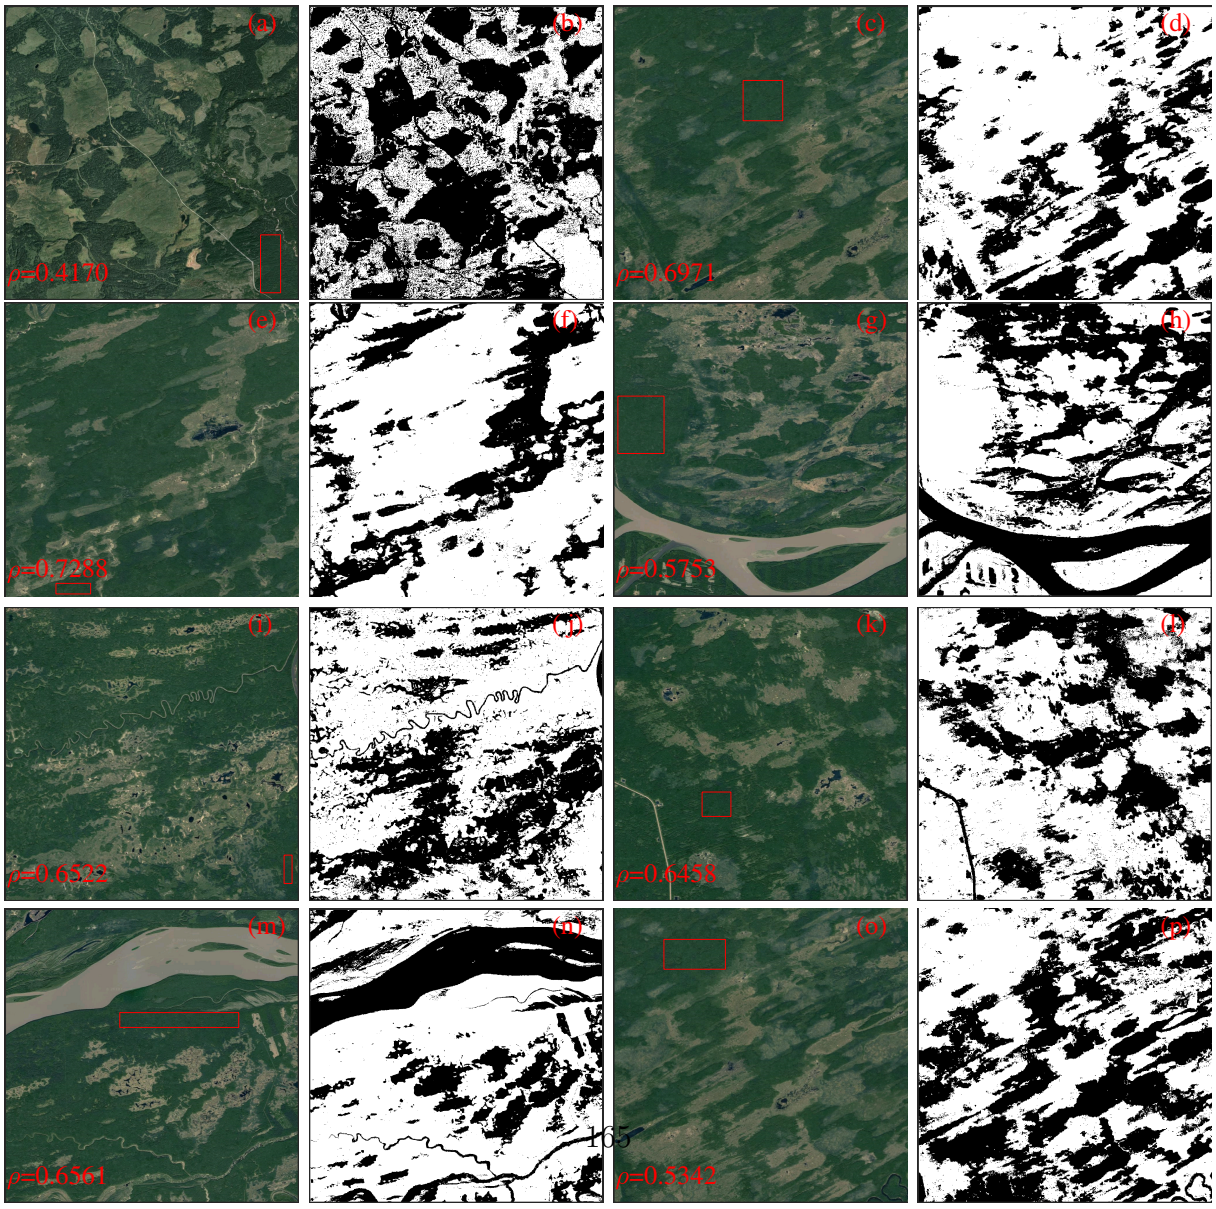

Figure C.1: Real images with the sampling boxes, and their binary matrices.

