

Chapter 14

Representing 3D Shape and Location

Andrew Glennerster

14.1 A Primal Sketch That Survives Eye Rotation

Many of the chapters in this book are concerned with 2D shape whereas this chapter discusses the representation of 3D shape. However, I will argue that there is a strong link between these. 3D shape may be better understood in terms of the 2D image changes that occur when an observer moves than 3D ego-centered or world-centered coordinates frames. The same applies to representations of 3D location. 3D shape and 3D location are properties that remain the same as an observer moves through a static world, despite rapidly changing images. Two different conceptions for visual stability emerge. One relies on generating a representation that is like the world and is stable in the face of observer movements. The other relies only on an ability to predict the sensory consequences of a movement. The implications for representation of 3D shape (and location) are quite different under these two frameworks.

Most of the literature on visual stability focuses on a situation that is relatively straightforward from a computational perspective, namely that of a camera (or the eye) rotating around its optic center [3–7]. In this case, all the light rays we wish to consider arrive at a single optic center from all possible directions (a panoramic view, what Gibson called the ‘optic array’ at a single point). In computer vision, the process of ‘mosaicing’ a set of such images is now standard [8, 9]. In principle, it requires only that the rays corresponding to each pixel in each image to be registered in a common 2D coordinate frame, or sphere, of visual directions from the optic center. Nevertheless, this is a sensible starting point for considering visual stability in general. If points in the scene are all very distant (take, as an extreme example, the stars at night), the optic array remains unchanged wherever you move. If these points are stable in the representation, we have a sound foundation for explaining visual stability in general.

A. Glennerster (✉)
School of Psychology and Clinical Language Sciences, University of Reading,
Reading RG6 6AL, UK
e-mail: a.glennerster@reading.ac.uk




Fig. 14.1 Hierarchical encoding of position. An image (*top left*) is bandpass-filtered to show regions that are darker than the local mean luminance, including finer scale features in one part of the image, such as the fovea (*top right* image) or across the whole scene, e.g., after many saccades (*bottom left*). Because the combination of filter outputs follows the MIRAGE algorithm [1], there is a natural hierarchical encoding of position as shown schematically in the *bottom right* image (see also Fig. 14.3)

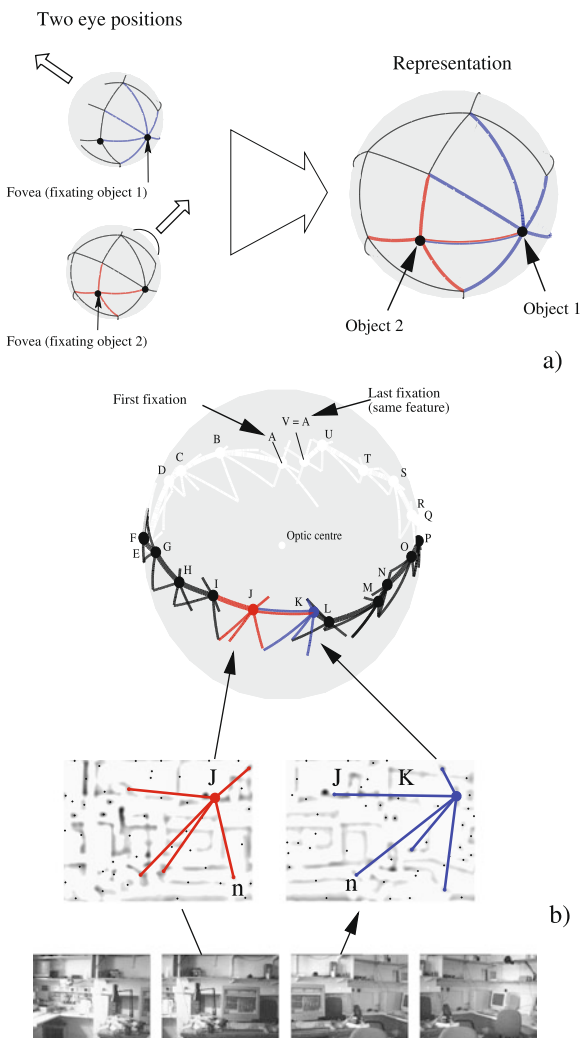
We are now in a position to consider translation of the optic center, either for a moving observer or the case of binocular vision. Translation of the optic center causes a change in the optic array. Two aspects of this change can be examined separately: first, the image change generated by a small patch in the scene and, second, the changes in the relative visual direction of objects that are separated by wide visual angles. The first is relevant for the representation of 3D surface shape; the second is relevant for encoding object location.

