At the retina, there is *low-level visio*n, including light adaptation and the center-surround receptive fields of ganglion cells.

At the other extreme is *high-level visio*n, which includes cognitive processes that incorporate knowledge about objects, materials, and scenes.

In between there is *mid-level visio*n. Mid-level vision is simply an ill-defined region between low and high. The representations and the processing in the middle stages are commonly thought to involve surfaces, contours, grouping, and so on.

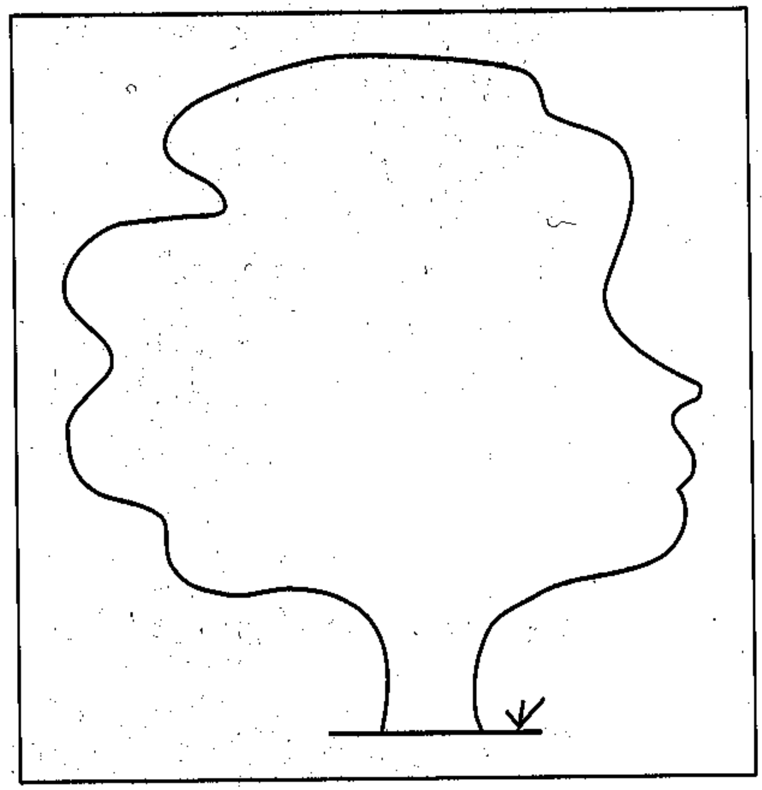

(a)

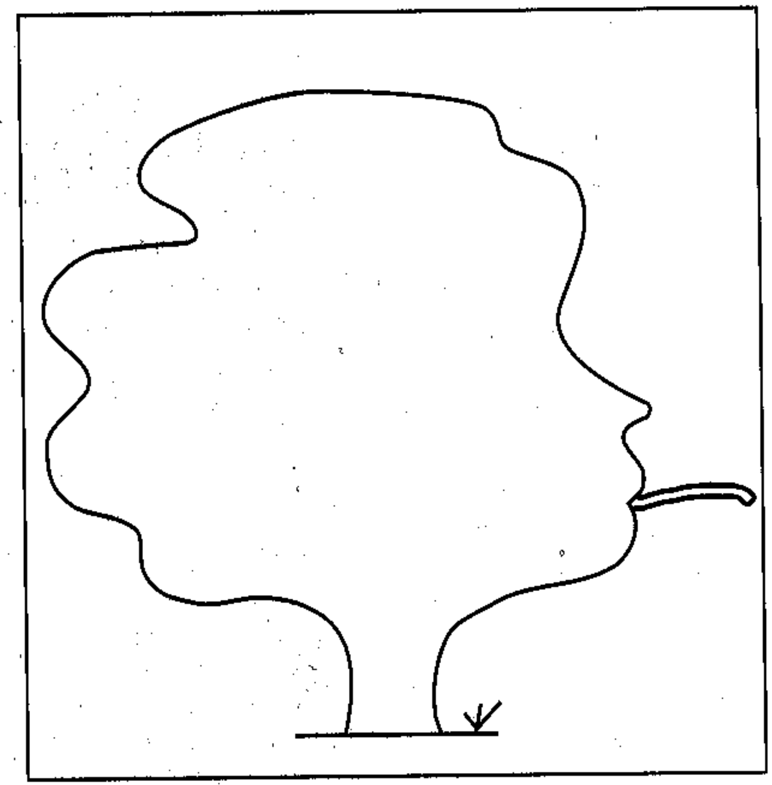

(b)

(a)

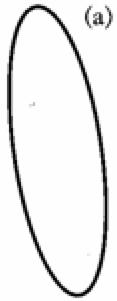

(b)

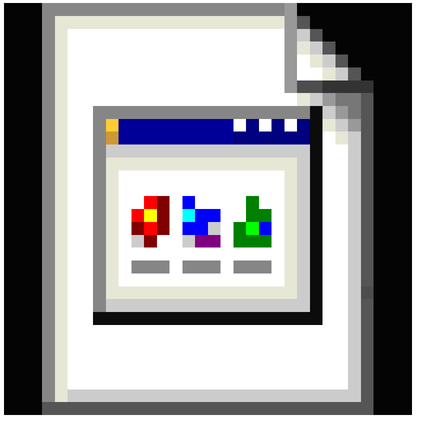

Ebbinghaus.exe

(A)

(B)

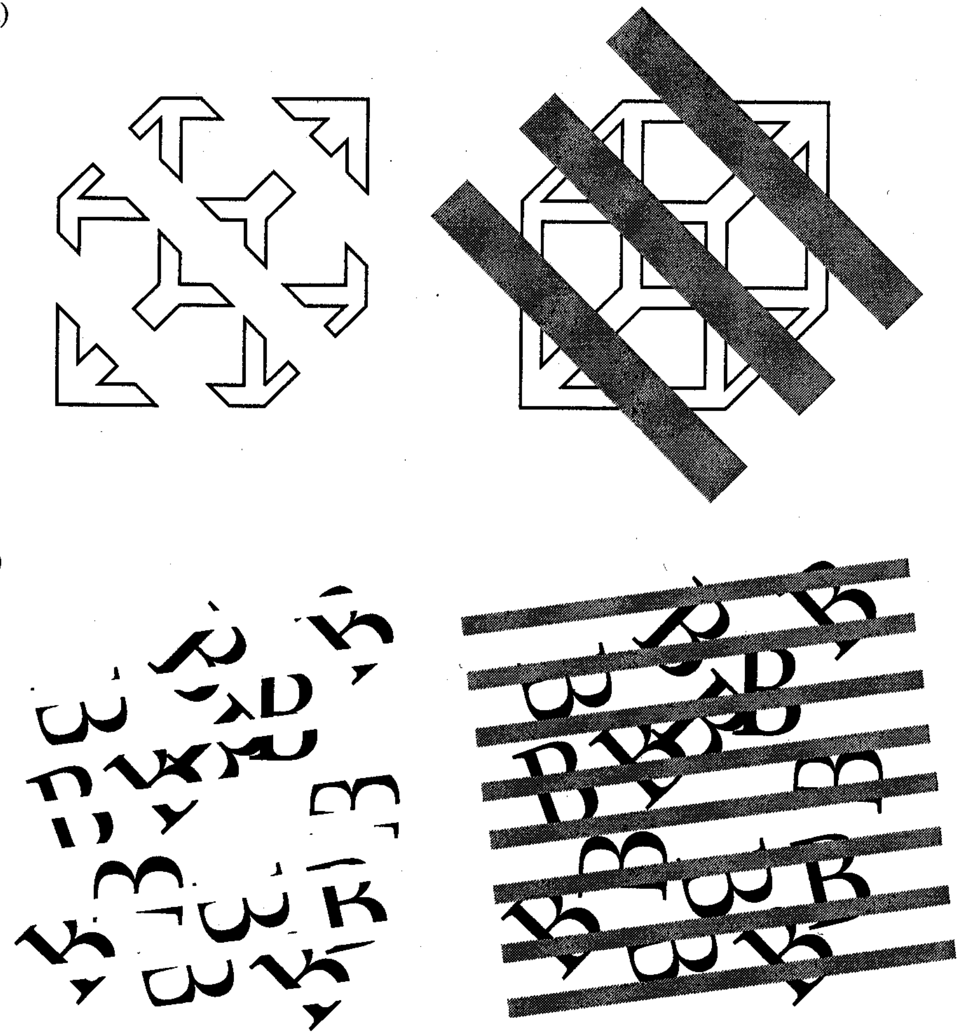

11.8 OCCLUSION AND OBJECT RECOGNITION. The presence of a clearly visible occluding surface helps us to integrate otherwise fragmentary image components. (A) When the line segments are seen without an occlusion cue, they appear as a set of uncorrelated two-dimensional patterns. By overlaying occluding boundaries, the pattern is seen as part of an object, namely, a three-dimensional cube. (B) When the pattern on the left is seen on its own, it appears as a jumble of unconnected curves and lines. By placing an occluding object over the white spaces, it is much easier to see that the occluded pattern is a collection of "B's." A after Kanizsa, 1979. B after Bregman, 1981.

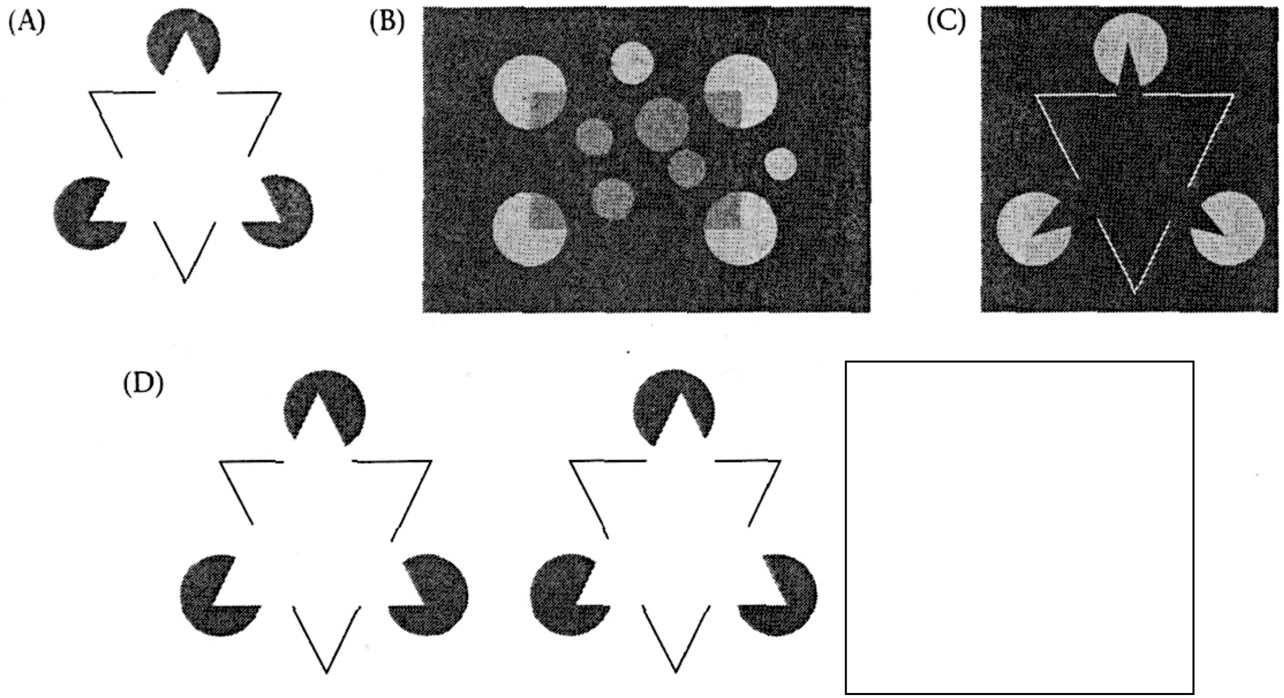

11.7 SUBJECTIVE CONTOURS. These subjective contours are inferred from occlusion and transparency cues in the images. (A) A triangle is suggested by occlusion, (B) a rectangle is suggested by transparency, and (C) a curved object is suggested by occlusion. (D) Stereo pairs of subjective contours. By diverging your eyes beyond the page, the image pair on the right (left) will fuse and you will see the subjective contours of a triangle in front (behind) of the circles. The subjective contour is somewhat more vivid when the depth cue is added. If you converge your eyes to fuse the images, the depth relationships will reverse. *A–C after Kanizsa, 1976. D after He and Nakayama, 1994b.*

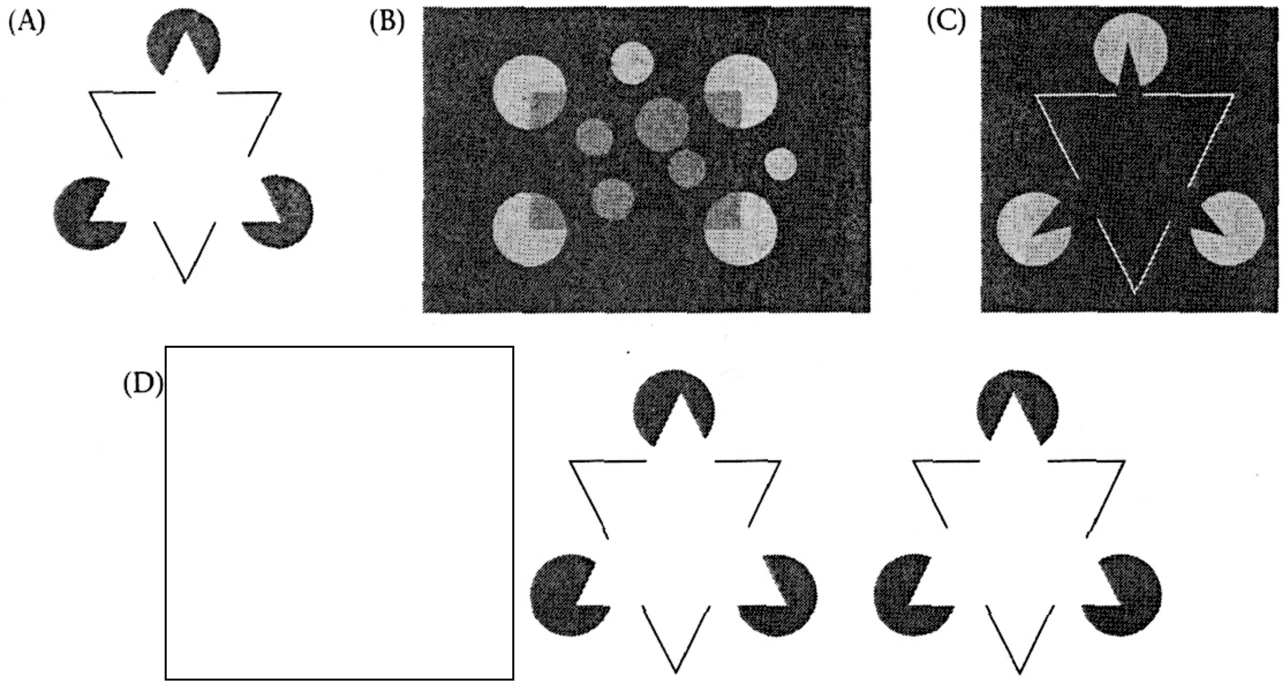

11.7 SUBJECTIVE CONTOURS. These subjective contours are inferred from occlusion and transparency cues in the images. (A) A triangle is suggested by occlusion, (B) a rectangle is suggested by transparency, and (C) a curved object is suggested by occlusion. (D) Stereo pairs of subjective contours. By diverging your eyes beyond the page, the image pair on the right (left) will fuse and you will see the subjective contours of a triangle in front (behind) of the circles. The subjective contour is somewhat more vivid when the depth cue is added. If you converge your eyes to fuse the images, the depth relationships will reverse. *A–C after Kanizsa, 1976. D after He and Nakayama, 1994b.*

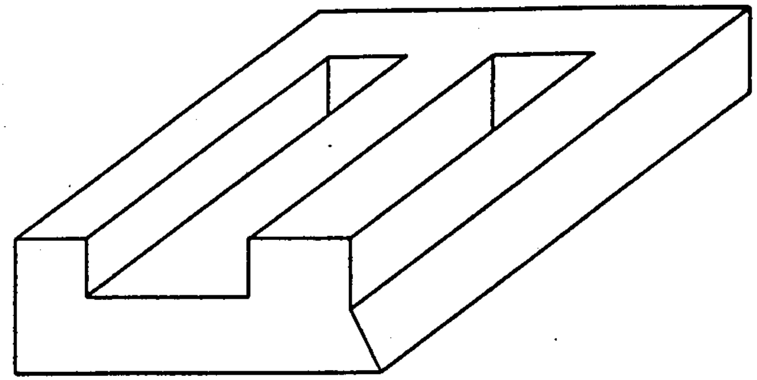

FIG.5.21. An impossible object. From Clowes (1971). Reproduced by permission of the North-Holland Publishing Company, Amsterdam.