For the real data, our grids are 10km x 10km. We managed to extract 15 different images for forest trees, which we show in Fig. C.1 and Fig. C.2.

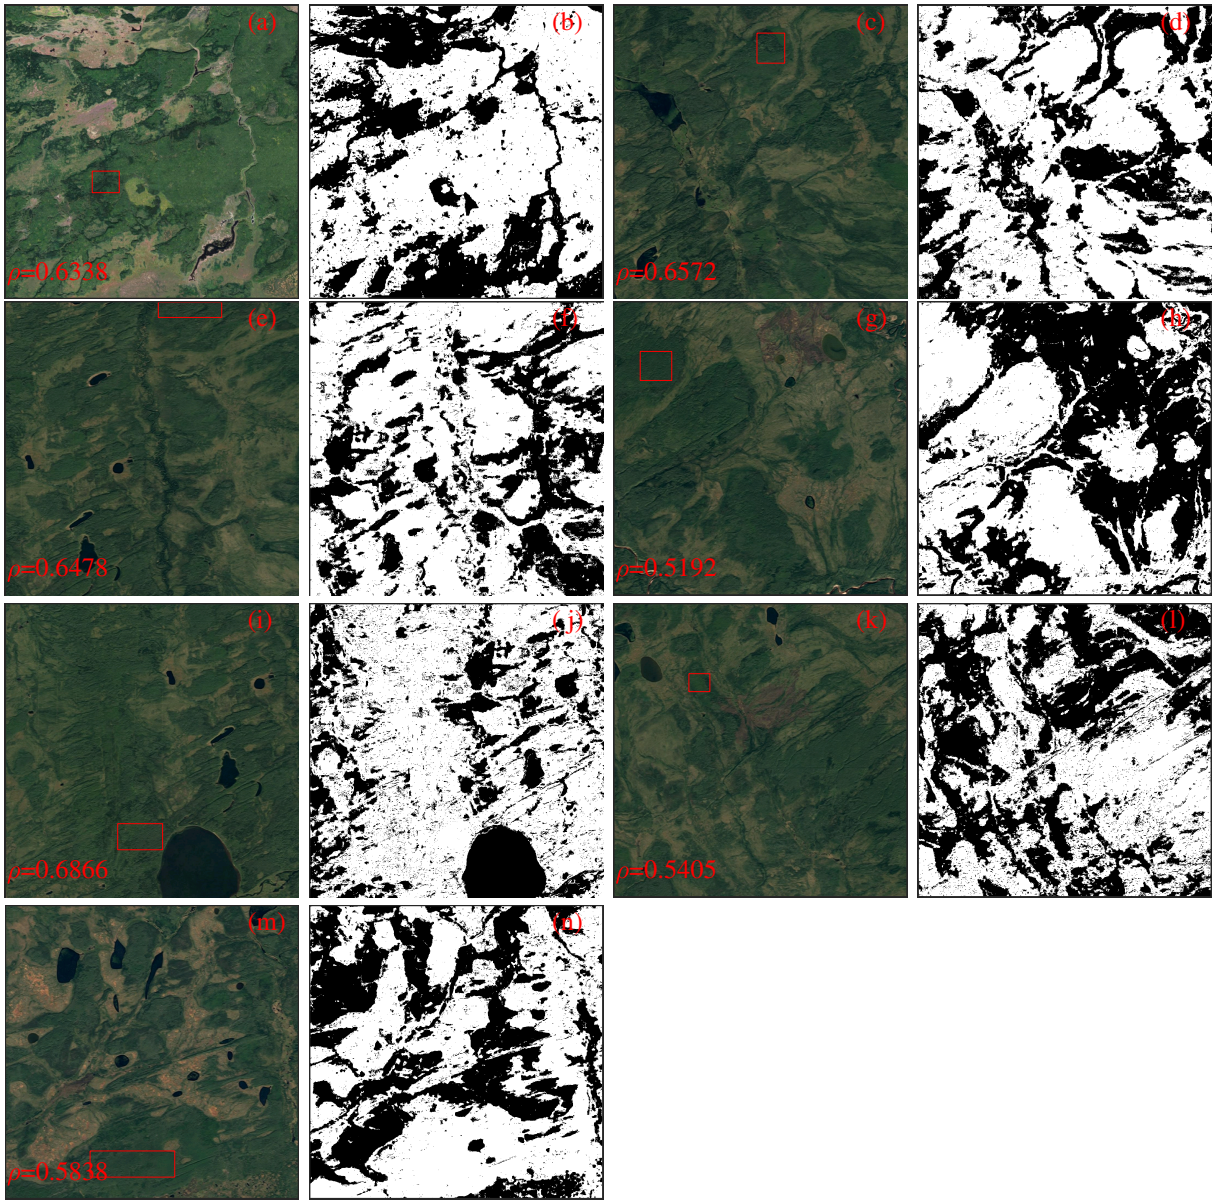

Figure C.2: Real images with the sampling boxes, and their binary matrices.

We now compare Besag's function for the individual real forestscapes with Besag's function for the associated synthetic data.

