Homeostasis and Dynamical Representations

Inman Harvey

Evolutionary and Adaptive Systems Group

EASy, Dept. of Informatics

University of Sussex

inmanh@cogs.susx.ac.uk

Two Parts

A talk in 2 disjoint parts:-

1. **(90%)** Homeostasis and the Dynamics of Daisyworld, based on Harvey 2004
2. **(10%)** The Dynamics of Representations – comments and open questions

Alife 9 Talk Title:

Homeostasis and Rein Control:

From Daisyworld to Active Perception

Note on Motivation: most people looking at Daisyworld are using ideas of homeostasis drawn from organisms as a way of understanding global climate issues – "geophysiology" …. My interest is the reverse!

Talk Plan

1. 4-page Quick Summary, defining all the words in the Title.
2. Original Daisyworld.
3. New Kindergarten Daisyworld.
4. Rein Control, general principles that can … ....
5. Extend to phototaxis in a Dalek-like simulated robot.
6. Conclude.

Quick Summary (1 of

Homeostasis

Organisms have feedback control mechanisms for maintaining conditions vital to their comfort and survival.

Too cold?	Shiver, warm clothes, go to Florida.
Too hot?	Sweat, strip off, air-conditioning

Many would argue that such *homeostasis* is central to the very concept of life.

E.g. *Autopoesis* is *homeostasis* of ones identity as an organisation.

Rein Control

Little-known principle in physiology, put forward by Manfred Clynes *(musician, neuroscientist, coiner of the term 'Cyborg', … …)*

"When a physiological variable is regulated against being both *too high* and *too low*, different mechanisms are used for each direction".

You need two reins to control a horse, one rein can only pull but not push.

Quick Summary (3')

Daisyworld

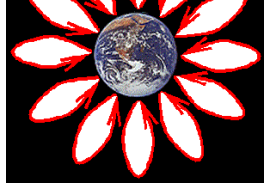

Gaia Hypothesis, Lovelock 1974 :- "the biosphere - atmosphere, oceans, climate, Earth's crust and biota, living organisms, is regulated as a homeostatic system in conditions comfortable for the living organisms"

How? Why? Teleology? Magic?

Daisyworld model, Lovelock 1983 :- Simple Artificial Life model presenting a possible Gaian mechanism, for e.g. temperature regulation.

This paper :- a new simplification of the Daisyworld model, showing how Rein Control leads to homeostasis. Confirming Lovelock, opening up new generalisations.

Quick Summary (4)

Active Perception

One generalisation will be the use of Rein Control and Homeostatic principles in a simple example of *Active Perception* in a light-seeking Animat (simulated robot)

Active Perception :- use of active movement of sensors in order to perceive

In Daisyworld, feedback and Rein Control keeps critical variable such as temperature within a viability range … …
In the Animat it keeps active sensors focussed on a light - *phototaxis*

Talk Plan

1. 4-page Quick Summary, defining all the words in the Title.
2. Original Daisyworld.
3. New Kindergarten Daisyworld.
4. Rein Control, general principles that can … ....
5. Extend to phototaxis in a Dalek-like simulated robot.
6. Conclude.

Context of the Original Daisyworld model

Gaia Hypothesis, Lovelock 1974 :- "the biosphere - atmosphere, oceans, climate, Earth's crust and biota, living organisms, is regulated as a homeostatic system in conditions comfortable for the living organisms"

One example :- Our Sun is heating up, it was say 30% less luminous 3.8bn years ago. By rights, it should have been far too cold for life then, and far too hot now (e.g. 290^0C)

But it seems the Earth's surface temperature has been maintained at around 20^0 C for aeons. A nice temperature!

HOW ?

Interactions between Planet Temperature and Li

Gaian Hypothesis :- somehow interactions between living organisms and the rocks / oceans / climate produce this homeostasis -- "let's model this"

Firstly, as temperature varies, Life has a *preferred temp* and a *viability zone,* such as this :- Similar to :-

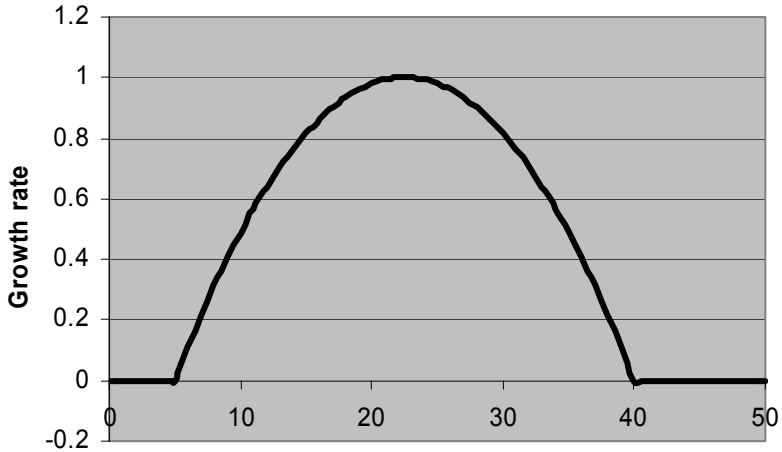

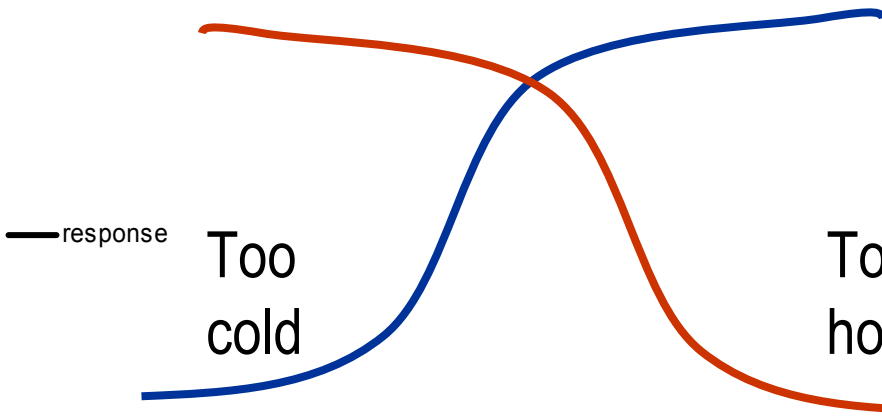

Interactions between Life and Planet Temperatur

Secondly, the existence of biota, of living things, affects the planet temperature.

E.g. on earth, phytoplankton in oceans generate a gas (DMS) which affects cloud cover which affects solar input.

Some of these interactions give **positive**, some **negative** feedback-components

*(-- in fact **both** will give homeostasis!)*

Terminology Alert !

Control theorists often use “positive (or negative) feedback” as shorthand for “positive (or negative) feedback **circuit**”

+ leads to runaway increase/decrease

— leads to stability and homeostasis

But here I refer to **+ve** or **–ve** feedback as just **one component** of a two-part circuit

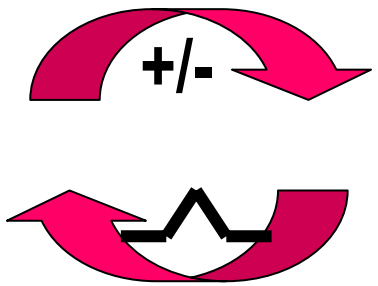

Original Daisyworld (Lovelock 1983)

To model this, we assume a grey planet can support **B**lack and/or **W**hite Daisies if their local temperature is right.

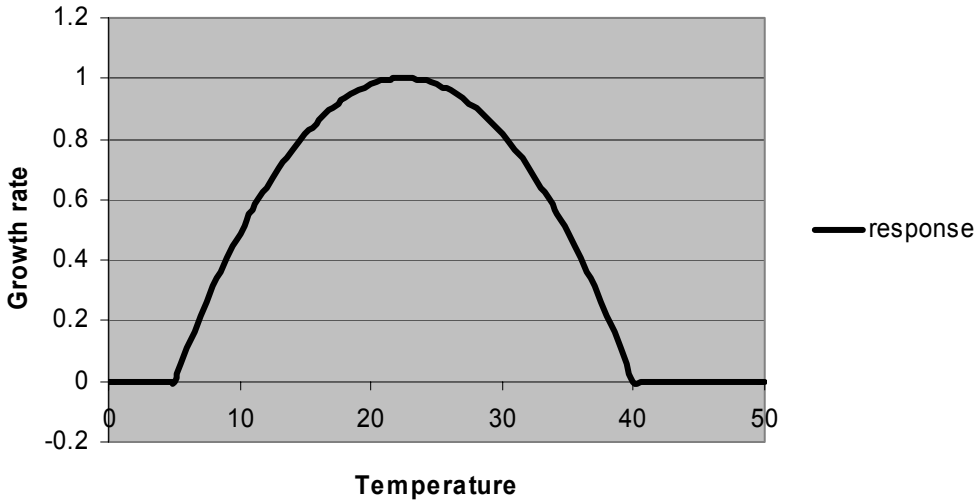

E.g. viable between 5^0C and 40^0C with preferred temp 22.5^0C

B and **W** have different albedos (reflectivity) and increase / decrease the local temperature (**+ve** or **–ve** feedback)

Homeostasis in the model

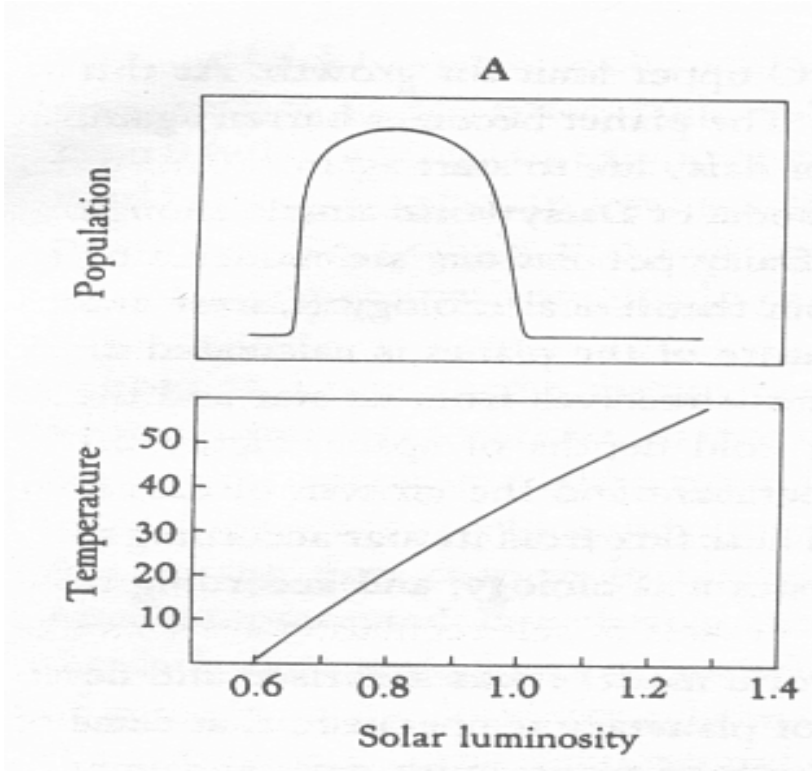

Consider what would happen as Solar luminosity increases →

In the absence of feedback, one expects planet temperature to increase smoothly, and Daisies to increase then die away

But, if you factor in the feedbacks, the result is very different!

Temperature Homeostasis

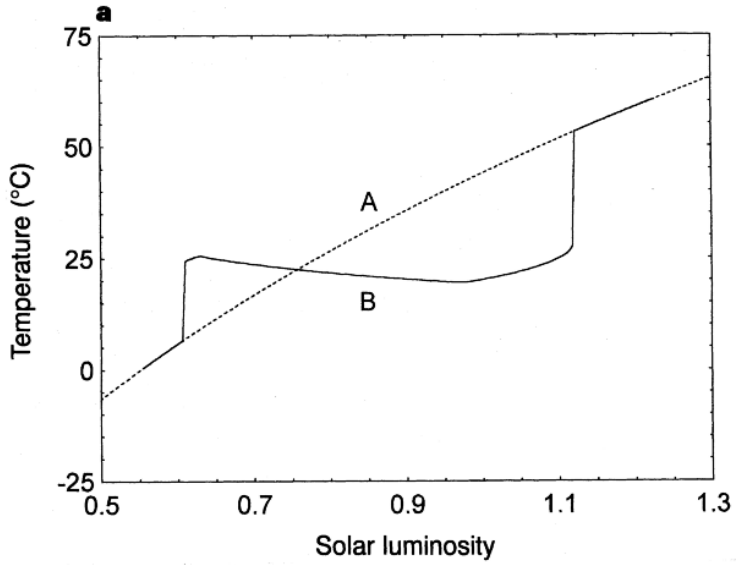

The planet temperature is maintained within the viability range as luminosity increases over a wide range – indeed it *decreases* slightly !

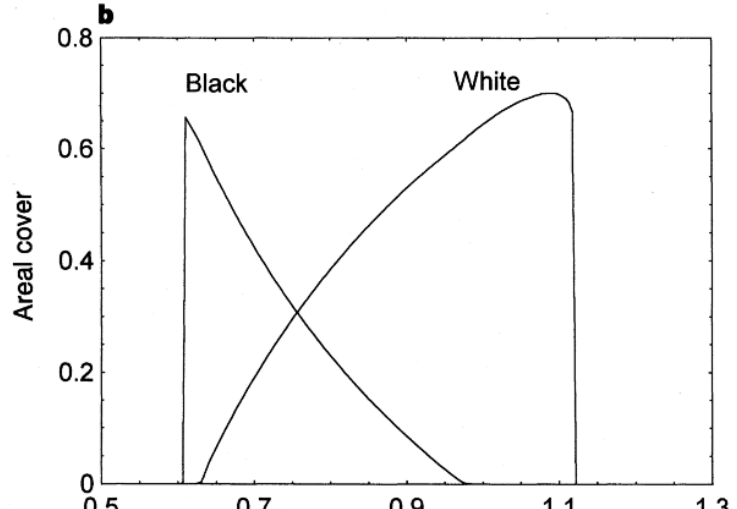

Black flourish at low luminosity so increasing temperature

White flourish at high luminosi so decreasing temperature

Underlying Maths of the model

The Lovelock Daisyworld model calculates heat flows according to Solar luminosity and the albedos of Black/White Daisies and Grey planet, using the Stefan-Boltzmann law for radiation absorption/emission.

