Stochastic models of plant growth and competition

Simon Antony Croft

Thesis submitted for the degree of
Doctor of Philosophy

University of York
Department of Biology
York Centre for Complex Systems Analysis

June 2012

For Isla

Abstract

Plants have been observed to show a range of plastic responses to environmental conditions. For example, the abundance and distribution of nutrients, as well as the presence and proximity of local competition, have been seen to result in changes in root proliferation and architecture. However, whilst some species have been witnessed displaying certain responses under given circumstances, experimental evidence suggests that responses to environmental factors can be far from simple, and sometimes counter-intuitive. Plant responses to components of the environment, and the benefit of such responses, are highly context sensitive.

This thesis explores some of the real world complexities that result in the observed responses to hierarchical sets of environmental factors, and presents a theoretical model that seeks to elucidate the interplay between different factors and their effects on "optimal" behaviour by both individuals and populations.

Starting with a simple one-dimensional model comprising a linearised approximation of a Gompertz growth function with nutrient patch dependent growth, the individual and combined effects of stochasticity in resource and competitor distribution are investigated. Complexity and functionality are progressively built up, with a resource dependent proliferation response, a scaling up into two-dimensions, and finally different intrinsic plant growth strategies trading growth rate against root system efficiency all introduced and investigated.

Throughout the work presented in this thesis, complex and subtle behavioural responses and patterns emerge from seemingly simple models. The importance of stochasticity on individual and population level performance is also highlighted, and the results demonstrate the inability for mean-field approximations and expected results to capture the emergent behaviour.

Contents

List of Figures

List of Tables

Acknowledgements

First and foremost I would like to thank my wonderful team of supervisors, Angela Hodge and Jon Pitchford, who it feels have gone far beyond the call of duty at times. It took a while, but we got there in the end! It was always a pleasure working with them, and I am immensely grateful to have had the experience. Without their knowledge, expertise and professionalism (I promise I am not taking the piss), not to mention their continued help, support, encouragement and occasional kicks up the arse (thankfully, mostly figurative), this thesis would not have been completed. You now know who to blame. An additional mention must go to the rest of the Pitchford clan for their continued warmth and hospitality (whisky, curry and toys: what more could a boy ask for?).

Thanks also go to Alastair Fitter and Richard Law for forming a helpful and insightful Thesis Advisory Panel. I did not necessarily always appreciate their thoroughness upon exiting a three hour grilling, but am aware my work is all the better for it. Likewise, I would like to offer sincere thanks to Calvin Dytham and Adam Kleczkowski for seeing the value in my work and providing a productive and enjoyable viva, and helpful and constructive feedback.

I would like to say a special thank you to Julie Knox for whom I have caused unspeakable quantities of paperwork over the years. Despite this, she always manages a warm smile and a cheery hello that never fail to brighten my day and have made a number of difficult circumstances far easier to face.

My thanks and a huge debt of gratitude go out to my Mum and my late Father for their continued encouragement and support throughout my education and life as a whole. Thank you for always being behind me (even though you do not have a clue what I do), and Mum, I hope you can now accept that I was correct when I said GCSE Religious Studies really was not important for my educational future.

A huge amount of thanks, love, hugs and fist bumps go to The Wests (simply the best) and The Hollands (hooroo!). Between you, you have helped, supported, encouraged, fed and housed me and, on one particularly special occasion, tended to my wounds (the magic of that moment will live with me forever, Chris). Chris and Emma have also been immensely helpful in assisting me to escape my perpetual studenthood and enter this weird and bizarre world known as the work place...

A very, very big thank you goes to Sarah Watling for everything she has done. It has not always been easy, but she has been behind me for much of my time working on this thesis, during which she has offered immeasurable help and support for which I cannot thank her enough. Not to mention the fact that during the course of my studies she gave me the greatest gift in the world (man I love my N64...).

Last, but by no means least, I would like to thank my beautiful daughter Isla. You have given me so much happiness and made me laugh and smile so many times that I cannot begin to quantify how much your perfectly round cheeks have helped me through the hard times. During the difficult final push to submit, when things sometimes seemed hopeless, it was a desire for you to look up to me and be proud of your Dad that provided the motivation I needed to finally succeed in completing this work. When you are old enough to understand this, I hope I will have succeeded in making you even a fraction as proud of me as I am of you. And if you do not behave, I will make you read it!

To everyone else I should thank but have not: thank you.

Declaration

This thesis is composed of original research conducted by myself, under the supervision of Dr Jon W Pitchford and Dr Angela Hodge. All information sources have been acknowledged. No part of this thesis has been submitted for a degree at any other university.

I performed all literature research, programming, analysis, and writing of the thesis chapters. The following chapter of the thesis has been published in a peer-review journal. I am the first author on the paper, but have received feedback and corrections from my co-authors.

Chapter 3 has been published as:

Croft SA, Hodge A, Pitchford JW (2012) Optimal root proliferation strategies: the roles of nutrient heterogeneity, competition and mycorrhizal networks. Plant and Soil 351:191206.

Chapter 1

Introduction

At the broad scale, plants can be thought of as constituting two parts: above and below ground. For obvious reasons, the above ground part of a plant's structure can be much more easily observed and measured, and consequently understood. Above ground, from the initial days as a shoot growing from a seedling through to a fully formed system of stemming branches adorned with leaves, buds or flowers, growth has evolved to allow the plant to best facilitate a number of functions, such as capturing light for photosynthesis or attracting pollinating insects. Survival for a plant means succeeding in adequately performing all of these functions to allow reproduction to take place.

Below ground, a plant's root system similarly performs a number of equally important tasks to enable the plant to grow and survive. The main functions of the root system are anchorage, nutrient acquisition and water uptake (Hodge et al. 2009). How an individual plants goes about performing these functions varies from species to species (Malamy 2005), and depends on both the environmental conditions in which the species have evolved to succeed and those which the individual plant finds itself growing within (Malamy 2005).

The focus of this thesis is on the below ground elements of plant growth, specifically the effects of stochasticity in environmental conditions and competition on the performance of an individual or population of plants. The aim is to further understanding of how their behaviour can maximise performance and, in turn, how different behaviour affects others in the environment. A better understanding of how resource distribution affects plant growth, at both the individual and population scale, could help to minimise the cost (financially and environmentally) incurred from excessive use of fertilisers (Cropper and Comerford 2004).

1.1 Root function

Observations of below ground activity of an individual or a plant community can be performed in many ways (Cahill and McNickle 2011), such as excavating root systems, the use of transparent growth chambers, or tracking the movement of tracer elements/compounds; the most appropriate method obviously depending on what exactly is being investigated. Whilst X-ray Computed Tomography (CT) scanning allows for the tracking and visualisation of root structures within the soil, and soil movement using isotopes, there are limitations to the accuracy and applicability of such approaches (Garbout et al. 2011). Although new techniques and approaches are leading to improvements in this area (for example, Mairhofer et al. 2012, 2013), consequently it is difficult to perform below ground observations in non-destructive ways without having to grow the plants in idealised lab conditions, and often it is plausible to only study one or two plants at a time. This is one of the key areas where theoretical models can help in understanding the growth and behaviour of plants.

The individual roots in a plant's root system directly affect a narrow volume around their surface, the rhizosphere. By secreting compounds the roots modify the immediate environment, stimulating microbial activity to aid nutrient capture, as well as helping to protect the roots from disease and drying out (Paterson 2003; Paterson et al. 1999; Hodge et al. 1998b; Kuzyakov 2002). Over a broader spatial range, the root system indirectly affects the environment through diffusion due to water uptake (Tinker and Nye 2000). The different nutrients required by a plant can have markedly different diffusion rates, varying by a number of orders of magnitude (Tinker and Nye, 2000). Consequently, optimal proliferation for different nutrient types requires different strategies. Nutrients with relatively high diffusion rates (such as NO_3^- ions) can cause issues of self competition amongst closely located roots (Casper and Jackson 1997) for example; conversely, relatively immobile nutrients (such as phosphate ions) experience little diffusion over time frames relavent to root growth (Tinker and Nye 2000). This immobility makes obtaining phosphate more difficult, and consequently many plants have evolved to form symbiotic relationships with mycorrhizal fungi, the most common form being that of the arbuscular mycorrhizal (AM) which can form on circa two-thirds of all land plant species (Smith and Read 2008). In exchange for carbon from the host plant, the AM fungal partner enhances the plant's nutrient capture by exploring a larger volume of soil and acquiring phosphurus (P) from beyond the phosphate depletion zone that quickly builds up around the plant root's surface (Smith and Read 2008; Tinker and Nye 1973).

1.2 Root system architecture

It is a challenge in itself to describe an existing root system, let alone model the growth and development of a system dynamically. In terms of root architecture there are two key elements: the shape of the system and its structure (Hodge et al. 2009). The shape of the root system refers to the spatial distribution of the roots, and can be represented by measurements such as root length densities (length of root per unit volume of soil) or the depth of soil in which roots are found. The structure describes the components of the system, and can be expressed by considering the root system as a collection of nodes (branching points) and links (root segments between nodes) to form different classifications (Fitter et al. 1991). Both shape and structure provide valuable information in terms of describing the physical properties of the root system, but both are needed to fully represent its architecture (two different root systems can have similar shape but different structure, or vise versa).

Structure classifications generally fall into one of two categories: developmental systems and topological systems (Fitter 1987; Fitter et al. 1991; Fitter and Strickland 1991). Developmental systems deal with the order of roots (starting with first order, and subsequent subsidiary roots taking higher order values) and how the structure grows and develops over time. Problems with this form of classification arise when trying to describe existing structures rather than developing ones, due to the need for decisions to be made over the ordering of the roots at existing branching points retrospectively.

Topological systems use a reverse ordering to describe an existing structure (exterior links are of first order, with interior links defined by the number of exterior links they supply). From this labelling system, topological parameters such as altitude (maximum path length to an external link) and external path length (sum of all path lengths to external links) can be defined (Fitter 1987). Due to the ordering employed, these models are not suited to growing/developing systems since they would require the re-ordering of existing roots as the system evolves over time.

Different species display markedly different root architectures (Hodge et al. 2009; Taub et al. 1996), but individual plants also see changes to their architecture through their life time. Taub et al. (1996) utilised the topological classification of Fitter (1987) to quantify the variations in root systems from numerous species when grown under different conditions. Classification within the topological model changed as a result of the plants altering their root system structure over time, reflecting shifts towards a more herringbone like structure under reduced resources.

Whilst root length density, branching angles, root depth and topological classifications can say much about the structure and spatial distribution of a root system, there are many other features that are important in attempting to fully define a root system. As stated, the root system as a whole performs a numbers of functions, and these functions are served as a combination of the functions of the individual roots. Different roots within a plant's root system perform different functions (Waisel and Eshel 1992; Hishi 2007). For example, it has been shown that as little as 30% of a plant's root system can be involved in water uptake, and only 10% involved in nitrate uptake, at a given time (Robinson 1991; Hodge et al. 2009).

1.3 Root responses to the environment

It is well established that localized root proliferation in nutrient-rich patches occurs in response to heterogeneity in resource supplies (reviewed by Hodge 2004, 2009). Aided by the modular nature of their root systems, plants can exhibit a significant degree of architectural flexibility in root deployment (de Kroon et al. 2009; Malamy 2005; Hodge et al. 2009).

Whilst such behavioural responses can seem quite logical, evaluating exactly why plants respond in the way they do to environmental cues is not straight forward. The marginal value theorem (MVT), more common in animal literature than that for plants, has been highly influential in the study of optimal foraging strategies (Charnov 1976). The MVT predicts that optimal foraging involves organisms investing foraging effort proportional to the quality of an encountered resource patch. When applied to plants, there is evidence that these predictions are true for root proliferation strategies (for example Gleeson and Fry 1997; Gersani et al. 1998), and there are arguments that optimality is important in assessing root foraging behaviour (McNickle et al. 2009)

However, observed results of the proliferation responses of different plants in varying scenarios suggest that notions of optimality are perhaps too simple. For example, it has been demonstrated that root proliferation is of little net benefit to either single plants or plants grown in monoculture as a means to acquire nitrogen (N) from N-rich zones (Hodge et al. 1998a; Fransen et al. 1998; van Vuuren et al. 1996). However, when plants are grown in interspecific competition for N from a complex organic patch (mixed N-sources) then root proliferation can confer a competitive advantage (Hodge et al. 1999a; Robinson et al. 1999). In short, the response itself to the available resources does not reflect the importance of the reaction in terms of nutrient capture, but rather the importance of the reaction (and subsequently obtained resources) are highly context sensitive (Hodge et al. 2009).

The presence/absence of competing plants obviously forms a part of this context. As well as plastic responses to resource heterogeneity, plants are also known to alter their behaviour in response to the presence of competitors (Cahill et al. 2010, Hodge 2009, Gersani et al. 1998). A commonly observed response is root segregation in order to avoid direct competition with the roots of other plants (reviewed by Schenk et al. 1999).

But just as root responses to resource heterogeneity are context sensitive, so too are the responses to competitors. For example, Cahill et al. (2010) observed that when *Abutilon theophrasti* seedlings are grown in isolation, they deploy a broad foraging strategy regardless of the level of resource heterogenity. When grown in the presence of competition, a more focussed foraging strategy was observed, which in turn was influenced by the resource distribution. Again, responses cannot simply be attributed to one factor or another, with an emerging hierarchical set of factors combining to determine the observed responses.

Responses to competitors also appear to be dependent on the identity of the competition. There is evidence to suggest that plants are able to recognise “kin”, and alter behaviour so as to reduce direct competition with sibling plants compared to non-related competitors (Dudley and File 2007; Bhatt et al. 2011; Murphy and Dudley 2009), though it has been observed that not all plants respond in the same way to self/non-self competition (see Hess and de Kroon 2007; Schenk 2006; Masclaux et al. 2010; Milla et al. 2009). Whilst such behaviour has evolutinary merit, it is again context sensitive and its effect on the outcome of competition compared to other factors could be limited (Milla et al. 2009)

The different contexts in which a plant can find itself growing can of course be hugely variable: soil impedance (affected by water content and soil type), resource abundance and distribution, competitor proximity and numbers, and a whole host of other factors will all affect what it means for a root system’s architecture and proliferation responses to be optimal.

Even when optimal behaviour can be identified and/or quantified, it does not necessarily mean that it is the preferred behaviour for an organism to display. Care must be taken when talking about optimality, as one has to be careful about what it actually means to be “optimal”; when talking about success and survival in an evolutionary context (which ultimately is what has shaped the form and behaviour of all organisms), it is essential to assess performance against evolutionarily relevant metrics rather than simply assess mean-field properties (Currey et al. 2007; James et al. 2010; Preston et al. 2010).

As well as changes to root architecture, varying population densities (and with it varying levels of competition) affect size hierarchies (distibution of inividual sizes amongst the population) and community make-up (Wijesinghe et al. 2005). In even-aged monocultures, as the population grows there is a tendency for a size hierarchy to quickly emerge comprising a large quantity of smaller individuals and just a few larger individuals (Hara, 1988; Purves and Law, 2002). As a plant grows, on the local scale it continually influences, and is influenced by, the growth of its neighbours (Schneider et al. 2006). Spatial population structure is highly localized, and mean-field dynamics of the population do not qualitatively capture information about local neighbourhood interactions and dynamics (Law and Dieckmann, 2000). Competition is both local (Schneider et al. 2006; Lv et al. 2008) and, on a paired-plant basis, strongly assymetric (Schwinning & Weiner, 1998; Schneider et al. 2006; Lv et al. 2008). Whilst summing these individual, local, pair-wise interactions can provide a model of population-level dynamics that successfully explain some of the observed variation in size-hierarchies, further work is needed to validate how accurate such methods are (Schneider et al. 2006).

Experiments performed to assess the compound effects of resource and neighbourhood heterogeneity are somewhat limited by the experimental constraints mentioned earlier. For example, a series of studies (Day et al. 2003; Wijesinghe et al. 2005; Hutchings and Wijesinghe, 2008) investigated the impact of the spatial pattern of nutrient supply on yield and community structure. In each of these studies, plants were grown in “chequerboard” environments, with squares of different sizes and nutrient concentrations. Such an approach allows for areas of varying nutrient concentrations to be created and observed. However, such organised, strictly defined resource heterogeneity is some way from the more random heterogeneity often experienced in nature (Fitter 1994).

In experiments carried out with *Cardamini hirsuta* (Day et al. 2003), it was observed that overall yield among populations grown with the same total levels of nutrient supply were similar irrespective of the distribution of nutrient supply, with total yield increasing with an increased total supply of nutrient. Whilst yield was unaffected by resource distribution, populations grown in heterogeneous populations responded by concentrating biomass into areas of higher nutrient concentration, though this allocation of biomass was not affected by patch scale. It was also observed that the mortality rate was greater under homogeneous conditions than heterogeneous conditions at all nutrient levels and patch scales, suggesting that nutrient heterogeneity can be an important component of plant survival.

When the population was comprised of a mixture of 20 herbaceous annual and perennial species (Wijesinghe et al. 2005), it was observed that whilst the pattern of nutrient distribution had few significant effects on a particular species or group of species, community biomass and elements of population composition were affected significantly. At greater concentration differences between patches, an increase in root biomass was observed, demonstrating that competing plants obtaining resources from limited sources increase root growth relative to those growing in homogeneous environments. Whilst different species demonstrated notably different responses to nutrient supply (for example *Dactylis glomerata* and *Galium verum* both saw significant increases in biomass with decreasing nutrient patch density, whereas *Rumex crispus* saw a corresponding decrease in abundance), the pattern of nutrient supply had little effect on species richness and diversity. In conclusion, this suggests the relationship between diversity and heterogeneity is not necessarily positive. Highlighting the inherent difficulties in performing observations of such experiments, it is worth noting that the separation of the roots of individual species was not possible when assessing biomass allocation within the environment.

When the impact of environmental heterogenity on the yield of the clonal plant

Glechoma hederacea was investigated (Hutchings and Wijesinghe, 2008), the yield at all scales was highly dependent on environmental context, with maximum yield attained under a limited range of heterogeneous conditions. Even though overall nutrient levels remained constant, the effects of nutrient heterogeneity on local and popolation-level yield was highly context sensitive, with yield strongly affected by the scale and relative concentrations of the patches. Yield significantly exceeded or fell short of that in the homogeneous treatments depending on these patch properties, with yield differing greatly between treatments with common patch contrast but different patch scales. The mean yield in high and low quality patches of the large scale heterogeneous treatments were observed to always be significantly higher or lower, respectively, than that achieved in the equivalent parts of the homogeneous treatments. However, at the greatest concentration contrasts between patches, overall yield in the heterogeneous treatments was significantly greater than in the homogeneous treatments, with the greatest net yield (compared to homogeneous conditions, and overall) realized when patches were large and with the greatest contrast between patches.

Despite the limitations imposed by the practicality of performing such experiments, the results of these studies (Day et al. 2003; Wijesinghe et al. 2005; Hutchings and Wijesinghe, 2008) help to cement the notion that responses to variability in resource supply are highly context sensitive. Again, a hierarchical response pattern is observed, with resource supply closely linked to competitor type and presence in determining the overall effects on individuals within a population. Also, they highlight that heterogeneity in resource supply (and the effects this has on plant growth) is not as simple as the existance of areas of presence and absence of a resourse, but that resource heterogeneity is also qualitatively and spatial-scale dependent.

Fitter (1994) defined a set of basic attributes to describe a patch. These key attributes fall into two scale categories: spatial and temporal. Within the spatial and temporal categories, there are sub-attributes for distribution (pattern and predictability, respectively), extent (size and duration, respectively) and number (abundance and frequency, respectively). These patch properties, and the corresponding effects they have on plant growth and competition, remain understudied.

With such a broad range of function and response, the modelling of plant root growth and competition can be something of a daunting task. Whilst partly defined by genetic information, in contrast to many modelled animals and organisms, plants do not have predefined, known structures, and instead determine their own structure during growth and development (Tardieu, 2010).

1.4 Modelling root systems

There are a numerous approaches to modelling root growth, each with their own relative strengths and weaknesses (reviewed by Dupuy et al. 2010). All modelling requires some level of compromise, and demands varying degrees of simplification or assumption. As with all forms of modelling it is important when making decisions about how to condense the complexity of the real world into a managable model that these assumptions are carefully considered. If not, it is quite plausible, and to an extent, inevitable, that a model will reflect nothing more than the assumptions made during its construction.

A relatively common simplified approach to the modelling of root growth and competition is to consider a root system's "zone of influence" (see for example Berger et al. 2002; Casper et al. 2003; O'Brien et al. 2007), the area or volume over which an individual alters its environment. Such an approach eschews explicit modelling of the root architecture and represents the root system as an area or volume (depending on the number of dimensions being modelled) exhibiting certain charactistics such as levels of resource uptake. Interaction and overlap between root systems are typically controlled by predefined rules (for example Berger et al. 2002), rather than root growth being determined by explicit processes and/or interactions.

An example of a different approach to the zone of influence methodology is taken by O'Brien et al. (2007). In their two-dimensional model, distance-dependent root growth costs and resource uptake proportional to resource levels results in a passive (from the plant's point of view) growth process where overlap and resource levels determine the outcome of competition and realised growth. Whilst this approach yields information about how the presence of local competition can alter a plant's realised performance in relation to its potential performance (in the absense of competition), its passive nature excludes a large part of a plant's responses and behaviour which are implicit in such occurances in reality (Hodge 2004, 2009; Cahill and McNickle 2011).

In contrast to these highly simplified representations of root system are architecturally explicit models (Dupuy et al. 2010; Fitter et al. 1991; Jakobsen and Dexter 1987; Jourdan and Rey 1997; see Vos 2009 for a review of structural plant modelling). The development of architecturally explicit models has been aided significantly in improvements in available computational power. While Fitter (1987) calculated nutrient uptake and movement by scanning the screen of an Acorn BBC microcomputer and counting coloured pixels, in more recent years advances in computational resources has led to increasingly more complex and sophisticated models (see for example Jourdan and Rey 1997).

An important facet of modelling the function (in terms of water and nutrient uptake) of root models is one of feedback loops between the root system and soil environment (Pierret et al. 2007, Leitner et al. 2010; see Bever 2003 for a review of conceptual and theoretical work on soil community feedback). This importance stems from the fact that whilst the physical properties of the soil can affect the plant's root structure and architecture (Hodge 2004; Fitter and Stickland 1991), similarly there is the effect of the plant's behaviour on the soil (Pierret et al. 2007, Leitner et al. 2010).

For an individual root, the uptake of water and nutrients are often modelled and described by partial differential equations. Two important examples of such equations are the Richards equation for the movement of water in unsaturated soils (Jury and Horton 2004) and the convection diffusion equation for solute transfer in soil (Barber 1995), which can be used to model the plant-soil dynamics of water and nutrient uptake/movement (Leitner et al. 2010). These partial differential equations can be solved numerically by using the finite element method (FEM) (Doussan et al. 2006; Javaux et al. 2008). An individual root element is often considered as a cylindrical volume of specified length and radius (Roose and Fowler 2004, Roose and Schnepf 2008), over which transport equations for solute and water movement can be solved to derive uptake along the root's length.

Explicit architectural models are often constructed by calculating uptake for an individual root segment and scaling this behaviour up to the whole-plant level (Roose and Schnepf 2008, Leitner et al. 2010). Due to the computational complexity of modelling water and solute uptake across an entire root system's architecture, and the resulting water and solute transport within the surrounding environment, it is often necessary to reduce the problem to a simpler form. This is often achieved by replacing a fully explicit model of resource uptake across a continuous root system with a sink term to describe uptake over the entire structure (Roose and Fowler 2004, Roose and Schnepf 2008). This introduction of such a sink term in effect scales uptake from the individual root to the root system as a whole, whilst simultaneously reducing the complexity from the root system as a whole to an individual point of acquisition. Whilst a number of different techniques exist to calculate this sink term (Roose and Schnepf 2008), these methods typically rely on simplifying assumptions about root structure (such as an even distribution of roots, or functional equivalence across all roots) and the environment (homogenious soil) (Roose and Schnepf 2008, Leitner et al. 2010). Such upscaling often results in a number of issues, not least the absence (or at least masking) of the effects of both intra- and inter-plant root-root competition for resources (Tardieu et al. 1992; Dupuy et al. 2010). As a result of these assumption, it follows that sink terms need validation on an experimental basis (Leitner et al. 2010).

Despite the continual increases in computational power, complex architectural models are still difficult to scale up to population levels (Dupuy et al. 2010). As such, whilst they can provide insight into certain aspects of plant growth, they can often not be applied to competitive scenarios, and accordingly provide direct information and insight that is highly constrained to specific contexts.

However, the many factors influencing plant growth and root architecture discussed here also give rise to one of the biggest issues with complex architecturally explicit models: parametrisation. Root growth models containing explicit root architecture require a large number of parameters to define growth, development and function (Dupuy et al. 2010). Measurements and observations of plant roots are not only difficult (Cahill and McNickle 2011) but, as has been stated a number of times within this discussion, highly context dependent (Hodge et al. 1999; Day et al. 2003; Wijesinghe et al. 2005; Hutchings and Wijesinghe 2008). The big question this raises is: how does one parameterise a model and evaluate its behaviour in different environmental conditions when the parameters themselves are so context sensitive? Not only is there a large degree of variability in observed behaviour dependent on the conditions a plant is grown in; there is a large degree of variability in growth and architecture in control conditions (Forde 2009). This is a significant obstacle to such models, and will likely persist to be so for some time to come.

It is possible for seemingly complex and realistic root system representations to be generated and visualised by relatively simple models and methods. For example, the fractal nature of plants' root and branch systems have long been observed and appreciated (Prusinkiewicz and Lindenmeyer, 1990), with simple algorithms able to generate apparently complex structures. These algorithms can also be coupled with environmental feedback (Mech and Prusinkiewicz, 1996), however this requires the specification of rules of interaction for the interface between the plant simulation and the environment model.

An alternative is a density based approach, where the density of a given property or group of properties is used to describe a given structure. For example, Dupuy et al. (2005) use branching density (the number of branching points per soil volume and angular inclination range) as the characteristic measurement, and use a model to generate a root structure with a branching density as close as possible to a specified value. Although using just a single general root branching density function was sufficient for the simulated root systems generated to show strong similarity to the real observed root systems, it is acknowledged that this is only possible for a static geometry. When linking such a density based structure to resource uptake (Leitner et al. 2010), the problems described previously regarding sink term validation still apply. A dynamic model (i.e. a growing, developing root system) would require much more consideration and parametrisation still.

It is generally excepted that the modelling of plants lags some way behind that of other ecological areas (Tardieu, 2010). Given this fact, and the level of complexity evident in the observations of plant growth and competition, it is not surprising that increasingly complicated and detailed models continue to be created. Whilst there appears to be a push for this trend to continue (Vos et al. 2010), it can be argued that the inherent difficulties and issues associated with parametrisation and defining behaviour in such models leaves plenty of scope for simplified models. As has been shown in other areas of research (for example fish recruitment and foraging strategies, see Brindley and Pitchford 2001; Brindley et al. 2003, 2005), models with minimal assumptions about growth and functionality can allow for effects of different factors to be investigated at the simplist level, and minimise concerns about complexity in one area clouding simplicity in another, or vise-versa. Such simplified models can potentially possess much greater flexability when it comes to scaling and applications to variable conditions than their more complicated counterparts.

1.5 Thesis overview

In this thesis an idealised one-dimensional model of plant growth and competition in heterogeneous environments is developed. By avoiding all non-essential features, minimal assumptions are required. With a focus on qualitative rather than quantative results, the need for “real world” parametrisation is avoided, and accordingly issues of context sensitive measurements and observations averted. By reducing the problems of growth and competition to their simplest conceptual form, behaviour can be observed and understood, as far as possible, at each stage before progressively adding more complexity.

In Chapter 2, the basic one-dimensional model is introduced and explained. Starting with a Gompertz growth curve, an environment consisting of a patchy resource and a neighbourhood of competitors is defined. Simulations are run for varying populations sizes and patch numbers under different combinations of plant and patch distributins, and for two different behaviours at neighbour-neighbour boundaries. A further simplified, linear growth model is defined and tested against the Gompertz growth model. It is shown that the linear model closely approximates the Gompertz model both qualitatively and quantitatively at the population and individual scale for nearly all scenarios.

In Chapter 3, the linear model from Chapter 2 is tested in environments with both statsitically uniform and aggregated patch distributions. The plants are equipped with a simple proliferation paramater that allows them to bias growth in the direction of the most recently found patch. Running the model within a genetic algorithm, context sensitive sensitive results emerge demonstrating a hierarchical response to environmental factors consistent with the literature.

In Chapter 4, the model is expanded into, and adapted for, two-dimensions. Initial tests are performed in analagous conditions to those in Chapter 3 to verify if the model demonstrates the same qualitative behaviour in two-dimensions as in one-dimension. Although subtle differences emerge in the results, similar trends are observed in the two-dimensional model as to for the one-dimensional model in Chapter 3.

In Chapter 5, different growth strategies are tested within the two dimensional model. Four different plant types displaying different growth rates, initial maximum size limits and root efficiencies are tested individually, in monocultures and in mixed competition to assess how different environmental factors can affect growth and competition among plants with different behaviours and properties.

The thesis concludes with a summary of the work and results presented, and discusses possible applications and developments for future research.

Chapter 2

Capturing the complexity: a one-dimensional model of plant growth and nutrient acquisition

2.1 Introduction

One of the key components of plant growth is the proliferation of roots into the surrounding soil to facilitate the capture of available nutrients (Taub and Goldberg 1996; Williamson et al. 2001; Hodge 2009). In the presence of competition between two or more plants, a race ensues to best capitalise on the available resources and hence attain an advantage over the competiton (Robinson et al. 1999). As such, the state of both the environment (Day et al. 2003; Wijesinghe et al. 2005; Hutchings and Wijesinghe 2008) (here taken to mean the spatial distribution of available resources) and the neighbourhood (Purves and Law 2002; Schneider et al. 2006; Lv et al. 2008) (the relative location of individual plants), and an individual plant's ability to react to these conditions (Robinson et al. 1999), have important and significant impacts on the outcome of competition.

Observing and understanding plants' reactions to, and the impact of, variations in the environmental and neighbourhood conditions requires a detailed knowledge of the conditions being experienced. An accurate understanding of both the physical make-up of the soil and knowledge of individual plants' positions in relation to both neighbouring plants and these physical features often requires heavily idealised experimental set ups (for example, Hodge et al. 1999a).

By creating a theoretical framework in which to investigate such scenarios, the environmental and neighbourhood conditions can be described exactly, either deterministically or in a rigorous statistical sense. As a result, the impact of different combinations of conditions along with that of different evolutionary strategies can be observed in models and, in turn, used to improve our understanding of the real biological systems.

In order to facilitate this, for the purposes of this chapter all other facets of plant growth and competition (see Chapter 1) are excluded, and the focus placed purely on the underground growth and interactions.

2.2 Modelling growth and competition in heterogeneous environments

2.2.1 Plant growth

Reducing the modelling of a plant's growth to its simplest form means representing a plant's size, L, as a function of time, with "size" usually taken to mean either length or dry weight (Erickson 1976). There are a number of different functions that perform this task depending on the stage of growth, part of the plant, or plant species.

The simplest such equation is a linear differential equation (Eqn. 2.1), with growth at a constant rate g. The growth of root segments (measured as length) have been observed to grow linearly over a substantial period of time (Kraus 1895).

$$\frac{dL}{dt} = g. \tag{2.1}$$

On a whole plant scale, a number of plants species have been measured (in terms of length and dry weight) to grow with expontential growth for considerable periods of time (Erickson 1948; Purves and Law 2002). In this instance, the rate of growth of an individual is proportional to its size; this proportion defined by a relative growth rate, g (Eqn. 2.2). That is, as a plant gets larger, in turn its rate of growth increases. This was famously proposed by Blackman (1919) as the “compound interest law of plant growth”.

$$\frac{dL}{dt} = gL. \tag{2.2}$$

Whilst linear and exponential growth have been observed to fit experimental data over certain time periods, it stands that no plant (or any organism for that matter) can keep growing with linear growth, and even less so exponential growth, without some sort of cap or interuption to this pattern.

Over more extensive periods of a plant's life, behaviour is more often consistent with a growth function that is sigmoid in form (Erickson 1976). An example of a proposed sigmoid function for plant growth is the “autocatalytic equation” (Erickson 1976), or logistic equation. This assumes that growth of an individual is proportional to both its current size (as with exponential growth) and also the difference between its current size and a specified upper limit, $L_{\max}$ (see Eqn 2.3).

$$\frac{dL}{dt} = gL(L_{\max} - L). \tag{2.3}$$

In this form, growth is initially exponential-like until the plant reaches size $L_{\max}/2$, at which point (the “inflexion point”) the relative growth rate begins to decrease as the plant gets bigger. The curve is symmetrical about the mid-point $L_{\max}/2$ and displays assymptotic behaviour at the lower and upper bounds (i.e. 0 and $L_{\max}/2$). However, observed plant growth often is not symmetrical around this point (see below), and so in these instances another kind of sigmoid function is required.

