

MOVEMENTS OF THE EYES

2ND EDITION
REVISED AND ENLARGED

P

PION · LONDON

RHS CARPENTER

P Pion Limited, 207 Brondesbury Park, London NW2 5JN

© 1988 Pion Limited

First published 1977

Second revised edition 1988

All rights reserved. No part of this book may be reproduced in any form by photostat microfilm or any other means without written permission from the publishers.

ISBN 0 85086 109 8

Preface to the second edition

A great deal seems to have happened in ten years. The 1200 or so new entries in the bibliography testify to one aspect of this, and not necessarily a very pleasing one. What *is* gratifying is the growing sense amongst neurophysiologists and clinicians that oculomotor physiology is not some kind of obscure and overspecialised academic backwater, but promises to hold the key to many of the most central problems in understanding the brain. Above all, it is an area where we seem to be *winning*. The inputs and outputs of the oculomotor system are circumscribed and easy to measure, and allow us to regard it *quantitatively* in a way that is scarcely possible elsewhere: in many cases, a few quick measurements of eye movements can provide as accurate a neurological diagnosis as a brain scan. Within this system the experimenter can find all the big issues of neural science—sensory processing, motor coordination, learning, attention—in a form compact enough to hope to get to grips with them. It is, in short, a microcosm of the brain itself.

It is a pleasure to thank those who have assisted with this second edition, particularly Mrs D Hughes (figures 7.14 and 9.16), Mr Peter Starling for his photographic expertise, and the editors, Dr Maggie Barker and Dr Jan Schubert for their herculean labours in bringing it to publication.

R H S Carpenter

Gonville and Caius, Cambridge

April 1988

It is not hard to see why there has recently been such an upsurge of interest in the systems controlling eye movements. The more we find out about the way in which our eyes seek out and capture objects of interest in the visual world, the more remarkable does its seemingly effortless precision and sophistication appear. At the same time—and this is perhaps why it is increasingly attracting the attention of neurophysiologists—it functions at a level of complexity somewhere between the banality of the spinal reflex and the inscrutability of the voluntary act. One may therefore hope that an understanding of how the eyes are controlled may help us when we seek to understand more complicated motor systems: and indeed it is now becoming apparent that many of what were formerly thought to be unique properties of the eye movement control system have direct parallels even in a system as complex as the control of the hand.

This book has been planned with three classes of reader in mind. First, the medical student who wishes to know rather more about the subject—one that has considerable importance as an aid to neurological diagnosis—than is available to him in more general textbooks. Second, the scientifically oriented physiology or psychology student who is interested in a field in which the application of techniques of systems analysis to neurophysiology has proved particularly fruitful; and lastly but perhaps chiefly, to the research worker in this field who, like the author himself, may well have felt a need for an up-to-date and not wholly clinical reference book that can also be used as a teaching text.

It is divided into three parts, in an attempt to reflect the clear distinction that ought to be made with physiological systems between *performance*, *structure*, and the notional *models* that are supposed to link structure to performance, and performance to function. Here, part 1 is intended as a description of the phenomena of eye movements, part 2 as an account of the anatomy of the system (including the mechanical properties of the eye ball itself), while part 3 attempts to relate this anatomy to the control mechanisms, and to assess the contribution of eye movements to visual function. Appendices to the text provide a review of techniques for measuring eye movements, and an introduction to linear systems analysis which it is hoped will enable the nonmathematician to follow the mathematical arguments presented in the text.

I wish to express my gratitude to the authors quoted in the text for the permissions received to reproduce material from their papers and books, and to the following publishers and organisations for allowing me to use their copyright material:

- Academic Press Inc, New York (Bach-y-Rita, Collins, and Hyde, Editors: *The Control of Eye Movements: Experimental Neurology*)
- Almqvist and Wiksell, Stockholm (*Acta Oto-Laryngologica*)
- The American Institute of Physics, New York (*Journal of the Optical Society of America*)
- The American Physiological Society, Bethesda, Maryland (*American Journal of Physiology; Journal of Neurophysiology*)
- The American Psychological Association, Washington, DC (*Journal of Comparative Psychology*)
- Edward Arnold (Publishers) Ltd, London (Matthews: *Mammalian Muscle Receptors and their Central Actions*)
- J F Bergmann-Verlag, Munich (*Experimental Brain Research*)
- Butterworth and Co (Publishers) Ltd, Sevenoaks, Kent (Roberts: *Neurophysiology of Postural Mechanisms*)
- Elsevier/North-Holland Biomedical Press, Amsterdam (*Archives Néerlandaises de Physiologie, Brain Research*)
- The Institute of Electrical and Electronics Engineers, Inc, New York (*IEEE Transactions on Human Factors in Electronics, IEEE Transactions on Systems and Sciences of Cybernetics*)
- S Karger AG, Basel (*Bibliotheca Ophthalmologica; Oto-Rhino-Laryngology*)
- Henry Kimpton Ltd, London (Duke-Elder: *System of Ophthalmology*)
- E and S Livingstone Ltd, Edinburgh (*Quarterly Journal of Experimental Physiology*)
- The C V Mosby Co, St Louis, Missouri (Moses, Editor: *Adler's Physiology of the Eye*)
- The New York Academy of Sciences, New York, NY (*Annals of the New York Academy of Sciences*)
- The Optical Society of America, Lancaster, Pennsylvania (*Journal of the Optical Society of America*)
- Pergamon Press Ltd, Oxford (Lennerstrand and Bach-y-Rita, Editors: *Basic Aspects of Ocular Motility; Vision Research*)
- The Physiological Society, Cambridge and London (*Journal of Physiology*)
- Plenum Publishing Corporation, New York (Yarbus: *Eye Movements and Vision*)
- The Psychonomic Society, Austin, Texas (*Perception and Psychophysics*)
- The Royal Society, London (*Proceedings of the Royal Society*)
- The Royal Society of Medicine, London (*Proceedings of the Royal Society of Medicine*)
- Springer-Verlag, Heidelberg (*Experimental Brain Research; Kybernetik*)
- Taylor and Francis Ltd, London (*Optica Acta*)
- The Wistar Institute Press, Philadelphia, Pennsylvania (*Anatomical Record; Journal of Comparative Neurology*)

In memoriam
D E C

Part 1 Function

1	The use of eye movements	
1.1	Consequences of velocity blur	1
1.2	Consequences of a limited field of view	5
1.3	Consequences of binocular vision	9
2	Vestibular eye movements	
2.1	The vestibular sense organs	12
	2.1.1 Gross anatomy	12
	2.1.2 The receptor cells and their innervation	14
	2.1.3 Electrical responses to angular accelerations	17
	2.1.4 Dynamic properties of the semicircular canals	19
	2.1.5 Electrical responses to linear accelerations	23
2.2	Static vestibulo-ocular responses	25
2.3	Dynamic vestibulo-ocular responses	29
	2.3.1 Vestibular nystagmus	29
	2.3.2 Natural head movements	34
	2.3.3 The nature of the quick phase	36
3	Optokinesis and smooth pursuit	
3.1	Optokinesis versus pursuit	39
3.2	Optokinesis	40
	3.2.1 Dynamic properties	43
	3.2.2 Nonlinearities	49
	3.2.3 The effect of field size	50
	3.2.4 The quick phase	52
	3.2.5 Asymmetries	53
3.3	Smooth pursuit	55
	3.3.1 What is being pursued?	56
	3.3.2 Responses to repetitive stimuli	58
	3.3.3 Open-loop dynamics	63
	3.3.4 Artificial feedback	65
	3.3.5 Smooth control of position	66
4	Saccades	
4.1	The time course of saccades	69
	4.1.1 Amplitude-velocity-duration relationships	70
	4.1.2 The termination of saccades	74
	4.1.3 Latency	76
4.2	The control of saccades in relation to the visual stimulus	83
	4.2.1 Correction saccades	83
	4.2.2 Saccadic intermittency	86
	4.2.3 Saccades to moving targets	94
	4.2.4 Selection of the saccadic goal	96

5	Vergence	
5.1	Sensory stimuli for vergence movements	103
	5.1.1 Disparity vergence	104
	5.1.2 Accommodation vergence	108
5.2	Dynamics of vergence movements	111
	5.2.1 Natural viewing	111
	5.2.2 Open-loop studies	113
5.3	Mixed vergence and version	120
6	Miniature movements	
6.1	General characteristics of miniature eye movements	124
6.2	Tremor	126
6.3	Drift	127
	6.3.1 The randomness of drift in the dark	128
	6.3.2 Slow visual correction of drift	130
6.4	Microsaccades	131
6.5	Statistical variation in the point of fixation	135
	6.5.1 The stationarity of fixation	137
	6.5.2 Stability of the gaze under natural conditions	137
Part 2	Structure	
7	Statics and dynamics of the eye	
7.1	Terminology	139
	7.1.1 Optical terminology	139
	7.1.2 Kinematic terminology	140
	7.1.3 Visual terminology	141
7.2	Measurements of the eye's functional geometry	142
	7.2.1 Determination of the optical and pupillary axes	142
	7.2.2 Determination of centres of rotation	145
7.3	Kinematics of eye rotations	147
	7.3.1 Stereographic representation of eye movements	148
	7.3.2 Torsion and Listing's Law	151
7.4	The action of the eye muscles	156
	7.4.1 Gross anatomy of the eye muscles	156
	7.4.2 Lines of action of the muscles	159
	7.4.3 Actual muscular cooperation	163
7.5	Kinetic properties of the eye	165
	7.5.1 The basic mechanical model	166
	7.5.2 Measurements of the mechanical properties	169
	7.5.3 Static behaviour	171
	7.5.4 Transient behaviour	174
	7.5.5 Behaviour during saccades	178

