# PLANTWORLD:
# The Evolution of Plant Dormancy in Contrasting Environments

**Jacqueline R. Dyer**
School of Oriental and African Studies
University of London
Thornhaugh Street
London
jd6@soas.ac.uk

**Peter J. Bentley**
Department of Computer Science
University College London
Gower Street
London
P.Bentley@cs.ucl.ac.uk

**Panash Shah**
Department of Computer Science
University College London
Gower Street
London

## Abstract

Evolutionary algorithms have been used to tackle many problems over the years. Here we describe a system designed to investigate ecological theories by modelling the population dynamics of evolving plants in a digital environment. The system is called PLANTWORLD.

## 1 INTRODUCTION

PLANTWORLD is a flat, two-dimensional landscape in which PLANTS grow from SEEDS, obtaining moisture, reproducing with similar neighbours, and *evolving* genomes that determine their survival strategies in the environment. Rainfall occurs in different spatial patterns, following real-world rainfall data. A water table maintains an underground reservoir of moisture, accessible to larger PLANTS. Figure 1 gives a graphical overview of the system.

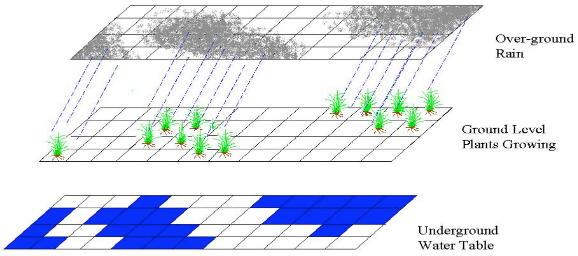

**Figure 1:** PLANTWORLD overview.

The PLANTWORLD model has been initially developed in order to examine the effects that the evolution of a functional response - in this case, dormancy - might have on the population dynamics of PLANTS. Each PLANT requires a single resource, *moisture*, which varies in availability both spatially and temporally. In addition, this implementation allows us to study the effects of two further strategies that can influence dynamics: (i) the effects of PLANT *storage capacity* ii) the effects of an alternative source of moisture, in the form of a *Water Table.*

Two objectives motivate the development of this system. The far-reaching objective is to attempt to develop systems that can integrate evolutionary and ecological dynamics in spatially extensive and temporally variable environments. Such an objective is prohibited in numerical models by its sheer complexity and is only recently becoming a realistic objective in computational models. PLANTWORLD represents only the initial stages in the development of such a system. As such, it only models one type of trophic agent, PLANTS, and a single resource, moisture. However, its modular design means that it is readily extendable and it is envisaged that other agents (herbivores, pests, etc) and variables (nutrients, light, fire, etc) will be added at later stages. In the meantime, the development of PLANTWORLD has a more immediate objective. One of the advantages of agent-based models over numerical models of population dynamics is that our agents can exhibit *behaviours.* Combined with evolutionary computation, such behaviours can evolve. Thus, we can examine how the evolution of traits in different environments affects the population dynamics in these environments. The immediate objective for building PLANTWORLD is therefore to examine the evolution and effects of plant dormancy on population dynamics in different spatially and temporally variable environments. The simulation is not intended to capture realistic behaviour of any specific flora but rather to test the veracity of predictions about population dynamics that arise from numerical models.

## 2 MOTIVATION

Numerical models of population dynamics in stable compared to stochastic environments make implicit assumptions regarding the functional types of populations under study. Based on these assumptions, certain predictions follow:

1. Populations in more stable environments are strongly driven by *biotic* factors (e.g. competition, herbivory) whereas populations in more stochastic environments are strongly driven *by abiotic* factors (e.g. rainfall, fire) (Wiens, 1984).
2. Populations in more stable environments experience *density-dependent* population dynamics whereas those in stochastic environments experience *density - independent* population dynamics (Lotka 1924).
3. Population in stable environments exhibit *asymptotic stability near equilibrium* whereas populations in stochastic environments exhibit strong fluctuations and lack of *equilibrium* (Wiens, 1984).
4. Organisms in stable environments are likely to be K-strategists whereas organisms in disturbed environments are likely to be r-strategists (MacArthur and Wilson, 1967).
5. *Niche overlap* is lower in more stable compared with more stochastic environments (Odum 1971).
6. Communities in more disturbed environments are likely to be less *complex* than communities in more stable environments (May, 1974).