The Black/White Daisies are also competing for space – in fact it is all rather complex to visualise.

So I have simplified like crazy, and produced my own new kindergarten version … … …

Talk Plan

1. 4-page Quick Summary, defining all the words in the Title.
2. Original Daisyworld.
3. New Kindergarten Daisyworld.
4. Rein Control, general principles that can … ....
5. Extend to phototaxis in a Dalek-like simulated robot.
6. Conclude.

New Kindergarten Daisyworld

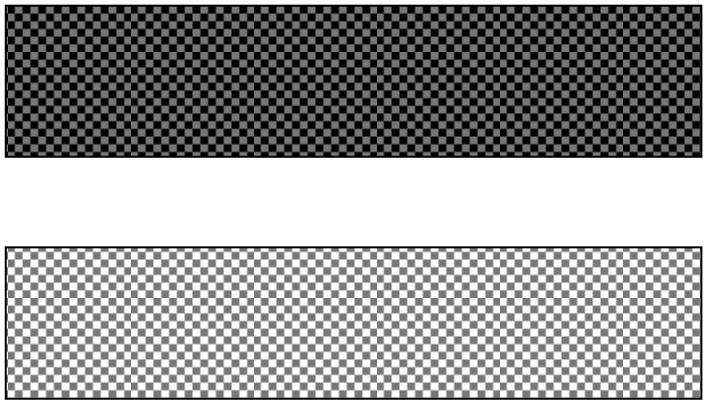

"Let's model the Black Daisies as in one Grey daisybed, the White in another, no longer competing for space"

"Let's assume heat flows depend **linearly** on temp diff **(S-T)**, modulated by Daisybed albedo, and on temp diff **(T-0)** to deep space"

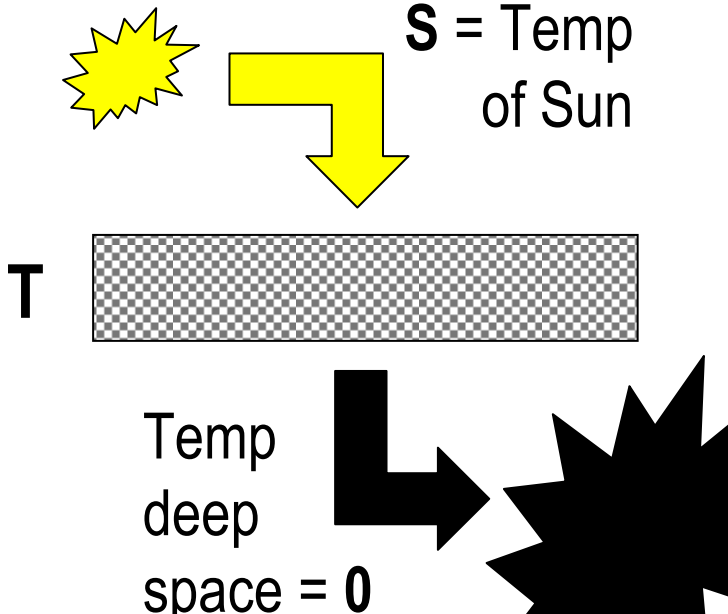

Hat Functions

Homeostasis goes with viability

"The Daisy viability function, as temperature varies, is a Hat-shaped function."

"Any Hat function will do"

"Let's use a Witch's Hat, directly for the number of Daisies, not Growth Rate"

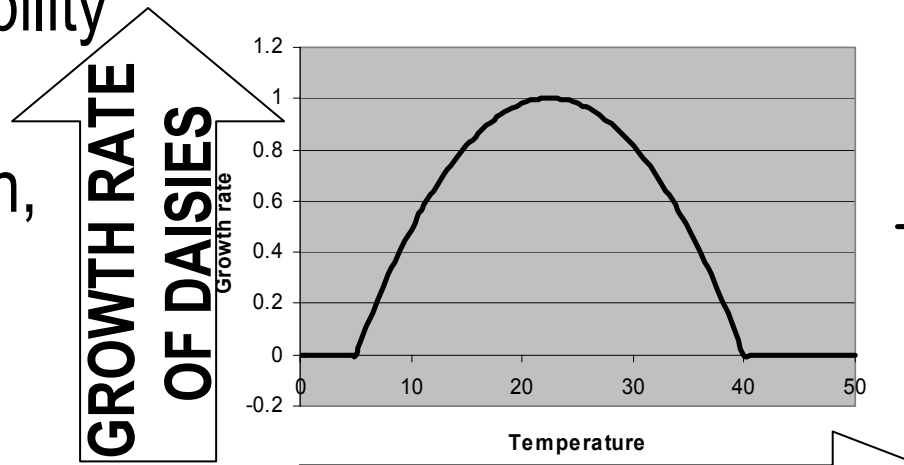

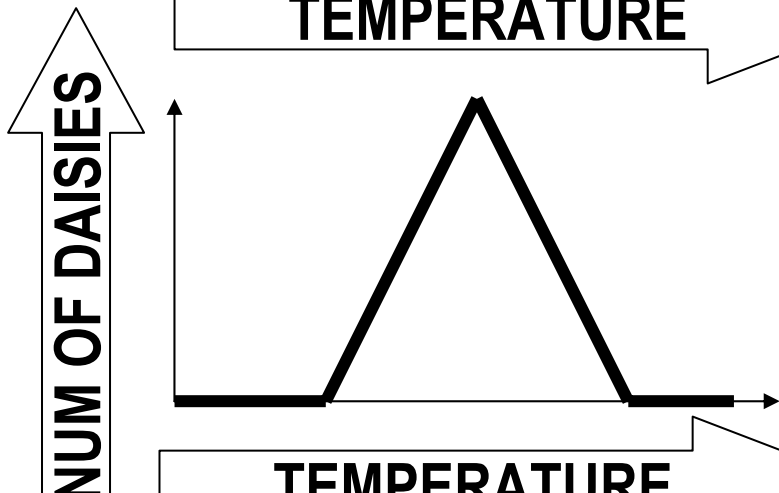

Consider just the Black Daisybed

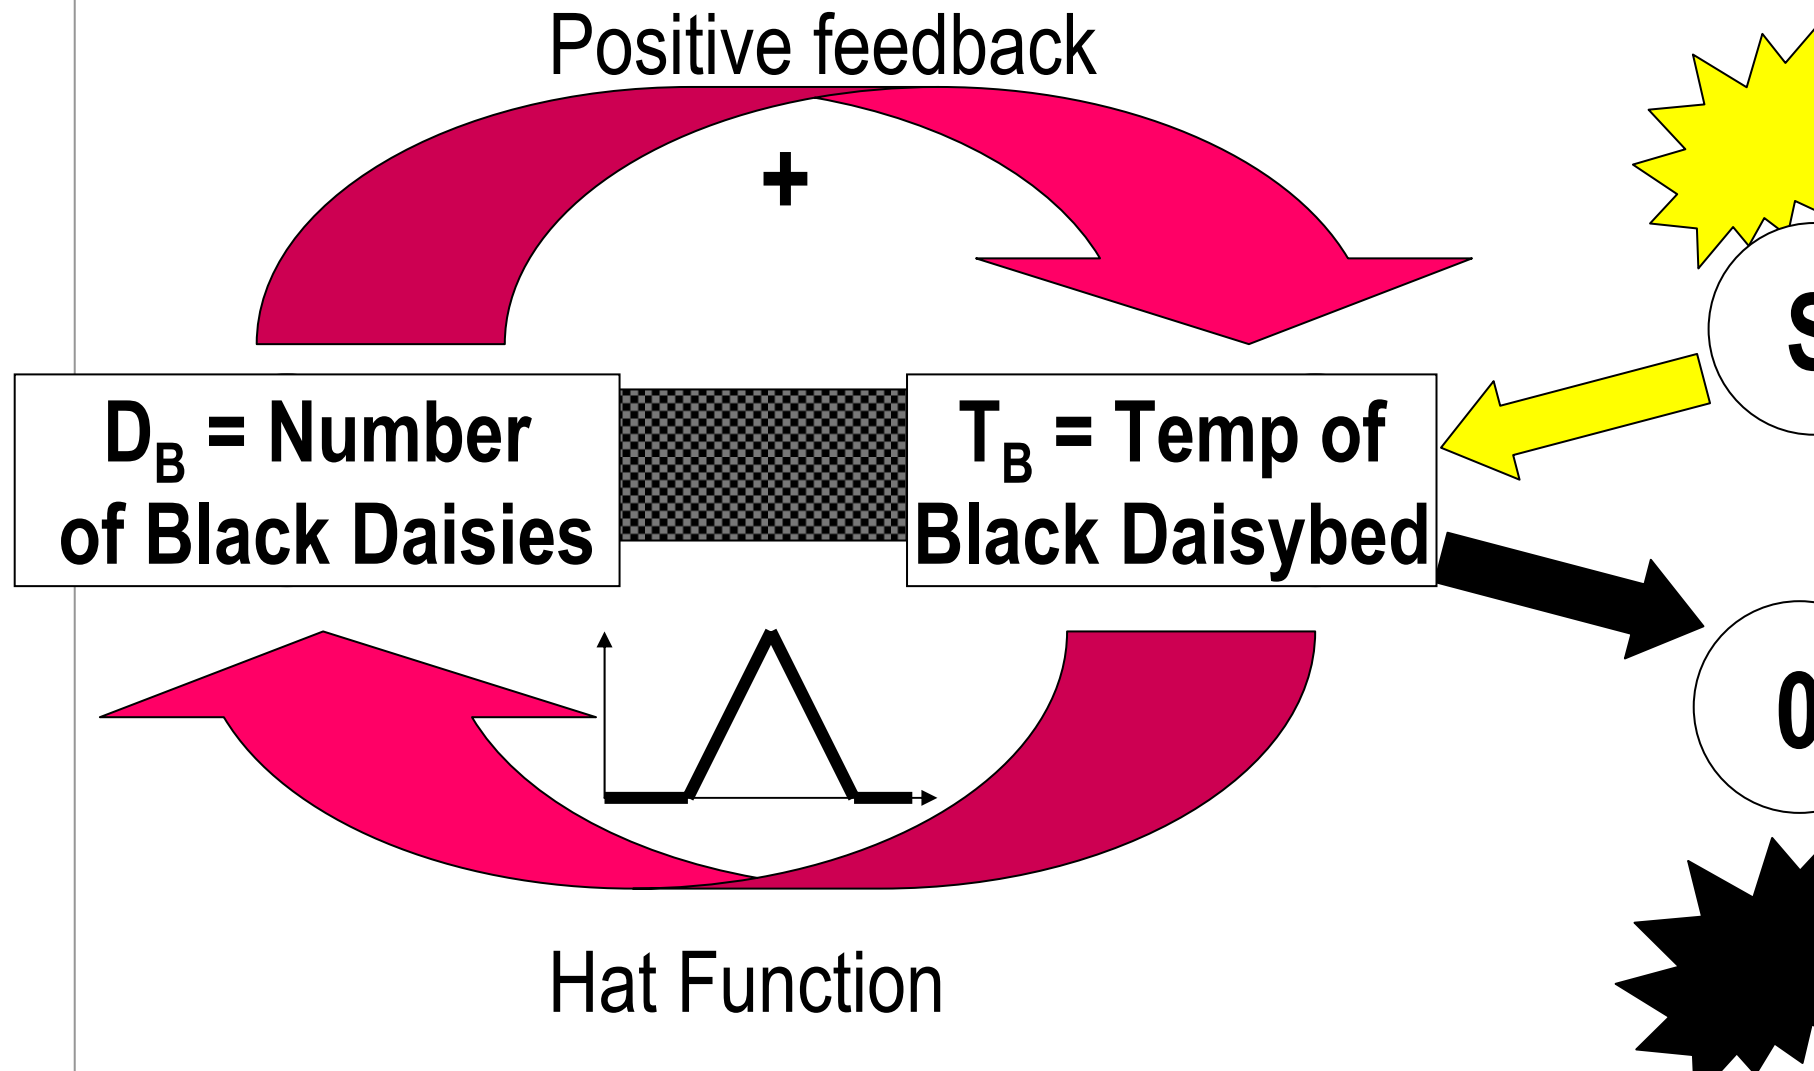

Just one Daisybed

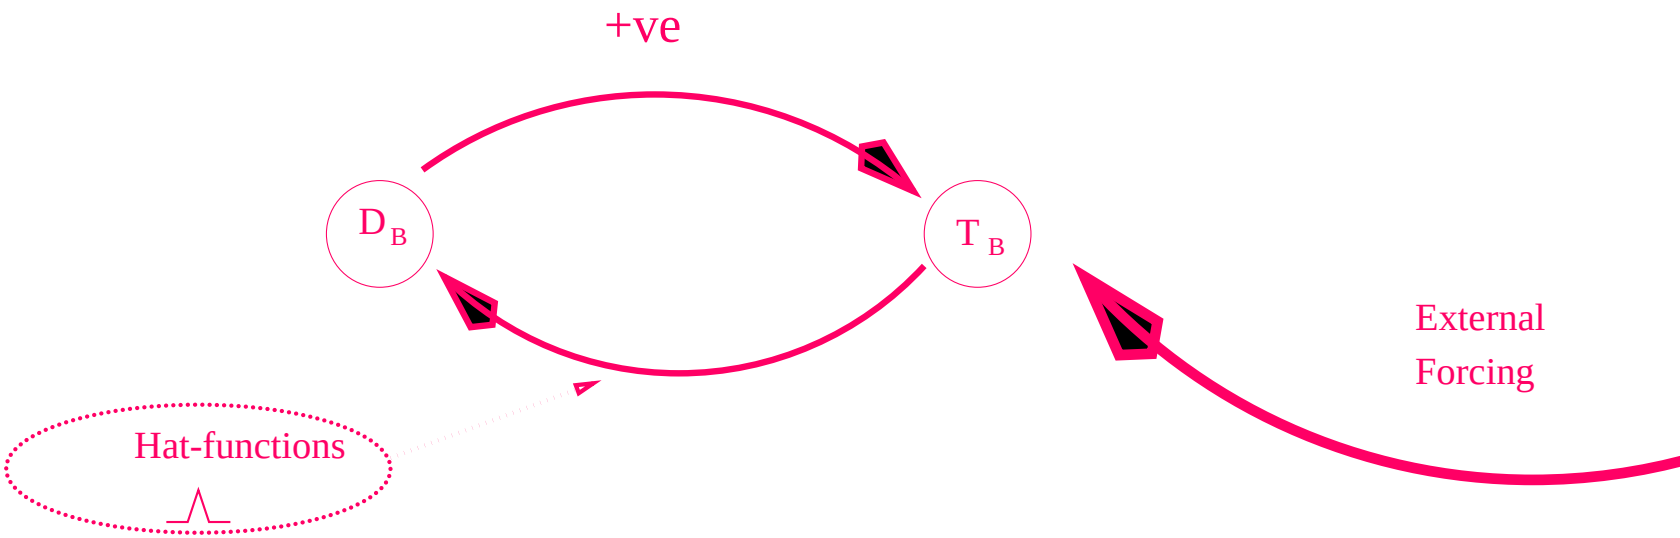

$$\frac{dT}{dt} = (1-\alpha)(S-T) - T + u.H(T)$$

Rate of change of Temp = albedo * (Suntemp – Temp) –Temp + Feedback-term * Hat-function(Temp)