14.2 Translation of the Optic Center

14.2.1 Representing Surface Slant and Depth Relief

When viewing a small surface patch, the rays reaching the eye can be considered to be parallel (orthographic projection). This means that the ways the image of the surface deforms when the optic center translates are relatively simple. For example, the component of eye translation along the line of sight causes expansion (or contraction) while the orthogonal component causes 1D shear or stretch. The *axis* of the shear/stretch depends on the tilt of the surface, corresponding to the intersection of the plane perpendicular to the line of sight with the plane of the surface. The *direction* of the shear/stretch depends on the direction of the observer translation. The *magnitude* of the shear/stretch is influenced by the slant of the surface away from

Fig. 14.2 A representation of visual direction. **(a)** An eye that rotates about its optic center (which is an approximation to the truth in most cases) provides information about the relative visual direction of objects. Fixating different objects provides different sets of relative visual directions (e.g., blue and red arcs) which can be combined across the entire sphere to provide a single, stable representation of relative visual directions. **(b)** An illustration of forming this type of representation from images taken using a camera that rotates about its optic center, including the same image and primitives as used in Fig. 14.1. Features *J* and *n* appear in two of the images allowing them to be registered with the correct orientation (adapted, with permission, from Glennerster et al. [2])



fronto-parallel. Figure 14.3 shows one ‘patch’ or blob that has been stretched as a result of observer translation. It also shows how a hierarchical encoding of spatial location could help to implement a method of recording image changes. Koenderink and van Doorn [10] have proposed that surface structure could be represented using an image-based coordinate frame that would not require the generation of a 3D object-based representation. Because the three basis vectors of the frame are image based, the coordinates of all points on a rigid object remain unaffected by changes in viewpoint, rather like the coordinates of points on a deformable rubber sheet. A similar approach can be applied to the deformation of the blob shown in Fig. 14.3. The centroids of the blobs at each scale are recorded in relation to the centroid of the




Fig. 14.3 Consequences of translating the optic center. The ‘blobs’ shown in Fig. 14.1 are repeated here with, in *grey*, the changes that would be caused by a movement of the observer or a change from the left to right eye’s view. The lower blob has shifted to the *left* without any change in width, size or the configuration of the finer scale blobs within it. This is compatible with the surface being fronto-parallel and at a different depth from the other blobs. The centroid location of the *top left blob* has not changed so it is at the same depth as the *top right blob*. However, the width of the blob has changed, compatible with these features being on a slanted surface. The *inset* shows that in this case all the relative visual directions of the features (*yellow and white lines*) have changed together, as if drawn on a rubber sheet. These features all lie in the same slanted plane

blob at a larger scale. If the coordinate frame for measuring these relative positions is inherited from the scale above, that is, the distance metric is not measured in minutes of arc at the eye but relative to the width and height of the blob at the next coarsest scale, this would lead to a representation of location with similar properties to those advocated by Koenderink and van Doorn [10]. Shear, stretch or expansion of an image region caused by moving laterally or closer to a planar surface patch (as shown in Fig. 14.3) would yield no change in the relative position of the finer scale features if positions are measured in this locally-defined, hierarchical coordinate frame. Similarly, any depth relief of points relative to the surface plane would give rise to a change in hierarchical position when the viewpoint changes but this would be independent of the slant of the surface and signal only the relief relative to the surface [11, 12].

One difference between this hierarchical scale-based scheme and that of Koenderink and van Doorn [10] concerns the basis vectors used. In Koenderink and van Doorn’s scheme, provided that the points defining the three basis vectors are not co-planar, the coordinate of every point on a rigid object is recorded using the same basis vectors. But in the hierarchical system illustrated in Fig. 14.3, the coordinate frame is local and scale-based. This means that the representation amounts to something like a set of planar patches at each scale, each patch having a location, depth, tilt and slant defined relative to the ‘parent’ patch at the scale above. With this proviso, the scale-based hierarchy is very similar to the object-based representation Koenderink and van Doorn proposed and has the advantage of avoiding an explicit 3D coordinate frame.