Many subsequent models make predictions that undermine one or other of those above (e.g. Abrams, 1986, Chesson, 1991; Law and Morton, 1996; Levin 1974; Martinez, 1992; Mattessi and Gatto, 1983; May 1975). The advantage of the approach undertaken in PLANTWORLD is that we can implement a bottom-up, synthetic, spatial model that can take into account many of the caveats implicit in numerical models and raised by subsequent researchers. One of the main caveats implicit in many numerical models is that populations in different environments do not differ in their functional response to these environments (MacNally, 1995). Although this assumption is not explicitly held by ecologists, it is implicit in numerical models that do not define the functional type of populations beyond that of relevance to the interaction concerned (i.e. predator or prey in predator-prey models). Such models predict that population dynamics in disturbed environments will fluctuate more strongly than those in stable environments, resulting in higher extinction rates, lower biodiversity and more simple community structure in disturbed, compared to stable environments.

Clearly, predictions concerning community structure (for example, 6. above) cannot be fully enlightened by modelling a single trophic level existing on a single resource. However, since plants are primary producers, the dynamics of higher trophic levels will be strongly influenced by plant dynamics. For example, if moisture variability is mitigated by a functional response such as dormancy, then plant populations in highly variable environments might not exhibit such strong population fluctuations as predicted by numerical models. Thus, we might have less reason to suspect such contrasts in community structure under stable and stochastic conditions as many models predict. A further advantage of PLANTWORLD is that we can examine the effects of small and large parameter changes, particularly the costs and benefits of evolving particular traits in particular environments. Finally, we hope to extend the model to include adaptive agents at other trophic levels and to include further variables, in order to examine aspects of community structure and dynamics in contrasting environments.

# 3 MODEL SUMMARY

In PLANTWORLD, the environment consists of a two-dimensional grid of cells that holds the resource, *Soil Moisture*. Inputs of moisture arrive via rainfall. Moisture is removed by PLANTS, SEEDS and by evaporation. The environment can also incorporate a *Water Table*. This underlies the *Soil Moisture* grid. PLANTS germinate, grow, reproduce and die on this landscape (Figure 2). PLANTS can evolve just two traits: *AdultSize* and *Dormancy Strategy*. *AdultSize* defines the size at which a PLANT can reproduce. The *Dormancy Strategy* defines the activity and dormancy of a PLANT over an annual cycle. Further parameters, such as the *Maintenance Parameter*, number of *Reproductive Events* and *Seed Dormancy*, are dependent on either *AdultSize* or *Dormancy Strategy*, hence these are all indirectly dependent on the genome. Other parameters, such as *Acquired Resources*, *Maintenance Requirements*, *Maximum Utilisation* and *Storage Effect* are dependent on the *CurrentSize* of the PLANT. PLANT parameters and behaviour are further described later in this paper.

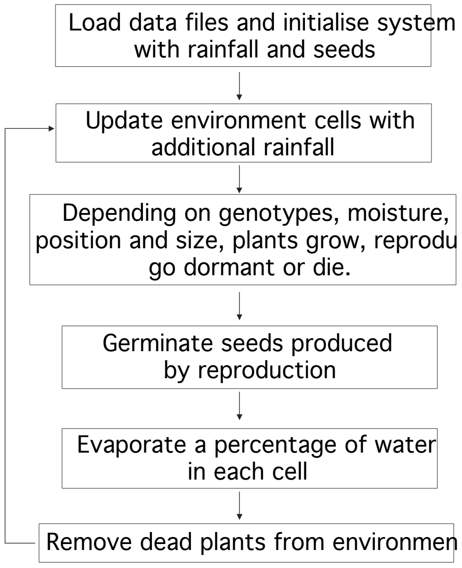

**Figure 2:** Flow diagram of PLANTWORLD.

## 3.1 THE ENVIRONMENT

Since we are interested in both spatial and temporal variability in resources, the model requires a spatial representation. The environment is modelled as a two-dimensional grid with 100 x 100 cells. The spatial representation afforded by the grid allows us to model a heterogeneous environment whereby moisture availability can be different in each cell. Moisture is available in two compartments of each cell: an upper one (given by the *Soil Moisture* grid) and a lower one (given by the *Water Table* grid). The *Soil Moisture* grid receives its moisture from rainfall and represents the main source of moisture for PLANTS. The *Water Table* grid is an extra feature added so that we can examine effects of this in certain experiments. Thus, the spatial variability in resource availability is given by the two moisture grids.

### 3.1.1 Rainfall Time-Series

In addition to spatial variability, moisture levels in the environment must vary with time. To achieve realistic fluctuations of rainfall, PLANTWORLD employs real time-series data describing actual rainfall in different regions of the world.

The modelling objectives require that two PLANT traits – *AdultSize* and *Dormancy Strategy* – evolve over time. We are therefore interested in the dynamics of populations over many generations. Since this implies that our model runs should extend for at least decades, we cannot expect to incorporate fine scale temporal phenomena into the model. We concluded that the appropriate temporal resolution in the model should be a single month. Thus, the rainfall time-series represent average rainfall receipts for each month. The data chosen for simulating the environment are taken from 100-year records of monthly rainfall receipts in various regions of the globe, gridded at 2.5° latitude by 3.75° longitude resolution (Hulme, 1998). In addition, we can generate our own data (for example, highly stable or highly stochastic temporal variation) for use in certain experiments, and to compare dynamics under real rainfall data and artificial data.