Equilibrium when …

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

LHS = 0 when a Linear function in T intersects a Hat-function of T

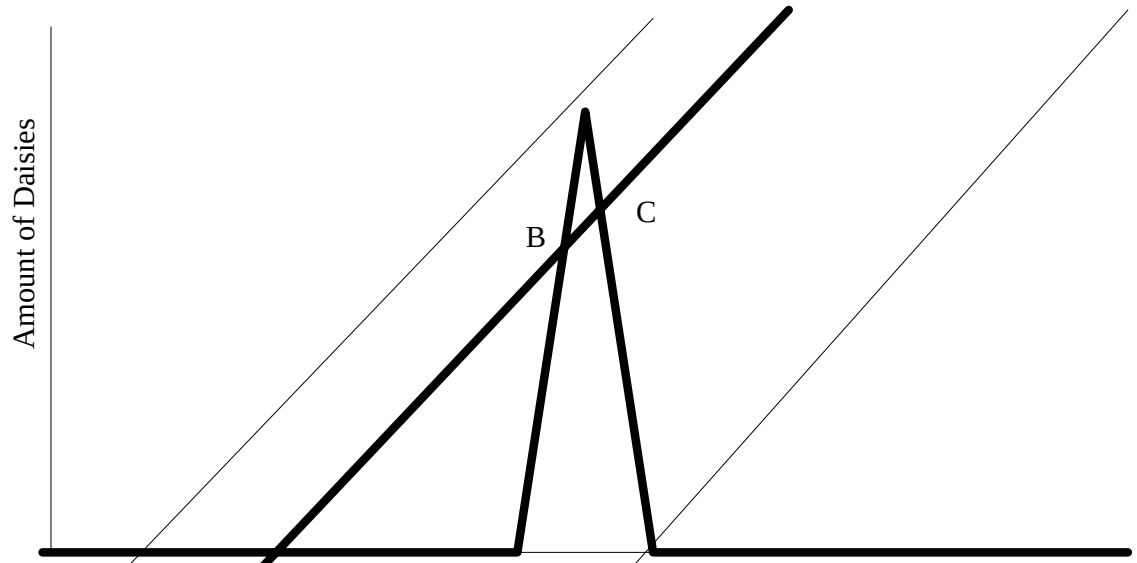

The Maths says Temperature settles at … …

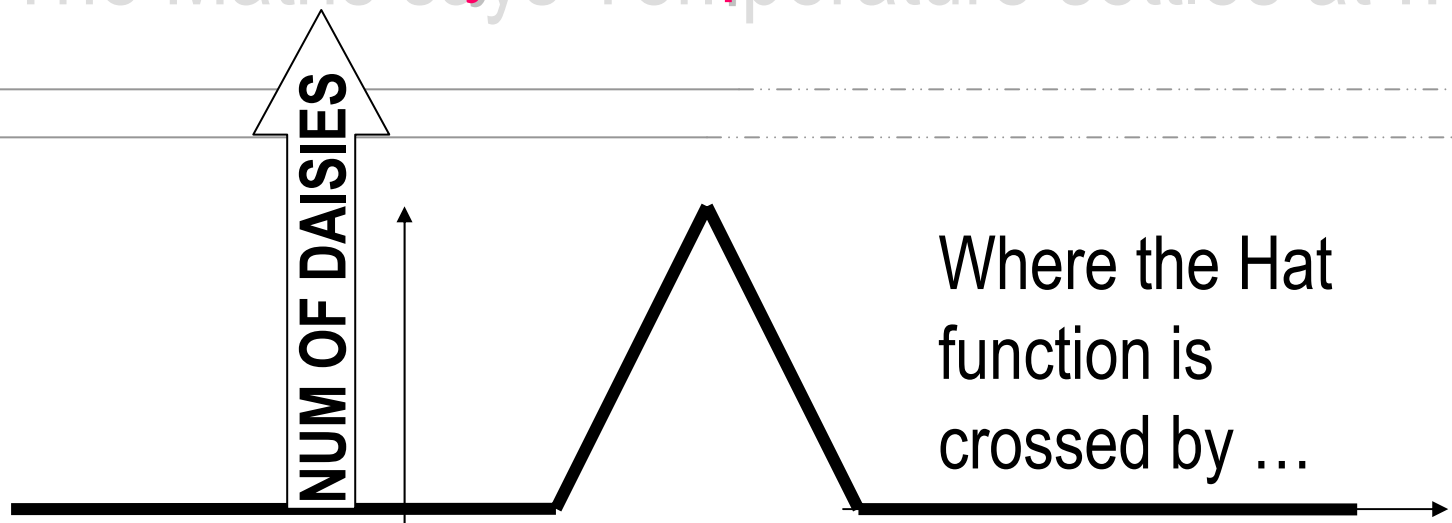

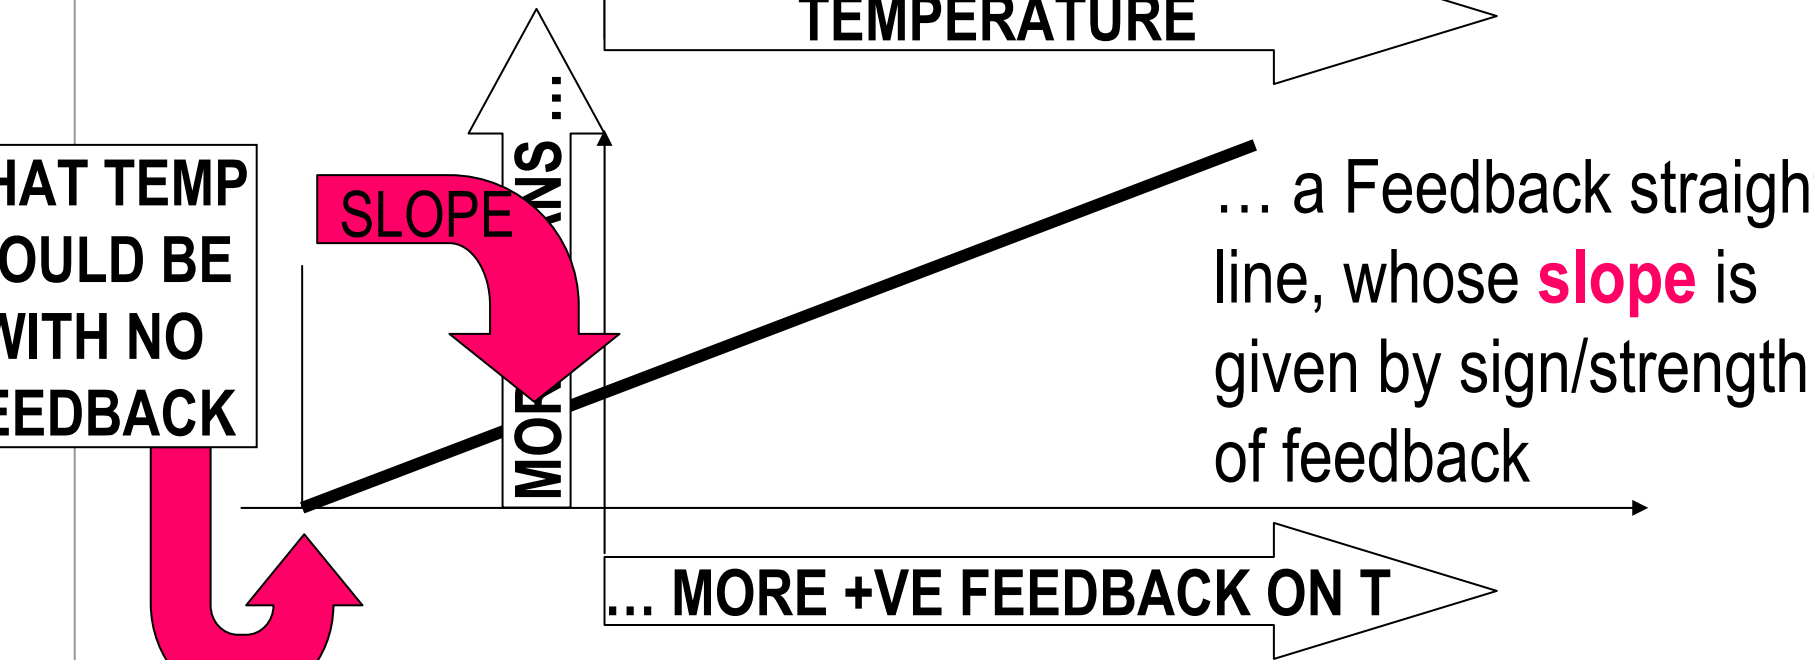

So Possible Temperatures are … …

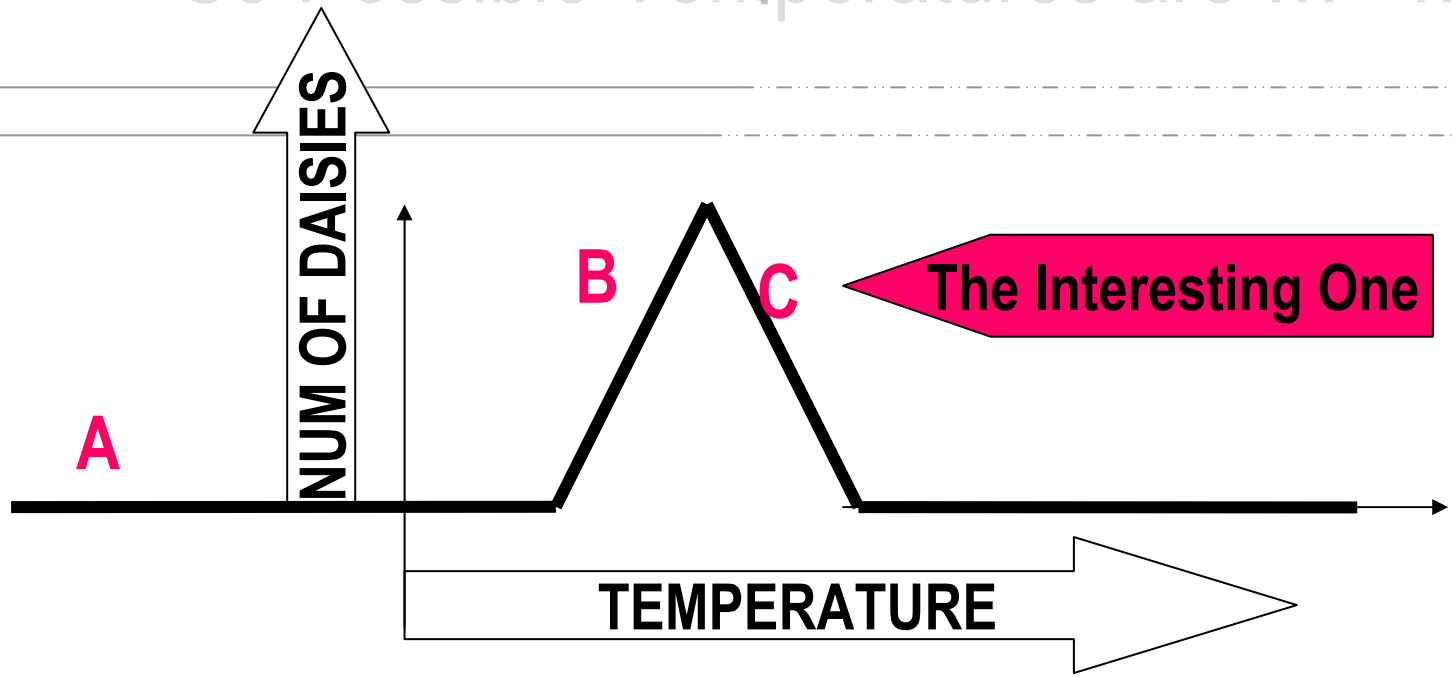

A is a low temperature, with no Black Daisies alive

B turns out to be an unstable equilibrium

C is a stable equilibrium, the +ve feedback from Black Daisies brings Daisybed temperature within the viable range

This extends the Zone of Viability

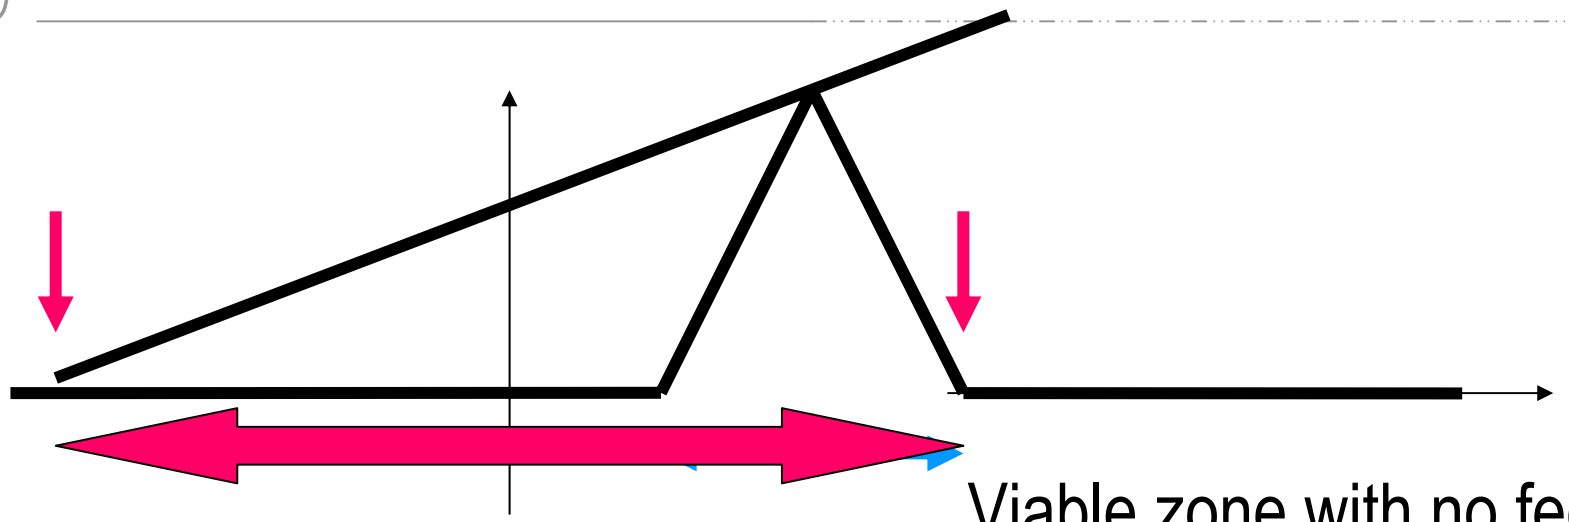

Viable zone with no feedbac

Shifting this slope left or right corresponds to changes in the external forcing of the Sun's luminosity – let's see how far it can shift and still intersect RHS of Hat Function … …

Viable zone **with** the positive feedback

There is a bigger range of sun luminosities (extended left) that can support viable daisy temperatures, because of the positive feedback from Black Daisies absorbing extra heat.