An example of a sigmoid growth curve that has been shown to fit experimental data for plant growth with this non-symmetry (Purves and Law 2002; Schneider et al. 2006; Lv et al. 2008) is the Gompertz growth curve (Eqn. 2.4).

$$\frac{dL}{dt} = gL\ln(\frac{L_{\max}}{L}). \tag{2.4}$$

Integrating equation 2.4, it follows that

$$L(t) = L_{\max}e^{-L(0)e^{-gt}}, \tag{2.5}$$

where $L_{\max}$ is the upper size limit as before, g is an intrinsic growth peramater and $L(0)$ is the initial size. The Gompertz function has been coupled with competition kernels to define a model for growth amongst competitors (Purves and Law 2002; Schneider et al. 2006), with an extension of this approach adding stochasticity into the framework (Lv et al. 2008) to account for observed variability.

There is no “one size fits all” model to describe plant growth, not least because different types of plant exhibit different types of growth (Taub et al. 1996). But also, as discussed in Chapter 1, “growth” is highly context sensitive. When plants grown in isolation as control tests (in idealised, homogeneous environments) demonstrate significant levels of stochasticity in their growth (Lv et al. 2008; Forde 2009), then plants growing in situ with all of the added factors and sources of stochasticity that it brings are inevitably going to demonstrate varied growth behaviour.

However, Purves and Law (2002), Schneider et al. (2006) and, in particular, Lv et al. (2008) demonstrate that, regardless of the stochasticity present from environmental and competition induced factors, the growth of *Arabidopsis thaliana* (one of the most widely studied plants) remains qualitatively Gompertz in nature. The Gompertz growth formula thus is a good selection as a base to build a highly idealised model upon: it gives a foundation for defining growth in a way that inherently includes processes that occur in plant growth without a need for explicity modelling them, such as growth due to uptake of background nutrient resources and loss due to metabolism (Purves and Law 2002). Furthermore, as complexity is built into the model, it can be retained as a reference point to check that “growth”, in its broadest sense, is adhering to observed experimental behaviour (i.e. it remains Gompertz-like).

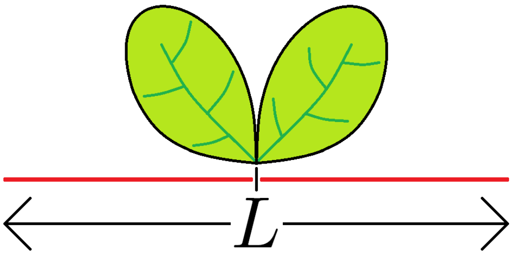

Figure 2.1: *Visualistation of basic plant growth and measurement of size. Growth is modelled purely at the root system (shown red in figure). Plant size, L, is taken as the measurement of the length of the extent of the root system (i.e. the sum of the size of the root structure growth in both directions from the plant's centre). Root growth will remain symmetrical around the plant's centre as shown unless neighbourhood conditions interfere (see Sections 2.2.4 and 2.2.5).*

2.2.2 Gompertz growth models

The Gompertz growth curve is applied to an idealised one-dimensional model in which a specified number of plants growing within a given environment compete for available nutrient resources. The environment is defined as a one dimensional length with periodic boundary conditions (i.e. the environment can be thought of as circular), with nutrient resources defined as simple patches (of point physical-size and uniform nutrient content) placed in the environment. In this model, a plant's roots grow in both directions (subect to neighbourhood conditions; see Sections 2.2.4 and 2.2.5) within the one-dimensional environment, and the size of the plant, L, is the total length of the roots (i.e. the sum of the growth in both directions; see Fig. 2.1). The "nutrient content", or quality, of a patch is defined by the marginal benefit to a plant from its acquisition, (p). Varying the number of plants, their physical distribution, and both the abundance and distribution of available resources allows the impact of different spatial distributions of plants, resources, and the joint impact of the two, on growth at the individual and population levels to be observed.

In the Gompertz model it is assumed that the size L of an individual at time t

is given by Eqn. 2.5, where g is an intrinsic growth parameter, $L(0)$ determines the initial size of an individual, and L_{max} is the maximum size limit (i.e. lim $L(t)$ as $t \to \infty$). As seen in Figure 2.2 and discussed in section 2.2.1, the Gompertz growth function forms a sigmoid curve, with growth initially exponential before tailing off, and non-symmetry around the inflection point. The growth parameter (g) determines the steepness of the curve, initial size ($L(0)$) the interception point of the axis, and the maximum size (L_{max}) the asymptotic limit.

2.2.3 Gompertz models with patch-dependent growth

In order to accommodate the influence of resource acquisition on plant growth, within this model it is assumed that L_{max} is not constant, but rather that its value increases with the acquisition of each resource patch by a value equal to the patch quality (in terms of the marginal benefit to the plant), p. At time t, the upper size limit L_{max} can be expressed as

$$L_{\text{max}}(t) = L_0 + N(t)p, \tag{2.6}$$

where $N(t)$ is the number of patches acquired by time t. Consequently the function in Eq 2.5 is only applicable for an individual before it encounters any patches. As patches are acquired, the formula for the size of $L(t)$ takes the form of a more complicated stochastic function incorporating plant and patch distribution; the size of an individual at a specified time depends not only on the number of patches it has acquired, but also importantly when these patches were acquired.

The rate of growth at a given time depends on current size and the number of patches acquired, and not explicitly when. It follows from equations 2.5 and 2.6 that the rate of growth at time t can be expressed as

(a)

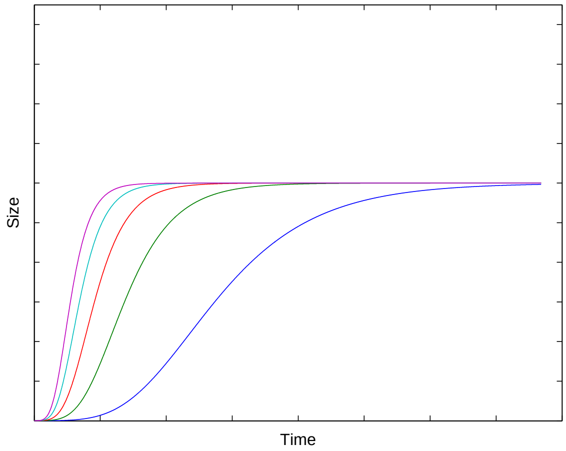

(b)

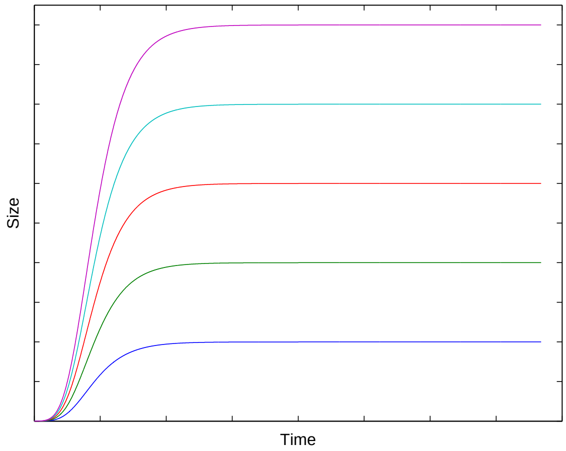

Figure 2.2: *Examples of Gompertz growth curves for different values of (a) growth parameter,* g*, and (b) maximum size limit,* L_{max}*. Initial size (as determined by* $L(0)$*) is fixed. Larger values of* g *result in quicker initial growth causing a steeper curve, whilst* L_{max} *defines the asymptotic growth limit, with larger values resulting in a larger final size.*

$$\frac{dL}{dt} = rL\ln\left(\frac{L_{\max}(t)}{L}\right), \tag{2.7}$$

$$= rL\ln\left(\frac{L_0 + N(t)p}{L}\right). \tag{2.8}$$

However, this is not an easily solvable differential equation since $N(t)$ depends on L, with L a stochastic variable.

It is important to check that the observed behaviour with the amended Gompertz growth function including patch acquisition (Eqn. 2.8) is still "Gompertz" in its nature. The previously mentioned examples of deterministic (Purves and Law, 2002; Schneider et al. 2006) and stochastic (Lv et al. 2008) Gompertz growth modelling have not explicitly modelled the actual functions and interactions that cause the observed behaviour, but rather they have been applied to observed behaviour both during and after the effects of neighbouring plants and environmental variability have taken effect. This means it is important to check that the fundamental Gompertz behaviour is not being radically altered by the amendment for patch acquisition. A number of tests were performed to investigate this, summarised in Figure 2.3. It shows a plot of three different growth curves: the blue line is an initial parameterisation of the Gompertz growth function; the red line takes the same parameterisation and adds growth "jumps" due to patch acquisition by increasing the asymptotic limit $L_{\max}$ (see equation 2.6)- note that the red line is not continuous, but rather piecewise continuous between patch acquisitions (that is, the function is continuous except at the instances of patch acquisitions); finally, the dashed black line shows a Gompertz growth curve reparametrised so as to attempt to fit the patch-amended version.

As can be seen, the addition of patch acquisition (red line) to the basic Gompertz growth curve (blue line) results in behaviour that is qualitatively consistent with a reparametrised Gompertz growth function (black line). This means that whether or not a plant is acquiring patches and the resulting growth benefit, growth can consistently be described as being Gompertz growth, consistent with data-fitting models from experimental observations (Purves and Law, 2002; Schneider et al. 2006; Lv et

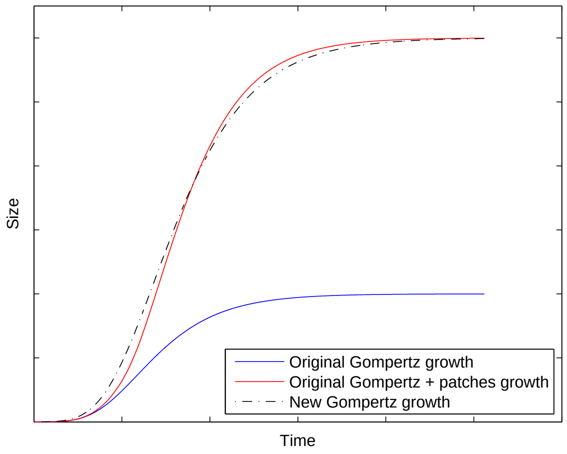

Figure 2.3: *Example of original Gompertz growth curve (blue), the same growth with the addition of patch acquisition and resulting growth benefit (red), and attempted fit of newly parameterised Gompertz growth curve to patch amended growth (dashed black line). Comparing the red and dashed black lines, it can be seen that the amended Gompertz growth with patch acquisition results in behaviour that is still "Gompertz" in nature.*

al. 2008).

Using this approach to control individual plant growth, freedom remains in how to implement boundary behaviour between neighbouring plants. In different spatial dimensions the implementation of overlap and interference in a model takes on different meanings. In the real, three-dimensional world, two plants obviously cannot occupy the same piece of physical space simultaneously. However, the roots of competing plants can, and do, overlap to create areas of interference and direct competition (Robinson et al. 1999; Wijesinghe et al. 2005). Consequently, in one dimension it is necessary to either completely deny overlap if a literal physical interpretation is to be taken, or otherwise to allow overlap if the one-dimensional environment is not considered a direct translation of real three-dimensional space. In the following models each plant grows and proliferates uniformly where permitted (see Fig. 2.4a); that is until any direct interaction occurs with a neighbour, the plants will distribute root growth equally in both directions. However, different assumptions are made as to how plants interact when their growth brings them into contact.

2.2.4 Non-overlapping model

In the first instance (the "**non-overlapping**" model), each neighbouring plant acts as an impenetrable physical barrier to its neighbours' growth. This means that plants will grow uniformly (i.e. equally in each direction) until a boundary (a neighbouring plant) is met. At this point all growth will be concentrated in the direction where space permits (Fig. 2.4b). The initial proximity of neighbouring plants, as well as the speed at which they can acquire nutrients and grow, relative to a given individual is what will ultimately dictate the final size of an individual.

2.2.5 Overlapping model

In the second case (the "**overlapping** " model), there is the potential for quite different results. Here plants can effectively "pass" through each other and uniform growth persists regardless of the presence of, and interaction with, neighbouring plants. The

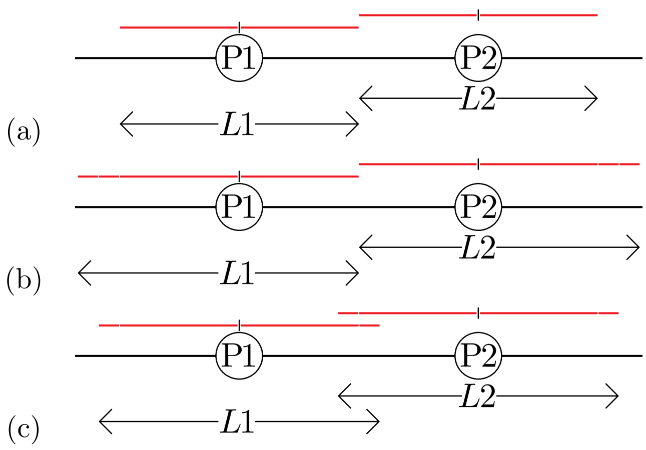

Figure 2.4: *The difference between the non-overlapping and overlapping models in appropriating growth: In scenario (a), two plants (P1 and P2) have grown (root growth shown as red lines) uniformly until they share a common boundary; (b) In the non-overlapping case, growth in the direction of the boundary ceases and instead all growth is channelled to the free space away from the boundary; (c) In the overlapping case, growth continues to be uniform, with each individual continuing to grow equally in both directions, overlapping at the common boundary. In scenarios (b) and (c), the sizes of the plants ($L1$ and $L2$) are unchanged; it is just the distribution of growth that is altered.*

interference here comes from that fact that an individual overlapping a neighbour will likely be penetrating into exhausted soil, with any nutrients occupying this area potentially already captured by other plants quicker to exploit the available resources (Fig. 2.4c). Conversely, an individual that is particularly quick at acquiring nutrients and grows relatively quickly has the possibility to “leap-frog” a neighbour and beat it to nutrients situated on the far side of said neighbour (Fig. 2.5).

2.2.6 Statistical structure of plant and patch placement

For each model, three different configurations of plant and nutrient distribution will be tested to assess the behaviour under varied environmental (patch distribution) and neighbourhood (plant distribution) conditions:

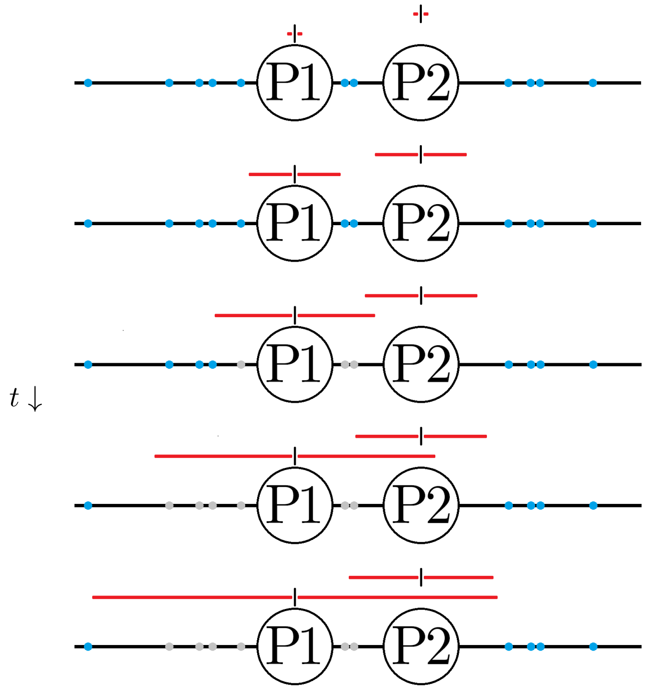

Figure 2.5: *An illustrative example of how an individual can "leapfrog" a competitor for resources situated beyond the neighbour in the overlapping model. Initially (top), both individuals (P1 and P2) have grown (root growth shown as red lines) at the same rate as no available patches (shown as blue dots) have been acquired by either plant. Through time (moving down the diagrams), Plant 1 obtains a growth advantage through acquiring patches (grey dots denoting exploited patches that have been removed) and a resulting increase in relative growth rate. By the bottom diagram, Plant 1 has grown past Plant 2 and is closer to the remaining patches to the right of Plant 2.*

- Fixed, regularly spaced plants with independently, uniformly randomly distributed nutrient patches;
- Independently, uniformly randomly distributed plants with fixed, regularly spaced nutrient patches;
- Independently, uniformly randomly distributed plants and independently, uniformly randomly distributed nutrient patches.

The first of these investigates the effect of simple stochastic variation in the distribution of nutrient patches. The second scenario investigates the impact of the variation in the neighbourhood in a patchy but regular environment. Finally, the third method shows the combined effect of stochastic variation in both the environment and neighbourhood. Since stochasticity within the model arises from variation within the plant and patch distributions, if both plants and patches are regularly distributed there will be no variation in results, and all individuals would perform identically. Consequently, the configuration with fixed, regularly placed plants and patches is omitted.

For each combination of model (non-overlapping and overlapping) and the three different configurations of plant and patch placement described above, a number of repetitions were run for different numbers of plants and patches (1000 repetitions for 2,3,4,5,6,7,8,9 and 10 plants, and 100 repetitions for 100 plants; each performed for 1,2,3,5 and 10 patches per plant). The size of the environment is scaled proportionally to the number of plants, meaning that relative space per plant remains constant across all of the experiments, and the difference in results reflects the effect of plant numbers and not population-level plant density. The patch numbers are defined on a variable patch-per-plant basis (since environmental space is fixed per-plant, this is equivalent to a specified spatial density of patches). The quality, p (in terms of nutrient content), of each individual patch is inversely proportional to the number of patches, since total nutrient content per-plant remains fixed.

2.2.7 An alternative approach: constant growth models

In the Gompertz model, patch acquisition is assumed to increase the maximum size limit of an individual ($L_{\max}$; see Eq 2.6) in the Gompertz equation (Eq 2.4) determining its growth. This has two effects: firstly, it increases the maximum achievable size (space/neighbour boundary conditions permitting); secondly, by the nature of the Gompertz equation, it increases the growth rate of the individual in the short term. That is to say acquiring nutrients both facilitates a temporary increase in the rate at which an individual grows, and also increases the potential upper size limit the plant can reach.

Whilst modelling plant growth according to a Gompertz growth equation makes sense since it has been shown to act as a good approximation of reality, there are issues to using this approach. Analytical progress with the problem is complicated by the fact that the different plants are often growing at different rates, and the size of a given individual at a given time t is defined by not just how many patches it has acquired, but also when it did so. The other problem is that stochastic simulations with large numbers of individuals are computationally intensive. With Gompertz growth demonstrating asymptotic behaviour, the later stages of growth take an increasingly long time to simulate, and it is necessary to impose a cut-off threshold to terminate each computational run. The larger this threshold, the bigger the time saving, but also the greater potential interference to results. This issue is affected by the number of plants, the number of patches, and also the distributions of both.

As a result, the approach as detailed in Section 2.2.3 has very real limitations for investigating the effect of stochastic environmental and neighbourhood conditions where repeated runs are required (especially when relatively complicated environmental descriptions are desired). With the Gompertz models this is simply not viable, and as such an alternative model is not just desirable but, realistically, essential.

Because of this, different approaches were tried and compared to see if simpler, quicker methods can yield comparative results. As an alternative to the two Gompertz models (non-overlapping and overlapping) detailed in Section 2.2.3, two different models have been created which address the issue of computation time and which also offer increased scope for progress in finding an analytical solution. These alternative models replace the variable growth rate of the Gompertz growth function with a constant intrinsic growth rate. The same initial size constraints as used for the Gompertz model are placed on each individual, and similarly patch acquisition results in the same increase in potential maximum size. However, patch acquisition now results in an instantaneous jump in size rather than the inherent increase in growth rate associated with increasing the upper limit in the Gompertz growth function.

The behaviour of the alternative constant growth model is much simpler to express. The size at a given time is defined by

$$L(t) = gt + N(t)p \tag{2.9}$$

until the upper size limit (L_{max}), is met. Provided the plant is not spatially constrained by a neighbour (non-overlapping model) its growth rate is fixed at g, except the instant when a patch has been encountered and an instantaneous jump occurs. For computational purposes, this speeds things up in two ways. Firstly, there is now no longer the asymptotic growth behaviour which arises in the Gompertz function, and instead the growth rate remains constant. Secondly, with all plants growing at a fixed, constant rate, it is possible to implement “short cuts” in the code. Since plant growth is constant and uniform across the population, it is possible to calculate the next “event” (i.e. patch encounter/acquisition or neighbour-boundary interaction) and jump forward to the time at which this occurs, adding the relevant level of growth to each individual where applicable.

The constant growth rate for all individuals makes analytical work much simpler. With a Gompertz growth function analytical progress means dealing with a set of coupled differential equations (one for each plant) whose parameters all depend on the number of patches acquired by an individual and the times at which they were acquired. With a constant intrinsic growth rate, the complexity is vastly reduced.

Simulations were run within the MATLAB computing environment to ascertain whether there were any systemic differences between the models formulated in Sections 2.2.3 and 2.2.7, and so establish if the constant growth models yields the same qualitative results as the more computationally intensive Gompertz approach.

2.3 Results

The results presented in this section are summarised to highlight the key results and findings. Plant numbers were seen to have limited effect on results for population sizes above two plants, and so results for the different scenarios are presented for when eight plants are grown in competition. Full results for all other plant numbers can be found in (Appendix A).

Results for when just two plants are growing are omitted since they demonstrate behaviour that is not consistent with larger population sizes. Under certain circumstances (especially at low patch numbers), results can be highly discretised across all models. For example, when there are just two plants with one patch per plant in the overlapping case, the final size of an individual can only take one of three fixed sizes: $L_{\max}(0)$ if no patches have been acquired, $L_{\max}(0)+p$ if one patch has been acquired, or $L_{\max}(0) + 2p$ if an individual has acquired both available patches. As such, the distribution of the final sizes achieved over multiple runs does not always provide a helpful point of reference for comparison across the different configurations of model and plant/patch numbers. For other results, the mean, variation and skew have been chosen to summarise the qualitative results (see Appendix A). These three statistical measurements provide insight into not just average performance, but also the level of inequality across the population as well as how this inequality is distributed. In turn, they reflect the trends occurring across the populations as population size, patch numbers, and plant and patch distributions are altered. It is important to capture such information as from an evolutionary perspective, statistical rarities can be the most significant results (Preston et al. 2010).

Across the different models, results are shown for three different configurations of plant and nutrient patch distributions: regularly placed plants with uniformly random patches, uniformly random plants with regularly placed patches, and uniformly random plants and patches. For each configuration, mean final plant size is presented, along with error bars showing the standard error of the mean (Rice, 2007). The standard error of the mean is calculated as the standard deviatation of the mean values across all repetitions (i.e. the standard deviation of the means), divided by the squareroot of the number of repetitions. The standard error of the mean provides the standard deviation of the sample mean's estimate of the population mean; in this case, this is the standard deviation of the simulation means' estimate of the actual mean (i.e. the mean of an infinite number of repetitions).

2.3.1 Gompertz growth models results

Non-overlapping Gompertz model

Fig 2.6 shows the mean of the final plant sizes for repetitions of the non-overlapping Gompertz model under the three different configurations of plant and patch distribution for a population size of eight plants and different choices of number of patches (per plant).

Regularly placed plants and uniformly random patches (Fig. A.1(a)) The numbers of plants has little effect on the results. With the plants placed in a fixed, regular distribution, all stochasticity arises from the distribution of the nutrient patches. As a consequence the results are non-scale (i.e. environment size or, accordingly within the model, plant number) dependent due to limited neighbour influence, and there is limited scope for variation across the population.

The effect of different patch numbers has a much more marked effect. Although the total level of nutrients within the environment remains fixed as the number of patches increases, the move towards a homogenised environment leads to both an increase in mean plant size, and also a reduction in the variation of the distribution of final individual sizes.

Uniformly random plants and regularly placed patches (Fig. 2.6(b)) Again

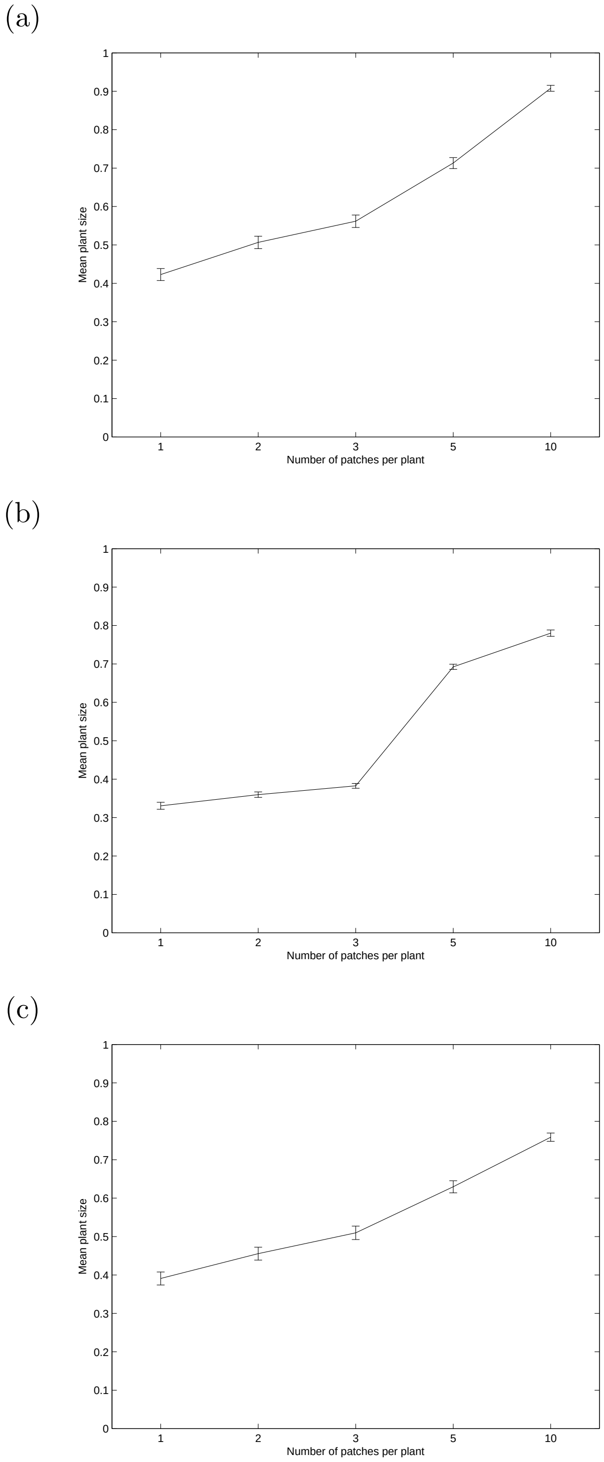

Figure 2.6: *Plots of mean individual sizes of the non-overlapping Gompertz model with a population size of eight plants for the three different environmental/neighbourhood configurations: (a) regularly placed plants and uniformly random patches; (b) uniformly random plants and regularly placed patches; and (c) uniformly random plants and patches. The error bars denote the standard error of the mean.*

an increase in plant numbers results in a small effect on the distribution of final plant sizes, whilsy increasing patch numbers result in a large increase in mean size.

Here all stochasticity is caused by the distribution of the plants, and with no overlapping permitted some plants are potentially highly constrained by immediate neighbours. Those plants which are not constrained by close neighbours have limited opportunity to exploit the extra space due to the regularly placed patches. As the patch numbers increase and the environment becomes increasinly homegenious, they are clearly better able to exploit the resources and space as their is a sharp increase in mean size at higher patch numbers, with a notable increase in mean size from three to five patches per plant.

Uniformly random plants and patches (Fig. 2.6(c)) As with the previous two cases, population size does not tend to have a large effect on the results, whilst greater patch numbers see a large increase in mean size. The performance response to patch numbers more closely resembles that of the regularly placed plants and uniformly random patches scenario (Fig. 2.6(a)) than that of the uniformly random plants and regularly placed patches case(Fig. 2.6(b)), though performance is typically not as good when both the plant and patches are uniformly randomly placed as when just the patches are uniformly randomly distributed.

Overlapping Gompertz model

Fig (2.7) shows the results for the same configurations with a population size of eight plants but for the overlapping Gompertz model. The permitted overlapping at the site of neighbour-neighbour interactions means that the population is better able to collectively take advantage of all patches acquired. In short: all patches acquired by any plant, regardless of neighbour locations, are turned into growth by the individual which acquired them. As a result of this, how the patch acquisition is distributed amongst the individuals in a population makes potentially negligible difference to the mean plant size.

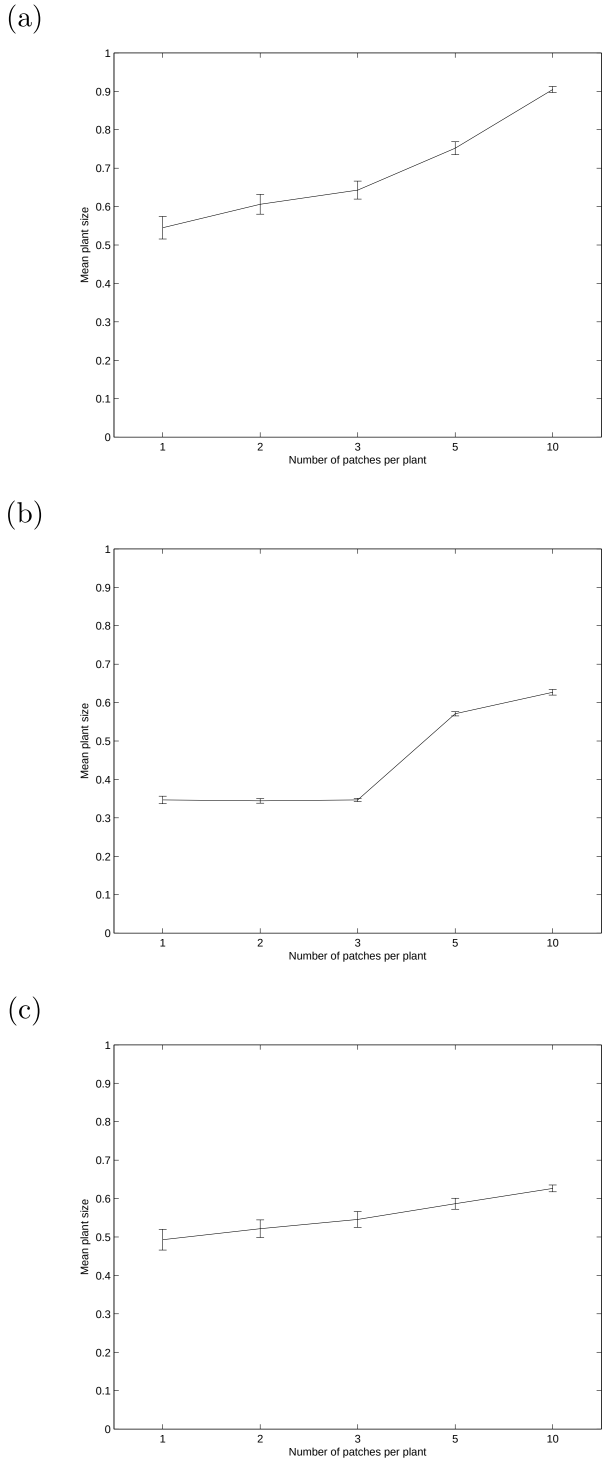

Figure 2.7: *Plots of mean individual sizes of the overlapping Gompertz model with a population size of eight plants for the three different environmental/neighbourhood configurations: (a) regularly placed plants and uniformly random patches; (b) uniformly random plants and regularly placed patches; and (c) uniformly random plants and patches. The error bars denote the standard error of the mean.*

Regularly placed plants and uniformly random patches (Fig. 2.7(a)) Like in the non-overlapping counterpart, population size has a limited effect on results. Increasing patch numbers lead to an increase in mean performance and a reduction in variation amongst results as the environment becomes more homogenious.

Uniformly random plants and regularly placed patches (Fig. 2.7(b)) The number of plants has little effect on performance, whereas as the number of patches grow, mean size increases. Like in the non-overlapping case, a sharp increase in average performance is observed when increasing from three to five patches per plant, with performance continuing to increase up to ten patches per plant.

Uniformly random plants and patches (Fig. 2.7(c)) Population size has a small effect on results. As patch numbers increase, mean size increases though at a far smaller rate than in the other scenarios and in the non-overlapping equivalent case, suggesting the neighbour-boundary constraint of the non-overlapping case actually increases average performance across the population.

2.3.2 Constant growth models results

The results for the non-overlapping constant growth model are shown in figure 2.8, and those for the overlapping constant growth model in figure 2.9.

It can be seen by comparing these results with those of the original Gompertz models that, in both non-overlapping (Fig 2.6) and overlapping (Fig 2.7) configurations, the simplified constant growth model offers very similar results across different patch numbers for a population size of eight plants. For the non-overlapping case the qualitative results as summarised by the mean, variance and skew (see Appendix A) are nearly indistinguishable, whilst for the overlapping case it is only at large plant numbers and low patch abundances in the configurations with uniformly random patch placement that important differences become apparent.