8	The extraocular muscles and their innervation	
8.1	The structure of the extraocular muscle fibres	182
	8.1.1 Gross anatomy of the muscles	182
	8.1.2 Types of motor fibre	183
	8.1.3 Sensory endings in extraocular muscles	187
8.2	The physiology of extraocular muscles	189
	8.2.1 Electrical properties of motor fibres	189
	8.2.2 Contractile properties of motor fibres	190
	8.2.3 Responses from stretch receptors	191
8.3	The nerve supply to extraocular muscles	193
	8.3.1 Efferent pathways	193
	8.3.2 Afferent pathways	195
8.4	The oculomotor nuclei	197
	8.4.1 The structure of the eye muscle nuclei	197
	8.4.2 Patterns of efferent discharge	200
	8.4.3 Stimulation experiments	207
9	Central eye movement pathways	
9.1	The visual proprioceptive nuclei	210
	9.1.1 Anatomy and connections	210
	9.1.2 Optokinetic functions	212
9.2	Vestibular nuclei	214
	9.2.1 The primary vestibular projection	214
	9.2.2 Other afferents to the vestibular nuclei	216
	9.2.3 Efferent pathways	218
	9.2.4 Neuronal activity	221
	9.2.5 The medial longitudinal fasciculus	229
9.3	The preculomotor nuclei of the brainstem	231
	9.3.1 The reticular formation: anatomy	231
	9.3.2 The reticular formation: saccade generation	233
	9.3.3 The accessory oculomotor nuclei	238
9.4	The cerebellum	240
	9.4.1 Structure	240
	9.4.2 Vestibular afferent systems	245
	9.4.3 Visual afferent systems	245
	9.4.4 Other afferents	249
	9.4.5 Efferent pathways	250
	9.4.6 Activity related to eye movements	251
	9.4.7 Electrical stimulation	255
	9.4.8 Lesions	259
9.5	Superior colliculus	263
	9.5.1 Anatomy and connections	263
	9.5.2 Activity related to eye movements	266
	9.5.3 Electrical stimulation	271
	9.5.4 Oculomotor function	275

9.6	Higher oculomotor areas	276
9.6.1	Basal ganglia	276
9.6.2	Thalamus	279
9.6.3	Cerebral cortex: anatomy	281
9.6.4	Cortical oculomotor functions	284
9.7	Summary of central oculomotor pathways	291
Part 3 The system as a whole		
10 The proprioceptive input		
10.1	Stretch reflexes	294
10.1.1	General considerations	294
10.1.2	Mechanical studies	298
10.1.3	Electrical recordings	300
10.1.4	The possibility of an intrinsic stretch reflex	303
10.2	Sensation	304
10.2.1	Conscious sensation	305
10.2.2	Cancellation of the visual effects of eye movements	306
10.2.3	Directing the eyes in the dark	310
10.3	Other possible functions	311
10.3.1	The quick phase of nystagmus	311
10.3.2	Integrative feedback	311
11 Eye movements and vision		
11.1	Saccadic suppression	313
11.1.1	Basic observations	313
11.1.2	Possible mechanisms	317
11.1.3	The lateral geniculate nucleus and other central areas	320
11.2	Miniature eye movements and vision	322
11.2.1	Methods of stabilising retinal images	322
11.2.2	The appearance of stabilised images	324
11.2.3	Movements of stabilised images	325
11.2.4	Are miniature eye movements useful?	327
11.3	Eye movements and spatial sense	330
11.3.1	The sense of direction and distance	330
11.3.2	The sense of motion	336
12 Towards a synthesis		
12.1	The final common path	343
12.2	The oculomotor integrator	346
12.2.1	The ubiquity of velocity control	346
12.2.2	Possible mechanisms of neural integration	349
12.2.3	Where is the integrator?	354
12.3	The velocity store	355
12.3.1	Models of velocity storage	355
12.3.2	Possible neuronal mechanisms	359

12.4	The saccadic pulse generator	360
12.4.1	The transformation from image to amplitude	360
12.4.2	The one-dimensional bang-bang model	362
12.4.3	Oblique saccades	366
12.4.4	Interactions with other movements	367
12.5	Saccadic procrastination	370
12.5.1	Saccadic delay	370
12.5.2	A model of target selection	371
13 Oculomotor adjustments		
13.1	Adaptation and adjustment	374
13.2	Matching the oculomotor plant	377
13.2.1	Dynamics	377
13.2.2	Nonlinearities	378
13.3	Gaze-holding	379
13.3.1	Vestibular damage	379
13.3.2	Habituation	380
13.3.3	Visual adjustment of vestibular responses	382
13.3.4	The development of gaze-holding	387
13.4	Gaze-shifting	388
13.4.1	Saccades	388
13.4.2	Smooth pursuit	390
13.4.3	Vergence	391
13.5	Neural mechanisms	393
13.5.1	Models of cerebellar plasticity	394
13.5.2	The cerebellar contribution to oculomotor adjustment	396
13.6	Sensory aspects of plasticity	399
13.7	Envoi	401

Appendices**Appendix 1 Methods of measuring eye movements**

A1.1	Direct viewing	405
A1.1.1	Photoelectric viewing	406
A1.2	Devices using reflection	413
A1.2.1	Reflection from the cornea	413
A1.2.2	Reflection from attachments to the eye	418
A1.3	Electrooculography	420
A1.4	Electromagnetic recording	421
A1.5	Other methods	424

Appendix 2 A brief introduction to systems analysis

A2.1	Functions from the point of view of the analyst	427
A2.2	Classes of function	431
A2.3	Some common functions	434
A2.4	Imaginary numbers	440

Function

A2.5	Operators and filters	444
	A2.5.1 Differentiation and integration	445
	A2.5.2 Transient functions	446
	A2.5.3 Derivatives of the elementary functions	448
	A2.5.4 Frequency transfer functions	449
A2.6	The linear first-order building blocks	453
A2.7	Linear second-order systems	456
A2.8	Transfer functions of actual systems	458
A2.9	Control systems and feedback	465
	A2.9.1 Transfer function of a system with a loop	465
	A2.9.2 Ballistic control systems	466
	A2.9.3 Feedback control systems	467
	A2.9.4 Adaptive control	470
Appendix 3 Oculomotor nomenclature		
A3.1	Symbols	472
	A3.1.1 The main oculomotor variables	472
	A3.1.2 Sign conventions	474
	A3.1.3 Other variables	475
	A3.1.4 Units	475
	A3.1.5 Terminology for saccadic tasks	475
A3.2	Abbreviations	477
	A3.2.1 General	477
	A3.2.2 Anatomical	478
References		481
Index		577

The use of eye movements

"The muscles were of necessitie provided and given to the eye, that so it might move on every side: for if the eye stode fast, and immoveable, we should be constrained to turne our head and necke (being all of one peece) for to see: but by these muscles it now moveth it selfe with such swiftnes and nimblenes, without stirring of the head, as is almost incredible..."⁽¹⁾

The eye is at once the master and the slave of vision. The six pairs of muscles that move our eyes dictate absolutely what we see; yet these muscles are also kept incessantly busy ensuring that the picture of the outside world that falls on the retina moves across it neither too little nor too much relative to the receptors, in order to comply with the demands of the visual system. The requirements of sight determine the major classes of eye movement: in this first chapter, we examine what kinds of movements are made and how they contribute to better vision, a matter considered in more detail in chapter 11. In the remaining chapters of part I, these classes of movement are each in turn discussed in more detail: but for the moment only the most general considerations will be presented.

1.1 Consequences of velocity blur

For a stationary object, the fineness of the detail that can just be resolved is almost entirely a function of such 'physical' factors as the size of the eye, the quality of its optics, the spacing of the receptors, and the degree to which the central pathways from them overlap in their connections. But as soon as the image of the object starts to move across the retina quite different factors come into play. The light falling on any particular receptor is now no longer constant, but fluctuates as different parts of the image pass over its receptive area. The fidelity of the pattern of activity in the receptors to the pattern of light and shade in the original image at any moment will be a function not just of the physical factors mentioned earlier, but also of how good the receptors are at *following* these fluctuations. It turns out that they are actually rather bad at it. Electrical responses in the retina to brief visual stimulation under the most favourable conditions typically show a time course of the order of tens of milliseconds, and flickering lights cannot be seen as such when the frequency of their flicker is more than some 60–80 Hz. The consequent degradation of visual acuity is quite striking: from appropriate published data (for example Green and Campbell, 1965; Burr and Ross, 1982) one may calculate that a retinal movement as slow as 1° s^{-1} —that is, in which any point on the target will take three minutes to cross the visual field—has roughly the same deleterious effect on grating resolution as nearly two diopters of myopia. More precisely expressed, the visual system has a band-pass characteristic both in the frequency domain

⁽¹⁾ The quotations at the head of each chapter are from Andreas Laurentius (1599).