Fig. 25 R.N. Shepard: The Doric Dilemma

FIGURE 7.24

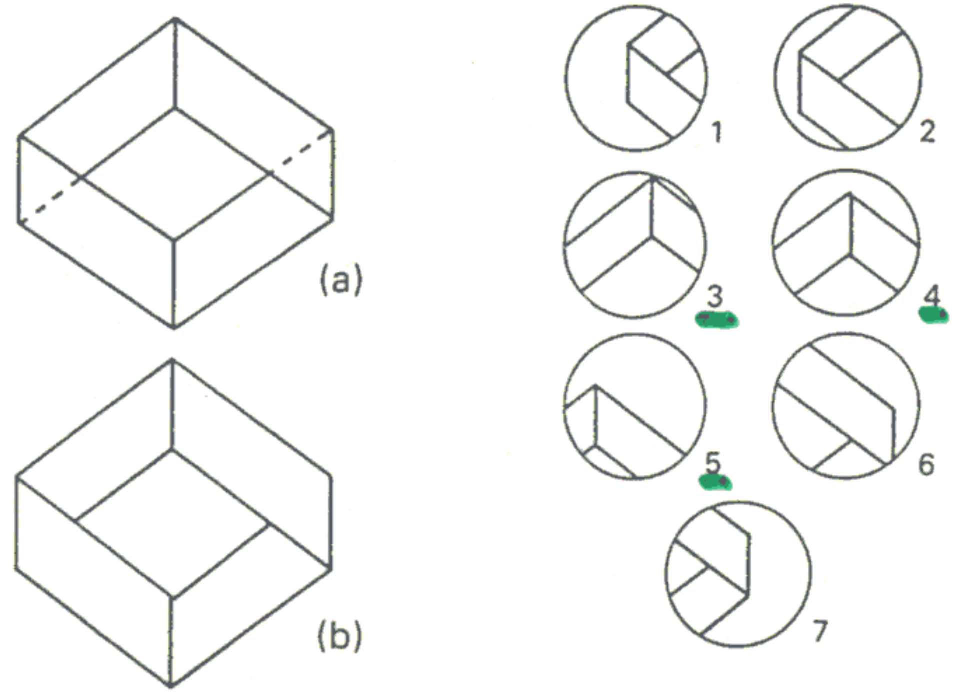

(a) shows a figure that is ambiguous in depth and (b) shows an impossible figure. Observers notice ambiguity and impossibility if a figure is viewed piece by piece through an aperture, as illustrated in frames 1–7. Adapted from Hochberg (1968).

FIGURE 6.20

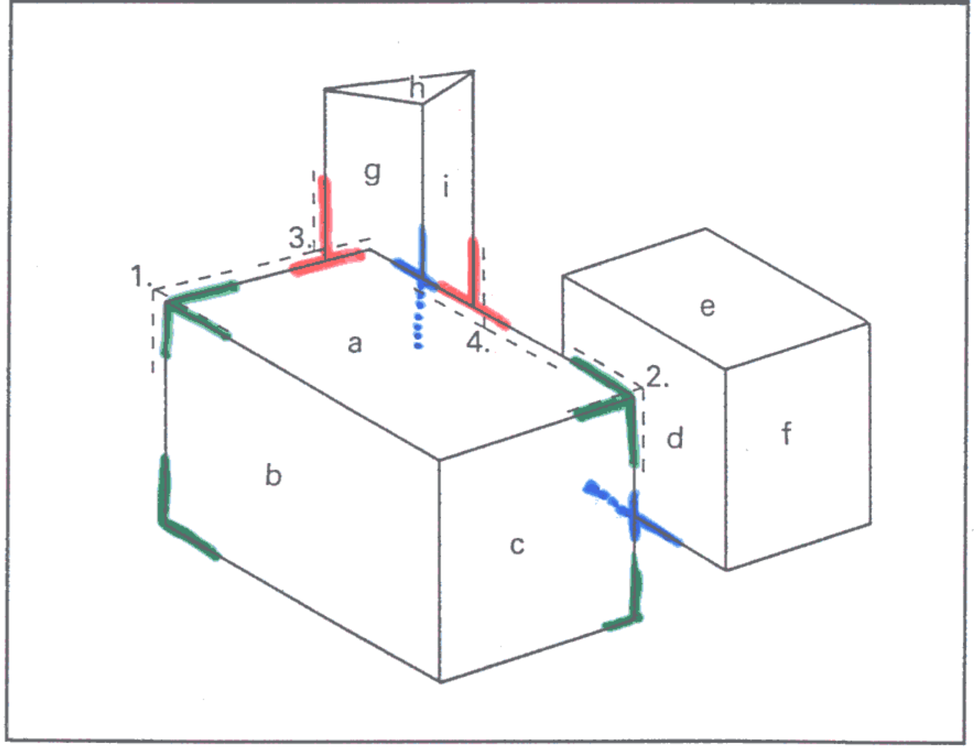

We have no difficulty in seeing that regions a, b, and c belong together (likewise d, e, and f; g, h, and i). Guzman's program SEE interpreted pictures like this by examining the junctions present. Examples of arrow junctions are shown at 1 and 2, and T junctions are shown at 3 and 4. What other kinds of junction are there in this picture?

T – occlusion → le pied du T appartient à l'objet caché

↑ – continuité de l'objet
L

X – Transparence

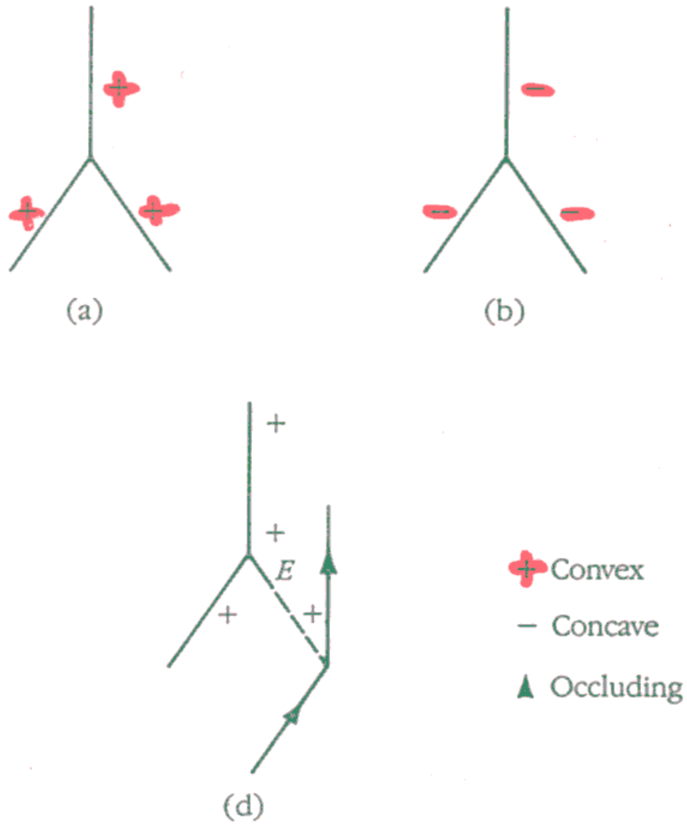

Figure 1–3. Some configurations of edges are physically realizable, and some are not. The trihedral junctions of three convex edges (a) or of three concave edges (b) are realizable, whereas the configuration (c) is impossible. Waltz cataloged all the possible junctions, including shadow edges, for up to four coincident edges. He then found that by using this catalog to implement consistency relations [requiring, for example, that an edge be of the same type all along its length like edge *E* in (d)], the solution to the labeling of a line drawing that included shadows was often uniquely determined.

Figure 3–1. The interpretation of some images involves more complex factors as well as more straightforward visual skills. This image devised by R. C. James may be one example. Such images are not considered here.

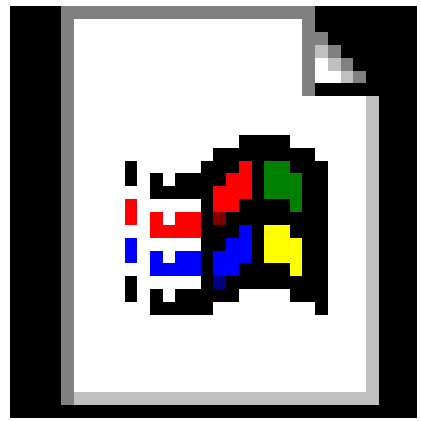

ThatcherFeynman.exe

HORIZON MOON

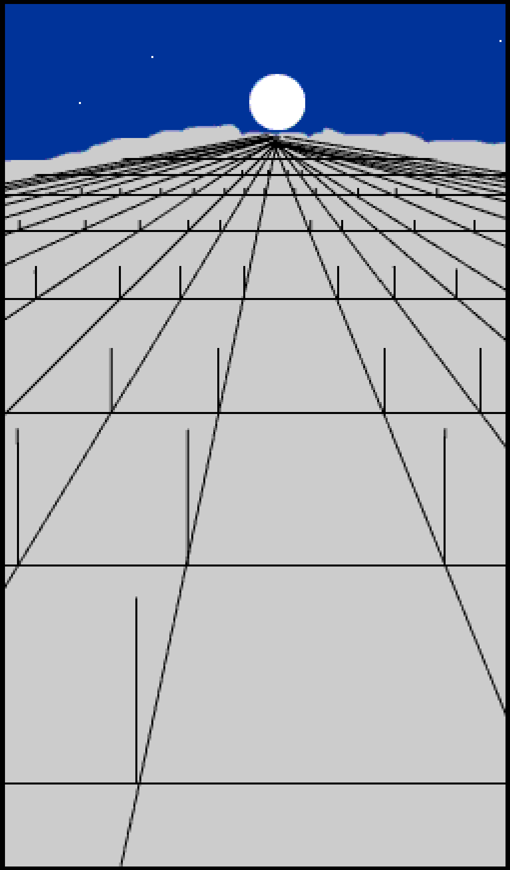

ZENITH MOON

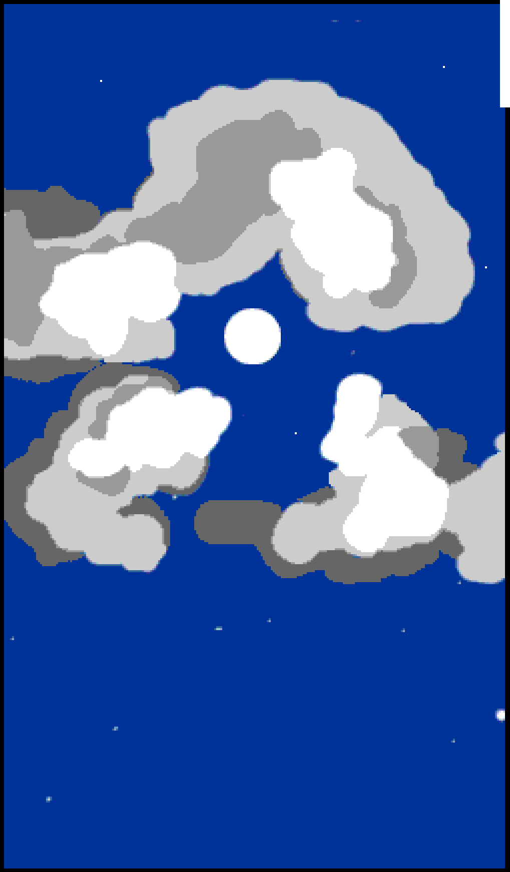

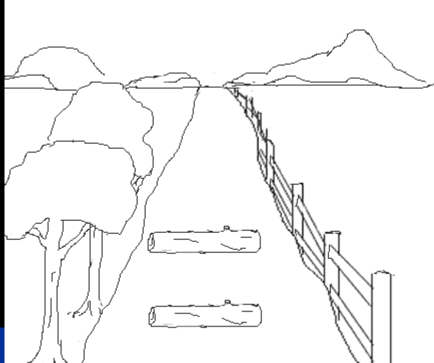

Illusion de Ponzo

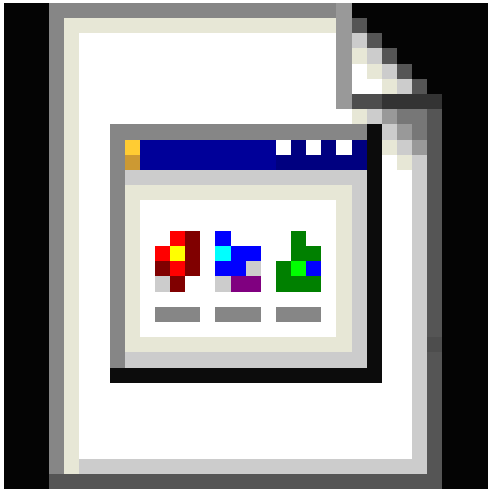

Planets.exe

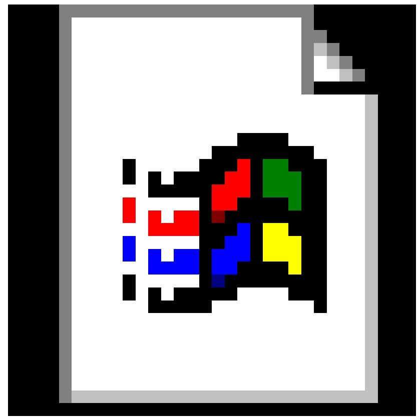

Perspective2.exe

LIGHTNESS ILLUSIONS

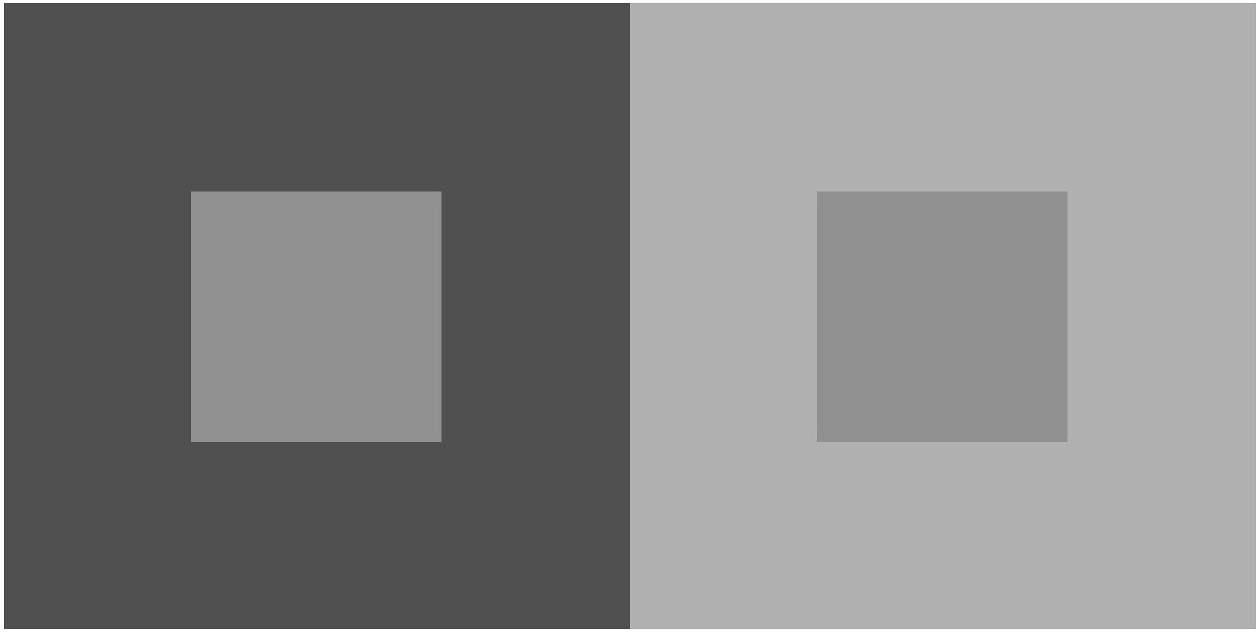

The simultaneous contrast effect.

