

Below is the unedited draft of:

Barlow, Horace (2001) The Exploitation of Regularities in the Environment by the Brain
Behavioral and Brain Sciences 24 (3): XXX-XXX.

This is part of a BBS Special Issue on the Work of Roger Shepard containing the following articles:

Shepard, R. N. (2001)

Perceptual-Cognitive Universals as Reflections of the World

Behavioral and Brain Sciences 24 (3): XXX-XXX.

Reprinted from Psychonomic Bulletin & Review, 1994, 1, 2-28

Barlow, H. (2001)

The Exploitation of Regularities in the Environment by the Brain

Behavioral and Brain Sciences 24 (3): XXX-XXX.

Hecht, H. (2001)

Regularities of the Physical World and the Absence of their Internalization

Behavioral and Brain Sciences 24 (3): XXX-XXX.

Kubovy, M. (2001)

Internalization: A metaphor we can live without [PDF version]

Behavioral and Brain Sciences 24 (3): XXX-XXX.

Schwartz, R. (2001)

Evolutionary Internalized Regularities

Behavioral and Brain Sciences 24 (3): XXX-XXX.

Tenenbaum, J. & Griffiths, T. L. (2001)

Generalization, Similarity, and Bayesian Inference

Behavioral and Brain Sciences 24 (3): XXX-XXX.

Todorovi , D. (2001)

Is kinematic geometry an internalized regularity?

Behavioral and Brain Sciences 24 (3): XXX-XXX.

This is the unedited draft of a BBS target article that has been accepted for publication (Copyright 1999: Cambridge University Press U.K./U.S. -- publication date provisional) and is currently being circulated for Open Peer Commentary. This preprint is for inspection only, to help prospective commentators decide whether or not they wish to prepare a formal commentary. Please do not prepare a commentary unless you have received the hard copy,

invitation, instructions and deadline information.

For information on becoming a commentator on this or other BBS target articles, write to: bbs@soton.ac.uk

For information about subscribing or purchasing offprints of the published version, with commentaries and author's response, write to: journals_subscriptions@cup.org (North America) or journals_marketing@cup.cam.ac.uk (All other countries).

The Exploitation of Regularities in the Environment by the Brain

Horace Barlow

Physiological Laboratory

Cambridge CB2 3EG

England

Email: hbb10@cam.ac.uk

URL: <http://www.physiol.cam.ac.uk/staff/barlow>

Horace Barlow, retired Royal Society Research Professor of Physiology at the University of Cambridge, has worked in Vision and Perception for 50 years. He contributed to the experimental discovery of bug-detectors and lateral inhibition in the frog retina, motion detection in the rabbit retina, and selectivity for disparity in the cat cortex. As a theoretician, he has argued that exploiting the redundancy of natural stimuli is important for effective perception and cognition, and that measuring absolute statistical efficiencies is useful for defining the work involved.

Short abstract

Past work on the exploitation of statistical regularities of the environment by sensory and cognitive mechanisms is briefly reviewed. Shepard proposed the related idea that the laws of kinematic geometry have been internalised, but these laws are mathematical relations rather than observable statistical regularities, and "internalising" does not specify how they bring advantages. The apparent trajectory of an object when seen successively at two positions could result from perceptual mechanisms that exploit observable statistical regularities as well as from internalising Chasles' rule. In either case, the observations suggest that cortical neurons might usefully be selective for twisting motions rather than pure translations.

Long abstract

Statistical regularities of the environment are important for learning, memory, intelligence, inductive inference, and in fact for any area of cognitive science where an information-processing brain promotes survival by exploiting them. This has been recognised by many of those interested in cognitive function, starting with Helmholtz, Mach and Pearson, and continuing through Craik, Tolman, Attneave, and Brunswik. In the current era many of us have begun to show how neural mechanisms exploit the regular statistical properties of natural images. Shepard proposed that the apparent trajectory of an object when seen successively at two positions results from internalising the rules of kinematic geometry, and although kinematic geometry is not statistical in nature, this is clearly a related idea. Here it is argued that Shepard's term, "internalisation", is insufficient because it is also necessary to derive an advantage from the process. Having mechanisms selectively sensitive to the spatio-temporal patterns of excitation commonly experienced when viewing moving objects would facilitate the detection, interpolation, and extrapolation of such motions, and might explain the twisting motions that are experienced. Although Shepard's explanation in terms of Chasles' rule seems doubtful, his theory and experiments illustrate that local twisting motions are needed for the analysis of moving objects and provoke thoughts about how they might be detected.

Keywords: Chasles' rule, evolution, geometry, perception, redundancy, statistics, twisting.

