

When she becomes dominant her body changes so that she becomes physiologically and morphologically distinct from other non-reproductive female ‘workers’, even her bone structure changes. Mole-rats were thought to be the only mammal to do this, but it has since been discovered that female meerkats also develop elongated vertebra when they become dominant breeders. As well as reproductives, there are also physically and behaviourally distinct dispersive morphs, only known since 1996. Dispersers are big, fat males that try to mate with strange animals from other colonies instead of attack them, as most mole-rats would do. These dispersers are very rare, however, and have never been seen above ground.

So how old is old? Another extraordinary aspect of mole-rat life is that they live to be absolutely ancient relative to their body size. Another rodent of similar size might expect to live for two years; mole-rats have been reported to live for 30 years. They have become of interest to science because of their longevity, as well as their fascinating social behaviour. Because of their extraordinary longevity, scientists expected to find reduced levels of cell-damage and higher levels of anti-oxidant activity in mole-rats, but actually their cell-damage is comparable with that found in other, shorter-lived species. How naked mole-rats survive this cell damage is, as yet, a mystery. It seems we still have a lot to learn from the naked mole-rat.

Where can I find out more?

- Buffenstein, R. (2008). Negligible senescence in the longest living rodent, the naked mole-rat: insights from a successfully aging species. *J. Comp. Physiol. B – Bioch. Syst. Environ. Physiol.* 178, 439–445.
- O’Riain, M.J., Jarvis, J.U.M., and Faulkes, C.G. (1996). A dispersive morph in the naked mole-rat. *Nature* 380, 619–621.
- Jarvis, J.U.M., O’Riain, M.J., Bennett, N.C., and Sherman, P.W. (1994). Mammalian eusociality – a family affair. *Trends Ecol. Evol.* 9, 47–51.
- Reeve, H.K. (1992). Queen activation of lazy workers in colonies of the eusocial naked mole-rat. *Nature* 358, 147–149.
- Sherman, P.W., Jarvis, J.U.M., and Alexander, R.D. (1991). *The Biology of the Naked Mole-Rat*. (Princeton University Press).

Institute of Evolutionary Biology, University of Edinburgh, West Mains Road, Edinburgh EH9 3JT, UK.
E-mail: a.griffin@ed.ac.uk

Primer

Stereopsis

Carlos R. Ponce^{1,2}
and Richard T. Born^{1,*}

“When I began to see with two eyes, my visual world completely transformed. Trees looked totally different. Consider a leafless tree in winter. Its outer branches enclose and capture a volume of space through which the inner branches permeate. I had no concept of this. Oh, I could infer which branches were in front of others using monocular cues such as object occlusion, but I could not perceive this. Trees to me looked somewhat like a drawing. Before my vision changed I would not have said that the tree looked flat, but I had no idea just how round a tree’s canopy really is... When I began to see with two eyes, everything looked crisper and much better outlined. As another formerly stereoblind person wrote to me, ‘Everything has edges!’ The world was not only flatter but less detailed and textured with my monocular vision.”

The above quotation, from an email contribution by Susan Barry to an on-line discussion of the benefits of stereopsis, eloquently captures many of the perceptual, aesthetic, and even philosophical aspects of ‘stereopsis’, a word derived from Greek that translates literally as: $\sigma\tau\epsilon\rho\epsilon\omicron\varsigma$, ‘solid’ and $\omicron\mu\omicron\varsigma$, ‘power of sight’. Because Susan Barry acquired this capacity as an adult, she was made vividly aware of what most of us take for granted, namely, a direct sense of the three-dimensionality of the visual world. The quotation also hints at some other critical features of binocular vision that are not related to depth perception *per se*. For example, the statement that “Everything has edges!” attests to the fact that seeing the world twice, as it were, from two slightly different perspectives, makes a number of important visual computations, such as the detection of edges, more robust.

In this primer, we shall mainly deal with stereopsis as narrowly defined, that is, as the use of *differences* in the images projected onto the retinas of the two eyes — so-called ‘binocular disparity’ — to reconstruct the third

visual dimension of depth. We will first describe the basic geometry of binocular disparity before discussing how the brain uses this information to compute depth. We would be remiss, however, if we did not clarify two vital distinctions that place this peculiar capacity within the larger context of vision. The first is that binocular comparisons also serve many other important functions in vision, apart from the computation of depth. The second is that there are many other cues to visual depth that do not require binocular comparisons. Thus, after discussing stereopsis proper, we will briefly treat these two issues.