(peaking around 10 Hz) and in the spatial domain [for us, around 4 cycles deg^{-1} , but an order of magnitude less in the cat (Campbell and Maffei, 1973)]. It follows that for a grating of a particular frequency there is a particular retinal image velocity at which the image will lose the least contrast during its transmission to the brain: for a high-frequency target this velocity will be relatively low, whereas a low-frequency target may actually be seen *better* if it moves quite rapidly (Burr and Ross, 1982; Kelly, 1985). Perhaps as a consequence, *pseudoacuity* tasks that involve a component of localisation—such as vernier acuity targets, or Landolt C's—may be affected surprisingly little by quite large amounts of retinal slip (Westheimer and McKee, 1975).

However, there is more to vision than *detection*, and one would certainly expect the visual system to have more difficulty in *recognising* objects whose retinal images are moving rapidly across the retina. At all events, it is clear that to see well we need an image that is moving neither too fast nor too slow. It turns out that the latter requirement appears to be sufficiently met by the *noise* inherent in any biological system: however hard we try to fixate an object, there are random movements of the eye and of the head that prevent the image from being absolutely stationary on the retina. These *miniature* or *fixational movements* are discussed in chapter 6, and their influence on vision is considered in chapter 11.

The need to prevent the image from moving too fast does, however, demand some positive action from the oculomotor system. It seems in fact that the first kinds of eye movements to evolve were *gaze-holding movements*, designed, paradoxically, not to move the gaze at all but to keep it still—still, that is, relative to the frame of reference provided by the outside world (see, for example, Walls, 1962). One can easily imagine a feedback system that would carry out this function automatically (figure 1.1). The visual system is richly endowed with directionally selective neurons, cells




Figure 1.1 How feedback of retinal slip can be used to stabilise the retinal image: optokinesis.

that are stimulated when the retinal image moves in a particular direction. Harness these cells to the eye muscles, so that retinal slip in a particular direction drives the whole eye in the same direction, and we have a system that will make the eye *track* the visual stimulus, its image remaining relatively stationary on the retina. Movements of this sort are well known throughout the animal kingdom, having surprisingly similar response characteristics in creatures as different as the crab, with its twenty-six pairs of eye-muscles (Horridge, 1966; Burrows and Horridge, 1968), and man: similar mechanisms control the head movements of insects (Land, 1969; 1973; Kien, 1974). Such movements, in response to motion of a large part of the visual field, are called *optokinesis*: they form the subject of chapter 3.

Important as optokinesis is, it has one severe drawback, inherent in all purely visual systems: it is very *slow*. It takes a fifth of a second or more for the eyes to begin to respond to a sudden movement of the retinal image. Fortunately, there are alternative sources of information that are a good deal quicker. In the natural world there is only one thing that can generate retinal slip over a wide area, and that is *head movement*. An animal can therefore improve its retinal stabilisation by one of two strategies: by trying to keep its head stationary relative to the outside world, or alternatively by using information about its head movement to control the eye muscles.

It turns out that animals vary in the extent to which they use one or other of these strategies. On the whole, animals such as birds and frogs with light heads and mobile necks tend to keep their heads still, whereas those with stiff necks and heavy heads tend to compensate with their eyes (Grüsser and Grüsser-Cornehls, 1976; Dieringer and Precht, 1982; Berthoz, 1985). If a frog's head is restrained, its eye movements are larger (Ashton et al, 1984a). Another factor determining the use made of the head in improving vision is whether the eyes point predominantly sideways or frontwards. The biggest threat to an animal's visual stability is of course its own locomotion: if its eyes, like ours, point forward then movement ahead will have only second-order effects on the retinal image. But if, as in a bird or rabbit, the eyes point sideways, locomotion will induce massive retinal slip, of a kind that cannot be compensated for by *rotation* of the eyes, since the angular extent of the slip will in general depend on the distance of the object (figure 1.2, see over). The only satisfactory solution is to stabilise not the eye but the *head*, which is capable of translation as well as rotation. The result is the slightly comical gait of such birds as chickens and ducks, in which the body walks on while the head is temporarily left behind, to be jerked forward again at the next step (figure 1.3, see over) (Dunlap and Mowrer, 1931). But for creatures with heavy heads and short necks this is not a feasible solution, and an animal like the rabbit responds to linear accelerations of the head with rotations of the eyes that are approximately such as to stabilise the horizon (section 2.2). In man, as we shall see, responses of the eyes solely due to translation of the head have all but disappeared.




Figure 1.2 Forward movement of an animal with side-pointing eyes leads to different retinal shifts for objects at different distances.




Figure 1.3 Sequential frames (32 s^{-1}) of a film of a pigeon walking, showing 'nystagmus' of the head (reprinted from Dunlap and Mowrer, 1931).

The second strategy, that of using information about head movement to help in stabilising the retinal image, has an important advantage: *speed*. The most important source of this information is the *vestibular apparatus*, the system of fluid-filled tunnels and caverns within the thickness of the skull that signals rotational and translational accelerations of the head. In the next chapter we shall see that exceptionally fast reflexes connect this sensory apparatus to the neurons that supply the eye muscles—the *vestibulo-ocular reflexes*. Other sensory pathways that can be used, though probably not to any great extent in man, come from proprioceptors in the muscles and articulations of the neck (section 2.2). These respond to the position of the head relative to the body, rather than to its movement in space. All these mechanisms are considerably faster in operation than the 150 ms or so that it takes the eye muscles to respond to visual stimulation, as a simple demonstration makes clear (Robinson, 1968b). Hold your hand out at arm's length, and shake it from side to side through a small angle. Compare the clarity with which you see it with what happens when you keep the hand still, but move your head at the same frequency through the same angle: the improvement is striking. The mechanisms by which head movements influence eye movement are considered in chapter 2.

1.2 Consequences of a limited field of view

Different species vary considerably in how much they can see at any given moment (Walls, 1942; Rochon-Duvigneaud, 1943; Hughes, 1977). The horse, for example, can see almost all round its head (figure 1.4), and, given




Figure 1.4 Maximum horizontal monocular (shaded areas) and binocular (black areas) fields, directions of optical axes (arrows), and location of the *area centralis* (inside ring) in different species (data from Walls, 1942; Hughes, 1977; Martinoya et al, 1981; Martin, 1984).

an adequate system for stabilisation of the retinal image, it is not obvious that it would ever need to make any other kinds of eye movements at all. The reason that it does is that it cannot see *equally well* in all parts of its field. In birds and mammals we find that vision comes in two grades: a luxury grade, with high acuity and often colour discrimination, and an economy grade of low acuity—though often of greater sensitivity—lacking the added refinement of colour. In the first case, the fine acuity is mostly due to a dense packing of the receptors, and a ratio of the connections between each receptor and the fibres in the optic nerve that approaches one-to-one: in the second, the receptors spread themselves out a little more, and pool their messages before sending them to the brain, which results in lowered thresholds but poor acuity. In many cases the first type of vision is associated with an abundance of cones, the second with rods.

The reason why the whole visual field is not provided with the luxury grade of vision is presumably partly that under nocturnal conditions it cannot function properly because of its relative insensitivity, and partly that the dense packing of the receptors and the one-to-one connections to the optic nerve would lead to the latter becoming unmanageably large. For example, if our entire retina were of the high-quality type, the cross-sectional area of our optic nerve would have to increase by a factor of over two hundred (and no doubt the size of the blind spot would also have to increase proportionately). Thus in mammals we generally find that high-quality vision is confined to a special central zone in the retina, whose position and configuration depends on how the animal normally uses its eyes, and what it particularly wants to see well: there may be more than one such zone. Walls (1942) has discussed the configuration of visual fields and specialised areas of the retina in different species with a wealth of interesting detail; more recently, Hughes (1977) has revised and extended Walls' work. Broadly, one finds that predators tend to have a front-facing, compact *area centralis*, often associated with a local depression in the retina, the *fovea*, whereas those of their prey often form long horizontal *visual streaks* that correspond roughly with the situation of the horizon when the eye is at its normal rest position (horse, figure 1.4). However this rule has many exceptions: the mouse has an *area centralis*, the crocodile a visual streak.

Whatever its shape, the size of this area in an animal that is not wholly diurnal is only a small fraction of the total field. In man, one can estimate its functional extent by measurements of visual acuity at different distances from the fixation point (taken to be the centre of the fovea): some observations of this type are shown in figure 1.5. By the time 1° from the centre has been reached, visual acuity has fallen off by a factor of two or three: the intercone spacing and the density of ganglion cells show corresponding changes (Jones and Higgins, 1947; rabbit: Hughes, 1971; cat: Hughes, 1975; Stone, 1978). In other words, only some ten-thousandth part of the visual field is seen with full clarity: the reason that we are not greatly aware of our loss is that our eyes are continually on the move, pointing this little tunnel of clear vision

now here, now there, at whatever we want to see. It is important not to underestimate the peripheral retina, which provides what has been described as *ambient* vision as opposed to the *focal* vision of the centre (Trevathan, 1968). Though poor at *resolution*, it can still *localise* well; localisation being a 'hyperacuity' task (Westheimer, 1979); and having alerted the brain to a threatening or interesting peripheral stimulus, can direct the fovea accurately enough to what is to be examined.