A series of psychophysical studies support the hypothesis that the visual system may use a surface-based coding system of this sort. Mainly, these studies have investigated the processing of binocular disparity but there is also some evidence from structure from motion experiments [13]. Mitchison and McKee [14] showed that binocular correspondences in an ambiguous stereogram were determined not by a nearest-neighbor rule using retinal coordinates to define proximity, as had always been supposed, but by proximity to an invisible ‘interpolation’ surface drawn between the edges of the patch. This is equivalent to the prediction of the hierarchical ‘rubber sheet’ representation outlined above, in which the metric for measuring the location of dots in the left and right eyes is determined by the shear/stretch of the patch in that eye. Like correspondence, perceived depth relief is also determined by the disparity of a point relative to a local surface even when observers are remarkably insensitive to the slant of the surface [15–17]. Finally, sensitivity to depth perturbations are determined not by the disparity of a point relative to neighboring points but instead by its disparity relative to an invisible interpolation plane [12, 18, 19], as a ‘rubber sheet’ model would predict.

As an aside, it is worth noting that the hierarchical encoding of blob location proposed here (following Watt and Morgan [1, 20]) brings some theoretical disadvantages but there is experimental evidence to suggest that the visual system may be prepared to pay this cost. In the coarse-to-fine stereo correspondence algorithm proposed by Marr and Poggio [21], the ‘coarse scale’ version of an image is always sparse, with large spacing between features (in their case, ‘zero-crossings’). This means that there will always be relatively wide gaps between true and false matches along any given epipolar line and hence a nearest-neighbor rule will yield correct correspondences over a wide range of disparities. In Watt and Morgan’s MIRAGE scheme, however, the ‘coarse scale’ representation is generated by summing the ‘on’ responses of filters at all spatial scales and, separately, the ‘off’-responses. While this has the merit that the fine scale features *always* lie within the boundary of coarse-scale blobs, the disadvantage is that in certain situations the ‘coarse scale’ representation can be much more densely packed with features than the pure low frequency channel output envisaged by Marr and Poggio. Figure 14.4 shows such a situation: a dense random dot pattern with, on the right, a MIRAGE ‘coarse scale’ output and a schematic version to illustrate how the ‘sea’ between the low frequency blobs have been ‘filled in’. A random dot pattern has much greater power at high frequencies than natural images and perceptually it appears far more crowded than most images. Glennerster [22] measured the ability of the visual system to find matches when random dot patterns were shifted (either in motion or by adding disparity) and showed that MIRAGE primitives predicted well the magnitude of shift that the visual system could tolerate before the perception of motion or stereo depth broke down. This price (a small D_{\max} for high density patterns) appears to be an acceptable sacrifice for the visual system. The positive benefit is that fine scale features always have a simple, hierarchical ‘address’ to define their location.




Fig. 14.4 A penalty for hierarchical encoding. If fine scale features are always to lie within the boundaries of coarse scale features, as they do in the MIRAGE algorithm [1] and illustrated in Figs. 14.1 and 14.3, then the ‘coarse scale’ representation must inevitably be more crowded than a low-pass version of the image. This is particularly evident in white noise images such as the random dot pattern shown here. In a D_{\max} task (see text), observers behave as if their representation of this type of image is quite crowded with features, as shown on the *top right* (reproduced, with permission, from [22]). The *white dots* mark the centroids of each blob measured along horizontal raster lines. The ‘coarse scale’ representation is crowded, as shown schematically in the *bottom right panel*, because blobs originating from different low, medium and high spatial frequency filters all contribute to the representation (see *bottom left panel*) and ‘fill in the sea’ between low spatial frequency ‘islands’

14.2.2 Representing Location

Having considered the effect of observer translation on a small patch of the visual field, we now turn to the consequences for widely separated features. There are strong similarities between these two scales but also important differences. In particular, disparity and motion of a small patch provide useful information about surface shape while changes in relative position of widely separated features, such those shown in Fig. 14.2a, provide information about object location.

Unlike the image changes in a small region of the visual field, the changes in relative visual direction of widely separated features do not suffer from the ‘bas relief ambiguity’. This refers to the fact that a small disparity or motion can be due either to the depth relief being small or to the patch being far away. By contrast, for two widely separated features, if the angle separating them does not change when the observer moves (or there is no change between the left and right eye’s view) then, in general, the points are distant: the bas relief ambiguity has disappeared (discussed

in detail by Glennerster et al. [2]). The tendency for the relative visual direction of two features to change as the observer moves gives useful information about whether those features belong to near or distant objects. The most distant points in a scene form a set whose relative visual directions (the angles separating each pair and triple of points) are the most stable when the observer translates. Against the background of these distant objects, nearer objects ‘slide around’ as the observer moves [8]. One could turn this around and propose, in Gibsonian fashion, that an observer moves themselves from one place to another by ‘grabbing’ an object (visually, by fixating it) and ‘pushing it’ one way or another against the background (by walking, say) until it is in the desired place relative to the background.