Whilst these results show that statistically the Gompertz and constant growth models are performing very similarly (except at the extreme parameter ranges men-

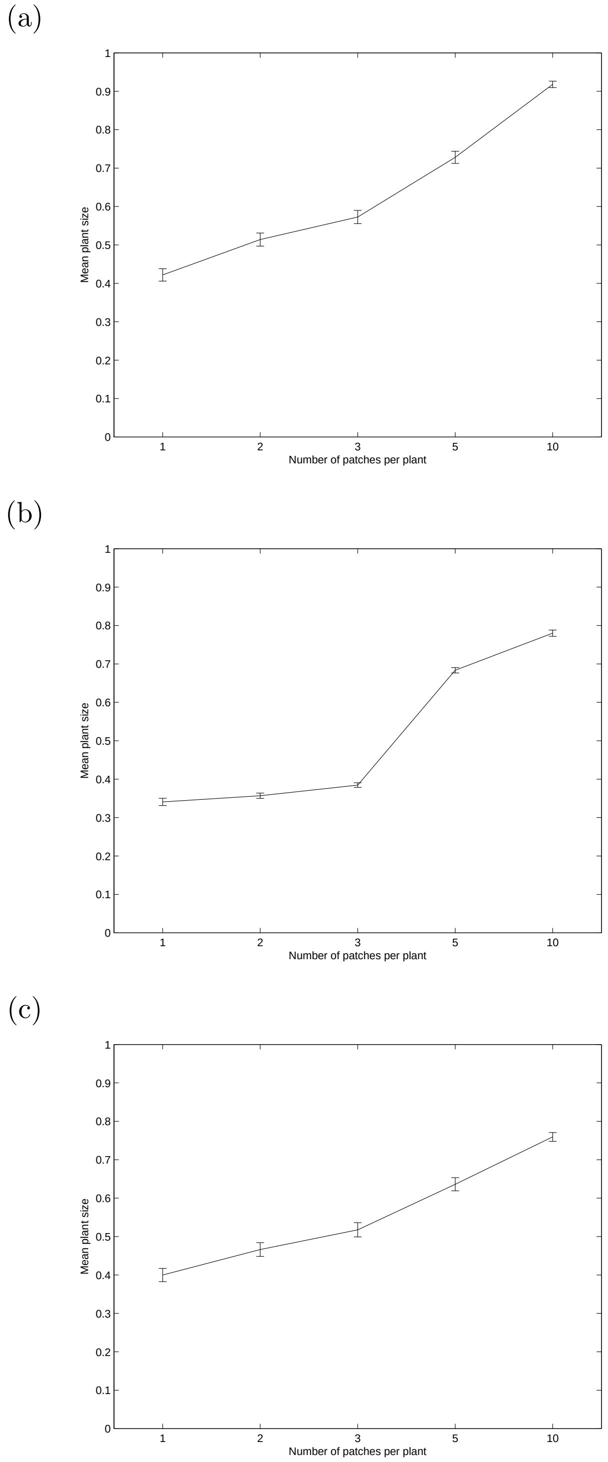

Figure 2.8: *Plots of mean individual sizes of the non-overlapping constant growth rate model with a population size of eight plants for the three different environmental/neighbourhood configurations: (a) regularly placed plants and uniformly random patches; (b) uniformly random plants and regularly placed patches; and (c) uniformly random plants and patches. The error bars denote the standard error of the mean.*

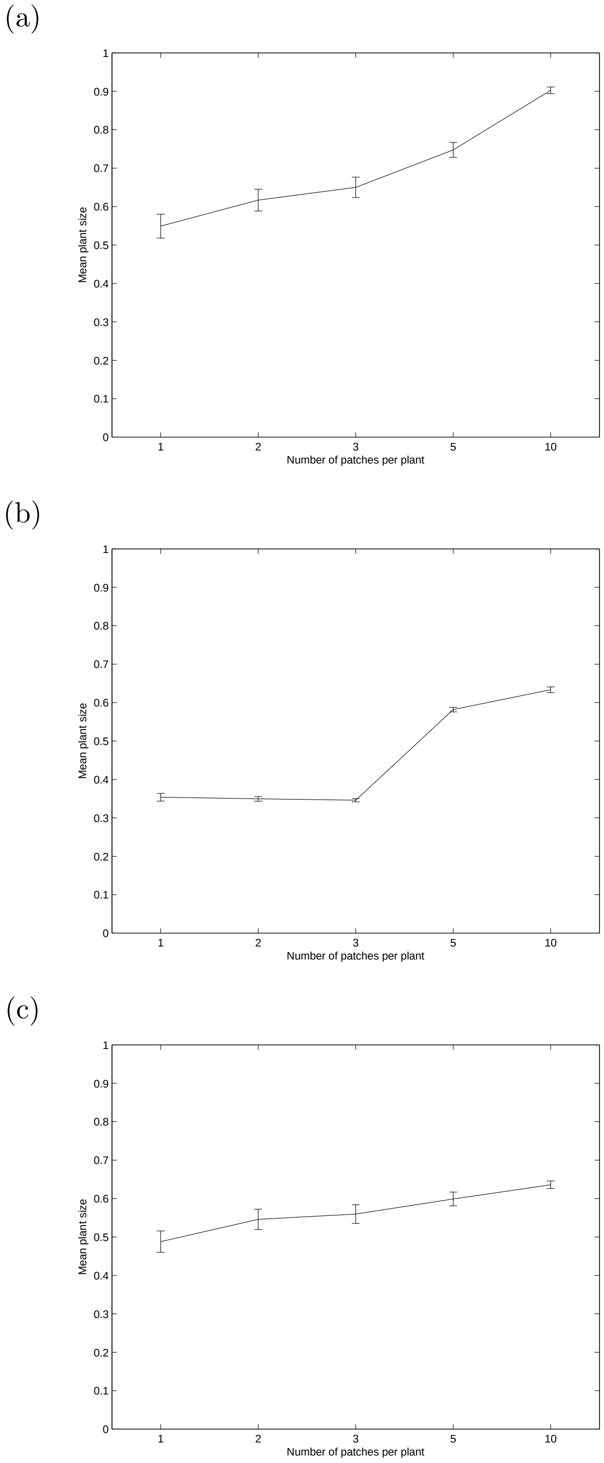

Figure 2.9: *Plots of mean individual sizes of the overlapping constant growth rate model with a population size of eight plants for the three different environmental/neighbourhood configurations: (a) regularly placed plants and uniformly random patches; (b) uniformly random plants and regularly placed patches; and (c) uniformly random plants and patches. The error bars denote the standard error of the mean.*

tioned above), they do not show how closely the two models match in terms of the behaviour of individual plants. In order to access this, a number of simulations were run across different permutations of plant and patch numbers, as well as plant and patch distribution, tracking the performance of each individual through time. In each of these comparisons, a plant and patch distribution were generated (according to the parameter choices) and the constant growth and Gompertz growth models both run with the same initial conditions (samples of comparisons shown in figure 2.10).

With the exception of large plant numbers combined with a low patch count (as reflected in the previous results), the constant and Gompertz growth models not only provide very similar performance on a population level, but also at the individual level. Even at the extremes of the parameter ranges, most individuals perform closely to their counterparts in the Gompertz model. The results can deviate during the growth process (though typically remain relatively close), however the constant growth model consistently performs well at reproducing the final results generated by the Gompertz growth model.

2.4 Discussion

Within the idealised models presented here, a number of factors are assessed. In terms of the stochasticity within the model (and consequently the results), there are two sources: plant distribution and patch distribution. As mentioned in section 2.2.6, if both plants and patches are distributed regularly throughout the environment, both potential sources of stochasticity are removed, and all individuals would perform identically regardless of population size or patch abundance. Consequently the results focus on the effect of stochasticity in plant distribution with regular patches, stochasticity in patch distribution with regular plant placement, and the compound effect of stochasticity in both plant and patch distribution. Within the framework of these stochastic factors, there are also deterministically defined properties: population size, patch frequency and the quality of each individual patch.

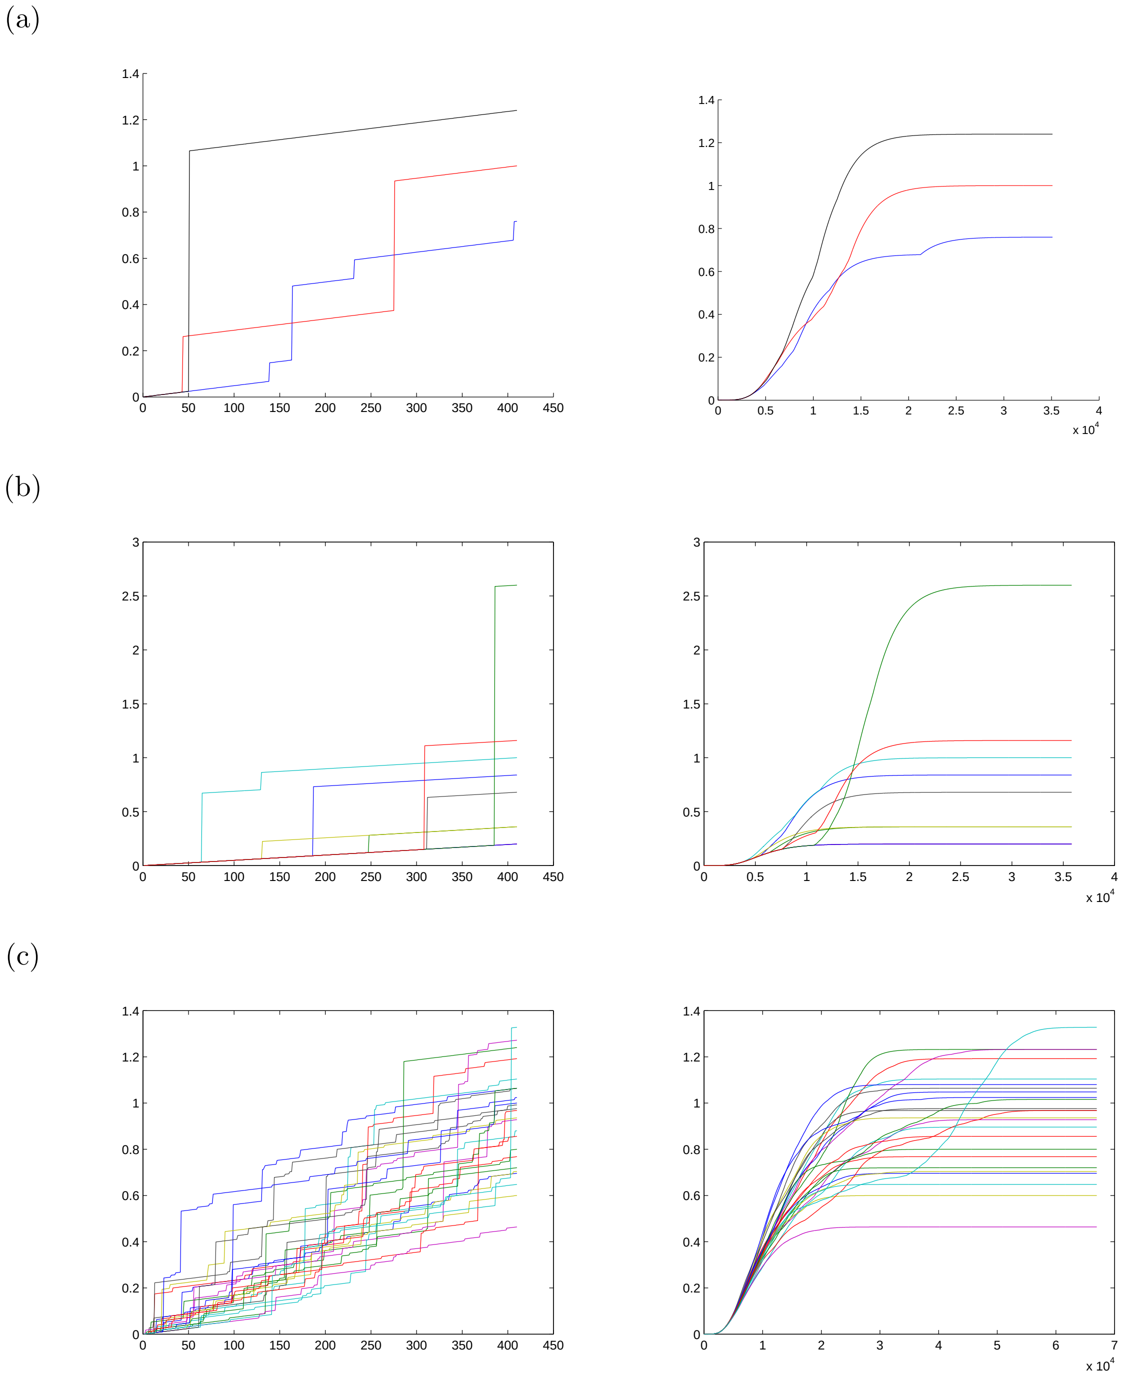

Figure 2.10: *Samples of comparisons between individual-level performance for constant growth (left column) and Gompertz growth (right column) models through time for: (a) three plants and ten patches per plant; (b) ten plants and five patches per plant; and (c) 25 plants and 100 patches per plant.*

2.4.1 Simulations with two plants

The reason for omitting the results for two competing plants is that at a population size of just two plants, certain characteristics offer the potential to make the results unique amongst varying population levels. Consequently they often stand out from the other results and do not fit in with the emerging trends. The first factor is that with just two plants, there can only be at most one neighbour-neighbour boundary of interaction. This means that even when the plants are situated very close together there cannot really be proper crowding effects, as even though the two plants might be tightly constrained on their nearest sides, they will conversely always have open space on their distant sides. Additionally, with just two plants there is a certain degree of inevitable asymmetry; with just two plants, for one plant to get especially lucky often requires that the second plant be especially unlucky.

It has been shown that in certain circumstances the behaviour on a population level can be built up as a series of coupled one-one interactions between paired plants (Schneider et al. 2006; Lv et al. 2008) (and that may very well also apply here), and that observed competition on a pair-wise scale is highly asymmetric (Purves and Law 2002; Schneider et al. 2006; Lv et al. 2008). The problem in this instance is that such an approach does not lend itself to the future plans for this model and the approach to be taken. The coupled interactions of local competitors (Schneider et al. 2006; Lv et al. 2008) describe observed behaviour encompassing all responses and interactions to the environment and neighbourhood. Here the model is explicitly simulating the capture of heterogeneous resources and (in Chapter 3) responses to this resource heterogeneity. Because of this, repeated runs of a two plant model cannot recreate the behaviour of a larger population, neither quantitatively nor qualitatively.

2.4.2 Results for larger population sizes

As the number of plants is increased, the inherent asymmetry of the two-plant competition quickly disappears. Even with just three plants, the size of one individual does not necessarily (depending on the model and configuration) provide information about the size of any of the other individuals beyond their potential combined size (as determined by the resource level). As the number of individuals increases further, the behaviour of one plant provides increasingly little information about another individual's performance in the remaining population. However, as the population size increases, the population-level behaviour converges and so the probability distribution for the size of an individual in a population also converges.

Regularly placed plants and uniformly random patches

With regularly placed plants and uniformly random patch distribution (Figs 2.6(a) and 2.7(a)), there is little sensitivity to population size other than at low patch frequencies, with increased sensitivity in the overlapping model. In these circumstances, increasing population size sees a small increase in mean performance, and an increase in variation. At low patch numbers, most plants do badly, and only a lucky few gain access to the patches. As the population numbers (and accordingly environment size and total patch numbers) increase, whilst average resources per plant do not change, the possibility for aggregations in resource distribution increases. When no overlapping is allowed, any of these lucky plants have their growth limited by their neighbours, and since the plants are regularly placed, the available space per plant is always (relatively) limited. With overlapping permitted, no such constraint applies, and an individual can fully exploit the available resources.

As the number of patches per plant increases, the probability that a patch lies within a plant's initial growth range also increases, and so too does the probability of an individual reaching further patches. Although an increase in patch numbers per plant results in lower quality patches, the greater access to the patches by an individual results in, on average, a larger quantity of nutrient being acquired. Accordingly, an increase in mean plant size is observed. Increased patch numbers per plant also result in a decrease in variability. This is to be expected since as patch numbers per plant increase, the resource distribution converges towards a homogeneous environment, removing the sole cause of variability in the model. Since at such a limit the regularly placed plants would not grow large enough to encounter each other, it follows that the behaviour of the overlapping and non-overlapping models converges.

There is little difference in average performance between the overlapping and non-overlapping models apart from at low patch numbers for larger population sizes. In these conditions, the overlapping model sees a slight increase in performance across the population.

Uniformly random plants and regularly placed patches

Swapping the distribution methods for plants and patches (Figs 2.6(b) and 2.7(b)), with patch positions fixed at regular intervals, there is limited scope for a plant to get lucky in patch acquisition. Instead, the "luck" (or lack of) comes from relative location to neighbouring plants. Variance remains relatively small since whilst the only stochastic factor is the plant distribution, any extreme variation in this distribution is largely damped by the fixed patch locations. In effect, plants can be extremely unlucky in being tightly neighboured on either side by competitors, with no patches within close proximity. However, positive luck can only extend so far as to being uninhibited by competitors, and does not include the scope to grow into a large aggregation of patches. A small decrease in mean performance is observed as plant numbers increase in the non-overlapping model. This is caused by the fact that as the numbers of plants increase, the potential for being unlucky (as described above) gets larger, and out-weighs (at a population level) the corresponding good luck for the other plants.

Patch numbers have a much larger impact on performance, with a large increase in mean performance as the number of patches per plant increases. Despite lower quality patches, the increased availability of access to an initial patch (and often subsequent patches) leads to an overall higher acquisition of nutrients; rather than having a lucky few plants in a position to grow larger with little chance of maximising this potential, as patch numbers increase more of the plants in this position are able to obtain patches and benefit from them. This also accounts for the increase in variation with plant numbers at higher patch frequency.

The overlapping and non-overlapping models show similar behaviour. Relatively unaffected by plant numbers, as patch numbers increase, plants on average do better with a decreasing numbers of plants doing particularly badly. When no overlapping is permitted, this is a result of those plants with little room between themselves and neighbours increasingly likely to acquire some patches as the frequency increases. With overlapping allowed, at low patch numbers, most plants do not obtain any patches. Consequently most only achieve the initial maximum size limit with only a select few growing bigger. As the number of patches increases a little (and their quality decreases), more of the plants that would not have obtained a patch are now able to do so. At the same time, fewer of the plants originally gaining a significant boost from one or two high quality patches achieve the same level of nutrient acquisition and growth. This trend is reversed as the number of patches increases further. The plants on the fringe of groups or that have space to themselves are able to obtain nearby patches and continue growing into space containing further patches whereas those within a group will obtain locally available resources but soon find themselves growing into depleted areas already exploited by competing neighbours. This causes the average plant to do better than when in a more heterogenious environment (fewer, higher quality patches), but still allows those lucky enough to be able to grow into empty space to make the most of the opportunity, hence the mean increases.

The significant increase in mean plant size between patch numbers of three and five patches per plant suggest some sort of “tipping point” occurs. With the increase from one to three patches per plant, there is only a modest increase in mean performance, but increasing patch frequency to five patches per plant results in more than a doubling of mean plant size. Increasing patch numbers from five to ten per plant does not see anything like as bigger gain.

Unlike with regularly placed plants and uniformly random patches (Figs 2.6(a) and 2.7(a)), there is a noticeable drop in average performance in the overlapping model compared to the non-overlapping model as patch numbers increase.

Uniformly random plants and patches

When both plants and patches are placed according to a uniformly random distribution (Figs 2.6(c) and 2.7(c)), the compound effects of stochasticity in both plant and patch distribution is observed. Like in the regularly placed plants and uniformly random patches case, there is little effect from population size other than at low patch numbers, again emphasised under the overlapping model, for the same reasons as discussed above. Indeed, the behaviour in general is quite similar, with increases in mean plant size as well as a decrease in variability as patch numbers increase. The decrease in variability is to be anticipated, as performance is levelling out across the population as the environment tends toward homogeneity. Those plants fortunate to be isolated from immediate competition see an increasing likelihood of obtaining patches and thus growing larger, but are less likely to get particularly lucky and obtain a large quantity of resources.

The overlapping model sees a reduceded average performance compared to the non-overlapping model at anything other than the lowest patch numbers.

Non-overlapping vs overlapping results

As discussed in Section 2.2.3, in representing the real three-dimensional world in an idealised one-dimensional model there are decisions that need to be made regarding interpretations of space. Specifically, when it comes to issues of root overlap, it needs to be decided what is permissible and what is not. Whilst in reality there has been an observed tendency for the roots of competing plants to favour segregation (Schenk, Callaway & Mahall 1999), competing plants also can, and often do, experience overlapping roots (Robinson et al. 1999; Wijesinghe et al. 2005), and this interaction cannot be dismissed when it comes to the importance of the effects of inter-plant competition. Whilst root interaction and overlap can occur, multiple roots (be they of the same individual, or two or more competing plants) cannot occupy the same region of space simultaneously. When space is represented in one-dimension though, any overlap means just that: both plants are occupying the same region of space at the same time.

Two approaches have been taken to try and accommodate both possibilities: "overlapping" and "non-overlapping" models. With the overlapping model, the root systems of different plants are permitted to grow into the same space. This allows for the phenomenon of root overlap to occur, and the corresponding consequences for the individuals involved. This approach can be justified by considering the root systems not as a solid impenetrable mass, but rather as a "loose" collection of roots growing in a given direction, with space and room for other roots to explore the same region simultaneously. Analogously, the plants can be thought of as directing the root growth in both directions through a tube (i.e. the environment is a torus), albeit with "perfect" precision in terms of acquiring resources along its length.

In the non-overlapping case, a more literal interpretation is taken of the one-dimensional physical space, and the presence of the root system of one plant act as an impassable barrier to that of a competitor. This removes the possibility of root overlap, and effectively enforces complete segregation.

Between these two approaches, it is intended to minimise assumptions about plant interactions and allow the respective behaviours in both circumstances to be evaluated and compared, so as to isolate the effects of direct interactions that occur with neighbour encounters.

Comparing the non-overlapping and overlapping results, with the exception of when plants are regularly placed, the overlapping results show a drop in mean plant size at higher numbers of patches per plant. This can appear counter-intuitive since in the non-overlapping case, individuals placed in close proximity to neighbours may find their growth stunted, and the more patches an individual acquires the more likely it is to encounter a neighbour and find growth constrained. In contrast, with overlapping allowed, every plant is able to grow to at least its initial size limit, and also to turn all resources acquired into growth. However, the absence of neighbour imposed boundaries can in fact both help and hinder an individual. One the one hand, any patches acquired by a plant or remaining intrinsic growth can always be fully utilised by a plant when there are no boundaries in place; a bounded plant may obtain resources or have remaining intrinsic growth that it is unable to use as it is constrained by neighbour imposed boundaries. On the other hand, a bounded plant is guaranteed to always be growing into unexploited soil, whereas a non-bounded plant could be growing into a region already explored and depleted by a competitor. To summarise: whilst an overlapping plant is always able to fully grow to the limits of its resources, a non-overlapping plant is always growing (when able to) into an area which potentially contains more resources to be acquired.

2.4.3 Constant growth models

The results show that the constant growth models return results that are very similar to the Gompertz growth models. The comparison runs between the two models when run under identical inital conditions (examples shown in fig. 2.10) demonstrate that not only does this occur on a population-level, but also at an individual level. The deviation between the Gompertz and constant growth models under the combination of large plant and small patch numbers can be easily explained. Just as in the Gompertz models, as the population grows (and accordingly so to the total physical space), in the constant growth models the possibility for plants to grow increasing large in the overlapping configurations continues to increase. Whilst the potential maximum size increases with available space, so too the probability of reaching it decreases. However, with instantaneous jumps rather than a temporary increase in relative growth rate, it is much more likely that a plant will be able to take advantage of these opportunities in the constant growth model. This probability is increasingly great at lower patch numbers, since the patches are high quality (p is relatively large) and so the instantaneous jumps are correspondingly so. If a plant obtains a single patch and consequently receives a large jump, it not only receives a large, immediate benefit over its competitors from the initial jump, it is also possible that it will, in the process, encounter another patch, and so on. Put simply: with large patches and a big area to grow into, the probability of being lucky is much bigger than in the

Gompertz model. As patch numbers increase (and therefore individual patch quality decreases), these jumps become much smaller. As these jumps decrease in size, the behaviour (in terms of final results) tends to that of the Gompertz models as the effect of the instantaneous jumps tend to a short-term increase in growth rate, as is the benefit to patch acquisition in the Gompertz models.

These deviations can be argued to be of limited relevance to the main aims of this study for the following reasons. One reason why these deviations at large plant numbers are of little consequence is that the whole reasoning for repeating these experiments for different plant numbers is to identify the trends as population size changes. With this in mind, it is not necessary to run the simulations for particularly large population sizes as the results at smaller populations can give the same qualitative information. That the constant growth model matches the Gompertz model at smaller population sizes means that it can be used in place of the Gompertz model at these levels, and the behaviour at larger scales for the Gompertz model can be extrapolated from the trends observed in the results above. Just as differing behaviour at large population sizes need not be an issue in using the constant model in place of the Gompertz model, likewise deviation at small patch numbers/large patch quality does not provide a problem. As the plan for these models is to test in more complicated environments, with larger numbers of (lower quality) patches, this will inherently remove the issue of low-frequency/large-quality patches.

It is worth stressing that as the quality of the individual patches (and more specifically the marginal benefit to an individual from acquiring a patch) converges to zero, the constant growth model does no converge towards the Gompertz growth model, but the results of the two models do converge. The constant growth model always demonstrates constant growth between patch acquisitions and the quality of the individual patches does nothing to affect this: it only effects the size of the patch-dependent jumps. As the quality of the patches tends to zero, so too do the size of these jumps, and so the constant growth model converges towards simple linear growth (i.e. it converges towards constant growth without any patches). Similarly, the Gompertz model tends towards its base behaviour (that is, its behaviour given no patches are acquired). If no patches are present/acquired (or equivalently the patches are of zero quality/benefit to the plants), the growth behaviour of the two models remains fundamentally different (namely constant and Gompertz, respectively). However, the relative performance of the different individuals within a population becomes the same, and so whether the plants are growing with constant growth or according to a Gompertz growth function has no effect on their relative sizes.

Future plans for these models consist of running them in environments consisting of a greater number of patches (allocated by more complicated distributions) and the implementation of genetic algorithms (for the investigation of evolutionary beneficial growth strategies). As comparison to the Gompertz model demonstrates, the constant growth model is a successful approximation; with the benefit of increased analytical tractability and reduced computational demands, it has been demonstrated to be suitable substitute in future work for the Gompertz model.

2.5 Conclusion

In the above work two different approaches (non-overlapping and overlapping) to modelling stochastic plant growth and resource acquisition amongst competition in one dimension have been demonstrated using an amended Gompertz function to define growth, and variation in plant locations and patch distributions providing the stochasticity. These models demonstrate strong trends in response to patch abundance/quality, as well as relative resilience to changes in population sizes (over a minimum population size and at relatively higher patch frequencies). These trends allow the use of modelling at lower population numbers to predict the behaviour at larger population sizes, and also offer a qualitative indication of the effects of stochasticity in plant location and/or nutrient distribution.

When the neighbourhood is defined with plants regularly spaced and patches independently uniformly randomly distributed, before environmental conditions are taken into account each plant has equal potential for succeeding regardless of whether under the non-overlapping or overlapping approach. It is purely the stochasticity in the patch locations that determines the variation in the final size hierarchy.

Conversely, when plant positions are varied under non-overlapping growth, some plants have the potential to grow much bigger than others before the patch distribution is even taken into account. In effect, the plant locations dictate maximum possible size, and the position of the patches how much of that potential is maximised and how quickly. When overlapping is permitted, these strict limitations are removed, however a plant in close proximity to its neighbours is going to have to get lucky in order to get big; either with a large, localised collection of patches it can reach before it begins to overlap with its competitors, or by getting enough of an advantage to "leap-frog" them in the race for available patches located beyond them.

This ability to get lucky depends on the patch distribution, and so when the patches are regularly distributed, it is unlikely that a plant in a crowded area is going to be able to attain any meaningful advantage over its immediate competitors. By comparison, when both plant and patch positions are independently uniformly distributed, there is a greater probability for an individual located in a densely populated area to get lucky and receive a significant advantage over its local neighbours.

In both the non-overlapping and overlapping models, maximum total yield is greatest when plants are regularly located and patches uniformly distributed. This suggests that avoiding inter-plant competition improves the average performance, and so competition cannot be strictly asymmetrical on a population-wide scale. The difference in yields between treatments is greater in the overlapping model than the non-overlapping model, demonstrating that enforced root segregation (by neighbour-neighbour boundaries) has a significant effect on population growth; again, reducing competition helps to maximise growth.

Also, in both non-overlapping and overlapping environments, when plants are uniformly placed, there is little difference in overall yield regardless of patch distribution, consistent with experimental observations (Day et al. 2003). However, whilst total yield remained constant, there is larger variation in individual size when resources were randomly distributed. This would imply that whilst resource distribution does not affect overall yield, it does have an impact on the uniformity of the yield of individuals within the population.

Comparing the results of the non-overlapping and overlapping approaches to modelling in one-dimension, the impact of neighbour-neighbour boundaries can be identified as in all other ways the models are identical. The impact of neighbour-neighbour boundaries has the potential to be more significant when the spatial distribution of neighbourhood is varied (i.e. plants are not regularly placed) and more pairs of plants will consequently be in close proximity to one-another, and conversely more individuals will be granted space to themselves. It is shown that whilst imposed neighbour-neighbour boundaries can limit an individual's growth by constraining the space available to them, on average it actually results in better performance due to forcing the plants to use their resources to proliferate into previously unexplored (and more importantly, undepleted) areas.

This interesting result touches upon on the potential benefits for an individual to selectively proliferate in response to signals from its environment, so as to maximise the potential return (in terms of the acquisition of resources) per unit growth of the root system. This ties in with real observations of plants' propensity to avoid direct competition (i.e. for competing plants to segregate their roots; Schenk, Callaway & Mahall 1999). Whilst in the non-overlapping case it is strict rules governing the model that force the plants to display this behaviour, going forward it is the non-bounded model but with the plants equipped to be selective in how they proliferate which is the area of interest for this work.

However, these Gompertz models have drawbacks that prohibit their use in more complicated, realistic environments and/or where many repetitions are required. In offering up a pair of alternative models using a constant intrinsic growth rate, rather than an asymptotic function, to determine growth, it has been demonstrated that these simpler, computationally quicker and less intensive approaches can yield results virtually indistinguishable in terms of qualitative behaviour across a range of environmental and neighbourhood configurations. Although the models diverge at certain extremes (under a combination of large population size and low patch numbers), it has been explained why these do not represent a problem and therefore do not stand in the way of using the simpler models for future work.

The constant growth models allow for more realistic and complicated environments to be defined and investigated, and these models can now be coupled with more sophisticated behaviour for the individuals to inspect the impact of changes in the environment on the performance of individuals and populations, but also how different growth strategies perform in these different conditions of resource and competitor abundance and distribution.

Chapter 3

Optimal root proliferation strategies: the roles of nutrient heterogeneity, competition and mycorrhizal networks

Plant and Soil February 2012, Volume 351, Issue 1-2, pp 191-206

Background and Aims Plants proliferate roots in order to acquire nutrients, typically contending with heterogeneous resources and competing neighbours. A mathematical model was developed to identify optimal root proliferation strategies in patchy nutrient environments. The impact of joining mycorrhizal networks was also assessed.

Methods A simple model of growth and competition in one spatial dimension was implemented within a genetic algorithm to obtain optimal proliferation strategies under different scenarios of resource distribution, and in the presence or absence of local competition and large-scale mycorrhizal networks.

Results A strong proliferation response emerged for isolated plants in heterogeneous environments with low resources, and also for plants growing in competition. Even in statistically homogeneous environments, the presence of competition con-

ferred a selective advantage to plants proliferating in the direction of the most recently acquired patch. In the presence of mycorrhizal networks, the optimal strategy switched from symbiosis to proliferation driven growth as the relative cost of acquiring resources through the networks increased.

Conclusions The optimal proliferation response in a given scenario was governed by a hierarchy of factors: resource levels and distribution; the presence or absence of competition; and the marginal benefit of obtaining resources via symbiotic relationships with mycorrhizas.