Basic geometry

Because the two eyes are separated horizontally, they see the same visual scene from two slightly different vantage points. When we ‘look at’ a particular feature in space, such as the black dot on the arrow in [Figure 1A](#), what we are doing is aiming the fovea — the tiny retinal region of highest visual acuity — of each eye at that feature. This act defines the ‘plane of fixation’. If we start at this point on the arrow and move a fixed distance along the gray circle, the projections of the new point in the two eyes will move by *exactly the same distance* on each retina. These points of projection on the two retinas are defined as ‘corresponding points’. The geometric horopter is the collection of all such points in the image that project to corresponding points on the retina. (To an observer these points would all appear to lie at roughly the same depth as the point of fixation. The collection of points that appear at *exactly* the same depth is referred to as the “empirical horopter” and is slightly flatter than the geometric horopter.)

Now consider parts of an object that lie either in front of or behind the plane of fixation, such as the head and tail, respectively, of the arrow in [figure 1B](#). These features will not project to corresponding points on the two retinas. That is, the retinal distance from f to h on the right retina will not be equal to the distance from f' to h' on the left retina. This difference $[(h - f) - (h' - f')]$ is called ‘binocular disparity’, and it is the basis for stereopsis. The projection lines for retinal points of *near* features cross in front of the horopter and thus produce what are referred to as *crossed* disparities (by convention,

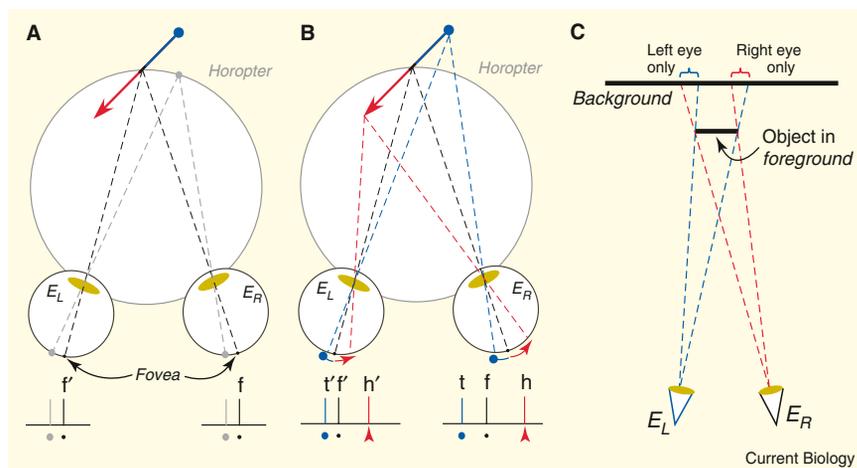


Figure 1. The geometry of stereopsis.

Each diagram represents a section through the horizontal equator of the eye as viewed from above. (A) Points along the horopter produce images on corresponding points of the two retinas. (B) Points on the arrow at different distances from the observer produce images at different distance from the fovea on the two retinas. This difference, called 'binocular disparity', is the basis for the classical variety of depth perception referred to as 'Wheatstone stereopsis'. (C) The basis of 'da Vinci stereopsis'. The presence of an object in the foreground creates regions of the background that can be seen by one eye only. These unpaired image points provide potent cues for depth and occlusion. (A,B) Adapted from Tyler 2004; (C) adapted from Nakayama and Shimojo 1990.

these are assigned *negative* values), and, conversely, *far* features, such as the tail of the arrow in Figure 1B, produce *uncrossed* disparities, which are assigned *positive* values.

If the retinal disparities produced by a feature fall within a limited range (about 0.5° in the vicinity of the fovea; slightly more at greater eccentricity), it will be seen as a single 'fused' feature that is offset in depth. Beyond this range, which is known as 'Panum's fusional area', the feature will appear double — the observer is aware of each retinal image independently, commonly known as 'double vision' — even though the visual system can still extract meaningful depth information up to binocular disparities of several degrees, depending on the size of the feature.

Note that the above discussion takes for granted that one knows which feature in the right eye to match with the same feature in the left eye. In fact, this can be a difficult thing to figure out, and it is known in computational circles as the "correspondence problem." The problem is quite formidable if one is trying to match single pixels in the two retinal images (because there are a huge number of possible matches), but becomes much more manageable if one uses local features — such as short line segments (see Figure 3) — to

do the matching. In addition, it is likely that higher order image properties, such as the global shape of an object, are used to constrain the local feature matches.

There is another type of stereopsis known as 'da Vinci stereopsis' (after Leonardo da Vinci who first described it), which involves the use of unpaired regions of the images in the two eyes that occur as the result of occlusion. Basically, because the eye cannot see around objects, a foreground object in the shape of a square will create a vertical strip of the background on the left side that is visible to the left eye only, and a similar strip on the right side that is visible to the right eye only (Figure 1C). This type of stereopsis is also useful for depth perception, in particular, for defining occlusive relationships, but will not be discussed further here. To distinguish the two, the more common form of stereopsis, treated here, is sometimes referred to as 'Wheatstone stereopsis' after its discoverer Charles Wheatstone.