The result is that so long as we are prepared to wait the fifth of a second or so that it takes to plan and execute these sighting movements, our vision is as good with our tiny fovea as it would be if most of our retina were packed with the luxury receptors. [However, it must be remembered that in man not all the visual field can be fixated by the eye. The human field of fixation is only some 90° in its horizontal and vertical extent (Hofmann, 1925; Burian and van Noorden, 1974) and in the cat, some 50° (Crommelinck and Roucoux, 1979)]. F W Campbell, in an unpublished communication, has put it well:

"The potential spatial information being continuously supplied to the eye is enormous. A hemisphere has an area 74×10^6 sq. min. Using the argument that a matrix of recognisable letters can be set up each in an area of 50 sq. min., it follows that 1.5×10^6 letters can be displayed in one hemisphere. Say an observer proceeded to read these letters sequentially, taking one second per letter, it would require just over 26 days to identify them, assuming 16 hours of observation were made each day. If the matrix is exposed for 1/10 second, to prevent eye movements, only the 14 letters falling on the central fovea could be identified: this is only 1/100,000 of the total available when eye movements are permitted."

Such ideas perhaps suggest the notion of deliberate scanning of the visual field by the eye, like a radar antenna, or the scanning spot in a television camera. Although this is indeed found in some species with a very small number of retinal receptors—for example the jumping spider (Land, 1969),

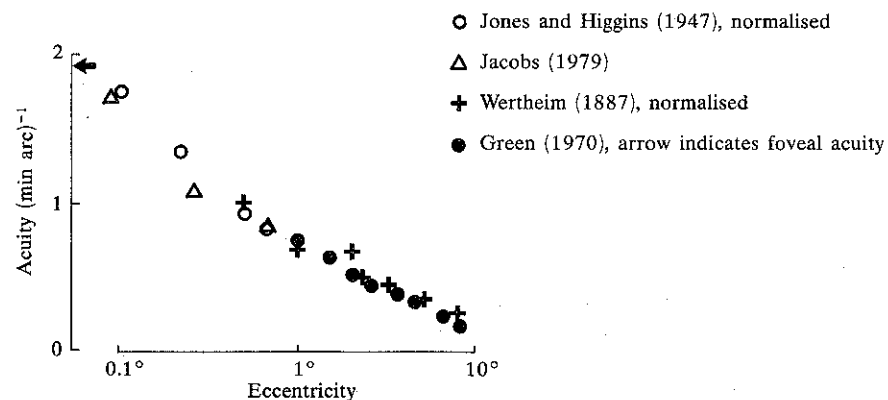


Figure 1.5 Visual acuity and cone spacing as a function of distance from the fovea (Wertheim, 1887; Jones and Higgins, 1947; Green, 1970; Jacobs, 1979).

or the extreme case of *Copilia* with its single retinal unit (Gregory et al, 1964)—it is an unnecessarily prosaic procedure when there are enough of the low-grade receptors surrounding the fovea to signal to the eye movement system the location of something interesting in the field that ought to be fixated: localised sounds or even touch can be used in the same way. The result is the *gaze-shifting movement*, which may be either visual or auditory or even purely voluntary, in which a new object of interest is seized by the fovea, and subsequently held there by the *holding* mechanisms introduced in the previous section.

It is apparent that these shifting and holding mechanisms are to some extent antagonistic. Shifting implies a movement of the eye relative to the visual field, and we saw earlier that a moving eye is an eye that cannot really see. How can these opposite requirements—of displacement without motion—be reconciled? The answer is, by making the gaze-shifting movements as *fast* as we can: although the consequent disruption of vision during the movement is that much greater, the period of time for which it occurs is correspondingly reduced. The resulting eye movements—called *saccades*—are actually very fast indeed, with velocities approaching $1000^\circ \text{ s}^{-1}$: they form the subject of chapter 4.

As might be expected, creatures with less distinction between the central and peripheral retina make movements of this kind less often. Frogs make very few, if any, spontaneous movements of the gaze (Walls, 1942; Berthoz, 1985). In the rabbit, saccades are considerably less common than in man, and are on average rather larger (Hughes, 1971; Collewijn, 1977b; Collewijn and Zuidam, 1977); presumably there is little point in a small movement that will barely improve the acuity. In many species, gaze-shifting is performed as much by the head as by the eyes [an owl's eyes, though supplied with the usual muscles, are tubular in shape and unable to move by more than a degree or so; but the head makes up for it by being able to turn through 270° (Walls, 1942; Steinbach and Money, 1973)], and restraint of the head may suddenly bring the oculomotor system to life (Stryker and Blakemore, 1972). One may also find that when the head moves, there is an associated suppression of the vestibular holding mechanism, which otherwise would defeat the object of the head movement (Robinson and Zee, 1981; van der Steen and Collewijn, 1984; Berthoz, 1985; Lauritis and Robinson, 1986). In a related fashion, children appear to make more head movements while reading than adults (Netchine et al, 1984); and the fact that human foetuses show holding eye movements some two months before they show saccades (Birnholtz, 1981) perhaps also suggests a general trend in development towards greater prominence of the fast shifting movements. Figure 1.6 shows comparative records of head and eye movements under natural conditions in cat, rabbit, and man.

Apart from their role as pure gaze-shifting movements, saccades also form an inseparable component of the complete holding response. If an animal's visual world is continuously moving in one direction, it is plain

that, even if the eye moves at an appropriate rate and direction to stabilise the retinal image, it cannot go on doing so indefinitely because of its limited range of travel. What happens in practice is that the slow following movement of the eye is interrupted at intervals by a rapid saccadic movement in the opposite direction, resulting in a sawtooth-like time course called *nystagmus*. Although the time course of these optokinetic saccades is in every respect similar to the saccades described above, many animals that are able to produce them as part of nystagmus are unable to use them as 'voluntary' gaze-shifting saccades: only those with some kind of central retinal area do so (Walls, 1962). Vestibular nystagmus—whose time course is indistinguishable from optokinetic nystagmus—can be evoked in the same way, by sufficiently prolonged rotation of the head.

Finally, if the target of interest is moving relative to the outside world, it may be *tracked* partly by a sequence of saccades, and partly by attempting to move the eye at the same *velocity* as the target. This second component is called *smooth pursuit*, and is considered in chapter 3.




Figure 1.6 Head and eye movements in the freely-moving rabbit, cat, and man: H = head, E = eye in space (gaze), e = eye relative to head (data from Collewijn, 1977a).

1.3 Consequences of binocular vision

With two eyes, it is almost inevitable that there will be at least some degree of overlap of their fields, even when they are essentially sideways-pointing (figure 1.4) [the only mammal without some binocular field is said to be the

whale (Hughes, 1977)]. But this does not necessarily indicate a need for much cooperation between the movements of the two eyes. In most lizards, for example, best vision occurs in the lateral direction; as a consequence, they are able to enjoy the ability of directing both eyes independently at will to objects of interest in their respective visual fields, although, as in all species, the involuntary holding movements are always fully coordinated in the two eyes. But as soon as there is overlap between the visual fields corresponding to the area centralis of each eye, there is an obvious advantage in binocular yoking of the two eyes in voluntary movements as well. If the foveas overlap, then there is the possibility of improving *detection*, especially in dim light, by pooling information from each eye (Jones and Lee, 1981). In addition, detailed comparison of the visual images in the two eyes, each of which sees from a slightly different viewpoint, permits the perception of *depth*. Both these functions, but particularly the latter, demand rather precise control of the relative positions of the eyes: they need in fact to move *conjunctly*, swivelling together as a single unit, with the two visual axes pointing in parallel directions. A recent comparison between two birds with different feeding habits and different patterns of oculomotor coordination (Wallman and Pettigrew, 1985) is of interest in this context: the little eagle is an aerial predator, whereas the tawny frogmouth is both predator and prey, hunts near the ground, and has the more divergent visual axes. Both are capable of aligning their eyes, but whereas the eagle's saccades are normally conjunct, those of the frogmouth are not. In addition, the frogmouth responds to a startling stimulus by divergence, increasing the extent of its cyclopean field; and it shows little sign of a vestibulo-ocular reflex, using head movements to direct the gaze.

Conjunct movements are fine so long as what we are looking at is sufficiently far away, but clearly inappropriate as soon as we try to look at objects at different distances from the eye. We then require the eyes to *converge* from parallelity to bring the two images of the object onto corresponding portions of the two retinæ. Oddly enough, this ability has only been acquired with difficulty: under natural conditions vergence movements are rather infrequent except in man and the primates. Even in man, these *disjunctive* or *vergence* movements are relatively undeveloped, and lack both the precise rapidity of saccades, and their complete voluntary control: it is significant that the diplopias and other disorders caused by malfunctioning of the disjunctive eye movement control system form the largest single group of oculomotor disturbances in clinical practice. Nevertheless, vergence can be observed in other animals, even those with essentially lateral eyes [for example, up to 18° in the rabbit (Zuidam and Collewijn, 1979), and some 10° in the pigeon (Bloch et al, 1981)]. Some aspects of the possible evolutionary history of binocular control have been presented in a thoughtful discussion by Walls (1962).