1.0 Introduction

Statistical regularities abound in the world around us, and many of them are actually, or potentially, important for our survival. Furthermore many of them are obviously exploited by our bodies, and the anatomy of the eye provides especially beautiful examples of evolutionary adaptations to different environments (Walls 1942). Shepard's papers (1984, 1994) were a major inspiration for the program at ZiF in Bielefeld that was the major source for this collection of articles. He evidently believed that the apparent motion trajectory of an object shown first at one position, then at another, resulted from the evolutionary adaptation of psychological mechanisms to the kinematic geometry of moving objects. Now kinematic geometry is concerned with non-probabilistic geometric relations, not statistics, but the way Shepard thought that perception was adapted to its rules parallels the way others have thought of perception adapting to statistical regularities. This article starts by giving a brief history of the development of these ideas, and then compares their predictions with Shepard's in the particular conditions he explored.

First we should perhaps note that all learning could be regarded as the internalisation of environmental regularities, for it is driven by the statistically regular occurrence of reinforcement following particular sensory stimuli or self-initiated actions. Furthermore it is well-recognised that statistical associations between sensory stimuli, as well as associations between the stimuli and reinforcement, influence learning (Rescorla & Wagner 1972; Mackintosh 1983). This makes it difficult to say where the use of statistical regularities in perception stops and their use in learning begins, but let us start with a brief historical review of the claims that have been made about their use in perception.

2.0 Helmholtz

Helmholtz flourished before Darwinian ideas about genetic adaptations to the environment were widely acknowledged, but he argued unremittingly that

perception results from the interaction of apperception - the immediate impact of sensory messages - with remembered ideas resulting from past experience. He wrote of perceptions that "...by their peculiar nature they may be classed as conclusions, inductive conclusions unconsciously formed" (Helmholtz 1925). Thus he held the view that experience of the environment was internalised or remembered and provided the basis, together with current sensory messages, for the statistical conclusions — mostly valid - that constitute our perception of the world.

3.0 Mach and Pearson

Ernst Mach (1886) and Karl Pearson (1892) also appreciated the importance of environmental statistics, but they viewed the matter somewhat differently. They argued that scientific concepts and laws simplify our complex experience of the world, and that they are important because they bring "Economy of Thought" to our mental processes. Although this idea has great appeal at an intuitive level, it can only be made convincing if economy can be measured, and that had to wait for the quantitative definition of information and redundancy (Shannon & Weaver 1949).

4.0 Craik

In a short book Craik (1943) developed the idea that the main function of the higher cognitive centres is to build symbolic working models. Such models must be based on the associative structure of objects and events in the environment, and are therefore expressions of environmental regularities. It is a more general form of Tolman's (1948) idea of cognitive maps.

5.0 Brunswik

Egon Brunswik (1956) seems to have been the first to suggest that the Gestalt laws governing grouping and segregation of figure from ground were more than empirical facts about perception: they were rules for using statistical facts about images to draw valid inferences from the scene immediately before the eyes. His work is not often quoted so it is worth describing in greater detail. He pointed out that objects have uniform properties compared with randomly selected regions

of an image, and hence if two patches have similar local characteristics they are likely to be derived from the same object in the external world; this is the reason why it is appropriate to group them together. By analysing stills from the Alec Guinness movie "Kind Hearts and Coronets" Brunswik and Kamiya (1953) were able to show that there was a tendency for the proximity of two parallel lines to indicate the presence of a manipulable object in the scene, though this was a disappointingly weak effect.

The methods available then were feeble compared with those available now, and recent work on the statistics of natural images (Ruderman 1997) has shown that correlations of straightforward luminance values are indeed much stronger between points that lie within the same object than they are between points lying in different objects. Within the brain, images are not represented just by luminance values but by neurons selective for features such as orientation, texture, colour, disparity, and direction of motion. It will be interesting to see if the difference between inter- and intra-object correlations are even greater for these features than for luminance; if this is so it would go a long way towards showing that the Gestalt laws of proximity, good continuation, common fate, and so on are rules for making valid statistical inferences from environmental regularities. Elder and Goldberg (1998) have recently studied the validity of various properties of edges for bringing about correct object segregation.

It is worth noticing that Brunswik's idea makes a lot of sense of the anatomical arrangement of primary visual cortex (V1) and the surrounding extra-striate areas (Barlow 1981). V1 has neurons selectively sensitive to those local characteristics of the image that cause grouping, namely orientation, colour, texture, disparity, and direction of motion, most of which had already been identified by the Gestaltists. V1 neurons then project topographically to surrounding extra-striate areas creating new maps (see Lennie 1998), but there is also a non-topographic component in the projection. Neurons in V1 and V2 that are selectively sensitive to a particular feature (eg movement in a certain direction) converge on to single neurons in these extrastriate areas, thus collecting together information about this feature from relatively large regions of the visual field. Assembling the information in this way is the crucial step that enables such a

feature to be detected at a low signal-to-noise ratio (Barlow and Tripathy 1997) even when it is spread over a fairly large patch of the image. Perhaps we are beginning to understand the physiology as well as the statistical logic of these first stages of object recognition.

