Asymmetries in ecological and sensorimotor laws: towards a theory of subjective experience

James J. Clark

Centre for Intelligent Machines
McGill University

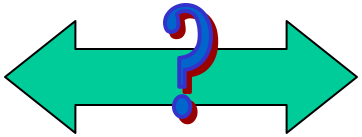

This talk will motivate an ecological approach to perception through a consideration of the philosophical problem of inverted spectra.

We will show that, at the very least, a restricted form of inverted spectra is not possible due to environmental asymmetries in the physical processes underlying image formation.

Our conclusion is that perception involves the ecological influences of surfaces on each other as well as the sensorimotor contingencies that exist between sensory data and observer self-motion.

Spectral Inversion - A Classic Philosophical Problem

(Locke, J. (1690/1987). An essay concerning human understanding.Oxford: Clarendon Press.)

When I see a RED object, do you see it as GREEN, but just call it RED??

Suppose your color detecting neurons were wired up backwards relative to my neurons (e.g. swap M and L photoreceptor connections).
Would this result in a spectral inversion?

Many philosophical arguments have ensued over this problem.

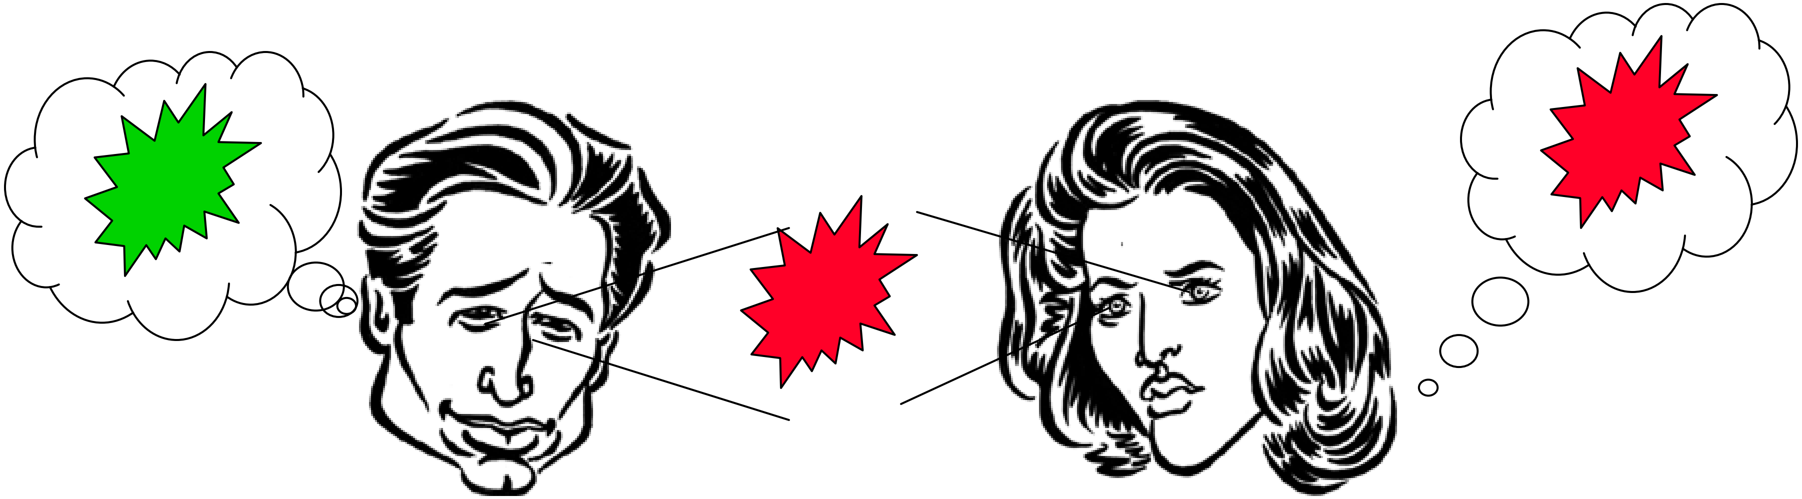

How could a spectral inversion occur in the brain?

One possibility: a Red-Green photoreceptor swap

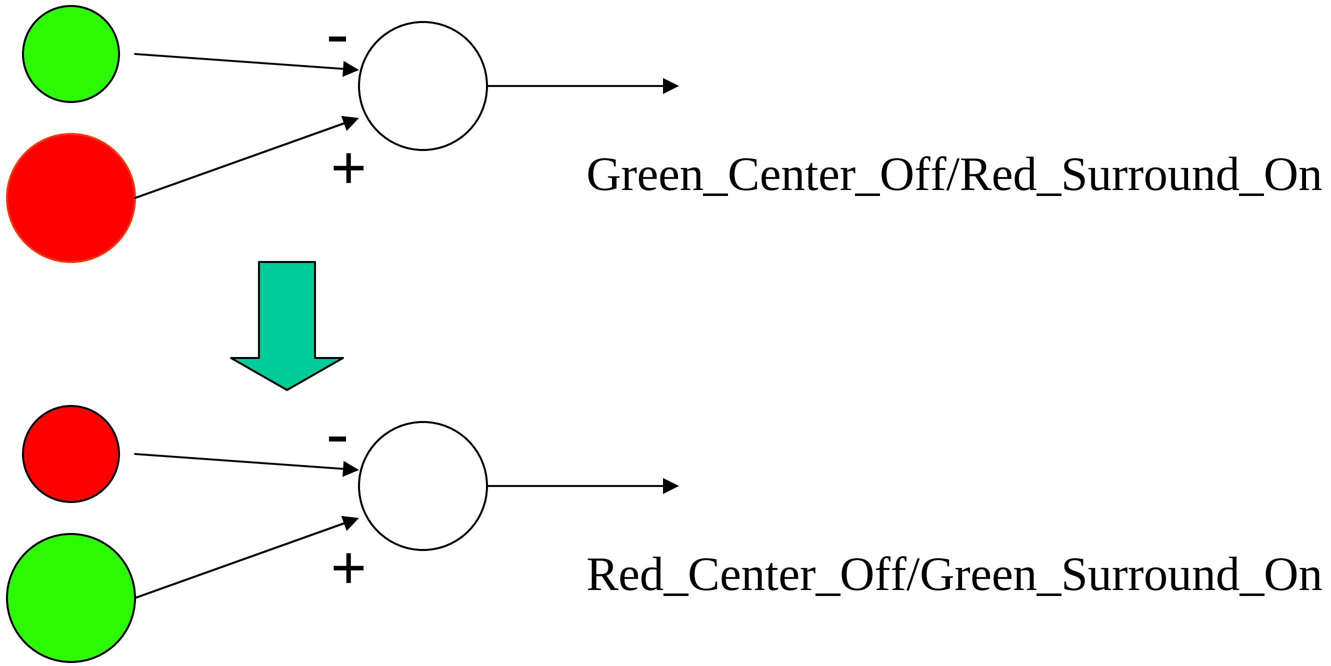

Such a swap could be caused by a rewiring of all of the L and M photoreceptor outputs, or by a mutation which causes the L receptors to have the same pigment as the (normal) M receptors (protanopia) *at the same time* as a mutation which causes the M receptors to have the same pigment as the (normal) L receptors (deuteranopia).