3.1 Introduction

Nutrient availability in soil is both spatially and temporally heterogeneous over relatively small distances and at scales relevant to plant roots (Cain et al. 1999; Jackson and Caldwell 1993; Farley and Fitter 1999; Gross et al. 1995). Plant roots have to cope with this heterogeneity or patchiness in resource supply and be able to exploit the nutrient-rich zones or patches upon encounter. Moreover, they have to capture nutrients from such patches both in competition with soil organisms and other plant root systems (Hodge et al. 2000a, 2000b; Kaye and Hart 1997). Plants are aided in this respect by the modular structure of their roots systems, which enables architectural flexibility in root deployment (de Kroon et al. 2009; Malamy 2005; Hodge et al. 2009). Localized root proliferation in nutrient-rich patches is a well established response to heterogeneity in resource supply (reviewed by Hodge 2004, 2009) although it has also been demonstrated that root proliferation is of little net benefit to either single plants or plants grown in monoculture as a means to acquire nitrogen (N) from N-rich zones (Hodge et al. 1998a; Fransen et al. 1998; van Vuuren et al. 1996). This is due to the high mobility of nitrate (NO_3^-) ions, which in moist soil have a diffusion coefficient in the region of 10^{-5}cm^{2}s^{-1} (Tinker and Nye 2000). However, when plants are grown in interspecific competition for N from a complex organic patch (mixed N-sources) then root proliferation does confer a competitive advantage (Hodge et al. 1999a; Robinson et al. 1999). In contrast to NO_3^- ions, phosphate ions are relatively immobile with a diffusion coefficient in moist soil of around 10^{-9}cm^{2}s^{-1} (Tinker and Nye 2000). Given the rate of growth and the average lifespan of the finer roots (see Fitter 1999; Eissenstat and Yanai 1997) which are responsible for much of the nutrient uptake, the relative volumes of soil able to be exploited for these different nutrients by a single root per unit root length can differ by orders of magnitude (Fitter et al. 2002). Consequently, whilst a plant's root system can capture NO_3^- from a relatively large surrounding area, successfully obtaining phosphate from the environment requires a plant either to proliferate roots directly within the phosphate sources, or to use another strategy e.g. forming a symbiotic relationship with mycorrhizal fungi.

The most common form of mycorrhizal symbiosis is that of the arbuscular mycorrhizal (AM) (Smith and Read 2008) association which can form on c. two-thirds of all land plant species. The key function of the AM symbiosis is to enhance nutrient capture for the associated host plant while, in return, the fungus obtains a supply of carbon (Smith and Read 2008). However, unlike fungi involved in both the ecto- and ericoid mycorrhizal associations (Hodge et al. 1995; Read and Perez-Moreno 2003), arbuscular mycorrhizal fungi (AMF) have no known saprotrophic capabilities and so are unlikely to play a direct role in organic matter decomposition (see Leigh et al. 2011). The fungal hyphae can however explore a large volume of soil and acquire phosphorus (P) beyond the phosphate depletion zone that rapidly builds up around the root surface (Smith and Read 2008; Sanders and Tinker 1973) and it has been shown that the arbuscular mycorrhizal fungi (AMF) may largely take over the acquisition of P for their associated host plant (Smith et al. 2009). In addition, a key role for AMF in N cycling has recently been identified (Hodge and Fitter 2010) and some of the N captured from nutrient patches may be passed to their associated host plant under some conditions (Leigh et al. 2009; Hodge 2003a; Barrett et al. 2011; Hodge and Fitter 2010), but not others (Hodge 2003b; Reynolds et al. 2005). Plant N:P biomass ratios (g N/g P) can vary markedly, with individual measurements ranging from approximately 1-100, with an average ratio among terrestrial plant species of 12-13 in their natural field sites (Güsewell 2004).

The results from numerous studies verify that the quality, type and distribution of nutrient patches influence the way in which a plant grows and the nutrients which it acquires (see for example Cahill et al. 2010; Hodge et al. 1999b; Fitter 1994; Shemesh et al. 2010). Hodge et al. (1999a) and Robinson et al. (1999) observed the reaction of the grass species *Lolium perenne* L. and *Poa pratensis* L. to the presence of a high-quality N patch. Both species saw an increase in root-length density within the patch and a proportional increase in N uptake. It is also important to note that Robinson et al. (1999) concluded that the results are driven by context sensitivity, with the need for reactive proliferation strongly linked to the presence of inter-specific competition. These conclusions are supported by a modelling study which demonstrated a theoretical basis for the weak affiliation in monocultures between proliferation and N capture. Williamson et al. (2001) and Linkohr et al. (2002) investigated the effect of the presence of a phosphate patch, the latter comparing directly with the results for the presence of an N patch. The different qualities of the two nutrients are reflected in contrasting responses to the N and phosphate patches: unlike in the response to a high quality N patch, an increased availability of phosphate caused an increase in primary root length and a decrease in lateral root density.

It is, however, challenging to disentangle the ecological processes driving the observations from experimental and field-based studies. In particular, the roles of environmental heterogeneity (nutrient patchiness), neighbourhood competition for resources, and the evolutionary context of root proliferation strategies, are hard to isolate.

With such a broad range of environmental factors influencing a plant's growth and so many physiological and morphological responses at a plant's disposal, modelling plant growth can be a daunting task. A number of different approaches have been taken in the past, focussing on different aspects of the problem (for example see Bever 2003, for a review of conceptual and empirical work on soil community feedback and competitor coexistence; Dupuy et al. 2010, for a review of root growth models; Vos et al. 2009, for a review of structural plant modelling).

Such different approaches help to yield insight into different areas of plant growth. For example, Cropper and Comerford (2005), coupled a mechanistic nutrient uptake model with a genetic algorithm to estimate the minimum addition of P required to meet the requirement of a 4-year growth demand of loblolly pine. They found that the amount of P required was doubled when using a low root length density input compared to that with a high root density input. O'Brien et al. (2007), created a spatially explicit model of below ground competition in plants, treating competition for space as an evolutionary game. They found that root spread can be predicted by the cost-benefit ratio for root production, and predict that in areas with no overlap root growth should match resource availability, whilst in areas of overlap they are predicted to display the 'Tragedy of the Commons' (Hardin 1968).

3.1.1 Overview of model

The aim of this work is to elucidate the processes driving root proliferation strategies by considering an idealised model of plant growth and competition in one spatial dimension. The growth of individual plants is modelled in a patchy nutrient environment, where both the quality and statistical distribution of immobile nutrient patches were known. The environmental scenarios included in this work represent plant growth with a growth limiting, fixed, immobile nutrient (such as P), requiring an individual plant to grow to it in order for acquisition, with no diffusion or movement of any other kind modelled. It is assumed that an individual plant can change its root proliferation only in response to local (temporally and spatially) information.

An idealised representation of mycorrhizal networks is then developed by grouping nutrient patches into "networks". Upon encountering a patch, a probabilistic choice is made by the plant to either acquire the individual patch, or to "join" the network and gain access to all patches within the network, albeit at a cost given that AMF are large sinks for plant assimilate (Johnson et al. 2002; Hodge 1996). Defining networks in this way, as a collection of connected patches, is an abstraction from reality (for example see Southworth et al. (2005) for a network theory analysis of mycorrhizal/plant networks), but this deliberately simplified approach allows the interplay between plant competition and mycorrhizal network dynamics to be assessed using a minimal set of assumptions, and also captures plant growth in heterogeneous environments in a way which is computationally tractable.

The basic model is simple and strategic, aiming to identify and broadly quantify the factors driving proliferation and competition in heterogeneous landscapes. Complexity is added to the model systematically. First the optimal proliferation response for an individual plant growing in isolation is obtained, in both a uniformly random and a statistically patchy environment, to show how spatial heterogeneity influences proliferation strategy. These baseline results are then compared to those arising from competition simulations (under the same average conditions). Finally the role of mycorrhizal networks with varying spatial structures is investigated.

Whilst in a given scenario a certain strategy may achieve the greatest fitness "on average", within an evolutionary context this may not be the strategy that best enables an individual to succeed against competitors (Currey et al. 2007). By coupling the model with a genetic algorithm (GA), it is possible to find the best strategy taking into account stochasticity within the environment (resource distribution), the neighbourhood (location of competitors), mycorrhizal network structure (where applicable) as well as evolution (population dynamics, selection and mutation). Such an approach therefore facilitates assessing fitness against evolutionarily relevant metrics rather than simple mean-field properties (Currey et al. 2007; James et al. 2010; Preston et al. 2010).

For an isolated plant in a uniformly random environment, a proliferation strategy driven by encountered nutrient patches should provide on average no advantage/disadvantage, since statistically any other given patch is equally likely to be found anywhere in the environment. However, when patches are distributed heterogeneously, finding a patch provides statistical information that the next patch is likely to be close by. Consequently, one might expect that a bias toward proliferating in the direction of the last found patch would statistically be beneficial in terms of maximising exploitation of the environment per unit growth.

In the presence of competition it is less clear what behaviour will be evolutionarily favoured; in a random environment the acquisition of a patch provides information that the individual is growing, at least in the short term, into unexploited soil. As such, the plant gains indirect information about its neighbours which could in principle be beneficially exploited. Within a patchy heterogeneous environment the acquisition of patches similarly provides indirect information about the relative proximity of neighbouring plants, but also provides statistical information about the location of other patches (at least in unexploited soil). It is not clear what effect this will have on the strength of the proliferation response in comparison to the control experiments, and the GA framework developed here provides a rigourous mechanism through which such questions can be answered.

When mycorrhizal networks are introduced, it is expected that as the cost of acquiring resources from the networks increases, the propensity for an individual to join the networks will decrease. Similarly, it is likely that responsive proliferation will become more important as the benefit from joining the networks decreases. What is not clear is what impact the spatial properties of the networks will have on when this “switch” occurs, and how it depends on the structure of the network itself.

3.2 Methods

Running in Matlab, the central model is conceptually simple: a growing plant proliferates roots in a one-dimensional patchy nutrient environment and receives a growth benefit from the acquisition of nutrient patches. This methodology is simplistic, but it is argued below that it captures the essential ingredients and allows biological, ecological, and evolutionary factors to be isolated.

3.2.1 Environment

The environment ($[0, d]$ on a horizontal axis, where $d = 1$ for an isolated individual, and scaled to $d = P$ for a population of P plants) contains a series of identical, discrete nutrient patches of point physical size (i.e. negligable physical size) and quality p. The quality of a patch reflects the marginal benefit to a plant from its acquisition (see Plant Growth subsection), with

$$p = p_{\text{tot}}/n, \tag{3.1}$$

where p_{tot} is the (specified) expected total nutrient content and n is the (again, specified) expected number of patches in the environment. Changes in p_{tot} and n therefore allow for environments with different quantities and/or qualities of nutrient patches to be defined. For consistency across the various simulations, as well a scaling

Table 3.1: Table of parameters used in the models. "Variable" implies parameter takes no fixed value (detailed in text), "fixed variable" means parameter takes different fixed values in different simulations (described in text) and "cont" and "comp" refer to control and competitive environments respectively.

Name	Property	Value	Dimensions
α	proliferation bias	[-1,1] (variable)	-
β	network preference	[0,1] (variable)	-
L_0	initial maximum size	0.2	mass
$L_{\max}$	maximum size at time t	variable	mass
P	number of plants	1 (cont) 100 (comp)	-
$p_{\text{tot}}\text{pp}$	mean total nutrient per plant	fixed variable	mass
p	individual patch quality	fixed variable	mass
n	mean no. of patches per plant	fixed variable	-
N	no. of patches obtained by individual	variable	-
g	growth rate	0.5	mass/time
dt	time step	10^{-3}	time
d	size of environment	1 (cont) 100 (comp)	length

the physical size of the environment, mean nutrient content and mean patch numbers are scaled on a patches per-plant (ppp) basis. For computational simplicity and to avoid boundary artefacts, the environment is taken to be periodic (i.e. circular).

Two methods of patch distribution are considered, differing in the statistical properties of the distribution of the distance x between patches:

Random nutrient distribution

In this case, n nutrient patches are independently distributed uniformly randomly throughout the environment. That is, each patch is placed independently according to a uniform distribution across the entire environment. The existence of a patch at a given location therefore contains no information about the locations of other patches. Throughout this work these environments will be referred to as "random environments".

Heterogeneous nutrient distribution

Here a Pareto distribution (specifically, a non-standard Pareto distribution of the second kind; see Johnson et al. 1994) is sampled to generate inter-patch distances, x, with the probability density function given by

$$f(x) = \frac{ra^r}{(a+x)^{r+1}}, (x > 0). \tag{3.2}$$

Reparameterising the Pareto distribution (as in James et al. 2005) by letting $a = \frac{r-1}{\lambda}$, it is possible to have two parameters, λ and r, which independently define the mean distance between patches and the "patchiness" of the overall distribution respectively. Starting at a random point, patches are placed across the environment spaced according to independent samples of these inter-patch distances. This results in an expected total of $P\lambda$ patches distributed heterogeneously across the interval $[0, d]$. Throughout this work, the parameter r remains fixed at 2.1 in order to generate heterogeneous environments whilst maintaining a finite variance in the Pareto distribution, whilst λ, which is equal to the expected number of patches per unit length, is varied to change the density of patches.

The choice of a Pareto distribution allows truly "patchy" environments to be generated, with individual patches aggregating into larger patches, and areas containing little or no nutrient emerging. Unlike in random environments, with this method the position of one patch provides statistical information about the location of neighbouring patches. Throughout this work these environments will be referred to as "heterogeneous environments". Figure 3.1 illustrates examples of both random and heterogeneous distributions.

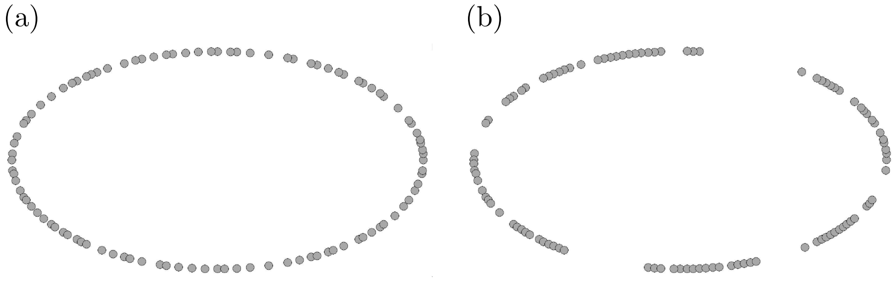

Figure 3.1: *Visual representations of examples of random (a) and heterogeneous (b) environmental resource distributions. Here 100 patches are distributed across the entire environment according to the two different methods.*

3.2.2 Mycorrhizal networks

The networks are defined according to three different sets of rules. The first method is random (Fig 3.2(a)), with each patch independently, randomly assigned to one of 10 networks. The second method uses a threshold parameter determined to provide on average 10 networks, where a patch lying within this threshold distance of a neighbour will belong to the same network as this neighbour, resulting in spatially-local networks (Fig 3.2(b)). The random and local networks are at the extremes of the spectrum of possibilities, suggesting totally uncorrelated spatial structure on the one hand, and strictly local interactions on the other. To bridge the gap between these extremes, the local networks were subjected to a small amount of random “re-wiring” (Fig 3.2(c)). This is achieved by each patch having a 0.001 chance of selection, with selection meaning every patch in the same network after the selected patch (inclusive) being assigned to another randomly chosen network. These generalised network structures allow the complexities involved in spatially extended plant-mycorriza interactions to be captured qualitatively without the need for the extra assumptions and complications required by a fully spatially explicit model.

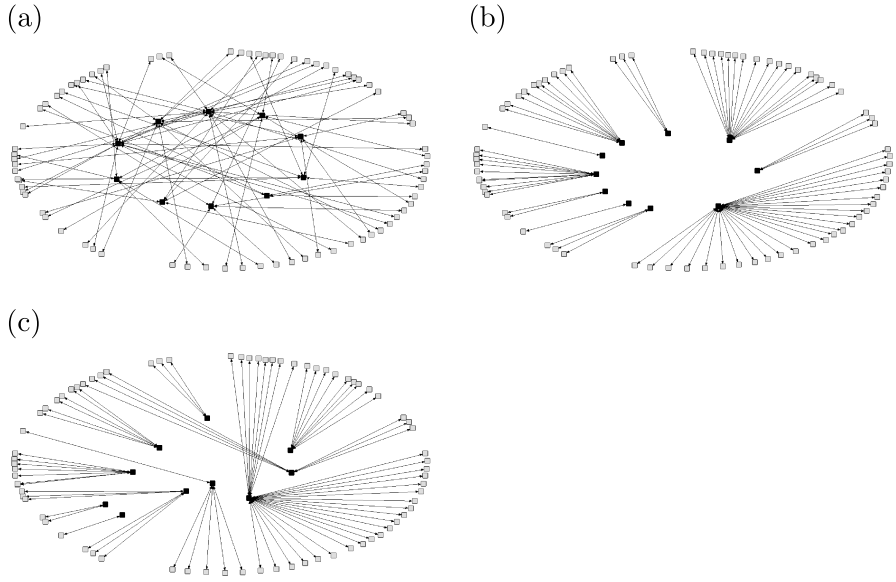

Figure 3.2: *Representative examples of* (a) *random,* (b) *local and* (c) *“rewired-local” networks in a heterogeneous nutrient environment. The black nodes represent the individual networks, whilst grey nodes represent the nutrient patches. The location of black network nodes are purely for visual clarity, whilst relative patch positions are represented by the distribution of the grey patch nodes.*

3.2.3 Plant growth

For the purposes of this study, a plant’s size, L, is regarded as equivalent to the size of its root system. Mass and length are interchangable as interpretations of size of an individual within the model, but here size shall be considered as mass. The plants are independently distributed uniformly randomly throughout the environment, and growth is assumed to occur at a constant rate, g, reflecting an unmodelled background homogeneous nutrient resource. Growth is initially uniform, with proliferation equal in both directions around the environment away from the "centre" of the plant. The choice of g is arbitrary and has no impact on the outcome of the model provided suitably small time steps are implemented (i.e. provided the amount of growth in a given time step, $g{*}dt$, is sufficiently small). If a plant encounters a nutrient patch, then the plant experiences a rapid (instantaneous) additional growth equal to the quality of the patch, p (Eqn. 3.1). Hence at time t, the plant is of size $L(t) = gt + N(t)p$, where $N(t)$ is the number of patches acquired by the individual at time t. The plant grows until it reaches a maximum size, L_{max}, which is dependent on the quantity of nutrient acquired with $L_{\text{max}}(t) = L_0 + N(t)p$ where L_0 is an arbitrary initial maximum size limit which can be thought of as representing a level of nutrient in the seed (fixed at 0.2 throughout this work for all individuals). Too small a value of L_0 would make it unlikely that an individual would reach an initial patch, whilst too high would make it likely an individual would easily exploit the environment. Consequently the value of $L_0 = 0.2$ is chosen as a suitable middle-ground within the confines of the rest of the model.

Previous work has shown that this method of modelling growth can be used as an accurate characterisation of a Gompertz growth function (as used by Purves and Law (2002); Lv et al. (2008)) with upper size limit equal to $L_{\text{max}}(t)$. It follows that the final size of an individual is $L = 0.2 + Np$ where N is the total nutrient acquired by the plant upon reaching its size limit. The final size of the plant is therefore a measure of the quantity of nutrient it has obtained, and is assumed to be a measure of its fitness and growth success.

Directional proliferation

When a plant encounters a patch, it is possible for the individual to proliferate roots directionally so as to potentially more efficiently exploit its environment. In order to minimise assumptions about a plant's ability to detect and "remember" its environment, a plant's information of the environment is limited to knowing in which direction its last acquired patch was located. A simple trait then uses this information, with each individual possessing a dimensionless parameter α between -1 and 1 which linearly dictates which proportion of its growth it proliferates in this direction, with a positive value representing a bias of growth towards the last found patch and a negative value resulting in a bias away. The parameter α is fixed for each individual, and is allowed to evolve between generations.

Growth with mycorrhizal networks

When the model is run for plants grown in the presence of mycorrhizal networks, expected nutrient per plant, p_{tot}pp, is fixed at 0.4 and expected patch numbers at 25 ppp. The decision to join a network is governed by a second dimensionless trait parameter, β, which takes value between 0 and 1 and equates to the probability that the plant will join the network at a given patch encounter. If the plant does join the network then it gains access to all patches within the network, with a "cost" parameter, c, determining the proportion of benefit the plant loses relative to if it were to acquire each patch independently of the network ($c = 0$ results in no cost; $c = 1$ implies full cost, i.e. no marginal benefit to the plant from acquiring resources from a network). The values of c tested were: 0.5, 0.8, 0.9, 0.95, 0.99 and 0.999. Whilst only one nutrient is explicitly modelled, the cost can be considered as a carbon cost to the plant in an exchange process whereby all P is obtained from the network. The parameters β and c thereby efficiently characterise the trade-off between gaining full access to local nutrients, and gaining access to distant nutrients (at a cost) via a mycorrhizal symbiotic association.

3.2.4 Population dynamics and evolution

In order to quantify the roles played by proliferation and network parameters α and β under different ecological scenarios, the model was coupled with a genetic algorithm

(GA) for both "control" (individuals grown in isolation in environment size $[0,1]$) and competitive (a population of 100 plants competing for available resources in environment size $[0,100]$) conditions. Using the final size of an individual as a measure of its fitness, the GA allows the relative success of plants with different values of α and β to provide evolutionarily derived optimum values for a given scenario.

The GA works as follows:

1. An initial population is created with each individual possessing randomly chosen proliferation (α) and, where applicable, network (β) parameter values.
2. The model is run and these individuals are then assessed for success within the generation as defined by the fitness function.
3. The most successful are identified and the next generation is created as their offspring, inheriting their α and β values.
4. The offspring then experience small, independent, individual mutations to their inherited parameter values.
5. Return to step 2.

This cyclic process continues until enough generations have been iterated for convergence (absolute or statistical) to occur. The GA uses a truncated selection process to select the fittest 10% of individuals in a generation, with each of these individuals producing 10 cloned offspring with mutation in the subsequent generation. Mutations, defined here as small random changes to the inherited parameter values, occur with 100% probability, with the mutations sampled from a uniform distribution with limits -0.05 and 0.05. See Supplementary material for further details of GA choice and implementation.

3.3 Results

3.3.1 Proliferation reponse for isolated plants, and plants in competition

The evolved mean values of the proliferation parameter α are shown in Figure 3.3, with bars indicating the variability in the outputs of the evolutionary algorithm. Explicitly, for each evolutionary simulation, we subsample at 1000 generation intervals so as to arrive at 100 pseudo-independent samples. Autocorrelation analysis establishes that this interval is sufficient for subsamples to be treated as statistically independent (see Supplementary material for details). The bars show the standard deviations for these subsamples. Treating the data as independent samples, significant differences in the mean from value 0 can be established by a t-test, indicated by the standard * notation. Results are shown for different permutations of patch numbers and total nutrient content, for control and competitive environments, and with random and heterogeneous patch distributions. These results summarise the long term averages of trait distributions across the modelled populations.

As anticipated, regardless of nutrient levels and patch density, the individuals grown in control tests within randomly defined environments demonstrate no proliferation preference (α remained close to 0).

In contrast, an isolated individual grown in a patchy heterogeneous environment demonstrates a propensity to proliferate towards the last found patch ($\alpha > 0$), though the value of α depends on patch density and, to a greater degree, global nutrient levels. In particular, at high nutrient levels and low patch density (Figure 3.3 (c)) there is no signal for a proliferation strategy.

In random environments, the introduction of competition causes a shift from no benefit in proliferation strategy to a preference to proliferate towards the last found patch ($\alpha > 0$). This demonstrates that the presence of competition allows an individual to improve its relative ability to exploit its environment by responding to encountered patches, even without any implicit statistical information about the en-

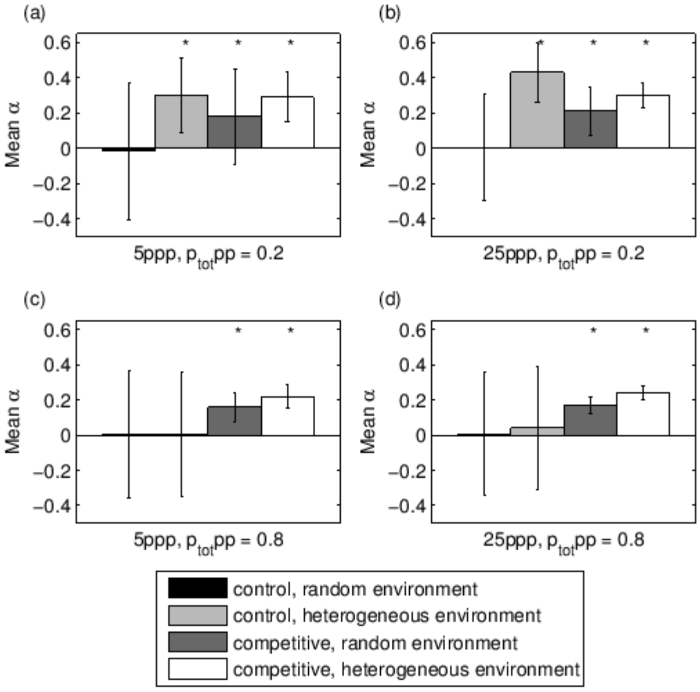

Figure 3.3: *Evolved mean values for proliferation preference, α, for plants grown in random and heterogeneous environments, in isolation (control) and among neighbours (competition). All tests were run for 5 ((a) and (c)) and 25 ((b) and(d)) patches per plant (ppp), and nutrient levels of p_{tot} equal to 0.2 ((a) and (b)) and 0.8 ((c) and (d)) per plant (p_{tot}pp). The bars show the standard deviations of 100 uncorrelated sub-samples from each simulation, see text for details.*

vironment.

At high nutrient levels, the introduction of competition in patchy heterogeneous environments causes an increased propensity to proliferate towards the last found patch (Figure 3.3 (c) and (d)). However, at lower nutrient levels where the control tests provide a strong signal for a positive α value, a weakening of the signal and a reduction in the benefit of possessing a positive proliferation value is observed (Figure 3.3 (a) and (b)).

With the exception of isolated individuals in random environments (where no proliferation strategy ever emerges), all results show that a higher patch density results in an increase in proliferation toward the last found patch (Figure 3.3 (b) and (d)). In contrast, at high nutrient levels there is a reduction in necessity to proliferate in such a way (Figure 3.3 (c) and (d)).

3.3.2 Additional heterogeneous control tests

In order to investigate the observed dependence on patch numbers and nutrient levels, further control tests within patchy heterogeneous environments were carried out for a larger set of nutrient levels (p_{tot}pp ranging from 0.05 to 0.8 in 0.05 increments) and a greater range of patch numbers (5, 10, 25, 50, 100 and 250 ppp). Figure 3.4 summarises the results from these tests, with standard deviations omitted for clarity (see Supplementary material for details).

At low patch numbers, proliferation response to patches (α) remains at a positive value as the total nutrient level increases, before beginning to decrease as the total nutrient content is further increased. This continues until α settles around 0.

As the number of patches increases (and accordingly individual patch quality decreases by Eqn. 3.1), both the rate at which α converges to 0 and the nutrient level at which this transition occurs are seen to change. With increased patch numbers, the transition takes place over a smaller range of nutrient levels (i.e. a steeper descent), and the transition occurs at higher nutrient levels. As well as these qualitative and quantitative changes to the observed transitions, with increasing patch numbers there

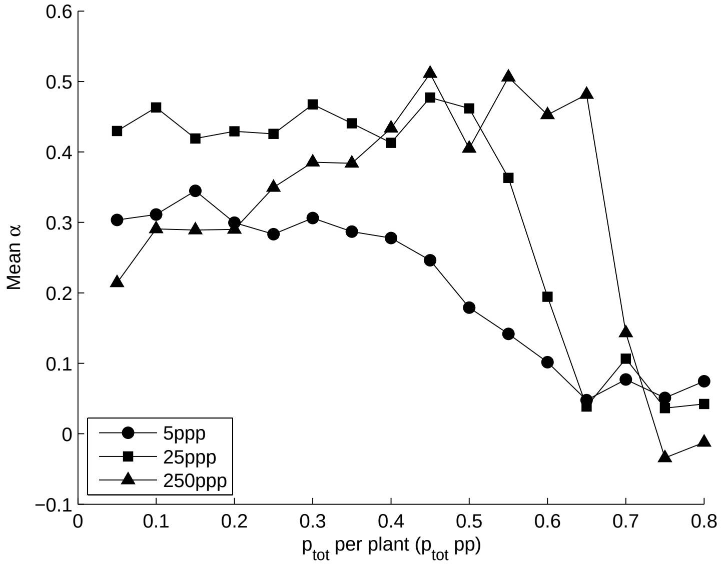

Figure 3.4: *Evolved mean values for proliferation preference, α, plotted against expected nutrient levels per plant ($p_{tot}pp$) for isolated individuals grown in heterogeneous environments. Expected patch numbers are: 5, 25 and 250 patches per plant (ppp). Standard deviations omitted for clarity (see Supplementary material for more details)*

is a lowering of the mean α value at lower nutrient levels.

From the lowest mean patch number (5 ppp) to the highest (250 ppp) tested, the evolved response of the proliferation bias, α, to an increase in total nutrient content shifts from steadily reducing and converging to value 0, to increasing steadily before sharply dropping to a value of 0. The results for 10, 50 and 100 ppp are consistent with these trends but are omitted from Figure 3.4 for clarity.

3.3.3 Proliferation and symbiosis responses for plants grown in competition in the presence of mycorrhizal networks

Figure 3.5 summarises the results for tests with competition in the presence of mycorrhizal networks. Figure 3.5(a) shows the effect of cost on proliferation strength (α) for different network types, and figure 3.5(b) shows the effect on propensity to join the network (β). Results are plotted against a rescaled cost $-\log(1-c)$; this rescaled cost increases monotonically with c, and the logarithm form allows a wide range of c values to be displayed.

In general (irrespective of network type) at relatively low cost to acquiring resources through the networks (small c value) there is no foraging preference (α remains close to 0) (Fig 3.5*a*)) and a preference to join the network ($\beta > 0.5$) (Fig 3.5*b*)). As the relative cost of acquiring resources from the network is increased ($c \to 1$), there is a switch from a plant's growth being driven by network symbiosis to being proliferation oriented, with it becoming less desirable to join the networks ($\beta \to 0$) and foraging strategy becoming relevant with a positive preference to proliferate towards the last acquired patch ($\alpha > 0$).

The propensity to join the random mycorrhizal networks at lower cost (c) is less than for the other types of mycorrhizal networks. As c increases, this propensity also increases, and carries on doing so past the cost at which under other network types a shift in preference for precision over symbiosis occurs.

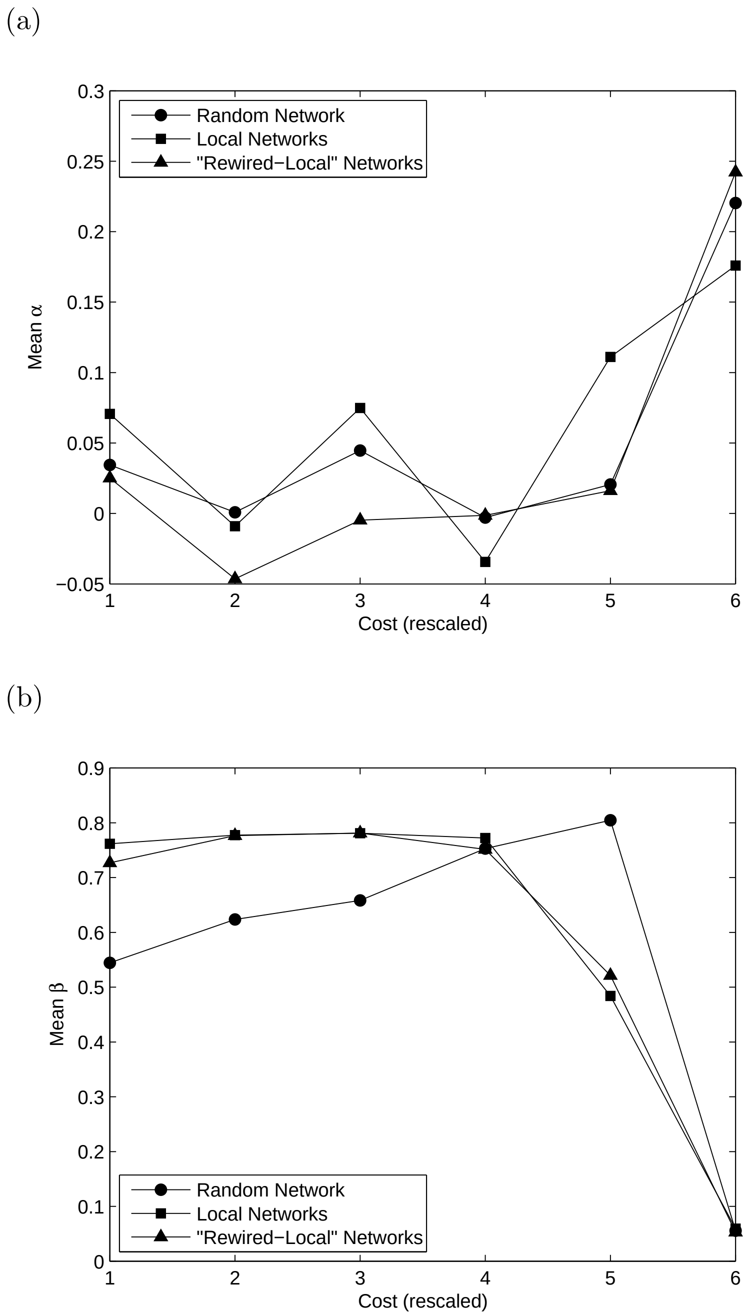

Figure 3.5: *Plots of evolved mean (a) proliferation preference, α, and (b) propensity to join a network, β, values. Plotted against rescaled cost $-\log(1-c)$, where c is the network associated cost (see text for details), for different network types. Standard deviations ommitted for clarity (see Supplementary material for more details).*

3.4 Discussion

Comparing results (Figure 3.3) from the control tests in random and heterogeneous environments demonstrates that no benefit exists from proliferating in response to patch acquisition when the acquisition of a patch provides no information about the environment. When the patches do provide information, this can be used to an individual's advantage by directing proliferation in response to acquiring the patch.

