

# Homeostasis and Rein Control: From Daisyworld to Active Perception

Inman Harvey

Evolutionary and Adaptive Systems Group  
Centre for Computational Neuroscience and Robotics  
Centre for the Study of Evolution

COGS/Informatics, School of Science and Technology, University of Sussex, BN1 9QH, UK  
inmanh@cogs.susx.ac.uk

## Abstract

Homeostasis refers to the ability of organisms to maintain vital properties, such as body temperature, within a zone of viability, or of comfort, and the Gaia Hypothesis proposes that the Earth with its biota acts as a homeostatic whole. The Daisyworld model was proposed as one possible mechanism for providing this homeostatic regulation. Here a new and much simplified version of this model is presented, demonstrating that the combination of any ‘*Hat function*’ with any feedback, positive or negative, can lead to homeostasis through ‘*Rein Control*’. This principle is so general that it can be extended to other domains such as active perception, here demonstrated in a simulated robot.

## Introduction

Living organisms have many physiological variables that must be maintained within upper and lower bounds for continued survival. Typically there are regulatory mechanisms that maintain these variables within these bounds even in the face of substantial environmental perturbations. These homeostatic systems can take many forms. For mammals such as humans that need to maintain their body temperature within fairly tight limits, the variety of mechanisms includes physiological, reflex behaviour such as shivering and sweating, and more considered behaviour such as moving towards or away from heat.

Such regulatory mechanisms are important for the survival of the organism, and often appear to be complex, subtle, and so crafted as to provide near-optimal conditions for the organism; their origins and maintenance are usually attributed to the power of selection over many generations of Darwinian evolution. So controversy was inevitable when Lovelock (1972) proposed the Gaia Hypothesis that the Earth with its biota acts as a homeostatic geo-physiological system that regulates global properties, e.g. temperature at the Earth’s surface, within a range that provides viable living conditions for the biota; for a review see (Lenton 1998). No immediate explanation was available for the origin or cause of any such regulation, as clearly the Earth as a whole had not evolved through successive generations of selection within a population of Earths. Further, as species within the biota, and individuals within each species, were competing with each other, it was difficult to see how selection at the individual level could favour behaviour that led to global cooperation in

regulating a global variable. Individual acts of behaviour typically have a cost to the individual; yet the net effect from just one individual on global temperature is insignificant. So selection at the individual level will favour those profligates who do not care to do their bit towards global homeostasis. These arguments appeared initially to give sound evolutionary reasons why the global regulation proposed in the Gaia Hypothesis could not be maintained, or indeed even have arisen in the first place.

## Daisyworld

The Gaia Hypothesis was originally restricted to the claim (based on observation) that such global regulation existed, without any theory or mechanism to explain how it might happen; it initially met with much skepticism. The ‘parable of Daisyworld’ was then proposed as a possible mechanism (Lovelock 1983, Watson and Lovelock 1983).

Whereas the surface temperature of a lifeless planet would have changed dramatically with the increase in luminosity of the sun over geological timescales, observation indicates that on our Earth it has remained remarkably constant around temperatures suitable for life. Daisyworld is a deliberately simplified model of an imaginary planet with just two species of daisies, black and white, that demonstrates how this could happen. The growth rate of the daisies depends only on their local temperature, but in turn the daisies modify this because of differences in the way they absorb radiation; black daisies have low albedo (or reflectivity) and heat up easily, whereas white daisies with higher albedo tend to reflect the sun’s radiation back.

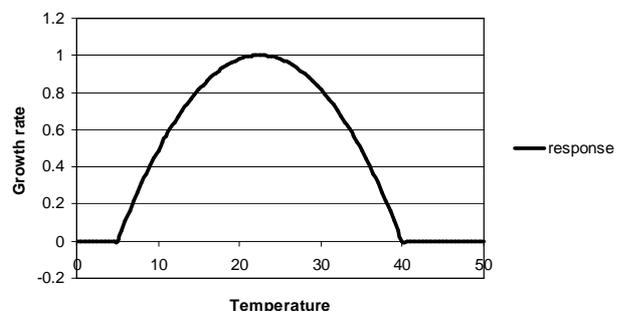


Figure 1. Hat function: daisy response to temperature

The dependency of growth rates on temperature is assumed to be a Hat-shaped function as in figure 1. The Daisyworld model demonstrates how these feedbacks via the environment, both positive for black and negative for white daisies, result in regulation of the planetary temperature. Watson and Lovelock (1983) demonstrate that with their particular parameter values and equations, the resulting close-coupled system regulates the temperature to within the viability zone (here  $5^{\circ}$  to  $40^{\circ}\text{C}$ ) over a far wider range of solar luminosities than would have been the case in the absence of any daisies.