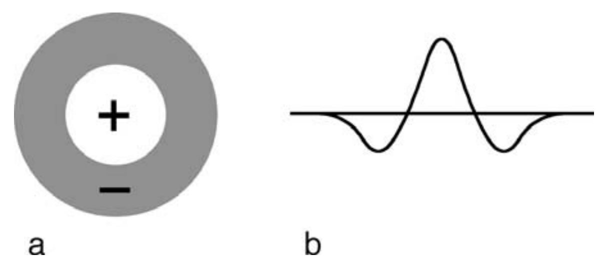

Figure 1 Static versions of the lightness illusions studied in our experiment (see also Supplementary Video 1). In a, the corresponding textured disks on the dark and light surrounds are physically identical, and in b the corresponding chess pieces on the two surrounds are identical. In both cases, the figures on the dark surround appear as light objects visible through dark haze, whereas the figures on the light surround appear as dark objects visible through light haze.

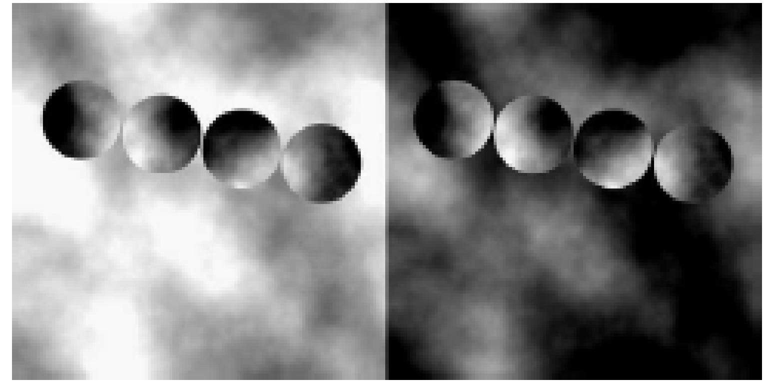

Figure 3 Transparency control experiment. The same targets and surrounds are used as in Fig. 1a, except that the surrounds have been rotated by 90 (see also Supplementary Video 3). This rotation destroys both the geometric and luminance conditions needed to evoke a percept of transparency, and also destroys the lightness illusion.

Anderson, B. & Winawer, J. (2005). Image segmentation and lightness perception. *Nature 434*, 79-83.

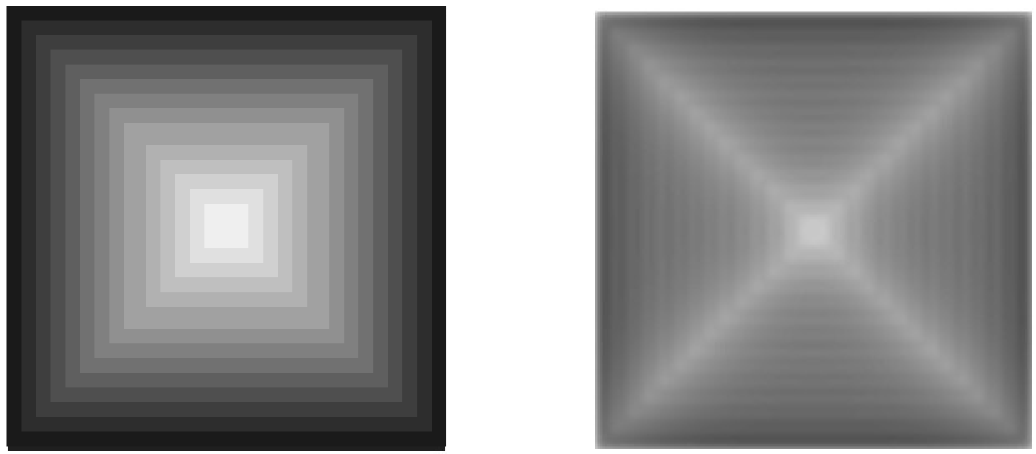

a b

An illusion by Vasarely, left, and a bandpass filtered version, right.

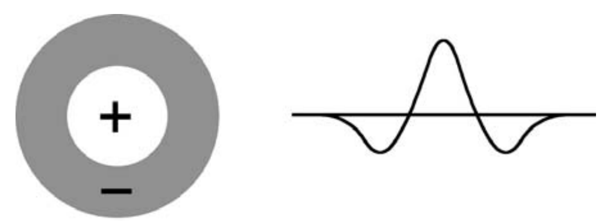

a b

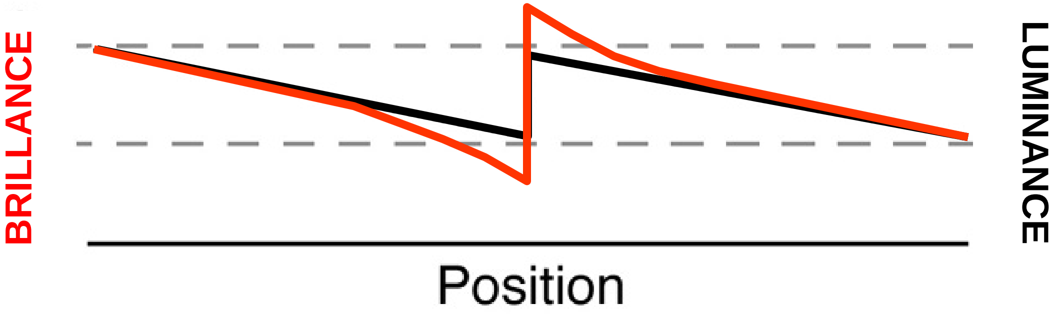

One version of the Craik-O'Brien-Cornsweet effect.

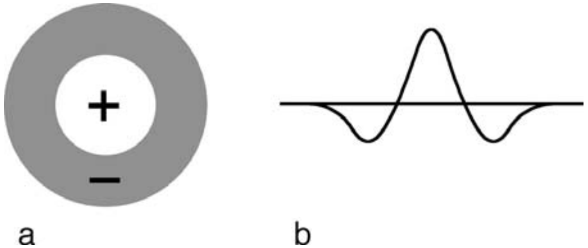

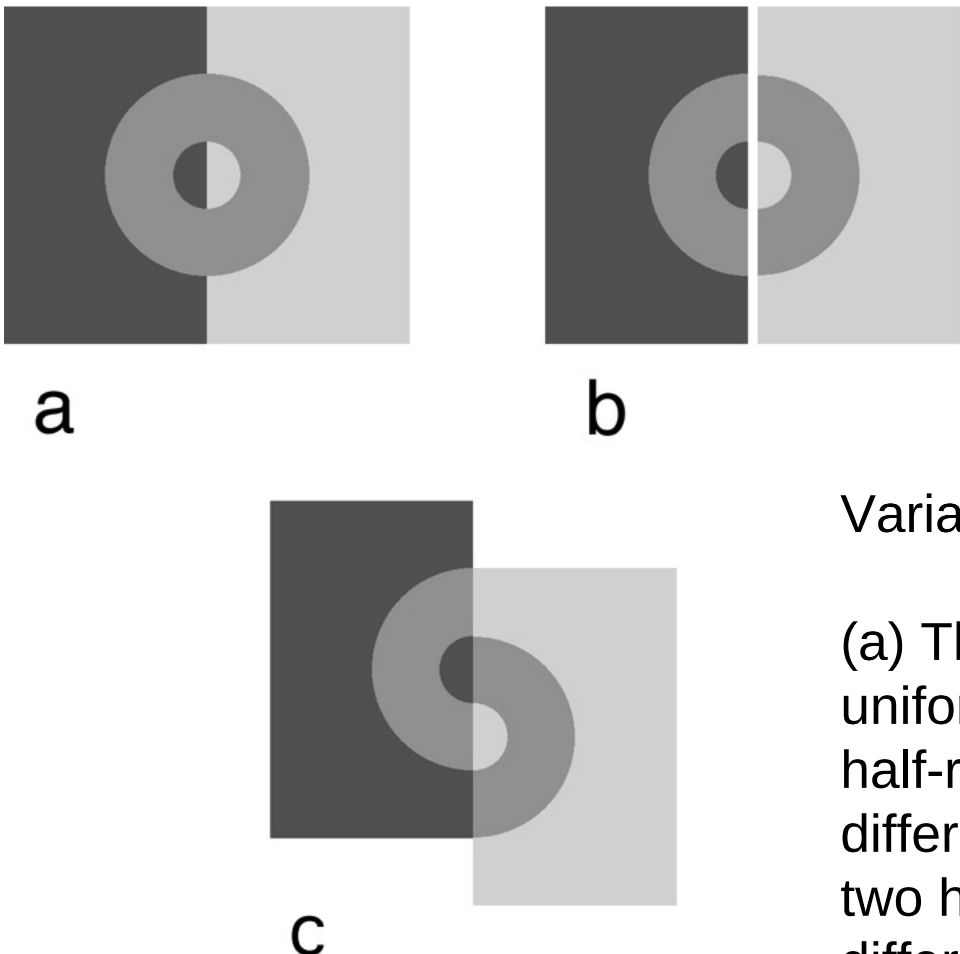

Variants on the Koffka ring.

(a) The ring appears almost uniform. (b) When split, the two half-rings appear distinctly different. (c) When shifted, the two half-rings appear quite different.

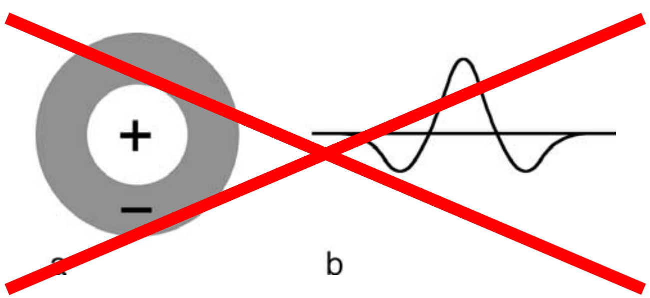

“Every light is a shade, compared to the higher lights, till you come to the sun; and every shade is a light, compared to the deeper shades, till you come to the night.”

John Ruskin (1879)

Cité par **E. H. Adelson** (2000) **Lightness Perception and Lightness Illusions,** (Chapter 24) in M. Gazzaniga, (Ed.), *The New Cognitive Neurosciences*, 2nd ed. Cambridge, MA: MIT Press, 339-351, 2000.

TOUT EST RELATIF

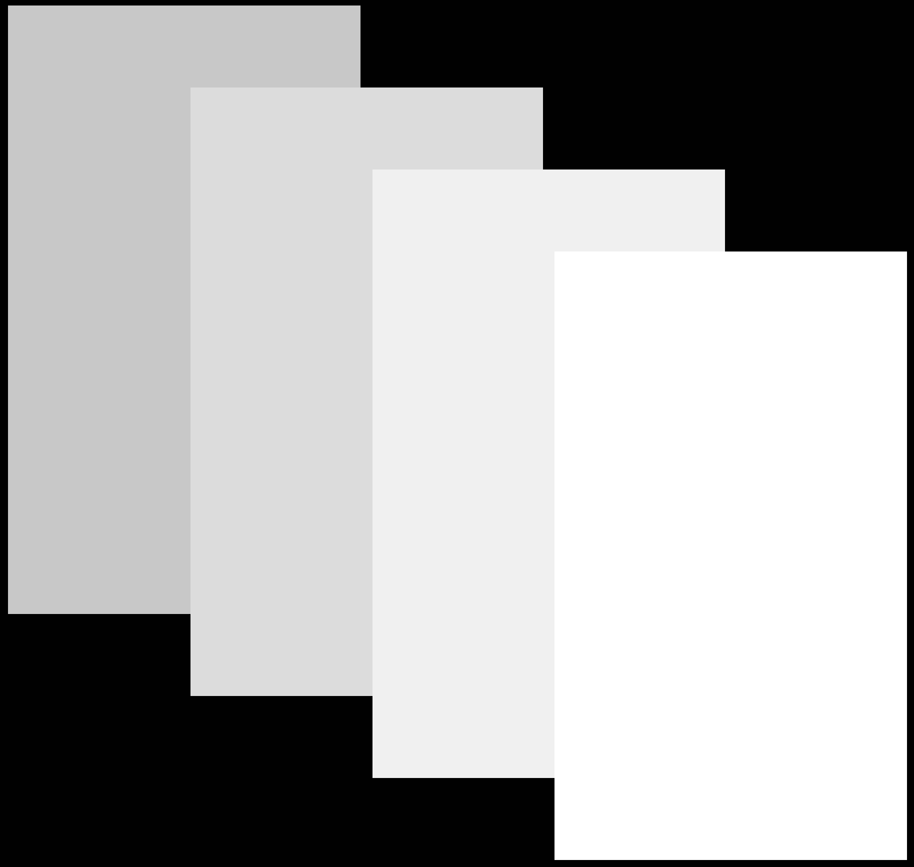

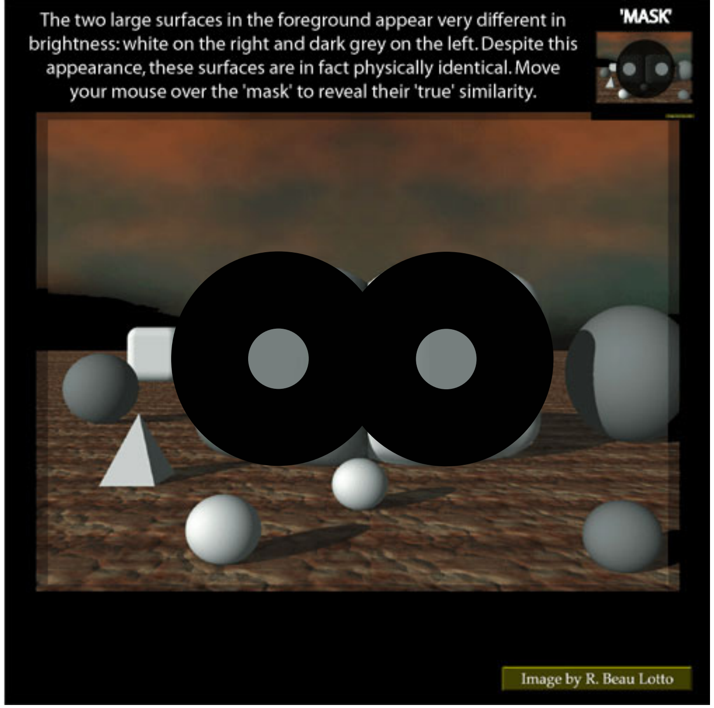
The two large surfaces in the foreground appear very different in brightness: white on the right and dark grey on the left. Despite this appearance, these surfaces are in fact physically identical. Move your mouse over the 'mask' to reveal their 'true' similarity.
'MASK'
Image by R. Beau Lotto