As would be anticipated, exceptions occur to this pattern when global nutrient levels are too low or too high (Figures 3.3(c) and 3.4). As patch quality tends towards zero, so too does the relative benefit of acquiring a patch, and thus the potential reward for proliferating in response to the presence of patches. If p is large enough, the relative benefit from finding such a patch becomes so great that there is no longer a need to be selective in proliferation, despite the increased reward in obtaining the undiscovered patches.

The contrasting results for control and competitive tests (Figure 3.3) within random environments confirm that, in the presence of competition, finding a patch imparts useful information to a plant. This is true even in an environment where a single patch provides no information about the distribution of the remaining patches. With no signal of neighbour presence/proximity contained within the model, an individual can never know that it is overlapping with a competitor and therefore proliferating into exploited soil. However, the acquisition of a patch does inform an individual that it is not at that moment overlapping with such a competitor. Information that proliferation in a given direction is definitely not in a region of overlap provides a strong enough advantage to bias growth in this direction.

In patchy heterogeneous environments where an individual patch potentially contains information about both the unexplored environment and also the presence (or, more precisely, absence) of competitors, a slightly more complicated picture emerges. References to control tests show that within a heterogeneous environment the strength of signal for a positive trait, when it exists (Figure 3.3 (a) and (b)), often decreases with the introduction of competition. Conversely, where little or no signal exists (Figure 3.3 (c) and (d)), the addition of competition leads to the emergence of a positive trait value. This, together with the results for the additional heterogeneous control tests (Figure 3.4), demonstrates that competition could reduce the effectiveness of responding to patches within low-mid nutrient level environments, but in nutrient rich environments could provide an additional pressure that necessitates selective proliferation. Further, the greater α values from heterogeneous competitive environments compared to random competitive environments demonstrates that not only can the addition of competition cause a positive signal to emerge, but part of this signal represents the emergence of a benefit in responding to the patch distribution.

Hodge (2009) observed that not all plants respond in the same way to self/non-self competition (see also Hess and de Kroon 2007; Schenk 2006; Masclaux et al. 2010; Milla et al. 2009). With no ability for an individual to directly detect neighbours, and the plants represented in this work being non-species specific, there was no scope for "kin" or competitor recognition (see for example, Dudley and File 2007; Bhatt et al. 2011; Murphy and Dudley 2009) and evolutionary selection pressure applies only at the individual level in the model.

However, the results shown here provide theoretical support for experimental results seen for particular species. Cahill et al. (2010) found that for *Abutilon. theophrasti* seedlings grown alone and in competition, in treatments combining different levels of resource heterogeneity, that root placement was driven by a "hierarchical set of decision rules dependent on presence or absence of a neighbour". As in the results shown here at high nutrient levels (Figure 3.3 (c) and (d)), regardless of resource heterogeneity, in isolation a plant would grow with a broad foraging strategy ignoring the resource distribution. In the presence of competitors, a more specific foraging strategy was adopted and was modified by resource distribution. More broadly, Cahill et al. (2010) conclude that plants "non-additively integrate information about both resource and neighbour based cues in the environment", which is consistent with the evolved results presented in this work.

When mycorrhizal networks were introduced, the general pattern of behaviour was easy to understand. At relatively low costs for acquiring resources via the networks (low c value) the plant had a strong pressure to join the networks. This is shown both by a propensity to join (large β value), but also by the lack of signal for a positive proliferation (α) value. Effectively a plant needed to join the networks in order to be successful, and consequently proliferation strategy became insignificant. As the cost increased, a point was reached where it no longer remained beneficial to the plant to join the network, and so β decreased to 0. At the same time, root proliferation became important again, hence the emergence of a signal for a positive α value.

At lower costs, the reduced pressure (lower β value) to join random networks compared to local and "rewired-local" networks does not reflect a relative lack in benefit in doing so, either directly to the individual or indirectly by depriving neighbours of resources. One explanation for the observed reduction in β could be that at such low costs, there was only scope for such a small number of individuals to monopolise the networks/resources that within the GA the selection pressure did not get fully expressed when using a truncated selection method.

Within an environment of random networks, the benefit to an individual from joining a network resulted in greater potential to join subsequent networks than with local and "rewired-local" networks. This made it possible for an individual to monopolise not only local resources, but also those over a greater spatial range. Consequently, there was greater potential to deprive the competition of available resources than with spatially structured networks. This explains why it was beneficial to continue joining the networks and resist switching to a proliferation-based growth strategy when the cost of doing so would have suggested otherwise. The lack of difference between results for local and "rewired-local" networks suggest that not enough rewiring took place to make a significant difference. With increasing amounts of rewiring the networks would tend towards random networks, suggesting a greater level of rewiring than was performed would see results in between those of the random and local network tests.

As discussed in the introduction, the implementation of mycorrhizal networks was necessarily idealised due to constraints imposed by the original model. Within such a framework, it was impossible to implement many of the dynamic features of plant/network interactions (Smith and Read 2008; Hausmann and Hawkes 2009). As such, this implementation acted to demonstrate the significance of the spatial distribution of networks among plants/resources on the cost/benefit relationship between plants and networks, and to inform future experiments and their research. The limitations also highlight the need for moving into a more temporally explicit model. Such a model would also allow the order of patch encounter (Duke and Caldwell 2000), temporal network dynamics (Hausmann and Hawkes 2009), and the physical size and temporal properties of patches (Fitter 1994; Hodge 2004) to be integrated.

Also, although it is recognised that roots can also modify their environment and the microbial decomposing community through rhizodeposition processes (Paterson 2003; Paterson et al. 1999; Hodge et al. 1998b; Kuzyakov 2002) this was not included in this model not least because rhizodeposition itself can be affected by many environmental factors (Hodge and Millard 1998; Hodge et al. 1997; reviewed by Jones et al. 2004, 2009; Hinsinger et al. 2009), and thus is often difficult to quantify.

The model is restricted to growth in one spatial dimension, but one can argue that the evolutionary effects on proliferation strategies will be amplified in higher dimensions. In one dimension, local proliferation in the "correct" direction happens 50% of the time when growth is random, and any evolved strategy can only improve upon this. In higher spatial dimensions there are more "incorrect" possible growth directions, and so the relative advantage of a directional proliferation strategy is increased. Similarly, as the number of dimensions is increased, so too is the complexity of the spatial distribution of patches. Consequently, a trait which maximises proliferation efficiency and allows more precise exploitation of space stands to be more beneficial in higher dimensions. As such, it is reasonable that the selection pressure to possess and utilise such a trait will be greater in higher dimensions and that a stronger signal would emerge for a parameter governing such behaviour. However, extension of the simulations to higher dimensions requires further assumptions, and is not considered here.

Robinson et al. (1999) suggested for a highly mobile and diffusive nutrient like N that it is the presence of inter-specific competition that drives the benefit for responsive proliferation. In that scenario, the mobility of the nutrient makes it unnecessary for the plant to be particularly selective in its proliferation of roots, but the added pressure of competition provides the benefit from response to the environment. The results presented here show that similar evolutionary forces can drive the emergence of reactive proliferation in response to randomly distributed immobile nutrients in environments where isolated plants do not benefit from such behaviour. With no direct information imparted to an individual about the presence and proximity of competitors in this model, the acquisition of immobile nutrient patches provided enough information for a plant to gain, on average, an advantage by responding to them, regardless of their distribution or quality.

Supplementary material

Genetic algorithm

The genetic algorithms (GAs) allow for evolved values of proliferation strength (α) and network symbiosis preference (β) to be obtained across different environmental and biological scenarios. The mean values of α (and β where applicable) across the population through time have are shown in Figures 3, 4 and 5 in the main text as they preserve both quantitative and qualitative properties of α and β (for α: < 0, ≈ 0 and > 0 representing bias away from last found patch, no preference and bias towards last found patch, respectively; for β: ≈ 0 and ≈ 1 demonstrates a weak and a strong propensity to join networks, respectively).

Within the genetic algorithm, one generation in the control experiments consists of 100 individuals grown in an independently defined environment, whilst in the competitive scenario 100 individuals grow and compete in the same environment simultaneously. The fitness measure of an individual is its final size, which in this model is equivalent to its success in obtaining nutrient patches.

The GA employs a truncated selection method, with the top 10% (10 in this case) of plants by fitness chosen and each producing 10 “offspring” in the subsequent generation, with each offspring being a clone of its single parent, with a small random mutation (i.e. the mutation rate is 100%). The mutation is drawn from a uniform distribution with mean zero, and limits -0.05 and 0.05. This approach was chosen as results were consistent with other tested implementations of mutation and selection algorithms, but required the fewest generations to obtain a clear signal (where one existed).

Due to the noisy nature of the results generated by the GA, it is important to be sure that enough results have been generated for a clear signal (if present) to be detected, i.e. that the recorded mean parameter values are not just random anomalies, but accurately reflect statistical convergence. The results shown in Figures 3,4 and 5 in the main text are consistent across different environment and parameter choices, with clear trends emerging. Also, all of the different selection methods and mutation rates tested have led to the same results; all that changes is the number of generations required for the signal to emerge. 100,000 generations were deemed sufficient to establish the distribution of the evolvable trait values. For selected parameter and environmental choices, the algorithm was run for 10,000,000 generations and confirmed that the chosen limit of 100,000 was sufficient.

The results from the GA show clear distributions emerging. As an example, Figure S 1 shows the distribution of α for 10,000,000 generations within the GA when no selection bias is imposed (i.e. the selection process is completely random), and so the emergent distribution is what one would see if there were truly no correlation between α and fitness. It is important to note that one would not expect to see a uniform distribution here: the peaks above the lower limit of -1 and below the upper limit of 1 are as a result of this being the mean of 100 bounded values. Consequently a value of 0 or 1 would require all 100 individuals to take this value, which is statistically highly unlikely.

Comparing this distribution with those for examples of random (Fig. S 2) and

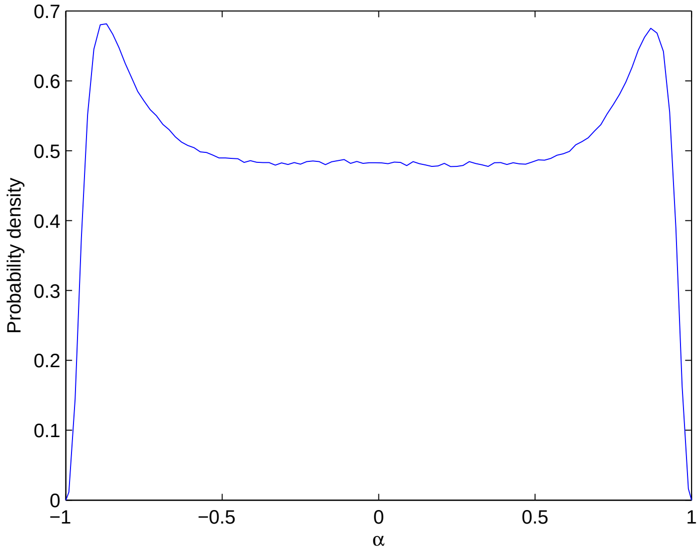

Figure 3.6: Distribution of proliferation parameter α value when tested for 10,000,000 generations with random selection method (i.e. no fitness measure is in place). This represents the expected distribution when the value of α has no effect on fitness.

patchy heterogeneous (Fig. S 3) control tests, it can be seen that in the random case there is no signal at all whilst in the heterogeneous case a biased distribution emerges leading to the recorded positive mean α value (see Fig. 3(a) in main text).

Due to the stochastic nature of the environments and the stochastic process of parent selection/offspring production and forced mutation within the GA, with successive generations a statistical convergence of the distributions of α and β emerge rather than single, optimal values.

Autocorrelation analysis and variability within results

In order to present a meaningful measure of variability, the 100,000 generations are split into 100 subsamples of 1,000 consecutive generations. This approach is supported by autocorrelation analysis which involves testing the data for correlation with itself over different time lags. This is achieved by time shifting the data and calculating

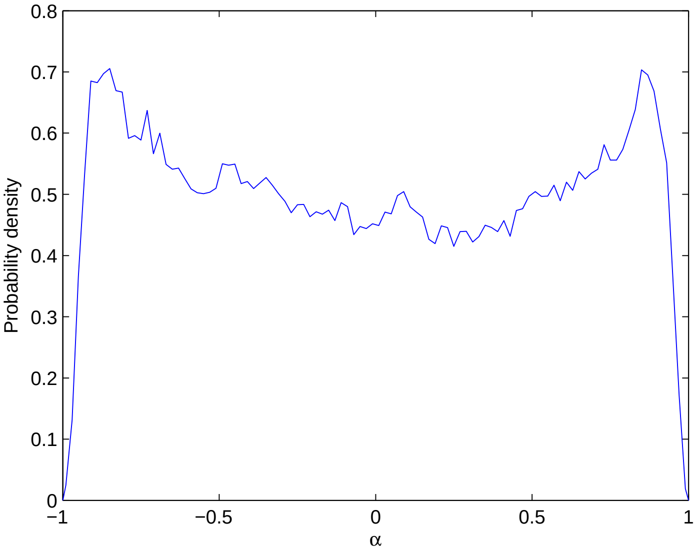

Figure 3.7: Distribution over 100,000 generations of mean of proliferation parameter α (across 100 individuals) for control test in random environments (5 patches per plant (ppp); nutrient level $p_{\rm tot}$pp =0.2).

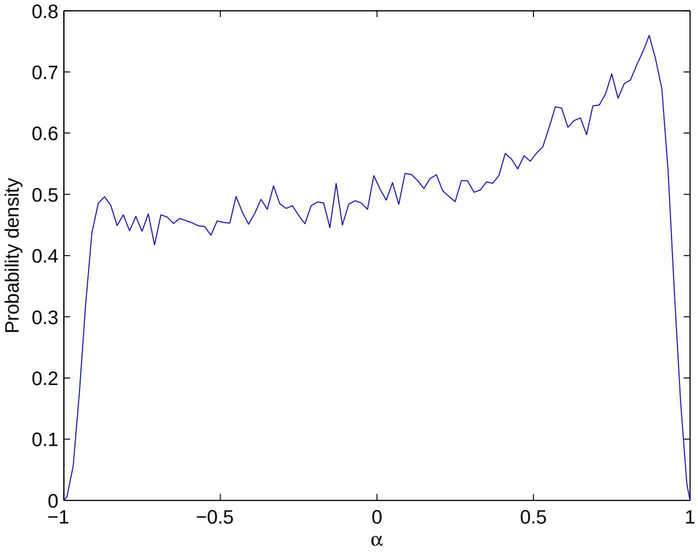

Figure 3.8: Distribution over 100,000 generations of mean of proliferation parameter α (across 100 individuals) for control test in heterogeneous environments (5 patches per plant (ppp); nutrient level p_{tot}pp =0.2).

the Pearson product-moment correlation coefficient (typically referred to as r). The forumula for finding the sample correlation coefficient r for samples $X_1, X_2, ..., X_n$ and $Y_1, Y_2, ..., Y_n$ of random variables X and Y respectively is given by (Canavos 1984):

$$r(X,Y) = \frac{\sum_{i=1}^{n}(X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\sum_{i=1}^{n}(X_i - \bar{X})^2}\sqrt{\sum_{i=1}^{n}(Y_i - \bar{Y})^2}},$$

Figure S 4 shows the correlation of data for control tests with 25 ppp in heterogeneous environments at different time lags (in generations). Each line represents data from simulations with different expected nutrient levels per plant (p_{tot}pp). Such tests were carried out for data from all types of simulation run, and 1,000 generations was consistently seen to be sufficient for correlation to be deemed negligible. As such, results for each test can be considered as 100 pseudo-independent repetitions over 1,000 generations, and the results can be tested for significance to verify recorded values are significantly different from those from a random distribution such as in Figure S 1.

Figure S 5 shows an extended version of Figure 4 (from main text) with bars showing standard deviations. This both confirms that where a positive proliferation perameter (α) is reported there is indeed a significant signal, but also highlights how variability tends to increase with a reduction in the strength of α. This makes sense since where a value close to 0 is reported, it is not because there is any pressure for this value, but rather it arises due to a lack of any signal, hence an increase in variability.

Similarly, Figure S 6 shows an extended version of Figure 5 (see main text) with bars showing standard deviations. It is worth noting that a positive proliferation parameter value ($\alpha > 0$) follows a lack of propensity to join the networks (β close to 0) as described in the main text. It follows that when there is a positive β value there is no clear signal for α and consequently there is a lot of variability. As the cost of joining the networks increases and consequently β tends towards 0, a positive α value emerges along with a reduction in variability. The higher variability evident in the random networks at lower costs also suppports the explanation in the main text that

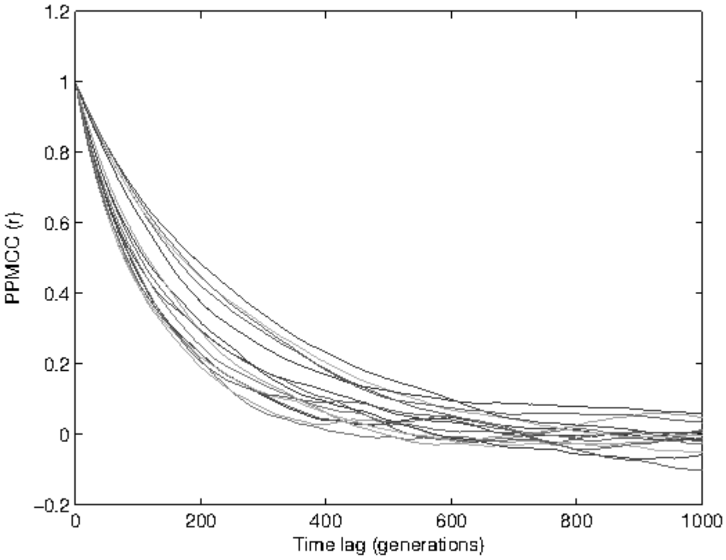

Figure 3.9: Results from control tests with 25 ppp in heterogeneous environments are time separateed and the Pearson product-moment correlation coefficient (PPMCC) calculated. Each line represents results from simulations performed with different expected nutrient levels per plant (p_{tot}pp).

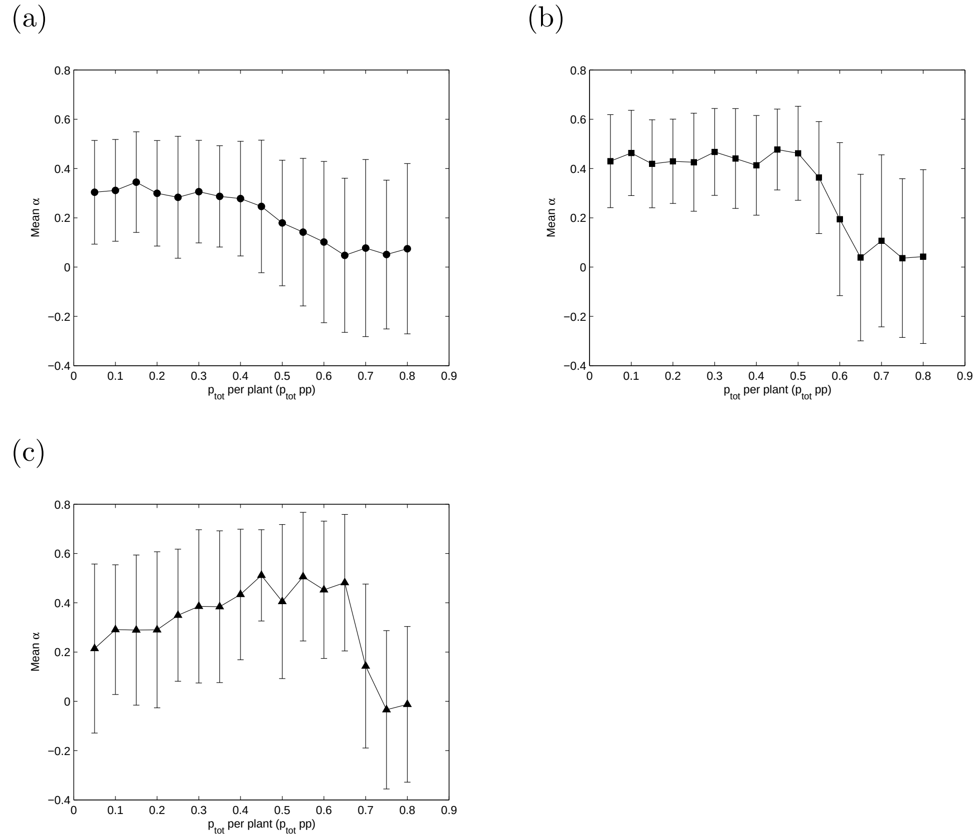

Figure 3.10: Extended version of Figure 4 within the main text. Figures (a), (b) and (c) show the mean proliferation parameter (α) values for 5, 25 and 250 patches per plant (ppp) respectively. Bars denote standard deviations.

the lower β values compared to other networks at low costs are caused by a lack of selection pressure within the GA.

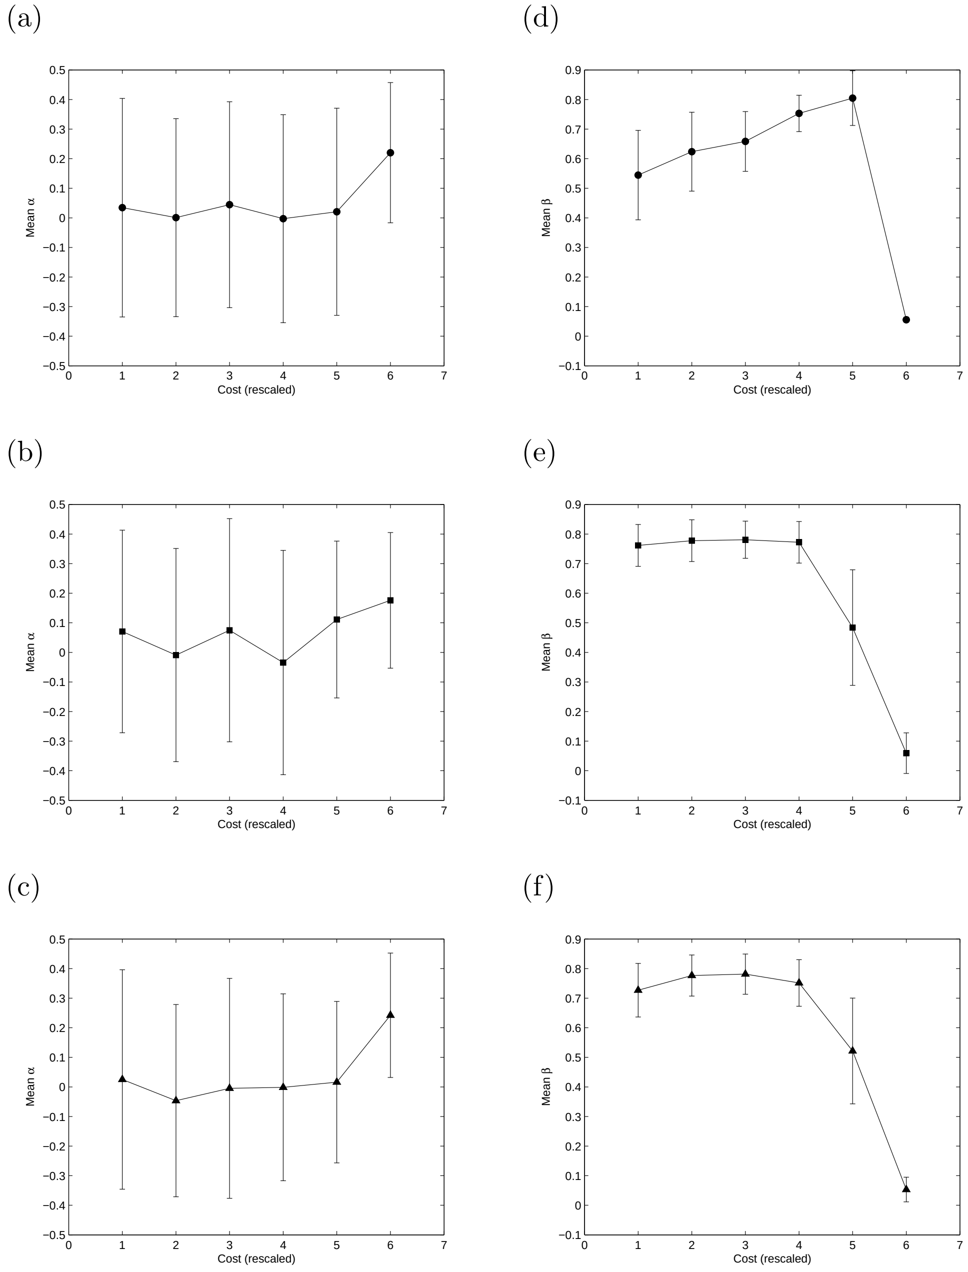

Figure 3.11: Extended version of Figure 5 within the main text. Figures (a), (b) and (c) show the mean proliferation perameter (α) values and (d), (e) and (f) the mean network parameter (β) values at different costs c for random, local and "rewired-local" networks respectively. Bars denote standard deviations.

Chapter 4

Higher dimensions: scaling root growth and proliferation to a two-dimensional model

4.1 Introduction

In Chapter 3, an idealised one-dimensional model of plant growth and competition was presented. Neighbourhoods were defined by independently uniformly randomly placing plants within a specified one-dimensional range, and the environment comprised of point-sized nutrient patches distributed either independently uniformly randomly ("random environments"), or by a random walk (a path composed of random steps; see Grimmet and Stiraker 2004) utilising a long-tailed Pareto distribution for inter-patch distances, causing aggregation of patches to occur ("heterogenous environments"). Individual plants were equipped with a proliferation response parameter (α) which governed their biasing of growth in reaction to the location of their most recently found patch. The model was run under different environmental conditions (varying total resource levels and patch abundance) to investigate the impacts of different effects of resource supply and distribution on plant proliferation strategies.

By running the model within a genetic algorithm, the effects of patch distribution, quality, abundance and overal nutrient content on the benefit of possessing a proliferation response to local (spatially and temporally) knowledge of the environment was assessed.

Overall a hierarchical set of factors emerged in controlling the pressure for a reactive proliferation response, with environmental conditions often only contributing to this pressure when competition was present, consistent with observed results (Cahill et al. 2010; Robinson et al. 1999) Whilst the one dimensional model achieved its aims of quantifying the impact of different factors (both individually and and collectively) on the benefit of plant proliferation responses to their environments, there are intrinsic limitations to the one dimensional model that are discussed in Chapter 3. Specifically the one-dimensional model and environments were deemed limited in being able to accomodate more complicated dynamics, with patch encounter being too strictly governed by initial conditions and more complex physical properties of patches (such as size and life-span) being difficult to include. Within one-dimension, the incorporation of mycorrhizal networks was also necessarily highly simplified.

For this reason the model has been expanded into two spatial dimensions. With a move into two dimensions it is possible to construct more realistic environments, and also to capture more realistic interaction between competing plants. The core functionality of the model largely carries over from the one-dimensional approach, with some changes being necessary to accomodate the move into an extra dimension.

4.2 Methods

4.2.1 Two-dimensional environments

The environment is defined as a square with sides of length S, and therefore if there are n patches within the environment, there is a patch density of $\frac{n}{S^2}$. It follows that the expected number of patches in an area K is $\frac{Kn}{S^2}$. This equates to a total nutrient content (in terms of marginal benefit to the plant) of Kx, where $x = \frac{np}{S^2}$ is the nutrient density within the environment.

As with the one-dimensional model (see Chapter 3), two types of environment are defined according to methods of resource distribution: “random” and “heterogeneous”.

Random nutrient distribution

In two dimensions, the random environments (see Fig 4.1) are defined in much the same way as within the one-dimensional model, with each patch being independantly distributed uniformly randomly within the environment. Whilst in one-dimension this means the location of each patch is determined by one sample from a uniform random distribution along the range of the environment, in two dimensions the analagous approach is to select two samples from a uniform random distribution to determine the x- and y-coordinates for each patch. As with the one-dimensional environments, these random environments see no statistical relationship between the locations of one patch and another (that is, the presence of one patch provides no information about the distribution of the other patches within the environment).

Heterogeneous nutrient distribution

Patchy heterogeneous environments (Fig again use the Pareto distribution described in Chapter 3 to sample inter-patch path lengths for a random walk. To account for the scaling up to two dimensions, the method used for the one-dimensional environments needs to be exapanded to account for filling a two-dimensional environment. This is achieved by for each path segment choosing not just a path length, but also a new orientation. By sampling a radian angle from a uniform [0,2π] distribution, a new orientation is sampled for each path step.

Plant distribution

In two dimensions, plants are independently distributed uniformly randomly throughout the environment. This is implemented in the same way as when creating a random

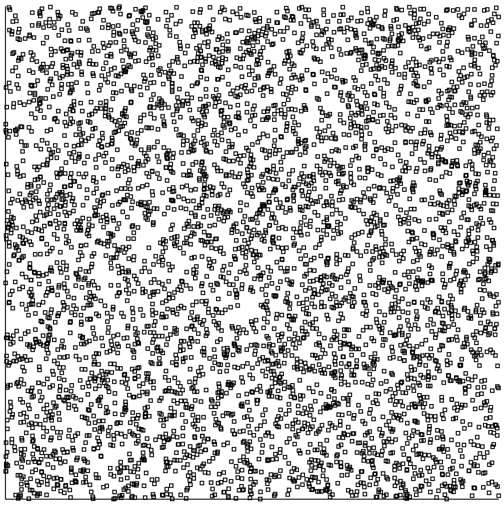

Figure 4.1: *Example of a "random" 2-dimensional environment. Here 5000 patches are independently distributed uniformly randomly across the environment, resulting in no statistical aggregation or spatial structure of patch locations.*

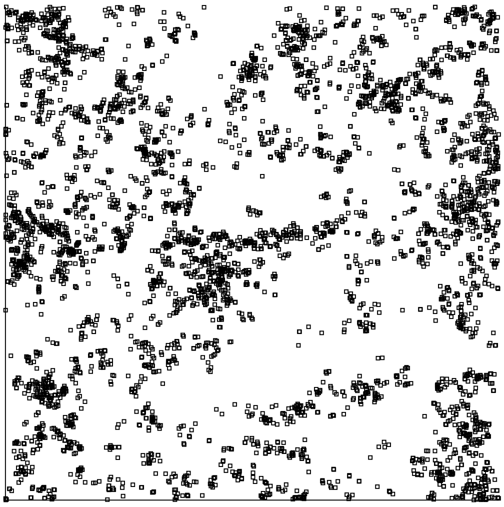

Figure 4.2: *Example of a "patchy heterogenious" 2-dimensional environment. Here 5000 patches are placed by starting in a random location and performing a random walk with step lengths sampled from a Pareto distribution. Between each path step, the direction of travel is chosen from a uniform distribution.*

nutrient distribution, with two samples taken from a uniform random distribution to determine the x- and y-coordinates for each plant.

4.2.2 Plant growth

As in Chapter 3, in the two-dimensional model intrinsic growth is of constant rate, g, until an upper size limit ($L_{\max0}$) is reached, that is

$$\frac{dL}{dt} = g \tag{4.1}$$

whilst $t < L_{\max0}/g$. Whilst in one dimension "size" is equivalent to length, in two dimensions it is taken as equivalent to area. Whilst the Gompertz growth model (of which the constant growth model is a simplification; see Sections 2.2.2 and 2.2.7) is often used to describe plant weight (Purves and Law 2002; Schneider et al. 2006; Lv et al. 2008), it is argued that is justifiable to apply the measure of "size" to area in this model. There is no implicit or implied third dimension in the two-dimensional model, therefore just as in one-dimension length, and in three-dimensions weight($\approx$volume), give a full measure of the plant's extent, so too in two dimensions does its area. Growth within the model is defined such that area grows at a constant rate, and so accordingly the rate of growth radially from the plant's centre decreases as the plant gets bigger. Assuming there is no proliferation bias (i.e. $\alpha = 0$, and so growth is uniform), radial size (L_{rad}) grows according to

$$\frac{dL_{\text{rad}}}{dt} = \frac{g}{2\pi L_{\text{rad}}}. \tag{4.2}$$

Patch acquisition again results in an instantaneous jump, p, in an individual's size (the marginal benefit to a plant from patch acquisition), and increases the upper size limit ($L_{\max}(t)$) by the same amount, so

$$L_{\max}(t) = L_{\max0} + pN(t), \tag{4.3}$$

where $N(t)$ is the number of patches acquired by the plant by time t.

4.2.3 Nutrient patch acquisition

When calculating the expected size of an individual, it is essential to factor in that the acquisition of a patch causes the plant to increase in size, and so acquiring a patch increases the size of the area explored and accordingly the expected number of patches found. This in turn results in more expected growth, leading to more patch acquisition, and so on. This becomes an infinite iterative-cycle, and the expected size of a plant at a given time can be expressed as an infinite sum.