### Comprehensibility of Daisyworld

Daisyworld is intentionally a simplified model; temperature is taken as just one example of an essential variable that can be regulated, and the lessons from the Daisyworld parable are meant to have far wider scope. One point made early on is that the precise form of the Hat function is unimportant, provided it has the general peaked character around an optimum temperature, with the brim of the hat representing here a zero growth rate outside the viability range. Nevertheless, the use of the Stefan-Boltzmann law to calculate absorbed and emitted radiation, and use of the equations governing the comparative and indeed competitive growth rates of the different daisies, means that analysis of this system is not a trivial problem.

I have therefore adopted the strategy of radically simplifying the usual Daisyworld model, to see how much can be left out whilst still retaining the homeostatic regulation. In doing so, firstly it becomes clearer how crucial is the difference in local temperatures between black and white daisies, something often obscured by the conventional graphs shown; secondly it becomes much easier to visualize the very simple underlying feedback interactions; thirdly it becomes plainer just how much can be generalized from this one example to other domains.

### Organisation of Paper

In the following sections I shall start by describing the cut-down version of Daisyworld. Visualisations of the conditions for steady state will be shown. Then results obtained through computer simulations integrating the equations to a steady state will be given. The conditions and parameter values will be manipulated to see just how far they can change whilst maintaining robustness.

I shall draw some very general conclusions, and to demonstrate their generality apply them to a very different domain of active perception. Here a simulated robot is supplied with oriented light-sensors that display a similar Hat function response to a light source. Feedback directly coupled to this response will change the orientation of the light-sensor in (a random choice of) either positive or negative direction. Collectively the coupling between many such individual light-sensors determines the global orientation of the robot. The result is homeostasis in the sense that collectively the system acts so as to maintain, as far as possible, the light-sensors oriented to the light so as

to stay within their sensitive regions; in other words, phototaxis despite the random nature of the feedback.

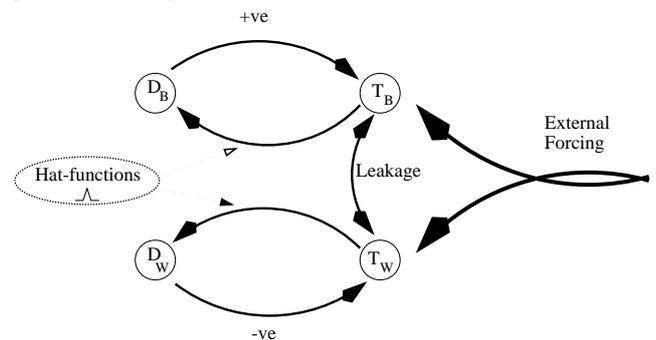


Figure 2. Interactions in the cut-down model: a black daisybed above, separate white one below. Both receive external forcing from the sun, and the only interaction between them is by 'leakage' or heat conductance.

### Cut-down Daisyworld

The simplifications are twofold: firstly, the model is idealized into a simpler form with fewer interactions; secondly, the Hat function and the putative underlying laws of radiation and heat conductance are simplified into linear or piecewise linear form.

The black and white daisies can be modeled as growing on separate daisybeds, in other words not competing for space. The interactions and feedbacks are then limited to those shown in figure 2. The Hat function can be replaced by a piecewise linear function of similar general form, which I call a 'Witch's Hat' function.  $T_B$  and  $T_W$  are the average temperatures of each daisybed,  $D_B$  and  $D_W$  are the proportional coverage of each daisybed by black and white daisies respectively, as determined by the Witch's Hat function. The temperature  $T$  of each bed (taken as uniform within the bed) is determined by a combination of factors: external forcing by the sun, feedback (positive or negative) proportional to  $D$  (the coverage of daisies in that bed), and a 'leakage' factor whereby some (parameterized) proportion of the heat flows from the hotter to the cooler daisybed. Unlike the original model, there is no direct interaction between  $D_B$  and  $D_W$ .