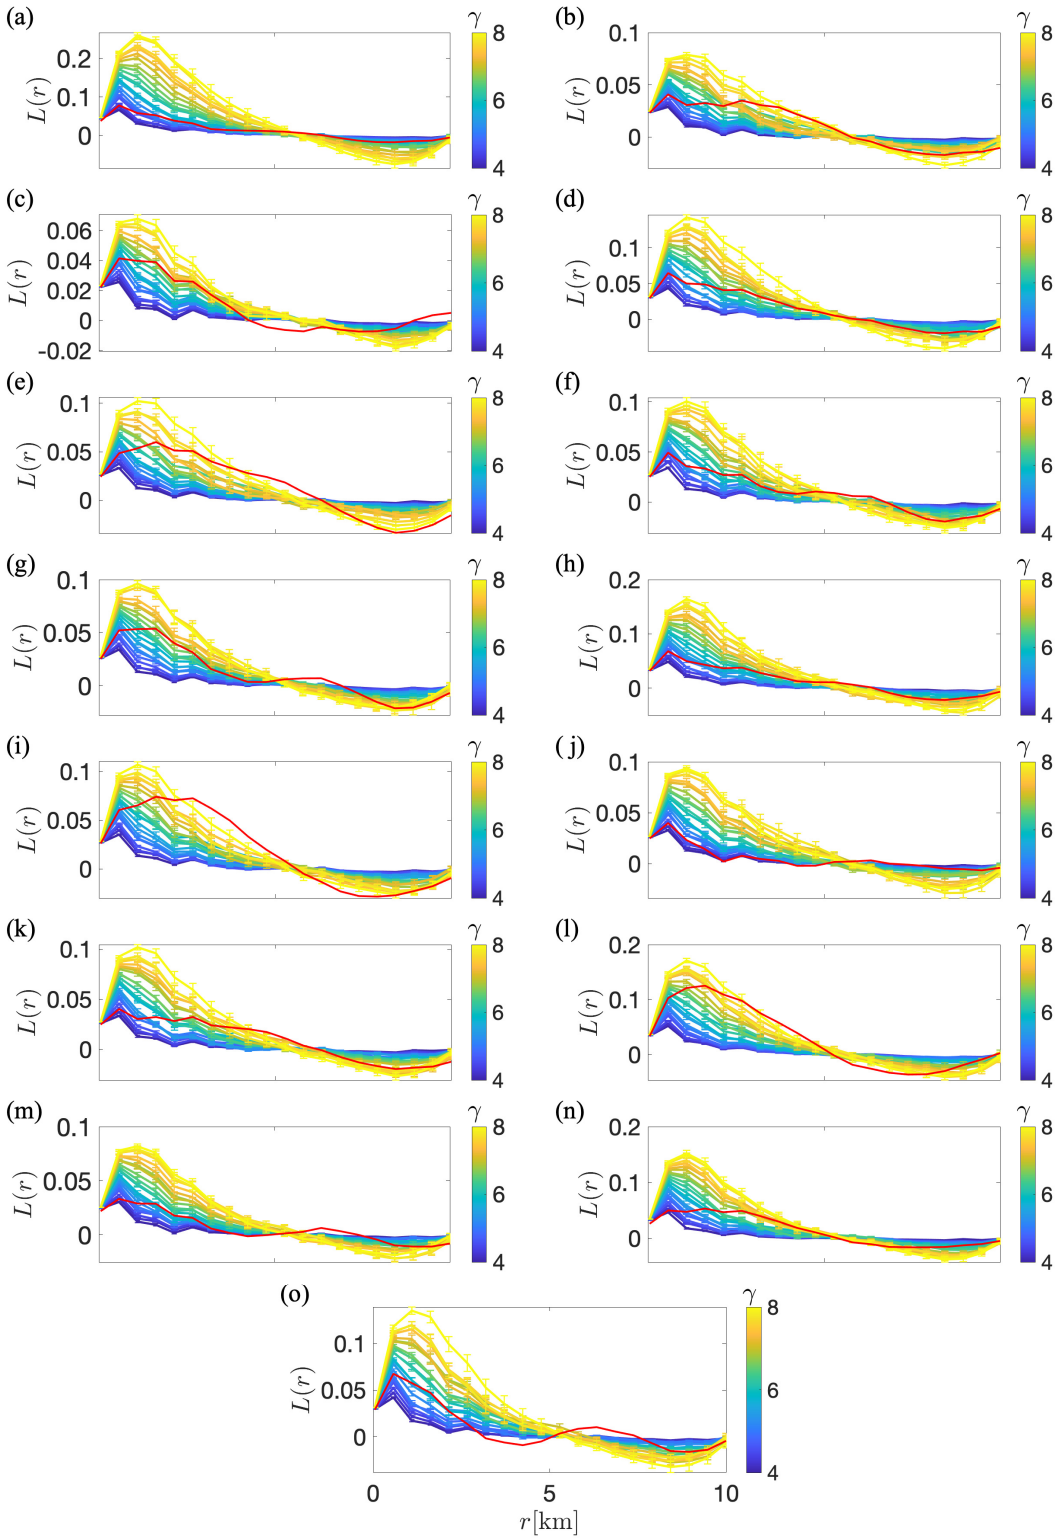

Figure C.3: The red curve represents the calculated Besag's function for a real image from Canada, which has a certain density, while the other curves represent Besag's function calculated for synthetic forestscapes with the same density, for different value of γ.

Bibliography

[1] L. Hill, G. Jones, N. Atkinson, A. Hector, G. Hemery, and N. Brown, "The £15 billion cost of ash dieback in Britain," *Current Biology*, vol. 29, no. 9, pp. R315–R316, 2019.

[2] M. Djerbi, "Diseases of the date palm: present status and future prospects," *Journal of Agricultural and Marine Sciences [JAMS]*, vol. 3, no. 1, pp. 103–114, 1998.

[3] T. P. Holmes, J. E. Aukema, B. Von Holle, A. Liebhold, and E. Sills, "Economic impacts of invasive species in forest past, present, and future," *In: The Year In Ecology and Conservation Biology, 2009. Ann. NY Acad. Sci. 1162: 18-38.*, vol. 1162, pp. 18–38, 2009.

[4] J. E. Aukema, B. Leung, K. Kovacs, C. Chivers, K. O. Britton, J. Englin, S. J. Frankel, R. G. Haight, T. P. Holmes, A. M. Liebhold, *et al.*, "Economic impacts of non-native forest insects in the continental united states," *PLOS ONE*, vol. 6, no. 9, 2011.

[5] I. Boyd, P. Freer-Smith, C. Gilligan, and H. Godfray, "The consequence of tree pests and diseases for ecosystem services," *Science*, vol. 342, no. 6160, p. 1235773, 2013.