To formulate this expression for a plant at time t (assuming $t < L_{\mathrm{max0}}/g$), we must first consider its size due to intrinsic growth:

$$L(t)_{\mathrm{intr}} = gt. \tag{4.4}$$

From above, it is expected that from this growth the plant receives a marginal benefit equal to gtx, and so growth due to intrinsic growth plus first order patch acquisition is given by

$$E(L(t)_{\mathrm{intr+O(1)}}) = gt + gtx. \tag{4.5}$$

However, this additional growth from first order patch acquisition in turn leads to more expected patch acquisition, and resulting growth due to second order patch acquisition equal to $(gtx)x$, making the expected size due to intrinsic growth plus first and second order patch acquisition

$$E(L(t)_{\mathrm{intr+O(1)+O(2)}}) = gt + gtx + gtx^2. \tag{4.6}$$

Continuing this process iteratively, the expression for the expected size of the individual due to intrinsic growth plus first to i^{th} order patch acquisition becomes

$$E(L(t)_{\mathrm{intr+O(1)+O(2)+...+O(i)}}) = gt + gtx + gtx^2 + ... + gtx^i, \tag{4.7}$$

which when expanded for infinite order patch acquisition results in the expression

$$E(L(t)) = gt\sum_{i=0}^{\infty} x^i. \tag{4.8}$$

Provided $x < 1$, the infinite sum in Eqn. 4.8 is equal to $\frac{1}{1-x}$, and so

$$E(L(t)) = \frac{gt}{1-x}. \tag{4.9}$$

Within the experiments described here, the nutrient content x (in terms of marginal benefit to the plant) is defined such that the expected size of a plant growing in isolation in a homogeneous environment is twice that of when grown with no patches. That is

$$\frac{L_{\mathrm{max0}}}{1-x} = 2L_{\mathrm{max0}}, \tag{4.10}$$

which gives a value of $x = 0.5$. L_{max0}, as in Chapter 3, is fixed at 0.2, and so expected size in control conditions is 0.4. This value lies within the mid-range of values tested in Chapter 3 and, along with a fixed number of 25 patches per plant, has been chosen as it allows for enough resources in the environment and marginal benefit from patch acquisition to have a significant effect on the experiments, without being so small so as to have no impact nor so large so as to dominate results and hide potential subtleties in results caused by other factors.

As in one-dimension, a mechanism is built into the model which allows a plant to potentially respond to its environment and influence root proliferation in response to its last acquired patch by means of a parameter α. In one dimension, growth is limited to two directions, therefore the distribution of growth can only be a weighting between these two directions. In two dimensions, however, there are an infinite number of directions within the plane for growth, and so any directional bias of growth can have an arbitrary degree of fidelity. In order to keep the model both relatively simple as well as analagous to the one-dimensional model, growth is now split into four “directions”, doubling the fidelity with the doubling of dimensions. This is achieved by splitting the plant’s radial growth into four sectors, each covering a $\pi/2$ radian range from the plant's centre (Fig 4.3(a)), with each plant being randomly oriented by an angle θ radians (chosen from a uniform [0,$\pi/2$] distribution; Fig 4.3(b)) so as to avoid any artifacts from enforced plant alignment. Now when a plant encounters a patch, the segment where the patch was found becomes the "direction" for any (potential) bias in growth (see Fig 4.3(c)). Whilst in reality plant exhibit a much higher fidelity of root response (Hodge et al. 1999, 2009; Hodge 2009; Cahill and McNickle 2011) it is argued that this approach is a sensible one to take. As stated above, this if for reasons of simplicity and keeping the two-dimensional model analagous to the one-dimensional model (Chaper 3). Also, experiments have been carried out investigating the response of root systems on the same sectoral basis (Campbell et al. 1991; Gao et al. 2012).

This has prompted a reworking of the proliferation bias value α. In one dimension, the value of α (-1:1) pertained to a linear shift of growth bias in the direction of the most recently found patch, with a value of 1 meaning a 100% bias of growth in the direction of the last found patch, a value of -1 meaning a 100% bias of growth in the opposite direction, and a value of 0 representing 50% of growth in each direction (i.e. no response to patch acquisition). When the growth bias relates to four sectors rather than two, an α value of -1 and 1 still results in 0 or 100%, respectively, of all growth being in the target direction. Now, however, a value of 0 must relate to 25% of total growth (that is, no growth bias in the target sector). In order to satisfy these conditions, a quadratic equation fitting the three points is solved giving the formula

$$\text{secgrowth}_{\text{target}} = 0.25\alpha^2 + 0.5\alpha + 0.25, \tag{4.11}$$

$$\text{secgrowth}_{\text{other}} = (1 - \text{secgrowth}_{\text{target}})/3, \tag{4.12}$$

where $\text{secgrowth}_{\text{target}}$ and $\text{secgrowth}_{\text{other}}$ are the growth ratios within the "target sector" (sector where the most recently acquired patch was located) and each of the remaining three sectors, respectively. This means that as the value of α changes, the percentage of overall plant growth in the target sector does not change linearly. For example, an α value of 0.4 in the one-dimensional model would see 70% of growth in

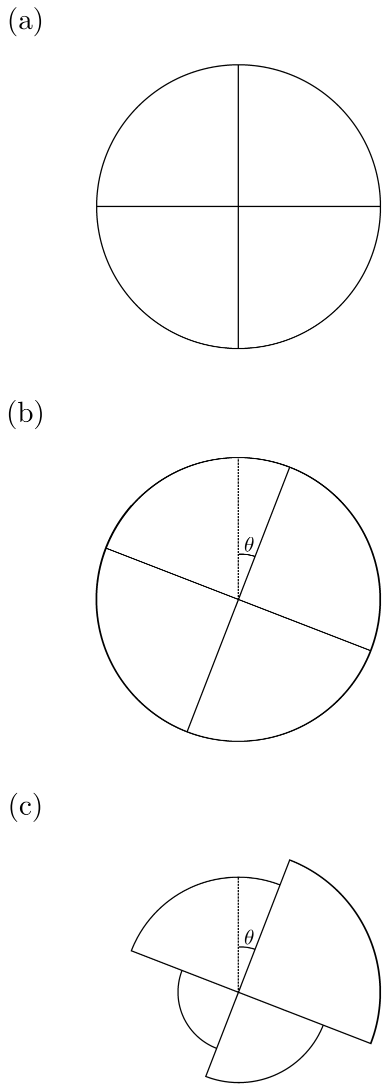

Figure 4.3: *(a) Plant split into four* $\pi/2$ *radian sectors and (b) rotated by angle* θ *(in radians), where* θ *is chosen from a uniform* $[0,\pi/2]$ *distribution. (c) Figurative example of how the individual in Fig 4.3(b) may look after experiencing growth and nutrient acquisition with a non-zero proliferation bias in response to nutrient patch acquisition. The proliferation bias leads to unequal growth in each of the four sectors of the root system.*

the target direction and 30% in the other direction. In the two-dimensional model, this value would result in 49% of the growth in the target sector, and 17% of the growth in each of the remaining three sectors. When α is not 0, growth in each of the four sectors is calculated using Eqn. 4.2 in conjunction with the individual sector radii and the percentage of growth assigned to each sector.

4.2.4 Experimental set up

Whilst in Chapter 3 the one-dimensional model was run within a genetic algorithm (GA) to assess the relative benefit of different growth strategies within an evolutionary context, in this chapter the two-dimensional development of the model is not run within a GA framework. Instead, simulations are run (again within the MATLAB computing environment) for a number of repetitions for the different configurations and the results analysed. This means that results returned will not be like for like with those in Chapter 3, and thus some caution must be displayed when comparing the results. It does, however, significantly speed up the process of running the simulations, and allows for the direct comparison of results for different strategies in terms of overall performance. Whilst in Chapter 3 the relative benefits and/or cost of different strategies are established in an evolutionary context, the simulations as performed in this chapter will still allow for the qualitative comparison of the relative merits of different growth strategies, and for comparison to results from the one-dimensional model.

4.3 Results

Figures 4.4-4.6 show the results for control tests (individuals grown in isolation), single type competition (competition with plants of same type, i.e. monocultures) and mixed competition (competition within a population between both types of plants), respectively. In control tests, the results given are for all individuals across all repetitions, whilst the competition results (monocultures and mixed) show the results for

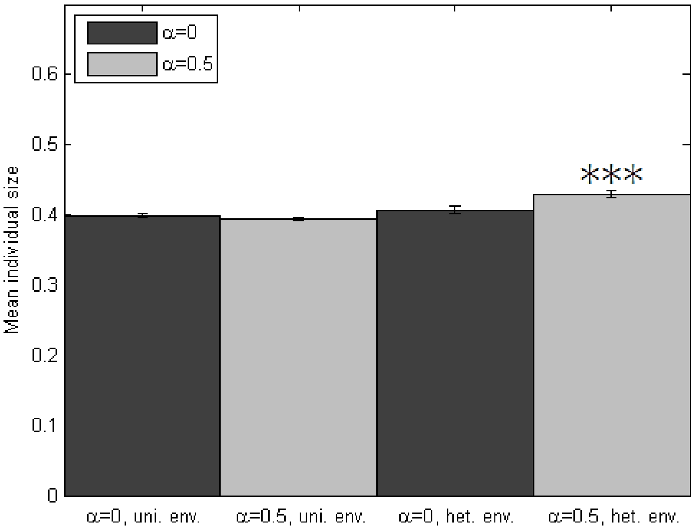

Figure 4.4: *Plant sizes from control (individual plants grown in isolation) tests. Results shown for plants with no ($\alpha = 0$; dark grey bars) and positive ($\alpha = 0.5$; light grey bars) proliferation bias, in uniformly random ("uni") and patchy heterogeneous ("het") environments ("env"). Error bars show +/- standard error of the mean. * denote significant difference from expected performace in control conditions in uniform environments (i.e. size 0.4), with * implying* $0.05 \geq P$, ** $0.01 \geq P$ *and* *** $0.005 \geq P$.

the populations across all repetitions. Standard errors of the mean are shown for all results. In Figures 4.4 and 4.5, significant difference from the expected peformance in control conditions in uniform environments (i.e. individual size of 0.4) is established by a t-test across all individuals from all replications, with significant difference indicated by the standard * notation. In Figure 4.6, a paired t-test was performed between the means of the two plant types over each replication of mixed competition, with significant difference again indicated by the standard * notation.

In control tests (individual plants grown in isolation in an individually generated environment) the results (Fig 4.4) are consistent with the trends seen in the results

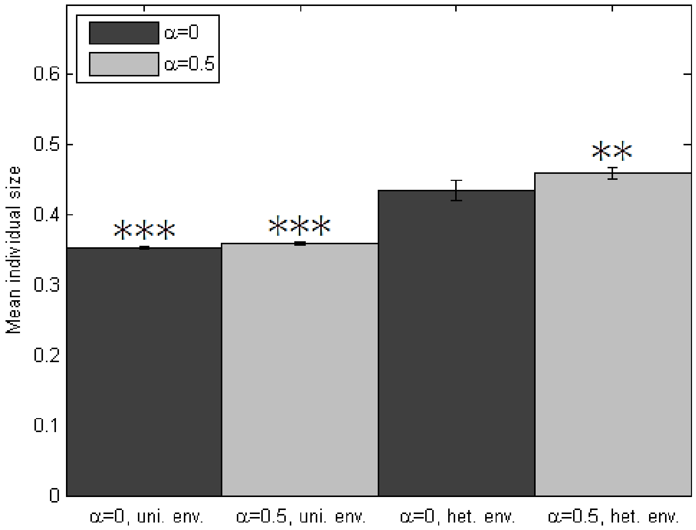

Figure 4.5: *Plant sizes from competitive (multiple plants grown together) tests with plants grown as monocultures. Results shown for monocultures with no ($\alpha = 0$; dark grey bars) and positive ($\alpha = 0.5$; light grey bars) proliferation bias, in uniformly random ("uni") and patchy heterogeneous ("het") environments ("env"). Error bars show +/- standard error of the mean. * denote significant difference from expected performace in control conditions in uniform environments (i.e. size 0.4), with * implying* $0.05 \geq P$, **** $0.01 \geq P$ *and **** $0.005 \geq P$.

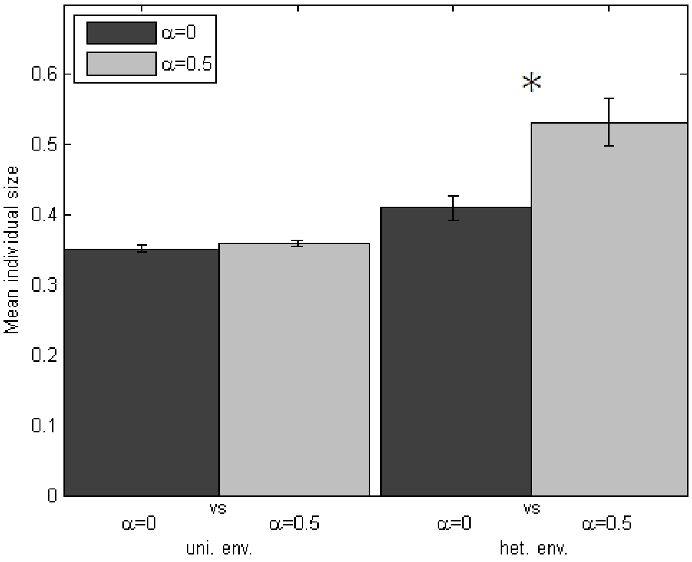

Figure 4.6: *Plant sizes from competivitive (multiple plants grown together) tests with mixed plant types. Results shown for competition between plant types with no* ($\alpha = 0$; *dark grey bars) and positive* ($\alpha = 0.5$; *light grey bars) proliferation bias, in uniformly random ("uni"; left two bars) and patchy heterogeneous ("het"; right two bars) environments ("env"). Error bars show +/- standard error of the mean. * denote significant difference in performance between plants grown with no proliferation bias and those grown with positive proliferation bias in mixed competition, with * implying* $0.05 \geq P$, *** $0.01 \geq P$ *and* *** $0.005 \geq P$.

from the one-dimensional model, though some subtleties emerge.

As was expected, and in line with previous results, in uniformly random environments no difference in performance was observed between plants with a positive proliferation bias and those with no bias. Both types of plant achieved an average size close to 0.4, in line with the chosen value of x (see Methods). When control tests were carried out in patchy heterogeneous environments, plants with no bias performed no better/worse than when in uniformly random environments. Plants with a positive bias performed slightly better than the expected performance in a uniform environment.

When competing as monocultures (Fig 4.5), both plants with a positive proliferation bias and those with no bias performed worse than in the control tests in uniformly random environments (with no difference between the two plant types). Both performed better than in the controls tests in patchy heterogeneous environments, with the positive bias plants performing better than the non-bias plants.

When both types of plants were competing with each other in mixed competition (Fig 4.6), in the uniformly random environments there was no difference in performance between the two (and no difference to when each plant type were grown as monocultures, see Fig 4.5). In patchy heterogeneous environments, the positive bias plants had an advantage over the non-bias plants. Compared to results when competing as monocultures, the non-bias plants performed less well, and the positive bias plants performed better.

4.4 Discussion

A key result to emerge from these tests is that competing populations (regardless of proliferation strategy) on the whole perform better within heterogeneous environments than in random environments. Whilst this increase in performance gain for plants with a positive proliferation bias can be explained by them maximising growth based on implicit and explicit information about their neighbourhood (see Chapter 3), that populations consisting solely of non-proliferation biasing plants see a similar (albeit slightly smaller) improvement in performance is more difficult to explain and in contrast to some observed behaviour. For example, Day et al. (2003) grew populations in environments with vary levels and scale of heterogeneity, and it was observed that population level yields when grown with the same total levels of nutrient supply were similar irrespective of the distribution of nutrient supply. It is perhaps worth noting that these "heterogeneous" environments were composed of chequerboard distributions of nutrient supply; whilst offering statistical levels and scales of heterogeneity, such environments are some way from the heterogeneity created in these tests and witnessed in reality. Conversely, Hutchings and Wijesinghe (2008) observed that population yields were highly context sensitive, with resource distribution having a distinct effect on overall yield. Again distibuting resources in a chequerboard formation, at the largest concentration contrasts between patches, overall yield in the heterogeneous tests were significantly greater than in homogeneous environemtns, with the greatest net yield observed when patches were large and with the greatest contrast between patches (Hutchings and Wijesinghe 2008).

A population of non-biased proliferating plants performing less well in random environments than control tests (in either environment type) can be attributed to the effects of overlap, and the resulting exploration of previously exploited soil. Such a negative effect from competition due to overlap fits with the observed tendency for plants to avoid the roots of competing plants in favour of segregation (Schenk, Callaway Mahall 1999). But the increase in performance as a population in the patchy heterogeneous environments (relative both to competition in random environments and, more importantly, control tests in both environment types) shows that the dynamics of plant interaction is causing the population to perform better as a whole. Given the idealised nature of the model, it is wholly plausible that such a phenomenon could have could occur, or at least have some impact, in real life.

With stochastic environments and interactions, it is possible (and indeed likely given enough repetitions) that "extreme" events may occur within the model. Specifically, there is a possibility that a very high quantity of patches may be located within a very small area (i.e. a localised high concentration of nutrient resource exists). It is worth noting that in reality "extremes" can occur and in a population is not unusual to see a few individuals perform significantly better than the rest of the population, even within experiments performed in homogeneous conditions (Schneider et al. 2006; Lv et al. 2008). The statistical testing confirm that any observed results can be verified as significant and are not simply skewed by anomalous results.

4.5 Conclusion

The results generated by the two-dimensional model are consistent with those seen for the one-dimensional model in Chapter 3, though subtle differences exist. The results in Chapter 3 demonstrated the different selection pressures for a positive proliferation response by means of running the model within a GA framework. In contrast, here the results demonstrate the relative benefit of such strategies via direct comparison of performance. However, both sets of results confirm the relative advantage of possessing such a trait under different circumstance.

Whilst competition itself causes a slight advantage to those with the ability to proliferate in response to the environment, it is the environment itself which seems to be the biggest factor in determining the benefit of this trait. Like the results in Chapter 3, it can be seen that when competition and heterogeneous environments are combined that the biggest advantage in posessing a proliferation response is observed. Again a compound effect of environmental and neighbourhood conditions provided the strongest selection pressure for such a strategy. This is consistent with experimental observations (Cahill et al. 2010).

Some of the trends emerging in the results presented here that were not evident in the results presented in Chapter 3 do not so much suggest that different behaviour is being observed, but rather can be attributed to fact that the tests themselves are not the same; in Chapter 3 the model was run within a GA and long term trends observed. In these tests, performance is observed as a series of direct comparisons of relative achieved size. The slightly stronger emphasis on the impact of competition within the one dimensional model is probably explained by the fact that a near neighbour in one-dimension has an effect of effectively closing off 50% of the possible area to be grown into. In two dimensions, a similar level of "crowding" would require a greater number of individuals located in close proximity to the target plant. In two-dimensions competition still provides an incentive to adjust proliferation in order to maximise resource yield per unit growth. However, the probability of such extreme crowding as in the one-dimensional model is much less, and also the ability to "out-compete" a neighbour within two dimensions is much greater than in one; instead of needing to "leap-frog" the competitor in order to beat them to resource located beyond them, in two-dimensions it possible to grow "around" a competitor.

Although there are subtle differences between the results, qualitatively the same trends are seen to emerge. These consistencies point to the fact that the one-dimensional model did a good job of qualitatively representing what occurs when the model scaled up to two-dimensions. The emergence of a significant benefit to all plant types when grown as a population (either as a monoculture or in mixed competition) within the patchy heterogenious environments relative to the expected (and observed) behaviour in uniformly random environments (both in control conditions and as competing populations) highlights the subtle results that can emerge from what remains a relatively simple model. These results are not easy to explain, but once again demonstrate that mean-field approximations can be of little relevance when dealing with stochastic, dynamic models (and as a consequence, real life).

Chapter 5

Different strategies: trading root growth rate and size against root system efficiency

5.1 Introduction

As discussed in Chapter 4, whilst the one dimensional model developed in Chapters 2 and 3 offers a good approximation of the qualitative results of the two dimensional model when investigating simple proliferation responses, its strict spatial constraints limit the growth (and subsequently emergent behaviour) of plants in a way that is not the case in higher dimensions. The move to two dimensions allows more flexability for plants to successfully grow and capture resouces in the presence of neighbourhood competition, and so should better allow different growth strategies to prove their relative values under different conditions.

It is well known that different plant species display markedly different root growth strategies (Malamy 2005; Hodge et al. 2009; Taub et al. 1996), and also high levels of architectural flexibility in root deployment (de Kroon et al. 2009; Malamy 2005; Hodge et al. 2009). In Chapters 3 and 4 the relative benefit of a plant's ability to focus root growth in response to the detection of available resources in one and two dimensions, respectively, was investigated. In this chapter, different plant types are specified by varying intrinsic growth rates and initial maximum sizes, and are balanced by limiting the different plants' abilities to successfully capture encountered nutrient patches. These different strategies are compared in control conditions (plants grown in isolation), competition when grown as a monoculture (a population of the same plant type grown simultaneously within the same environment) and in mixed competition (populations made up of two different plant types growing in competition). The relative effects of the positive proliferation responses to acquired patches seen in Chapters 3 and 4 are also investigated.

By defining plants in this way, fundamentally different growth strategies are evaluated across a range of conditions. It has been proposed (Campbell et al. 1991) that different plants exhibit a trade-off between scale and precision when it comes to foraging for resources. Campbell et al. (1991) suggest that plants can be loosely categorised as “dominants” and “subordinates”, with the dominant plant types displaying large scale at the expense of precision, whilst subordinates demonstrate greater foraging precision at a smaller scale. It is hypothesised that this balance of trading one trait off against another contributes towards diversity in community structure within populations of competing plants.

The idea of a scale-precision trade-off is not wholly accepted however, with anlyses of published data suggesting there exists no empirical evidence to support such a claim (Kembel and Cahill 2005; Kembel et al. 2008). Kembel et al. (2008) propose that scale and precision need to not be looked at in isolation, but rather to be considered as part of a broader range of traits, finding that “root foraging precision seems to form part of a suite of traits related to rapid growth rate and the fast set of resource-economic leaf and root traits”. There is also evidence to suggest that many observations of apparent scale-precision trade-offs are in fact an artefact of experimental practice and the time periods over which observations are made (Kembel et al. 2008; Fransen et al. 1999; Berendse et al. 2007).

Due to the nature of model developed in the previous chapters and adopted here, by varying growth rates and the ability of an individual to uptake available resources there is some freedom in interpreting what these traits represent in a “real” context. Whilst increasing growth rate and size is fairly self explanatory, the relative ability of an individual to successfully acquire available resources (see Section 5.2 for details of implementation) could represent precision as described by Campbell et al. (1991), but could equally represent properties of individual root function (Waisel and Eshel 1992; Hishi 2007) within the root system or uptake capacity of the system as a whole.

5.2 Methods

The core model was carried over from Chapter 4, with the environments (random and heterogeneous) and plant placement defined in the same way (see Section 4.2.1). Proliferation strategies were implemented in the same way, with the plants’ root systems (and therefore growth) split into four sectors, and the ability to focus growth in the sector of the last acquired patch (see Section 4.2.2).

5.2.1 Different growth strategies and root system efficiency

Different plant types and growth strategies were implemented by introducing different growth rates (g) and initial maximum sizes ($L_{\max 0}$). In order to impose some sort of compromise on a strategy favouring quicker growth to a larger size, the new feature added to the model is different qualities of root system, or a semblance of root efficiency. This in implemented as a probabilistic success rate of acquiring the encountered nutrient patches. Each plant’s root system has a parameter f which is the probability that a patch encountered by the plant is captured. The average of this behaviour is equivalent to deterministically equating f to the proportion the resources encounted that are captured, and so there is some flexability in how to interpret the “efficiency” as it is implemented.

The value of f can be fixed, or can decay with time (linearly at rate k). By combining this with different growth rates, different types of root system can be represented. For example, a quicker growing plant with a lower f value (and/or a higher k value) could be regarded as a “woody” plant; fast growing and covering a large area, but not particularly good at acquiring all avilable resources (a pseudo-“dominant” plant in the categorisation of Campbell et al. (1991)). Conversely, a slower growing plant with a higher f value (and/or lower k value) could be seen as a “weedy” plant; slow growing and covering a less expansive area, but better equipped for depleting this area of available resources (a pseudo-“subordinate” plant in the categorisation of Campbell et al. (1991)). Without any real-world parametrisation to determine relative growth rates, it is deemed unhelpful to relate the different plant types to specific real-world counterparts. However, by offering these contrasting growth strategies, it is possible to see how different trade-offs reward the plant under different circumstances. The intial value of f is denoted f_0, and so over time the value of f is given by

$$f(t) = f_0 - kt, \tag{5.1}$$

provided $t < \frac{f_0}{t}$. If $t \geq \frac{f_0}{t}$, then $f(t) = 0$. Following the same derivation as for Eqn. 4.8, it follows that if $k = 0$ then

$$L(t) = \frac{gt}{1 - f_0 x}. \tag{5.2}$$

If $k \geq 0$ then

$$L(t) = \frac{g \log(1 - f_0 x + ktx)}{kx} - \frac{g \log(1 - f_0 x)}{kx}. \tag{5.3}$$

5.2.2 Different plant types

Four different plants types are considered: two quicker growing with lower root efficiency, and two slower growing with higher root efficiency. Within both of these pairs, one plant has constant root efficiency, whilst the other decays over time with an aver-

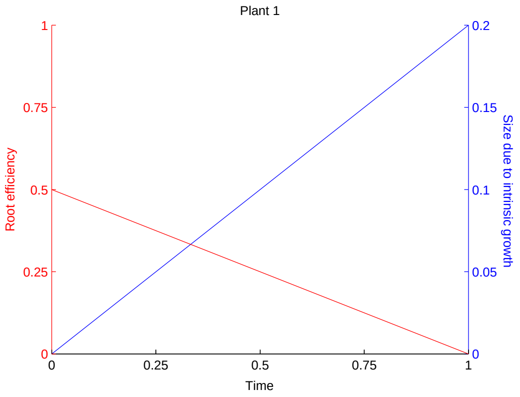

Figure 5.1: *Plant type 1 intrinsic (i.e. non-patch dependent) growth (blue) and root efficiency (red) against time.*

age efficiency equal to the constant variant. Figures 5.1-5.4 show the (deterministic) growth and root efficiency for the four different plant types over time.

5.2.3 Performance normalisation and experimental set up

In order to investigate the relative benefits and disadvantages of each growth strategy, it was necessary to some how "normalise" the performance so that under certain conditions performance was equal (or as close as possible). Since survival amongst competition is the driving force behind evolving strategies, it was decided to normalise performance across the different plant types by attempting to normalise their behaviour within monocultures (i.e. when grown in competition with only their own kind).

By running the model within MATLAB for each of these plant types as a monoculture within uniformly random environments, the environment was parametrised (by varying the total nutrient content of the environment, p_{tot}, and consequently the quality of each individual patch, p) so as to normalise population growth between the four types. Whilst they are not exactly the same (ignoring stochasticity, it is impossible to do so), they are close enough so as to provide a relatively equalised "control" performance to compare other results against (fig 5.5) (normalised value: $x = 0.825$).

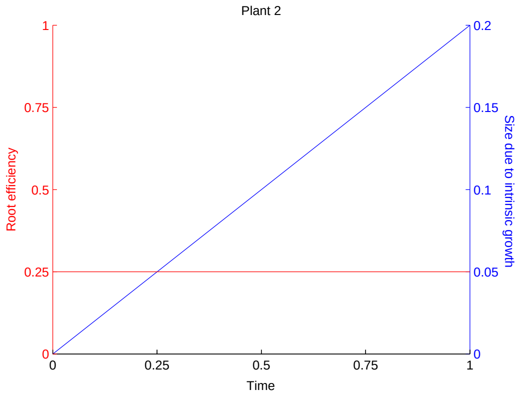

Figure 5.2: *Plant type 2 intrinsic (i.e. non-patch dependent) growth (blue) and root efficiency (red) against time.*

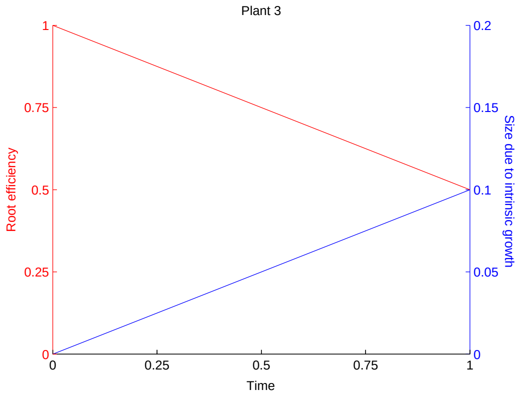

Figure 5.3: *Plant type 3 intrinsic (i.e. non-patch dependent) growth (blue) and root efficiency (red) against time.*

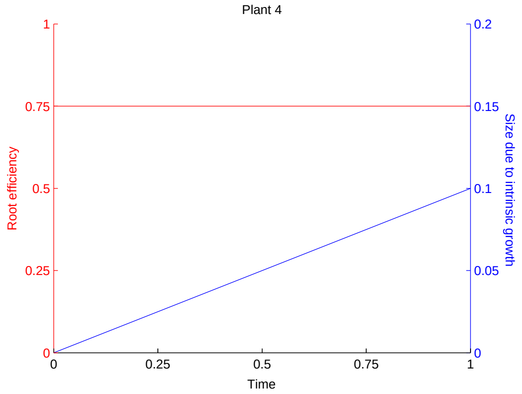

Figure 5.4: *Plant type 4 intrinsic (i.e. non-patch dependent) growth (blue) and root efficiency (red) against time.*

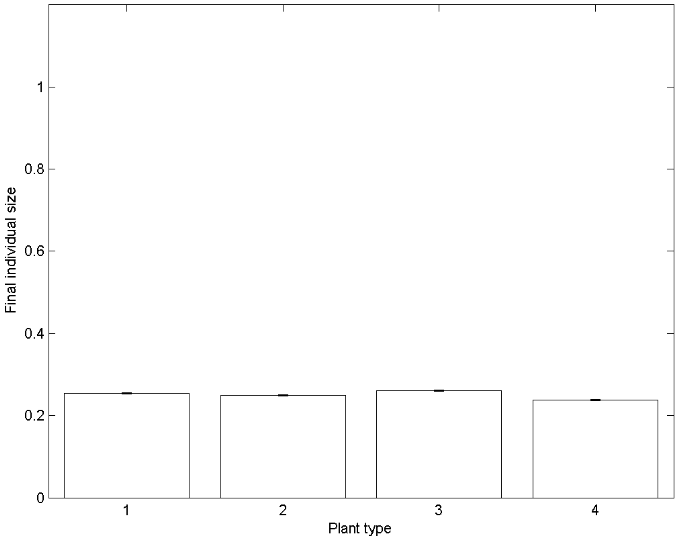

Figure 5.5: *Results for monocultures of each of the four defined plant types (with no proliferation bias) in uniformly random environments. This forms the "base" result for comparison to, and demonstrates that performance across the different plant types is relatively equal. Error bars denote +/- standard error of the means.*

The population was then split in two, with each half made up of one of the four plants types (including pairing plants with their own kind). Like in Chapter 4, simulations were run as independent iterations, and not within a genetic algorithm (GA) as in Chapter 3. Simulations were run in MATLAB for both uniformly random and patchy environments, for each permutation of either one, the other or both having a positive proliferation bias to the last found patch (resulting in 72 distinct combinations in total, since in monocultures there is symmetry). In total, 100 repetitions were then performed for each of these combinations, and the results compared to see how the different strategies fared in different conditions, and in which circumstances each plant type performed best and worst.

5.3 Results

The results for the different configurations of simulations are summarised in Figures 5.6-5.12. In each case, mean values of all individuals across all repetitions are shown, with error bars denoting the standard errors of the mean (across all individuals for control tests, and all populations for competition tests). In Figures 5.6 and 5.7, significant difference in performance by each plant type from its peformance in monoculture competition conditions in uniform environments (i.e. the initial normalised performance) is established by an unpaired t-test across all individuals from all replications, with significant difference indicated by the standard * notation. In Figures 5.8-5.12, an unpaired t-test was performed between all individuals across all replications of the paired results, with significant difference between the two again indicated by the standard * notation.