White Daisies give Negative Feedback

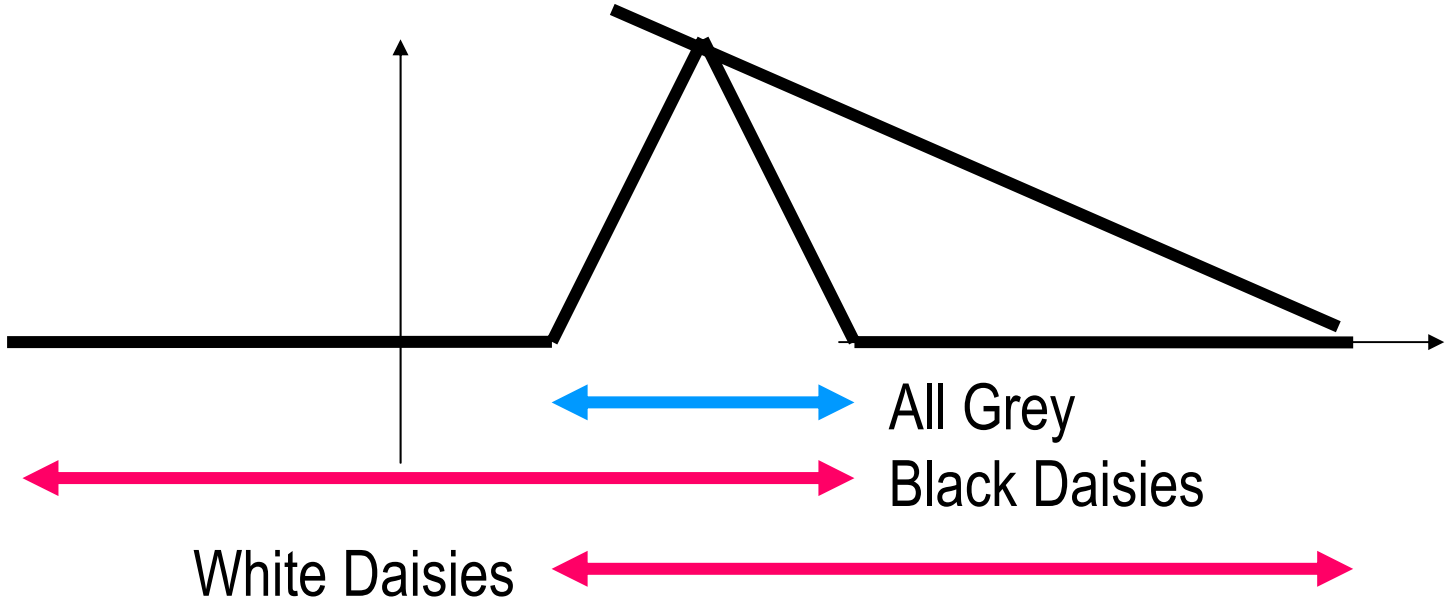

Similarly, on a White Daisybed, the more White against the Grey background, the more **negative feedback**.

This gives a line with a negative slope, but similarly extends (now to the right) the range of viability of a White Daisybed.

So both Positive and Negative Feedback works

There is no need to suppose God, or Evolution (in the real world), or some Trickery (in the Daisyworld model) has cunningly put in the **"right kind of feedback"** to make this homeostasis work.

Because **ANY** kind of feedback-response, positive **or** negative, combined with a Viability Hat-function, gives this type of homeostasis :- extends the range of viability beyond what it would be without any feedback.

Terminology: "Positive and Negative Feedback"

Within each Daisybed, temperature **T** affects Daisy quantity **D** via a Hat-function. In turn, there is an effect **feeding back** from **D** to **T** that is either +ve (Black) or –ve (White daisies)

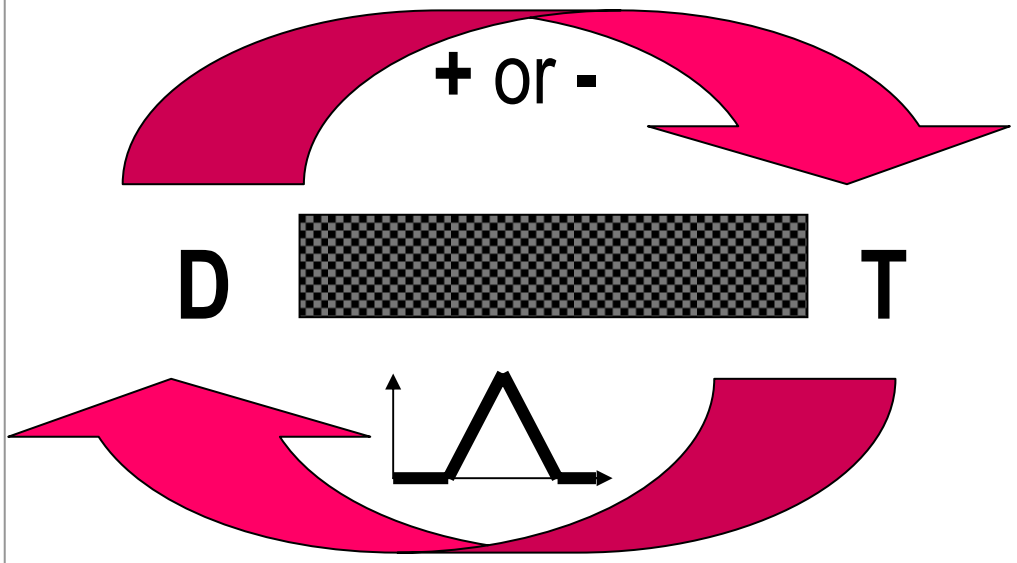

But this **doesn't** mean that this circuit as a whole is a (+ve or –ve) feedback control circuit – because the Hat-function is a crucial part !

Daisyworld ≠ Negative Feedback Circuit

Conventionally you need a “Negative feedback control circuit for homeostasis – using a **Set Point** (eg “desired temp”) and **Negative Feedback** to compensate for any Error …

… and Positive feedback leads to instability

This Daisyworld homeostasis is very different – for a start, there is no Set Point, only a viability range !

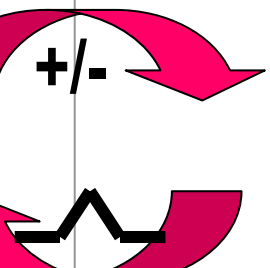

And both “Hat plus Positive feedback” and “Hat plus Negative feedback” work, to give regulation.

Talk Plan

1. 4-page Quick Summary, defining all the words in the Title.
2. Original Daisyworld.
3. New Kindergarten Daisyworld.
4. Rein Control, general principles that can … ....
5. Extend to phototaxis in a Dalek-like simulated robot.
6. Conclude.

Rein Control

As the Sun (or other external perturbing factor) threatens to push the Temperature (or other critical variable) too high or too low, this mechanism (Hat+feedback) automatically resists – homeostasis.

But note :- one mechanism counters the threat of being **too hot (White Daisies)**, a different one the threat of being **too cold (Black Daisies)**

Two “reins” of **Rein Control** (Clynes 1969) – each can **pull but not push**, you need both for regulation in both directions

How do Black and White Daisies interact?

So far, we have just been looking at an isolated Black Daisybed **or** White Daisybed.

What happens if we have both together, with some transfer of heat or **"Leakage" = L** between them ?

And in particular, what happens as we vary **L** from zero, no leakage, through intermediate values to maximum – where **B** and **W** daisybeds will have the same temperature **?**

Two Daisybeds

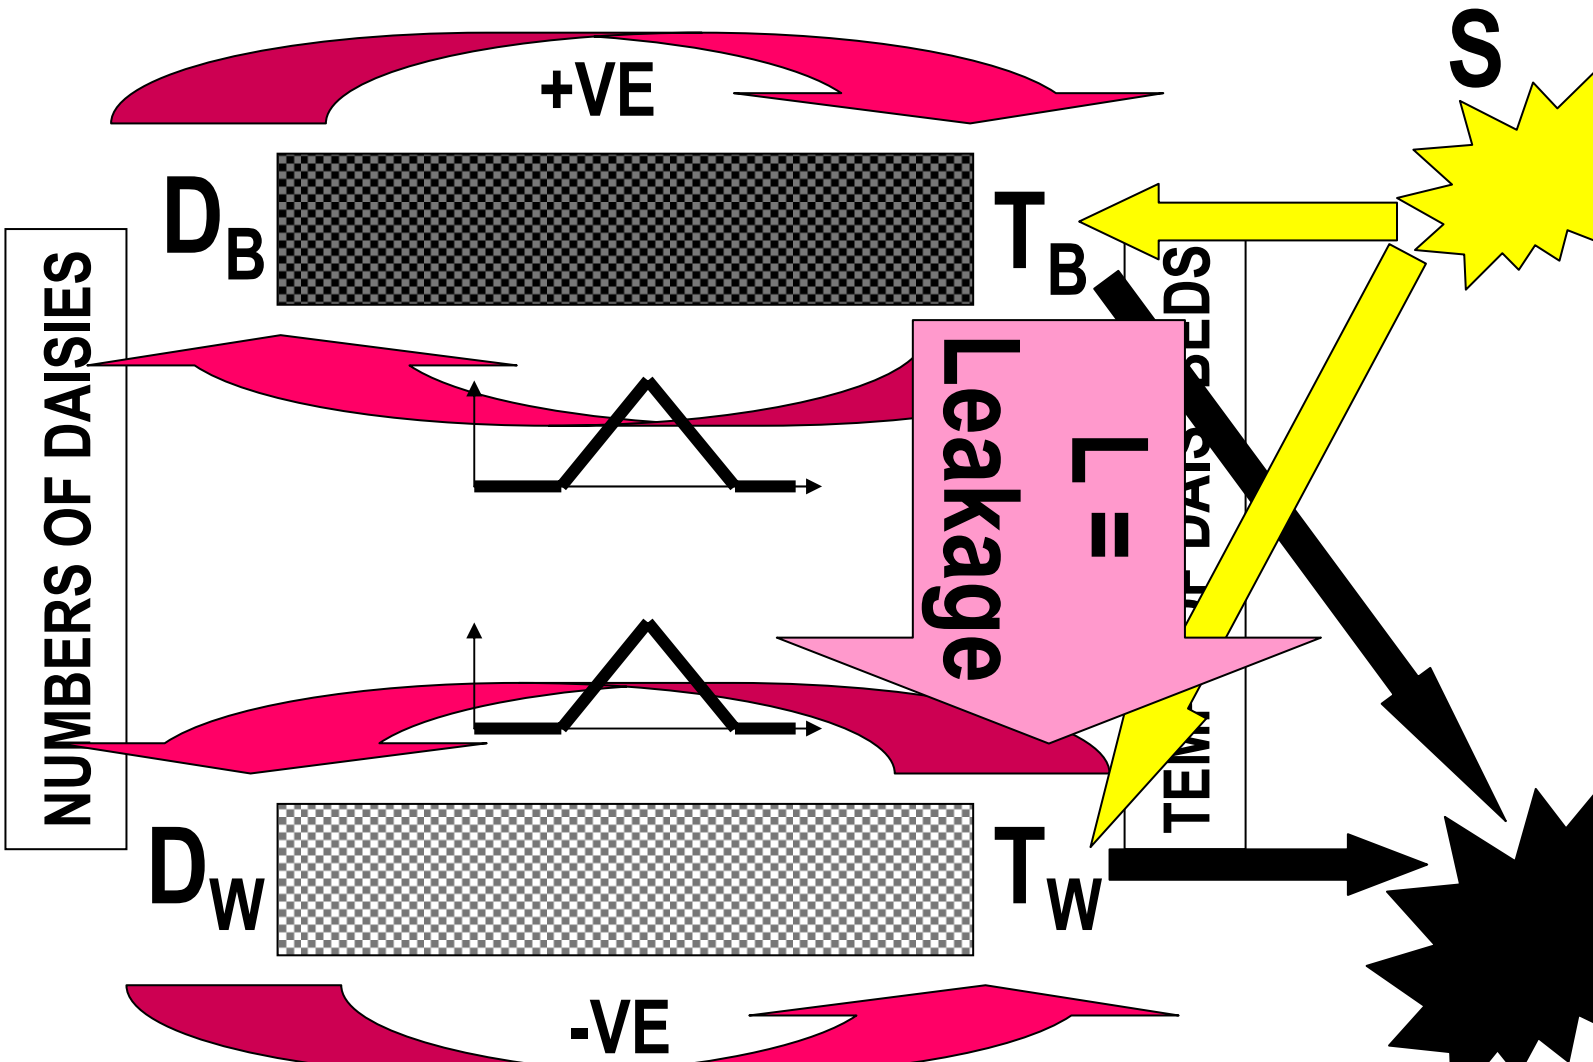

What Happens as we vary Leakage?