### A Single Daisybed

Initially we can simplify still further by just looking at the behaviour of a single daisybed; consider one half only of figure 2.  $T$  is the temperature of a bed with albedo  $\alpha$ ,  $S$  the temperature of the Sun, and deep space is at zero temperature. In a simplified, distorted version of physics, heat flow into the bed from the Sun is  $(1 - \alpha)(S - T)$ , and out of the bed into deep space is  $(T - 0)$ , i.e.  $T$ .

$D$  is now the quantity of daisies (rather than growth rate), which varies according to a Hat function of the local temperature  $T$ :  $D=H(T)$ . In the simplest version, where we assume there is feedback linearly proportional to  $D$  to raise

(or lower) the local temperature  $T$ , this feedback is  $uD$ : for black daisies  $u$  is a positive feedback to increase the temperature, for white daisies  $u$  is negative.

The rate of change of flowerbed temperature is

$$\frac{dT}{dt} = (1-\alpha)(S-T) - T + uH(T)$$

Equilibrium is when the rate of change is zero:

$$0 = (1-\alpha)S - (2-\alpha)T + uH(T)$$

$$H(T) = \frac{(2-\alpha)T - (1-\alpha)S}{u}$$

For fixed  $S, u, \alpha$ , this is linear in  $T$  with zero value when

$$T = \frac{(1-\alpha)S}{(2-\alpha)}$$

The line has slope  $(2-\alpha)/u$ . The equilibrium points are where this straight line crosses the Hat function.

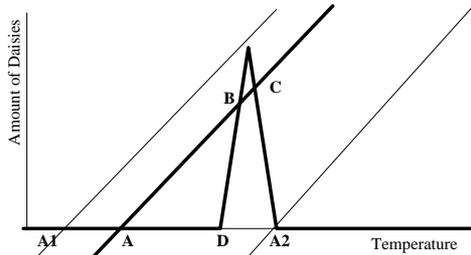


Figure 3. Slope meets Witch's Hat function at A, B, and C.

In figure 3 the heavy sloping line has a positive slope (black daisies for positive feedback) and crosses  $H(T)$  at A, B, and C. A implies no daisies, B proves to be an unstable equilibrium, but C is a stable non-zero equilibrium. For different amounts of external forcing from changing luminosity of the sun, the sloping line shifts, in parallel fashion, along the temperature axis. There will be a stable equilibrium point C, with a positive quantity of daisies, for any such line between the lighter sloping lines through A1 and A2 in the figure; these indicate the limits for intersecting the RHS of the Hat. This implies that the range of viability that allows some daisies to survive extends all the way from A1 to A2, rather than the limited range D to A2 available if there is no feedback. The slope of the line reflects the degree of feedback, with a vertical line corresponding to  $u=0$ . The stronger the feedback, the larger  $u$  is, the further away the slope is from vertical; and hence the further away to the left the viability range is extended. A1 lies at a distance  $u/(2-\alpha)$  to the left of the central optimum temperature of the Witch's Hat, if one takes the maximum height of that hat to be scaled to 1.0. If the line has negative slope (white daisies for negative feedback) then the mirror image case holds, and the range of viability is extended out to the right instead. So regardless of the sign of the feedback, the range of viability is extended.

### Rein Control

In a thoroughly mixed metaphor, in figure 3 the line AC can be thought of as a rein pulling the zone of viability

towards A, rather than its default left limit of D. Clynes (1969) put forward the notion of *Rein Control*, in commenting that biological systems typically have (at least) two channels for sensing and regulating variables: one (or more) in one direction from the norm, another in the other direction. This notion has received relatively little currency, although it is taken up in recent work drawing ideas from Daisyworld theory and applying them in modified form to physiological control (Saunders et al. 1998). The rein metaphor is appropriate as a rein can only pull, not push. Hence for control in both directions we need a further feedback loop, as in the following extension to the simulation; we need both reins.

### Two Daisybeds

For the simulation, we assume 2 daisybeds whose bare ground is grey with albedo 0.5. One bed can support only black daisies with a lower albedo (typical value used 0.0); the other can only support white daisies (typical albedo 1.0). For each bed, the average albedo depends on the proportion of cover by black daisies ( $0 \leq D_B \leq 1$ ) or white daisies ( $0 \leq D_W \leq 1$ ). The consequent temperatures, assumed to be uniform across each bed, are  $T_B$  and  $T_W$ . These temperatures are then potentially modified by heat transfer between the beds from the hotter black one to the white. This transfer is parameterised by a factor  $0 \leq L \leq 1$ . When  $L=0$ , no heat transfer takes place, but if  $L=1$  then the beds each have their temperature modified to the mid-temperature  $(T_B + T_W)/2$ ; for intermediate values of  $L$  the temperatures are scaled linearly between these two extreme cases.