Piantanida gives an estimate of 14 in 10,000 males for the incidence of such mutations

Piantanida, T. P. 1974. A Replacement Model of X-linked Recessive Colour Vision Defects. *Annals of Human Genetics* 37: 393-404.

Normal

Red-Green Swap

Would a simultaneous protanope/deuteranope experience the Left image or the Right image?

We propose that, even if these types of neural re-wirings were to occur, people will still experience colour in the ***same*** *way!*

Our argument against spectral inversion will follow that of *Erik Myin* (whose arguments follow those of *Susan Hurley*…)

(Myin, Erik (2001), 'Color and the duplication assumption', *Synthèse*, **129,** No. 1. pp. 61-77)
(Hurley, Susan (1998), Consciousness in Action, Harvard University Press)

Let us consider the so-called "*Inverted-Earth*" arguments of Ned Block for inverted spectra (a variant of Hilary Putnam's "*Twin-Earth*"). This argument relates to a restricted version of the inverted spectrum, making the following assumptions:

- ***The set of phenomenal colour states is the same for each observer.***

- ***For each phenomenal colour state in this set there is a physically realizable world configuration that gives rise to that state in an observer.***

- ***Different observers may have different worlds that are needed to give rise to a given phenomenal colour state.***

In other words, each observer is assumed to have the same ***palette*** of colour experiences, and there always exists something in the world which will provide a given experience in the palette.

Possible spectral inversions therefore correspond to a re-ordering (e.g. swapping of red and green) of the elements in the common palette of colour states.

While there is something in the world which observer ***A*** will experience as red, and something in the world which observer ***B*** will experience as red, these two somethings conceivably might not be the same.