[6] P. H. Freer-Smith and J. F. Webber, "Tree pests and diseases: the threat to biodiversity and the delivery of ecosystem services," *Biodiversity and Conservation*, vol. 26, no. 13, pp. 3167–3181, 2017.

[7] G. Allard, S. Fortuna, L. S. See, J. Novotny, A. Baldini, and T. Coutinho, "Global information on outbreaks and impact of major forest insect pests and diseases," in *XII World Forestry Congress, Quebec City, Canada*, 2003.

[8] Forests Restoration Act, "The healthy forests initiative and healthy forests restoration act," 2004.

[9] D. Cappaert, D. G. McCullough, T. M. Poland, and N. W. Siegert, "Emerald ash borer in North America: a research and regulatory challenge," *American Entomologist. 51 (3): 152-165.*, vol. 51, no. 3, 2005.

[10] D. J. Crook, A. Khrimian, A. Cossé, I. Fraser, and V. C. Mastro, "Influence of trap color and host volatiles on capture of the emerald ash borer (coleoptera: Buprestidae)," *Journal of Economic Entomology*, vol. 105, no. 2, pp. 429–437, 2012.

[11] W. S. Klooster, D. A. Herms, K. S. Knight, C. P. Herms, D. G. McCullough, A. Smith, K. J. Gandhi, and J. Cardina, "Ash (fraxinus spp.) mortality, regeneration, and seed bank dynamics in mixed hardwood forests following invasion by emerald ash borer (agrilus planipennis)," *Biological Invasions*, vol. 16, no. 4, pp. 859–873, 2014.