To find through computation any non-zero stable equilibrium point for fixed values of  $S$  and  $L$ , we initialise  $D_B$  and  $D_W$  to 0.5, and then iterate this loop:

1. Calculate albedo for each bed from  $D_B, D_W$ .
2. Calculate  $T_B$  and  $T_W$  from  $S$  and these albedos.
3. Adjust these temperatures by the between-bed heat transfer, or 'leakage', parameterised by  $L$ .
4. Use the Hat function to calculate  $D'_B = H(T_B)$  and  $D'_W = H(T_W)$ .
5. Adjust  $D_B$  and  $D_W$  a small proportion of the way towards these new values  $D'_B$  and  $D'_W$  by:
$$D \leftarrow (1-\delta)D + \delta D'$$
for a small value of  $\delta$ .
6. Go back to 1.

$\delta$  should be chosen small enough to ensure that the values change smoothly over successive iterations of this loop, and then the loop must be repeated sufficiently many times until the changes in values at each iteration are vanishingly small. In practice it was found, for the range of parameters used here, that  $\delta = 0.0001$  and 200000 iterations of the loop made further changes in the variables invisible at the level of double precision floating point numbers.

Results are shown in graphical form for various values of the parameter  $L$ . In each case the resulting equilibria are

shown across the full range of external forcing by the sun, as it varies from excessively cool to excessively hot.

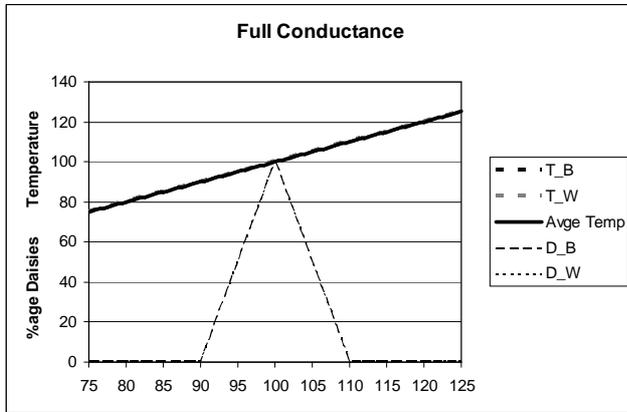


Figure 4. Conductance or Leakage  $L=1$ . The superimposed lines for  $D_B$  and  $D_W$  indicate the Witch's Hat function, with a viability zone between 90 and 110 on the lower scale. Vertical axis indicates temperatures ( $T_B$ ,  $T_W$  and average, here all superimposed) and also percentage of Daisies in each bed ( $D_B$ ,  $D_W$ , here also superimposed).

**Maximum heat conductance.** If  $L=1$ , the temperatures of black and white beds are identical. Since they follow the same Hat function, there is always the same number of black daisies in one bed as white daisies in the other. Hence the overall average effect is that of grey, albedo 0.5. In other words, there is no net feedback, and (regardless of how many daisies there are) the temperature is the same as if there were none. Maximum conductance means minimum, or zero, homeostasis. See figure 4, where the horizontal axis indicates the sun's output, scaled according to the corresponding temperature of a lifeless planet; here the temperatures  $T_B$  and  $T_W$  are the same as this.

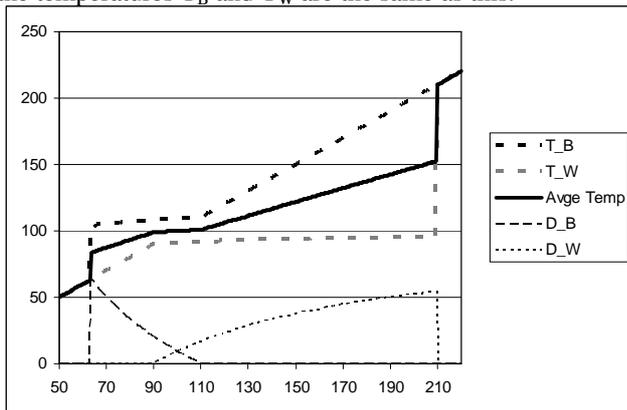


Figure 5.  $L=0$ . Daisybeds are independent, and only in the range 90-110 on horizontal axis do both daisies coexist.