A possible spectral inversion scenario -

a person with swapped L and M photoreceptors

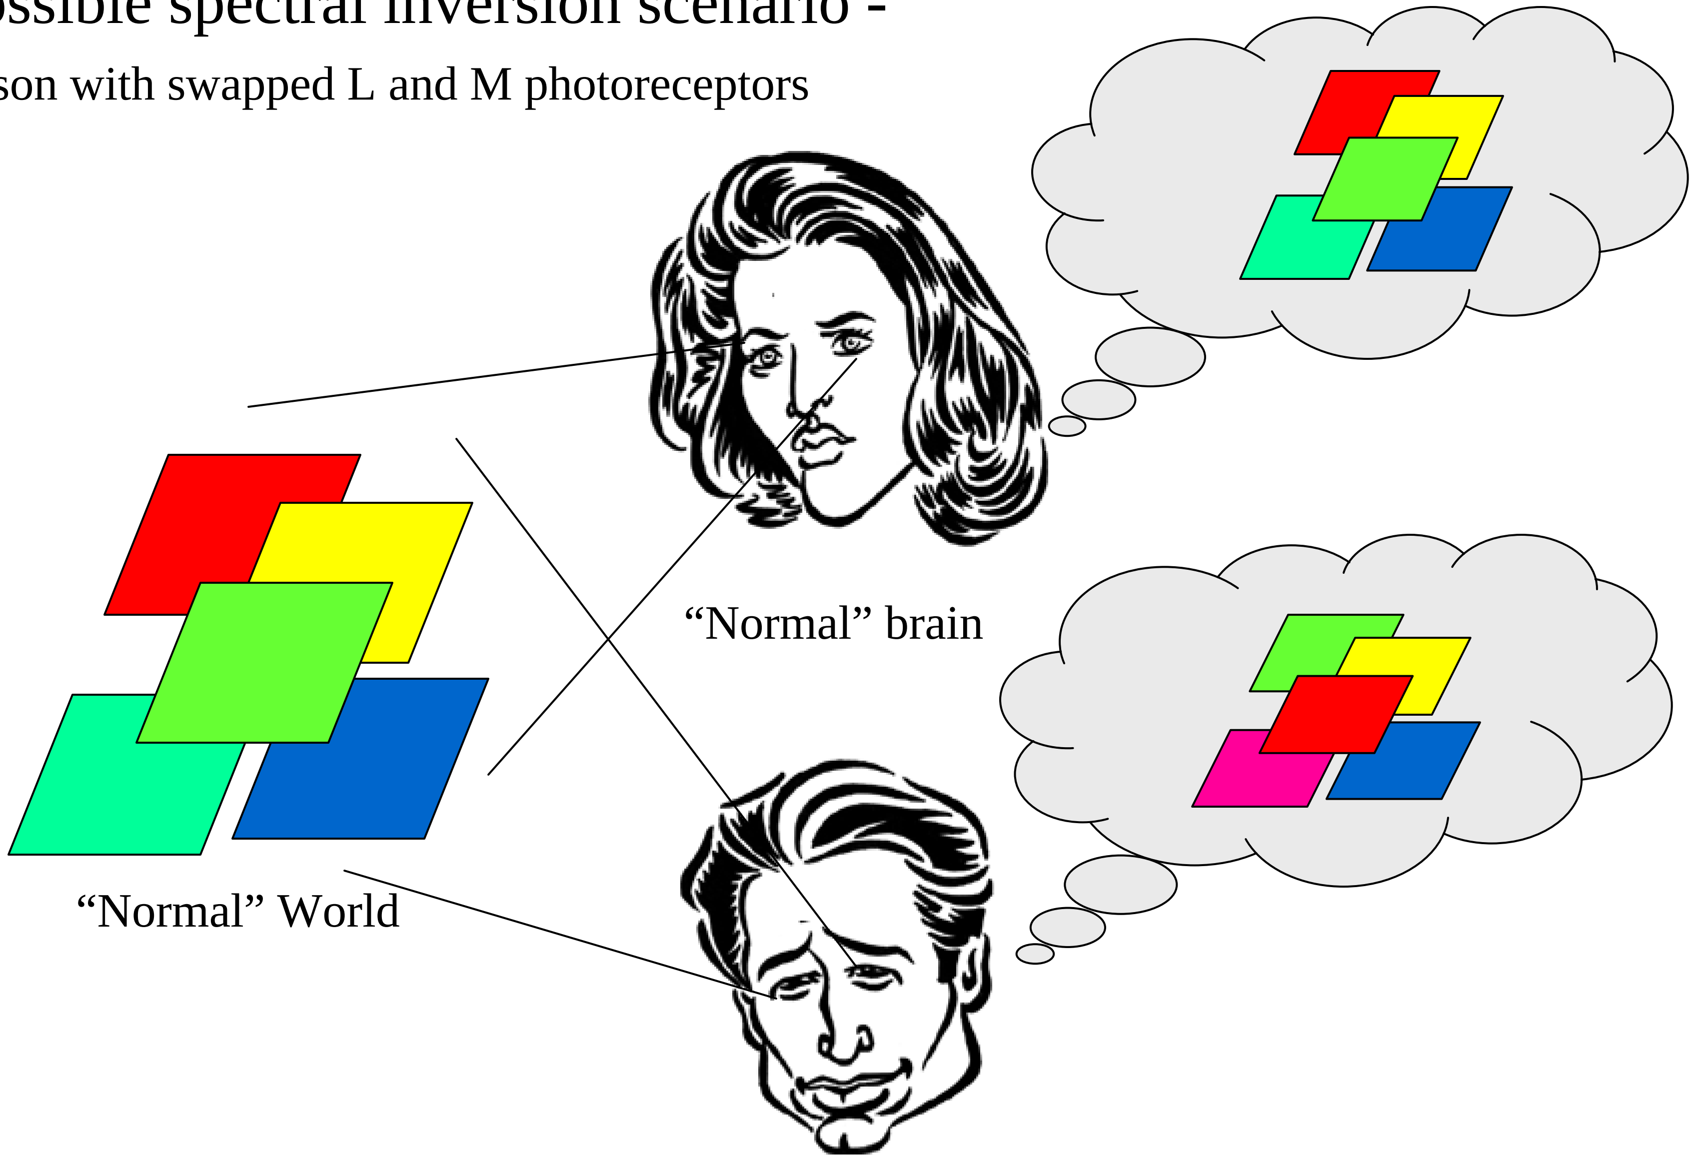

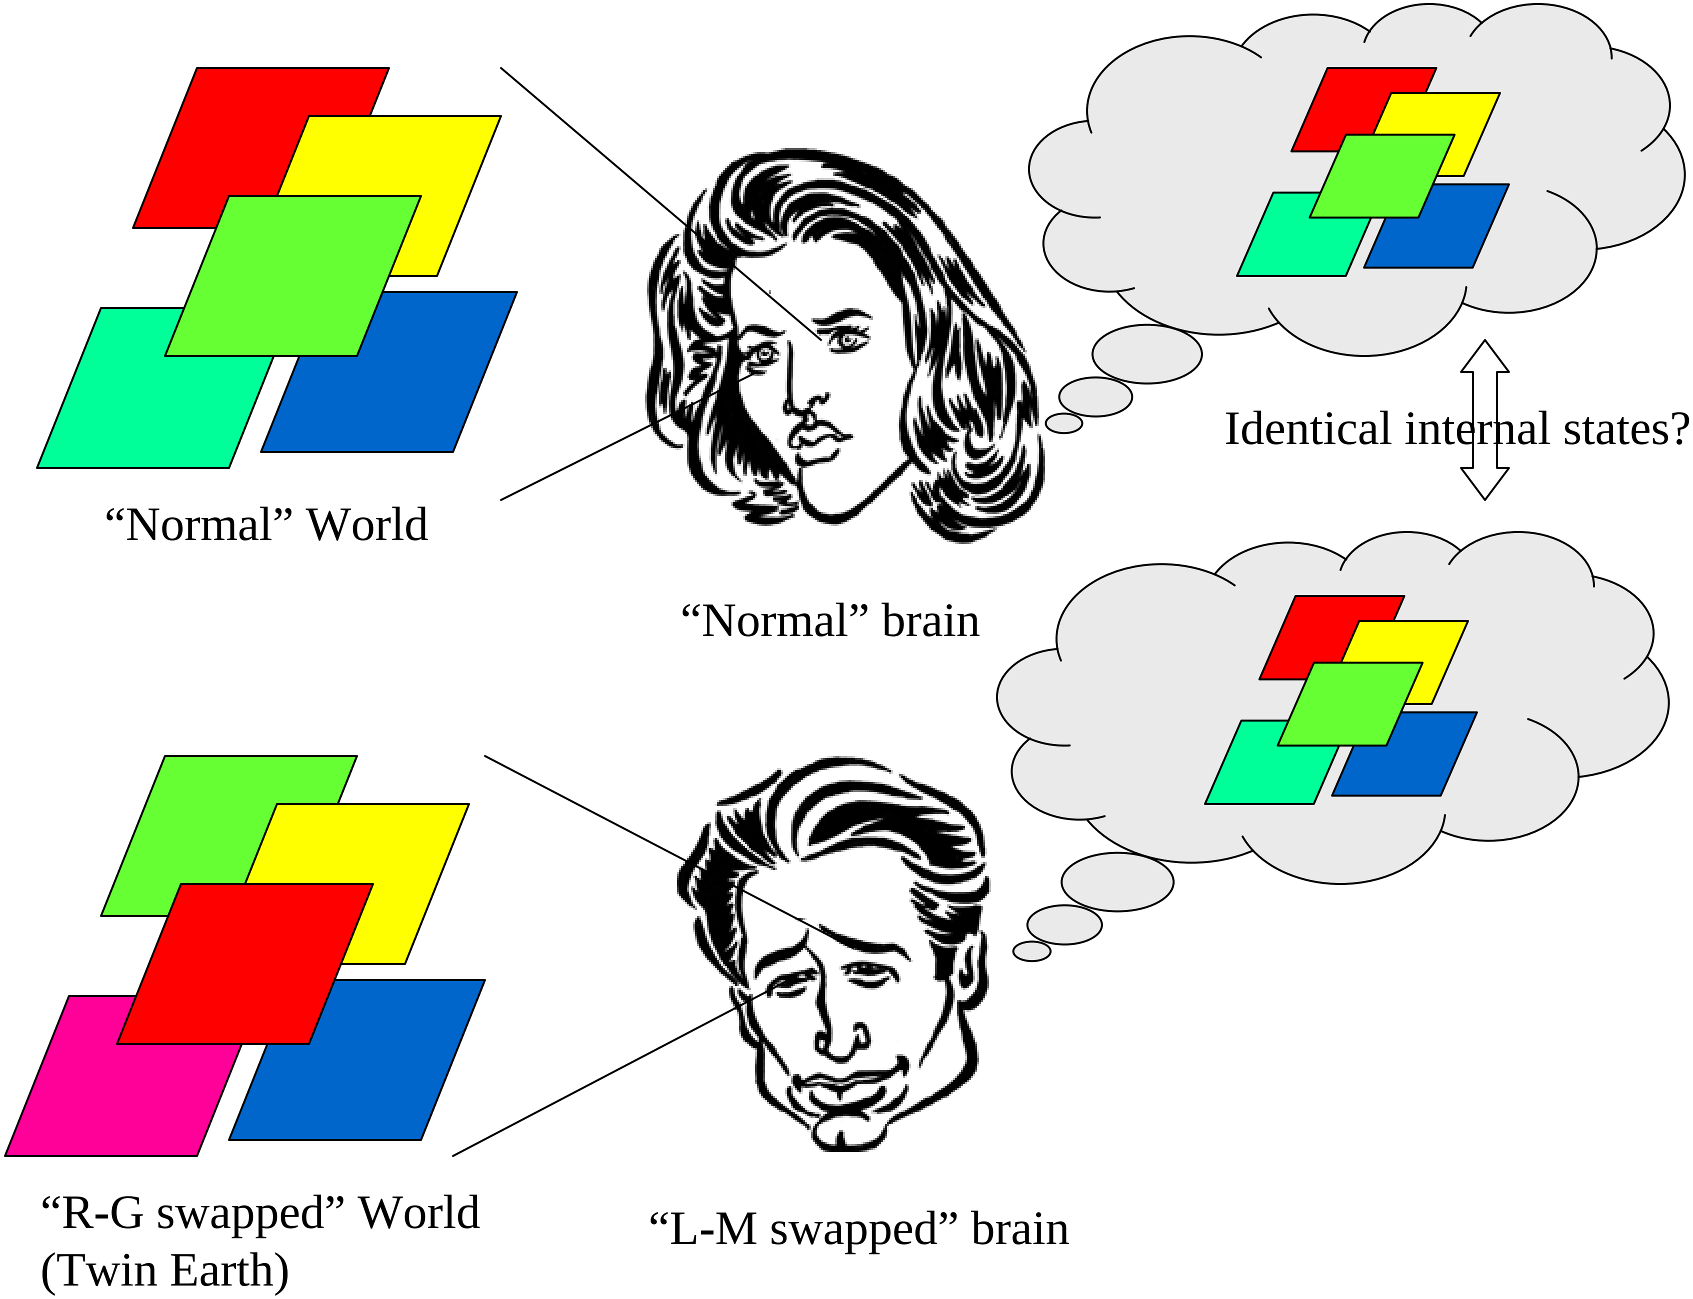

"Normal" World
"Normal" brain
Identical internal states?
"R-G swapped" World
(Twin Earth)
"L-M swapped" brain

Is such a “Twin-Earth” scenario possible?

To answer this we must investigate the assumptions inherent to the argument.

Susan Hurley considers the Twin-Earth arguments in terms of a ***duplication assumption****:*

*"Consider the **traditional view** that the **contents of the mind are in principle independent of the world**, so that we might be systematically deluded. If dualism is ruled out, it is natural to understand this view as holding, first, that **internal physical states fix mental contents** (Internalism) and, second, that it makes sense in principle to suppose that **internal physical states can be duplicated in systematically different environments** (the Duplication Assumption). This yields the possibility of systematic delusion:*

***mental contents can be duplicated in different environments**."*

*Let **A** be the operator which takes internal physical states (e.g. neural signals) into mental contents (subjective or phenomenal states) **I**. Assume that **A** is the same for all people. Let **R** be an operator which takes world states **W** (e.g. colour) into internal physical states. Assume that **R** can vary between people, and does not depend on any internal state.*

The duplication assumption then says that

*For any physically realizable world, W, given a specific systematic information preserving (or lossless) distortion of the internal states, **R->R'**, there exists some other physically realizable world, **W',** (the Twin Earth) for which*

$$\boldsymbol{A \,.\, R' \,.\, W' = A \,.\, R \,.\, W = I}$$

The duplication assumption is essentially a statement about the existence of *symmetry*:

Consider the set of all observer-world transformations, $\mathscr{T}$*, where*

$$\mathscr{T} : \{\mathscr{R}, \mathscr{W}\} \Rightarrow \{\mathscr{R}, \mathscr{W}\}$$

We define "Twin-Earth" transformations $\boldsymbol{T_t}$ *to be any element of* $\mathscr{T}$

$$\boldsymbol{T_t} \in \mathscr{T} : (\boldsymbol{R}, \boldsymbol{W}) \Rightarrow (\boldsymbol{R'}, \boldsymbol{W'})$$

for which, given any $\boldsymbol{W} \in \mathscr{W}$ *and* $\boldsymbol{R} \in \mathscr{R}$,

$$\boldsymbol{A} \,.\, \boldsymbol{R'} \,.\, \boldsymbol{W'} = \boldsymbol{A} \,.\, \boldsymbol{R} \,.\, \boldsymbol{W}$$

This definition of a Twin-Earth transformation implies that phenomenal states ***I*** *are symmetric with respect to "Twin-Earth" transformations.*

