

Evolution and the Regulation of Environmental Variables

Hywel Williams and Jason Noble

School of Computing, University of Leeds, Leeds UK
{hywelw, jasonn}@comp.leeds.ac.uk

Abstract. The idea that the biota can regulate the abiotic components of their environment to levels suitable for life has attracted criticism from neo-Darwinian theorists but is still a viable hypothesis. Here we present a model, similar to Daisyworld [1] but more general, which allows for a more extensive study of the compatibility of biotic regulation with evolutionary theory. Results obtained highlight the importance of constraints on the evolutionary process for the emergence of regulation, and set the scene for more comprehensive future study.

1 Introduction

The Gaia hypothesis, initially proposed by Lovelock and Margulis in 1973 [2] and developed further by subsequent debate in the literature, postulates the regulation of the abiotic environment by the biota so that the biosphere is maintained in conditions suitable for life. Although the initial idea of biotic feedback concerned the chemical composition of the atmosphere, the concept has been extended to other features of the Earth system, such as the carbon and nitrogen cycles or the N:P ratio in the oceans [3]. A significant amount of empirical evidence has been gathered on various facets of the Earth system, but the scale of the system under study and the sample size (of one!) has meant that no conclusive evidence has been found to either prove or disprove the hypothesis. Debate of the Gaia theory has therefore been largely theoretical, concerning the plausibility of such planet-wide regulatory mechanisms and the compatibility of regulation with Darwinian evolution.

The neo-Darwinian critique has shifted focus over the years. Early criticisms of the Gaia theory claiming teleology were rebuffed by the Daisyworld model [1], a numerical thought experiment which demonstrated how the temperature of a simulated planet might be regulated in the face of increasing solar radiation by a process of ecological competition between black and white daisy species. Since then the Daisyworld model has become the focus of debate; lacking a more approachable target in the real Earth system, protagonists in the theoretical debate have been forced to fight many of their battles over a model that was originally intended as a simple proof of concept.

While it should always be remembered that a Daisyworld is not a real world, a number of key conceptual clarifications have been made by its use. However,

the simplicity of the original Daisyworld model means that it cannot answer questions about the compatibility of biotic regulation and evolution. To address these questions a number of extensions to the model have been presented [4,5,6,7,8,9,10], but it seems likely that the Daisyworld model has now answered as many questions as it is able to.

Biotic regulation of abiotic variables is a general and widely applicable concept (see also the extensive literature concerning the related concepts of niche construction [11] and the ‘extended phenotype’ [12]), that extends far beyond daisies and temperature regulation, but it is not yet fully understood. Nowadays not many scientists would dispute that life affects the physical environment (itself a radical claim when Lovelock and Margulis first stirred up the Gaia debate), and vice versa, but the nature of the feedback between the two is unclear (witness the ongoing Gaian debate in the climate change literature [13]). It is difficult to study such processes in the real world, because of the aforementioned problems of scale and sample size, but simulation modelling can offer a useful tool. While models cannot prove anything about the real world, they can focus debate, generate hypotheses and allow the testing of assumptions as ‘opaque thought experiments’ [14].

This paper presents the first step in a longer-term project to model and explore how biotic regulation of abiotic environmental variables may evolve. The bulk of the paper will be concerned with developing a simplified version of Daisyworld that captures all known results. We further simplify the ‘cut-down Daisyworld’ model presented by Harvey [15] and extend it to a 2-dimensional cellular automata model amenable to the inclusion of Darwinian evolution. Existing results from earlier Daisyworld models are reiterated before presentation of some new results concerning the necessity of constraints on evolution if regulation is to emerge.