Table 1.1 summarises the classification of the major types of eye movement presented above; an extremely complete analysis of the categories of eye

movements may be found in Hallett (1986). For completeness eye movements that do *not* serve a visual purpose should perhaps be mentioned: not only the winks and *oeillades* that enliven social intercourse but the repulsively substantial voluntary convergences cultivated by small boys, and the odd facility of voluntarily vibrating the eyes at about 20 Hz—'voluntary nystagmus' (Westheimer, 1954c)—that is enjoyed by a small number of otherwise normal subjects. Finally, the commonest nonvisual eye movement of all is that which occurs each time we blink. It is commonly said that the eyes turn upwards when the eyelids are closed (Bell, 1823), but more recent experiments (Evinger et al, 1984; Collewijn et al, 1985d) have shown that this is not normally the case. During blinking in man there is a small downward and nasal movement accompanied by retraction, which has a shorter duration than the lid movement and is not simply the result of mechanical linkage; with more prolonged closure there is a prolonged tonic deviation, that is upwards in some subjects but downwards in others.

Table 1.1 A classification of types of eye movements.

<i>Gaze-shifting movements</i>	
fast:	saccades microsaccades quick phase of nystagmus
slow:	smooth pursuit vergence
<i>Gaze-holding movements</i>	
	vestibular optokinetic
<i>Fixational movements</i>	
	tremor drift

Vestibular eye movements

"... for is it not a wittie exploit of nature to close up in so small a hole, a drumme hard laced, having on the hinder part two small strings, and three little bones, resembling a forge, a hammer and a stirrop, three small muscles, and a labyrinth contayning the inward ayre?"

2.1 The vestibular sense organs

2.1.1 Gross anatomy

As its name suggests, the labyrinth of the inner ear is a complex structure. It seems to have evolved as an elaboration of the apparatus used by fish to signal movement of the water that surrounds them. In its most primitive form there are sensory cells with hairs that project directly into the external medium and are bent when it flows past them; in higher fishes, these cells lie in protected channels under the scales and form the lateral line organs. So their function is partly *exteroceptive* (responding to eddy currents set up by the movement of things in the outside world) and partly *proprioceptive*, (responding to the fish's own movements through the water). In the course of evolution, and in adapting to terrestrial life, closed fluid systems containing *endolymph* developed, that were no longer in direct communication with the surrounding water. Of these systems, some parts came to lie sufficiently close to the surface to become essentially exteroceptive, responding to vibrations in the outside world, while others became more deeply buried and specialised as proprioceptors: these are the cochlear and vestibular organs of land vertebrates, responding respectively to sound vibrations and to movements of the head. At the same time, the vestibular organs themselves differentiated on the one hand into closed *semicircular canals* capable of signalling angular rotation, and on the other hand into the *otolith organs*, hollow sacs containing heavy accretions resting on the hairs, whose bending thus signals linear acceleration or the direction of gravity.

So the labyrinth can be divided into three functionally distinct parts: the cochlea (with which we will not be concerned), the three semicircular canals, and a pair of endolymphatic sacs (the *utricle* and *sacculle*, forming the otolith organs), all of which are in communication (figure 2.1). In man, one of the canals is near the horizontal when the head is in a normal erect position, whereas the planes of the others form an angle close to 45° with the sagittal plane (figure 2.2), so that the plane of the right posterior canal is parallel to that of the left anterior one, and vice versa. They thus form a set of three balanced and opposed pairs in mutually perpendicular planes (actually not *quite* perpendicular: Blanks et al, 1975), whose orientations define the three coordinates in which the body senses angular rotation. Comparative aspects of the vestibular organs have recently been summarised by Simpson and Graf (1985).

The receptor cells in the labyrinth are gathered together in special regions of the epithelium, the *maculae* of the utricle and saccule, and the *cristae* of the canals. In each case, their hair-like processes project into a mass of gelatinous substance. In the canals, this aggregation of jelly and




Figure 2.1. The human labyrinth and its innervation (Hardy, 1934).




Figure 2.2. The position of the semicircular canals in the head: the arrows show the directions of fluid movement that are excitatory for each canal (s, superior; p, posterior; h, horizontal).

Figure 2.3. Diagrammatic section through a single semicircular canal, showing the membranous canal containing endolymph (shaded), surrounded by perilymph within the bony canal. At top left, the cupula can be seen effectively blocking the flow of endolymph through the ampulla: associated with it are the hair cells and their sensory innervation.

hairs forms a flap (the *cupula*) that projects across the canal, widened at this point to form the *ampulla*. The flap appears to be watertight, so that any displacement of fluid in the canals makes the cupula bend to one side, stimulating the receptor cells whose hairs are embedded in it (figure 2.3). Normally, such displacements are the result of *rotation* of the head, when the fluid tends to get left behind: the detailed mechanics of this process are considered in section 2.1.4 below: In the utricle and saccule the jelly contains heavy calcareous granules. Again, during movements of the head the whole thing tends to get left behind because of its inertia, stimulating the receptor cells by the bending of their hairs; but in this case they respond essentially to *linear* rather than rotational acceleration. These same cells also respond to static tilt of the head, because the weight bends them in different directions at different orientations of the head. In fact their response to gravity is exactly the same as to any other linear acceleration, and if the two act together they add vectorially to give a response equivalent to some particular angle of tilt (figure 2.4). This is only to be expected from a simple inertial system that lacks anything like a gyroscope.

The structures described so far are contained within the membranous labyrinth: this in turn lies within a corresponding arrangement of bony tunnels and caves, the bony labyrinth. Between the two is a gap that is filled with a fine network of connective tissue and bathed in *perilymph*. Presumably the perilymph also moves about in response to movements of the head: but there is no equivalent to the cupula to impede its flow, and it is thought not to contribute substantially to the responses of the hair cells in the cristae (Dohlman and Kuehn, 1973; McCabe and Ryu, 1973; Wilson and Melvill Jones, 1979).




Figure 2.4. The action of the otolith receptors: (a) at rest, (b) with head tilted, and (c) under horizontal linear acceleration. The last two conditions are indistinguishable in their effects.

2.1.2 The receptor cells and their innervation

The sensory cells of the maculae and cristae are closely similar in appearance: their sensitivity to different types of motion is thought to be entirely a function of their relation to the larger structures of which they form a part.

In mammals and birds there seem to be two basic types of hair cell: a flask- or bottle-shaped *type I cell*, and a roughly cylindrical *type II cell* (figure 2.5); intermediate types have also been described (Wersall, 1956; Engström, 1965). Both have the same pattern of hairs projecting from the epithelial surface. Each cell has a single flexible *kinocilium*, near its edge, and some sixty to a hundred stiff *stereocilia*, arranged in a regular pattern and graded in such a way that the longest originate near the kinocilium, and the shortest furthest away (figure 2.5) (Engström et al, 1962).

This gradation of size defines a direction of polarisation for each cell in the epithelial plane: bending of the stereocilia in the direction of the kinocilium causing excitation, and in the opposite direction, inhibition. A striking feature of the cristae is that the vast majority of the cells show exactly parallel orientation, facing the utricle in the horizontal canal and away from it in the vertical canals. The arrangement in the macula is more complex (Flock, 1964). The directions of individual cells lie roughly parallel with their immediate neighbours, but on a larger scale there are systematic variations of this 'morphological polarisation' in different regions of the macula (figure 2.6) (Spoendlin, 1965). These differences reflect the fact that the semicircular canals are individually one-dimensional receptors, whereas the utricle and possibly the saccule are essentially two- or three-dimensional, signalling the *direction* of an acceleration as well as its magnitude.

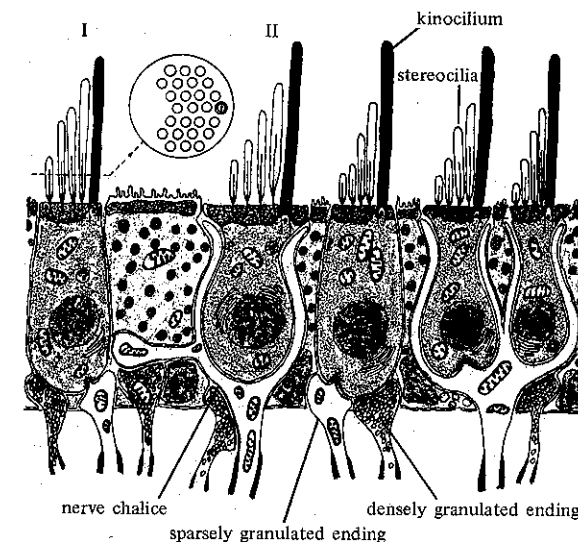


Figure 2.5. Diagrammatic cross section of vestibular sensory epithelium showing type I and type II cells and their innervation. The inset shows a typical arrangement of the sensory hairs.