### 3.1.2 Soil Moisture Grid

The *Soil Moisture* grid stores moisture that arrives through rainfall. Each cell can have different initial amounts of moisture at the beginning of each timestep. Moisture is added to each cell according to the rainfall given by the time-series and scaled by the cell's *Receipt Value*, which is given by a 100 * 100 pattern (see below). Moisture is removed from cells by PLANTS and by evaporation. The input of moisture to *Soil Moisture* grid is described first, followed by a description of how moisture is removed.

***Input of moisture to the soil moisture cells***

Rainfall each timestep is given by the time-series data. Rainfall received in each cell each timestep is given according to the amount given by the time-series, scaled to the cell's *Receipt Value*. The cell's *Receipt Value* is a number from either 1 to 5, 1 to 10, or 1 to 15. *Receipt Values* are given according to a pattern stipulated for use in a given timestep. (These are generated beforehand using a cellular automaton pattern-generator written for this task.) PLANTWORLD calls a different pattern each timestep, so that each cell may have a different *Receipt Value* each timestep, unless the pattern called is set to be the same. The program can cycle through up to 12 patterns, then repeats. Figure 3 shows an example of a 10 * 10 grid that provides *Receipt Values* for each cell.

| | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| 1 | 3 | 1 | 4 | 7 | 7 | 8 | 9 | 10 | 10 |
| 2 | 2 | 3 | 5 | 6 | 7 | 9 | 10 | 10 | 10 |
| 1 | 3 | 4 | 6 | 7 | 8 | 10 | 10 | 10 | 10 |
| 2 | 4 | 5 | 6 | 8 | 7 | 9 | 10 | 10 | 10 |
| 3 | 4 | 5 | 7 | 8 | 8 | 9 | 9 | 10 | 10 |
| 3 | 4 | 6 | 7 | 7 | 9 | 9 | 8 | 9 | 10 |
| 4 | 5 | 6 | 7 | 8 | 7 | 9 | 7 | 8 | 8 |
| 3 | 4 | 5 | 5 | 5 | 4 | 6 | 6 | 7 | 6 |
| 2 | 3 | 3 | 3 | 2 | 3 | 5 | 4 | 6 | 6 |
| 1 | 1 | 3 | 2 | 2 | 2 | 2 | 2 | 3 | 5 |

**Figure 3**. 10 * 10 Grid pattern of different *Receipt Value* in different cells.

A pattern like this (100 * 100) is used each timestep to set the *Receipt Value* for each cell. The rainfall given by the timeseries (an integer) is then scaled between the minimum and maximum *Receipt Values*, in this case 1 to 10. So a *Receipt Value* of 1 implies no rainfall that timestep and *Receipt Value* 10 implies the full rainfall given by the timeseries. Thus, the Soil Moisture grid can reflect the following spatial conditions:

***Homogeneous landscape and rainfall pattern***

If all cells are set to receive the same rainfall, e.g. level 15 at each timestep, then there is no variation either on the landscape or in the spatial rainfall pattern. Each cell always receives the rainfall given by the time-series.

***Heterogeneous landscape, homogeneous rainfall pattern***

If cells are set to receive different rainfall amounts (Levels 1-15) but the pattern used is the same at each timestep, then the landscape is heterogeneous but rain falls homogeneously over it.

***Homogeneous landscape, heterogeneous rainfall pattern***

If the initial moisture availability in the cells is homogeneous throughout the cells, but subsequent patterns of rainfall are heterogeneous with patterns changing through time, then the landscape is homogeneous but rainfall is heterogeneous and its spatial pattern varies through time.

***Heterogeneous landscape and rainfall pattern***

If moisture availability is initially heterogeneous and different heterogeneous patterns are used to determine rainfall receipts per timestep, then both the landscape and the spatial pattern are heterogeneous through time.

Although this implementation allows considerable flexibility in generating spatio-temporal variability in moisture availability patterns, there is no way of knowing if particular simulations correspond to actual variability in any given (real) environment. There is no empirical data with which to calibrate the model because, even though we are able to record variation in rainfall receipts over time and space, the actual moisture availability at any location is also a function of several complex landscape attributes, such as gradient and soil type at the location and at other neighbouring locations. While small-scale studies of, for example, watersheds, may have recorded actual moisture variability across a landscape, the scales involved are of little use for our modelling objective. For this, we simply want to compare the effects of different kinds and levels of heterogeneity and the current implementation is therefore adequate for our needs. However, we need to be sure that the choice of patterns for each simulation does reflect the sort of variability we expect. This can be confirmed by carrying out some statistical analysis on each cell through a simulation of changing moisture availability, before PLANTS are introduced to the environment. Thus, we can measure the *coefficient of variation* in moisture availability for each cell and use these indices for classifying the different environments given by different rainfall time-series and spatial patterns.