Spatial range of stereopsis

Because the basic data that the brain has to work with are binocular disparities, the smallest depth difference that can be distinguished is ultimately set by the minimum retinal disparity that can be resolved. In human observers, this difference

approaches 5 arc-seconds (0.0014 degrees) at the fovea — a truly remarkable precision given that the spacing between cones in the fovea is on the order of 30 arc-seconds.

The binocular disparity produced by a given depth difference varies as the inverse square of the viewing distance, which means that, at close viewing distances, exceedingly small depth differences can be discriminated — as tiny as $25\ \mu\text{m}$ (about the width of a fine human hair). Even at relatively large viewing distances, stereopsis can play a role provided the depth dimension is large enough. For example, at a distance of 100 meters, the minimum resolvable depth is about 4 meters.

Physiological basis of stereopsis

The raw data for stereopsis are binocular disparities, but in order to actually calculate the depth of an image feature, the signals from the two retinas must be compared by a neural circuit. This occurs in the visual cortex. In order to understand how the brain computes stereoscopic depth, we must first review some general properties of neurons in the visual pathways. A given neuron generates electrical impulses (called 'action potentials' or 'spikes') in response to visual stimuli that occur within a small region of the visual field. This region is the neuron's *receptive field*. The output cells of each eye, the retinal ganglion cells, project to the lateral geniculate nucleus (LGN) of the thalamus, which, in turn, projects to the primary visual cortex (V1). It is in V1 where single neurons that receive inputs from both eyes are first observed — a necessary, but not sufficient, condition for stereopsis.

The key property for such a neuron is that it is sensitive to *specific differences* between the receptive fields in the two eyes. This can be tested for a given neuron by presenting the same visual stimulus *independently* to each eye while systematically varying the position of the stimulus within each receptive field. Visually-evoked spikes are counted and plotted as a function of the difference in image position between the two eyes, that is, the binocular disparity. Many neurons in V1 are remarkably sensitive to the binocular disparity, showing high firing rates for some disparities and lower rates for others. This type of firing rate modulation, shown graphically as a 'disparity tuning

curve', is the signature of binocular disparity-selective neurons, which comprise the biological substrate for stereopsis. Disparity tuning curves occur in many shapes, which have been traditionally categorized by the location of the peak firing response (Figure 2).

What is the underlying mechanism behind this diversity of disparity tuning functions? While this question is still being actively explored, several key findings have contributed to an elegant model for the computation of binocular disparity. The key mechanisms behind this model depend on the internal shape of each neuron's interocular receptive field. Many cells in the recipient layers of the primary visual cortex have receptive fields which can be mapped out using small spots of light. These so-called *simple cells* will fire action potentials in response to a small spot of light appearing in some parts of their receptive fields (ON regions), but are inhibited when the same spot appears in other regions of the receptive field and only respond upon stimulus disappearance (OFF regions). This method of receptive field mapping, pioneered in the visual cortex by David Hubel and Torsten Wiesel, has revealed that most simple cells have receptive fields consisting of elongated regions of ON excitation flanked by elongated regions of OFF inhibition (Figure 3A; or the inverse, OFF regions flanked by ON). These receptive field shapes are well described by a simple mathematical function (Figure 3B), which is useful for the kinds of models discussed below. Neurons with such receptive fields are optimized for the detection of lines or edges of a particular orientation.

The specificity of stimulus position and orientation required to elicit responses from V1 simple cells suggests that binocular disparity could be coded by introducing small differences between the left and right receptive fields. That is, one could simply introduce a spatial offset between the retinal positions of each receptive field to capture binocular disparity cues, so-called 'position disparity'. Alternatively, both receptive fields could share the exact same retinal position but contain differences in the relative positions of the excitatory and inhibitory regions within each receptive field, known as 'shape disparity'.