In other words, phenomenal states are not altered by a Twin-Earth transformation.

Do such transformations exist?

For simple, *impoverished*, worlds, one can probably find such symmetries. e.g. worlds that consist of a single planar surface illuminated by white, isotropic diffuse light, with fixed observers.

But such symmetries ***do not generally exist*** in the complex physical world that we live in.

There are a number of significant asymmetries. We can group these into two classes:

- *Sensorimotor*
- *Ecological*

Sensorimotor Contingencies:

(J.K.O'Regan & A. Noë, A sensorimotor account of vision and visual consciousness, in Behavioral and Brain Sciences, 2001, 24(5))

Changes in the sensory input as a result of observer self-actions (e.g. eye or head movements, or hand manipulations of objects)

Examples of Asymmetries induced by motor acts:

- spectral shifts resulting from eye movements
- shading variations resulting from manipulation of surfaces
- geometric variations resulting from egomotion (e.g. looming)
- spectral shifts due to chromatic aberration

In complex worlds the *information* contained in the sense data includes more than that of a particular phenomenal quantity.

For example, in a world where the observer is not fixed, in addition to the colour of a surface patch, we also have the position of the patch relative to the observer.

Let $\boldsymbol{W_1}$ be a given world state, and let $\boldsymbol{W_2}$ be the world state obtained by moving the observer to a new position.

For a Twin-Earth transformation to exist both of the following two equalities must hold:

$$\boldsymbol{A} \,.\, \boldsymbol{R'} \,.\, \boldsymbol{W_1'} = \boldsymbol{A} \,.\, \boldsymbol{R} \,.\, \boldsymbol{W_1}$$
$$\boldsymbol{A} \,.\, \boldsymbol{R'} \,.\, \boldsymbol{W_2'} = \boldsymbol{A} \,.\, \boldsymbol{R} \,.\, \boldsymbol{W_2}$$

While you can come up with an observer-world transformation which makes one of these two equalities true, we will show that, in the case of colour perception in humans, you cannot make both of them true, and hence there is no Twin-Earth transformation possible.

Example -

Consider the sensorimotor contingency associated with the variation across the retina of photoreceptor spectral sensitivity.

There is a yellowing of the incoming spectra in fovea as compared with periphery. This is due to absorption of short wavelength (blue) light in the lens and macula.

This yellowing affects blue more than green, and green more than red. Thus changes in the signal from a blue patch across eye movements will be large compared to those from green patches, which will in turn be larger than those from red patches.

A sensorimotor contingency or asymmetry with respect to colour...