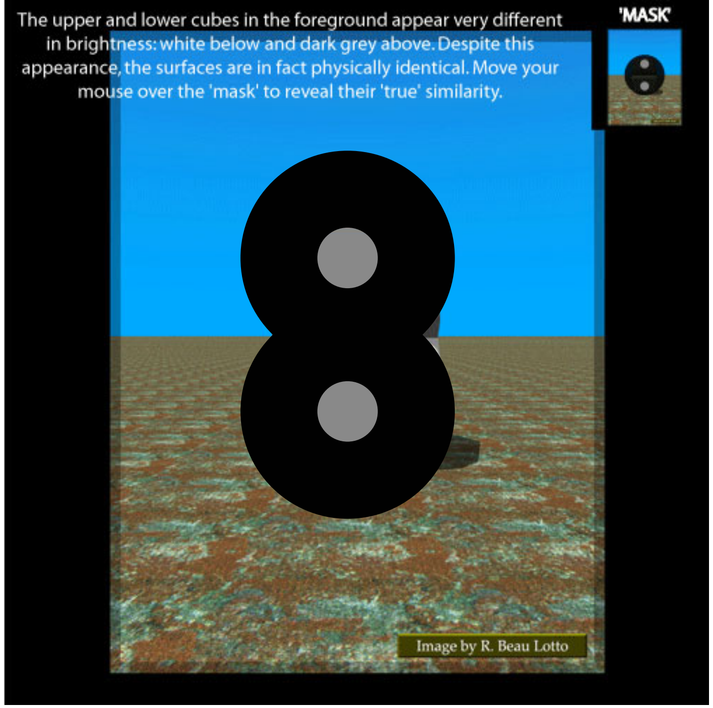
The upper and lower cubes in the foreground appear very different in brightness: white below and dark grey above. Despite this appearance, the surfaces are in fact physically identical. Move your mouse over the 'mask' to reveal their 'true' similarity.
'MASK'
Image by R. Beau Lotto

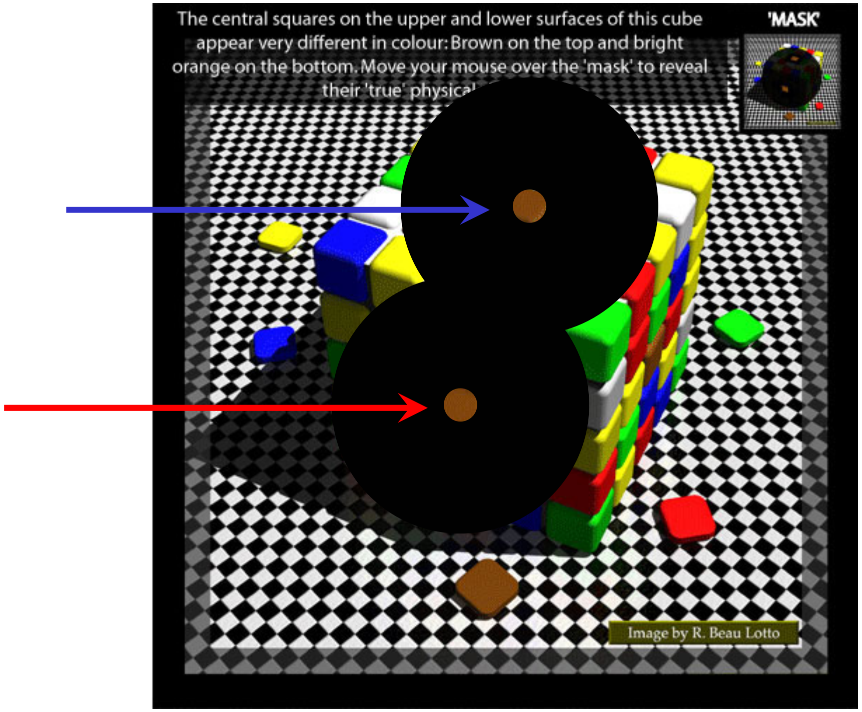
The central squares on the upper and lower surfaces of this cube appear very different in colour: Brown on the top and bright orange on the bottom. Move your mouse over the 'mask' to reveal their 'true' physical
'MASK'
Image by R. Beau Lotto

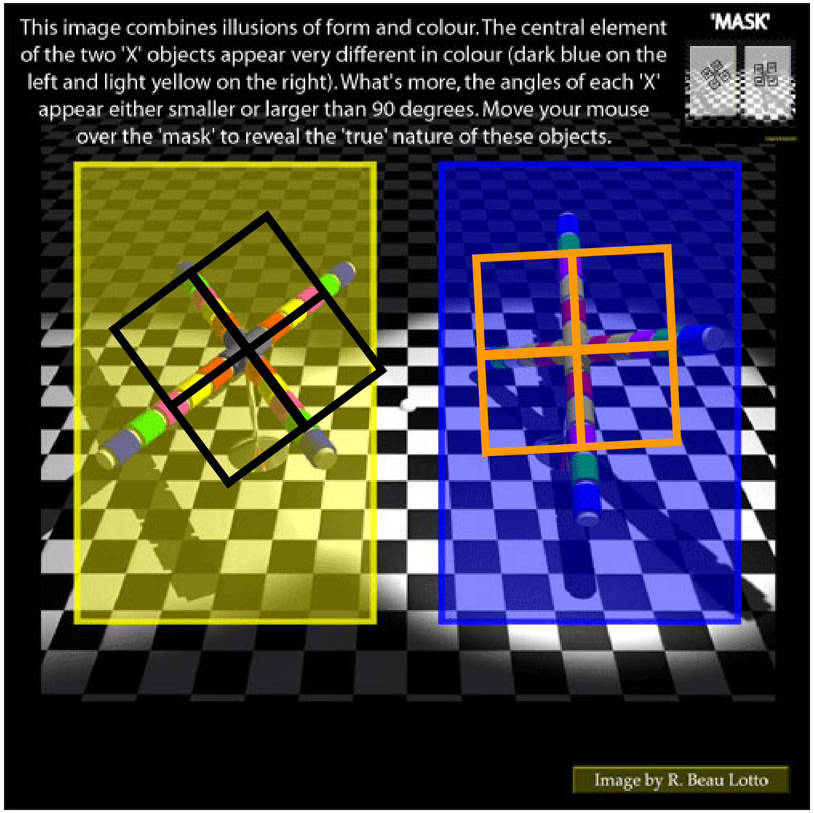
This image combines illusions of form and colour. The central element of the two 'X' objects appear very different in colour (dark blue on the left and light yellow on the right). What's more, the angles of each 'X' appear either smaller or larger than 90 degrees. Move your mouse over the 'mask' to reveal the 'true' nature of these objects.
'MASK'
Image by R. Beau Lotto

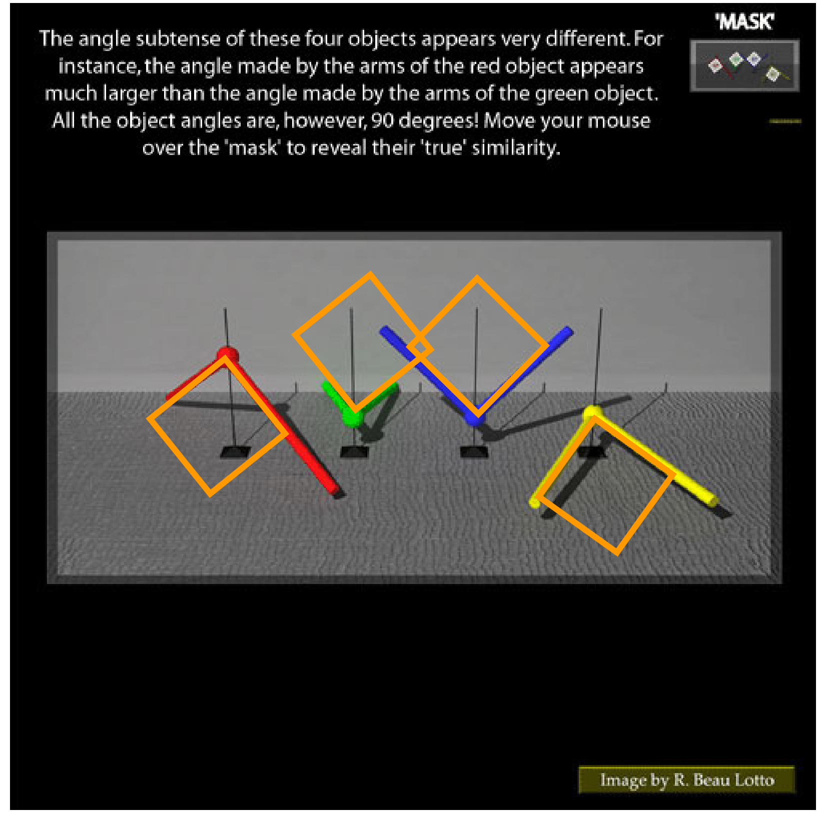
The angle subtense of these four objects appears very different. For instance, the angle made by the arms of the red object appears much larger than the angle made by the arms of the green object. All the object angles are, however, 90 degrees! Move your mouse over the 'mask' to reveal their 'true' similarity.
'MASK'
Image by R. Beau Lotto

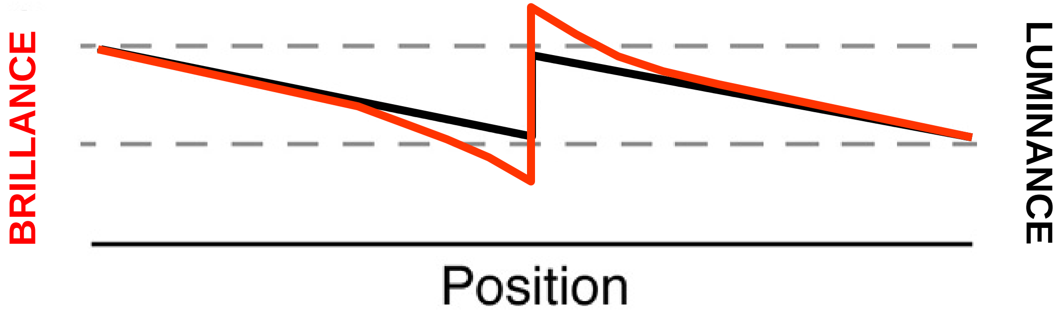

One version of the Craik-O'Brien-Cornsweet effect.

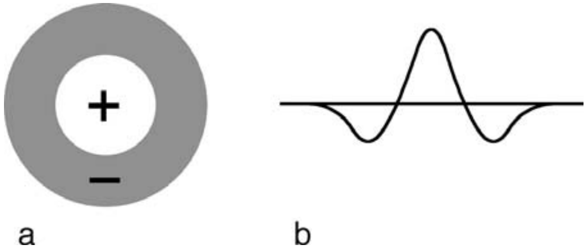

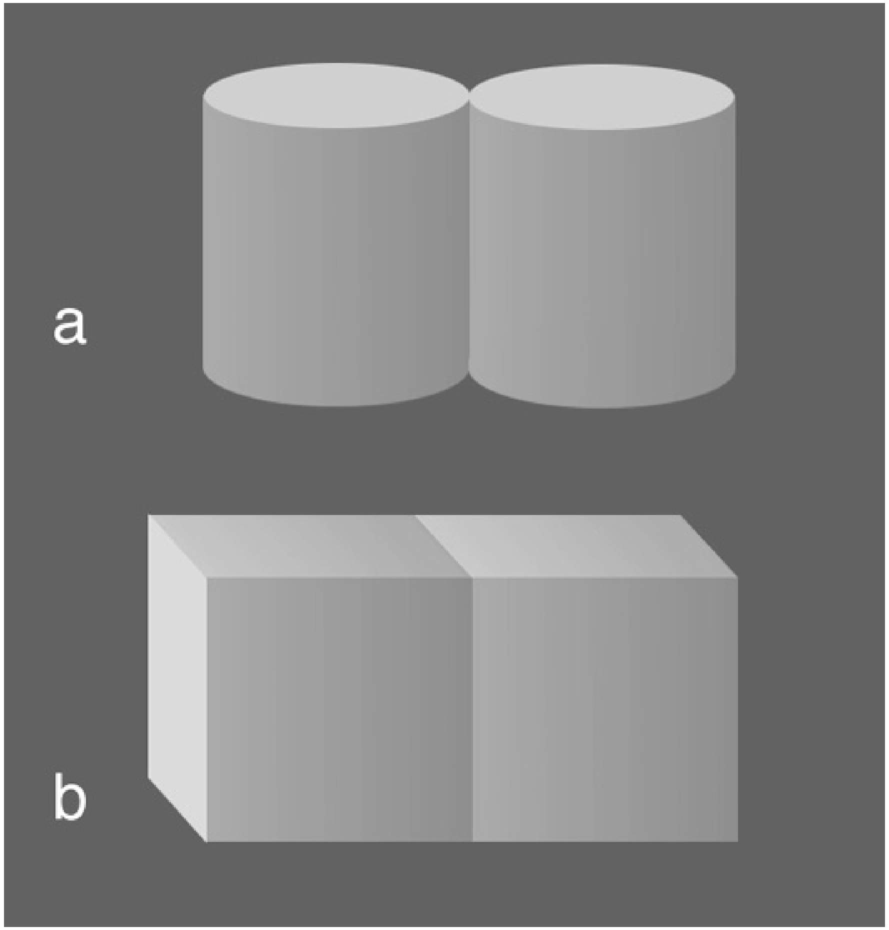

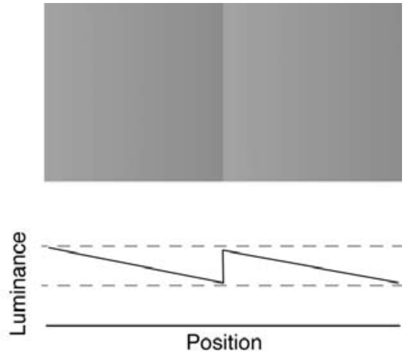

Knill and Kersten's illusion. Both figures contain the same COCE ramps, but the interpretations are quite different.

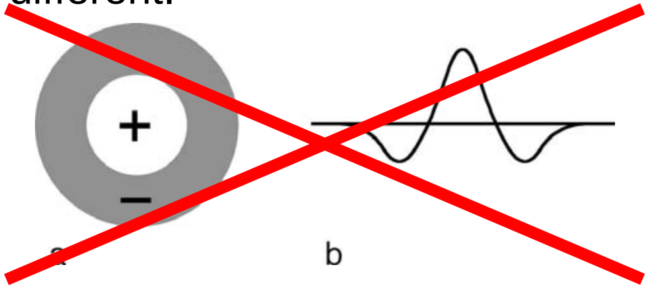

Lightness Perception and Lightness Illusions

Edward H. Adelson

Department of Brain and Cognitive Sciences
Massachusetts Institute of Technology

Chapter 24 in M. Gazzaniga, (Ed.),
The New Cognitive Neurosciences, 2nd ed.

Cambridge, MA: MIT Press, 339-351, 2000.

"Every light is a shade, compared to the higher lights, till you come to the sun; and every shade is a light, compared to the deeper shades, till you come to the night."

—John Ruskin, 1879.

Mid-level vision is simply an ill-defined region between low and high. The representations and the processing in the middle stages are commonly thought to involve surfaces, contours, grouping, and so on. Lightness perception seems to involve all three levels of processing [low-, mid- and high-level vision].

Adelson E.H. (2000). Lightness Perception and Lightness Illusions. In M. Gazzaniga, (Ed.), *The New Cognitive Neurosciences*, 2nd ed., Cambridge, MA: MIT Press, 339-351, 2000.

Vision is only possible because there are constraints in the world, i.e., images are not formed by arbitrary random processes.