**Minimum conductance.** When  $L=0$ , the two daisybeds are completely unconnected, and behave as if they were separate planets each regulating itself; the black bed extends its viability only towards lower sun temperatures, the white bed only towards higher. See figure 5.

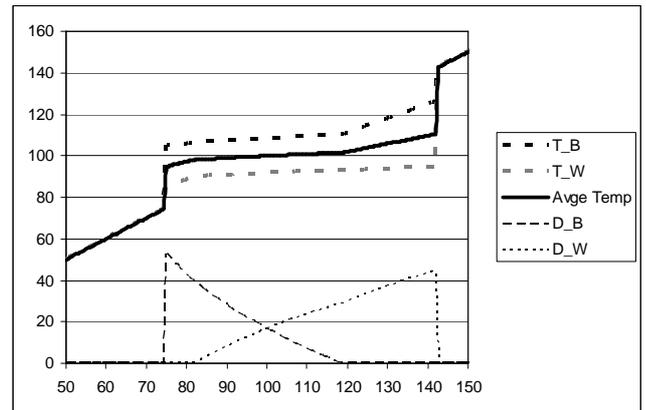


Figure 6. Conductance  $L=0.5$ . Coexistence of both daisies between 82 and 118.5

**Intermediate conductance.** Figure 6 shows the more general picture, where although the extension of black daisies left, white daisies right, is not as far as in figure 6, the range of coexistence of both daisies is greater. Note that at all times that either type of daisy is viable, the black daisies are hotter than their optimal temperature of  $100^{\circ}$ , and the white daisies are cooler than this.

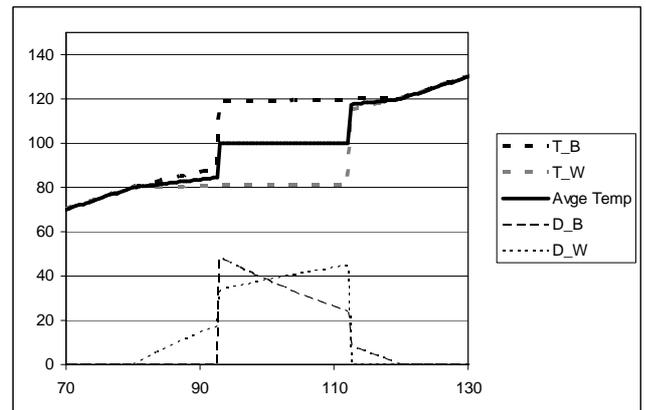


Figure 7. The daisies are given different, narrower, Hat functions: black's from 115-120, white's from 80-85. There is coexistence of both daisies between 93 and 112.

**Moving the Hat functions.** Figure 7 shows that there is still homeostasis when the viability zones (Hat functions) of the daisies are shifted relative to each other. Note from the figures in the caption, black is here shifted towards the hotter end, white towards the cooler end.

## Moving to Active Perception

We have seen above how the simple yet powerful combination of a Hat function with a feedback loop (either positive or negative) produces *Rein Control*, and in the context of Daisyworld, homeostasis; the Hat function directly relates to the idea of a zone of viability, and Rein Control tends to regulate a system to stay within it. The cut-

down Daisyworld has reproduced the basic homeostatic results of the usual version; though the simpler equations used means that it does not reproduce the phenomenon whereby the average planetary temperature actually decreases slightly as solar luminosity increases.

Such a powerful principle can be extended to other domains, and here it is demonstrated with active perception in a simple simulated 2-dimensional robot. Despite the very different domain, the underlying principles are identical.

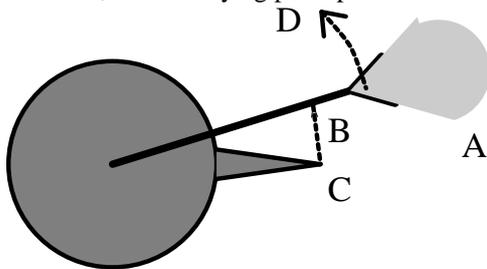


Figure 8. Circular agent can only rotate about its centre; orientation is indicated by nose C, currently facing East. One tentacle is shown, with sensor at end A.