The spectrum of light falling on the central retina is yellowed (reduced in amount of blue) as compared with that falling on the peripheral retina

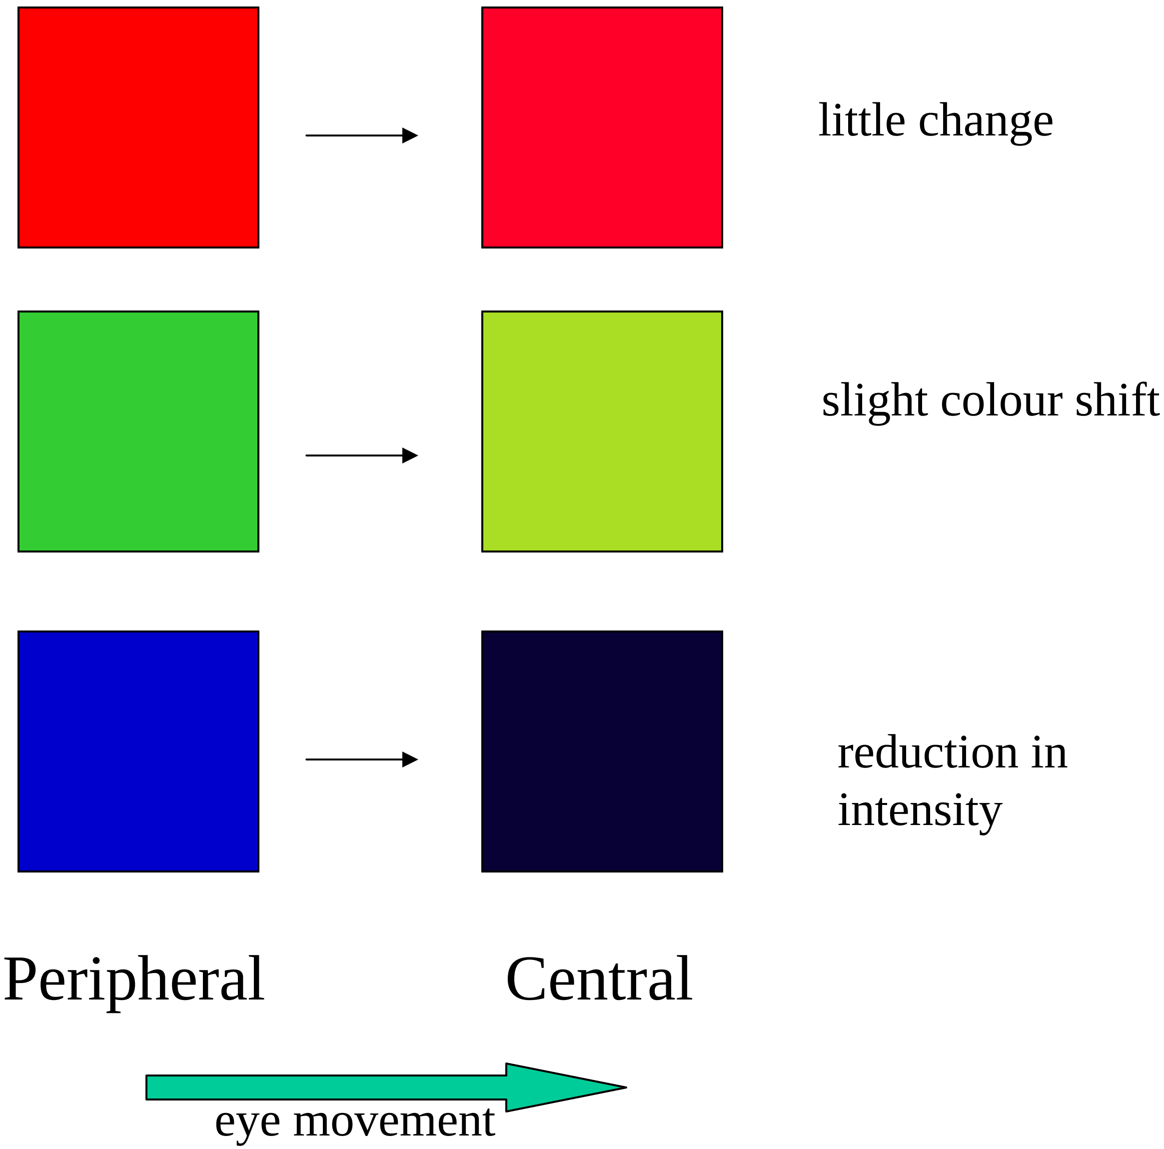

Viewing a White Patch across Eye Movements

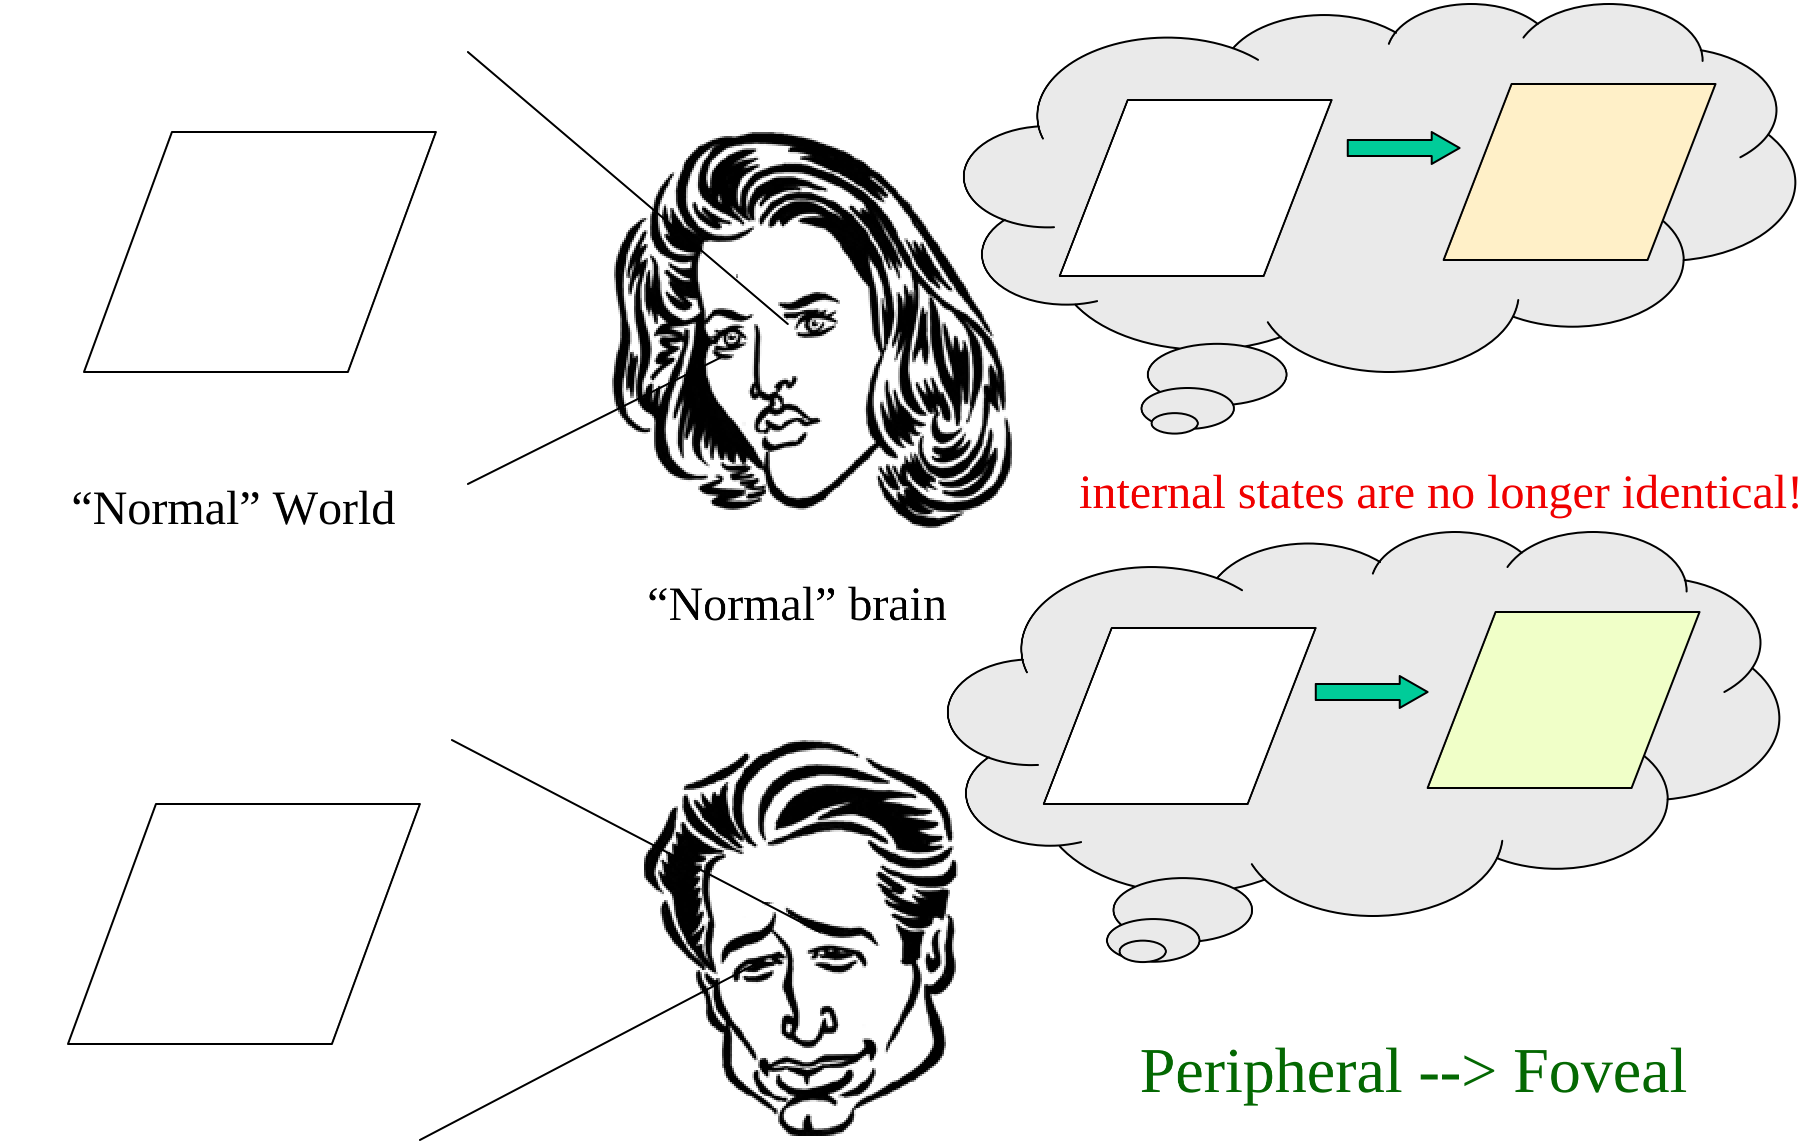

Even though we could arrange things so that the internal states of observer A and observer B are initially the same when viewing the white patch, their internal states are different after the eye movement.

One could think of applying some sort of *inverse filter* in the Twin-Earth to compensate, and reverse, the yellowing action (to turn it into a blueing action, for example). This would provide the same internal states for the white patch in both worlds, but would not give the same internal states for viewing a blue patch (since the blue neurons are unchanged, but the surface blue component has been changed).

Ecological Contingencies:

Changes in the sensory input as a result of environmental changes (e.g. changes in relative position of other scene structures)

Examples of Asymmetries induced by the environment:

- surface inter-reflection (indirect illumination)
- occlusions of one surface by another
- wavelength dependent absorption/scattering of light by the atmosphere
- asymmetry in visibility of objects between lit and shadow regions

Ecological asymmetry arises when considering the interaction of two surface patches. In such a world, in addition to the colours of the two surface patches, we also have the position of the patches relative to each other.

Let (W_1 , W_2) be a given world state corresponding to patch 1 and 2 both being present but widely separated so that they do not influence each other. Let $(W_1 \otimes W_2)$ denote the world state in which both patches are present and near to each other.

For a Twin-Earth transformation to exist both of the following two equalities must hold:

$$A \,.\, R' \,.\, (W_1 , W_2)' = A \,.\, R \,.\, (W_1 , W_2)$$
$$A \,.\, R' \,.\, (W_1 \otimes W_2)' = A \,.\, R \,.\, (W_1 \otimes W_2)$$