2 The Model

2.1 Overview

The original Daisyworld model [1] incorporates two species of daisy, identical except that one is black (with low albedo) and the other is white (with high albedo). Daisy albedo alters the local temperature of each daisy patch, with daisies assumed to live in single-species clumps large enough to maintain their own local temperature. This in turn alters the growth rate of the daisies, which varies as a function of temperature. Since the albedo of bare earth lies between the albedos of black and white daisies, population dynamics allow global temperature to move away from that expected of a dead planet. Competition between black and white daisies led to global temperature regulation around the optimal temperature for daisy growth; deviations away from this point were counteracted by negative feedback engendered by the selective advantage gained by one of the daisy species away from this point. Black daisies out-compete white daisies at low temperatures because of their ability to increase local temperature, and vice versa at high temperatures. Regulation was observed for a significant range of

solar luminosity, outside which the planet was too cold or too hot to support daisies of any colour.

The Daisyworld model used reasonably accurate approximations of the real-world phenomena on which it was based and thus incorporated quite complicated mathematical formulations of (for example) the interaction between solar luminosity and the level of heat radiation emitted by the planet. Harvey [15] developed a simplified model, his ‘cut-down Daisyworld’, that used much simpler approximations but conserved the essential regulatory behaviour of the system. Following Harvey and simplifying even further, we present the very basic model described below.

Our model is a cellular automata model in which patches are arranged in a 2-dimensional toroidal lattice (another CA Daisyworld model was presented by von Bloh et al [6], but in a different form and with different aims). Each patch may be barren (bare earth) or may contain a single species of daisy. Barren patches can be colonised by daisies from neighbouring patches, while living patches may die. Each patch has a local temperature that changes in relation to solar luminosity (applied at an equal level to all patches) and to its albedo (determined by the presence of daisies). The global temperature of the planet is taken as the mean of all the local patch temperatures. This scheme is covered in more detail below.

2.2 Daisies

A daisy species is represented by an albedo and a growth function. Albedo is drawn from the range $[0.25, 0.75]$ (representing a continuum from black to white), while the growth function is a piecewise linear function of local temperature that has the qualitative form shown below in figure 1. This hat-shaped function can be represented by the location of its centre point H_{mid} . In all the experiments reported here the hat function reached zero at $H_{mid} \pm 15$.

Daisies do not grow as such, since they are assumed to either fully occupy a patch or not to be present, but the growth rate of a daisy species determines its likelihood of colonising a neighbouring bare patch. High growth rates lead to increased colonisation.




Fig. 1. An example growth function. The growth rate of all daisy species varies from 0 to 1 as a piecewise linear function of temperature.

2.3 Seeding

An empty patch may be seeded with a new daisy species with low probability (0.03 in the simulations described below). When seeding occurs an entirely new daisy type is randomly generated from the set of permissible values for albedo and growth function parameters. Seeding allows new genetic stock to gain a foothold in the world and takes the place of mutation in the evolutionary process.

2.4 Colonisation

Empty patches may be colonised by daisy species living in neighbouring patches. Each neighbour species has a chance to colonise that is proportionate to its growth rate. This is implemented by assigning a probability $P(C_i)$ to the event that the empty patch is colonised by the i^{th} neighbouring patch (alive or dead), such that $P(C_i) = \frac{G_i}{N}$ (where G_i is the growth rate of the i^{th} neighbouring patch and N is the total number of neighbours) and noting that the growth rate of a dead patch is zero. Thus daisies with a higher growth rate have a higher likelihood of colonisation. Also, a daisy species occupying multiple neighbouring patches has a higher likelihood of colonisation due to having more ‘tickets in the lottery’.

2.5 Death

If a daisy species living in a patch has a growth rate of zero, it is assumed not to be able to survive and the patch becomes empty. Also, daisies living in a patch will die (and the patch become empty) with a probability of 0.1 at each timestep. This may be seen as a simple instantiation of death by natural causes and serves to promote selection and competition.

2.6 Calculation of Patch Temperature

Local patch temperature depends on the current temperature of the Sun (traditionally taken in Daisyworld models as a monotonically increasing value), the albedo of the patch (determined by daisy growth), and heat loss to space. The rate of change of local patch temperature is therefore given by equation 1 below, where T_P is the patch temperature, T_S is the temperature of the Sun, and α is the patch albedo.

$$\frac{dT_P}{dt} = (1 - \alpha)(T_S - T_P) - T_P \quad (1)$$

Patch temperatures are integrated numerically using Euler’s forward method. The global temperature of the planet is taken as the mean of all the patch temperatures.