To function in this world**, the visual system must exploit the ecology of** images—it must “know” the likelihood of various things in the world, and the likelihood that a given image-property could be caused by one or another world-property**.**

This world-knowledge may be hard-wired or learned, and may manifest itself at various levels of processing.

Objective/physical quantities

Subjective quantities

- ***p*** and ***q***: same reflectance, but different luminances.
- ***q*** and ***r***: different reflectances and different luminances; but same illuminance.
- ***p*** and ***r*** happen to have the same luminance (the lower reflectance of *p* is counterbalanced by its higher luminance).

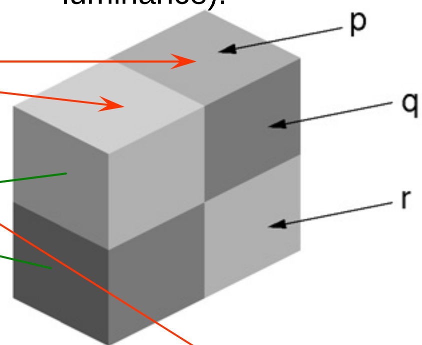

Illuminance: the amount of light incident on a surface. It varies from 0 to ∝.

Luminance: the amount of visible light that comes to the eye from a surface. It varies from 0 to ∝.

%

Brightness: the perceived intensity of light coming from the image itself, rather than any property of the portrayed scene. Brightness is sometimes defined as perceived luminance.

Reflectance (or *albedo*): the *proportion* of incident light that is reflected from a surface. It varies from 0 - ideal black - to 1 - ideal white.

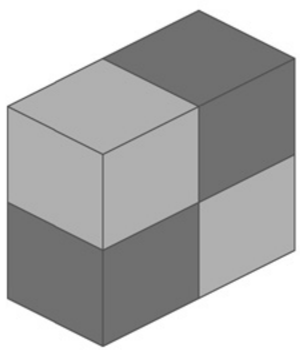

Reflectance image

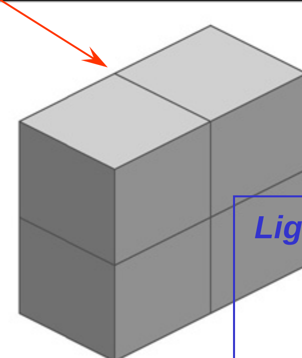

Illuminance image

Lightness: the perceived reflectance of a surface. It represents the visual system's attempt to extract reflectance based on the luminances in the scene.

Objective/physical quantities

Subjective quantities

***Illuminance**:* the amount of light incident on a surface. It varies from 0 to $\propto$.

Luminance: the amount of visible light that comes to the eye from a surface. It varies from 0 to $\propto$.

%

Reflectance (or *albedo):* the *proportion* of incident light that is reflected from a surface. It varies from 0 - ideal black - to 1 - ideal white.

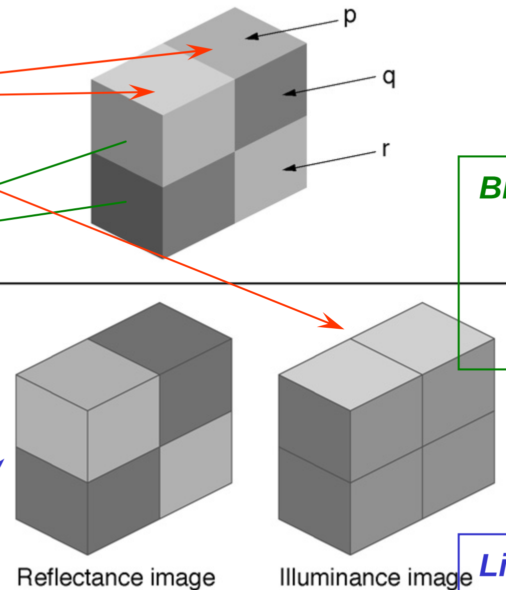

***Brightness**:* the perceived intensity of light coming from the image itself, rather than any property of the portrayed scene. Brightness is sometimes defined as perceived luminance.

***Lightness**:* the perceived reflectance of a surface. It represents the visual system's attempt to extract reflectance based on the luminances in the scene.

The problem of lightness constancy

From a physical point of view, the problem of lightness constancy is as follows. An illuminance image, $E(x,y)$, and a reflectance image, $R(x,y)$, are multiplied to produce a luminance image, $L(x,y)$:

$$L(x,y) = E(x,y)R(x,y).$$

An observer is given L at each pixel, and attempts to determine the two numbers E and R that were multiplied to make it. Unfortunately, unmultiplying two numbers is impossible. If $E(x,y)$ and $R(x,y)$ are arbitrary functions, then for any $E(x,y)$ there exists an $R(x,y)$ that produces the observed image. The problem appears impossible, but humans do it pretty well. This must mean that **illuminance and reflectance images are not arbitrary functions**. **They are constrained by the statistical properties of the world**.

Land and McCann argued that **reflectance tends to be constant across space except for abrupt changes at the transitions between objects or pigments**. Thus a reflectance change shows itself as step edge in an image, while illuminance will change only gradually over space.

By this argument one can separate reflectance change from illuminance change by taking spatial derivatives: *high derivatives are due to reflectance and low ones are due to illuminance*.

The Retinex model applies a derivative operator to the image, and thresholds the output to remove illuminance variation. The algorithm then reintegrates edge information over space to reconstruct the reflectance image.

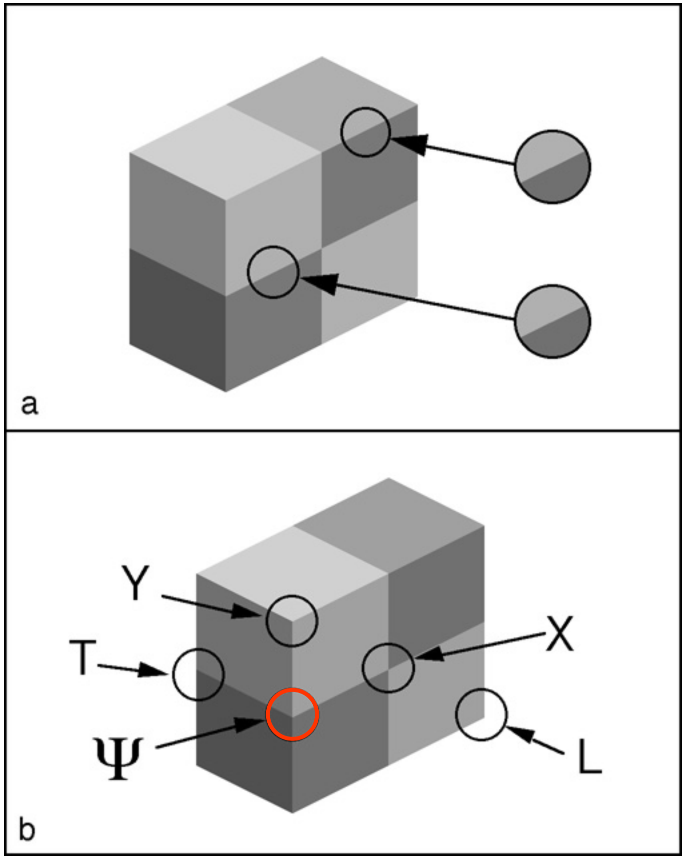

The above constraints fail when applied to the checker-block image. Figure (**a**) shows two light-dark edges. They are exactly the same in the image, and any local edge detector or filter will respond to them in the same way. Retinex will classify both as reflectance steps. Yet they have very different meanings. One is caused by illuminance (due to a change in surface normal); the other is caused by reflectance.

To interpret the edges, the visual system must consider them in a larger context. One good source of information is the junctions, such as those labeled in figure (**b**). The configuration of a junction, as well as the gray levels forming the junction, can offer cues about the shading and reflectance of a surface. Particularly strong constraints are imposed by a ψ-junction:

The vertical spine appears to be a dihedral with different illuminance on the two sides.

The angled arms appear to represent a reflectance edge that crosses the dihedral. The ratios of the gray levels and the angles of the arms are consistent with this interpretation.

L’influence d’une jonction en ψ peut se propager le long des contours qui se rencontrent en ψ. L’interprétation d’un seul bord clair-sombre, ambigü en lui-même peut-être biaisée vers une interprétation particulière par les ψ-s voisins.

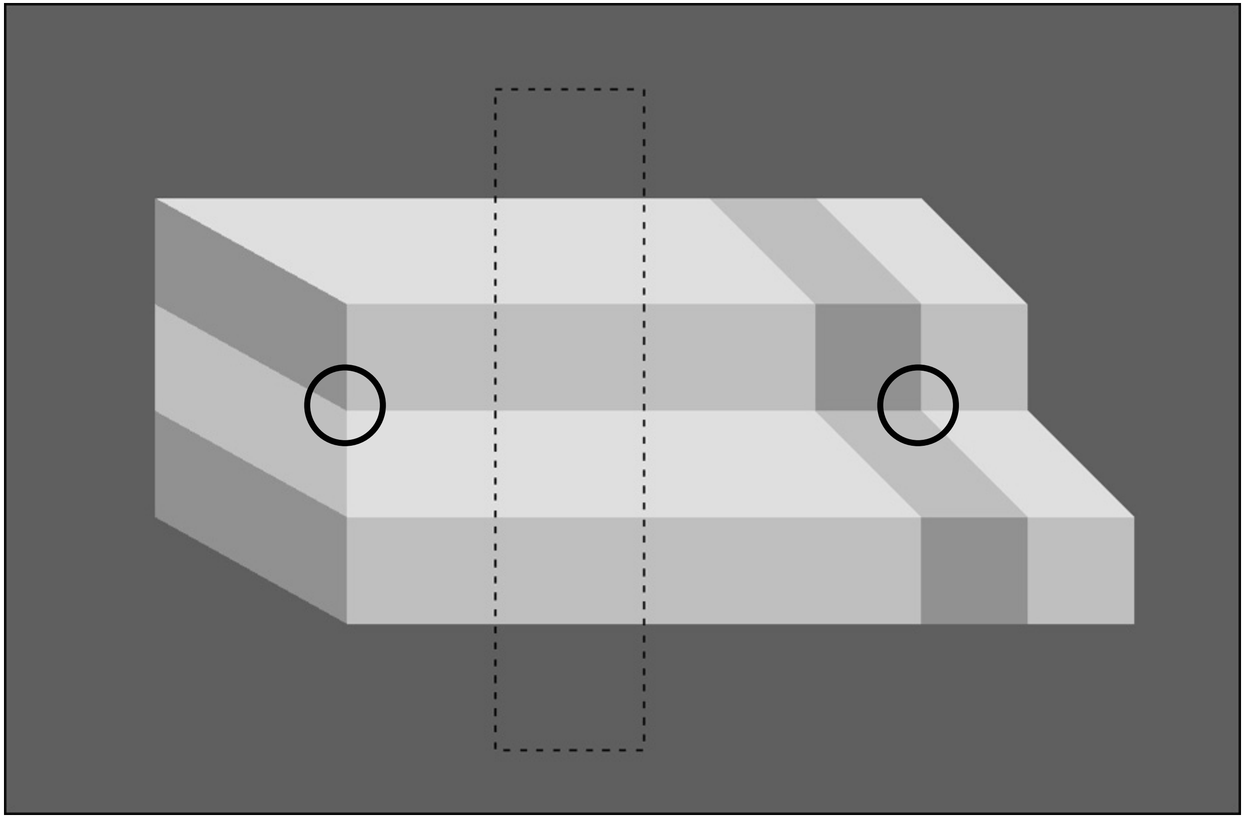

The impossible steps.

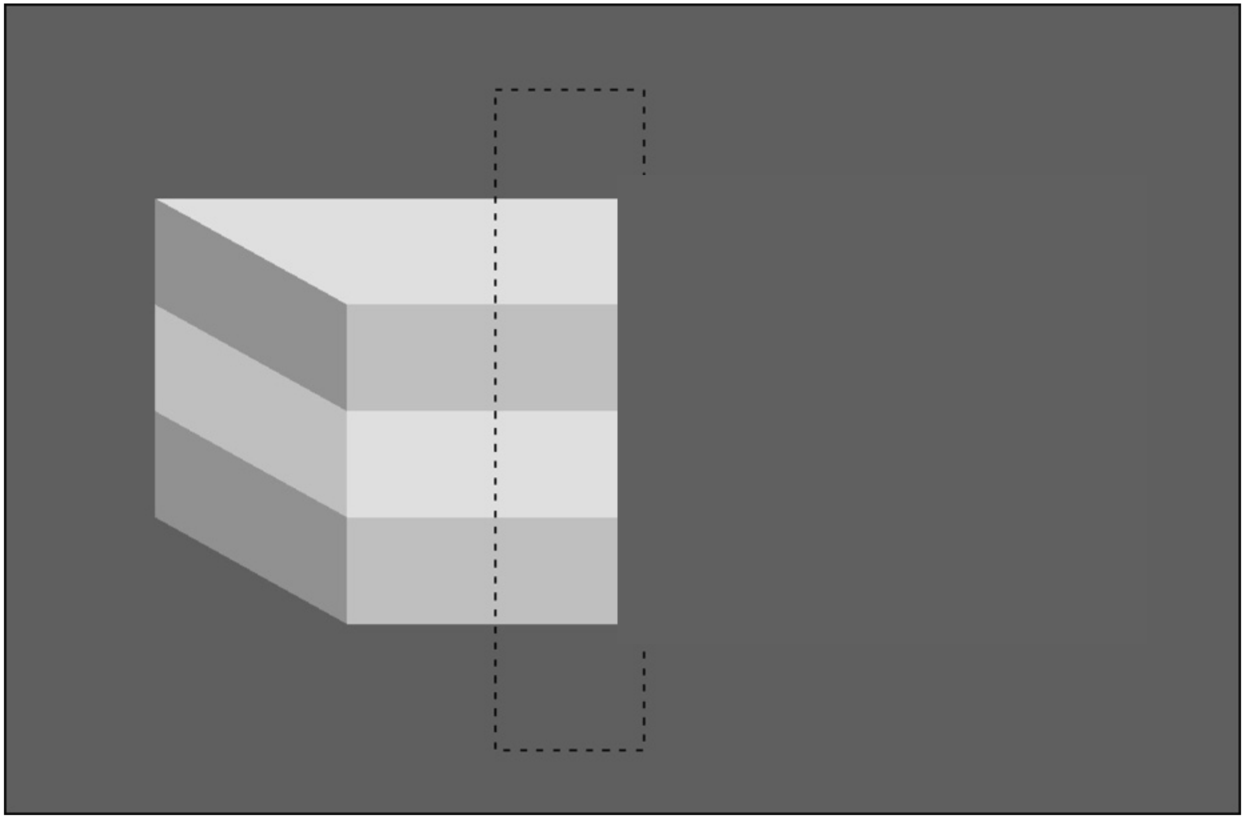

The horizontal strips appear to be due to paint.