The kinocilium shows a higher degree of internal organisation than do the relatively simple stereocilia, with an arrangement of nine internal longitudinal filaments on the outside and two in the middle that is typical of motile cilia. Nevertheless, there is increasing evidence that it is the *stereocilia* that are the source of the ionic currents that generate the efferent response of the cell. The transducer mechanism is still uncertain.

Apart from shape, the main difference between the two types of hair cell lies in their innervation. Type I cells are closely enveloped in a single nerve calyx (figure 2.5), from which the cell is separated by a gap of some 25–30 nm; sometimes two or three receptors may be found within the same enveloping chalice. Synaptic vesicles and other accessory structures are evident at the junction between the sensory cell and the afferent fibres. However, at some points the gap is very much reduced (often in association with synaptic bars): this fact, together with the very large potential synaptic area, has led to the suggestion that the functional connection between receptor cell and nerve fibre may be partly electrical in nature. The type II cells are innervated by at least two different kinds of ending. The first variety is sparsely granulated, and may be equivalent to the chalice endings of the type I cells, since collaterals of this type can sometimes be seen innervating both types of receptor. The second type is less common, and forms button-like endings having a dense granulation: these endings are found not only on the bodies of type II cells, but also on the surface of nerve chalice and their collaterals (figure 2.5) (Ades and Engström, 1965), and appear to be efferent (centrifugal) in nature. Each type II cell receives many

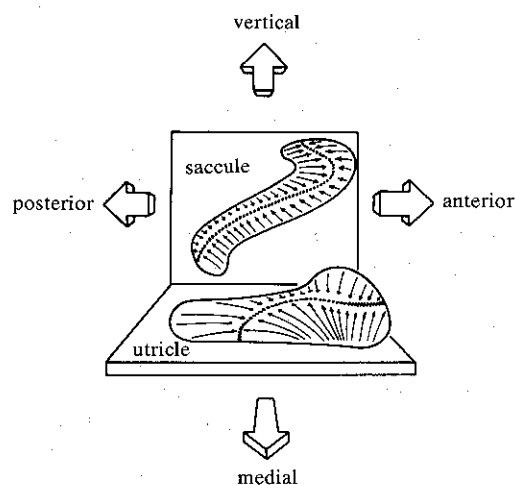


Figure 2.6. Diagrammatic representation of the maculae of the utricle and saccule, in the relative orientation that they occupy in the head. The arrows indicate the direction of functional polarisation of the receptor cells at different points on the surface (data from Spoendlin, 1965).

endings from independent branching fibres. Counts of afferent and efferent fibres in the vestibular nerve (Gacek, 1961) show that the number of efferent neurons is very small, perhaps as few as two hundred in the cat; they originate in a region lying between the abducens nucleus and superior vestibular nucleus (Gacek and Lyon, 1974; Goldberg and Fernandez, 1980). Thus, either these fibres must branch very considerably, or alternatively some of the afferent fibres may send back collaterals that terminate as apparent efferents. True centrifugal fibres, despite their small number, have nevertheless been shown to have marked effects on the discharge of afferent fibres in the vestibular nerve, normally excitatory but in some cases inhibitory (Klinke and Schmidt, 1968; Dichgans et al, 1970; Klinke, 1970; Russell and Roberts 1972; Schmidt et al, 1972; Precht, 1974; Goldberg and Fernandez, 1980). In fish, their activity is modulated by active movement in such a way as to discount the stimulation of the lateral line receptors that is caused by the fish's own movements (Roberts and Russell, 1972; Russell and Roberts, 1974). It has been suggested that they may have a similar function in mammals during large head movements, shifting the response curves of the receptors away from the cutoff point so as to match their dynamic range to that of the stimulus (Goldberg and Fernandez, 1980). There is no evidence that the efferents are influenced by visual or oculomotor activity per se (Keller, 1976; Büttner and Waespe, 1981).

The cell bodies of afferent fibres innervating vestibular receptors lie in the superior and inferior *vestibular ganglion* (Scarpa's ganglion) situated close to the labyrinth. The central course of these fibres constitutes the *vestibular nerve*, which together with the auditory nerve forms the eighth cranial nerve (nVIII). These fibres project to the vestibular nuclei and also directly to the cerebellum: the projections are discussed in chapter 9.

2.1.3 Electrical responses to angular accelerations

Lowenstein and Sand (1940) appear to have been the first to record the relation between movements of the head and activity in the vestibular fibres. Adrian (1943) showed that the otolith organs of the cat gave a slowly adapting and somewhat nonlinear response to tilt of the head, and characterised the directions of rotational acceleration that were excitatory for the different canals (figure 2.7). It can be seen that these excitatory directions correspond with the direction of polarisation of the hair cells of the cristae. Groen et al (1952) measured the frequency of firing of individual fibres immediately after impulsive decelerations of different magnitudes, and related these frequencies to the corresponding calculated cupular deflections. They found that individual fibres gave S-shaped response curves as a function of deflection, and that the curves for different units, although similar in shape, were typically displaced by different amounts along the deflection axis, thus accounting for their differing rates of firing at rest (figure 2.8). A functional advantage of this arrangement is that the dynamic range of the nerve as a whole is thus considerably greater than that of any

one fibre in it, so that information is carried by *how many* of the fibres are firing as well as by their individual frequencies—a kind of recruitment. In fact it is difficult to see otherwise why so many sensory fibres are needed to carry information about a single one-dimensional variable.

Vestibular fibres show a certain amount of variation in their resting rate of discharge, some being spontaneously active and responding to rotational acceleration in different directions by an increase or decrease in their firing rate, whereas others are normally silent and fire only at some threshold degree of stimulation in the excitatory direction; one may also divide them




Figure 2.7. Directions of head movement that are stimulatory for the semicircular canals of the left side.

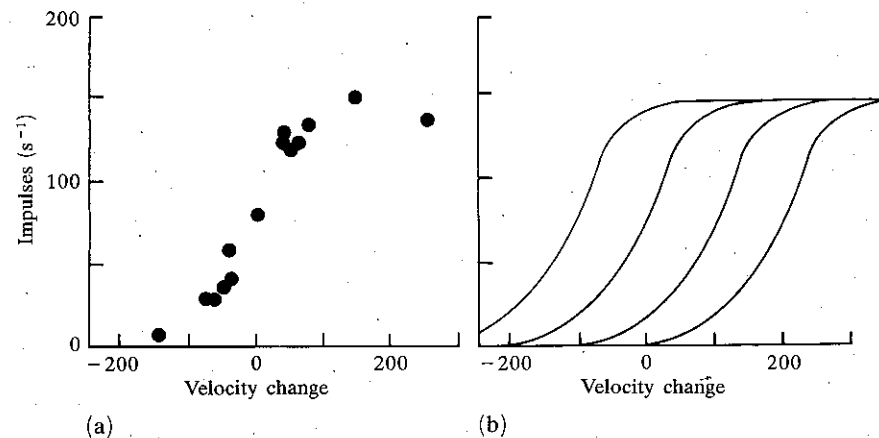


Figure 2.8. Peak firing frequency of vestibular fibres in the ray to sudden angular velocity changes of different magnitudes. One can assume that maximum cupular deviation under these conditions is proportional to the change in velocity, so that (a) there appears to be an S-shaped relationship between cupular deflection and firing frequency. Units of this type lying at different points along a scale of cupular deflection show variations in firing rates at rest (b); a set of such units will have a wider dynamic range than any one of them (data from Groen et al, 1952).

into those that tend to fire at a regular rate, with low conduction velocity, and those that fire irregularly (Goldberg and Fernandez, 1977; Ezure et al, 1978). However, there is as yet no compelling evidence linking such a twofold division of fibres on the basis of their electrical responses with the two morphological classes of receptor cell described in section 2.1.2. It will be recalled that the same sensory fibres are often seen to innervate both types of receptor, so that any differences in their responses are likely to be blurred. Possibly the spread of characteristics that is observed between fibres as to their thresholds for firing, regularity and degree of adaptation is simply due to variations in the proportion of each type of receptor that particular fibres innervate.

2.1.4 Dynamic properties of the semicircular canals

As in so many mechanoreceptors, it turns out that the dynamic response of the receptors in the semicircular canals is largely determined by the special properties of the mechanical system of which they form a part rather than by any intrinsic property of the receptor cells themselves. Thus an analysis of the mechanical behaviour of the endolymph and cupula can provide quite a good description of the dynamic characteristics of the neural response to rotations. Steinhausen (1931) was the first to make such an analysis. If an individual canal is taken to be a circular tube of uniform cross section, filled with an ideal fluid of known density and viscosity, and the cupula is assumed to be perfectly elastic (producing a restoring force proportional to its displacement from the rest position) and watertight, then the system can be perfectly represented by a simple second-order mechanical model—the *torsion-pendulum* model—shown in figure 2.9.