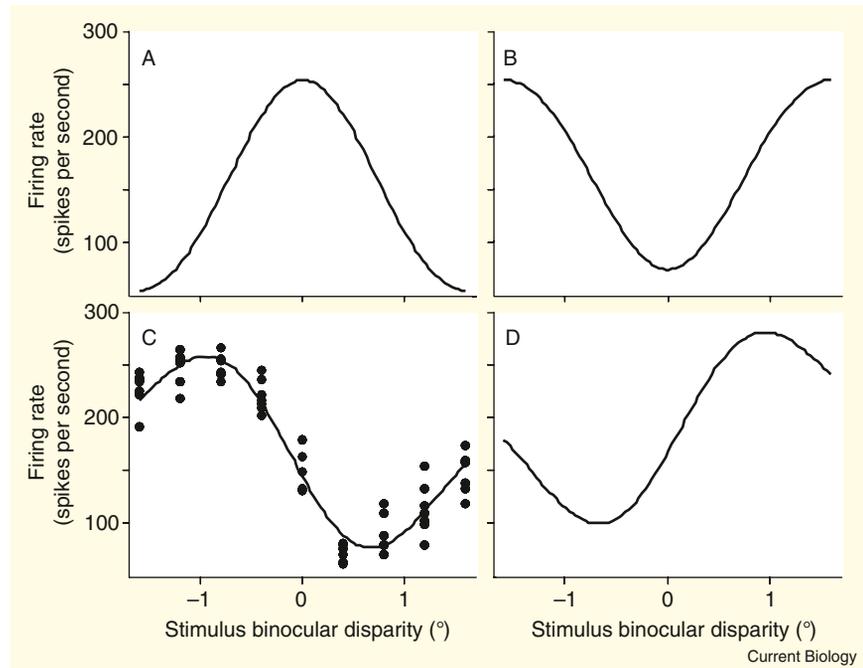


Figure 2. Binocular disparity tuning functions.

Spike rates evoked by visual stimuli of differing binocular disparity. For the 'near' cell (C), actual data are shown — each filled circle represents the spike rate evoked during a single stimulus presentation. Solid lines show the mathematical function (called a 'Gabor function') that best fit the data. The other panels represent idealized versions of their respective cell types. (A) The most common type of disparity tuning curve found in V1 peaks at binocular disparity values near zero, corresponding to a visual stimulus on the horopter, and such neurons are termed 'tuned excitatory'. The hallmark of these disparity tuning functions is that they are symmetric around zero (responding equally well to very small near or far disparities). (B) Another type of symmetric curve belongs to 'tuned inhibitory' cells, which are suppressed by a small range of disparities around zero and thus appear to be an inverted version of the tuned excitatory cells. (C,D) The third and fourth types of disparity tuning functions have their response peaks at either near or far (crossed or uncrossed, respectively) disparities. These tuning curves are generally asymmetric about zero. While it is convenient to think in terms of these four different classes, in reality one finds a continuum of different tuning curves spanning the gamut from near-, to zero-, to far-tuned neurons.

The majority of orientation-selective cells in the primary visual cortex have more complicated functional properties, however. While simple cells require that a given stimulus be placed in a specific excitatory region within its receptive field, most neurons in V1 respond optimally to oriented lines appearing anywhere within their receptive fields. This position invariance is a hallmark of 'complex cells', and it is believed to arise by the combination inputs from multiple simple cells. Not surprisingly, the disparity tuning properties of complex cells are also more nuanced: complex cells can respond to their preferred binocular disparities no matter where they may be presented within their receptive fields.

A more elaborate 'energy model' has been developed to explain the responses of complex cells. This

model describes disparity-sensitive complex cells as receiving inputs from multiple simple cells, each having the same binocular disparity and orientation preference but differing in the relative position of the peak excitatory domain within each receptive field (Figure 3). In the simplest case, four simple cells with the same elongation axis and phase disparities would have their spatial frequency component shifted by one quarter of a cycle. In reality, many more simple cells — on the order of several hundred — are likely to serve as inputs to complex cells.

The disparity energy model accounts quite well for the observed responses of disparity-sensitive cells. In turn, the diversity of tuning curves emerging from this mechanism provides a neural toolkit for stereopsis. How these