6.0 Attneave

Attneave (1954) imported into psychology mathematical concepts that had been developed by Shannon and Weaver (1949) to quantify the transmission of information down communication channels. The most important of these from the present point of view are information, channel capacity and redundancy. A communication channel can only transmit information at rates up to a finite limit called its capacity, but the messages actually transmitted often contain less than this amount of information; the difference is the redundancy of the messages. The importance here is that any form of regularity in the messages is a form of redundancy, and since information and capacity are quantitatively defined, so is redundancy, and we have a measure for the quantity of environmental regularities.

Attneave pointed out that there is much redundancy in natural images and suggested that the subjective prominence of borders provides an example of a psychological mechanism that takes advantage of this fact: you can represent an object more economically by signalling transitions between object and non-object because these are the unexpected, and therefore information-bearing, parts of the image. He illustrated with his famous picture of a sleeping cat that the same rule applies to the orientation of boundaries, for the picture was produced simply by connecting the major transition points in the direction of the border that outlines it.

7.0 Barlow

I became interested in the importance of statistical structure in sensory messages as soon as I came across Shannon's definitions of information, capacity, and redundancy. It seemed to me (Barlow 1959, 1961) that redundancy must be

important throughout our sensory and perceptual system, from the earliest coding of physical messages by sensory receptors, right through to the intelligent interpretation of the patterns of excitation that occurs at the highest cognitive levels (Barlow 1983).

There has been one major change in my viewpoint. Initially I thought that economy was the main benefit to be derived from exploiting redundancy, as it is for AT&T and British Telecom. But, as explained in greater detail below, the physiological and anatomical facts do not fit the idea that the brain uses compressed, economical, representations, and one can see that these would be highly inconvenient for many of the tasks it performs, such as detecting associations. Therefore I now think the principle is redundancy exploitation, rather than reduction, since performance can be improved by taking account of sensory redundancy in other ways than by coding the information onto channels of reduced capacity.

My initial idea was similar to Attneave's, described above, but what excited me was the fact that one could point to physiological mechanisms, such as the accommodation of sensory discharges to constant stimuli, light and dark adaptation, and lateral inhibition, that actually put the principles to work. I first wrote about it in 1956 (though the article was not published until 1961), and made predictions about the coding of motion that have subsequently been confirmed. If I had been smart enough I would have predicted the orientational selectivity of cortical neurons that Hubel and Wiesel (1959) discovered, for it has been shown that this fills the bill for redundancy exploitation (Ohlshausen & Field, 1996; van Hateren & van der Schaaf, 1998). An attractive feature of the idea is that a code formed in response to redundancies in the input would constitute a distributed memory of these regularities - one that is used automatically and does not require a separate recall mechanism.

The original article (Barlow 1961) suggested sparse coding, ie that the economy is brought about by reducing the frequency of impulses in neurons carrying the representation rather than by reducing the number of neurons involved. Barlow (1972) is mainly concerned with experimental evidence showing

that single neurons in sensory pathways are highly sensitive and selective in their response properties; hence perceptual discriminations can be based very directly upon their activity and may characteristically depend upon only a few of the most active neurons. The article also develops the idea of sparse coding, where the activity of a small number of neurons selected from a very large population forms a distributed representation of the sensory input (see also Field 1994). The elements of this type of distributed representation are called "cardinal cells" to indicate their partial resemblance to Sherrington's "pontifical neurons". They signal directly the occurrence of messages belonging to subsets of the possible sensory inputs that it would be useful for an animal to learn about. The elements of distributed representations are often assumed to represent random or arbitrary subsets of input states, whereas a cardinal cell representation has some of the merits of grandmother or mother cell representations (see Lettvin's note in Barlow 1995), as well as those of sparse distributed representations.

Barlow (1989) argued for the general importance of the associative structure of sensory messages and proposed factorial coding, in which representative elements are formed that are statistically independent of each other, as a means of storing knowledge of these environmental regularities. Barlow (1990) suggested that motion and other after-effects result from adaptive mechanisms that tend to make representational elements independent of each other, and Carandini et al (1997) provided some experimental evidence for the predicted contingent adaptation in neurons of monkey V1. Barlow (1996) reviewed some of this work and attempted to bring it up to date in the general context of Bayesian inference and perception.