### 3.1.3 Removal of moisture from soil moisture cells

Moisture is removed from every cell each timestep by the following means:

a) When PLANTS maintain, grow and reproduce, and when SEEDS germinate, they use moisture units. The amounts and circumstances in which PLANTS use resources are detailed in the section regarding PLANT behaviour, which follows. The amount of resources (moisture) left in the cell after the PLANTS and SEEDS have used units is then subject to *evaporation*.

b) *Evaporation* occurs after the PLANTS and SEEDS have used resources. The *Evaporation rate* can be set by the user. After evaporation, the cell updates its resource units.

The cell updates its resources after removal by PLANTS, SEEDS and evaporation, and this amount is then added to the new rainfall input (the number given by the time-series and scaled by the cell's *Receipt Value*) to give the cell's *Soil Moisture* at the beginning of the next timestep.

### Water Table Grid

As an addition to the main requirements of the model, we would like to examine the effects on dynamics of making an alternative source of moisture available to PLANTS. The simplest implementation we could think of was to examine the effect of an underlying *Water Table* grid, that indicates whether or not a *Water Table* is available in a particular cell. PLANTS have access to the *Water Table* when they reach a *ThresholdSize*, defined by the user. If the *Water Table* is present and the PLANT has reached the *ThresholdSize*, the PLANT has access to the *Water Table*. The effect of the *Water Table* on the PLANT that can utilise it, is to lower its *Maintenance Parameter* (*k*) by 0.1, implying that the PLANT requires less resources from its *Soil Moisture* cell in order to meet its *Maintenance Requirements* (see Section 3.2.3).

The *Water Table* is represented by another grid that underlies the *Soil Moisture* grid. Each cell has a value of 0 or 1. Hence, the *Water Table* is available or not in a particular cell. The pattern does not change throughout a given simulation. By modelling the *Water Table* as either present or not, we are not giving it any variable value, and in effect, we are not removing any resource from the *Water Table* when a PLANT uses it. This is permissible because of the way we have modelled PLANT interaction with the *Water Table*. If the PLANT reaches a certain size, and a *Water Table* (cell value 1) does underlie its cell, it can use the *Water Table*. The *Water Table* has the effect of lowering the maintenance demands of the PLANT, so that it requires less resources from the *Soil Moisture* compartment for maintenance and more of its *Acquired Resources* can be spent on growth. Thus, we can examine this simple instantiation of the effect of an alternative moisture resource on PLANT strategies, without having to deal with the dynamics of actual amounts of the resources held in the *Water Table* (which would be much more complicated). In order to examine the effects of the *Water Table*, we need to be able to run experiments both with and without *Water Tables*. This is easily done: complete absence of the *Water Table* can be implemented by using a pattern of all zeros for the *Water Table* grid.

### 3.1.4 Seed Germination

When PLANTS reproduce, SEEDS are produced and dispersed 360° in a random Gaussian distribution centred around one of the parents. The number of SEEDS produced per reproduction event can be varied, but currently is a random number between 50 and 100. SEEDS therefore 'arrive' in cells. Different SEEDS can remain dormant for different lengths of time (given by their genome). They can germinate within this time, but should they fail to germinate they die. Germination is a function of the amount of resources left in the cell after PLANTS have utilised resources and before evaporation occurs. If the decision to germinate was determined by SEED methods, each SEED would have to check out resources and decide whether it could germinate every timestep. This instantiation would require considerable processing time, since SEEDS can rapidly build up in cells. It was therefore decided that SEED germination would be carried out by the cell, and the following behaviour has been implemented.

Once the PLANTS in the cell have removed their *Acquired Resources* (see Section 3.2.5), the cell calculates how many units of moisture are left over. A percentage of this (currently 10%, but can be changed) is available for SEED germination. The units of moisture available is then equal to the number of SEEDS that can germinate, since each germination event requires one unit of resource. The actual number of SEEDS that germinate depends on the number in the cell. If the amount of moisture units available exceeds the number of SEEDS in the cell, then all SEEDS germinate. If the amount is below the number of SEEDS in the cell, then the SEEDS that can germinate are chosen at random. When each SEED germinates it grows spontaneously to a PLANT of *CurrentSize1*, and in the following timestep it carries out the behaviour stipulated by its PLANT genome.