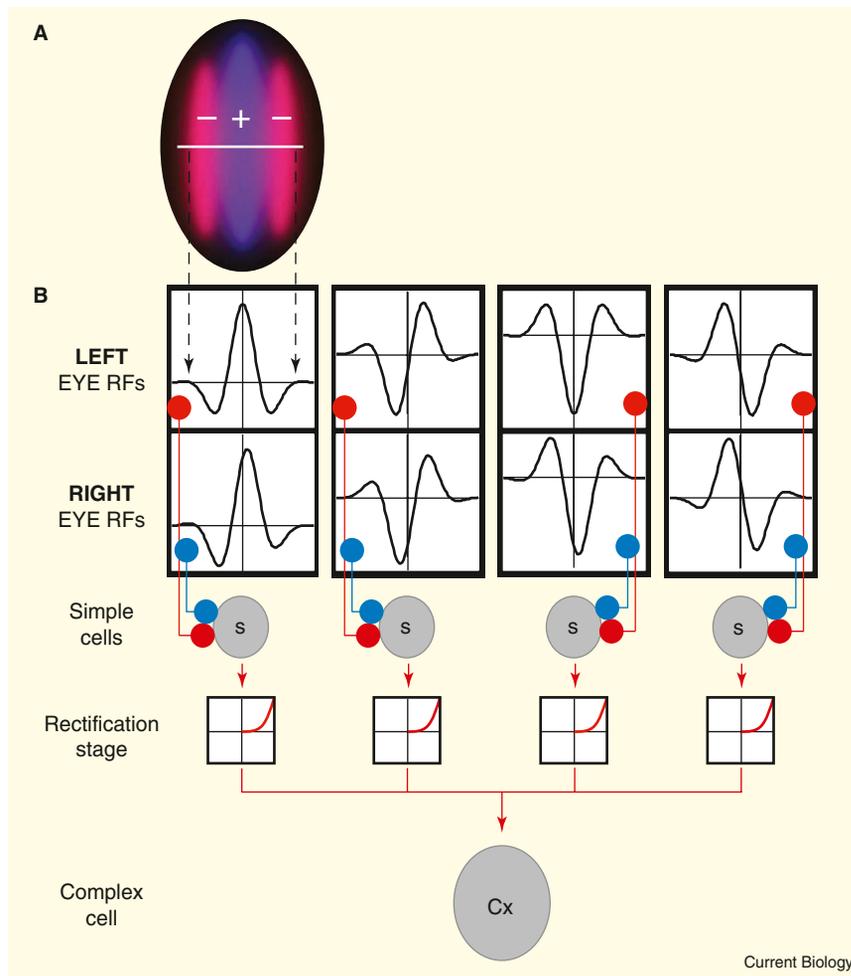


Figure 3. Encoding of binocular disparity by cells in the primary visual cortex.

(A) Typical structure of a V1 simple cell receptive field. The blue region shows an area of ON responses (+); magenta indicates regions eliciting OFF responses (-). The white line shows the level of a cross-section of this two-dimensional map shown below. (B) Complex cells respond to orientation and binocular disparity signals everywhere in their receptive fields. The energy model proposes that each complex cell (Cx) receives inputs from a number of simple cells (s), each having the same orientation and disparity preference (left and right receptive field profiles shown in black squares), but differing in the position of the excitatory domains within their own receptive fields. To make sure that the outputs of the different simple cells do not cancel each other, a rectification stage (to make all responses excitatory) is added between the simple complex cell stages. Adapted from Cumming and DeAngelis (2001).

disparity-sensitive cells contribute to the actual *perception* of depth has been explored using a variety of techniques, including the analysis of correlation between single-neuron activity and perceptual decisions and electrical activation of disparity-sensitive cell clusters.

The most informative approaches have used monkeys trained to report the depth of a visual stimulus (Figure 4A). Investigators then record the activity of disparity-selective neurons while the animal is performing the task. Using this approach, it has been shown that single neurons are not only very

sensitive to the depth signals in the visual stimulus, but they can be highly predictive of the choices that the animal will make on a given trial. This predictive power, or 'choice probability', constitutes strong evidence that the signals from these neurons are actually being used by the monkey's brain to solve the perceptual task. Further evidence for a causal role in perception is that electrically activating small clusters of similarly tuned binocular disparity-selective cells can directly bias a monkey's decisions in favor of the preferred disparity represented by those neurons (Figure 4B,C).

Other benefits of having two eyes
Stereopsis is certainly a good reason for having two eyes, but it is not the only one. Some of the other benefits of binocularity have nothing to do with making binocular comparisons. For example, given how important vision is for many animals, it is clearly useful to have a spare eye in the event that one becomes damaged. Another potential advantage is that, if the two eyes are placed on opposite sides of the head, the combined fields of view can effectively cover the entire 360° of the visual field. This ability to 'have eyes in the back of one's head' might be particularly useful if one is being stalked and, indeed, is frequently found in prey, such as gazelles. Of course, this arrangement precludes stereopsis, or at least limits it to a very narrow strip where there is overlap of the monocular visual fields.

This trade-off appears to play out in the animal kingdom as strategies adopted by the *hunter* — forward facing eyes with a large degree of binocular overlap and good stereopsis — versus the *hunted* — lateral eyes for panoramic vision. The hunter with mobile eyes can also make use of the angular displacement required to aim the two foveae on an object in order to compute its absolute distance. This range-finding ability would obviously be useful during encounters in which the predator must pounce on or strike at its prey. Another, often overlooked, binocular advantage is that combining the outputs of n independent detectors decreases noise by a factor of \sqrt{n} , so using two eyes improves the visual system's signal-to-noise ratio 1.4-fold. At very low light levels, such as those experienced by our nocturnal mammalian ancestors, this might have proven critical for survival.

Finally, there is a rather large category of visual computations that make use of binocular disparity information, but that are not used for determining depth *per se*. Perhaps the most important of these is 'breaking camouflage'. Imagine an animal with a coat that is perfectly matched in color and texture to that of its surroundings (the visual background against which it is viewed). As long as this animal remains stationary, it would be invisible to a single eye, but a mechanism exquisitely sensitive to binocular disparity would cause the animal's outline to 'pop

out' from the background, because the animal and the background are at different depths. This use of disparity processing for figure-ground segregation was vividly demonstrated by the random dot stereograms created by Bela Julesz and can be seen in the once popular 'Magic Eye' images (Figure 5A). In these cases, the figure is visible *only* when viewed stereoscopically, so it is not so much a matter of perceiving the depth metrically but rather of grouping the relevant regions of the image so that a shape can be recognized.