Suppose we can adjust the Leakage between Zero and Max?

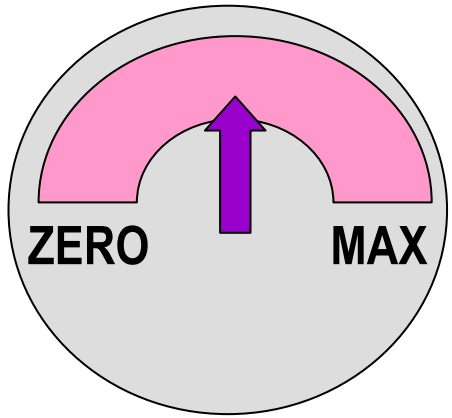

It will turn out that it is Intermediate values that give the interesting results – **loose coupling** between Daisybeds

But let's look at the extreme values of Leakage first

Suppose Maximum Leakage

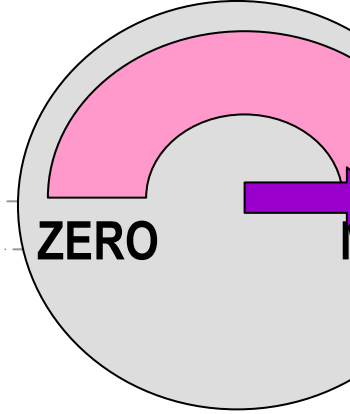

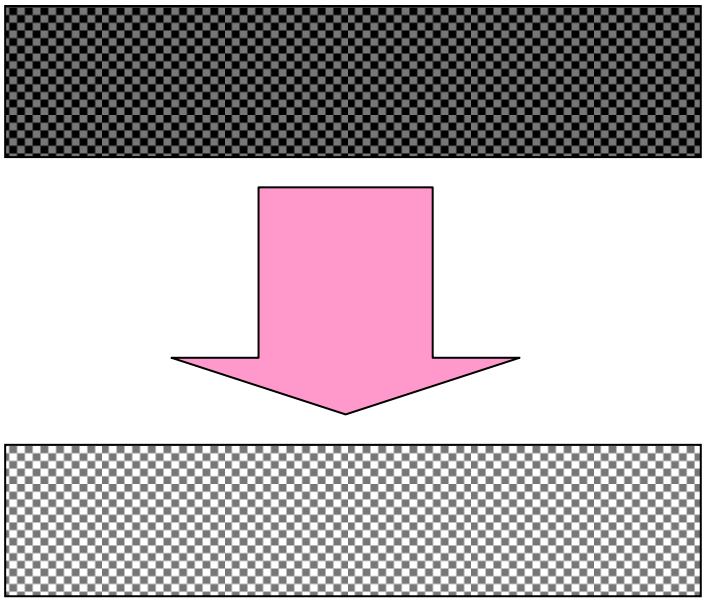

Then both Daisybeds equalise at the same temperature, hence equal numbers of **B** and **W** daisies … …

… but **B + W = GREY**

Suppose Maximum Leakage

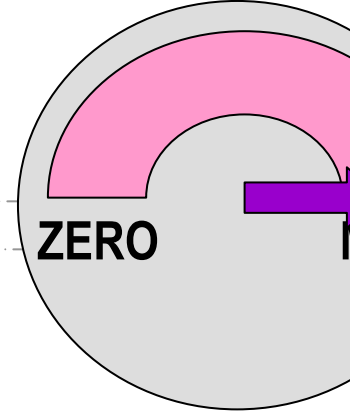

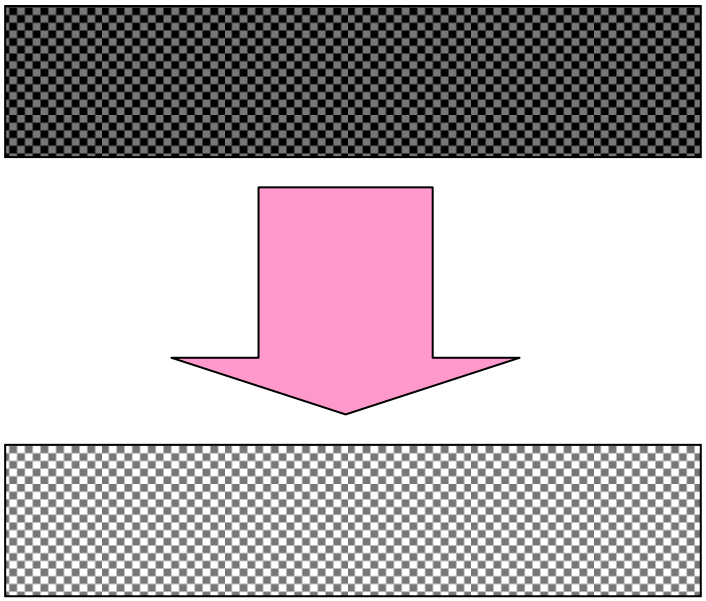

Suppose Maximum Leakage

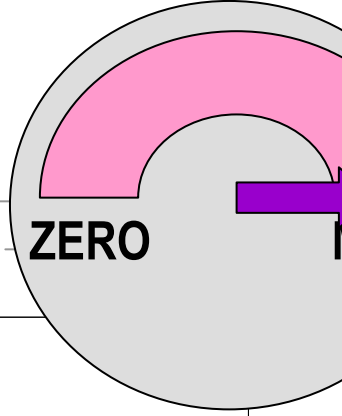

In this model, 90-110 is the viable range, with 100 as the optimum

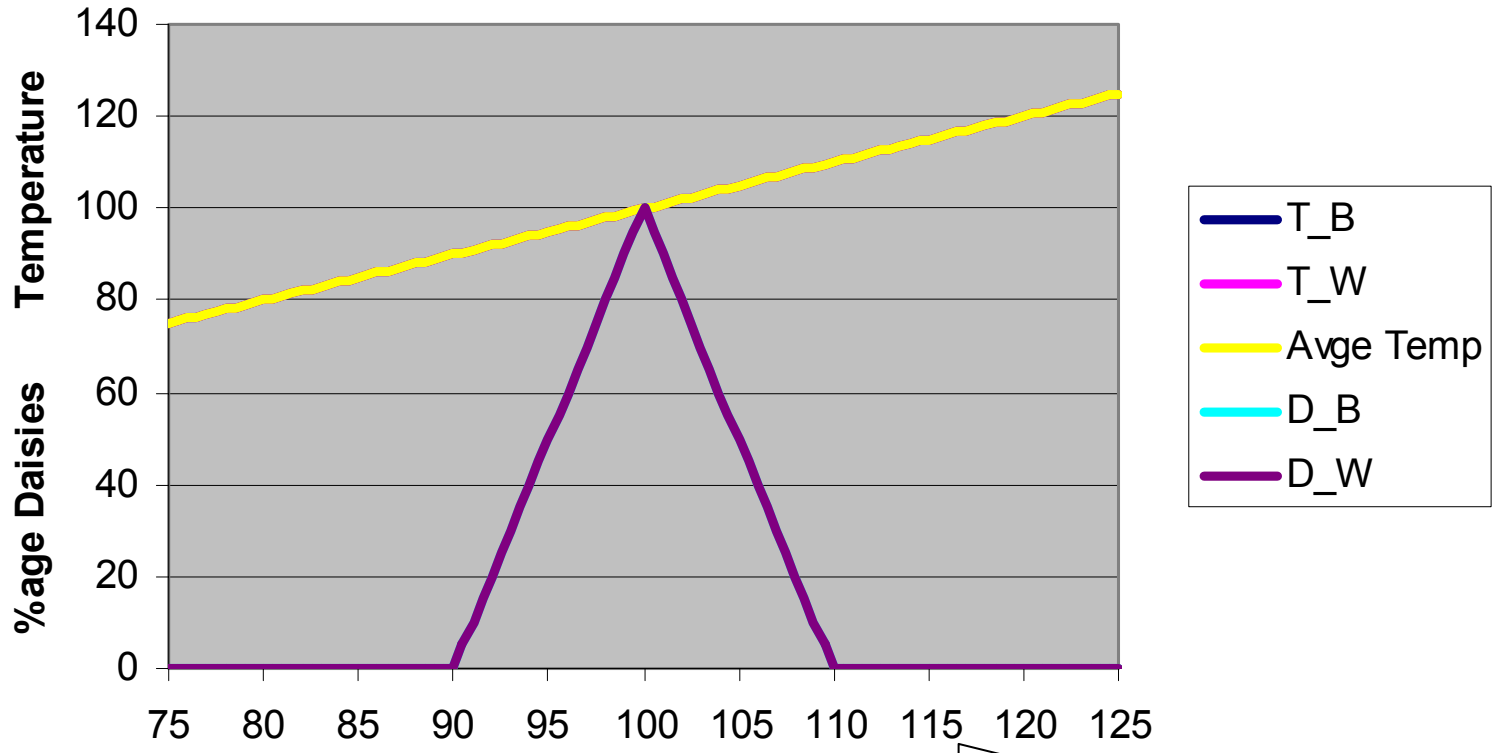

SOLAR LUMINOSITY INCREASES

Suppose ZERO Leakage

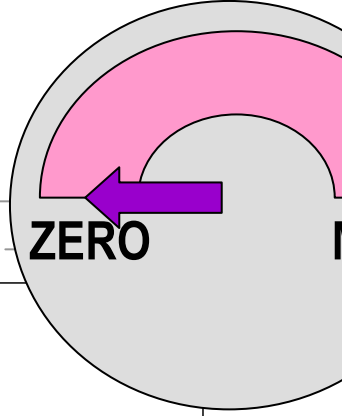

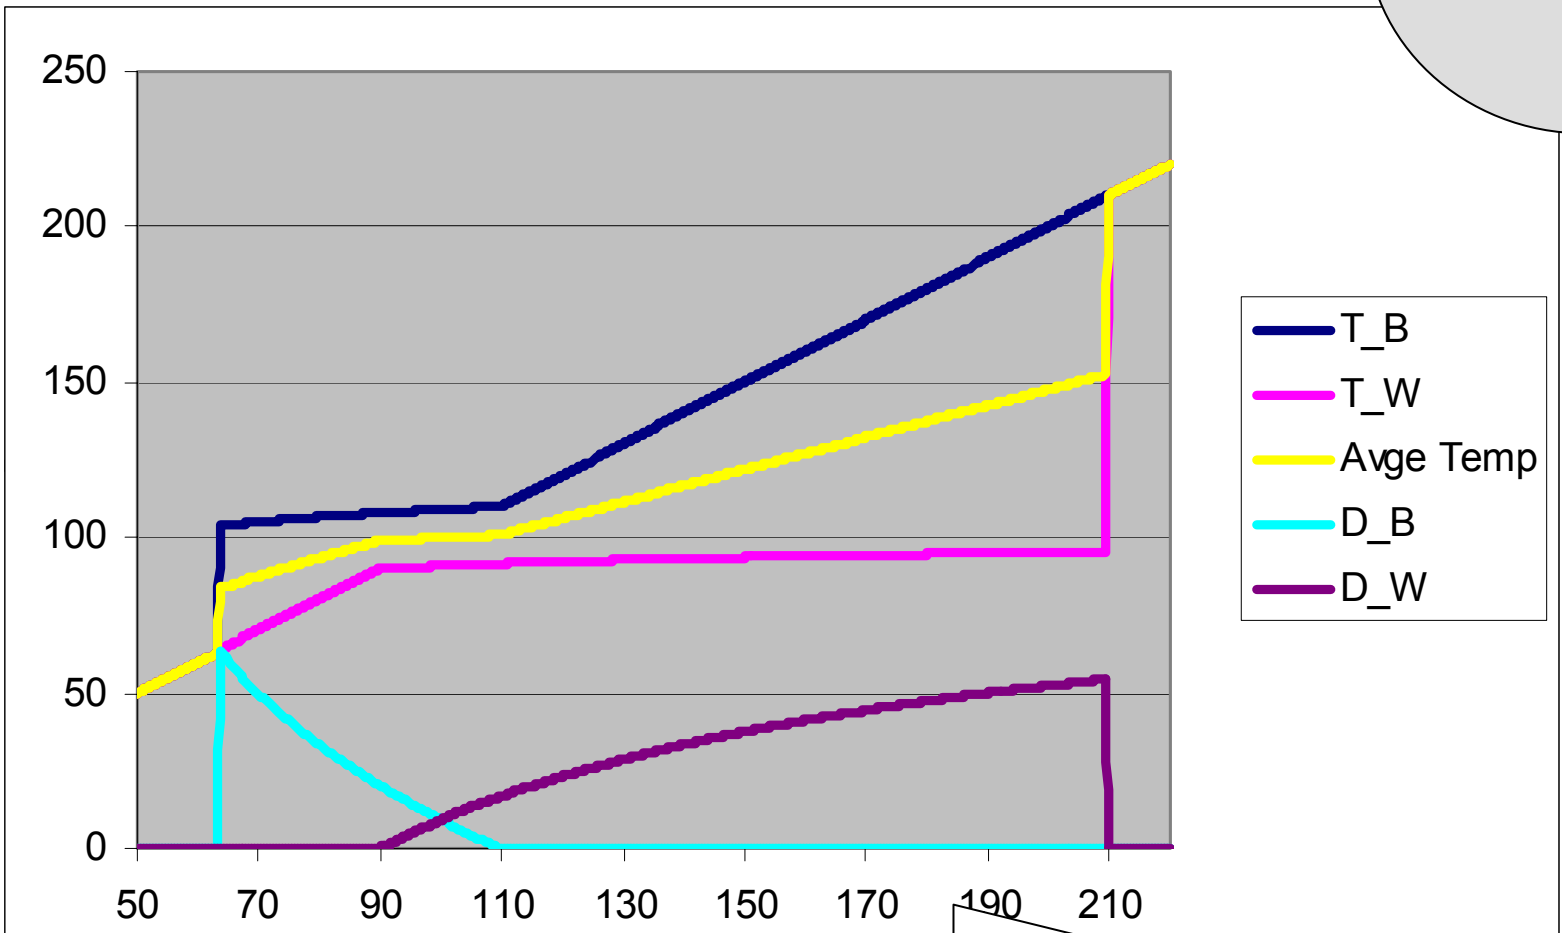

SOLAR LUMINOSITY INCREASES

No “crossover effects”