The results for each of the plant types when grown as monocultures (without any proliferation bias) in patchy heterogeneous environments are shown in figure 5.6. Compared to the corresponding results for uniform environments (Fig 5.5), it can be seen that performance is markedly different both between growth in the two types of environment, but also between the different plant types when grown in the patchy environments. The fast/large growing plants with lower root efficiency (plant types

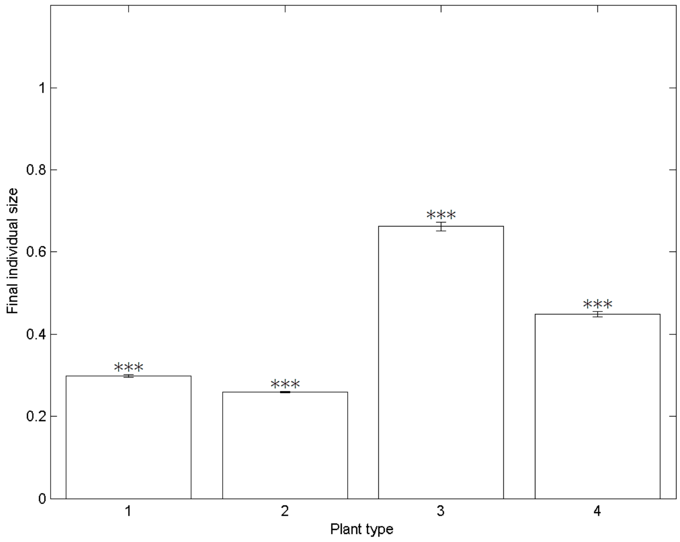

Figure 5.6: *Results for monocultures of each of the four defined plant types (with no proliferation bias) in patchy environments. Comparing to Fig 5.5, it is evident that the "equal" performance seen across the different plant types in uniformly random environments does not carry over to patchy environments. Error bars denote +/- standard error of the means. * denote significant difference in performance from when grown as monocultures in uniform environments (fig 5.5), with * implying* $0.05 \geq P$*, *** $0.01 \geq P$ *and **** $0.005 \geq P$*.*

1 and 2) see no or modest (though significant) improvements on the performance from uniform environments, whilst the slower/smaller growing plants with higher root efficiency (plant types 3 and 4) see large significant improvements in final size.

Within each of these pairs of plant types, those with initially higher root efficiency that then deteriorates (plant types 1 and 3) see a greater increase in performance than those with root efficiency fixed as the average of the former (plant types 2 and 4, respectively). Variation in results increases with average size.

Figures 5.7 and 5.8 show the results for plants grown in control tests (plants grown in isolation) in uniformly random and patchy heterogenious environments, re-

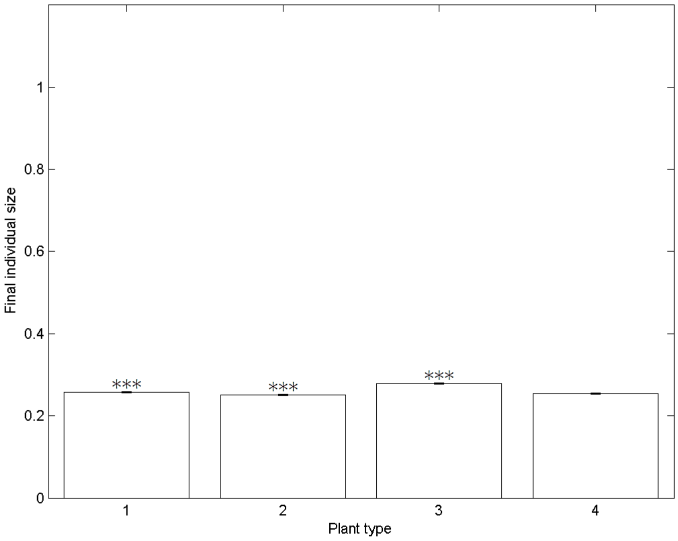

Figure 5.7: *Comparison of performance when each plant type is grown in control conditions in uniformly random environments with no proliferation bias ($\alpha = 0$). Error bars denote +/- standard error of the means. * denote significant difference in performance from when grown as monocultures in uniform environments (fig 5.5), with * implying* $0.05 \geq P$*, ** * $0.01 \geq P$ *and **** $0.005 \geq P$.

spectively. To check the effect of proliferation response, α, figure 5.8 compares the results from control tests within patchy environments when plants with no proliferation bias ($\alpha = 0$) and a positive proliferation bias ($\alpha = 0.5$). As discussed in Chapter 3, there can be no advantage to an individual plant grown in a uniform environment from any proliferation bias, so figure 5.7 only shows results with no proliferation bias ($\alpha = 0$).

In uniformly random environments there is a small but statistically signiicant decrease in performance for plant types 3 and 4 when grown as a monoculture compared to when grown in control tests (compare with Fig 5.6), with little or no difference for plant types 1 and 2. In contrast, in patchy environments (Fig 5.8) all plants types see a small increase in performance, with significant improvements when grown with

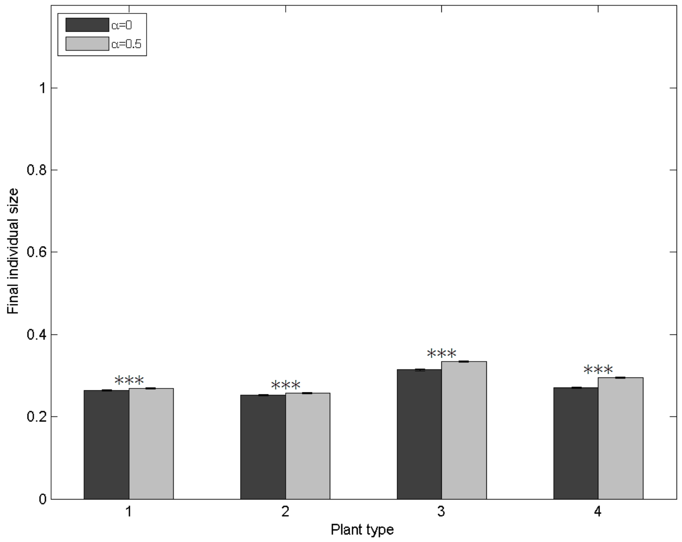

Figure 5.8: *Comparison of performance when each plant type is grown in control conditions in patchy environments with no (*$\alpha = 0$*; dark grey bars) and positive (*$\alpha = 0.5$*; light grey bars) proliferation bias. Error bars denote +/- standard error of the means. * denote significant difference in performance between plants grown with no proliferation bias and those grown with positive proliferation bias, with * implying* $0.05 \geq P$, **** $0.01 \geq P$ *and **** $0.005 \geq P$.

a positive proliferation value. Plant types 3 and 4 see a bigger increase than types 1 and 2.

In the control tests, plant types 1 and 2 see no real difference in performance between uniform and patchy environments, as expected, whilst types 3 and 4 see a small increase in performance when grown in patchy environments compared to in uniform environments.

Figures 5.9 and 5.10 compare the results for plants grown in monocultures with no proliferation bias ($\alpha = 0$) and a positive proliferation bias ($\alpha = 0.5$) in uniformly random and patchy environments, respectively. There is no difference in performance between the two strategies in uniform environments, and a very small improvement for all plant types in patchy environments, with only plant type 2 demonstrating a significant improvement.

Across all permutations of mixed competing plant types, in uniformly random and patchy environments, only very small patterns of response to proliferation bias (α) emerged (results not shown). Consequently the results for mixed competition are just shown for growth with no proliferation bias ($\alpha = 0$). The best and worst performance for each plant type when grown in uniform environments is shown in figure 5.11. The results show that there is very little (although often statistically significant) difference in performance across the four plant types, regardless of which other plant types they are grown with. The results remain consistent with those for when the different plant types were grown as monocultures (fig 5.5), i.e. performance is effectively equal across all four plant types, irrespective of competition type.

The best and worse performance for each plant type when grown in patchy environments are shown in figure 5.12. Compared to the results from the uniformly random environments (Fig 5.11), their is much greater (and strongly significant) variation between both the best performances of each of the plant types (as expected after the results when grown as monocultures in patchy environments), but the difference between the best and worst performances for the individual plant types is also greater.

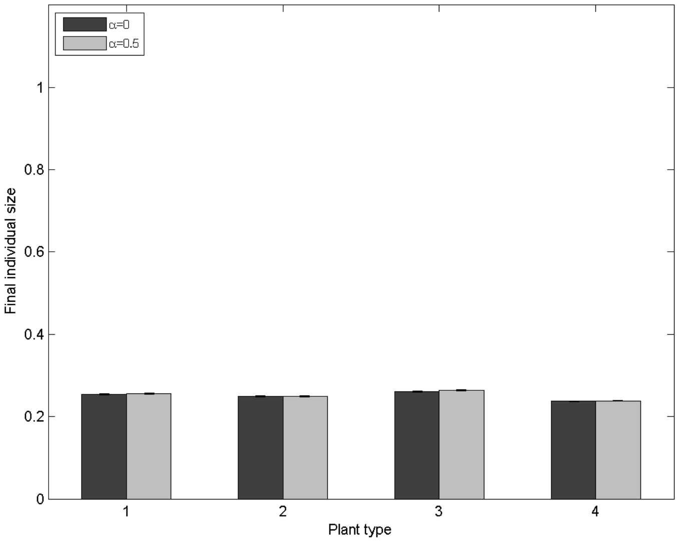

Figure 5.9: *Comparison of performance when each plant type is grown in competition as a monoculture in uniformly random environments with no proliferation bias* ($\alpha = 0$; *dark grey bars) and positive proliferation bias* ($\alpha = 0.5$; *light grey bars). Error bars denote +/- standard error of the means. No significant difference was observed in performance between plants grown with no proliferation bias and those grown with positive proliferation bias.*

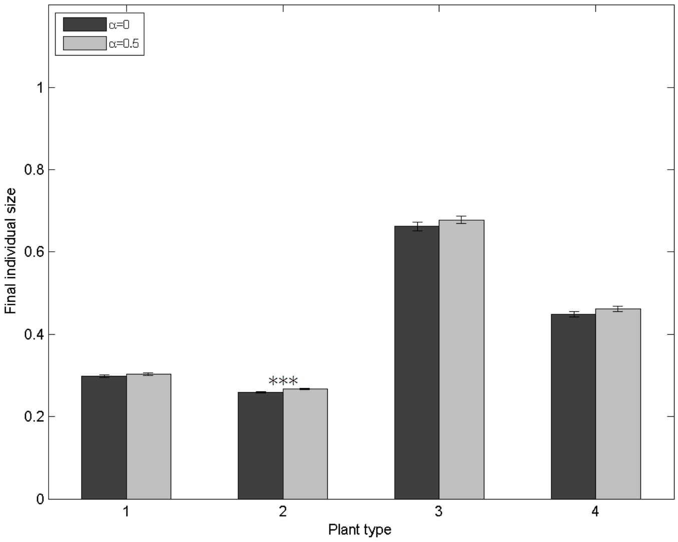

Figure 5.10: *Comparison of performance when each plant type is grown in competition as a monoculture in patchy environments with no proliferation bias* ($\alpha = 0$*; dark grey bars) and positive proliferation bias* ($\alpha = 0.5$*; light grey bars). Error bars denote +/- standard error of the means. * denote significant difference in performance between plants grown with no proliferation bias and those grown with positive proliferation bias, with * implying* $0.05 \geq P$*, *** $0.01 \geq P$ *and **** $0.005 \geq P$.

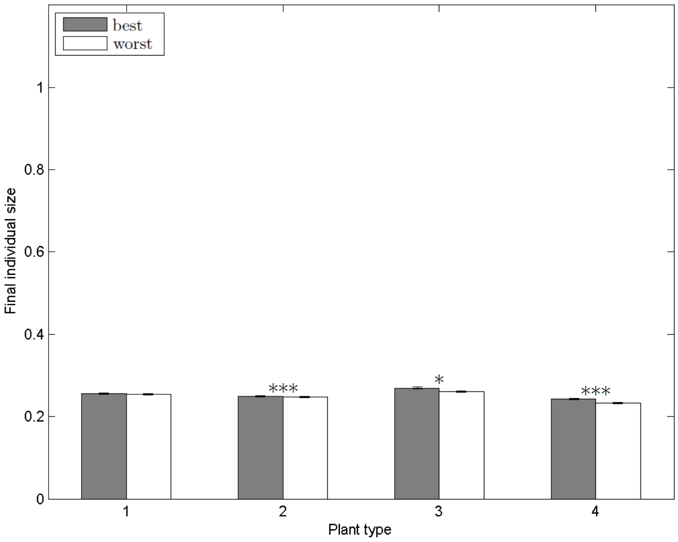

Figure 5.11: *Best (dark grey) and worst (light grey) performance for each plant type from across all of the competition tests in uniformly random environments. Error bars denote +/- standard error of the means. * denote significant difference in performance between best and worst performance in mixed competition, with * implying* $0.05 \geq P$*, *** $0.01 \geq P$ *and **** $0.005 \geq P$*.*

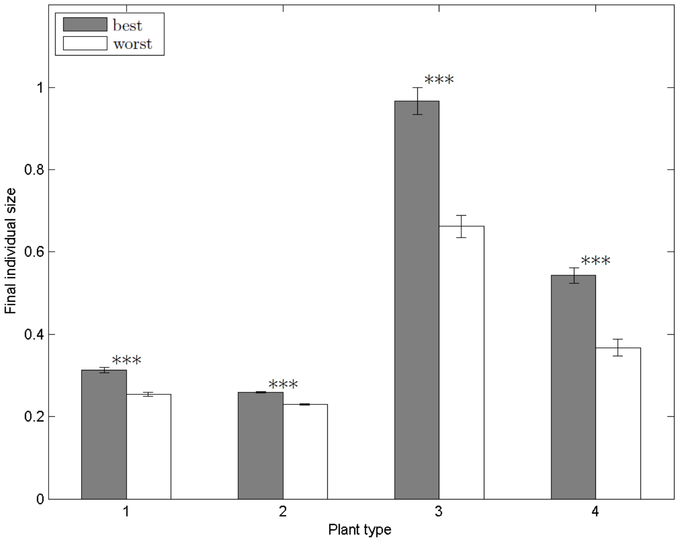

Figure 5.12: *Best (dark grey) and worst (light grey) performance for each plant type from across all of the competition tests in patchy environments. Error bars denote +/- standard error of the means. * denote significant difference in performance between best and worst performance in mixed competition, with * implying* $0.05 \geq P$*, *** $0.01 \geq P$ *and **** $0.005 \geq P$.

5.3.1 Summary of performance of different plant types

5.3.2 Plant 1

Plant type 1 sees a small improvement in performance in patchy environments compared to uniformly random environments when grown as a monoculture (Figs 5.5 and 5.6). It experiences little or no difference in performance regardless of environment type when it is grown in control conditions (Figs 5.7 and 5.8). In mixed competition with other plant types there is no real variation in performance in uniformly random environments (Fig 5.11). There is some variation in patchy environments (Fig 5.12), with best performance against plant type 2, and worst against type 3.

5.3.3 Plant 2

Plant type 2 typically sees little difference in performance between uniformly random and patchy environments, or when grown in control conditions or as a monoculture (Figs 5.5-5.10). There is very little difference in mixed competition regardless of competition in uniformly random environments (Figs 5.11), though significant (but still relatively small) differences in performance in patchy heterogeneous environments (Fig 5.12).

5.3.4 Plant 3

Plant type 3 sees a large improvement in performance in patchy environments compared to uniform environments when grown as a monoculture (Figs 5.5 and 5.6), and average plant performance when grown in a monoculture is slightly worse than control in uniform environments (Figs 5.5 and 5.7) and markedly better in patchy environments (Figs 5.6 and 5.8).

In mixed competition, there is little difference in performance regardless of competitor in uniform environments (Fig 5.11) whilst a large variation is seen in patchy environments (Fig 5.12), with best performance against plant type 1, and worst when grown as a monoculture (i.e. competing with itself).

5.3.5 Plant 4

Like plant type 3, when grown as a monoculture plant type 4 sees a large increase in performance when grown in a patchy environment comparded to a uniformly random environment (Figs 5.5 and 5.6), and comparing growth as a monoculture to control tests there is a slight decrease in performance in uniformly random environments (Fig 5.7)) and a large improvement in patchy environments (Fig 5.8). Again as with plant type 3, the introduction of mixed competition results in little difference in performance against all plant types in uniformly random environments (Fig 5.11) but a larger deviation in performance is seen in patchy environments (Fig 5.12), with, like plant type 3, best performance against type 1 and worst performance against type 3.

5.4 Discussion

Across all permutations of competing plant types, in uniformly random and patchy environments, no strong pattern of response to proliferation bias emerged (results not shown). This is likely due to the increased quality, in terms of marginal benefit to the plants, of the individual patches. Although a statistally significance difference in performance was demonstrated by plant type 2 when grown as a monoculture with and without a positive proliferation parameter (Fig 5.10), the actual performance difference is very small. Caution must be taken when interpreting results not to confuse statistical significance with importance. When running simulations, in principle it is possible for an arbitrarily large number of replicate “experiments” to be performed, and so using statistical tests to establish “significance” can be deceptive as with enough replicates an arbitrarily small P-value can be arrived at for effects which are of no practical significance (Currey et al. 2009).

The lack of a prominent benefit from the existence of a positive proliferation characteristic can potentially be explained by a number of factors. As with the results obtained in Chapter 4, the tests here are evaluating the relative performance of different individuals based on the outcome of experimental simulations, rather than as in Chapter 3 where the model is run within a genetic algorithm to obtain an "optimum" strategy based on evolutionary processes. It is perhaps worth noting that whilst there isn't a strong indication that a positive proliferation bias offers a significant advantage, in all like-for-like tests the plants possessing a positive proliferation bias perform at least as well as those that possess no bias, and often better (if not by a large amount). As such it would not be unreasonable to expect that were the model to be run within a GA like in Chapter 3, that ultimately a positive proliferation bias would emerge as a favourable strategy.

As discussed within Chapter 3, in one-dimension, the benefit of a positive proliferation bias is highly context sensitive, and it isn't simply the presense/absense of competition or the resource distribution that determines the level of benefit, but a hierarchical combination. The benefit of a positive proliferation response is also tightly linked to the quality of the individual patches. No benefit exists if the individual patches are of too high or too low a quality, and it is only in a "sweet spot" in between that a benefit was observed. The quality of the patches in these experiments (as determined by the balancing of the performance across the four different plant types when grown as monocultures in uniform environments) lies beyond this region, and there is a good chance that the same behaviour observed in the one-dimensional model applies here, and that the patch quality is beyond that for which a positive proliferation bias can yield a benefit.

The four different plant types show markedly different reactions to both environment and competition. At one extreme plant types 1 and 2 see only a small effect on performance from changes to either patch distribution or the construction of the neighbourhood competition, whilst at the other extreme plant type 3 shows a large positive performance response to a heterogeneous environment when grown as a monoculture, and a large increase in performance depending on the type of plants it is growing in competition with.

In between these extremes, plant type 4 shows a smaller response to environment type than plant type 3 when grown as a monoculture, but sees a similarly large positive effect from growing in competition with types 1 and 2, but a decrease in performance against type 3.

As described earlier, plant types 1 and 2, and types 3 and 4, can be paired up as sharing the same growth rates and average root efficiency over the growing period. Plants types 1 and 2 show a smaller reaction to environment type and less ability to exploit a patchy environment than types 3 and 4. However, they also show a high resilience to environmental and neighbourhood effects, and whilst they never perform as well as plant types 3 and 4 do in their ideal circumstances, they do not seem to suffer from the inclusion of higher performing plant types (3 and 4), and so seem quite capable of co-existing with neighbours who far better exploit the environment.

The lack of significant (in the practical rather than statistical sense) benefit to the individuals of reactive directional proliferation in response to acquired patches means that any semblance of benefit due to the "precision" of the plants must be implicit within the different root system efficiencies of the different plant types. If precision is considered in this way, then the above results are inconsistent with the broad scale findings of Grime et al. (1991), who found "results indicate that the most successful competitors under the productive conditions of the conventional competition experiment exploited patches more completely because of higher growth rates and larger size, not because of greater flexibility within the leaf canopy and root system", whilst they do potentially support the assertion that different strategies help to support co-existance within the population. However, in relating the results in this chapter to experimental results, it is worth emphasising that the different root system effiencies desribed for the different plant types potentially embody a lot more properties than simply scale and precision. These other properties implicitly contained within the model's simplified nature could well support the suggestion that scale and precision in isolation are not sufficient to desribe relative performance, and that rather they are part of a larger "suite" of important traits (Kembel and Cahill 2005; Kembel et al. 2008).

Plants types 3 and 4 experience high levels of performance in patchy environments when grown either as monocultures or in mixed competition. This ability to successfully exploit the environment when grown in competition highlights that one can't simply extrapolate average performance in such conditions from performance when grown in isolation. When grown as monocultures with no proliferation bias, some drop in average performance compared to when grown as isolated individuals is expected since overlap and competition (specifically depletion of resources by competitors) reduces the expected available resources available to an individual per unit growth of the root system. Consequently, as with the results seen in Chapter 4, the increased performance when grown in patchy heterogenious environments as monocultures (and to a lesser extent, in mixed competition) compared to in isolation is somewhat counter-intuitive.

5.5 Conclusion

By adding another level of complexity to the model that has been developed over the previous chapters (in this chapter, it is the addition of "root efficiency") results emerge that are both consistent with those seen previously, as well as offering new insight into the complex interactions that occur between competing plants.

The four different plant types defined in this chapter can each be categorised by two characteristics: a relatively high (plants types 1 and 2) or low (plant types 3 and 4) intrinsic growth rate (and, respectively, relatively large or small initial upper bounds, and relatively low or high root efficiency), and decreasing (plant types 1 and 3) or constant (plant types 2 and 4) root efficiency. With performance normalised for growth as monocultures in control conditions, the relative effects of different environment types and mixed competition are assessed for the different plant types.

Across all tests in uniformly random environments (controls, monocultures and mixed competition) there is little to separate the different plant types. Performance remains close between the different plant types when control tests are performed in patchy heterogeneous environments too, with plant types 3 and 4 seeing a small advantage over plant types 1 and 2. However, when grown as monocultures, this advantage is drastically increased. Although plant types 1 and 2 see a small increase in performance, plant type 4 nearly doubles its average size, and plant type 3 sees an even greater increase in performance. The increase in performance between environment types is consistent with that seen in Chapter 4, but the disparity of the levels of improvement reflect the intrinsic differences of the four plant types.

If it were known that more patches would be available to an individual, it would be expected that plant types 3 and 4 would perform better than plant types 1 and 2 since they have a higher root effiency (manifesting itself as a higher conversion rate of patch encounters into successful patch acquisitions). Yet it is also true that plant types 1 and 2 would expect to encounter more patches based on their higher intrinsic growth rate and initial maximum sizes. The “balancing” of these two factors are what allow all of the plants to perform relatively equally in uniform environments. In patchy heterogeneous environments, there is clearly a shift in balance favouring the improved acquisition rates over the initial higher encounter rates, with the compound effect of patch acquisition (see Chapter 4) ultimately outweighing the benefit of an increased initial encounter rate.

These differences in performance are further highlighted when the different plant types are grown in mixed competition in patchy heterogeneous environments. When grown in competition against plant type 1, plant type 3 achieves an average size nearly four times larger than when it is grown as a monoculture in a uniformly random environment. But conversely, in the same mixed competition, plant type 1 performs nearly identically to its performance as a monoculture in a uniformly random environment; the significant increase in performance by plant type 3 appears to not occur at the expense of plant type 1.

As with much of the results shown in the previous chapters, a key message to again take from these results are that the expected mean field performance cannot be used to predict the performance of the plants when grown under stochastic conditions. Whilst in chapter 4 it was seen that the plants perform better as monocultures in patchy heterogeneous environments compared to uniformly random environments, in this chapter it is shown that additonal properties of the plants that have little obvious effect on expected performance can compound these effects greatly.

Chapter 6

Concluding Remarks

6.1 Summary of work and results

Whilst a far from exhaustive review of the subject, Chapter 1 of this thesis gives some insight into the intricacies and complexities of the world of plant roots. After decades of experiments and observations, there remain many areas of plant growth and competition that remain far from fully understood. Whilst the models and results presented in Chapters 2-5 barely scratch the surface, it is believed that they do shed some light onto certain aspects of plant growth and competition, and they do offer some new insight.

Beyond the inherent difficulties involved in making measurements and observations of plant root activity (Cahill and McNickle 2011), a recurring issue is that of highly context sensitive responses (see for example: Hodge et al. 1999a; Robinson et al. 1999; Cahill et al. 2010; Wijesinghe et al. 2005). Whilst experiments performed under laboratory conditions can reveal interesting observations of behaviour, of what relevance they are when it comes to plants growing in nature is debatable.

The goal of Chapter 2 was to derive a “simple” model of plant growth and competition, with a mechanism for nutrient acquisition, which could be understood and taken forward to apply in different conditions. The starting point was the Gompertz growth function. This function is an experimentally justified choice for desribing generic plant growth both in isolation and in competition (Purves and Law 2002; Schneider et al. 2006; Lv et al. 2008), and can be considered to implicitly account for processes that occur in plant growth such as growth due to uptake of background nutrient resources and loss due to metabolism (Purves and Law 2002), removing the need for explicit modelling. The Gompertz fucntion not only implicitly captures many mechanism and features of plant growth, but when coupled with an explicit modelling of resource capture and resulting growth it scales to retain Gompertz-like behaviour. This idealised one-dimensional model, consisting of a deterministic growth function coupled to an explicit, stochastic model of nutrient resource distribution, can simulate the growth of a number of individuals growing and competing for a scarce, patchily distributed resource.

However, the model is computationally (relatively) slow, and analytically difficult to make progress with. A constant growth model provides a computationally and analytically simplified alternative that is demonstrated to accurately capture the behaviour of the Gompertz model, both at the population and individual level.

Although the aim in Chapter 2 was to provide a model to utilise in more intensive and complicated settings and scenarios, results of interest are apparent from the validation tests. Through comparisons of the non-overlapping and overlapping models, it can be seen that a strict mechanism of non-overlap between neighbouring plants can significantly improve average performance. Whilst this might not come as much of a surprise considering the propensity for the segregation of roots among many plant species (Schenk et al. 1999), it is surprising how effective it is at increasing overall yield when it is remembered that these enforced neighbour-neighbour boundaries constrain the growth of those within them, and potentially limit the ability of an individual to turn captured resources into resultant growth.

Similarly, the relative invariance of results to population size (apart from at extreme environmental parameter choices) suggests that a population of, say, ten plants can be successfully used to simulate the behaviour of a significantly larger population. The spatial limitations of the one-dimensional framework will no doubt help in contraining growth and behaviour, and may well contribute to this. Likewise, the "stochastic" distributions of plants and patches as defined by uniformly random distributions have limited ability to simulate consistantly aggregated resource distributions.

In Chapter 3 the contstant growth model derived previously is applied to patchy heterogeneous environments. The simulations are performed for generic plants with minimal assumptions about growth and fuction. A simple trait allowing for the proliferation in the direction of the last acquired patch allows for response to environmental information, and provides the plants with an ability to actively proliferate in response to local resources. Tested within the evolutionary context of genetic algorithm, the results support observations of a percieved hierarchy of influences on the benefit of proliferation response to resources (Cahill et al. 2010).

Whilst the results in Chapter 3 highlight that a simple model can generate interesting and intricate results, it can no be denied that the one-dimensional environments impose highly limiting restrictions on what can be done with the model. In Chapter 4, the model is expanded into two dimensions. The model largely carries over directly from the one-dimensional framework developed previously, with choices made to deliberately keep the expanded two-dimensional model as analogous as possible to the one-dimensional case in order to assess the effect of dimensionality. Overall consistent behaviour was observed with the one-dimensional model, though subtlies emerged in the results and the different ways in the which the models were assessed (within a GA in Chapter 3 to assertain evolutionary benefits of behaviours, and here as replicated experiments to assess relative performance) make it impossible to do direct like-for-like comparisons. Whilst the observed reduction in performance when grown in competition compared to in isolation when grown within uniformly random environments can easily be explained as due to overlap, the increase in performance when grown in patchy heterogeneous environments under competition compared to in isolation is less intuitive. As a population, the plants are better able to benefit from exploitation of the environment when resources are aggregated in distribution.

This can be understood when plants are able to respond to their environment and so focus root growth into areas of higher resource concentration, but that it occurs when plants grow passively with no response is a far from obvious or trivial result. Whilst a change in overall yield with resource distribution is not consistent with some experimental results (Day et al. 2003), it is with others (Hutchings and Wijesinghe 2008). With no proliferation response, this is an interesting result breaking from intuitively expected behaviour, and arguably all the more interesting that is generated by such a relatively simple model.

Similarly interest results emerge in Chapter 5, where different plant types with contrasting growth strategies are introduced by trading root growth rate (and size) against root system efficiency. With root system efficiency implemented as a probabalistic chance of patch acquisition given an encounter, there is some freedom as to how to interpret this in terms of real world behaviour. By normalising behaviour across monoculture tests in uniformly random environments, a “base” level is created where the relative strengths and weaknesses of the different plant types level out, and performance is very similar. Whilst little difference in performance persists across all tests in uniformly random environments, large differences in relative performance are observed in competition with patchy heterogeneous environments. When growing as a competing population (either as a monoculture or in mixed competition), the two plant types which favour root system efficiency over growth rate display significantly better performance than the faster growing counterparts. These results and their deviation from expected mean-field results are not easy to explain. Whilst they seem to be at odds with some observations (Grime et al. 1991), it is necessary to remember that in a model so idealised, a nunber of different facets of plant growth could be implicitly affected by changes to certain properties without being apparent. “Root system effiency” is easily defined and understood within the confines of the models, but when relating such a model to real life, it could cover a myriad of related factors.

The models and experiments developed and discussed in this thesis are highly conceptual, idealised representations of incredibly complicated real-world processes and interactions. Whilst it can easily be pointed out that in the whole thesis not a single real-world parameter is sourced or implemented, it would be arguably unfair to see this as a criticism. Whilst the visual representations of one- and two-dimensional root structures presented in this thesis are a long way from the complexity and realism seen generated by architectually explicit models (Jourdan and Rey 1997; Leitner et al. 2010), they allow for tests and comparisons to be performed that would be all but pointless with more complicated models. In order to assess the impacts of different strategies on performance and under different environmental and neighbourhood conditions, it is necessary to have a model flexible enough to be applied to all of these different scenarios.

The simple models presented here not only allow for the application to a range of different conditions, but also allow for results to be considered within the confines of a model whose behaviour is well understood. When counterintuitive behaviour emerges like that witnessed under competition within heterogeneous environments in Chapters 4 and 5, with a simple model it is much easier to decide if such behaviour is of interest and deserving of further investigation than if a model with a plethora of parameter choices and imposed assumptions (Roose and Schnepf 2008, Leitner et al. 2010) returns a similarly puzzling result. This is not to dismiss the merits of such complicated models, but rather to stress the continued importance and scope of simplifying approaches. The results mentioned above are not all easy to explain, but given the nature of the model which generated them, it is hard to dismiss them and not conclude that they could hold interesting insight into the underlying processes being modelled.

6.2 Further research

It is considered the work presented in this thesis could be taken forward in two distinct ways: applying the existing models to new applications, and developing the models further.

With some of the results generated thus far no easily understood, it would appear sensible to first take the existing model and apply it in new ways. Perhaps the most obvious example of this would be to take existing studies and attempt to replicate the experimental set ups and then compare results. One series of studies (Day et al. 2003; Wijesinghe et al. 2005; Hutchings and Wijesinghe 2008) studied the relationship between the spatial pattern of nutrient supply on yield and community structure. Spatial patterns of nutrient supply were created by grids of squares of different sizes and nutrient concentrations. Whilst obviously an experimental ideal rather than reflection of naturally occurring resource heterogeneity, such environments could easily be replicated within the two-dimensional model defined in Chapter 4. By defining resource and plant distributions in line with these studies, it may be possible to reveal and understand why behaviour evident in the results of Chapters 4 and 5 seem consistent with results observed by Hutchings and Wijesinghe (2008), but at odds with those by Day et al. (2003).

Another area that the model could be applied to investigate is the effects of patch sizes. Fitter (1994) defined a set of basic attributes to describe a patch, with attribute falling into two scale categories: spatial and temporal. The individual patches described in Chapters 2-5 have no physical size within the model, but they can be aggregated construct larger patches. The resource distributions described in this thesis typically fall under “uniformly random” and “patchy heterogeneous” classifications. By varying the distributions and number of different environment types and properties could be defined.

Finally, the temporal properties of patches could be investigated. This would require the introduction of patch turnover and duration, as well as creation. By combining such temporal properties along with the above mentioned spatial-scale properties, a large variety of resource patterns could be defined and tested within the model.

Appendix A

Results for different population sizes

In this appendix the full results from the simulations run in Chapter 2 are shown for different population sizes

Across the different models, results are shown for three different configurations of plant (P) and nutrient patch (N) distributions: regularly placed plants with uniformly random patches (rPuN), uniformly random plants with regular patches (uPrN), and uniformly random plants and patches (rPrN).

A.0.1 Gompertz growth models results

Non-overlapping Gompertz model

Fig A.1 shows the mean, variance and skew of the final plant sizes for repetitions of the non-overlapping Gompertz model under the three different configurations of plant and patch distribution for different choices of population size and number of patches (per plant).