The input to this system is the angular position of the head at any moment, θ_{in} (resolved in the plane of the canal), whereas the output is taken to be the corresponding deviation of the cupula, θ_{out} ; the whole system behaves like a second-order displacement divider (see appendix 2), and has properties analogous to a damped pendulum. If we use the symbols r , k , and M to represent respectively—in appropriate units—the viscous, elastic, and inertial constants of the system, we can derive the transfer function of the whole arrangement:

$$\frac{\theta_{out}}{\theta_{in}} = \frac{(rD + k)^{-1}}{(rD + k)^{-1} + (MD^2)^{-1}} = \frac{D^2}{D^2 + (r/M)D + k/M}$$

Thus only two parameters are needed to describe its behaviour: the ratio r/M , and the ratio k/M . [There is some variation of notation amongst different authors: thus von Egmond et al (1949) and many others use Π , Δ , Θ for r , k , M ; Melvill Jones and Spells (1963) use Φ , Δ , I ; and Melvill Jones and Milsum (1971) prefer b , k , J . The reader can take his pick! Although an admirably rationalised scheme of notation for this and other aspects of vestibular stimulation was proposed by Hixson et al (1966), it is not in general use.] It turns out that the semicircular canals are very highly

damped, and so behave like two first-order high-pass units in series:

$$\frac{\theta_{\text{out}}}{\theta_{\text{in}}} = \left(\frac{D}{\lambda + D} \right) \left(\frac{D}{\mu + D} \right),$$

where $\lambda\mu = k/M$, and $\lambda + \mu \approx \mu = r/M$; thus $\lambda = (K/M)(M/r) = k/r$. The two associated time constants are thus given by M/r (the inertial time constant) and r/k (the elastic time constant). For the posterior canal of the squirrel monkey, these constants have values of 3 ms and nearly 6 s, respectively, as determined by electrical recording from the vestibular nerve (Fernandez and Goldberg, 1971); in the cat, M/r is somewhat smaller, around 1.7 ms (Fernandez and Valentinuzzi, 1968). For reasons that are discussed below, animals with smaller heads tend to have shorter elastic time constants: 4 or 5 s in the cat, 3 s in the frog, 2.5 s in the gerbil (Blanks and Precht, 1976; Schneider and Anderson, 1976; Ezure et al, 1978); however, some care must be taken with comparative figures since time constants measured in different individual fibres in the same species may show a spread of a factor of two or more (Curthoys, 1982), related to whether their firing is of the regular or irregular type. In man, the value for the long (elastic) time constant of around 10 s originally reported by von Egmond et al (1949) was probably too great, as were estimates by subsequent workers who assumed that the time constant of the decay of the vestibulo-ocular reflex (VOR) during rotation at constant velocity was the




Figure 2.9. The mechanical action of the semicircular canal. Behaviour of the cupula, and of a simple mechanical analogue, (a) at rest; (b) immediately after, and (c) some time after, a rotation at constant velocity.

same as that of the canals themselves (see section 2.3.1 below). Differences between the time constants of the different canals are also found, and again may well be related to the dynamics of typical head movements in different planes (Melvill Jones et al, 1964; Fernandez and Valentinuzzi, 1968; Gilson et al, 1973).

The functional significance of these dynamic properties can perhaps best be appreciated by considering the response of the semicircular canals to sinusoidal rotations of the head. It is a simple matter to calculate from the transfer function already derived how the gain and phase of the vestibular response will be related to the frequency of such an input (appendix 2). Figure 2.10 shows the frequency response of the mechanical model, with the use of Fernandez and Goldberg's (1971) experimental parameters. At very high frequencies the gain is constant, and there is only a small phase difference between head and cupular displacement, so that in this region the cupula is responding essentially to angular position of the head. At intermediate frequencies (between about 0.07 and 20 Hz) the response is

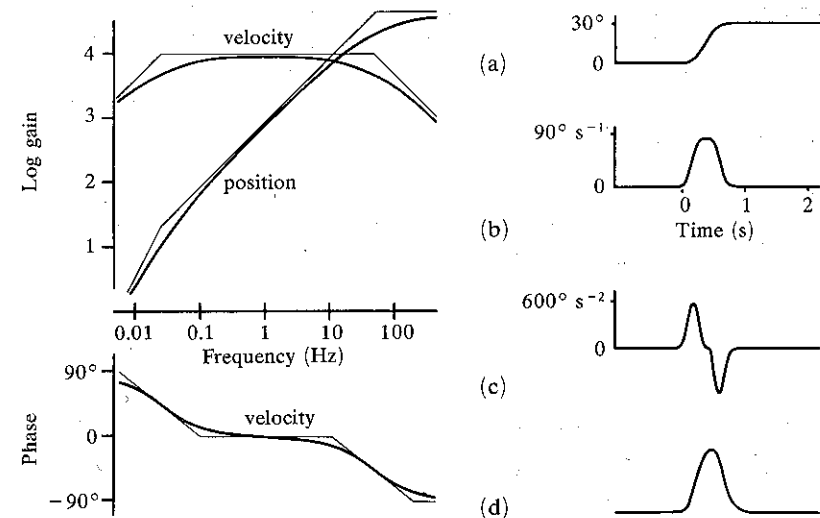


Figure 2.10. Gain and phase of the torsion-pendulum model of the semicircular canal, calculated with the use of Fernandez and Goldberg's (1971) experimental parameters. The 'position' curve is for conditions in which the head is rotated sinusoidally at different frequencies and with constant amplitude of movement: in the 'velocity' case it is the transfer function relating head angular velocity to cupular displacement that is plotted.

Figure 2.11. The canals essentially signal angular velocity under natural conditions. (a) Represents the time course of a typical natural head movement; (b) and (c) show respectively the time course of its velocity and its acceleration; (d) shows the calculated cupular displacement resulting from the same head movement: its similarity to the velocity curve is evident (courtesy of Dr T D M Roberts).

advanced in phase by some 90° with respect to head position and the gain increases linearly in proportion to the frequency, implying that here it is angular velocity that determines the response.

We can see this more clearly if we plot cupular displacement relative to head *velocity* rather than position (figure 2.10), giving a response which is seen to be flat over the same range of frequencies, with only small phase differences. Finally, at frequencies lower than some 0.01 Hz, the cupula displacement leads head velocity by around 90° , and the corresponding portion of the velocity transfer function increases linearly with frequency; in this region the cupula is responding essentially to the angular acceleration of the head. Since frequencies above 10 Hz are of little relevance to the control of eye movements, it is best to consider angular *velocity* rather than position as the primary vestibular input, and the transduction process as a high-pass filter with a time constant of a few seconds. Thus the statement often found in general textbooks, that the semicircular canals are for signalling rotational acceleration, is only true for very slow movements. Over some two and a half decades of frequency, effectively covering the 'physiological range', the canals signal not acceleration but velocity.

This can be appreciated by considering the movement of the cupula during a natural spontaneous turn of the head (Roberts, 1967), when the deflection of the cupula follows quite faithfully the time course of the head's instantaneous velocity (figure 2.11). The lower-frequency effects only make themselves felt under the very unnatural conditions of ordinary clinical examination, when the subject is rotated in one direction for a prolonged period. The cupular response then dies away in a roughly exponential manner (figure 2.12), and an *after-response* of opposite sign occurs at the end of the rotation.

Some evidence that the range of frequencies for which the canals act as velocity transducers is indeed 'tuned' to the dynamics of natural movements of the head comes from Melvill Jones and Spells' (1963) investigation of the dimensions of the canals in different species. The long time constant of the semicircular canal is in part a function both of its radius of curvature and of its diameter: the authors find that these two quantities show slight but significant changes (corresponding to a lengthening of the time constant) that are related to the mass of the animal to roughly the extent that might be predicted by a consideration of the natural movements of its head. No doubt similar factors underlie the differences between the time constants of the horizontal and the vertical canals, that have already been remarked on.

The simple mechanical model discussed so far accounts pretty well for the way in which vestibular fibres respond to angular rotations, but not quite exactly. At very high frequencies, vestibular firing frequency shows a greater phase lead than would be expected from the torsion-pendulum model, suggesting that like many mechanoreceptors the vestibular hair cells have an *intrinsic* adaptational mechanism as well as what is provided by the mechanics of the canal (Fernandez and Goldberg, 1971; Goldberg and

Fernandez, 1971; Ezure et al, 1978). However, it manifests itself at frequencies somewhat removed from those that are significant in the control of eye movements in real life, at least as far as gain is concerned. In the case of the squirrel monkey, the observed discrepancy in phase for horizontal movements amounts to some 37° at 8 Hz, equivalent to about 13 ms of prediction, against which must be offset delays associated with passage of information through the vestibulo-ocular system and in initiating muscular contraction, which also amount to some 13 ms. Uncertainties in these latter factors make the interpretation of phase measurements above 3 Hz or so somewhat metaphysical.