One sensory tentacle is shown in the plan view of figure 8. For fans of *Doctor Who*, think in terms of a *Dalek*. The tentacle rotates around the centre of the robot, and has a sensory angle of acceptance as indicated at the end A. If a light source passes across this receptive field, the sensor response is given by a Witch's Hat function, with maximum response when the tentacle points directly at the source. This response produces a torque D on the tentacle A, in this case shown as left or counter-clockwise. D is counteracted by a restraining spring B attached to the nose C. Two parameters, specific to this tentacle, modulate the torque response D and the spring modulus B.

Now consider 100 such tentacles, each with different randomly chosen angles of acceptance, different directions and torque parameters for D, different spring constants for B. As a light source passes in front of them, each tentacle will respond independently. But collectively they will be held together via the springs attached to the nose, some pulling in one direction, some in the other. In the absence of any light, all these tentacles will be drawn together over the nose, but in the presence of any stimulation they will rotate apart independently in different directions, restrained only by the springs. The resultant balance of these spring forces on the nose will rotate the robot as a whole in one direction or the other, about its centre.

The translation from Daisyworld is direct as far as the underlying equations go, even though conceptually it is quite a leap. Each tentacle corresponds to a daisy species; on average half will provide feedback in one direction and half the other. The Hat function on sensor response corresponds to daisy dependence on temperature. The springs correspond to the leakage or conductance of heat between daisybeds, and the resulting direction of the nose corresponds to the average planetary temperature. The dynamics of the motion of each tentacle, and consequent

rotation of the robot, are calculated in simulation with similar liberties and simplifications used for modeling the physics as were used in modeling Daisyworld, whilst respecting the general principles. The end result is that, despite the random parameterization of all the tentacles, this robot efficiently performs phototaxis.

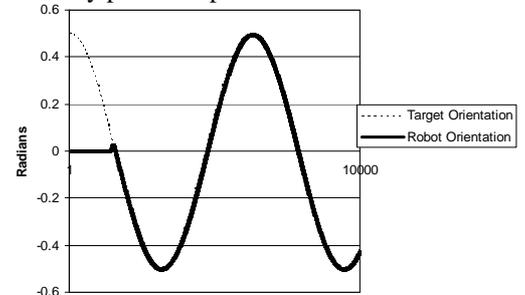


Figure 9. Time runs horizontally across this graph from the left, and the heavy line indicates the orientation of the nose in radians. The lighter line indicates a light source passing in wide, sinusoidal fashion across the front of the robot.

As indicated in figure 9, the robot will immediately pick up on a passing light source, and track it so accurately that thereafter the plots coincide. Further testing shows that this behaviour is exceptionally robust to changes in the allowed ranges for the randomly chosen parameters. The maximum angle of acceptance can be allowed to vary over 3 orders of magnitude, out as far as 3 radians (or nearly  $180^\circ$ ) each side. The upper limit on torque parameters and spring constants can be allowed to vary over more than 2 orders of magnitude; phototaxis is still reliable.

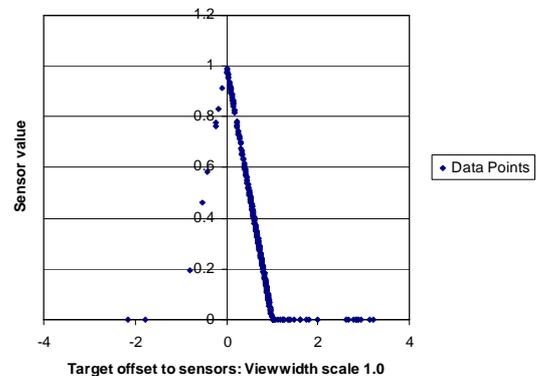


Figure 10. Samples from all the left-moving sensors whilst the robot is performing phototaxis. Each data point indicates on the vertical scale the sensor response, calibrated to a maximum of 1.0, and on the horizontal scale the angular offset of target source to tentacle direction. Although different tentacles have different angles of acceptance, this offset is here rescaled so as make all the individual Hat functions coincide on this graph.

**Analysis.** Although each tentacle can move independently, and each response to sensory stimulus is in a random direction, their collective coupling means that almost all the tentacles will stay approximately oriented towards the light source nearly all the time. While unqualified teleological

language is just as inappropriate here as it is in Daisyworld, we can carefully say: "Although this robotic system only functions this way as a whole, through multiple feedbacks, it can seem to a casual observer **as if** the tentacles are trying to maintain their sensory stimulation; just as in Daisyworld it might seem **as if** the daisies are trying to regulate the temperature so as to stay within the viability zone."