Suppose Intermediate Leakage

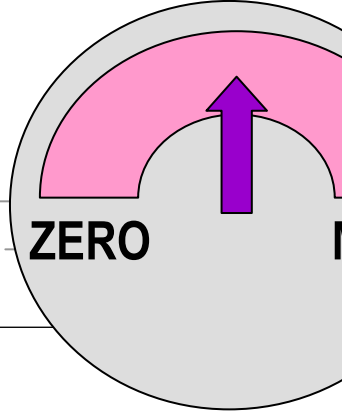

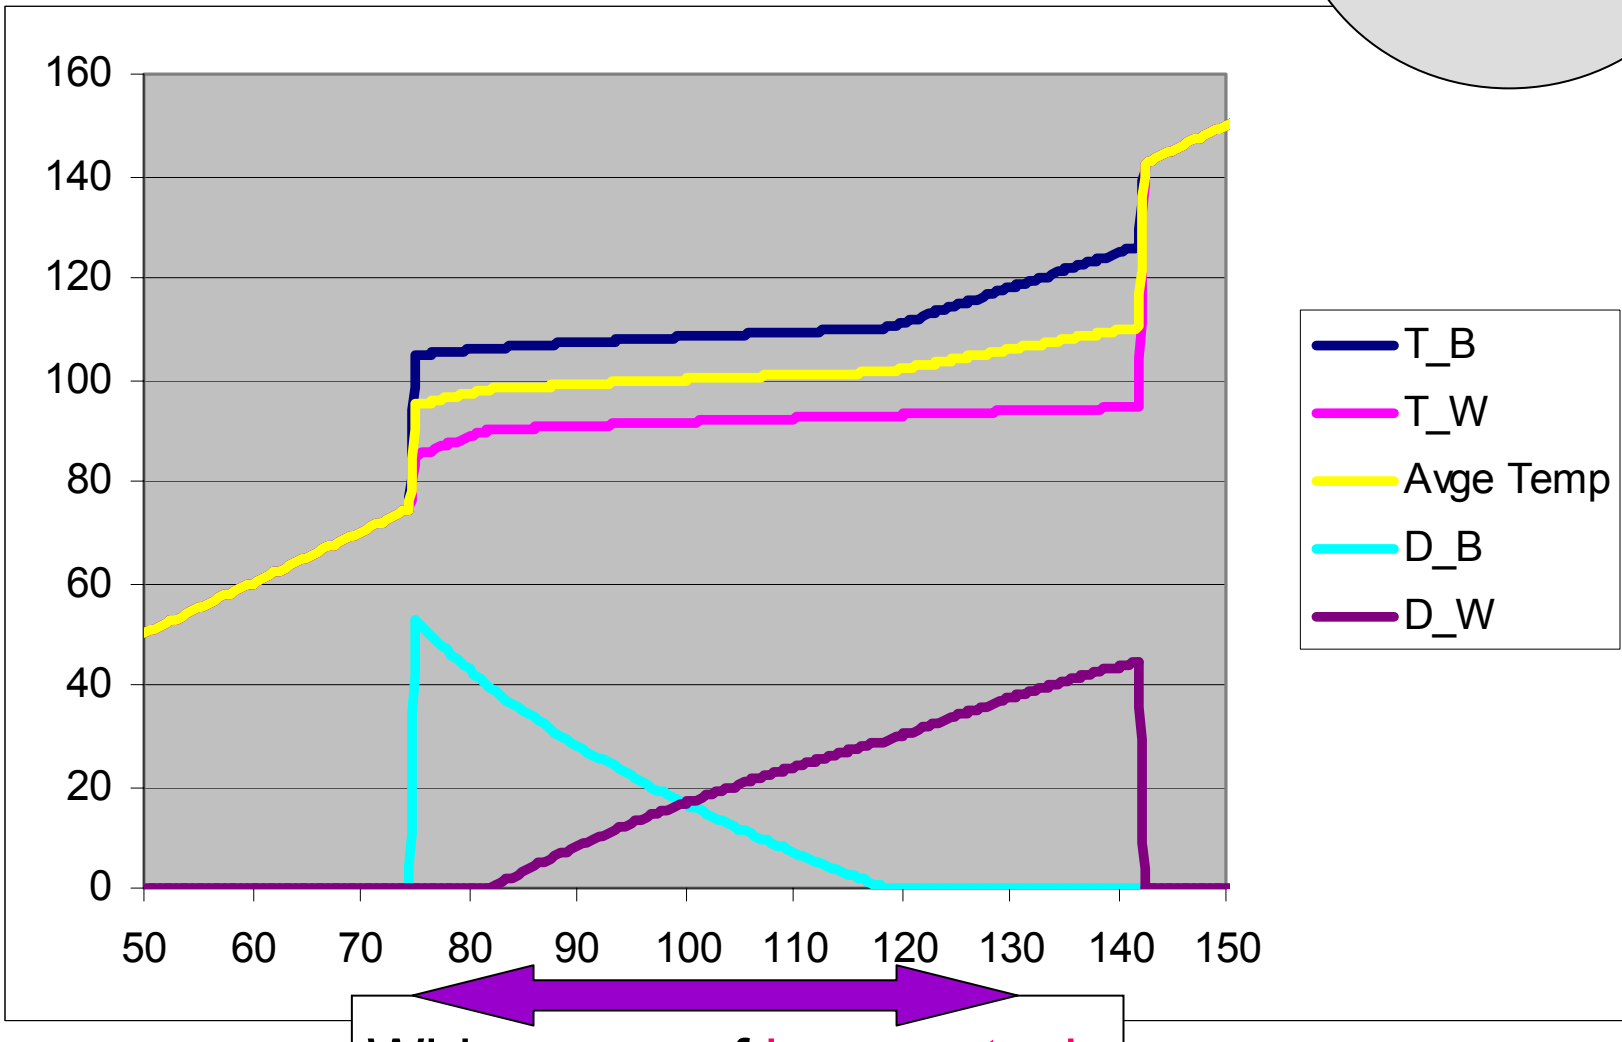

Rein Control and Loose Coupling

So the lessons are:-

1. Hat-function plus **any** feedback-response gives homeostasis, regulation against perturbation in **one** direction
2. To get regulation in **both** directions, you need feedback-repsonses in **both** directions – Rein Control
3. For the different regulations to interact for greater common benefit, you need Loose Coupling

Daisyworld Summary

This is a **parable**, where temperature stands for *any* critical parameter affecting viability, and the Sun for a perturbing external influence that threatens to take this parameter outside the viable range.

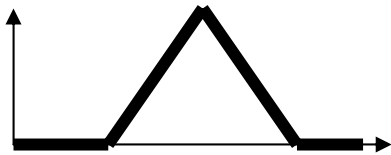

"Viable range" must imply some kind of Hat Function

Combining this with **any** kind of feedback-response leads to some degree of homeostasis, and the stronger the feedback-response the more the viability range is extended.

New Kindergarten Daisyworld

This new simplified Daisyworld, presented for the first time here, just looks at the overall shapes of Hat functions, and the signs of feedback-responses, ignoring any complexities of the underlying physics.

And it emphasises for the first time the significance of **Rein Control** in the Daisyworld model (cf Saunders' work), and the significance of **loose coupling**

Daisyworld and Rein Control Summary

To get regulation in **both** directions, you need both reins for Rein Control – and they need to be **loosely coupled**

Current work, not yet published, investigates how much coupling (here **'leakage'**) maximises range of homeostasis

This phenomenon, of individual interactions between Hat Functions and Feedback-responses of any direction (the stronger the better), loosely coupled with other such interactions, is simple and can be expected to be widespread.

Maximum Entropy Production Principle

1. Paltridge (1975) noted that **if** one hypothesised a *demon* who manipulated the climate to adjust the heat flow between the equator and the poles in such a fashion as to **Maximise the Rate of Entropy Production (MEP)**,..... **then** this constrained the heat flow equations so as to reproduce actual earth temperatures remarkably

2. Dewar (2003) gives a theoretical basis for why an open system, given enough degrees of freedom, can be expected to **MEP**, ……4th Law of Thermodynamics

MEP and the Daisyworld model

Dyke, using Harvey's Daisyworld model, showed (2004) that if one applied the same principles here, the same heat flows that maximised EP **also** had the property of maximising the range of (solar) perturbations under which the daisies remain viable.

MEP implies Maximum Homeostasis

Support Pujol 2002

Generalisation

Let's give just **one** example of how these principles can be generalised – here to a **very** different domain of

Active Perception

It's going to look very different, but *trust me*, the underlying principles are the same!

Talk Plan

1. 4-page Quick Summary, defining all the words in the Title.
2. Original Daisyworld.
3. New Kindergarten Daisyworld.
4. Rein Control, general principles that can … ....
5. Extend to phototaxis in a Dalek-like simulated robot.
6. Conclude.

An Animat – a Simulated Agent

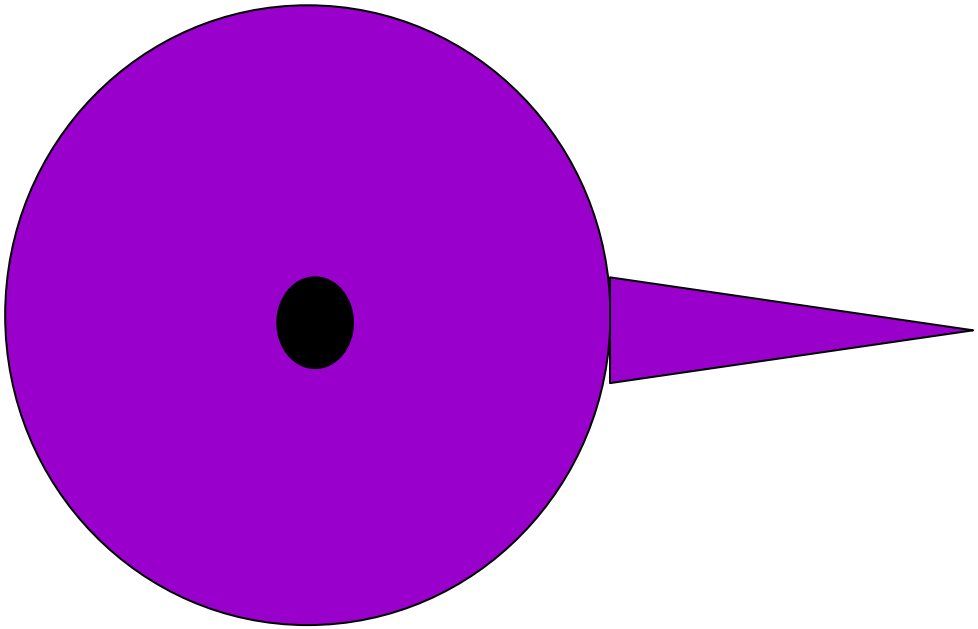

View from above – nose shows which way it is facing, all it can do is rotate about its centre.

An Animat – a Simulated Agent

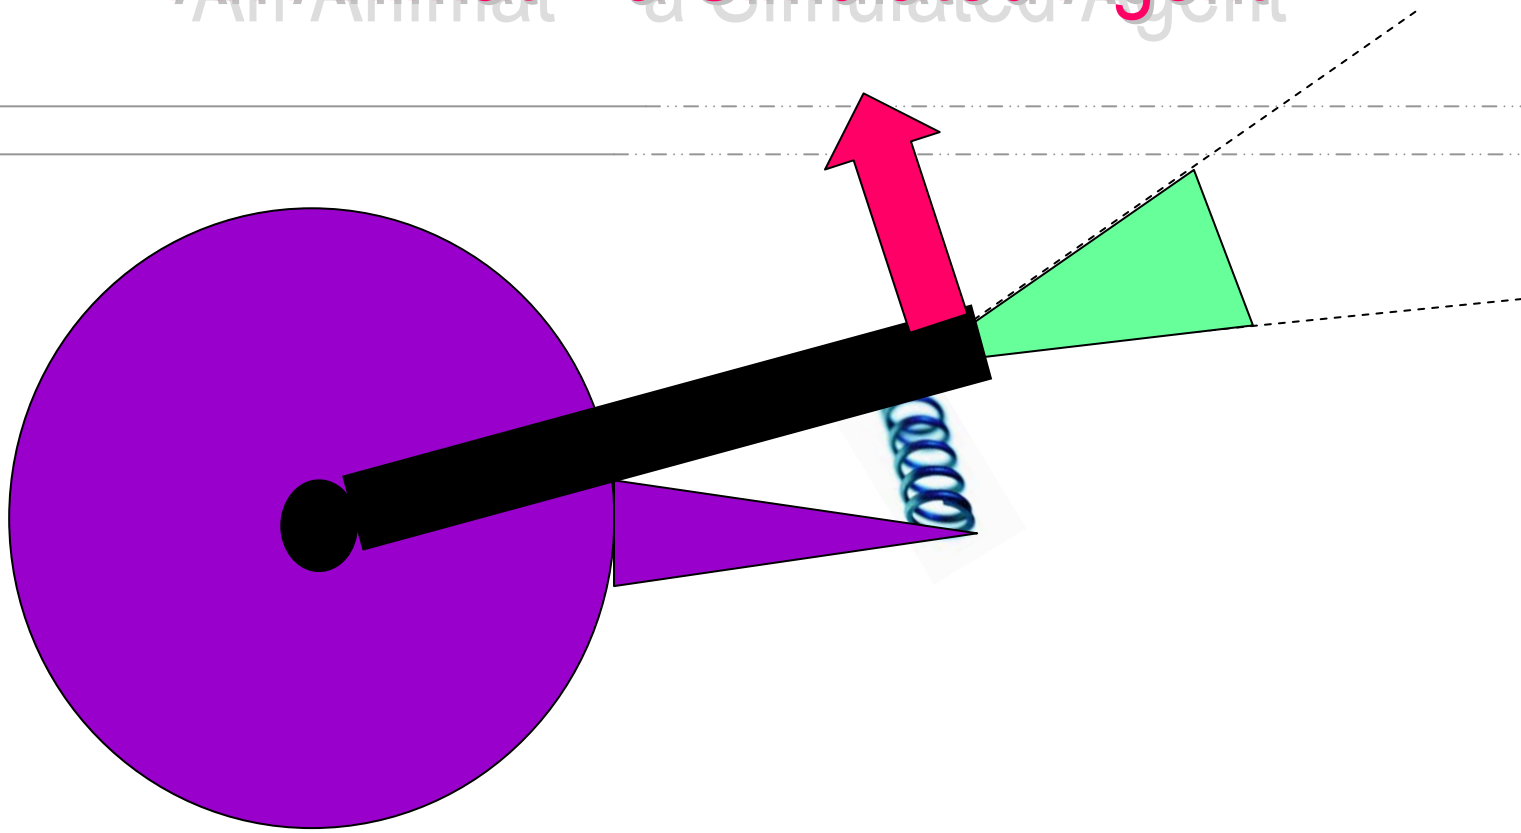

Add a **tentacle**, that can also rotate around the centre.