The idea of economy in representation that Mach and Pearson proposed and Attneave and I recast in terms of Shannon's redundancy provides a key to understanding much in sensation, perception and cognition, but the problem whether channel capacity decreases at higher levels in the brain needs to be faced. Initially it would seem that, if redundancy is to be reduced, the transformations in sensory pathways would have to generate very compact sensory representations with a reduced number of channels, each active for a high proportion of the time. In fact almost the opposite occurs: at higher levels in the brain there are vastly

more channels, though it is true that each is active at a lower rate. The increase in cell numbers is enormous, with more than a thousand times as many neurons concerned with vision in the human cortex as there are ganglion cells in the two retinas. The average frequency of impulses certainly becomes lower in the cortex, so coding does become sparser, but even if the capacity is deemed to be limited by this reduced mean firing rate, the increased number of cells dominates: on any plausible assumptions the capacity of the cortical representation is vastly greater than that of the retinal ganglion cells, so redundancy appears to be increased, not reduced. But as mentioned earlier, economising the capacity of the central representation may not be the factor of importance in the brain.

Information theory has always assumed that cost is proportional to channel capacity, and the commercial value of redundancy reduction lies in the reduction of costs it can bring about. Linking cost with channel capacity was appropriate for man-made communication channels, but costs and benefits may be quite different in the brain, and A R Gardner-Medwin and I have been looking at the "cost" for efficient learning that results from the use of distributed representations (Gardner-Medwin & Barlow, 2000). Efficient learning requires the ability to count the occurrences of different attributes of sensory stimuli with reasonable accuracy, but unavoidable errors occur in distributed representations, where different attributes activate the same neuron. This overlap causes an increase in the variance of frequency estimates, which means that learning for features or events represented in a distributed manner must require the collection of more evidence than is necessary for features or events represented directly, in a localist manner. To reduce this loss of efficiency it is necessary to use many more neurons in distributed representations - ie to increase their redundancy. This reduces the extent to which representational elements are active in events other than the one that is being counted.

Distributed representations allow an enormous number of different input patterns to be distinguished by relatively few neurons, and no-one doubts that they are used in the brain. But for learning associations it is not sufficient just to distinguish different input patterns - one must also estimate how often they occur - and to do this with reasonable efficiency requires highly redundant representations.

Perhaps this is a reason why there are so many neurons in the cortex.

Note that redundancy is a measure of any kind of statistical regularity, and there is no necessary relationship between the redundancy that can be exploited in the input, and the redundancy added by using a number of elements that is unnecessarily large for representation alone. Sensory redundancy is important because knowledge of regularities in the environment is advantageous for many purposes, such as making predictions. Redundancy in the representation has a quite different role: it reduces the extent to which elements are necessarily active for more than one type of input event, which is what hinders accurate counting. There is no guarantee that the redundancy in the input would achieve the latter purpose, though it might be possible to transform input redundancy directly into a form that would reduce overlaps in the way required. Perhaps the "repulsion" between frequently co-active elements postulated to account for pattern-contingent after-effects (Barlow 1990) represents such a mechanism.

8.0 Watanabe

Watanabe (1960) drew attention to the similarity between inductive inference and recoding to reduce redundancy, a theme also taken up by Barlow (1974). If a particular type of regularity is identified in a mass of data, then it is possible to represent those data more compactly by exploiting the regularity. Carrying this argument to its logical conclusion, inductive inference is a matter of using statistical regularities to produce a shorter, more compact, description of a range of data, an idea that is carried further in the minimum description length (MDL) approach to these problems.

9.0 Minimum Description Length

Solomonoff (1964a,b) and Wallace and Boulton (1968) suggested that the computer code with the minimum length necessary to reproduce a sequence of data provides the shortest, and therefore least redundant, representation of that data. This is obviously related to Occam's razor, to Mach and Pearson's ideas about economy of thought, and to later ones on economy of impulses. The idea has been related to Bayesian interpolation (MacKay 1992), to the problem of simplicity and

likelihood in perceptual organisation (Chater 1996), and to the general problem of pattern theory (Mumford 1996). Like redundancy exploitation, it uses regularities in the data to provide a basis for induction and prediction.

10.0 Recent work in this area

These ideas have become very much alive in the last decade or so. Linsker (1986a,b,c) applied information theory to understand the properties of neurons in the visual pathway, and Field (1987) tried to relate the properties of cortical neurons to the amplitude distribution for spatial frequencies in natural images. The statistics of such images were investigated by Tolhurst, Tadmor & Chao (1992), Ruderman (1994), Field (1994), Baddeley (1996) and Baddeley et al. (1997). Atick (1992) reviewed the redundancy reduction idea, and Atick & Redlich (1990, 1992) argued that visual neurons were adapted to deal with the statistics of natural images. The idea that the form of the receptive field of V1 neurons is specifically adapted to the regularities of natural images has been around for a long time (eg Barrow 1987, Webber 1991) but early attempts were not very successful in using this principle to derive receptive fields like the real ones from the statistical properties of natural images. More recent attempts (Bell & Sejnowski 1995; Olshausen & Field 1996; Hyvarinen & Oja 1996) have used non-linear methods and somewhat different principles. For instance van Hateren & van der Scharff (1998) ran a program on natural images that performed Independent Component Analysis. This determines what receptive fields would be expected if the goal was to produce a limited number of descriptors of image patches that would, when added in the right proportions, generate accurate replicas of the range of images it was trained on. They showed that the predicted receptive fields match those determined experimentally in some, though not all, of their properties. More remains to be done along these lines, but it seems probable that the receptive fields of V1 neurons are indeed adapted to the regularities of natural images.