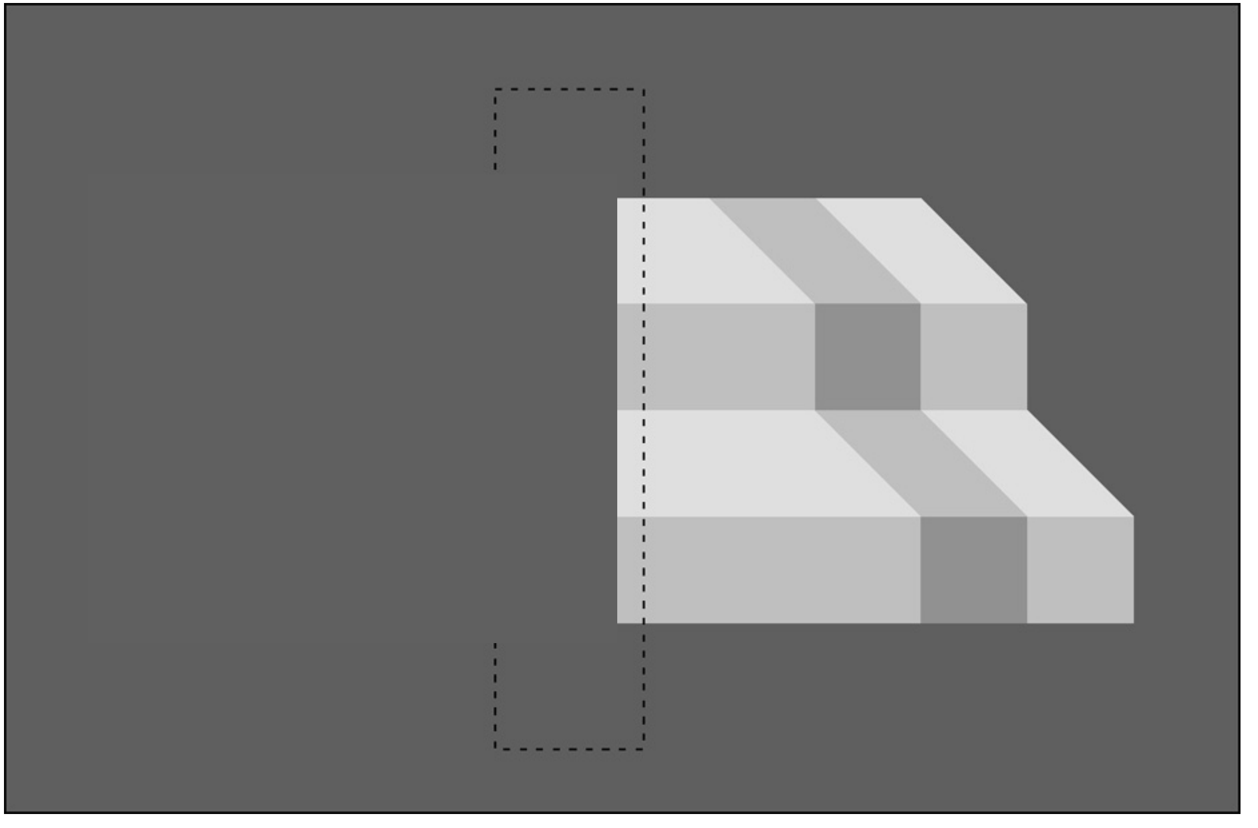

On the right, the horizontal strips appear to be due to shading.

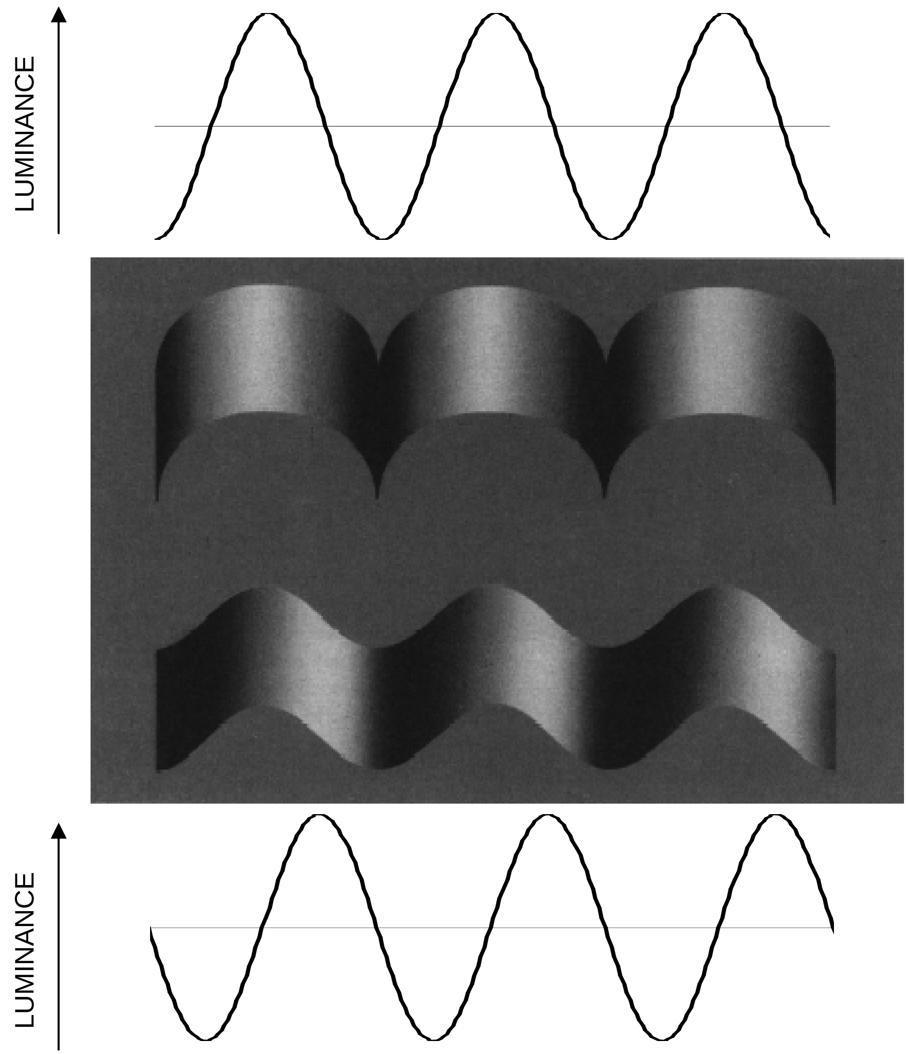

These two images are identical in their shading information but only differ in their boundary edge.
Despite being so similar the percept of these images is very different.
The upper image looks like three cylinders next to each other while the bottom shape looks like corrugated metal.

Les configurations des jonctions en ψ peuvent moduler la brillance (intensité perçue).

Pas/peu d'illusion:
- même plan de profondeur
- même environnement dans ce plan (contrastes locaux)

Illusion:
- plans de profondeur différents, donc
- présupposition d'éclairage différent (ombres), et
- contrastes locaux différents

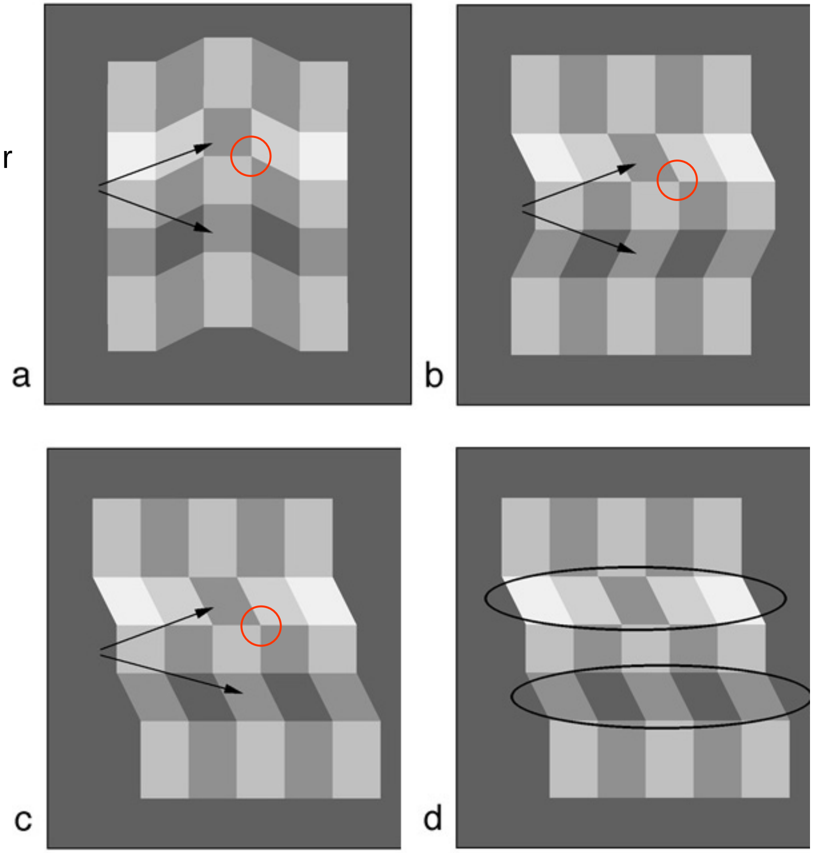

Illusion:
- plans de profondeur différents, donc
- mêmes présuppositions d'éclairage
- Mais contrastes locaux différents

Possible groupings induced by junctions

The corrugated plaid

In figure (**a**) the two marked patches are the same shade of gray. The upper patch appears slightly darker. In (**b**) & (**c**) only the geometry has been changed, parallelograms having been substituted for squares and vice-versa. The illusion is much enhanced. A low-level filtering mechanism, or a mechanism based on local edge interactions, cannot explain the change in the illusion.

Interpretation: the change in ψ-junctions causes a change in the perception of 3D surface orientation and shading. In (**a**) the two test patches appear to be in the same illumination, but in **(b**, **c**) they are differently illuminated. A brightly lit patch of dark gray looks quite different from a dimly lit patch of light gray.

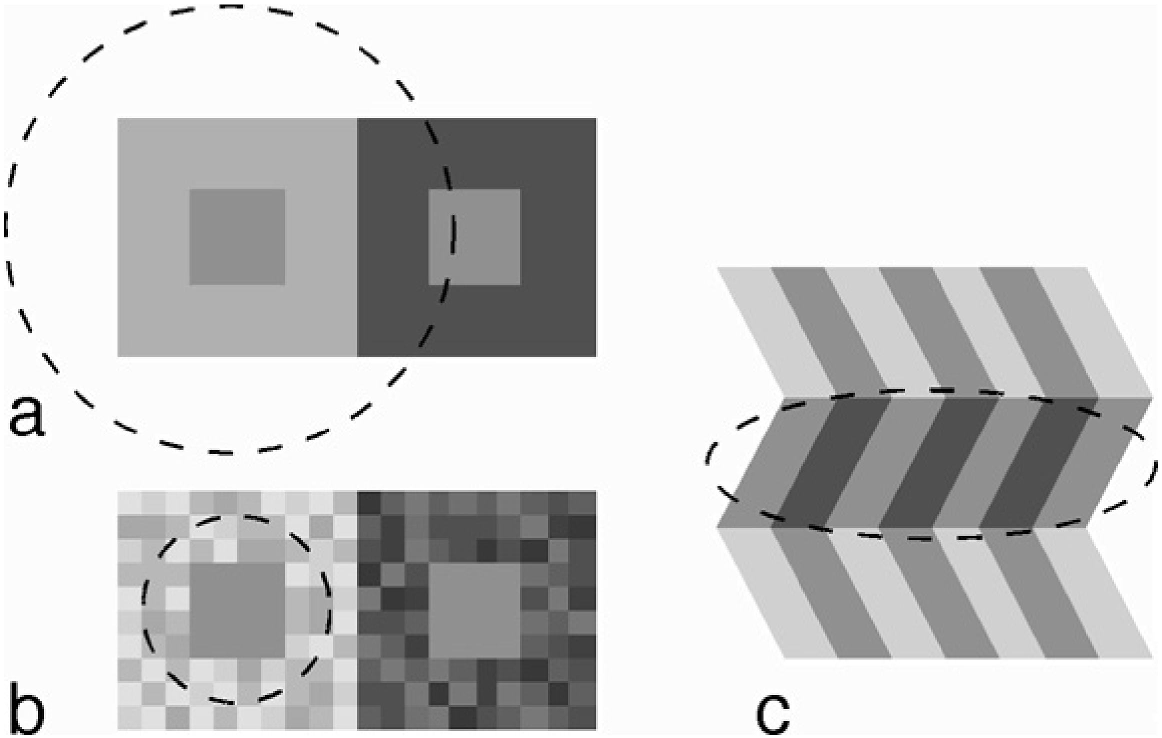

Lightness computations may depend on *articulation* & may employ adaptive windows. Simultaneous contrast is enhanced with *articulated* surrounds.

articulation = nombre de surfaces/patches distincts dans une région (telle qu'isolée par la fenêtre adaptable (plus d'articulation, plus de lightness constancy).

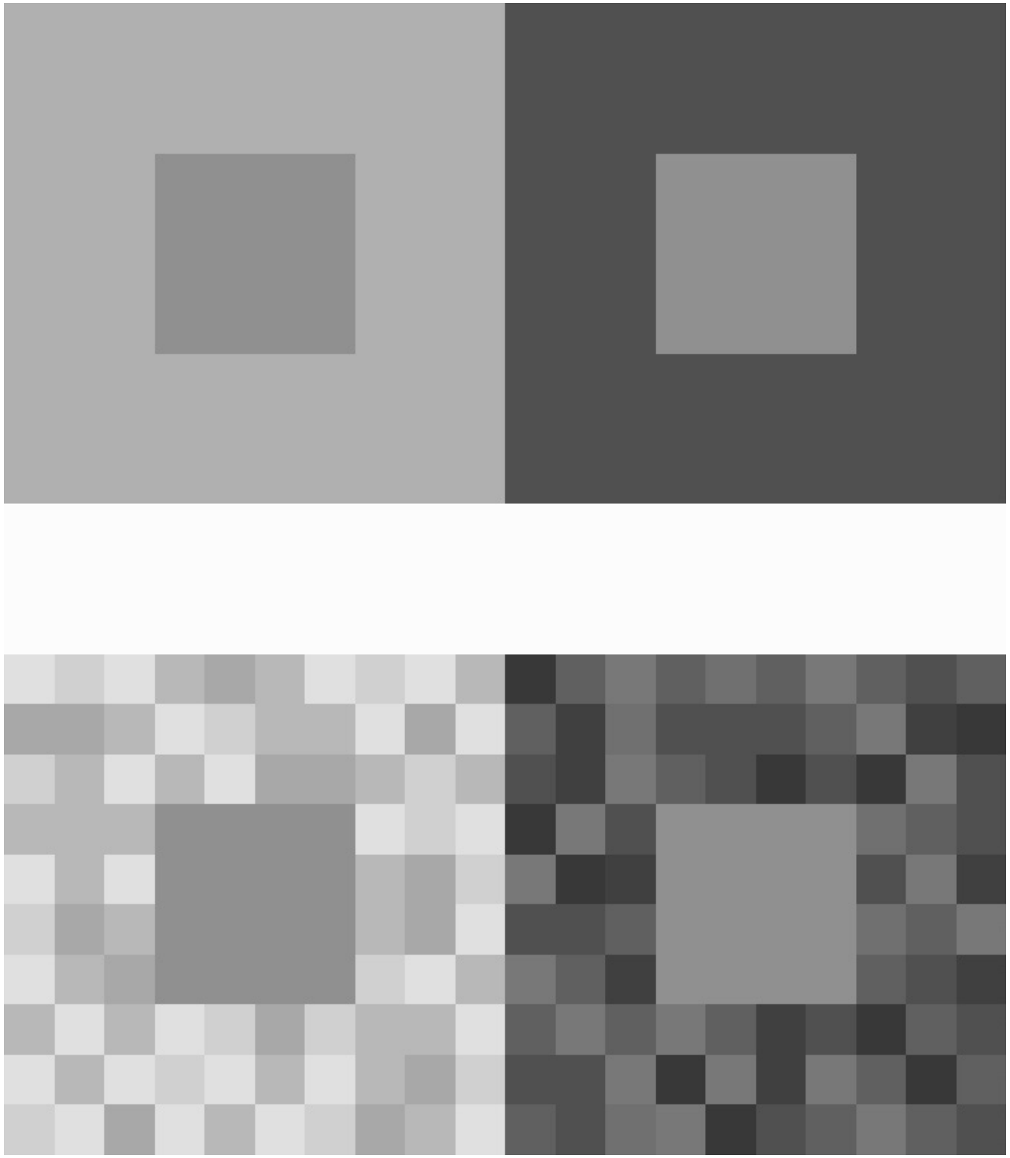

Simultaneous contrast is enhanced with *articulated* surrounds.