A related role is the determination of border ownership, which has profound effects on the linking of features that may be partially occluded. In this case, binocular comparisons aid in distinguishing real object boundaries from those that are created by occlusion. This effect was dramatically demonstrated by Ken Nakayama and colleagues, and can be appreciated even in a non-stereo version of their task: The face in Figure 5B is much more difficult to recognize when it appears in front of the occluding strips rather than behind them. A host of visual stimuli are perceived differently depending on the relative depth they are assigned. For example, the same small bright patch on a curved surface is interpreted differently depending on its relative binocular disparity: as a reflection of the light source from a shiny surface if it appears behind the surface or as paint if it is at the same depth as the surface. In sum, binocular disparity processing is not *just* for depth perception, but rather permeates many aspects of the act of seeing.

Other cues for depth

A significant proportion of the human population — anywhere from 5% to 10% — lacks stereoscopic depth perception, just as Susan Barry did prior to her therapy. Given the discussion thus far, we might expect these people to be severely visually impaired, but, remarkably, they are not. Some, in fact, have managed to succeed at careers that would seem to place a high demand on good depth perception, such as surgery, professional skiing and dentistry. The reason this is possible is that there are a number of other methods by which depth information can be obtained. One of the most useful is to move the head in order to sample multiple views over time, producing 'motion parallax'.

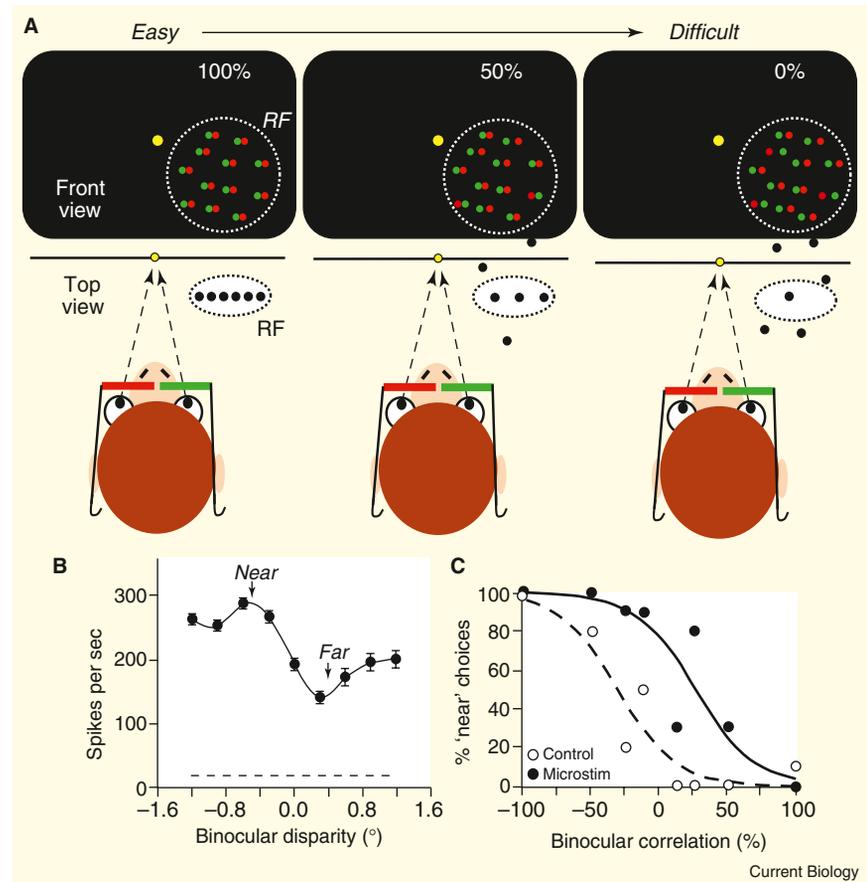


Figure 4. Microstimulation of disparity-tuned neurons can bias a monkey's perceptual decisions.