## 3.2 THE PLANTS

Since we are interested in the evolution of dormancy strategies and their effects on population dynamics and organisation in contrasting environments, a minimal requirement for the model is that PLANTS can evolve different dormancy strategies that imply different costs in terms of maintenance requirements. A further requirement is that we are able to model plants that reach maturity at different stages. The latter is required because plants that mature quickly may not require dormancy to persist, if they complete their lifecycles when resources are available (r-strategists). Supporting these requirements in turn requires various other behaviours. PLANTS must germinate from SEEDS, maintain themselves and grow, reproduce and die.

### 3.2.1 Evolution of Dormancy Strategies

We are interested in two main dormancy strategies: *periodic dormancy* and *spontaneous dormancy*. In periodic dormancy, the PLANT is deterministically dormant or deterministically active. In spontaneous dormancy, the state (dormant/active) of the PLANT depends on the availability of resources. The PLANT can therefore be in any of three states in a given timestep: deterministic activity (*Kpon*), deterministic dormancy (*Kpoff*) and spontaneous activity/dormancy (*Ks*).

In PLANTWORLD, the active state implies that the PLANT must be able to gain resources equal to or in excess of its *Maintenance Requirements*. Its *Maintenance Requirements* (*m*) are governed by its *Maintenance Parameter*, *k* in the following equation:

*Maintenance Requirements (m) = k . CurrentSize*

Thus, when the PLANT is active, $k > 0$ and when the PLANT is dormant, $k = 0$.

If the PLANT is in state *Kpon* (deterministic activity) in a given timestep, it requires its maintenance resources. If these are not available, it may be able to survive due to the *StorageEffect* (see below), otherwise it dies with a given probability (set currently at 50%).

If the PLANT is in state *Kpoff* (deterministic dormancy) in a given timestep, it does not require resources for maintenance and cannot use resources for growth or reproduction whether or not they are available.

If the PLANT is in state *Ks* (spontaneous activity/dormancy) in a given timestep, then whether or not it is active depends on resource availability. If maintenance resources are available, the PLANT is active; if not it is dormant.

The PLANT's genome dictates the dormancy strategy of the PLANT. This is given by a schedule of dormancy over 12 timesteps. Each slot in the schedule can be occupied by one of the three dormancy states, *Kpon*, *Kpoff* or *Ks*.

| 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 |
|---|---|---|---|---|---|---|---|---|---|---|---|
| Kpon | Kpon | Kpoff | Kpoff | Ks | Ks | Ks | Kpon | Kpon | Kpon | Kpoff | Kpoff |

Note that for every PLANT, if it does not receive its *Maintenance Requirements* at least once in 12 timesteps, it dies with 50% probability, hence preventing it from surviving forever in a dormant state.

The PLANT's genome is created via crossover from its parents' genomes. The three phenotypic states: *Kpon*, *Kpoff* and *Ks* are encoded as tertiary genotypes that can mutate with a certain probability (currently 1/1000).

### 3.2.2 Costs and benefits of Dormancy Strategies

The ability to go dormant has obvious benefits in any environment in which resource shortage occurs. We would like to examine the evolution of dormancy strategies under certain costs. In PLANTWORLD, dormancy has a cost in terms of the maintenance requirements of the PLANT, by effecting the *Maintenance* parameter, *k*. The scheme is as follows:

*Kpon* – no cost since Kpon implies that the PLANT is deterministically active;

*Kpoff* – costs are associated with the ability to go deterministically dormant;

*Kpon* – higher costs are associated with the ability to be spontaneously active or dormant.

In the current implementation these costs are set as follows:

*Kpon* – no effect on *k*;

*Kpoff* – increases *k* by 0.01 for every timestep given by the genome for which the PLANT's state is Kpoff;

*Ks* – increases *k* by 0.02 for every timestep given by the genome for which the PLANT's state is Ks.

The parameters 0.01 and 0.02 can be varied for different experiments. In the current implementation, the base value of *k*, $k_o$ is set at 0.5. Thus, *k* is raised by 0.01 for each timestep out of 12 that the PLANT is deterministically dormant, and by 0.02 for each timestep out of 12 that it is spontaneously active/dormant.

**Example:** for the genotype in 3.2.1 (above),

*Kpon* = 5/12 timsteps ; no cost

*Kpoff* = 4/2 timesteps; cost = 4 * 0.01 = 0.04

*Ks* = 3/12 timesteps; cost = 3 * 0.02 = 0.06

Costs of this dormancy strategy = 0.04 + 0.06 = 0.1, thus, $k = 0.5 + 0.1 = 0.6$.

Thus, the maintenance requirements of this PLANT, $m = 0.6 * CurrentSize$

### 3.2.3 Benefits of reaching the Water Table

The PLANT's *Maintenance Requirements* can also be effected by the *Water Table*. If a PLANT reaches the *ThresholdSize* defined by the user, then it has access to the *Water Table* should one be present in its cell. If the *Water Table* is present, the PLANT's *Maintenance Parameter, k,* is reduced by 0.1.