The advantage of this representation is that the 3D origin of the coordinate frame is never defined. This makes sense. If you are star-gazing and see only stars, their relative visual directions do not change as you move and hence they provide no information about where you are on earth. The location of the 3D origin is impossible to define. Distant mountains allow your location to be defined more precisely, nearby trees even more so. The closer the objects in view, the more it becomes possible to pinpoint the location of the origin. Only with near objects in view would it make sense to distinguish between the origin of a coordinate system being at the eye, head, body or hand. If, however, the goal is not to build a 3D coordinate frame at all but instead to build an image-based representation, then the stars, the mountains, trees and very near objects provide a hierarchical method of locating the current image in that representation. These ideas are discussed in detail by Glennerster, Hansard and Fitzgibbon [2, 23].

In summary, both 3D shape and 3D location can be considered as properties derived from the changes in relative visual direction of features produced by observer translation. The way that each of these are encoded in the visual system should leave traces when we test psychophysical performance, as we have discussed. Two further examples are described in the final section (Sect. 14.4).

14.3 Implementation of a Universal Primal Sketch

There is no pretence that the suggestions raised in this chapter are anything like a recipe for implementation, but they do provide some useful pointers. The case of a camera rotating around its center is an exception. In that case, a solution was described by Watt 25 years ago [1, 24], with the location (visual direction) of features defined hierarchically across scale space for the entire optic array. But once the optic center of the camera or observer translates, practical issues emerge that are considerably more tricky.

One example is the matching process that must link data structures describing the same surface seen from different view points. For example, if a surface is viewed from two distances, the spatial frequency of the filters responding to features on the surface will be higher for the farther viewing distance but if scales, like positions, are defined relative to one another, then the data structure recording fine scale features and a coarse scale outline of the object might be relatively unchanged by this

alteration in viewing distance. Relative measures are likely to be a prominent aspect of the primal sketch. Of course, in the real world, with real images, complex changes occur with changes in viewpoint due to cast shadows, occlusions and specularities. The suggestions made in this chapter provide no quick fix for these problems.

It is also worth questioning the extent to which a view-based representation could underlie *all* visual tasks, not just the ones described here. One particularly problematic class of tasks involves imagining you are at a different location and making responses as if you were there. In a familiar environment, the observer may have visited that location in the past, in which case it is possible that an observer could ‘run the tape’ instead of actually walking to the new location and solve the task that way. But people are able to imagine being on the other side of a room that they have never seen before and to make judgements as if from that location. In our lab, we are currently exploring ways to model behaviors of this type using view-based methods, without relying on the assumption that the brain generates a Cartesian representation of the scene. In general, it is not yet clear what the limits will be to the set of tasks that could be carried out using a primal sketch or view-based framework.

14.4 Apparent Paradoxes in the Representation of 3D Shape and Location

The primal sketch outlined in this chapter is a source of ‘raw’ visual information that could be used for many different tasks. We discuss here two experiments that show how participants’ performance appears paradoxical if we assume the visual system uses a 3D representation but both experiments are readily explained if we suppose that the visual system extracts ‘raw’ visual information once the task is defined [25]. In one case, the task is a judgement of object shape and in the other it is a judgement of object location.

Figure 14.5 illustrates the shape task. We know that under rich-cue conditions, people show good size constancy and good depth constancy when they compare the size or depths of similar objects across different distances [26, 27] but exhibit large biases when asked to make a judgement of the metric shape of a surface such as comparing the depth to the half-height of a horizontal cylinder [27–29]. In the case shown in Fig. 14.5, the visual system must apparently estimate four values, namely the depths and half-heights of two semi-cylinders presented at two distances: d_1, h_1 and d_2, h_2 . If these values were all available to the visual system, independent of the task the participant was set, then it would not be possible for participants to judge $d_1 \approx d_2, h_1 \approx h_2$ and yet, under the same viewing conditions, $d_1 > h_1, d_2 < h_2$ (i.e. d_1 judged as reliably larger than h_1 but d_2 judged to be reliably smaller than h_2). Yet, this is what observers see. If they built a single consistent representation of the scene and accessed the values d_1, h_1, d_2 and h_2 from this representation for all tasks, then the data would present a paradox. However, comparisons of height (h_1 versus h_2) can be done with other short-cuts, such as comparing the retinal size of test objects to other objects in the scene and the same is true of the comparisons of