Figure 10 shows samples of sensory inputs from all the counterclockwise or left-moving sensors during a run. The outline of the Witch's Hat function is clearly visible, with almost all the data points on the right-hand slope, which is where Rein Control is acting for regulation in that direction. The data from right-moving sensors is the mirror image, giving the second of the pair of metaphorical reins.

## Discussion

Many people have been mystified as to how homeostatic regulation is achieved in the Daisyworld model. How could such regulation have arisen, since surely it requires some care in setting up the feedback structure and the parameters? An evolutionary origin appears unrealistic. Indeed any ongoing system that includes biota and *seems* to require global collaboration *seems* susceptible to exploitation by evolution of sub-groups towards cheating.

The cut-down version of Daisyworld presented here makes several useful pedagogical points. The analysis of a single daisybed (see figure 3) shows that extension of the range of viability in one direction arises from the simple interaction of a feedback of any sign (the sloping heavy line) and a Hat function. The feedback need not be linear, though normally it should be monotonic; the Hat function can be anything to suit your millinery tastes, although the Witch's Hat seems near-ideal from a mathematical stance. Simple feedbacks are universal in natural systems, and Hat functions are also widespread; the zone of viability associated with any homeostasis automatically implies a Hat function, and so does the typical response of any active sensor. So no special design process needs to be postulated for the basic phenomenon illustrated here with a single daisybed. We should observe such systems everywhere.

The idea of Rein Control (Clynes 1969) deserves wider currency. This phenomenon, like a rein, can only 'pull' and not 'push'. So for homeostatic regulation in both directions we need feedbacks in both directions. In the active perception example above, the 100 simulated sensor feedbacks in random directions would average out at roughly half in each direction, so any system that includes Hat functions and many arbitrary feedback loops is likely to incorporate both reins of Rein Control.

In cut-down Daisyworld, the interaction between black and white daisybeds is limited to 'leakage' and its role is made clear. Too much coupling means the opposing homeostatic tendencies will nullify each other, whereas being uncoupled would imply, in effect, separate planets. So some intermediate loose coupling between the different systems is essential, but no further global organisation is needed. In this simple version we can observe directly the

(uniform) temperature of the black bed, and likewise that of the white bed. By plotting these separately for didactic purposes, it is easier to appreciate the importance of the difference between these temperatures, something obscured in much of the previous literature where typically only average planetary temperatures have been displayed.

The original Daisyworld model has extra layers of complexity on top of this cut-down version, so it is of interest to see what is common to both. Homeostatic regulation is already apparent in the simpler version, but the phenomenon whereby average planetary temperature can actually decrease slightly as solar luminosity increases is only seen in the more complex version.

The plots shown in the various figures indicate that the black daisies are almost always living in a hotter climate than their optimum temperature. They are mostly on the right slope of the Witch's Hat; and vice versa for white daisies. Figure 10 shows the equivalent for the perception example. What if daisy evolution allowed either species to modify their metabolism, and so 'shift their viability zone' (subject to underlying physical constraints) in the direction of the climate they actually experience? Then the black Hat function would shift to the right, and the white one to the left, as shown in figure 7.

The principles shown here are very simple and of wide applicability. They do not require an evolutionary origin or explanation, but may be quite compatible with evolution.

**Acknowledgment** This work was motivated and stimulated by the participants in the series of EPSRC-funded workshops *Daisyworld and Beyond*, organised by Tim Lenton and myself in Sussex and Edinburgh, 2001-03.

## References

- Clynes, M. 1969. Cybernetic implications of rein control in perceptual and conceptual organization. *Ann. NY Acad. Sci.* 156: 629-670
- Lenton, T. M. 1998. Gaia and natural selection. *Nature* 394: 439-447
- Lovelock, J. E. 1972. Gaia as seen through the atmosphere. *Atmos. Environ.* 6: 579-580.
- Lovelock, J. E. 1983. Gaia as seen through the atmosphere. In: P. Westbroek and E. W. d. Jong. *Biomineralization and Biological Metal Accumulation*. Dordrecht: D. Reidel Publishing Company, 15-25.
- Saunders, P. T., Koeslag, J. H., and Wessels, J. A. 1998. Integral rein control in physiology. *J. Th. Biol.* 194: 163-173.
- Watson, A. J. and Lovelock, J. E. 1983. Biological homeostasis of the global environment: the parable of Daisyworld. *Tellus* 35B: 284 -289.