### 3.2.4 Evolution of AdultSize

The main stages in plant life history are growth and reproduction. While in many species these stages depend on the organism's age, in plants the reproductive stage is often more a function of size than of age (Bazzaz, 1996). In PLANTWORLD, the reproduction stage is therefore dependent on reaching a certain size, defined by the genome, called *AdultSize*. Depending on *AdultSize*, a given PLANT will have different experiences of environmental variability. A PLANT that becomes reproductive at small *AdultSize* may complete its life cycle in a short period of time. Hence, ephemeral PLANTS may adapt to variability by timing their life cycle to fit in with resource-abundant periods, such that they do not experience strong resource variability and do not require dormancy strategies for coping with such periods. Since we are interested in the evolution of dormancy, and since dormancy strategies will be highly dependent on the 'window' of environmental variation experienced by PLANTS, we would like to examine evolution in PLANTS that can also evolve different *AdultSizes*.

Thus, PLANTS have another gene that determines their *AdultSize*. Its *AdultSize* is given by a binary string of 12 characters, thus *AdultSize* ranges between 1 and 4096. Again, the PLANT receives its gene for *AdultSize* via crossover between its parents, with a certain mutation rate.

*AdultSize* also defines the maximum number of *Reproductive Events* a PLANT can achieve. PLANTS die as soon as they complete these, hence their lifespan is also defined by their *AdultSize*. The number of *Reproductive Events* is high for large PLANTS and low for small PLANTS, hence the lifespan of large PLANTS is concomitantly longer.

*AdultSize* is also used to enable speciation or niching since it determines who can reproduce with whom. PLANTS can only reproduce with other PLANTS that are some percentage larger or smaller than them. This percentage is currently set at +/- 5% but can be changed by the user.

### 3.2.5 PLANT Behaviour

PLANT behaviour is briefly summarised here. Detailed explanation of the methods and equations can be provided on request by the authors. PLANT behaviour includes maintenance, growth, dormancy, reproduction and death. In order to maintain, grow or reproduce, PLANTS must gain resources from the environment, i.e. the cell that the PLANT occupies. Each cell's resources are shared amongst the PLANTS in the cell. For simplicity, a cell's resources are not available to PLANTS in other cells. A small portion of the cell's resources are partitioned equally amongst all PLANTS (currently 10%) while the large portion (currently 90%) is partitioned in proportion to the PLANT's *CurrentSize*. These factors can be changed by the user. The share of resources available to the PLANT is called its *Available Resources*. PLANTS are also limited in how much moisture they can acquire, according to their *CurrentSize,* given by their *Maximum Utilisation Rate*. If *Available Resources* > *Maximum Utilisation Rate*, the PLANT acquires the amount of resources defined by its *Maximum Utilisation Rate*. If not, it acquires its *Available Resources*.

Once the PLANT has acquired resources, it must check out its dormancy status. If the PLANT is in state *Kpoff*, it does nothing this timestep. If it is in state *Kpon*, then it compares its *Acquired Resources* with its *Maintenance Requirements*. If these are met, it uses resources (left over from maintenance) for growth or reproduction, depending on its life-history stage. If *Maintenance Requirements* are not met, the PLANT may resort to its *Storage Capacity*. The PLANT's *Storage Capacity* depends on its *CurrentSize* and allows the PLANT to survive when in state *Kpon* without meeting its *Maintenance Requirements*, for a given number of consecutive timesteps. If the PLANT is too small to store moisture, or has used up its *Storage Capacity*, it dies with 50% probability (N.B. the probability of death can be defined by the user). If the PLANT is in state *Ks*, then it may be dormant or active depending on whether it can meet its *Maintenance Requirements*. If *Maintenance Requirements* are met, the PLANT grows or reproduces, depending on its life-history stage. If not, the PLANT remains dormant for this timestep.

If the PLANT is below *AdultSize*, it grows according to the amount of resources it acquires and its current *Growth Rate*. The *Growth Rate* is a function of the *CurrentSize* of the PLANT. If the PLANT is equal to or larger than *AdultSize*, then it can reproduce once every twelve timesteps. When the PLANT reaches the maximum number of *Reproductive Events*, defined by its *AdultSize* (see section 3.2.4) it dies.

Note that the model can be initialised with either SEEDS or PLANTS. When PLANTS are chosen, initial CurrentSizes are randomly chosen between 1 and the *AdultSizes* defined by the user. The user can define many parameters, some of which are listed below.