Fig. 14.5 Paradoxical representations of shape. Observers are good at size constancy ($h_1 = h_2$) and depth constancy ($d_1 = d_2$) but, under essentially identical viewing conditions, they make systematic errors when judging the shape of objects ($d_1 > h_1$ while at the same time $d_2 < h_2$). The solution to the apparent paradox is to assume that in each case, once the task is defined, the visual system acquires the relevant information and computes the solution. One task depends on an estimate of viewing distance (e.g., D_1) while the other requires only an estimate of the ratio of viewing distances to the two objects (D_1/D_2) [27]




Fig. 14.6 Paradoxical representation of location. In virtual reality, observers judged the relative depth of two squares presented in separate intervals. Sometimes the room expanded between intervals (A to B and C to D), although the participants never noticed a change in room size [30]. On the other trials, the room stayed still (small room: A to C or large room B to D). It is impossible to determine a single location of D relative to A that is compatible with all the pairwise settings observers make. However, similar to Fig. 14.5, there is no paradox if the visual system acquires the relevant information for any given comparison once the task is defined

depths. By contrast, comparing d_1 to h_1 or d_2 to h_2 requires an estimate of absolute (not relative) viewing distance which means that these estimates are open to a source of bias that does not affect the other judgements [27]. The important point is that these data provide compelling evidence that the visual system uses information in a more ‘raw’ form than the metric values d_1 , h_1 , d_2 and h_2 when carrying out these judgements of 3D shape.

For 3D location, a good example of an apparent paradox is the case illustrated in Fig. 14.6 from Svarverud et al. [30]. Several experiments using immersive virtual reality have shown that moving observers fail to see a room changing in size around them, by as much as a factor of four in all directions, provided that looming cues are eliminated [31–33]. This is compatible with earlier evidence on observers’ poor sensitivity to change in disparity in the absence of looming cues [34] and raises

interesting questions about the type of representation that observers must be building of the scene. Svarverud et al. [30] measured subject's biases when they judged the relative depth of objects either with or without an expansion of the room between the presentation of the two objects. Observers did not notice any difference between these two types of trial. As Fig. 14.6 illustrates, although their perception of the room was stable throughout, their pairwise depth matches cannot be explained by a single, consistent 3D representation. There is, therefore, no one-to-one mapping between a participant's internal representation of the room and a single static 3D room. It does not matter that the stimulus is an unusual one. The point is that the observer's perception is one of an ordinary, stable room so the conclusions we draw from probing the representation underlying that perception should apply to other ordinary, stable scenes.

These examples raise questions about what the minimum requirements are for a useful representation of the scene. It is no use claiming, as Gibson often appeared to [35], that an internal representation is unnecessary. More recent accounts emphasise the importance of information stored 'out in the world' rather than in the head [25], but these still require a coherent set of rules that will allow the information 'out there' to be accessed. The stored information must remain useful even if the object or visual information is not within the current field of view. This chapter outlines a possible primal sketch of blob location that is an example of a representation of 'raw' visual information. Something like this might, with further elaboration, fulfil the criteria for a store that could be used to access information 'out there'. Such a representation must store sufficient information to allow the observer to turn their gaze to any object they remember and, if necessary, walk in the right direction until the object comes into view. It must also contain information about the slant of surfaces and the depth relief of points compared to local surfaces. These requirements fall short of the attributes of a full 3D reconstruction, but psychophysical evidence suggests the same is true of human vision.

Acknowledgement Supported by the Wellcome Trust.

References

1. Watt RJ (1987) Scanning from coarse to fine spatial scales in the human visual system after the onset of a stimulus. *J Opt Soc Am A* 4:2006–2021
2. Glennerster A, Hansard ME, Fitzgibbon AW (2001) Fixation could simplify, not complicate, the interpretation of retinal flow. *Vis Res* 41:815–834
3. Duhamel JR, Colby CL, Goldberg ME (1992) The updating of the representation of visual space in parietal cortex by intended eye-movements. *Science* 255:90–92
4. Zipser D, Andersen RA (1988) A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature* 331:679–684
5. Bridgeman B, van der Heijden AHC, Velichovsky BM (1994) A theory of visual stability across saccadic eye movements. *Behav Brain Sci* 17:247–292