(A) Depth discrimination task. The monkey must report whether the dots appear in front of (near) or behind (far) the plane of fixation. The difficulty of the task is varied by adjusting the proportion of binocularly correlated dots. In the 100% condition (left), all of the dots appear at a single depth, making the task very easy. By adjusting the proportion (% binocular correlation) of dots that have random binocular disparities (noise dots), the task is made progressively more difficult. In the limit, all of the dots appear at random depths (0% binocular correlation) and there is no correct answer. In this case, the monkey must guess, and he is rewarded randomly on half of the trials. (B) In this microstimulation experiment, the pool of neurons being stimulated prefers near (negative) binocular disparities. Based on this tuning curve, two binocular disparities are chosen for the behavioral task (arrows): the near stimulus is a crossed binocular disparity of -0.4° and the far stimulus is an uncrossed binocular disparity of $+0.4^\circ$. While the animal is performing the task shown in (A), on half of the trials the neurons are electrically stimulated during the period when the animal is viewing the visual stimulus. (C) Psychometric functions. The x-axis shows the signal strength for near (negative) vs. far (positive) binocular disparities, and the y-axis is the percentage of trials on which the monkey reported a near stimulus. On stimulated trials, the animal is much more likely to report a near depth, as shown by the upward shift of the filled circles (trials with microstimulation) compared to the open circles (no microstimulation). Since the animal is only rewarded for making correct decisions, the most reasonable interpretation is that the stimulated neurons play a causal role in the decision process. It is as if stimulating neurons tuned to near binocular disparities causes the monkey to see more near dots. (A) Adapted from Uka and DeAngelis (2004); (B,C) adapted from DeAngelis *et al.* (1998).

When one moves one's head in this manner — a number of animal species and some monocular humans use this technique of 'head bobbing' — objects *nearer* than the point of fixation will move in a direction *opposite* to the direction of head motion and those *farther away* will move in the *same* direction. The speed of the motion

indicates the magnitude of the depth difference. When tested under similar conditions, depth discrimination from motion parallax can be almost as good as that from stereopsis, differing by only a factor of two or so.

In addition to motion parallax, there are a number of other 'pictorial' cues to depth. These are the cues that artists

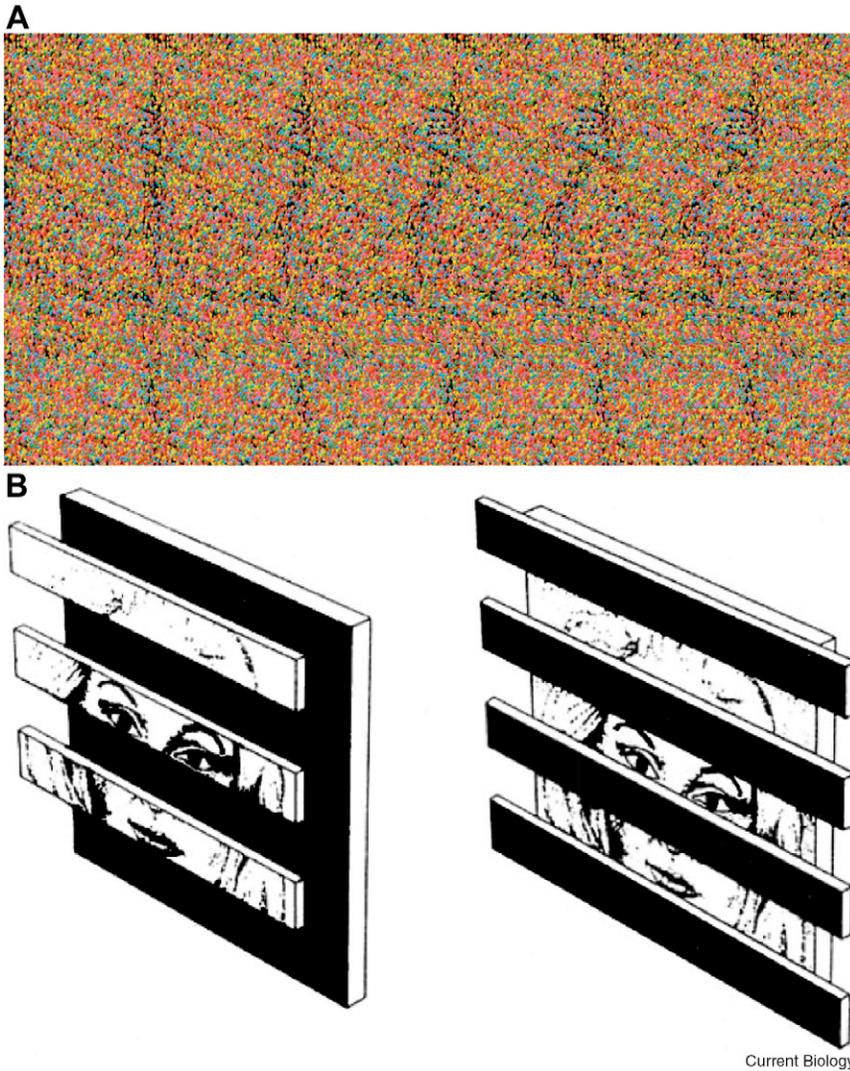


Figure 5. Binocular comparisons are not just for computing depth.