The work reviewed above suggests important roles for neurophysiological mechanisms that exploit the redundancy of sensory messages resulting from statistical regularities of the environment. For example it has given us an idea why sensory nerves accommodate, why lateral inhibition occurs, why neurons are

selectively sensitive to movement, why cortical neurons have the receptive fields they do, why and how the Gestalt segregation of figure from ground occurs, and why the striate and extrastriate visual cortex are organised the way they are. We now need to look more closely at Shepard's ideas to see how they are related to the adaptation of pre-perceptual mechanisms, through evolution and experience, to handle statistical regularities in the input with improved effectiveness. The next section argues that such adaptations could make perception expert in handling the images of moving objects, and that Shepard's idea of the internalisation of kinematic geometry, which emphasises non-probabilistic geometric rules instead of statistical regularities, is too vague to describe the process and does not explain its advantages.

11.0 The problems of internalisation and kinematic geometry

Shepard's choice of the word "internalisation" is curious. In reading his original article it was not quite clear whether he thought this was straightforward evolutionary adaptation or not, but in his later (1994) article he clarifies this point by mentioning genes in the opening sentence, though he still frequently refers to internalisation. Now to understand the process as an evolutionary adaptation it is not sufficient just to copy the regularity internally, which is what the term internalisation implies. In addition the regularity must be turned to some advantage, for without this the mechanisms would have no survival value. This is obvious in an example Shepard uses himself - diurnal rhythms - for a diurnal animal exploits the rhythm to become active by day and sleep at night while a nocturnal animal does the reverse, but both can be described as internalisation of the rhythm. In that case only the appropriate phase has to be found in order to gain advantages, but for evolutionary adaptation to other environmental regularities the mechanisms required to gain advantages are likely to be more complicated and much less obvious. Let us try to apply this to Shepard's experiments.

When an object is shown successively in two positions, subjects experience it moving along a path between these positions, and Shepard claims that his experiments show the path to be close to that dictated by Chasles' rule. There are three problems here. First, Chasles' rule provides a concise way of describing how

a 3-D object can move from the first to the second position, but it does not say that the object has to move along the path corresponding to simultaneous translation and rotation, as Todorovi'c (2000) explains in greater detail. As will be described below, it would be advantageous to have mechanisms adapted to respond to the types of motion that actually occur in the features of moving objects, and such mechanisms would be predicted on the redundancy exploitation hypothesis; but it is hard to see any basis for expecting these adaptations to correspond to the internalisation of Chasles' rule, because this does not necessarily specify the motions that actually occur.

Second, it is not clear whether the subject's judgements of intermediate positions are accurate enough to distinguish Shepard's predictions from ones based on mechanisms that respond to the different types of motion that occur frequently in the images of moving objects, as the redundancy exploitation hypothesis predicts.

The third problem is even more basic: we do not understand the neural basis for subjective experiences of moving objects, so it is risky to try to relate the experience to mechanism. To make sense of Shepard's claim we would have to accept a framework of the following sort. Assume the views of 3-D objects are represented in our brains by symbols for the coordinate values of prominent features in those objects. The views at the two positions would create two such representations, and we need to assume that the brain can interpolate between the corresponding features in these two representations to create representations of intermediate values. Shepard's claim would be that these interpolated representations correspond to positions along the screw transformation path that connects the two seen positions. With such a model the predictions of the hypothesis can at least be clearly stated, but it is a most implausible model because it assumes that experience is based on a temporal sequence of static representations, whereas we know there are neurons that represent movement: Where do these fit in? And how could the interpolated representations of static positions be formed early in the movement, before the object has appeared at its second position?

This model ignores modern knowledge of the neurophysiology of sensory

systems and is very unconvincing, but one must have some model before one can make predictions. To go to other extremes, it is clear that there are some ways of representing image information in the brain for which it would be meaningless to talk of applying the rules of kinematic geometry. What if the 3-D object is described in our brains as "like a carrot bent in the middle", or in terms of the muscle activations required to place our finger on its various features? How could one apply the rules to these representations? The way perceptions are represented in our brains is far from settled, but the rules of kinematic geometry could not be applied to some possibilities.