Atmospheres

Illuminance is only one of the factors determining the luminance corresponding to a given reflectance. Other factors: *interposed filters* (e.g., sunglasses), *scattering*, *glare* from a specular surface such as a wind-shield, etc.

Most physical effects of this kind lead to linear transforms (of luminance). Therefore the combined effects can be captured by a single linear transform (characterized by two parameters). This is what we call an atmosphere. The equation we use is,

$$L = m R + e,$$

where L and R are luminance and reflectance, m (≥ 0) is a multiplier on the reflectance, and e (≥ 0) is an additive source of light. The value of m is determined by the amount of light falling on the surface, as well as the proportion of light absorbed by the intervening media between the surface and the eye.

An ***atmosphere*** may be thought of as a single transparent layer, except that it allows a larger range of parameters. It can be amplifying rather than attenuating, and it can have an arbitrarily large additive component. (Putting on sunglasses or dimming the lights has the same effect on the luminances, and so leads to the same effect on atmosphere.)

The ***Atmospheric Transfer Function* (*ATF*)**

the mapping between reflectance and luminance

- Objective -

The ***Lightness Transfer Function* (*LTF*)**

the mapping between luminance and reflectance

- Subjective -

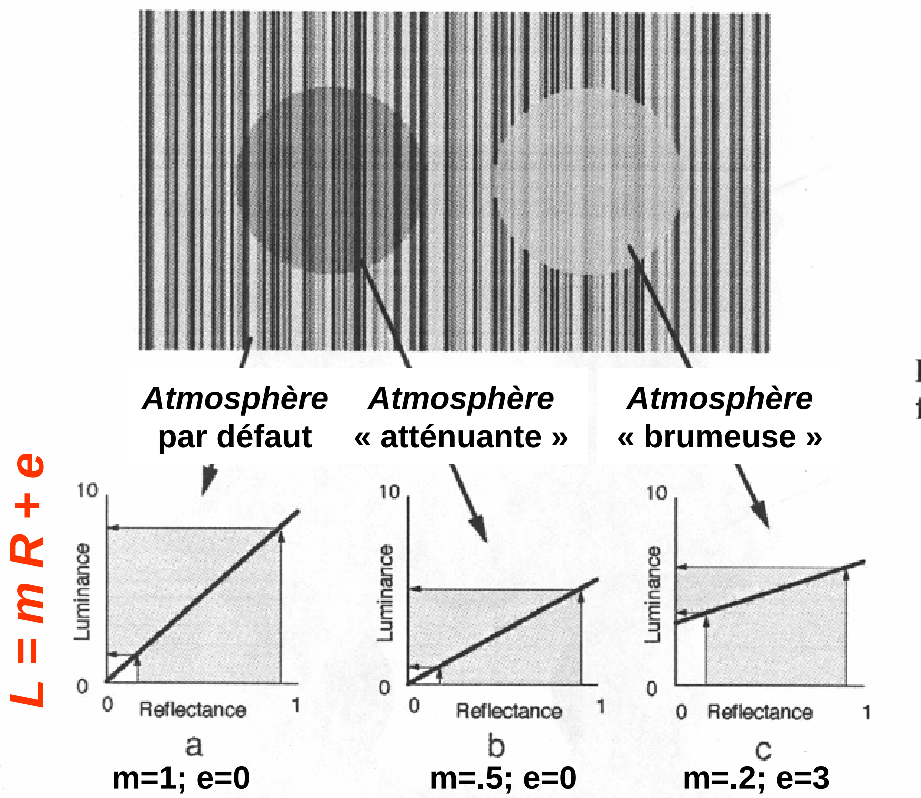

FIGURE 24.14 Lines of random gray, viewed under three different atmospheres. The ATF's, shown below, determine the mapping from reflectance to luminance.

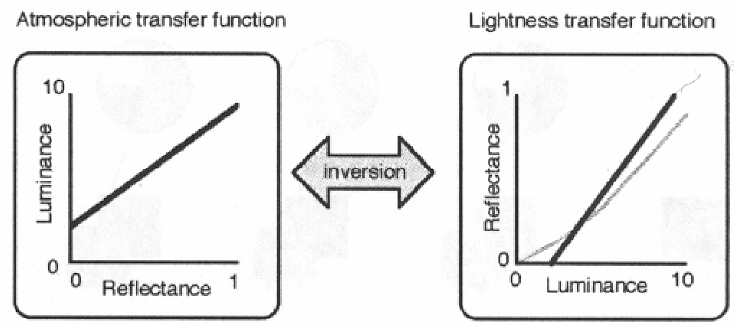

FIGURE 24.15 The inverse relation betweent he atmospheric transfer function and the ideal lightness transfer fuction.

Since the atmosphere maps a reflectance to a luminance, the observer must implicitly reverse the mapping, turning a luminance into a perceived reflectance, as illustrated in figure 24.15. The inverting function, for a given observer in a given condition, may be called the lightness transfer function, or LTF. The **LTF is subjective**; it needs not be linear and needs not be the correct inverse of the ATF. For a given observer it must be determined empirically.

Atmosphères et jonctions X

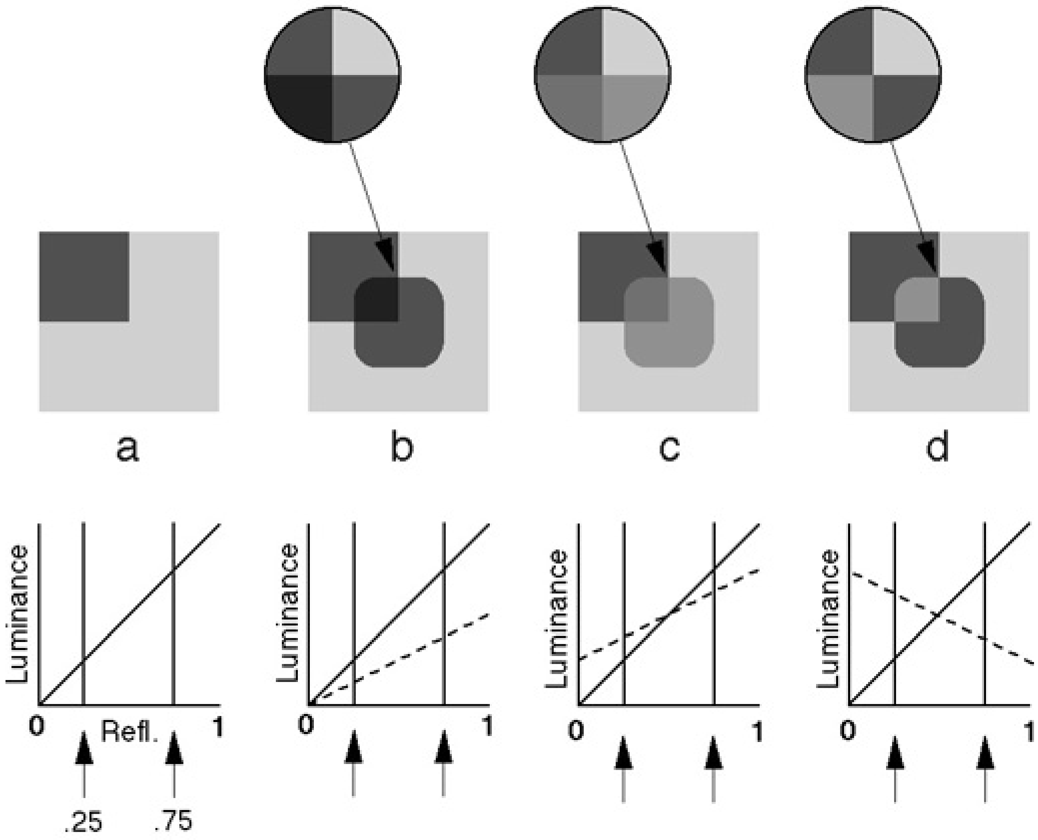

Des types différents d'atmosphère conduisent à des catégories différentes de jonctions en X:

a. *Atmosphère par défaut* (pas de jonction X); bords de réflectivité;
b. *Atmosphère atténuante* (filtre, ombre); jonction X avec le signe préservé;
c. *Atmosphère « brumeuse »* (fumée, fenêtre sale); *jonction X avec un signe inversé* (processus additif + atténuation) conduisant à une frontière entre atmosphères;
d. Atmosphère conduisant à une *jonction X avec une double inversion de signe*; de telles jonctions signalent un bord de réflectivité et non pas d'atmosphères

Transparency involves the imposition of a new atmosphere. The resulting X-junction category depends on the *atmospheric transfer function*.

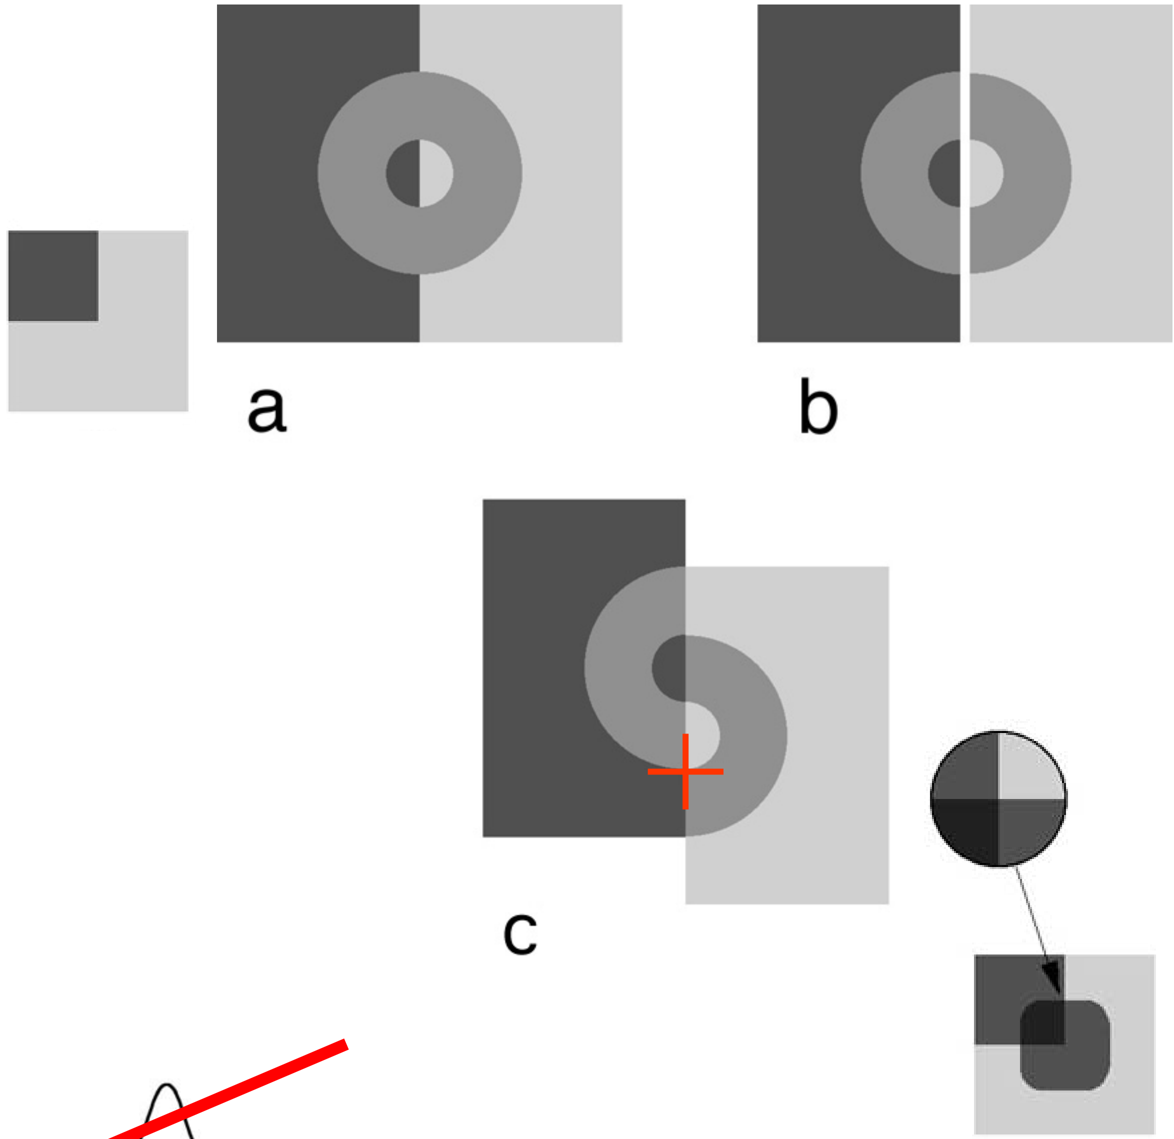
a
b
c

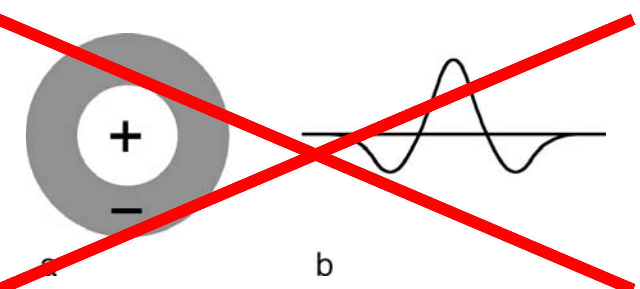
+
−
b

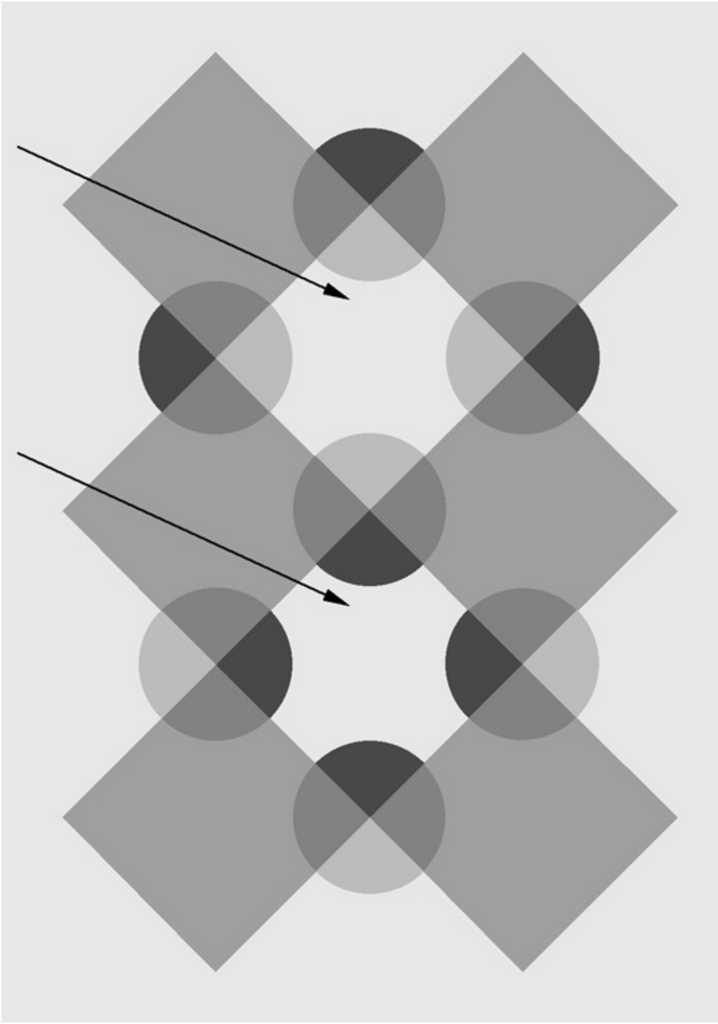

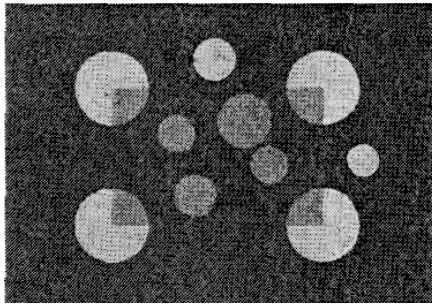

An illusion of haze. The two marked regions are identical shades of gray. One appears clear and the other appears hazy.