***User-definable parameters:***

Rainfall time series, spatial patterns, *Water Table* patterns

Initialising genotypes: *AdultSize* and *Dormancy Strategy*

Costs of *Dormancy Strategies*

*ThresholdSize*

Mating distances, seed number and dispersal distances, SEED Dormancy

*Evaporation Rate*

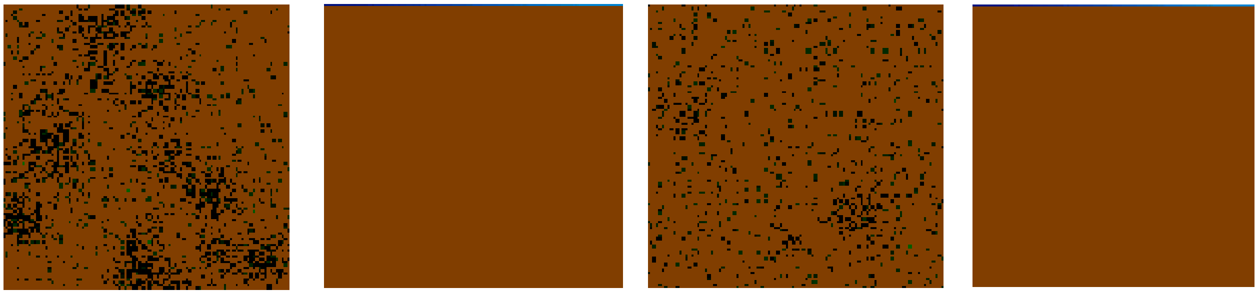

a) GENOTYPE 1
(i) Timstep 17
Number of PLANTS
2035

(ii) Timestep 210
Number of PLANTS
0

b) GENOTYPE 2
(i) Timstep 17
Number of PLANTS
1777

(ii) Timestep 206
Number of PLANTS
0

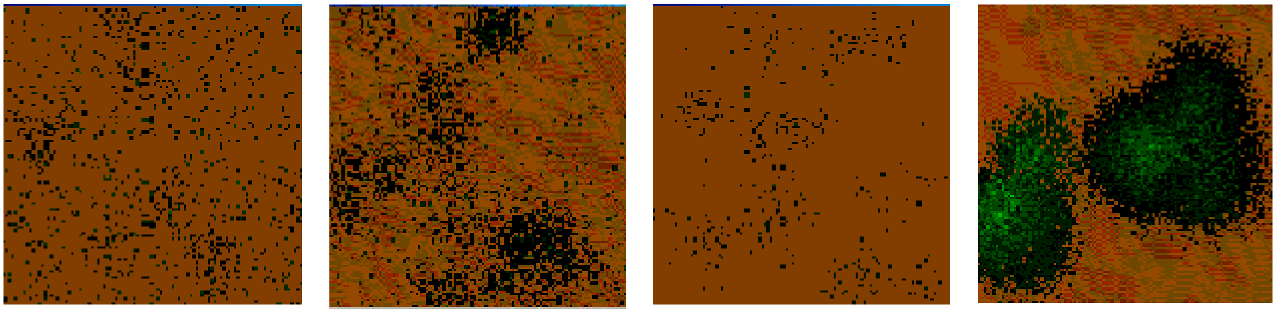

c) BOTH GENOTYPES
(i) Timestep 7
Number of PLANTS
1411

(ii) Timestep 21
Number of PLANTS
2269

(iii) Timestep 140
Number of PLANTS
273

(iv) Timestep 389
Number of PLANTS
23,648

**Figure 4.** The evolution of PLANTS using the Botswana rainfall data.
a) PLANTWORLD is initialised with GENOTYPE 1: *AdultSize* 100; *Dormancy Strategy*: 1 1 2 2 3 3 1 2 2 3 1 1
where 1= Kpon (active); 2 = Kpoff (dormant); 3 = Ks (condtionally active/dormant)
PLANTS go extinct, (i) → (ii)

b) PLANTWORLD is initialised with GENOTYPE 2*: AdultSize* 101; *Dormancy Strategy*: 1 1 2 2 3 1 3 2 2 1 1 1
PLANTS go extinct, (i) → (ii)

c) PLANTWORLD is initialised with both GENOTYPES:
Initially total PLANT population rises as older PLANTS reproduce, (i) →(ii), however, by timestep 140, populations have fallen dramatically, (ii) → (iii), recovering and expanding by timestep 389, (iii) → (iv). Note that simulations using only single populations of each genotype reached extinction over several (all) runs. It seems that while mutation rates were not high enough to produce viable phenotypes in single populations, when the environment is intialised with two genotypes, variation via crossover eventually produces a viable genotype, though until this point numbers fall dramatically. This simulation used Botswana rainfall data, and the environment was intialised with PLANTS of random size between 1 and the user-defined *AdultSize* (100, 101). The evaporation rate was set at 100%.