Now consider what advantages could be derived from having mechanisms evolutionarily adapted for signalling the motions of moving objects, for we need to understand these advantages if we are to relate the mechanisms to evolutionary survival. The results of van Hateren & van der Scharff (1998), briefly described above, suggest that there are cortical neurons that respond selectively to the spatial patterns that occur commonly in static natural images. If there are also ones that respond selectively to the spatio-temporal patterns of moving features in the images of moving objects, this would bring definite advantages. First, these neurons would act as matched filters for these patterns, and would therefore be optimal for detecting them at low signal-to-noise ratios, or as early as possible in the course of the movement. Second, they have the potential for extrapolating these movements into the future, that is for signalling a spatio-temporal pattern in its early stages, before it has been completed. And third, they could also interpolate, that is signal the whole motion when only its first and last parts were actually visible; this capacity is obviously important in considering the interpretation of Shepard's experiments. Analogous advantages might possibly be obtained by applying kinematic geometry, but Shepard does not suggest these potential benefits. Furthermore, as Todorovi'c (2000) points out, moving objects do not necessarily follow the helical path described by Chasles' rule, so its predictions might be misleading; the ordinary redundancy exploitation hypothesis makes more sense.

12.0 Geometric rules, the rigidity assumption, and experience

Shepard bravely proposed that a geometric rule, not a statistical regularity, is internalised, and that this leads directly to predictable subjective experiences. Sinha and Poggio (1996) have recently described experiments that show how subjective experience is influenced by interactions between the mathematical rules of perspective transformations, recent experience of particular motions, and a tendency to interpret motions that occur together as motions of a rigid object. Though the details are quite different, the important factors in these experiments are sufficiently similar to those in Shepard's experiments for their results to be relevant here.

Their subjects look for a few minutes at the computer-simulated silhouette of a 3-D figure made of straight wire segments being rotated to and fro about a horizontal axis. This is normally perceived as a rigid 3-D body being rotated, and if tested with the same figure and the same motions within a few minutes, the impression of rigid rotation is retained. But if instead a silhouette that is identical at its mid-position is moved in a way corresponding to a wire frame figure that has a different 3-D shape, then the subjects frequently perceive non-rigid motion: that is the object appears to deform as it moves. If those same motions had been seen without the training experience, they would have been perceived as the rigid rotation of a different 3-D body. Therefore the subjects certainly remember or internalise something as a result of their initial experience, but it is not the laws of kinematic geometry: it is that particular 3-D shape of the wire-frame object seen in the initial adapting experience that would, when rotated, generate the set of images that was actually experienced.

I think Helmholtz would have been delighted by this experiment, not only because it illustrates so well the relation between his "apperception" and "remembered experience", but also because it brings out something that would have been new to him. It is the initial assumption of rigidity that makes it logically possible to infer a 3-D shape from the rotating image, and this gives new insight into the role of such a "default assumption" in perception. Furthermore this type of assumption is presumably genetically determined and is the consequence of evolutionary selection, both of which would I think have been further new ideas for Helmholtz.

If I have understood Shepard correctly, he thinks that perceptions somehow embody as a whole the regularities constituting kinematic geometry. In contrast, this experiment shows how the laws of perspective transformation are used, together with the rigidity assumption, to form one particular detail of the percept, namely its 3-D shape. This fits the ideas reviewed in previous sections of this article: regularities in the motions of the wire segments are detected and used to construct a rigid 3-D shape that is compatible with them and, when possible, the rigidity assumption; it is this shape that we experience and which influences subsequent interpretations of motion. A hierarchy of operations occurs in the visual pathways, and at least for the early ones the evidence is now strong that they conform to the principle of exploiting redundancy. In this light, the idea that perception has internalised the rules of kinematic geometry seems vague and implausible. Furthermore it is doubtful if Shepard's experimental tests of his specific idea distinguish it from the more general hypothesis about exploiting statistical regularities.

13.0 Conclusions

The principle that the redundancy in sensory messages resulting from regularities in the environment are exploited in sensory pathways illuminates a host of sensory phenomena, such as accommodation, light and dark adaptation, lateral inhibition, the form of feature detectors in the cortex, their relation to the Gestalt laws, the organisation of extrastriate areas, the functional role of figural and contingent after effects, and possibly the nature of intelligence itself. The principle of adaptation to regularities has a very respectable past, it is a fertile inspiration for current research, and looks set for a prosperous future. But Shepard has certainly drawn attention to an interesting phenomenon, and there may be an important lesson for neurophysiologists to learn from it.

Shepard claimed that when an object is presented first in one position, then another, "...one tends to experience that unique, minimum, twisting motion prescribed by kinematic geometry". There may be doubts about the role of kinematic geometry, but there is certainly a rotary component to the motion experienced, and once it has been pointed out it is clear that such rotations must

play an important part in the interpretation of the images of moving objects. Furthermore these rotations need to be tightly localised to particular image features when using them in this way, and V1 is the high resolution area of the visual cortex. Hence one must ask "Do neurons in V1 detect rotation directly?" A recent analysis of their responses at different delays after the presentation of an oriented stimulus found that, for some of them, the favoured orientation does in fact change with the delay (Ringach et al, 1997), so the optimum stimulus would be a twisting motion. Since Hubel and Wiesel we have known that V1 neurons signal the orientation of edges in the visual field: Do they also signal change of orientation with time? This hint that some V1 neurons are tuned to twisting movements, not pure translations, urgently needs to be followed up.