[12] D. A. Herms and D. G. McCullough, "Emerald ash borer invasion of north america: history, biology, ecology, impacts, and management," *Annual Review of Entomology*, vol. 59, pp. 13–30, 2014.

[13] M. Pautasso, G. Aas, V. Queloz, and O. Holdenrieder, "European ash (fraxinus excelsior) dieback–a conservation biology challenge," *Biological Conservation*, vol. 158, pp. 37–49, 2013.

[14] "Assessing the impacts of ash dieback." "`https://www.britishecologicalsociety.org/assessing-the-impacts-of-ash-dieback/`", 2013. [Online; accessed 14-October-2020].

[15] S. L. Anagnostakis, "Chestnut blight: the classical problem of an introduced pathogen," *Mycologia*, vol. 79, no. 1, pp. 23–37, 1987.

[16] G. G. Parker, S. M. Hill, and L. A. Kuehnel, "Decline of understory American chestnut (Castanea dentata) in a southern Appalachian forest," *Canadian Journal of Forest Research*, vol. 23, no. 2, pp. 259–265, 1993.

[17] R. K. Meentemeyer, N. J. Cunniffe, A. R. Cook, J. A. Filipe, R. D. Hunter, D. M. Rizzo, and C. A. Gilligan, "Epidemiological modeling of invasion in heterogeneous landscapes: spread of sudden oak death in california (1990–2030)," *Ecosphere*, vol. 2, no. 2, pp. 1–24, 2011.

[18] N. J. Cunniffe, R. C. Cobb, R. K. Meentemeyer, D. M. Rizzo, and C. A. Gilligan, "Modeling when, where, and how to manage a forest epidemic, motivated by sudden oak death in california," *Proceedings of the National Academy of Sciences*, vol. 113, no. 20, pp. 5640–5645, 2016.

[19] D. M. Rizzo and M. Garbelotto, "Sudden oak death: endangering California and Oregon forest ecosystems," *Frontiers in Ecology and the Environment*, vol. 1, no. 4, pp. 197–204, 2003.

[20] `https://www.gov.uk/government/organisations/forestry-commission`. [Online; accessed 14-October-2020].

[21] "Lutèce, a resistant variety brings elms back to Paris." `https://web.archive.org/web/20061125162935/http://www.international.inra.fr/research/some_examples/lutece_r_a_resistant_variety_brings_elms_back_to_paris`, 2006. [Online; accessed 14-October-2020].

[22] S. K. Abdullah, L. V. Lopez-Lorca, and H. B. Jansson, "Diseases of date palm (Phoenix dactylifera l.).," vol. 9, no. 2, 2010.

[23] C. Killian, R. Maire, *et al.*, "'Baïoud', a disease of the date palm.," *Bulletin de la Societe d'Histoire Naturelle de l'Afrique du Nord*, vol. 21, no. 6-7, 1930.

[24] S. Abdullah, L. Asensio, E. Monfort, S. Gomez-Vidal, J. Palma-Guerrero, J. Salinas, L. Lopez-Llorca, H.-B. Jansson, and J. Guarro, "Occurrence in Elx, SE Spain of inflorescence rot disease of date palms caused by Mauginiella scaettae," *Journal of Phytopathology*, vol. 153, no. 7-8, pp. 417–422, 2005.

[25] K. Al-Hassan and B. Waleed, "Biological study on Mauginiella scaettae Cav. the cause of inflorescence rot of date palms in Iraq," *Yearbook Plant Prot. Res. Min. Agric. Agrar. Ref. Iraq*, vol. 1, pp. 223–236, 1977.

[26] E. Arias-Jiménez, *Date palm cultivation*. No. 156, FAO, 2002.

[27] G. M. Lovett, M. Weiss, A. M. Liebhold, T. P. Holmes, B. Leung, K. F. Lambert, D. A. Orwig, F. T. Campbell, J. Rosenthal, D. G. McCullough, *et al.*, "Nonnative forest insects and pathogens in the United States: impacts and policy options," *Ecological Applications*, vol. 26, no. 5, pp. 1437–1455, 2016.

[28] M. Daugherty and K. Hung, "Sudden Oak Death." `https://cisr.ucr.edu/invasive-species/sudden-oak-death`. [Online; accessed 14-October-2020].

[29] D. E. Davis, "Historical significance of American chestnut to Appalachian culture and ecology," *Restoration of American Chestnut to Forest Lands*, pp. 53–60, 2006.

[30] R. Ganley and L. Bulman, "Dutch elm disease in New Zealand: impacts from eradication and management programmes," *Plant Pathology*, vol. 65, no. 7, pp. 1047–1055, 2016.