(A) An autostereogram. The figure of the shark is visible only when the image is binocularly ‘free fused’ to make the shape stand out in depth. In the absence of stereopsis, the shark is perfectly camouflaged. (B) The face is much more easily recognized when the parts appear to be behind the occluding strips. This is even more apparent when binocular disparity is used to define the relative depths of the face and the strips. (A) from http://en.wikipedia.org/wiki/Image:Stereogram_Tut_Random_Dot_Shark.png; (B) adapted from Nakayama, Shimojo and Silverman (1989).

exploit to create the illusion of depth on a flat canvas. They include linear perspective (the tendency for parallel lines to converge as they recede towards the horizon), size and texture gradients (objects shrink and textures become finer with distance), occlusion (an object that partially occludes another is perceived as being nearer), atmospheric effects (distant objects become partially blurred as the result of particulates in the air, such as water vapor or pollution), and luminance (brighter objects appear closer). While none of these cues can produce the fine depth discrimination achievable by stereopsis, they provide an overall

sense of the three-dimensional spatial structure of the environment, and, when combined with motion parallax, they can almost completely compensate for a lack of stereopsis. This appears to be especially true in individuals who have lacked stereopsis from a very young age and thus had to learn to navigate and perform tasks using only monocular cues.

Concluding remarks

Thus, while not essential for vision, binocular disparity processing is remarkably useful for a host of visual computations, only one of which is the determination of depth, or stereopsis.

It is also an excellent example of the amazing precision that neural circuits can achieve — differences in relative retinal position can be detected that are nearly a factor of ten smaller than the spacing between photoreceptors, perhaps the most dramatic example of ‘hyperacuity’. Because stereopsis is a highly derived measure — requiring the precise comparison of signals from the two eyes — it has also provided an excellent test-bed for models of how neurons compute and how these computations are related to perception.

Acknowledgments

We thank Greg DeAngelis, Marge Livingstone, Ken Nakayama, and members of the Born lab for helpful comments.

Further reading

- Cumming, B.G., and DeAngelis, G.C. (2001). The physiology of stereopsis. *Annu. Rev. Neurosci.* 24, 203–238.
- CVnet discussion on the uses of stereopsis can be found at: http://www.sas.upenn.edu/~wilmer/CVNet_StereopsisConsequences_Syntheses.pdf
- DeAngelis, G.C., Cumming, B.G., and Newsome, W.T. (1998). Cortical area MT and the perception of stereoscopic depth. *Nature* 394, 677–680.
- Howard, I.P., and Rogers, B.J. (1995). *Binocular Vision and Stereopsis*, (Oxford University Press).
- Hubel, D.H. (1988). *Eye, Brain, and Vision*. New York, NY (Scientific American Library).
- Julesz, B. (1971). *Foundations of Cyclopean Perception*, (University of Chicago Press).
- Nakayama, K., Shimojo, S., and Silverman, G.H. (1989). Stereoscopic depth: its relation to image segmentation, grouping, and the recognition of occluded objects. *Perception* 18, 55–68.
- Nakayama, K., and Shimojo, S. (1990). Da vinci stereopsis: depth and subjective occluding contours from unpaired image points. *Vision Res.* 30, 1811–1825.
- Parker, A.J. (2007). Binocular depth perception and the cerebral cortex. *Nat. Rev. Neurosci.* 8, 379–391.
- Pettigrew, J.D. (1986). The evolution of binocular vision. In *Visual Neuroscience*, J.D. Pettigrew, K.J. Sanderson, W.R. Levick, eds. (Cambridge University Press), pp. 208–222.
- Roe, A.W., Parker, A.J., Born, R.T., and DeAngelis, G.C. (2007). Disparity channels in early vision. *J. Neurosci.* 27, 11820–11831.
- Sacks, O. (2006). “Stereo Sue.” *The New Yorker*. June 19, 2006; pp 64–73.
- Tyler, C.W. (2004). Binocular vision. In *Duane’s Foundations of Clinical Ophthalmology*, Vol. 2, W. Tasman and E.A. Jaeger, eds. (J.B. Lippincott Co.)
- Uka, T., and DeAngelis, G.C. (2004). Contribution of area MT to stereoscopic depth perception: choice-related response modulations reflect task strategy. *Neuron* 42, 297–310.
- Wheatstone, C. (1838) Contributions to the Physiology of Vision—Part the First. On some remarkable, and hitherto unobserved, Phenomena of Binocular Vision. *Phil. Trans. R. Soc. Lond.* 128, 371–394.

¹Department of Neurobiology, Harvard Medical School, 220 Longwood Avenue, Boston, Massachusetts 02115, USA.

²Harvard-MIT Division of Health Sciences and Technology, Harvard Medical School, 260 Longwood Avenue, Boston, Massachusetts 02115, USA.

*E-mail: richard_born@hms.harvard.edu