Acknowledgment.

The author was supported by Grant No.046736 from the Wellcome Trust during the preparation of this paper.

References

- Atick, J.J. (1992) Could information theory provide an ecological theory of sensory processing? *Network*, 3 213-251.
- Atick, J.J. & Redlich, A.N. (1990) Mathematical-model of the simple cells of the visual-cortex. *Biological Cybernetics*, 63 99-109.
- Atick, J.J. & Redlich, A.N. (1992) Convergent algorithm for sensory receptive field development. *Neural Computation*, 5 45-60.
- Attneave, F. (1954) Informational aspects of visual perception. *Psychological Review*, 61 183-193.
- Baddeley, R.J. (1996) An efficient code in V1? *Nature (London)*, 381 560-561.
- Baddeley, R., Abbott, L.F., Booth, M.C.A., Sengpiel, F., Freeman, T., Wakeman, E.A. & Rolls, E.T. (1997) Responses of neurons in primary and inferior temporal

visual cortices to natural scenes. *Proceedings of the Royal Society, Series B*, 264 1775-1783.

Barlow, H.B. (1959). Sensory mechanisms, the reduction of redundancy, and intelligence. In *The Mechanisation of Thought Processes* (pp. 535-539). London: Her Majesty's Stationery Office.

Barlow, H.B. (1961). The coding of sensory messages. Chapter XIII. In *Current Problems in Animal Behaviour*, Thorpe and Zangwill (Eds), Cambridge University Press, pp. 330-360.

Barlow, H.B. (1972) Single units and sensation: a neuron doctrine for perceptual psychology? *Perception*, 1 371-394.

Barlow, H.B. (1974). Inductive inference, coding, perception and language. *Perception* 3,123-134.

Barlow, H.B. (1981) Critical limiting factors in the design of the eye and visual cortex. The Ferrier lecture, 1980. *Proceedings of the Royal Society, London, B*, 212 1-34.

Barlow, H.B. (1983) Intelligence, guesswork, language. *Nature*, 304 207-209.

Barlow, H.B. (1989) Unsupervised learning. *Neural Computation*, 1 295-311.

Barlow, H.B. (1990). A theory about the functional role and synaptic mechanism of visual after-effects. In C.B.Blakemore (Eds.), *Vision: coding and efficiency* Cambridge, U. K.: Cambridge University Press.

Barlow, H.B. (1995). The Neuron Doctrine in Perception. In M. Gazzaniga. (Eds.), *The Cognitive Neurosciences* (pp. Chapter 26 pp 415-435). Cambridge, Mass: MIT Press.

Barlow H B (1996) Banishing the Homunculus. pp 425-450 In D. Knill & W. Richards (Eds.), *Perception as Bayesian Inference* Cambridge: Cambridge University Press.

Barlow, H.B. & Tripathy, S.P. (1997). Correspondence noise and signal pooling as

factors determining the detectability of coherent visual motion. *Journal of Neuroscience* 17, 7954-7966

Barrow, H.G. (1987). Learning receptive fields. In *IEEE First Annual Conference on Neural Networks*, 4 (pp. 115-121).

Bell, A.J. & Sejnowski, T.J. (1995) An information maximisation approach to blind separation and blind deconvolution. *Neural Computation*, 7 1129-1159.

Brunswik, E. (1956). *Perception and the representative design of psychological experiments*. Berkeley: University of California Press.

Brunswik, E. & Kamiya, J. (1953) Ecological cue-validity of "proximity" and of other Gestalt factors. *American Journal of Psychology*, 66 20-32.

Carandini, M., Barlow, H.B., O'Keefe, L.P., Poirson, A.B. & Movshon, J.A. (1997) Adaptation to contingencies in macaque primary visual cortex. *Proceedings of the Royal Society, Series B*, 352 1149-1154.

Chater, N. (1996) Reconciling simplicity and likelihood principles in perceptual organisation. *Psychological Review*, 103 566-581.

Craik, K.J.W. (1943). *The Nature of Explanation*. Cambridge: Cambridge University Press.

Elder, J.H. & Goldberg, R.M. (1998). The statistics of natural image contours. In *IEEE Workshop on Perceptual Organisation in Computer Vision 1998*.

Field, D.J. (1987) Relations between the statistics of natural images and the response properties of cortical cells. *Journal of the Optical Society of America A*, 4 2379-2394.

Field, D.J. (1994) What is the goal of sensory coding? *Neural Computation*, 6 559-601.

Gardner-Medwin, A.R. & Barlow, H.B. (2000) The limits of counting accuracy in distributed neural representations. *Neural Computation* (accepted).