[31] S. Yao and K. Al-Redhaiman, "Date palm cultivation in Saudi Arabia: Current status and future prospects for development," in *Hortscience*, vol. 49, pp. S139–S140, American Society of Horticultural Science, 2014.

[32] "Volume of date production in Saudi Arabia from 2010 to 2014 (in metric tons)." `https://www.statista.com/statistics/693947/saudi-arabia-date-production/`, 2020. [Online; accessed 14-October-2020].

[33] "Tree health resilience strategy," *Department for Environment Food & Rural Affairs*, 2018.

[34] J. A. Lampietti and J. A. Dixon, *To see the forest for the trees: A guide to non-timber forest benefits*. The World Bank, 1995.

[35] Z. Munzbergova and D. Ward, "Acacia trees as keystone species in Negev desert ecosystems," *Journal of Vegetation Science*, vol. 13, no. 2, pp. 227–236, 2002.

[36] A. D. Manning, J. Fischer, and D. B. Lindenmayer, "Scattered trees are keystone structures–implications for conservation," *Biological Conservation*, vol. 132, no. 3, pp. 311–321, 2006.

[37] K. Stagoll, D. B. Lindenmayer, E. Knight, J. Fischer, and A. D. Manning, "Large trees are keystone structures in urban parks," *Conservation Letters*, vol. 5, no. 2, pp. 115–122, 2012.

[38] A. Fuentealba, R. Alfaro, and É. Bauce, "Theoretical framework for assessment of risks posed to Canadian forests by invasive insect species," *Forest Ecology and Management*, vol. 302, pp. 97–106, 2013.

[39] A. M. Bate, G. Jones, A. Kleczkowski, A. MacLeod, R. Naylor, J. Timmis, J. Touza, and P. C. White, "Modelling the impact and control of an infectious disease in a plant nursery with infected plant material inputs," *Ecological Modelling*, vol. 334, pp. 27–43, 2016.

[40] M. F. Macpherson, A. Kleczkowski, J. R. Healey, C. P. Quine, and N. Hanley, "The effects of invasive pests and pathogens on strategies for forest diversification," *Ecological Modelling*, vol. 350, pp. 87–99, 2017.

[41] G. Wurzer, K. Kowarik, and H. Reschreiter, *Agent-based modeling and simulation in archaeology*. Springer, 2015.

[42] E. Bonabeau, "Agent-based modeling: Methods and techniques for simulating human systems," *Proceedings of the National Academy of Sciences*, vol. 99, no. suppl 3, pp. 7280–7287, 2002.

[43] L. Schulman and P. Seiden, "Statistical mechanics of a dynamical system based on Conway's game of life," *Journal of Statistical Physics*, vol. 19, no. 3, pp. 293–314, 1978.

[44] P. Bak, K. Chen, and C. Tang, "A forest-fire model and some thoughts on turbulence," *Physics Letters A*, vol. 147, no. 5-6, pp. 297–300, 1990.

[45] T. Beer and I. Enting, "Fire spread and percolation modelling," *Mathematical and Computer Modelling*, vol. 13, no. 11, pp. 77–96, 1990.

[46] P. Grassberger, "On a self-organized critical forest-fire model," *Journal of Physics A: Mathematical and General*, vol. 26, no. 9, p. 2081, 1993.

[47] S. Orozco-Fuentes, G. Griffiths, M. Holmes, R. Ettelaie, J. Smith, A. Baggaley, and N. Parker, "Early warning signals in plant disease outbreaks," *Ecological Modelling*, vol. 393, pp. 12–19, 2019.

[48] M. Blomhøj, T. Kjeldsen, and J. Ottesen, "Compartment models," 2014.

[49] W. O. Kermack and A. G. McKendrick, "A contribution to the mathematical theory of epidemics," *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, vol. 115, no. 772, pp. 700–721, 1927.

[50] X. Shi and G. Song, "Analysis of the mathematical model for the spread of pine wilt disease," *Journal of Applied Mathematics*, vol. 2013, 2013.

[51] J. Murray, "Models for interacting populations," *Mathematical Biology: I. An Introduction*, pp. 79–118, 2002.

[52] Y. Hosono and B. Ilyas, "Traveling waves for a simple diffusive epidemic model," *Mathematical Models and Methods in Applied Sciences*, vol. 5, no. 07, pp. 935–966, 1995.

[53] A. Ducrot and H. Matano, "Plant disease propagation in a striped periodic medium," in *Applied Analysis in Biological and Physical Sciences*, pp. 121–164, Springer, 2016.