2.7 Cellular Automata Update

The results presented below were gathered from a 10x10 toroidal CA where each patch has 4 neighbours at top, bottom, left and right. The CA is synchronously updated at each timestep by testing for colonisation, seeding and death in that order. T_S is typically increased from 100 to 500 in increments of 2, and the CA is updated for 1000 timesteps for each increment in T_S to allow the daisy population to stabilise for the new level of external forcing.

3 Repetition of Existing Daisyworld Results

First of all we compared the results generated from our model with known results generated from existing Daisyworld model. In all of the following experiments the albedo of bare earth was set to 0.5 and the world was initialised with all patches bare.

The primary Daisyworld phenomenon, that of temperature regulation by competition between black and white daisies [1], was considered first of all. We set the albedo of black daisies to 0.25 and the albedo of white daisies to 0.75. Results are shown in figure 2, which displays global temperature regulation occurring by competition between the daisy species as it does in the original work.

The next significant result to be repeated is that allowing albedo to mutate does not affect regulation, and may in some cases actually increase its range [4,5,6,7]. For this scenario we allowed albedo to take any value in the range [0.25, 0.75], corresponding to the full range from black to white. Temperature was regulated as before, although in this case it is by a steady shift in the albedo of the dominant daisy species to maintain the global temperature close to the optimal level, rather than competition between black and white daisies. The overall effect is the same at a global level; temperature regulation in this case and in the previous case is achieved by keeping the mean global albedo close to the level which keeps temperature optimal. This in turn is a result of selection for the daisy species with the highest growth rates.

Having shown that temperature regulation is not affected by mutation of albedo, the next result is to show that unconstrained mutation of the growth




Fig. 2. Daisy population, global albedo and global temperature for a world with both black (albedo = 0.25) and white (albedo = 0.75) daisies




Fig. 3. Daisy population, global albedo and global temperature for a world where daisy albedo is allowed to mutate freely between the levels for black and white daisies, i.e., within the range $[0.25, 0.75]$




Fig. 4. Daisy population, global albedo and global temperature for a world where daisy albedo is allowed to mutate freely between the levels for black and white daisies, i.e., within the range $[0.25, 0.75]$ and where the centre of the growth function is allowed to mutate freely

function causes the regulation to break down [9]. Here we do this by allowing H_{mid} to vary freely in the range $[70, 130]$. No regulation is observed, although the daisy population flourishes. The mutation in growth function simply tracks the solar forcing; the daisies adapt themselves to the environment rather than adapt the environment to themselves. The tracking is not precise, as the mutating albedo allows brief periods of quasi-regulation when the population becomes fixated on a particular growth function and uses the albedo to maintain the global temperature this value of H_{mid} requires. In this scenario albedo can be seen as a free variable, since a suitable growth function can be found to give optimal growth for any albedo level.

Lenton and Lovelock [10] showed that when there is some constraint on the mutation of the growth function, regulation will again emerge. They set up a Daisyworld model where the centre of the hat function was mutated towards the current ambient temperature, but where the maximum growth rate (i.e., the maximum height of the growth function) fell away to zero with distance from some optimal value, in a way supposed to be analogous with the decline in maximum achievable photosynthesis rate varies with temperature in plants. We implemented this by letting the maximum growth rate decline linearly to zero with distance from an optimal temperature of 100. We observed similar results to Lenton and Lovelock [10], in that regulation was observed to occur, but with a more gradual tailing in and tailing out than with the non-evolvable growth function.




Fig. 5. Daisy population, global albedo and global temperature for a world where daisy albedo is allowed to mutate freely between the levels for black and white daisies, i.e., within the range $[0.25, 0.75]$ and where the centre of the growth function is allowed to mutate freely. Maximum achievable growth rate declines linearly with distance from $T = 100$, reaching zero at $T = 100 \pm 30$.