It should also be noted that in the current implementation, the only traits that vary amongst populations are the evolvable traits of *Dormancy Strategies* and *AdultSize*. The traits associated with these, i.e. *SEED Dormancy, lifespan,* number of *Reproductive Events* and the *Maintenance Parameter, k*, do not vary independent of the genome, being completely determined by the two evolvable traits. Other traits, such as *Growth Rate, Maximum Utilisation Rate* and *Storage Capacity*are fully determined by the size of the PLANT at a given timestep (*CurrentSize*). Thus, these traits are the same for all PLANTS and effectively, out populations are all the same *species* with regard to all parameters, except *AdultSize* and *Dormancy Strategy*. This implementation is legitimate because it allows the analysis of the evolution of dormancy strategies for different *species* (defined by *AdultSize)* uncomplicated by other variables that would considerably effect PLANT dynamics. Different implementations of PLANTWORLD could be used in the future to study dynamics under different evolvable traits or to increase the number of evolving traits as the model is developed.

## 4 PRELIMINARY RESULTS

Since we have yet to write statistical output programs for PLANTWORLD, to date we have not rigorously tested its full range of behaviour. However, graphical output and some basic statistics have been partially analysed for a range of input parameters. Environmental behaviour appears to meet requirements with cell moisture being updated according to rainfall, PLANT uptake and evaporation. PLANT and SEED behaviour also appears to meet requirements, with SEEDS germinating and PLANTS maintaining, growing, going dormant, reproducing and dying, as envisaged in the specification.

The patterns in Figure 4 show the graphical output from the PLANTWORLD system using Botswana rainfall data. Being a semi-arid region, Botswana's rainfall is expected to be highly heterogeneous over time and space, since coefficients of rainfall variation tend to increase as annual average rainfall decreases (Tyson, 1986). The environment is first initialised with one of the two single genotypes (1 and 2). In both cases (a and b), PLANT numbers initially increase but go extinct around timestep 210. In c) the environment is initialised with both genotypes. Initially ill-adapted PLANTS die out, but eventually PLANT numbers increase and begin to colonise the landscape. Since we have no statistical output to analyse yet, we can only assume that PLANTS are evolving better adapted *Dormancy Strategies* and *AdultSizes*, allowing them to survive resource fluctuations so that dynamics are somewhat stabilised. Thus, the evolutionary aspects of the model appear to meet requirements, since PLANTS seem to be evolving strategies to aid their survival in particular environments.

## 5 CONCLUSIONS

PLANTWORLD is a system designed to investigate ecological theories by modelling the population dynamics of adaptive agents in a digital environment. To date, only a prototype has been developed and preliminary results obtained. This is a work in progress, however, and an optimised version of the system will be completed shortly. It is anticipated that this system will help inform ecological theory by allowing us to incorporate adaptive behaviour in populations evolving in a spatially explicit landscape. Such approaches are likely to challenge standard ecological theory and provide new insights hitherto inaccessible to population and community models that could not incorporate evolutionary dynamics or adaptive behaviour.

### Acknowledgments

Thanks to Panash Shah for his implementation of the system.

### References

Abrams, P A (1986) The nonlinearity of competitive effects in models of competition for essential resources. *Theor. Pop. Bio* **32**, 40-65.

Bazzaz, FA (1996) *Plants in changing environments*, Cambridge University Press, Cambridge.

Chesson, P (1991) A need for Niches? *TREE* **6** 26 – 28.

Hulme, M (1998) *Monthly precipitation dataset for global land areas from 1900 to 1998*, gridded at 2.5° latitude by 3.75° longitude resolution.

Law, R and R D Morton (1996) Permanence and the assembly of ecological communities. *Ecol.* **73**: 567-578.

Levin, S A (1974) Dispersion and population interactions. *Amer. Nat.* **108**, 207-228.

Lotka, A J (1924) *Elements of physical biology*. Baltimore, Williams and Wilkins.

MacArthur, R H and E O Wilson (1967) *The theory of island biogeography*. Princeton, NJ, Princeton Univ. Press.

MacNally, R C (1995) Ecological versatility and community ecology. Cambridge, Cam. Univ. Press.

Martinez, N D (1992) Constant connectance in community food webs. *Am. Nat.* **139**, 1208 – 1218.

Mattessi C and M Gatto (1984) Does K-selection imply prudent predation? *Theor. Pop. Bio.* **25**, 347-363.

May, R M (1974) *Stability and complexity in model ecosystems.* Princeton, NJ, Princeton Univ. Press.

May R M (1975) Nonlinear aspects of competition between three species. *SIAM J. Appl. Math.* **29**, 243-253.

Odum, E P (1971) *Fundamentals of ecology.* 3rd ed. Philadelphia, Saunders Press.

Tyson, P D (1986) *Climatic change and variability in southern Africa*. Cape Town, Oxford Univ. Press.

Wiens, J (1984) On understanding a nonequilibrium world: myth and reality in community patterns and processes. In D R Strong (ed), *Ecological communities: conceptual issues and the evidence.* Princeton, NJ, Princeton Univ. Press.