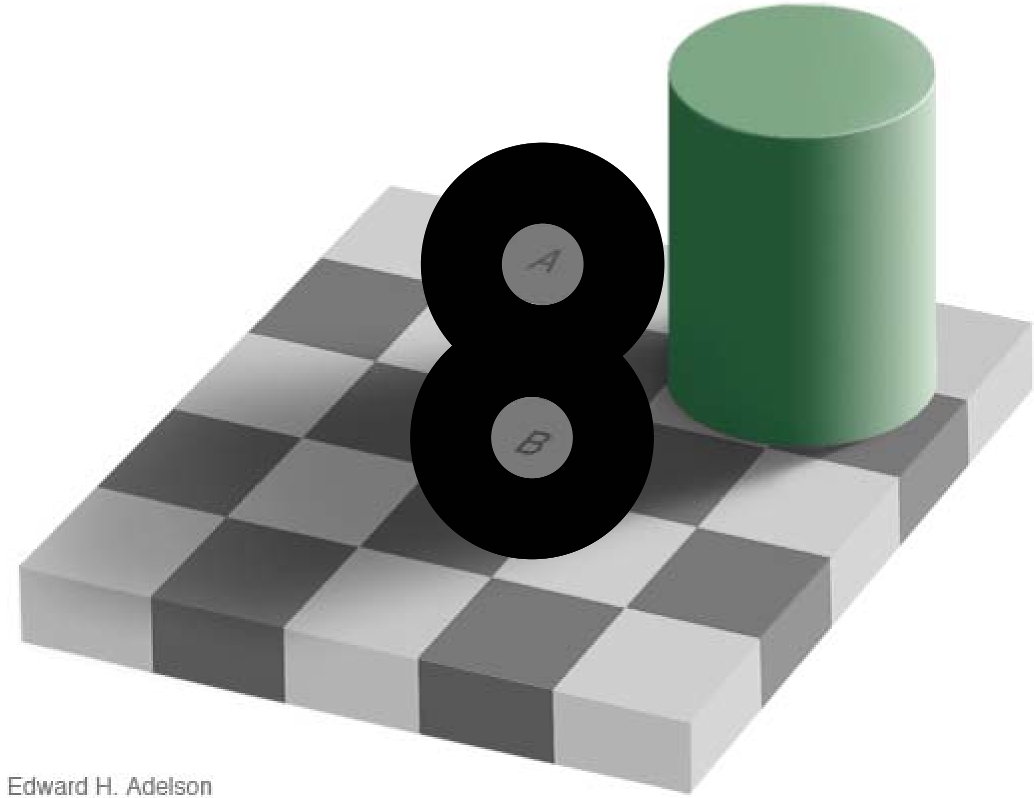

Les carrés A et B ont strictement la même luminance ; ils apparaissent pourtant comme étant très différents du fait de leur contexte, à savoir l'ombre du cylindre.

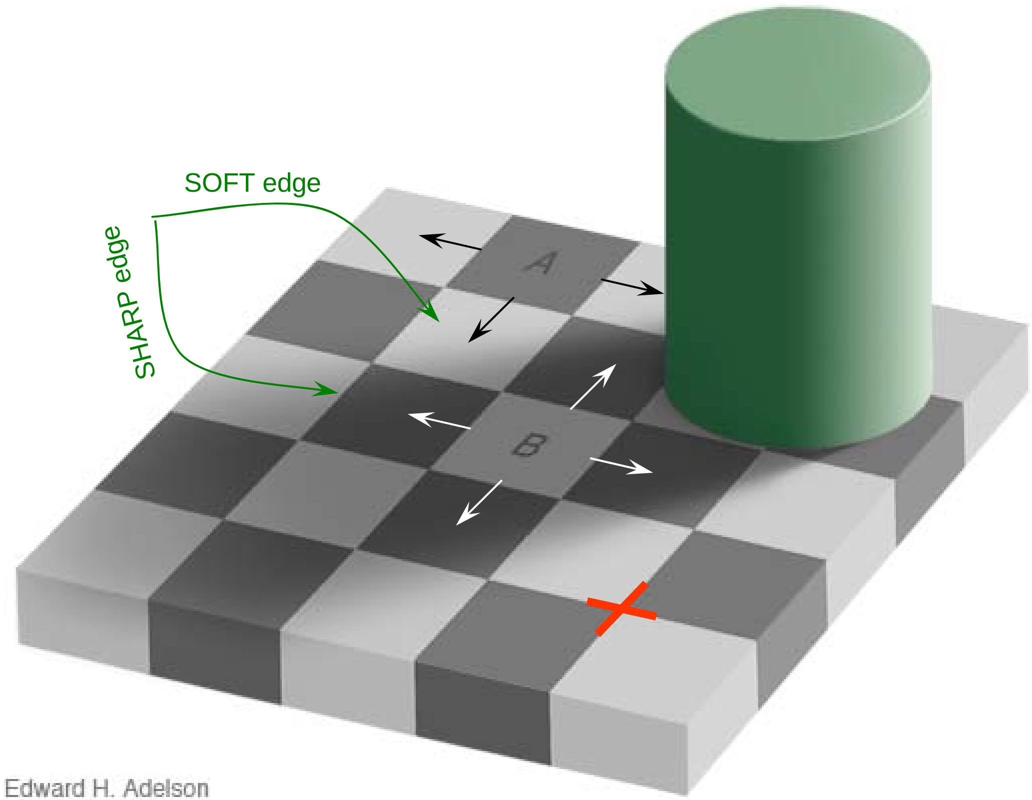

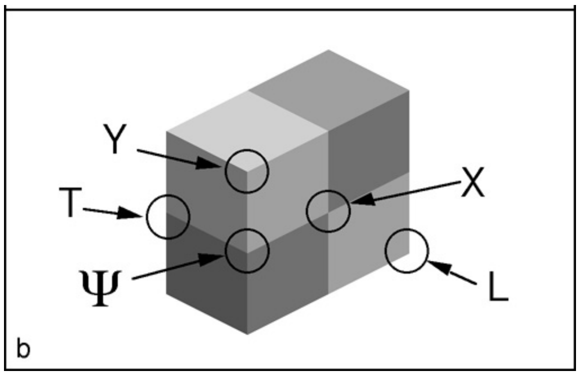

The visual system needs to determine the color of objects in the world. In this case the problem is to determine the gray shade of the checks on the floor. Just measuring the light coming from a surface (the luminance) is not enough: a cast shadow will dim a surface, so that a white surface in shadow may be reflecting less light than a black surface in full light. The visual system uses several tricks to determine where the shadows are and how to compensate for them, in order to determine the shade of gray "paint" that belongs to the surface.

The first trick is based on local contrast. In shadow or not, a check that is lighter than its neighboring checks is probably lighter than average, and vice versa. In the figure, the light check in shadow is surrounded by darker checks. Thus, even though the check is physically dark, it is light when compared to its neighbors. The dark checks outside the shadow, conversely, are surrounded by lighter checks, so they look dark by comparison.

A second trick is based on the fact that shadows often have soft edges, while paint boundaries (like the checks) often have sharp edges. The visual system tends to ignore gradual changes in light level, so that it can determine the color of the surfaces without being misled by shadows. In this figure, the shadow looks like a shadow, both because it is fuzzy and because the shadow casting object is visible.

The "paintness" of the checks is aided by the form of the "X-junctions" formed by 4 abutting checks. This type of junction is usually a signal that all the edges should be interpreted as changes in surface color rather than in terms of shadows or lighting.

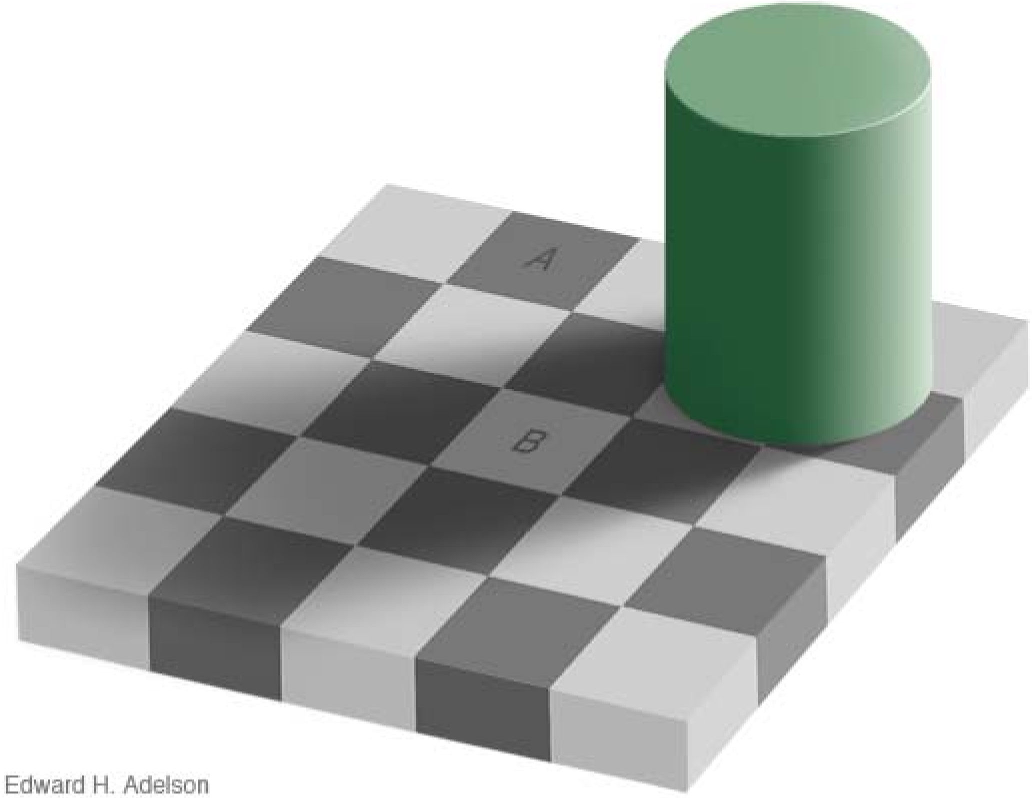

As with many so-called illusions, this effect really demonstrates the success rather than the failure of the visual system. The visual system is not very good at being a physical light meter, but that is not its purpose. The important task is to break the image information down into meaningful components, and thereby perceive the nature of the objects in view.

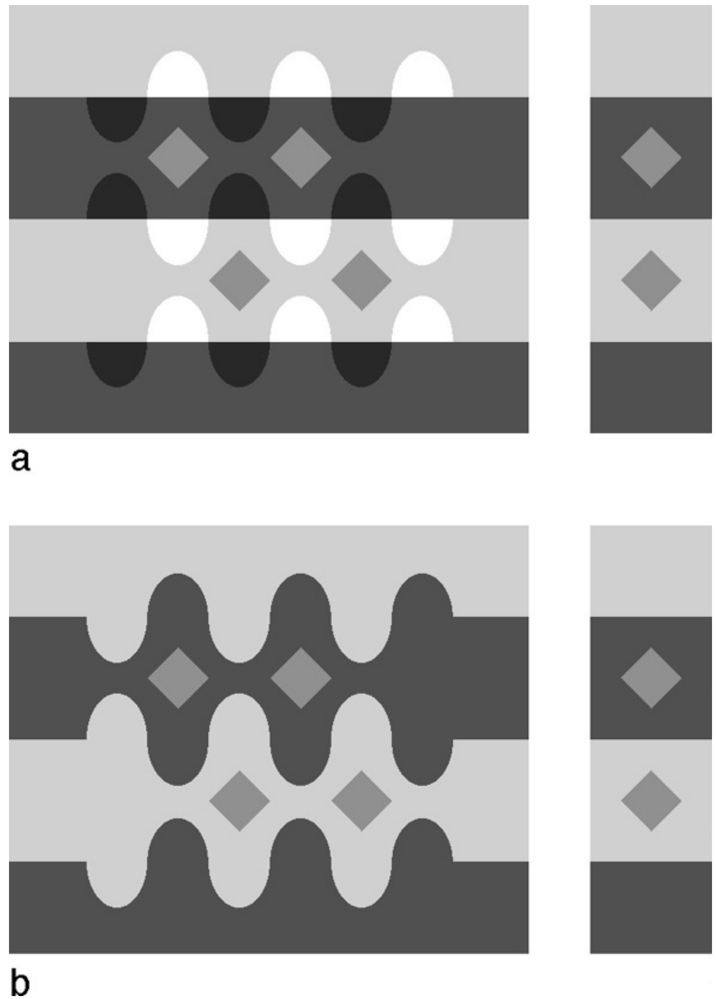

The snake illusion. All diamonds are the same shade of gray. (a) The regular snake: the diamonds appear quite different. (b) The "anti-snake:" the diamonds appear nearly the same. The local contrast relations between diamonds and surrounds are the same in both (a) and (b).

Motion from Shadow

Kersten, D., Knill, D. C., Mamassian, P. & Bülthoff, I. (1996). Illusory motion from shadows. *Nature*, *379*, 31.

Ball in a Box

Kersten, D., Mamassian, P. & Knill, D. C. (1997). Moving cast shadows induce apparent motion in depth. *Perception*, *26*, 171-192.

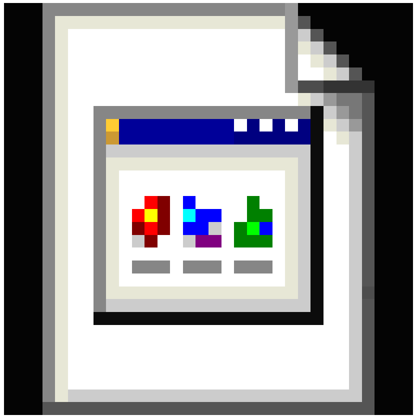

Cafewall.exe

The original paper by Gregory and Heard provides a thorough treatment of the way in which this illusion arises. It is based primarily on a concept called *border locking* that involves edge detection in the context of simultaneous spatial and colour registration in the human visual system.

You may also be interested in a somewhat simpler (and therefore modestly less accurate) explanation. If you look at the boundary between two dark tiles (in the default configuration), the mortar line is plainly evident. At the boundary between two light tiles it can also be seen clearly. At the boundary between a light and dark tile, however, your visual acuity simply isn't sharp enough to resolve the mortar line as a separate object. Nevertheless, it still occupies some space on the screen and your brain must somehow interpret that ``missing'' space. It therefore simply interprets the mortar as part of the tile above or below it (depending on which one is nearest the center of your field of view). When you look at a single tile, then, it appears taller at one end than the other by twice the width of a mortar line, giving it that characteristic wedge shape.

But this is only half of the story. If all of the tiles looked like wedges, then the boundary between them should appear jagged. Your brain, however, is looking for the simplest explanation that fits the evidence that is presented to it. In this case, the evidence supports the theory that the rows of tiles are separated by simple straight lines (which is reasonable because this is in fact true). The best compromise between the incompatible notions of straight lines forming the boundary between a succession of wedges is the interpretation that the lines are in fact straight, but neither horizontal nor parallel. When the tiles are made quite small, more evidence is available to refute this theory, so the appearance is qualitatively changed. Depending upon where you focus your attention, both conflicting perceptions (straight lines and rows of wedges) can be seen independently.

In light of this explanation it should be clear why the various controls alter the nature or strength of the illusion as they do.

Motion Induced Blindness.exe