Once again, while you can come up with an observer-world transformation which makes one of these two equalities true, we will show that, in the case of colour perception in humans, you cannot make both of them true, and hence there is no Twin-Earth transformation possible.

Example -

Consider the ecological contingency associated with *inter-reflection* between surfaces.

Gilchrist and Ramachandran point out that the image of a red room in white light looks different than a white room in red light. This difference is due entirely to the asymmetry in surface inter-reflection in the two cases.
(AL Gilchrist and VS Ramachandran, "Red Rooms in White Light Look Different than White Rooms in Red Light", ARVO (1992))

A similar study was done by *Langer and Gilchrist* who showed that subjects could easily distinguish a green room illuminated with white light from a white (gray) room illuminated with green light. The key being that the green surfaces altered the spectrum of the incident light on each reflection, whereas the white surfaces did not.

The darkest regions contain information about the surface color, while the lightest regions contain information about the illuminant color.

Purple painted room under green illumination

Green painted room under purple illumination

Langer, M. and Gilchrist, A., “Color Perception in a 3-D scene of one reflectance”, ARVO 2000.

Consider the simple case of a green surface patch next to a blue surface patch. In R-G swapped Twin Earth, this will be a red surface patch next to the blue surface patch.

A green surface patch will contain more energy in blue wavelengths than a red surface patch, and hence a blue patch nearby will appear brighter next to the green patch than next to the red patch (and the green patch will be somewhat brighter as well than the red patch, due to the illumination from the blue patch).

A blue sensitive photoreceptor can therefore provide information about the red-ness or green-ness of a surface patch, independent of the red and green photoreceptors. Swapping the red and green photoreceptor connections will in no way alter this behaviour of the blue receptors.