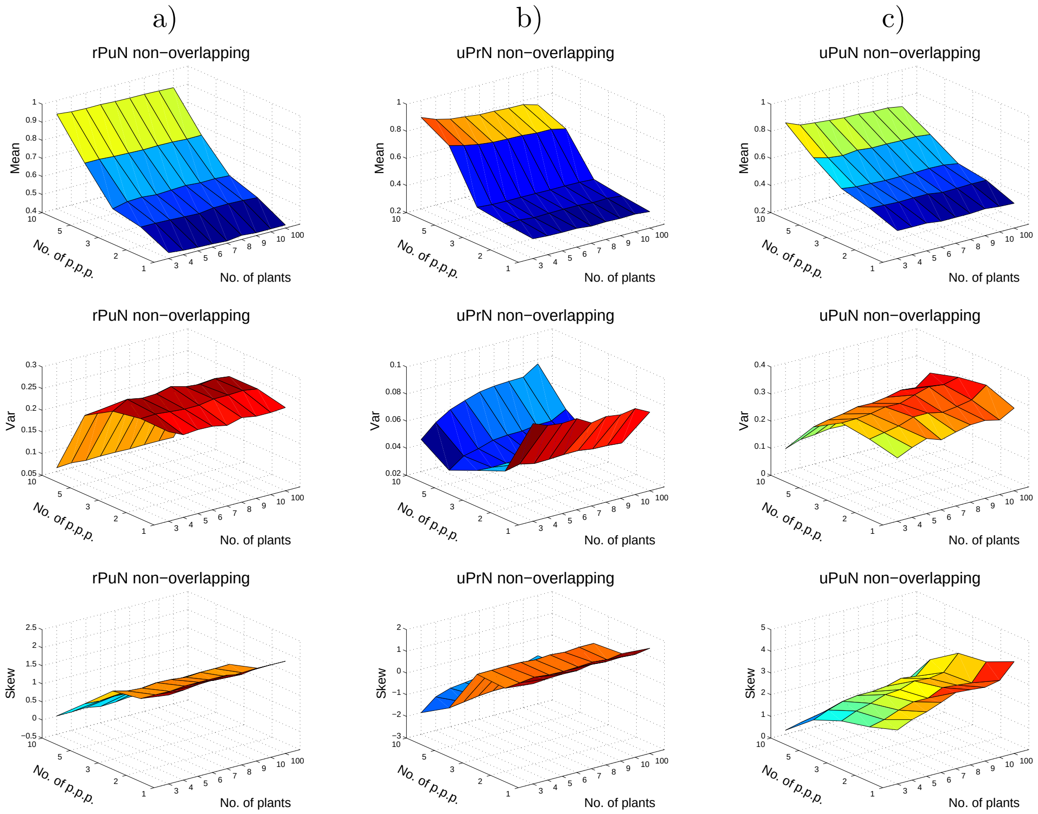

Figure A.1: Plots of mean, variance and skew of the non-overlapping Gompertz model for the three different environmental/neighbourhood configurations: regularly placed plants and uniformly random patches (rPuN); uniformly random plants and regular patches (uPrN); and uniformly random plants and patches (uPuN). Note that different scales are used on the z-axis across the different figures.

Overlapping Gompertz model

Fig (A.2) shows the results for the same configurations but for the overlapping Gompertz model. The permitted overlapping at the site of neighbour-neighbour interactions means that the population is better able to collectively take advantage of all patches acquired. In short: all patches acquired by any plant, regardless of neighbour locations, are turned into growth by the individual which acquired it. As a result of this, how the patch acquisition is distributed amongst the individuals in a population makes negligible difference to the mean plant size.

A.0.2 Constant growth models results

The results for the non-overlapping constant growth model are shown in figure A.3, and those for the overlapping constant growth model in figure A.4.

It can be seen by comparing these results with those of the original Gompertz models that, in both non-overlapping and overlapping configurations, the simplified constant growth model offers qualitatively very similar results. For the non-overlapping case the qualitative results as summarised by the mean, variance and skew are nearly indistinguishable, whilst for the overlapping case it is only at large plant numbers and low patch abundances in the rPuN and uPuN configurations that differences become apparent.

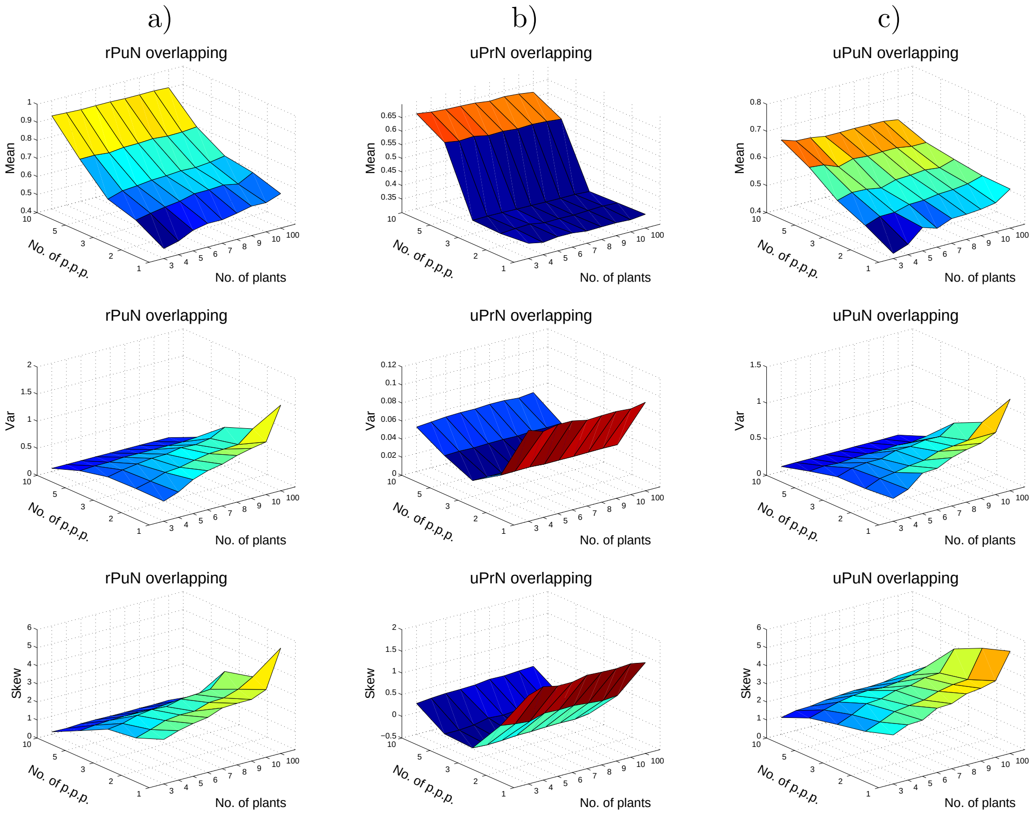

Figure A.2: Plots of mean, variance and skew of the overlapping Gompertz model for the three different environmental/neighbourhood configurations: regularly placed plants and uniformly random patches (rPuN); uniformly random plants and regular patches (uPrN); and uniformly random plants and patches (uPuN). Note that different scales are used on the z-axis across the different figures.

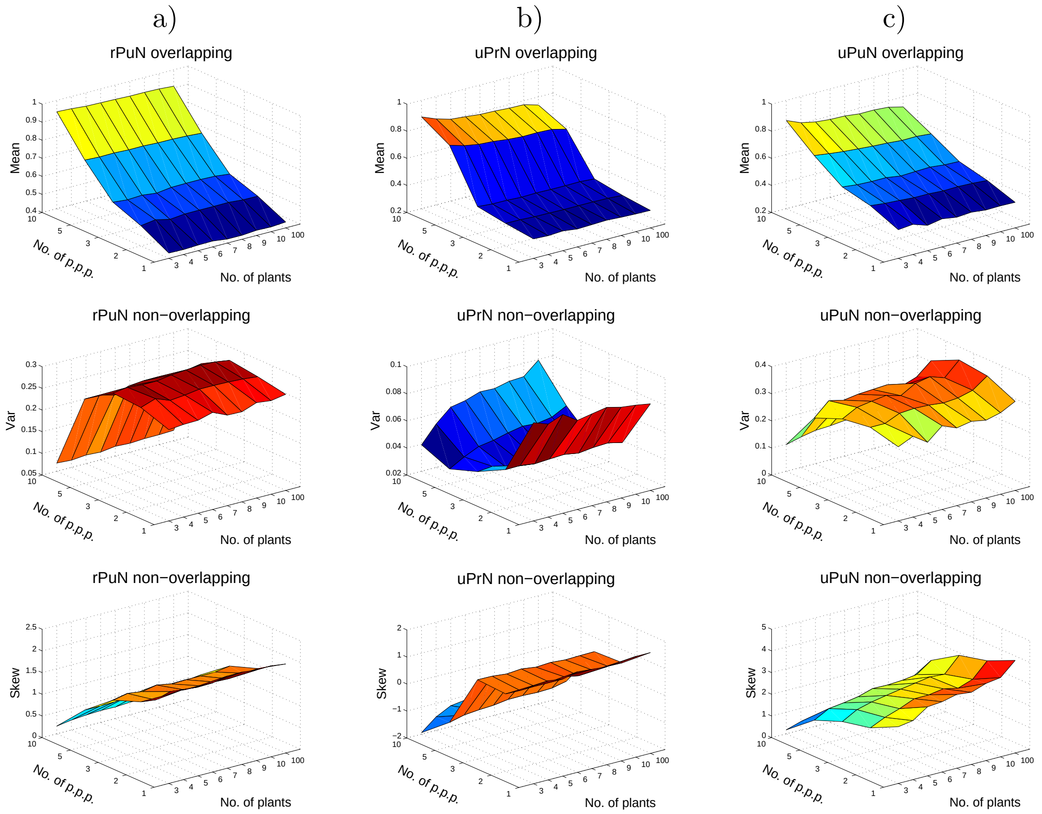

Figure A.3: Plots of mean, variance and skew of the non-overlapping constant growth model for the three different environmental/neighbourhood configurations: regularly placed plants and uniformly random patches (rPuN); uniformly random plants and regular patches (uPrN); and uniformly random plants and patches (uPuN). Note that different scales are used on the z-axis across the different figures.

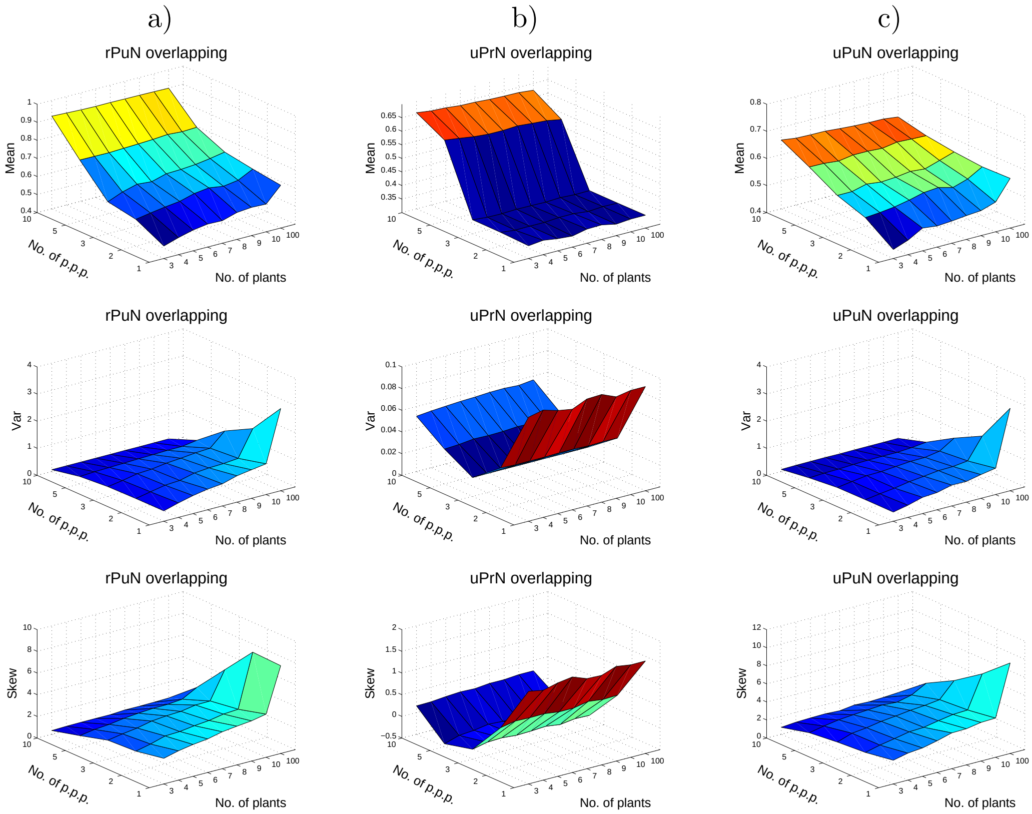

Figure A.4: Plots of mean, variance and skew of the overlapping constant growth model for the three different environmental/neighbourhood configurations: regularly placed plants and uniformly random patches (rPuN); uniformly random plants and regular patches (uPrN); and uniformly random plants and patches (uPuN). Note that different scales are used on the z-axis across the different figures.

Bibliography

[1] Barber SA (1995) Soil nutrient bioavailability: a mechanistic approach. Wiley, New York

[2] Barrett G, Campbell CD, Fitter AH, Hodge A (2011) The arbuscular mycorrhizal fungus *Glomus hoi* can capture and transfer nitrogen from organic patches to its associated host plant at low temperature. Appl Soil Ecol 48:102-105
berendse2007 Berendse F, de Kroon H, Braakhekke WG (2007) Acquisition, use and loss of nutrients. In: Pugnaire FI, Valladares F. (eds) Functional plant ecology. Atlanta, GA

[3] Berger U, Hildenbrandt H, Grimm V (2002) Towards a standard for the individual-based modelling of plant populations: self-thinning and the field of neighbourhood approach. Nat Res Mod 15:39-54

[4] Bever JD (2003) Soil community feedback and the coexistance of competitiors: conceptual frameworks and empirical tests. New Phytol 157:465-473

[5] Bhatt MV, Khandelwal A, Dudley SA (2011) Kin recognition, not competitive interactions, predicts root allocation in young *Cakile edentula* seedling pairs. New Phytol 189:1135-1142

[6] Brindley J, Pitchford JW (2001) Prey patchiness, predator survival and fish recruitment. Bull Math Biol 63:527-546

[7] Brindley J, James A, Pitchford JW (2003) Optimal foraging in patchy turbulent environments. Mar Ecol Prog S 256:99-110

[8] Brindley J, James A, Pitchford JW (2005) Quantifying the effects of individual and environmental variability in fish recruitment. Fish Oceanog 14:156-160

[9] Cahill JF Jr, McNickle GG (2011) The behavioural ecology of nutrient foraging by plants. Annu Rev Ecol Evol Syst 42:289-311

[10] Cahill JF Jr, McNickle GG, Haag JJ, Lamb EG, Nyanumba SM, St. Clair CC (2010) Plants integrate information about nutrients and neighbours. Science 328:1657

[11] Cain ML, Subler S, Evans JP, Fortin MJ (1999) Sampling spatial and temporal variation in soil nitrogen availability. Oecologia 118:397-404

[12] Campbell BD, Grime JP, Mackey JML (1991) A trade-off between scale and precision in resource foraging. Oecologia 87:532538.

[13] Canavos GC (1984) Applied probability and statistical methods. Little, Brown and Company, Boston p443

[14] Casper BB, Jackson RB (1997) Plant competition underground. Annu Rev Ecol Syst 28:545-70

[15] Castper BB, Schenk HJ, Jackson RB (2003) Defining a plant's belowground zone of influence. Ecol 84:2313-2321

[16] Charnov EL (1976) Optimal foraging, marginal value theorem. Theor Popul Biol 9:129-136

[17] Cropper WP, Comerford NB (2005) Optimizing simulated fertilizer additions using a genetic algorithm with a nutrient uptake model. Ecol Model 185:271-281

[18] Currey JD, Baxter PD, Pitchford JW (2007) Variability of the mechanical properties of bone, and its evolutionary consequences. J R Soc Interface 4:127-135

[19] Day KJ, Hutchings MJ, John EA (2003) The effects of spatial pattern of nutrient supply on yield, structure and mortality in plant populations. J Ecol 91:541-553

[20] Dudley SA, File AL (2007) Kin recognition in an annual plant. Biol Lett 3:435-438

[21] Duke SE, Caldwell MM (2000) Phospthate uptake kinetics of *Artemisia tridentata* roots exposed to multiple soil enriched-nutrient patches. Flora 195:154-164

[22] Dupuy L, Fourcaud T, Stokes A, Danjon F (2005) A density-based approach for the modelling of root architecture: application to Maritime pine (*Pinus pinaster* Ait.) root systems. J Theor Biol 236:323334

[23] Dupuy L, Gregory PJ, Bengough AG (2010) Root growth models: towards a new generation of continuous approaches. J Exp Bot 61:2131-2143

[24] Eissenstat DM, Yanai RD (1997) The ecology of root lifespan. Adv Ecol Res 27:1-60

[25] Farley RA, Fitter AH (1999) The responses of seven co-occurring woodland herbaceous perennials to localized nutrient-rich patches. J Ecol 87:688-696

[26] Fitter AH (1987) An architectural approach to the comparative ecology of plant root systems. New Phytol 106:61-77

[27] Fitter AH (1994) Architecture and biomass allocation as components of the plastic response of root systems to soil heterogeneity. In: Caldwell MM, Pearcy RW (eds) Exploitation of environmental heterogeneity by plants. Academic Press, San Diego, pp 305-323

[28] Fitter AH (1999) Roots as dynamic systems: the developmental ecology of roots and root systems. In: Press, MC, Scholes JD, Barker MG (eds) Plant physiological ecology. British Ecological Society symposium, No. 39. Blackwell Scientific Publications, Oxford, pp 115-131

[29] Fitter AH, Strickland TR (1991) Architectural analysis of plant root systems 2. Influence of nutrient supply on architecture in contrasting species. New Phytol 118:383-389

[30] Fitter AH,Strickland TR, Harvey ML, Wilson GW (1991) Architectural analysis of plant root systems 1. Architectural correlates of exploitation efficiency. New Phytol 118:375-382

[31] Fitter AH, Williamson LC, Linkohr B, Leyser O (2002) Root system architecture determines fitness in an *Arabidopsis* mutant in competition for immobile phosphate ions but not for nitrate ions. Proc R Soc Lond 269:2017-2022

[32] Forde BG (2009) Is it good noise? The role of developmental instability in the shaping of a root system. J Exp Bot 60:3989-4002

[33] Fransen B, de Kroon H, Berendse F (1998) Root morphological plasticity and nutrient acquisition of perennial grass species from habitats of different nutrient availability. Oecologia 115:351-358

[34] Fransen B, de Kroon H, De Kovel CGF, Van den Bosch F (1999) Disentangling the effects of root foraging and inherent growth rate on plant biomass accumulation in heterogeneous environments: a modelling study. Ann Bot-London 84:305311

[35] Garbout A, Munkholm LJ, Hansen SB, Petersen BM, Munk OL, Pajor R (2011) The use of PET/CT scanning technique for 3D visualization and quantification of real-time soil/plant interactions. Plant Soil 352:113127

[36] Gersani M, Abramsky Z, Falik O (1998) Density-dependent habitat selection in plants. Evol Ecol 12:223-234

[37] Gleeson SK, Fry JE (1997) Root proliferation and marginal patch value. Oikos 79:387-393

[38] Grimmett G, Stirzaker D (2004) Probability and random processes (third ed.) Oxford University Press, Oxford

[39] Gross KL, Pregitizer KS, Burton AJ (1995) Spatial variation in nitrogen availability in three successional plant communities. J Ecol 83:357-367

[40] Güsewell S (2004) N:P ratios in plants: variation and functional significance. New Phytol 164:243-266

[41] Hara T (1988) Dynamics of size structure in populations. Trends Ecol Evol 3:129-132

[42] Hardin G (1968) The tragedy of the commons. Science 162:1243-1248

[43] Hausmann NT, Hawkes CV (2009) Order of plant host establishment alters the composition of arbuscular mycorrhizal communities. Ecology 91: 2333-2343

[44] Hess L, de Kroon H (2007) Effects of rooting volume and nutrient availability as an alternative explanation for root self/non-self discrimination. J Ecol 95:241-251

[45] Hinsinger P, Bengough AG, Vetterlein D, Young IM (2009) Rhizosphere: biophysics, biogeochemistry and ecological relevance. Plant Soil 321:117-152

[46] Hodge A (2003a) Plant nitrogen capture from organic matter as affected by spatial dispersion, interspecific competition and mycorrhizal colonization. New Phytol 157:303-314

[47] Hodge A (2003b) N capture by *Plantago lanceolata* and *Brassica napus* from organic material: the influence of spatial dispersion, plant competition and an arbuscular mycorrhizal fungus. J Exp Bot 54:2331-2342

[48] Hodge A (2004) The plastic plant: root responses to heterogeneous supplies of nutrients. New Phytol 162:9-24

[49] Hodge A (2009) Root decisions. Plant Cell Environ 32:628-640

[50] Hodge A, Alexander IJ, Gooday GW (1995) Chitinolytic enzymes of pathogenic and ectomycorrhizal fungi. Mycol Res 99:935-941

[51] Hodge A (1996) Impact of elevated CO_2 on mycorrhizal associations and implications for plant growth. Biol Fert Soils 23:388-398

[52] Hodge A, Paterson E, Thornton B, Millard P, Killham K (1997) Effects of photon flux density on carbon partitioning and rhizosphere carbon flow of *Lolium perenne.* J Exp Bot 48:1797-1805

[53] Hodge A and Millard P (1998) Effect of elevated CO_2 on carbon partitioning and exudates release from *Plantago lanceolata* seedings. Physiol Plant 103:280-286

[54] Hodge A, Stewart J, Robinson D, Griffiths BS, Fitter AH (1998a) Root proliferation, soil fauna and plant nitrogen capture from nutrient-rich patches in soil. New Phytol 139:479-494

[55] Hodge A, Paterson E, Grayston SJ, Campbell CD, Ord BG, Killham K (1998b) Characterisation and microbial utilisation of exudate material from the rhizosphere of *Lolium perenne* grown under CO_2 enrichment. Soil Biol Biochem 30:1033-1043

[56] Hodge A, Robinson D, Griffiths BS, Fitter AH (1999a) Why plants bother: root proliferation results in increased nitrogen capture from an organic patch when two grasses compete. Plant Cell Environ 22:811-820

[57] Hodge A, Stewart J, Robinson D, Griffiths BS, Fitter AH (1999b) Plant, soil fauna and microbial responses to N-rich organic patches of contrasting temporal availability. Soil Biol Biochem 31:1517-1530

[58] Hodge A, Stewart J, Robinson D, Griffiths BS, Fitter AH (2000a) Plant N capture and microfaunal dynamics from decomposing grass and earthworm residues in soil. Soil Biol Biochem 32:1763-1772

[59] Hodge A, Stewart J, Robinson D, Griffiths BS, Fitter AH (2000b) Spatial and physical heterogeneity of N supply from soil does not influence N capture by two grass species. Funct Ecol 14:645-653

[60] Hodge A, Berta G, Doussan C, Merchan F, Crespi M (2009) Plant root growth, achitecture and function. Plant Soil 321:153-187

[61] Hodge A, Fitter AH (2010) Substantial nitrogen acquisition by arbuscular mycorrhizal fungi from organic material has implications for N cycling. P Natl Acad Sci USA 107:13754-13759

[62] Hutchings MJ, Wijesinghe DK (2008) Performance of clonal species in patchy environments: effects of environmental context on yield at local and whole-plant scales. Evol Ecol 22:313-324

[63] Jackson RB, Caldwell MM (1993) Geostatistical patterns of soil heterogeneity around individual perennial plants. J Ecol 81:683-692

[64] Jakobsen BF, Dexter AR (1987) Effect of soil structure on wheat root growth, water uptake and grain yield. A computer simulation model. Soil Tillage Res 10:331-345

[65] James A, Baxter P, Pitchford J (2005) Modelling predation as a capped rate stochastic process, with applications to fish recruitment. J R Soc Interface 2:477-487

[66] James A, Pitchford JW, Plank MJ (2010) Efficient or inaccurate? Analytical and numerical modelling of random search strategies. B Math Biol 72:896-913

[67] Johnson D, Leake JR, Ostle N, Ineson P, Read DJ (2002) In situ CO_2-C-13 pulse labeling of upland grassland demonstrates a rapid pathway of carbon flux from arbuscular mycorrhizal mycelia to the soil. New Phytol 153:327-334

[68] Johnson NL, Kotz S, Balakrishnan N (1994) Continuous univariate distributions-1, Second Edition. New York: Wiley-Interscience, pp 574-575

[69] Jones DL, Hodge A, Kuzyakov Y (2004) Plant and mycorrhizal regulation of rhizodeposition. New Phytol 163:459-480

[70] Jones DL, Nguyen C, Finlay RD (2009) Carbon flow in the rhizosphere: carbon trading at the soil-root interface. Plant Soil 321:5-33

[71] Jourdan C, Rey H (1997) Modelling and simulation of the architecture and development of the oil-palm (*Elaeis guineenis Jacq.*) root system. Plant Soil 190:217:233

[72] Kaye JP, Hart SC (1997) Competition for nitrogen between plants and soil microorganisms. TREE 12:139-143

[73] Kembel SW, Cahill JF (2005) Plant phenotypic plasticity belowground: a phylogenetic perspective on root foraging trade-offs. Am Nat 166:216230.

[74] Kembel SW, de Kroon H, Cahill JF, Mommer L (2008) Improving the scale and precision of hypotheses to explain root foraging ability. Ann Bot-London 101: 12951301

[75] Hishi T (2007) Heterogeneity of individual roots within the fine root architecture: causal links between physiolgical and ecosystem functions. J For Res 12:126-133

[76] Jury WA, Horton R (2004) Soil physics (sixth ed.). Wiley, New York

[77] de Kroon H, Visser EJW, Huber H, Mommer L, Hutchings MJ (2009) A modular concept of plant foraging behaviour: the interplay between local responses and systemic control. Plant Cell Environ 32:704-712

[78] Kuzyakov Y 2002. Review: factors affecting rhizosphere priming effects. J Plant Nutr Soil Sci 165:382-396

[79] Law R, Dieckmann U (2000) A dynamical system for neighbourhoods in plant communities. Ecol 81:2137-2148

[80] Leigh J, Hodge A, Fitter AH (2009) Arbuscular mycorrhizal fungi can transfer substantial amounts of nitrogen to their host plant from organic material. New Phytol 181:199-207

[81] Leigh J, Fitter AH, Hodge A (2011) Growth and symbiotic effectiveness of an arbuscular mycorrhizal fungus in organic matter in competition with soil bacteria. FEMS Microbiol Ecol 76:428-438

[82] Leitner D, Klepsch S, Bodner G, Schnepf A (2010) A dynamic root system growth model based on L-Systems. Plant Soil 332(1-2):177192

[83] Linkohr BI, Williamson LC, Fitter AH, Leyser OHM (2002) Nitrate and phosphate availability and distribution have different effects on root system architecture in *Arabidopsis*. Plant J 29:751-760

[84] Lv Q, Schneider MK, Pitchford JW (2008) Individualism in plant populations: using stochastic differential equations to model individual neighbourhood-dependent plant growth. Theor Popul Biol 74:74-83

[85] Mairhofer S, Zappala S, Tracy SR, Sturrock C, Bennett MJ, Mooney SJ, Pridmore TP (2012) RooTrak: automated recovery of three-dimensional plant root architecture in soil from x-ray microcomputed tomography images using visual tracking. Plant Physiol 158:561569

[86] Mairhofer S, Zappala S, Tracy SR, Sturrock C, Bennett MJ, Mooney SJ, Pridmore TP (2013) Recovering complete plant root system architectures from soil via x-ray μ-Computed Tomography. Plant Methods 9:8

[87] Malamy JE (2005) Intrinsic and environmental response pathways that regulate root system architecture. Plant Cell Environ 28:67-77

[88] Masclaux F, Hammond RL, Meunier J, Gouhier-Darimont C, Keller L, Reymond P (2010) Competitive ability not kinship affects growth of *Arabidopsis thaliana* accessions. New Phytol 185:322-331

[89] McNickle GG, St. Claire CC, Cahil JF (2009) Focusing the metaphor: plant foraging behaviour. Trends Ecol Evol 24:419-426

[90] Milla R, Forero DM, Escudero A, Iriondo JM (2009) Growing with siblings: a common ground for cooperation or for fiercer competition among plants? Proc R Soc B 276:2531-2540

[91] Murphy GP, Dudley SA (2009) Kin recognition: competition and cooperation in Impatiens (Balsaminaceae). Am J Bot 96:1990-1996

[92] O'Brien EE, Brown JS, Moll JD (2007) Roots in space: a spatially explicit model for below-ground competition in plants. Proc R Soc B 274:929-934

[93] Paterson E (2003) Importance of rhizodeposition in the coupling of plant and microbial productivity. Eur J Soil Sci 54:741-750

[94] Paterson E, Hodge A, Thornton B, Millard P, Killham K (1999) Carbon partitioning and rhizosphere C-flow in *Lolium perenne* as affected by CO_2 concentration, irradiance and below-ground conditions. Glob Change Biol 5:669-678

[95] Pierret A, Doussan C, Capowiez Y, Bastardie F, Pages L (2007) Root functional architecture: a framework for modelling the interplay between roots and soil. Vadose Zone J 6(2):269281

[96] Preston MD, Pitchford JW, Wood AJ (2010) Evolutionary optimality in stochastic search problems. J R Soc Interface 7:1301-1310

[97] Prusinkiewicz P, Lindenmayer A (1990) The algorithmic beauty of plants. Springer, Berlin

[98] Purves DW, Law R (2002) Experimental derivation of functions relating growth of *Arabidopsis thaliana* to neighbour size and distance. J Ecol 90:882-894

[99] Read DJ, Perez-Moreno J (2003) Mycorrhizas and nutrient cycling in ecosystems - a journey towards relevance? New Phytol 157:475-492

[100] Rice JA (2007) Mathematical statistics and data analysis (third ed.) Thomson Higher Education, Belmont

[101] Reynolds HL, Hartley AE, Vogelsang KM, Bever JD, Schultz PA (2005) Arbuscular mycorrhizal fungi do not enhance nitrogen acquisition and growth of old-field perennials under low nitrogen supply in glasshouse culture. New Phytol 167:869-880

[102] Robinson D (1991) Roots and resource fluxes in plants and communities. In: Atkinson D (ed) Plant root growth: an ecological perspective. Special publication of the British Ecological Society n10. Blackwell Scientific Pub., London, pp 103-130

[103] Robinson D, Hodge A, Griffiths BS, Fitter AH (1999) Plant root proliferation in nitrogen-rich patches confers competitive advantage. Proc R Soc Lond 266:431-435

[104] Roose T, Fowler AC (2004) A mathematical model for water and nutrient uptake by plant root systems. J Theor Biol 228(2):173184

[105] Roose T, Schnepf A (2008) Mathematical models of plant soil interaction. Philos Trans R Soc A 366:45974611

[106] Sanders FE, Tinker PB (1973) Phosphate flow into mycorrhizal roots. Pestic Sci 4:385-395

[107] Schenk HJ (2006) Root competition: beyond resource depletion. J Ecol 94:725-739

[108] Schenk HJ, Callaway RM, Mahall BE (1999) Spatial root segregation: are plants territorial? Adv Ecol Res 28:145-180

[109] Schneider MK, Law R, Illian JB (2006) Quantification of neighborhood-dependent plant growth bt Bayesian hierarchical modelling. J Ecol 94:310-321

[110] Schwinning S, Weiner J (1998) Mechanisms determining the degree of size asymmetry in competition among plants. Oecologia 113:447-455

[111] Shemesh H, Arbiv A, Gersani M, Ovadia O, Novoplansky A (2010) The effects of nutrient dynamics on root patch choice. PloS ONE 5:1-6

[112] Smith FA, Grace EJ, Smith SE (2009) More than a carbon economy: nutrient trade and ecological sustainability in facultative arbuscular mycorrhizal symbioses. New Phytol 182:347358

[113] Smith SE, Read DJ (2008) Mycorrhizal Symbiosis, 3rd Edn. Academic Press, London

[114] Southworth D, He X-H, Swenson W, Bledsoe CS, Horwath WR (2005) Application of network theory to potential mycorrhizal networks. Mycorrhiza 15:589-595

[115] Tardieu F (2010) Why work and discuss the basic principles of plant modelling 50 years after the first plant models? J Exp Bot 61:2039-2041

[116] Tardieu F, Bruckler L, Lafolie F (1992) Root clumping may affect the root water potential and the resistance to soil-root water transport. Plant Soil 140:291-301

[117] Taub DR, Goldberg D (1996) Root system topology of plants from habitats differing in soil resource availability. Func Ecol 10:258-264.

[118] Tinker PB, Nye PH (2000) Solute Movement in the Rhizosphere. Oxford University Press, Oxford, p 84

[119] Vos J, Evers JB, Buck-Sorlin GH, Andrieu B, Chelle M, de Visser PHB (2010) Functional-structural plant modelling: a new versatile tool in crop science. J Exp Bot 61:2101-2115

[120] van Vuuren MMI, Robinson D, Griffiths BS (1996) Nutrient inflow and root proliferation during the exploitation of a temporally and spatially discrete source of nitrogen in soil. Plant Soil 178:185-192

[121] Waisel Y, Eshel A (1992) Differences in ion uptake among roots of various types. J Plant Nutr 15:945-958

[122] Williamson LC, Ribrioux SPCP, Fitter AH, Leyser O (2001) Phosphate availability regulates root system architecture in *Arabidopsis*. Plant Physiol 126:875-882

[123] Wijesinghe DK, John EA, Hutchings MJ (2005) Does pattern of soil resource heterogeneity determine plant community structure? An experimental investigation. J Ecol 93:99-112