… with a **photosensor** with an ‘angle of acceptance’

… and a **jet** that converts light input into sideways force

… and a **spring** restraining the tentacle to the nose

Tentacle Responds Actively to Light

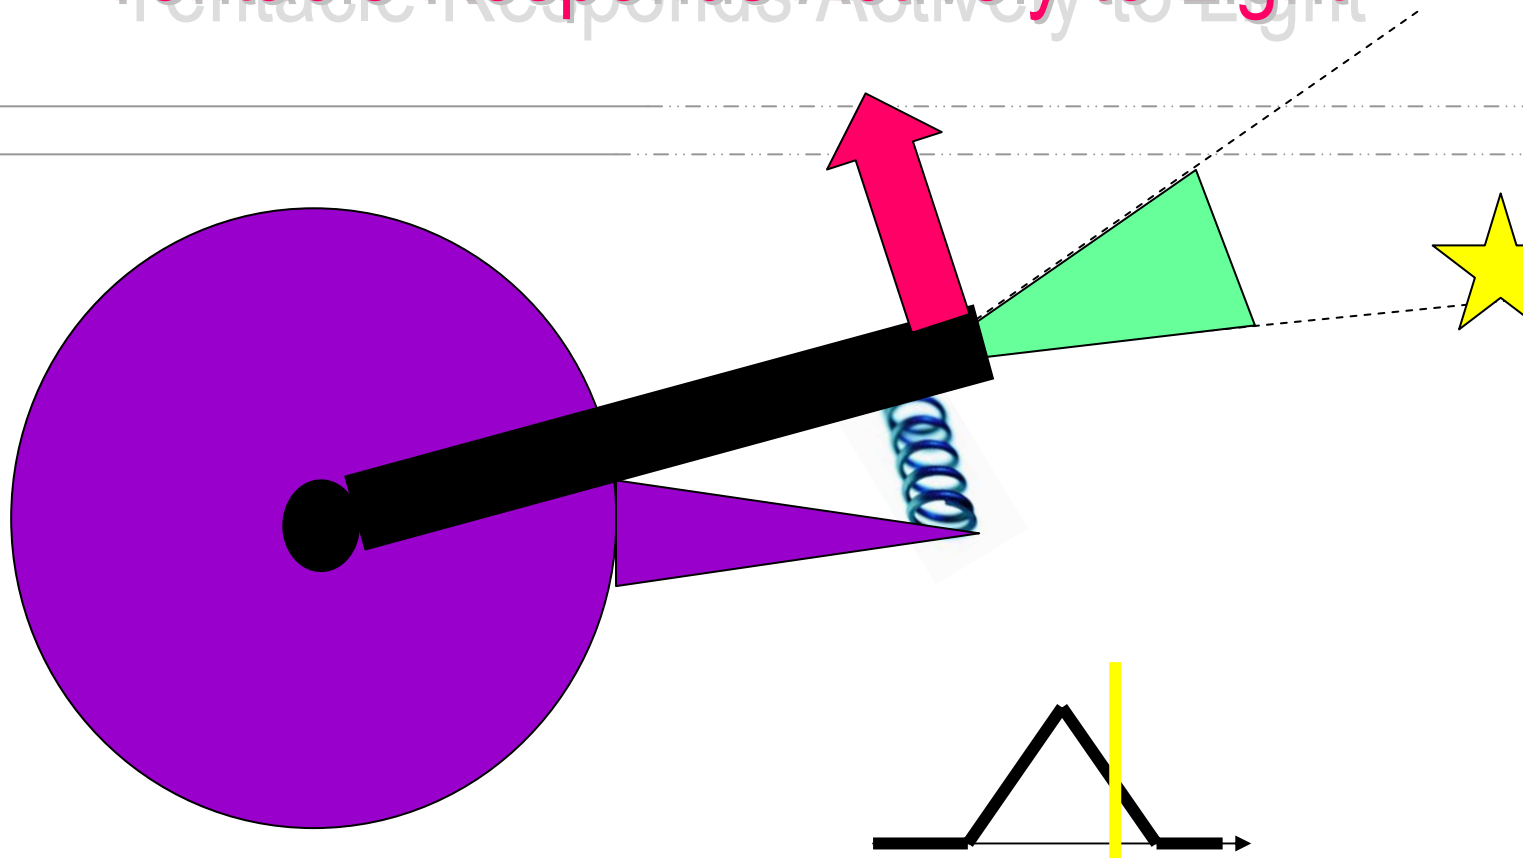

The **photosensor** responds with a **Hat Function** to light, maximum sensitivity when the tentacle points directly at a light

So a light **off-centre** means a medium **jet** force

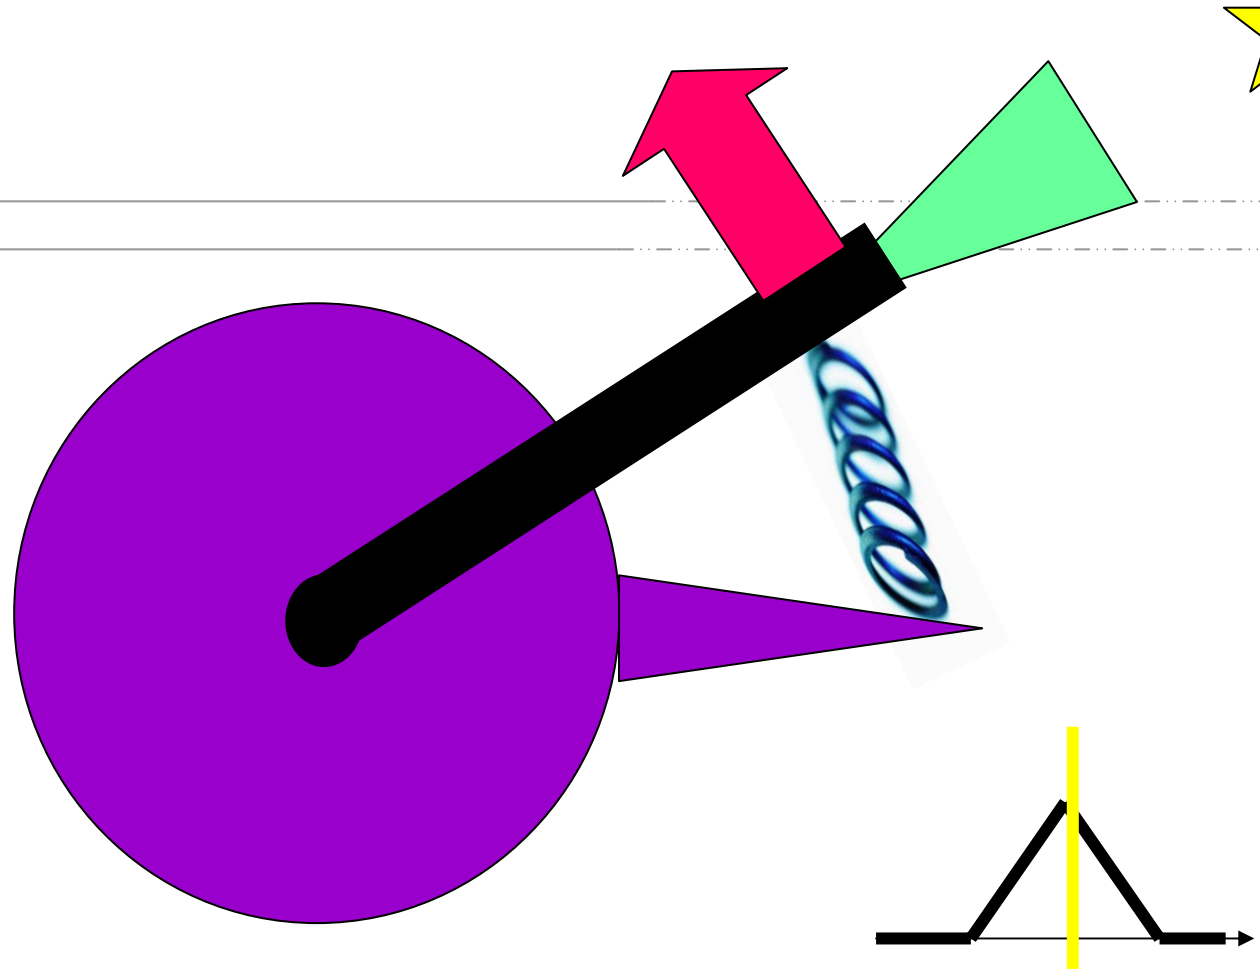

Light **central** in the **photoreceptor** produces maximum **jet** force, extending the **spring**

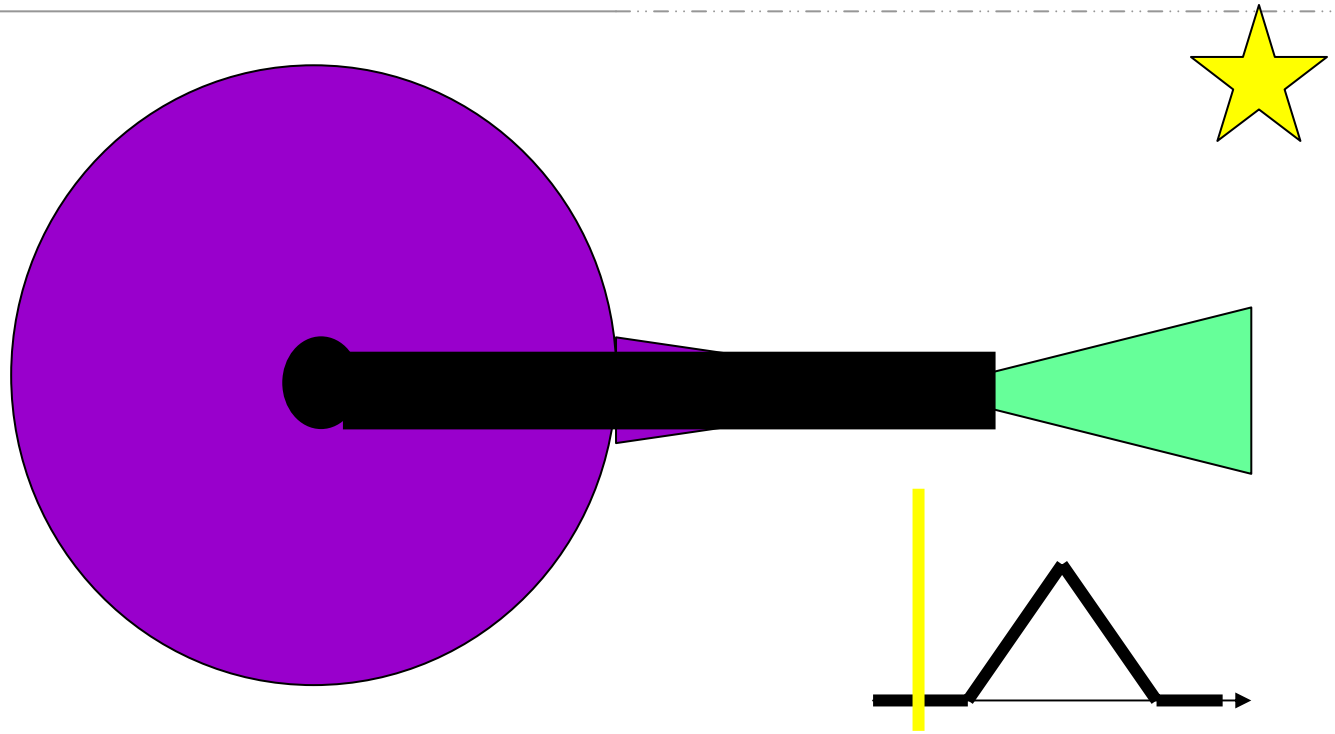

But if the **photoreceptor** can see **no** light, there is zero **jet** force, and the tentacle **springs** back over the nose