[54] M. J. Jeger, M. Pautasso, O. Holdenrieder, and M. W. Shaw, "Modelling disease spread and control in networks: implications for plant sciences," *New Phytologist*, vol. 174, no. 2, pp. 279–297, 2007.

[55] C. Rhodes and R. Anderson, "Dynamics in a lattice epidemic model," *Physics Letters A*, vol. 210, no. 3, pp. 183–188, 1996.

[56] S. R. Abades, A. Gaxiola, and P. A. Marquet, "Fire, percolation thresholds and the savanna forest transition: a neutral model approach," *Journal of Ecology*, vol. 102, no. 6, pp. 1386–1393, 2014.

[57] T. Manabe, N. Nishimura, M. Miura, and S. Yamamoto, "Population structure and spatial patterns for trees in a temperate old-growth evergreen broad-leaved forest in japan," *Plant Ecology*, vol. 151, no. 2, pp. 181–197, 2000.

[58] L. Li, Z. Huang, W. Ye, H. Cao, S. Wei, Z. Wang, J. Lian, I.-F. Sun, K. Ma, and F. He, "Spatial distributions of tree species in a subtropical forest of China," *Oikos*, vol. 118, no. 4, pp. 495–502, 2009.

[59] H. Du, F. Hu, F. Zeng, K. Wang, W. Peng, H. Zhang, Z. Zeng, F. Zhang, and T. Song, "Spatial distribution of tree species in evergreen-deciduous broadleaf karst forests in southwest China," *Scientific Reports*, vol. 7, no. 1, pp. 1–9, 2017.

[60] `https://www.waldwissen.net/waldwirtschaft/waldbau/bergwald/wsl-rottenaufforstung/index-EN`. [Online; accessed 05-April-2018].

[61] `http://hickoryclusterassociation.blogspot.co.uk/p/trees.html`. [Online; accessed 16-October-2020].

[62] V. Timmermann, I. BØrja, A. M. Hietala, T. Kirisits, and H. Solheim, "Ash dieback: pathogen spread and diurnal patterns of ascospore dispersal, with special emphasis on Norway," *EPPO Bulletin*, vol. 41, no. 1, pp. 14–20, 2011.

[63] A. Chandelier, M. Helson, M. Dvorak, and F. Gischer, "Detection and quantification of airborne inoculum of Hymenoscyphus pseudoalbidus using real-time PCR assays," *Plant Pathology*, vol. 63, no. 6, pp. 1296–1305, 2014.

[64] N. J. Cunniffe, B. Koskella, C. J. E. Metcalf, S. Parnell, T. R. Gottwald, and C. A. Gilligan, "Thirteen challenges in modelling plant diseases," *Epidemics*, vol. 10, pp. 6–10, 2015.

[65] S. Riley, K. Eames, V. Isham, D. Mollison, and P. Trapman, "Five challenges for spatial epidemic models," *Epidemics*, vol. 10, pp. 68–71, 2015.

[66] T. Toffoli and N. Margolus, *Cellular automata machines: a new environment for modeling*. MIT press, 1987.

[67] A. Ben-Menahem, *Historical encyclopedia of natural and mathematical sciences*. Springer Science & Business Media, 2009.

[68] L. Gray, "A mathematician looks at Wolfram's new kind of science," *Notices-American Mathematical Society*, vol. 50, no. 2, pp. 200–211, 2003.

[69] T. C. Bailey and A. C. Gatrell, *Interactive spatial data analysis*, vol. 413. Longman Scientific & Technical Essex, 1995.

[70] N. Kenkel, "Modeling Markovian dependence in populations of Aralia nudicaulis," *Ecology*, vol. 74, no. 6, pp. 1700–1706, 1993.

[71] I. Prentice and M. Werger, "Clump spacing in a desert dwarf shrub community," *Vegetatio*, vol. 63, no. 3, pp. 133–139, 1985.

[72] C. Skarpe, "Spatial patterns and dynamics of woody vegetation in an arid savanna," *Journal of Vegetation Science*, vol. 2, no. 4, pp. 565–572, 1991.

[73] R. W. Sterner, C. A. Ribic, and G. E. Schatz, "Testing for life historical changes in spatial patterns of four tropical tree species," *The Journal of Ecology*, pp. 621–633, 1986.

[74] J. E. Besag, "Discussion of the paper by Ripley," *Journal of the Royal Statistical Society Series B*, vol. 39, pp. 193–195, 1977.

[75] P. Haase, "Spatial pattern analysis in ecology based on Ripley's K-function: Introduction and methods of edge correction," *Journal of Vegetation Science*, vol. 6, no. 4, pp. 575–582, 1995.