Helmholtz, H.von (1925) *Physiological Optics. Volume III. The Theory of the*

- Perceptions of vision (Translated from 3rd German Edition, 1910)*. Ch 26. Washington: Optical Society of America. (Also Dover Publications, New York, 1962)
- Hubel, D.H. & Wiesel, T.N. (1959) Receptive fields of single neurones in the cat's striate cortex. *Journal of Physiology*, 148 574-591.
- Hyvarinen, A. & Oja, E. (1996) Simple neuron models for independent component analysis. *International Journal for Neural Systems*, 7 671-687.
- Lennie, P. (1998) Single units and visual cortical organization. *Perception*, 27 889-935.
- Linsker, R. (1986a) From Basic Network Principles to Neural Architecture - emergence of spatial opponent cells. *Proceedings of the National Academy of Sciences, USA*, 83 7508-7512.
- Linsker, R. (1986b) From Basic Network Principles to Neural Architecture - emergence of orientation selective cells. *Proceedings of the National Academy of Sciences, USA*, 83 8390-8394.
- Linsker, R. (1986c) From Basic Network Principles to Neural Architecture - emergence of orientation columns. *Proceedings of the National Academy of Sciences, USA*, 83 8779-8783.
- Mach, E. (1886). *The Analysis of Sensations, and the Relation of the Physical to the Psychological (Translation of the 1st, revised from the 5th, German Edition by S. Waterlow)*. Chicago and London: Open Court (Also Dover reprint, New York 1959).
- MacKay, D.J.C. (1992) Bayesian interpolation. *Neural Computation*, 4 415-447.
- Mackintosh, N.J. (1983). *Conditioning and associative learning*. Oxford: Oxford University Press.
- Mumford, D. (1996). Pattern theory: a unifying perspective. In D. C. Knill & W. Richards (Eds.), *Perception as Bayesian inference* (pp. 25-62). Cambridge: Cambridge University Press.

- Olshausen, B.A. & Field, D.J. (1996) Emergence of simple-cell receptive-field properties by learning a sparse code for natural images. *Nature*, 381 607-609.
- Pearson, K. (1892). *The Grammar of Science*. London: Walter Scott.
- Rescorla, R.A. & Wagner, A.R. (1972). A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and non-reinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: current research and theory* (pp. 64-99). New York: Appleton-Century-Crofts.
- Ringach, D.L., Hawken, M.J. & Shapley, R. (1997) Dynamics of orientation tuning in macaque primary visual cortex. *Nature(London)*, 387 281-284.
- Ruderman, D.L. (1994) Statistics of natural images. *Network: computation in neural systems*, 5 517-548.
- Ruderman, D.L. (1997) Origins of scaling in natural images. *Vision Research*, 37, 3385-3398.
- Shannon, C.E. & Weaver, W. (Ed.). (1949). *The mathematical theory of communication*. Urbana: Univ. Illinois Press.
- Shepard, R.N. (1984) Ecological constraints on internal representation: Resonant kinematics of perceiving, imagining, thinking and dreaming. *Psychological Review*, 91 417-447.
- Shepard, R.N. (1994) Perceptual-cognitive universals as reflections of the world. *Psychonomic Bulletin & Review*, 1 2-28.
- Sinha, P. & Poggio, T. (1996) Role of learning in three-dimensional form perception. *Nature*, 384 460-463.
- Solomonoff, R.J. (1964a) A formal theory of inductive inference. Part I. *Information and Control*, 7 1-22.
- Solomonoff, R.J. (1964b) A formal theory of inductive inference. Part II. *Information and Control*, 7 224-254.
- Todorovi'c, D. (2000). Is kinematic geometry an internalised regularity? *BBS -this*

volume.

Tolhurst, D.J., Tadmor, Y. & Chao, T. (1992) Amplitude spectra of natural images. *Ophthalm. Physiol. Opt.*, 12 229-232.

Tolman, E.C. (1948) Cognitive maps in rats and men. *Psychological Review*, 55 189-208.

van Hateren, F.H. & van der Schaaf, A. (1998) Independent component filters of natural images compared with simple cells in primary visual cortex. *Proceedings of the Royal Society, Series B*, 265 359-366.

Wallace, C.S. & Boulton, D.M. (1968) An information measure for classification. *Computing Jopurnal*, 11 185-195.

Walls, G. (1942). *The vertebrate eye and its adaptive radiation*. Bloomfield Hills, Michigan: Cranbrook Institute of Science.

Watanabe, S. (1960) Information-theoretical aspects of Inductive and Deductive Inference. *I.B.M. Journal of Research and Development*, 4 208-231.

Webber, C., J. StC. (1991) Competitive learning, natural images, and the self-organisation of cortical cells. *Network - computation in neural systems*, 2 169-187.