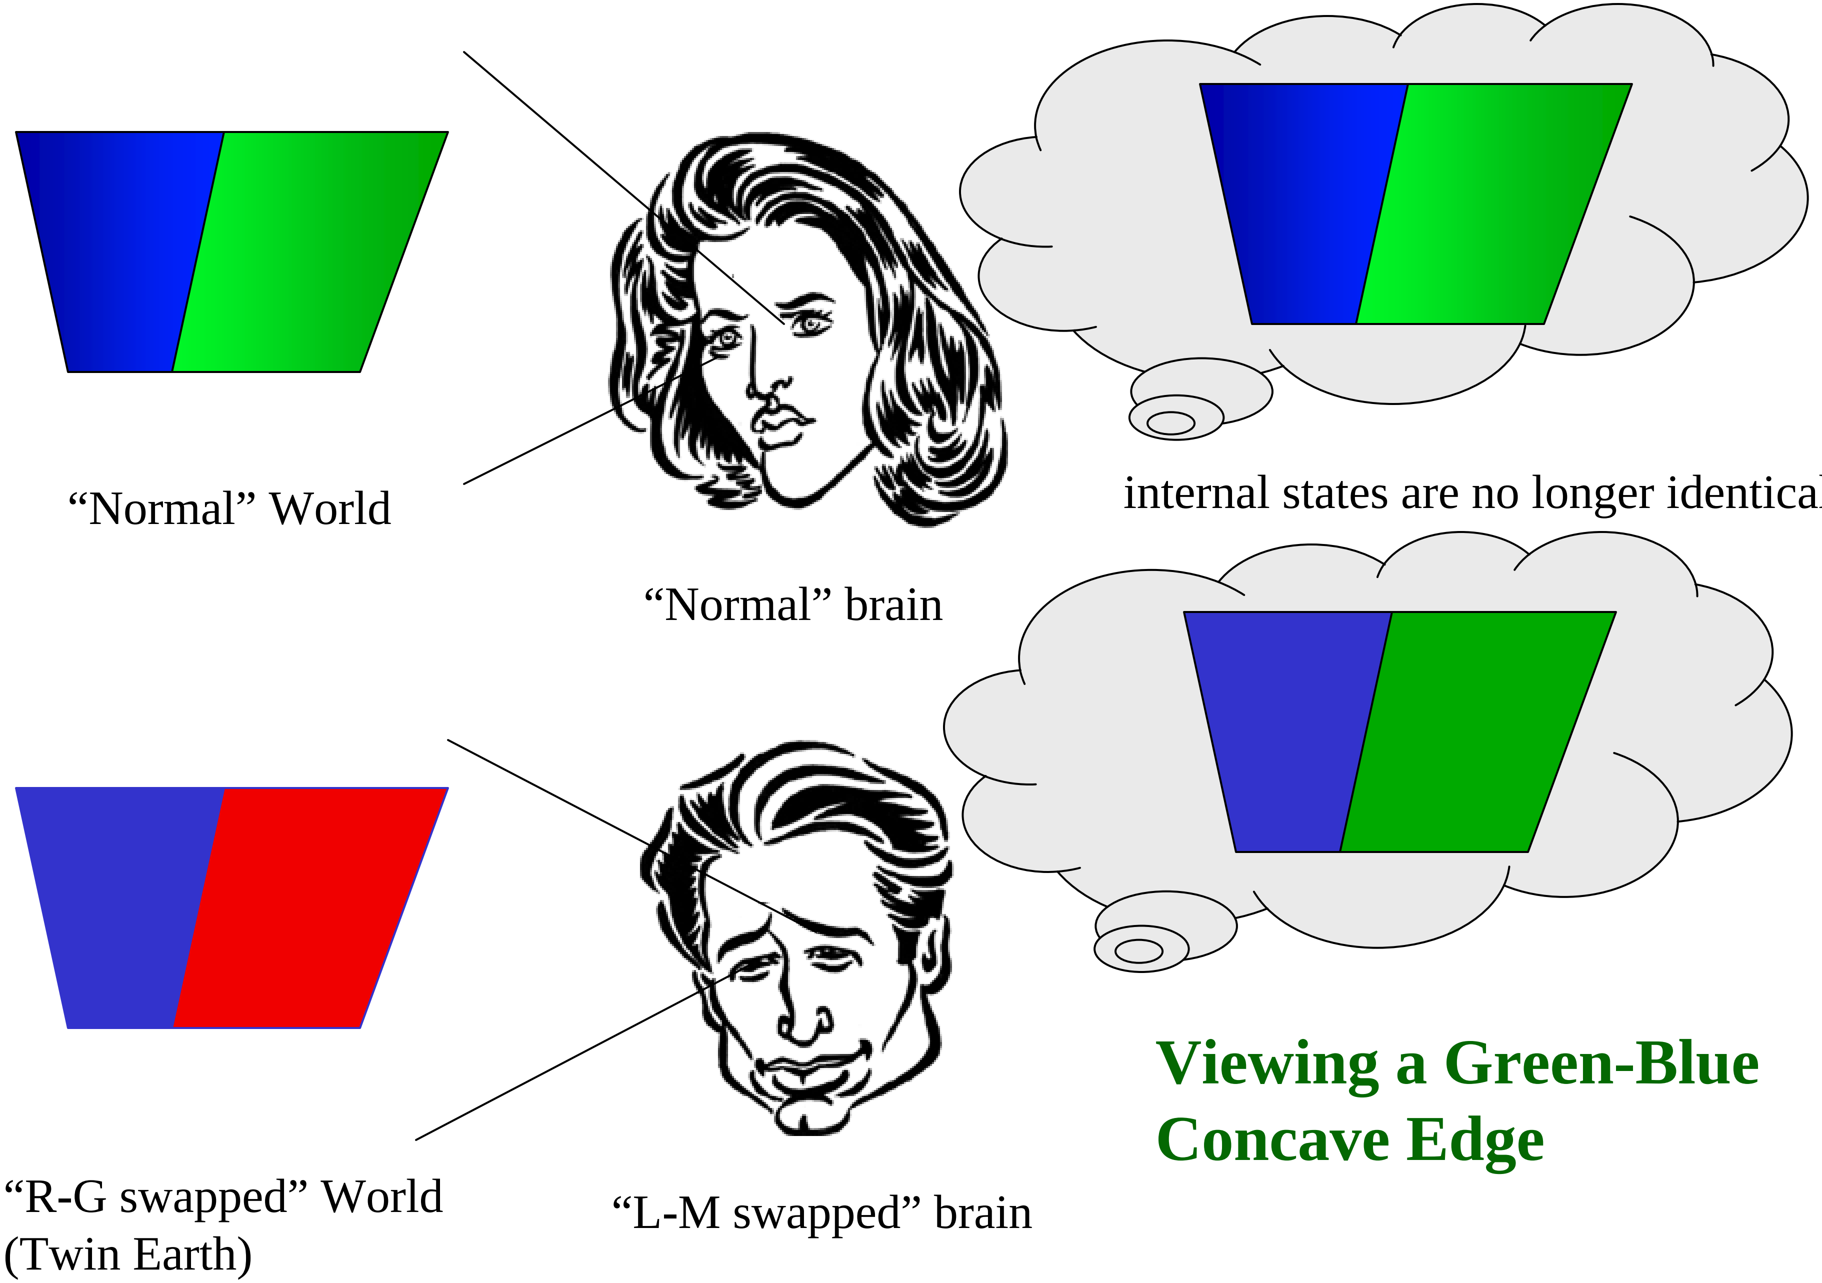
"Normal" World
"Normal" brain
internal states are no longer identical
"R-G swapped" World
(Twin Earth)
"L-M swapped" brain
Viewing a Green-Blue
Concave Edge

Once again, the internal states of observer A and observer B are no longer identical.