[76] D. Stauffer and A. Aharony, *Introduction to percolation theory*. Taylor & Francis, 2018.

[77] W. H. Stockmayer, "Theory of molecular size distribution and gel formation in branched polymers II. General cross linking," *The Journal of Chemical Physics*, vol. 12, no. 4, pp. 125–131, 1944.

[78] S. R. Broadbent and J. M. Hammersley, "Percolation processes: I. crystals and mazes," in *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 53, pp. 629–641, Cambridge University Press, 1957.

[79] R. M. Ziff and C. R. Scullard, "Exact bond percolation thresholds in two dimensions," *Journal of Physics A: Mathematical and General*, vol. 39, no. 49, p. 15083, 2006.

[80] K. Malarz and S. Galam, "Square-lattice site percolation at increasing ranges of neighbor bonds," *Physical Review E*, vol. 71, no. 1, p. 016125, 2005.

[81] A. A. Saberi, "Recent advances in percolation theory and its applications," *Physics Reports*, vol. 578, pp. 1–32, 2015.

[82] M. Scheffer, J. Bascompte, W. A. Brock, V. Brovkin, S. R. Carpenter, V. Dakos, H. Held, E. H. Van Nes, M. Rietkerk, and G. Sugihara, "Early-warning signals for critical transitions," *Nature*, vol. 461, no. 7260, pp. 53–59, 2009.

[83] M. Scheffer, *Critical transitions in nature and society*. Princeton University Press, 2009.

[84] S. R. Carpenter and W. A. Brock, "Rising variance: a leading indicator of ecological transition," *Ecology Letters*, vol. 9, no. 3, pp. 311–318, 2006.

[85] I. O. Morales, E. Landa, C. C. Angeles, J. C. Toledo, A. L. Rivera, J. M. Temis, and A. Frank, "Behavior of early warnings near the critical temperature in the two-dimensional Ising model," *PloS One*, vol. 10, no. 6, p. e0130751, 2015.

[86] C. A. Boulton, P. Good, and T. M. Lenton, "Early warning signals of simulated Amazon rainforest dieback," *Theoretical Ecology*, vol. 6, no. 3, pp. 373–384, 2013.

[87] R. Corrado, A. M. Cherubini, and C. Pennetta, "Early warning signals of desertification transitions in semiarid ecosystems," *Physical Review E*, vol. 90, no. 6, p. 062705, 2014.

[88] B. M. Rogers, K. Solvik, E. H. Hogg, J. Ju, J. G. Masek, M. Michaelian, L. T. Berner, and S. J. Goetz, "Detecting early warning signals of tree mortality in boreal North America using multiscale satellite data," *Global Change Biology*, vol. 24, no. 6, pp. 2284–2304, 2018.

[89] L.-G. Stener, "Clonal differences in susceptibility to the dieback of Fraxinus excelsior in southern Sweden," *Scandinavian Journal of Forest Research*, vol. 28, no. 3, pp. 205–216, 2013.

[90] A. M. Hietala, V. Timmermann, I. Børja, and H. Solheim, "The invasive ash dieback pathogen Hymenoscyphus pseudoalbidus exerts maximal infection pressure prior to the onset of host leaf senescence," *Fungal Ecology*, vol. 6, no. 4, pp. 302–308, 2013.

[91] S. P. Brand, M. J. Tildesley, and M. J. Keeling, "Rapid simulation of spatial epidemics: a spectral method," *Journal of Theoretical Biology*, vol. 370, pp. 121–134, 2015.

[92] T. Boutell and T. Lane, "PNG (Portable Network Graphics) specification version 1.0," *Network Working Group*, pp. 1–102, 1997.

[93] P. C. Mahalanobis, "On the generalized distance in statistics," National Institute of Science of India, 1936.

[94] R. De Maesschalck, D. Jouan-Rimbaud, and D. L. Massart, "The Mahalanobis distance," *Chemometrics and Intelligent Laboratory Systems*, vol. 50, no. 1, pp. 1–18, 2000.

[95] G. J. McLachlan, *Discriminant analysis and statistical pattern recognition*. John Wiley & Sons, 2004.

[96] S. W. Laffan, E. Lubarsky, M. P. Ward, and L. D. Highfield, "A geographic automata system for modelling disease outbreaks in wild and unfenced animal populations,"

[97] M. E. Craft, E. Volz, C. Packer, and L. A. Meyers, "Disease transmission in territorial populations: the small-world network of serengeti lions," *Journal of the Royal Society Interface*, vol. 8, no. 59, pp. 776–786, 2011.