Now let's have **LOTS** of these tentacles

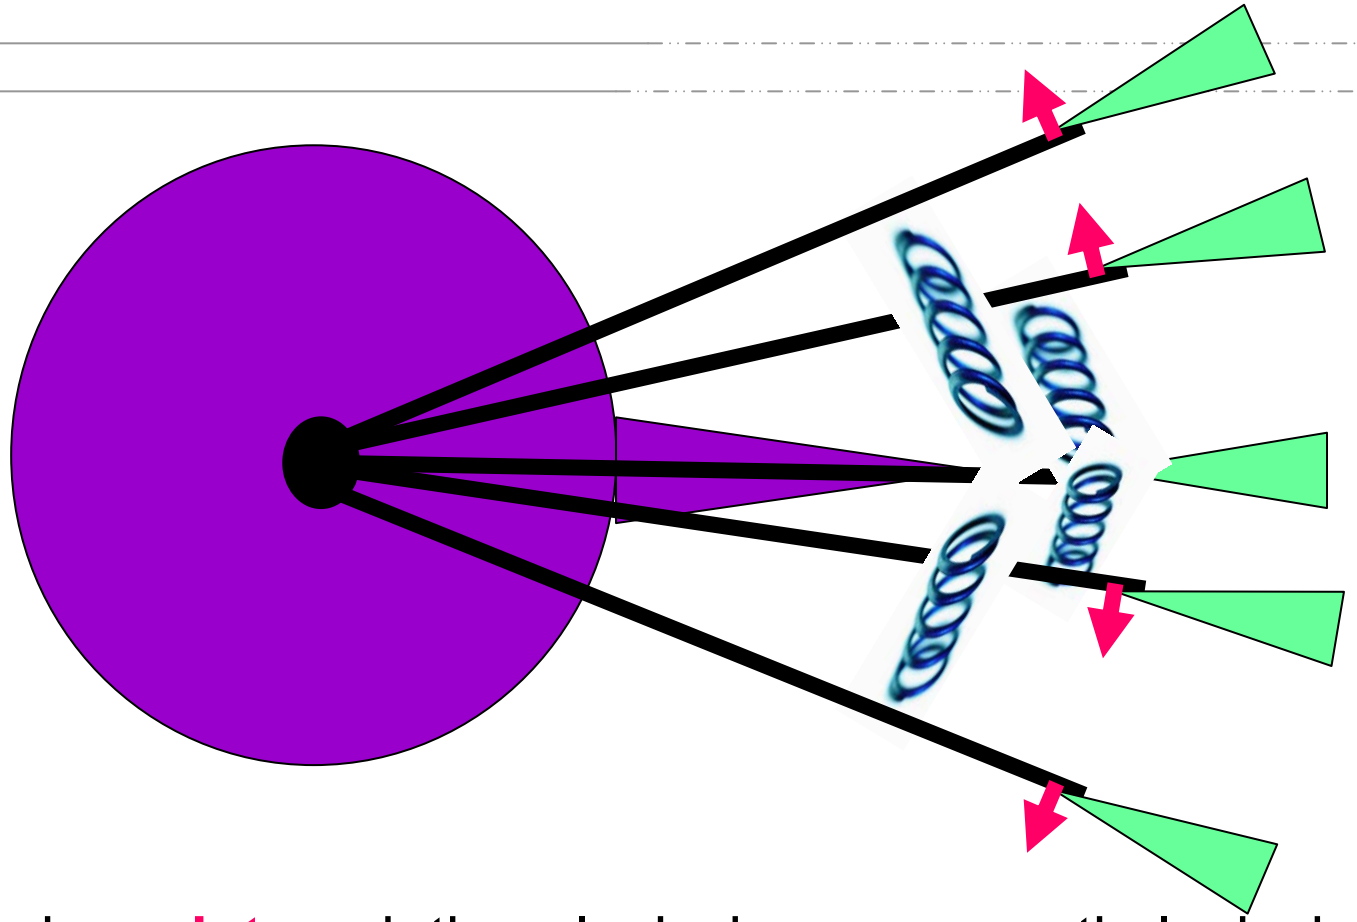

Some have **jets** pointing clockwise, some anti-clockwise, at random. All are connected by **springs** to the nose, but are otherwise independent of each other.

Parallels with Daisyworld?

The **angles** correspond to **temperatures**

The **photoreceptors**' range of **sensitivity** corresponds to Daisies' range of **viability** – Hat Functions

The **jets**, one direction or other, correspond to temperature feedbacks from **B** and **W** daisies, **+ve** or **–ve** responses

The **springs**, all coupled to the nose, correspond to the **loose coupling** between Daisybeds.

What Happens?

Let the balance of forces on the nose, from the randomly connected tentacles, rotate the robot around its centre (corresponds to the average planetary temperature)

Just as in Daisyworld, the effect is **as if** the Daisybeds were **'trying'** to stay within their zones of viability …

… so here, the effect is **as if** the tentacles are **'trying'** to stay within their zones of sensitivity, i.e. pointing near to the light.

So with a moving light, we get **PHOTOTAXIS**

Phototaxis

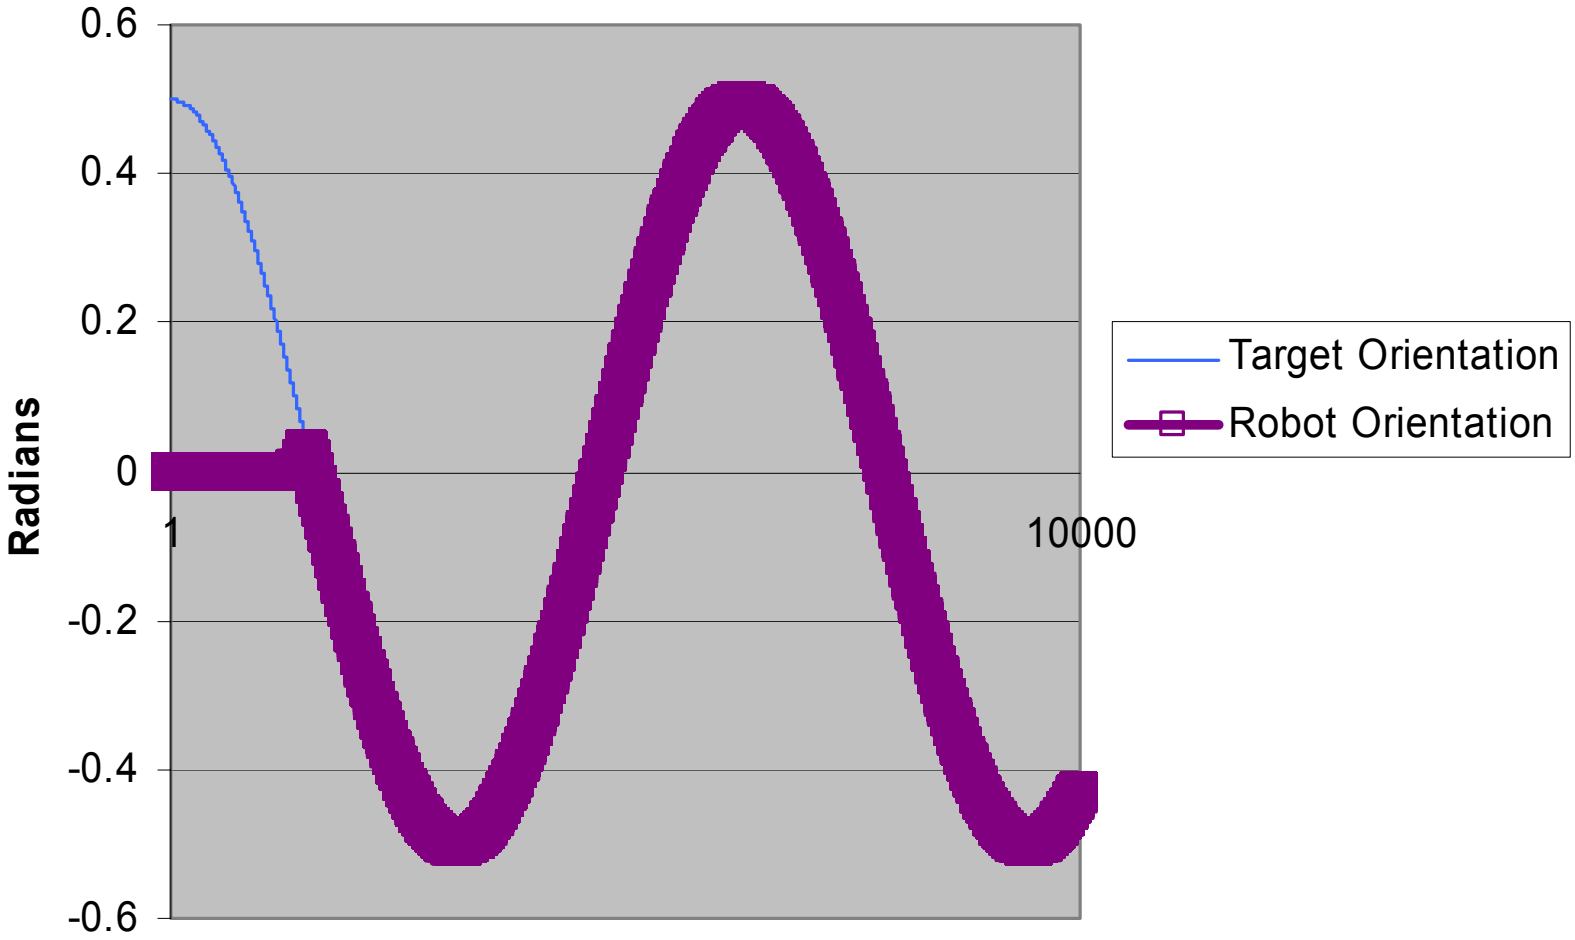

It picks up the light and tracks it ...

Successful Translation

So we have translated the simple mechanisms underlying homeostasis in Daisyworld

into Active Perception in an Animat – the underlying Maths is the same

Simple mechanisms, randomly wired up, loosely coupled

Talk Plan

1. 4-page Quick Summary, defining all the words in the Title.
2. Original Daisyworld.
3. New Kindergarten Daisyworld.
4. Rein Control, general principles that can … ....
5. Extend to phototaxis in a Dalek-like simulated robot.
6. Conclude.

Conclusion (1)

Homeostasis

Organisms have feedback control mechanisms for maintaining conditions vital to their comfort and survival.

Many would argue that such *homeostasis* is central to the very concept of life.

E.g. *Autopoesis* is *homeostasis* of ones identity as an organisation.

An understanding of basic mechanisms of *homeostasis* is crucial both for Biology and for Artificial Life.

Rein Control

Little-known principle in physiology, put forward by Manfred Clynes *(musician, neuroscientist, coiner of the term 'Cyborg', … …)*

"When a physiological variable is regulated against being both *too high* and *too low*, different mechanisms are used for each direction".

You need two reins to control a horse, one rein can only pull but not push.

Daisyworld

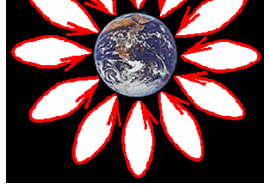

Gaia Hypothesis, Lovelock 1974 :- "the biosphere - atmosphere, oceans, climate, Earth's crust and biota, living organisms, is regulated as a homeostatic system in conditions comfortable for the living organisms"

Daisyworld model, Lovelock 1983 :- Simple Artificial Life model presenting a possible Gaian mechanism, for e.g. temperature regulation.

This paper :- a new simplification of the Daisyworld model, showing how Rein Control leads to homeostasis. Confirming Lovelock, opening up new generalisations.

Active Perception

One generalisation is the use of Rein Control and Homeostatic principles in a simple example of *Active Perception* in a light-seeking Animat (simulated robot)

In Daisyworld, feedback and Rein Control keeps critical variable such as temperature within a viability range … …
In the Animat it keeps active sensors focussed on a light - *phototaxis*

New principles – many opportunities for further research!

Two Parts

A talk in 2 disjoint parts:-

1. **(90%)** Homeostasis and the Dynamics of Daisyworld, based on Harvey 2004
2. **(10%)** The Dynamics of Representations – comments and open questions

The Dynamical Systems Approach to Representations

The Dalek example shows how an agent can home in on, and then track a target, without making any comparison between “direction target is” and “direction I am facing”

Likewise, van Gelder’s discussion of the Watt Governor shows how it regulates speed without making any comparison between “current speed of steam engine” and “desired speed of steam engine”

The regulation arises from the dynamics, “without Int Reps”

So is a DS approach Anti-Representationalist?

Absolutely not! We should reject this misleading term.

We should be clarifying unambiguous usage of the term **representation**, and in particular distinguishing between

1. **Internal representations** as posited **mechanisms** within cognitive systems
2. **Representation-using** as something we humans do all the time

(1) Mechanism **(2) Behaviour**

Representations as Mechanisms

Many explanations of mechanisms use a **homuncular metaphor:** “the wire carries the signal from the thermostat to the central heating boiler”

This is a legitimate metaphor where it is useful. It draws on the metaphor of: thermostat and boiler are ‘people’, current level in the wire is like a ‘letter’ or ‘telegram passing info’

‘Representations’ are the **‘Billiard Balls’** of GOFAI Cognitive Science

The Grounding of this Metaphor

But this metaphor draws on the common understanding of how we use representations, letters, words, signs in the real world (the **external!?!** World) every day.

We are so familiar with this that we see no mystery in it.

Representations as **explanans** rather than **explanandum**

‘**Representations**’ are the ‘**Billiard Balls**’ of GOFAI Cognitive Science

Two Consequences

1. People who use ‘internal representations’ as explanations for mechanisms are **extremely** reluctant to define what they mean by the term
2. They often find DS explanations disturbing, because they rarely fit neatly into this homuncular metaphor.

The Representation Wars I

One strategy: List the dozens of different incompatible ways in which the term (internal) representation is used, and ask for clarity and unambiguity

Typical problem met: unwillingness to define – Billiard Ball problem

Working usages of "Representation"

1. Everyday usage re-presentation: a picture of a cat is a re-presentation of a real cat.
2. A stand-in: A Member of Parliament represents his/her constituents
3. The act of representing (as opposed to the image/picture etc)
4. A variable that correlates with another variable.
5. As 4, but also needing also some causal correlation.
6. As 5, but also implicitly using homunculi.
7. A representation needs a consumer.
8. A representation does not need a consumer.

… more …

9. Representations are in the brain.

10. Representations are in the mind.

11. Internal representations = mental representations

12. Representations are in the head, and I reserve the right to use head=brain or head=mind at will.

13. Reps are in the head when you imagine a cat, not when you see it

14. Reps are in the head both when you see *and* imagine a cat

15. To try and define representations is a mistake.

16. No need to define our usage of the term, because it is obvious.

17.

The Representation Wars II

One DS strategy has been to ask Cognitivists for examples of what they think are **representation-hungry** problems, and get characterisations in operational terms

Then design a ‘minimal cognition’ experiment that fits the bill, evolve a CTRNN to do it, and see whether this challenges the preconceptions.

Typical problem met:- moving goal posts

The Representation Wars III

Another strategy: just ignore those who deal in internal representations – there are enough sensible people around to talk to

Problem: newcomers to the field are likely to be misled by the orthodox camp

The Representation Wars IV

My current preferred strategy: to reclaim the term Representation, to reject the label anti-representationalist, to make it clear that **understanding** how we humans use representations is one core goal of Cognitive Science that a DS approach aims at tackling.

However a DS explanation goes, I expect it to be of the form that **representing** (in words, in images on paper…) is **something we do**, rather than **representations** are **something we have.**

To conclude

I used to be genuinely puzzled at being called an anti-representationalist; now I am very much annoyed by the term.

The use of representations is **something to be explained.** Using *representations* as an *explanans* is OK with the homuncular metaphor -- but the one place where the homuncular metaphor is **completely illegitimate** is in the project to understand what it **is** to be an agent, an organism, a human – or a homunculus!

The End