Fig. 6. Regulation of global temperature for a world where there are two permissible growth functions and where daisy albedo is allowed to mutate freely between the levels for black and white daisies, i.e., within the range $[0.25, 0.75]$. Plots are shown for well-separated growth functions (centres at $T=65$ and $T=135$) 6(a) and for growth functions with overlapping ranges (centres at $T=85$ and $T=115$) where solar forcing increases 6(b) and decreases 6(c) over time.

For completeness, we have also run the model with heat transfer between neighbouring patches [6], and found that the qualitative nature of the results is unchanged for all of the above scenarios. The exception is where the heat transfer is so efficient that there is no perceptible difference in temperature between patches; in this case there is no possibility of a particular daisy species gaining a selective advantage over its competitors by altering its local temperature. When this occurs the regulation of global temperature is lost.

4 Constraints on Evolution and Their Implications for Environmental Regulation

It seems that the key criteria for regulation of global temperature to emerge are *need* and *ability*. Unless there is some reason for the daisies to alter their local environment, i.e., some selective advantage to be gained from doing so, then regulation will not occur. If evolution is added to the model, then the only cases in which daisies have a reason to alter their environment are those in which the evolutionary process is constrained in some way so that the daisy population cannot evolve to prefer the environment as it is. Selective advantage is gained

by improving the fit between daisy and environment; this can be achieved by changing the daisy or by changing the environment, and evolution will generally opt for the easiest method available.

Constraints on evolution are an inevitable feature of any real-world biological system, due to the existence of physical and chemical laws that no system may violate. Chemical laws constrain metabolism, the rate of which typically depends on a number of parameters as some bell-shaped curve. This idea is captured simply in Daisyworld as a growth function that depends on temperature by a Gaussian function, and in the current model by a piecewise linear hat function. While evolution cannot alter the chemical reactions involved in metabolism, it may tinker with the conditions under which those reactions operate to maximise their rate and efficiency, or it may select between different sets of reactions, that is, between different types of metabolism. However, once a metabolism has been chosen during the course of evolution it may often be easier to regulate the environment to suit this metabolism than to switch to a new metabolism entirely.

Different metabolic types may be more successful at different ranges of an environmental variable. In our simplified Daisyworld model, consider a situation where there are two growth functions with centres at different temperatures. The different growth functions may be well-separated, leading to independent regulatory epochs (figure 6), or have overlapping ranges, leading to competitive exclusion (figures 6(b) and 6(c)). When ranges overlap there will usually be one dominant metabolic type around which the environment is regulated, with a flip from one to the other at some critical level of solar forcing. The level at which this occurs depends on the history of the system. Whichever metabolic type becomes abundant first will stop the late-comer from getting a foot-hold in the ecology by holding temperature close to its own optimal level, and thus delay the onset of an ecology (and regulation) based around the other type. This is demonstrated by figures 6(b) and 6(c) which show competition between two growth functions with overlapping ranges in the face of increasing and decreasing solar forcing respectively (i.e., time flows to the right in 6(b) and to the left in 6(c), although forcing is plotted increasing left-to-right in both).

Another way in which evolution may be constrained and create an opportunity for regulation to evolve is if evolution operates at different rates on different phenotypic traits. Consider the case where the daisy growth function is free to mutate so that it can operate at any temperature and where daisy albedo may also mutate freely to any level between those for black and for white daisies. If both types of mutation occur at the same rate, then the growth function simply tracks the increasing solar forcing and regulation is lost (figure 4). However, if the mutation rate for the growth function is very slow compared to mutation rate of albedo the differential creates an opportunity for regulation. It is easier for a daisy species to evolve a new albedo than a new growth function. This can be observed in figure 7, in which the world is started with a viable daisy population that is then allowed to mutate. At each daisy reproduction (each colonisation of an empty patch), the daisy species may mutate its growth