One could argue that it is possible to *paint* the blue and red surfaces in twin earth to give the *illusion of interreflection*, but this would only work, in general, for one configuration of the surfaces.

Relative surface motion would result in internal states that differ between observer A in normal earth and observer B in twin earth.

It could be that two asymmetric versions of world states nonetheless give the same internal physical states.

For example, the operator R could be such that the shaded and unshaded versions of the red patch both result in neural signals equivalent to that of an unshaded patch. But this would require that the neural signals be dependent on information about the relative position of the two patches in the world (since the shading depends on the relative position). There is evidence that the human visual system does this to some extent (e.g. colour constancy, or discounting of illuminants

Bloj *et al* demonstrated that subjects discount inter-reflections if they have evidence as to the 3D layout of surfaces.

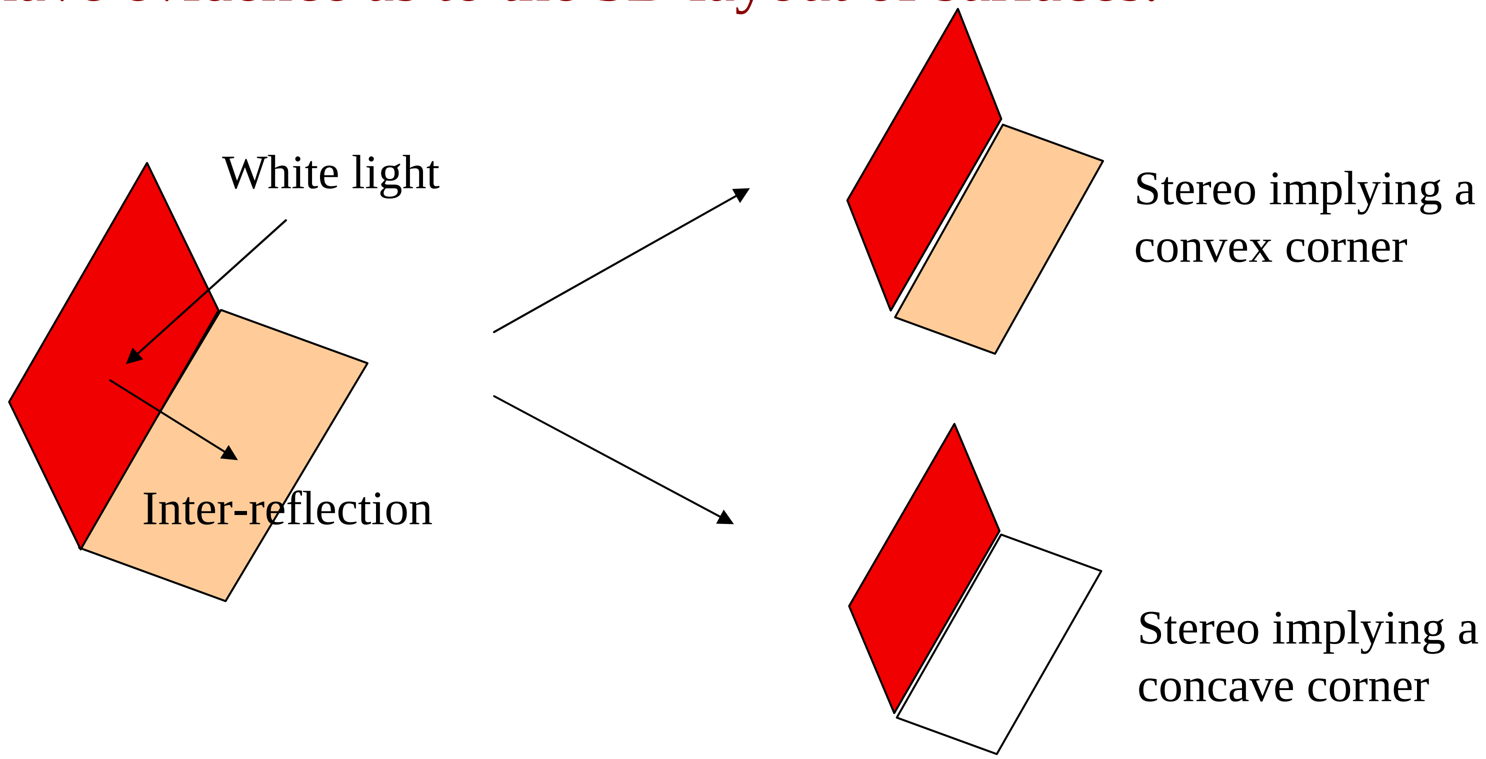

But this will not save the argument since the required information can be blocked by closing one eye (to remove stereo) or by occluding one of the patches from view.

Bloj, M. G., Kersten, D. & Hurlbert, A. C. (1999), Perception of three dimensional shape influences colour perception through mutual illumination. Nature, ***402****, 877-879.*

It could be that two asymmetric versions of internal physical states nonetheless give the same mental contents.

For example, the operator A could be such that the shaded and unshaded versions of the red patch both result in mental contents equivalent to that of an unshaded patch. But this would require that the mental contents be dependent on information about the relative position of the two patches in the world (since the shading depends on the relative position). This information is unlikely to be available to the organism.

Another solution is to relax the requirement that the operator A is the same for everyone. This is entirely possible, but …

If we accept the impossibility of a suitable twin world we are still left with the degenerate case where the twin world is the world itself. Thus it is possible that:

For any physically realizable world, W, given a specific systematic distortion of the neural signals, ***R->R'****, we have*

$$\boldsymbol{A \,.\, R' \,.\, W = A \,.\, R \,.\, W = I}$$

i.e.

$$\boldsymbol{A \,.\, R' = A \,.\, R}$$

This admits the possibility that even if the neural signals are distorted or scrambled, the operator *A* is such that the subjective experience, *I*, remains the same.

Are such operators conceivable?

Yes, the ***sensorimotor contingency theory*** of O’Regan and Noë proposes exactly this type of operation.

Are such operators implementable (in a machine)?

- In this conference we have heard the work of Philipona et al, who search to perceive spatial attributes independent of neural representations.
- Koenderink, among others, has also considered the problem of learning spatial relationships from neural signals with arbitrary codings.

Koenderink, J. J. (1984). The concept of local sign. In A. J. van Doorn, W. A. van de Grind, & J. J. Koenderink (Eds.), *Limits in perception* (pp. 495-547). Zeist, Netherlands: VNU Science Press.

How could we develop such algorithms?

- Scientific study and analysis of sensorimotor and ecological contingencies (this is already the basis of much work in computer vision).
- On-line learning of contingencies, through neural networks or formation of statistical models.

We propose that both of these approaches are useful for computer vision, but that the human brain probably uses the latter.
That is, the visual cortex is learning the sensory contingencies associated with self-action as well as (perhaps to a lesser extent) those associated with ecological interactions.

Our research program is to:

- discover and enumerate the various ecological and sensorimotor asymmetries and contingencies which are associated with scene qualities (such as colour and lightness)

- develop statistical learning algorithms for learning the ecological and sensorimotor contingencies from experience and exploration of the world