6. Melcher D (2007) Predictive remapping of visual features precedes saccadic eye movements. *Nat Neurosci* 10(7):903–907
7. Burr DC, Morrone MC (2011) Spatiotopic coding and remapping in humans. *Philos Trans R Soc Lond B, Biol Sci* 366(1564):504–515
8. Irani M, Anandan P (1998) Video indexing based on mosaic representation. *Proc IEEE* 86:905–921
9. Brown M, Lowe DG (2007) Automatic panoramic image stitching using invariant features. *Int J Comput Vis* 74(1):59–73
10. Koenderink JJ, van Doorn AJ (1991) Affine structure from motion. *J Opt Soc Am A* 8:377–385
11. Mitchison G (1988) Planarity and segmentation in stereoscopic matching. *Perception* 17(6):753–782
12. Glennerster A, McKee SP (2004) Sensitivity to depth relief on slanted surfaces. *J Vis* 4:378–387
13. Hogervorst MA, Glennerster A, Eagle RA (2003) Pooling speed information in complex tasks: estimation of average speed and detection of non-planarity. *J Vis* 3:464–485
14. Mitchison GJ, McKee SP (1987) The resolution of ambiguous stereoscopic matches by interpolation. *Vis Res* 27:285–294
15. Mitchison GJ, McKee SP (1990) Mechanisms underlying the anisotropy of stereoscopic tilt perception. *Vis Res* 30:1781–1791
16. Cagenello R, Rogers BJ (1993) Anisotropies in the perception of stereoscopic surfaces—the role of orientation disparity. *Vis Res* 33:2189–2201
17. Bradshaw MF, Rogers BJ (1999) Sensitivity to horizontal and vertical corrugations defined by binocular disparity. *Vis Res* 39:3049–3056
18. Glennerster A, McKee SP, Birch MD (2002) Evidence of surface-based processing of binocular disparity. *Curr Biol* 12:825–828
19. Petrov Y, Glennerster A (2004) The role of a local reference in stereoscopic detection of depth relief. *Vis Res* 44:367–376
20. Watt R, Morgan M (1983) Mechanisms responsible for the assessment of visual location: theory and evidence. *Vis Res* 23:97–109
21. Marr D, Poggio T (1979) A computational theory of human stereo vision. *Proc R Soc Lond B, Biol Sci* 204:301–328
22. Glennerster A (1998) D_{\max} for stereopsis and motion in random dot displays. *Vis Res* 38:925–935
23. Glennerster A, Hansard ME, Fitzgibbon AW (2009) View-based approaches to spatial representation in human vision. *Lect Notes Comput Sci* 5064:193–208
24. Watt RJ (1988) Visual processing: computational, psychophysical and cognitive research. Erlbaum, Hove
25. O'Regan JK, Noë A (2001) A sensori-motor account of vision and visual consciousness. *Behav Brain Sci* 24:939–1031
26. Burbeck CA (1987) Position and spatial frequency in large scale localisation judgements. *Vis Res* 27:417–427
27. Glennerster A, Rogers BJ, Bradshaw MF (1996) Stereoscopic depth constancy depends on the subject's task. *Vis Res* 36:3441–3456
28. Johnston EB (1991) Systematic distortions of shape from stereopsis. *Vis Res* 31:1351–1360
29. Tittle JS, Todd JT, Perotti VJ, Norman JF (1995) A hierarchical analysis of alternative representations in the perception of 3-d structure from motion and stereopsis. *J Exp Psychol Hum Percept Perform* 21:663–678
30. Svarverud E, Gilson S, Glennerster A (2012) A demonstration of 'broken' visual space. *PLoS ONE* 7:e33782. doi:[10.1371/journal.pone.0033782](https://doi.org/10.1371/journal.pone.0033782)
31. Glennerster A, Tcheang L, Gilson SJ, Fitzgibbon AW, Parker AJ (2006) Humans ignore motion and stereo cues in favour of a fictional stable world. *Curr Biol* 16:428–443
32. Rauschecker AM, Solomon SG, Glennerster A (2006) Stereo and motion parallax cues in human 3d vision: can they vanish without trace? *J Vis* 6:1471–1485

33. Svarverud E, Gilson SJ, Glennerster A (2010) Cue combination for 3d location judgements. *J Vis* 10:1–13. doi:[10.1167/10.1.5](https://doi.org/10.1167/10.1.5)
34. Erkelens CJ, Collewijn H (1985) Motion perception during dichoptic viewing of moving random-dot stereograms. *Vis Res* 25:583–588
35. Gibson JJ (1950) *The perception of the visual world*. Houghton Mifflin, Boston