Fig. 7. Daisy population, global albedo and global temperature for a world where daisy albedo is allowed to mutate freely between the levels for black and white daisies (i.e., within the range $[0.25, 0.75]$) with a probability of 0.2 at each reproduction and where daisy growth function can mutate freely with a probability of 0.002 at each reproduction.

function with probability 0.002 and its albedo with probability 0.2 (so albedo mutates two orders of magnitude faster than the growth function). As can be seen from figure 7, this results in regulatory epochs where the daisy population regulates the global temperature around the optimum for some growth function. Eventually the albedo can mutate no further and mutants with a more suitable growth function can out-compete the existing population to become established as the new dominant metabolic type around which regulation occurs.

5 Conclusion

We have presented a model that is derived from Daisyworld, but is simplified and extended to allow for a more comprehensive study of the compatibility of biotic environmental regulation with evolutionary theory. The model has been described here using the language of Daisyworld (daisies using albedo to regulate temperature in the face of solar forcing), but the mathematical formulation of the model is actually very general, allowing its possible use to study the concept of biotic regulation as a general phenomenon. Our model shows that what is needed for regulation to emerge are constraints on the evolutionary process and the possibility of organisms creating some local buffer against the global environment, criteria that we feel are plausible in a wide variety of biological systems. In future work we hope to move away from the daisy metaphor and look at multi-dimensional regulation in a more general sense.

References

1. Watson, A., Lovelock, J.: Biological homeostasis of the global environment: the parable of daisyworld. *Tellus* **35B** (1983) 284–289
2. Lovelock, J., Margulis, L.: Atmospheric homeostasis by and for the biosphere: the gaia hypothesis. *Tellus* **26(2)** (1973) 2–9
3. Volk, T.: *Gaia's Body: Toward a Physiology of Earth*. Springer-Verlag, New York (1998)
4. Lovelock, J.: A numerical model for biodiversity. *Philosophical Transactions of the Royal Society of London: Series B* **338** (1992) 383–391

5. Stocker, S.: Regarding mutations in daisyworld models. *Journal of Theoretical Biology* **175** (1995) 495–501
6. von Bloh, W., Block, A., Schellnhuber, H.: Self-stabilization of the biosphere under global change: a tutorial geophysiological approach. *Tellus* **49B** (1997) 249–262
7. Lenton, T.: Gaia and natural selection. *Nature* **394** (1998) 439–447
8. Saunders, P.: Evolution without natural selection: further implications of the daisyworld parable. *Journal of Theoretical Biology* **166** (1994) 365–373
9. Robertson, D., Robinson, J.: Darwinian daisyworld. *Journal of Theoretical Biology* **195** (1998) 129–134
10. Lenton, T., Lovelock, J.E.: Daisyworld is darwinian: constraints on adaptation are important for planetary self-regulation. *Journal of Theoretical Biology* **206** (2000) 109–114
11. Laland, K., Odling-Smee, F., Feldman, M.: Niche construction, biological evolution and cultural change. *Behavioral and Brain Sciences* **23** (1999) 131175
12. Dawkins, R.: *The Extended Phenotype*. Oxford University Press, Oxford (1982)
13. Volk, T.: Toward a future for gaia theory. *Climatic Change* **52** (2002) 423–430
14. Di Paolo, E., Noble, J., Bullock, S.: Simulation models as opaque thought experiments. In Bedau, M., McCaskill, J., Packard, N., Rasmussen, S., eds.: *Artificial Life VII: The Seventh International Conference on the Simulation and Synthesis of Living Systems*, Cambridge, MA, MIT Press/Bradford Books (2000) 497–506
15. Harvey, I.: Homeostasis and rein control: From daisyworld to active perception. In Pollack, J., Bedau, M., Husbands, P., Ikegami, T., Watson, R., eds.: *Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems, ALIFE'9*, Cambridge, MA, MIT Press (2004) 309–314