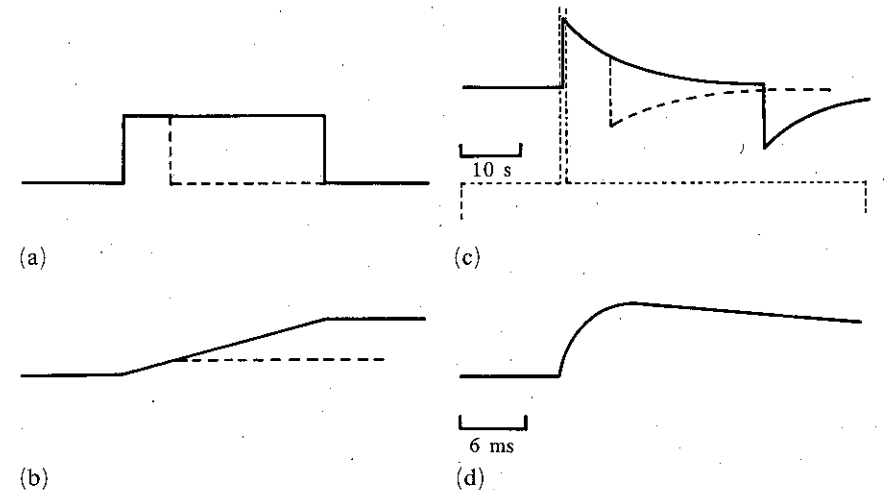


Figure 2.12. Cupular deflection during rotations at a constant angular velocity; (a) shows the velocity, and (b) the position of the head during a prolonged rotation: the dashed lines indicate a different case, when the movement is stopped relatively soon after it has started; (c) shows the time course of the cupular deviation on the same time scale, and (d) is an expanded representation of the beginning of the response, to show the course of the initial deviation (which depends on the short time constant). In (c) the dependence of the size of the after-response on the duration of the rotation can be seen.

2.1.5 Electrical responses to linear accelerations

The utricle and saccule have on the whole received rather less attention from investigators than the semicircular canals. The saccule in particular is still something of a mystery: Ashcroft and Hallpike (1934) could find no responses from saccular fibres in the frog during rotation or tilting, although the saccule appeared to be sensitive to low-frequency vibration (see also Lewis et al, 1982); this has been confirmed more recently by Hess and Precht (1984), although others have found clear evidence for saccular responses to head tilt as well as vibration (Gallé and Clemens, 1973;

Lannou and Cazin, 1976). But in mammals there is agreement that saccular fibres respond to head tilt as well as to vibration, and can generate ocular reflexes (Fluur and Mellstrom, 1970; 1971; Hwang and Poon, 1975; Young et al, 1977; Chan et al, 1979; Tomko et al, 1981).

Lowenstein and Roberts (1949), recording from utricular fibres in the ray, found two classes of response. One was essentially static or tonic, showing no appreciable adaptation in firing frequency for fixed angles of tilt, whereas the other class was dynamic or phasic, and greater for changes in the angle of tilt, declining to a lower resting level when the position of the head was fixed. Units of the first kind were found to be preferentially sensitive to different angles of tilt, corresponding no doubt to the variations in the direction of morphological polarisation that is seen in receptors in different parts of the macula. A similar mixture of tonic and adapting units is found in the cat, and most units are maximally sensitive when the head is in the normal upright position (Loe et al, 1973; Anderson et al, 1978) (figure 2.13). The rate of discharge for different directions of linear acceleration is a simple function of the component of the acceleration resolved in the direction of maximum response: the relationship may indeed be a linear one (Fernandez et al, 1973).

In these experiments the rate of tilt was sufficiently slow that one can exclude possible effects from the semicircular canals. With faster movements of the head one cannot exclude the possibility that, because of the continuity of the endolymph between the otolith organs and the canals, some kind of mechanical interaction may take place between them. A further possibility is that linear accelerations may distort the semicircular canals and thus lead to modification of the cupular response. Although interactions between linear and rotational accelerations have often been




Figure 2.13. The rate of firing of a single fibre from the vestibular nerve of a cat, at different angles of tilt in the sagittal plane (data from Loe et al, 1973).

described in the VOR (for example Guedry, 1965; Benson and Bodin, 1966a; 1966b; Niven et al, 1966; Benson, 1970; Bles and Kapteyn, 1973), it is much more probable that these are the result of more central interactions, since it has been clearly demonstrated that convergence occurs in the vestibular nuclei between fibres from the different parts of the peripheral vestibular apparatus (see section 9.1.4).

All the same, Benson et al (1970) were able to show that horizontal canal units in the vestibular nuclei, that were not responsive to linear accelerations in the horizontal plane (and were therefore assumed not to be influenced by otolith organs), could nevertheless be made to fire periodically in response to a rotating *linear* acceleration vector without rotational acceleration. If the canals do indeed distort significantly under linear acceleration, a rotating vector of this type is the best way to demonstrate the fact, since the distortion is likely to push the endolymph round the canals in a peristalsis-like manner as it moves in response to the changing direction of acceleration. However, the results of experiments by Correia and Money (1970) in which the response to a rotating linear acceleration vector could still be obtained after blockage of the canal ducts suggest that this cannot, in fact, be the true explanation. Whatever the mechanism of such interactions, they probably in any case play only a minor role in the vestibular control of eye movements under natural conditions.

A simpler kind of interaction is the *positional nystagmus* that can often be observed on tilting the head while 'suffering' from alcohol poisoning. It has been suggested (Money and Myles, 1974) that this may result from a change in the relative density of cupula and endolymph brought about by the direct physical effect of the alcohol: the cupula would then act, in effect, like an otolith organ. They observe, in support of this theory, that ingestion of heavy water in suitable quantities can cancel out the density change and prevent the positional nystagmus. It is quite possible that the well-known nauseating aftereffects of alcohol are due to a motion sickness arising from illusory stimulation of this kind.

2.2 Static vestibulo-ocular responses

It is convenient to divide reflex vestibular movements into two classes—*static* and *dynamic*—according to whether they are produced by *linear* or *rotational* accelerations, stimulating respectively the otolith organs and the canals. In man and other species with front-facing eyes, the static ocular reflexes are rather weak and tend to be dominated both by dynamic vestibular reflexes and by vision, but in animals like the rabbit they are easily demonstrated (figure 2.14): if the skull is tilted, the eyes tend to move in such a way as to keep their former orientation in space. A complicating factor in such demonstrations is the presence of reflex eye movements which occur in response to stimulation of proprioceptors in the neck if the head is moved relative to the body (the *cervico-ocular reflex*, or COR: see below).

(In crayfish a somewhat analogous response to stimulation of the limbs can be observed: Olivo and Jazak, 1980). Thus it is necessary to move head and body together if one wishes to isolate the vestibular component. In the rabbit, the vestibular contribution to movements of the eyes with the body held still is about 70% (de Kleijn, 1921b), and almost full compensation can be achieved for tilt angles of up to 70° (figure 2.15).

In man, the only static VOR that is not normally completely dominated by vision is the *torsion* of the eye (that is, rotation about the anterior-posterior axis) that can be observed on lateral tilt of the head (Merton,




Figure 2.14. Oculomotor responses to tilt of the head compared in the rabbit and in man (after Duke-Elder and Wybar, 1973).

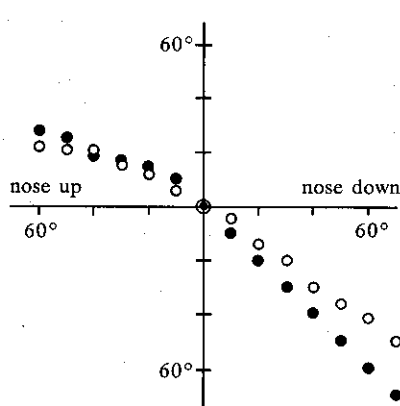


Figure 2.15. Compensatory rotation of a rabbit's eye for different degrees of tilt in the sagittal plane (filled circles): the open circles represent observations made when the dorsal roots of the first and second cranial nerves have been cut, that is, when the influence of neck reflexes has been removed to leave a response that is almost wholly vestibular in origin (data from de Kleijn, 1921b).




Figure 2.16. Ocular counterrolling in man: rotation of the eye about a torsional axis measured for different angles of lateral tilt of the head (data from Miller, 1962).

1956; Miller, 1962; Belcher, 1964; Collewijn et al, 1985c; in the monkey: Krejcova et al, 1973; and in the owl: Steinbach et al, 1974). These counter-rolling movements are small, [typically less than 7° or so—as in figure 2.16—although Petrov and Zenkin (1973) report vestibular torsion of up to 18°] and can hardly be said to be compensatory to any significant extent, except perhaps for very small angles of tilt, where the slope of the relationship between stimulus and response is steeper (Belcher, 1964). The limitation is not in the eye itself, for very large lateral linear accelerations can generate much larger angles of torsion: the signal from the otolith organs appears to depend on the size of the horizontal acceleration component rather than the angle of the gravitational vector (Woellner and Graybiel, 1959; see also Young, 1985). Possibly these movements represent a vestibular reflex originally intended for animals with lateral eyes. As in the rabbit, otolith and neck proprioceptor components can be isolated, and there is a further phasic component dying away with a time constant of some two seconds that may well be due to the canals, and is considerably larger in amplitude [40–70% in the dark (Collewijn et al, 1985c)].

This last component can be isolated by having the subject lie on his back, and measuring eye torsion during movements around a vertical axis (Merton, 1956); one cannot assume that this procedure eliminates all otolith contribution, since there is no reason in principle why the maculae of the utricle and saccule could not respond transiently to rotational as well as to linear accelerations. Conversely, as we have seen, any responses to a change in the direction of linear acceleration may possibly be due to canal as well as otolith function. Using sinusoidal acceleration along a linear track, Baarsma and Collewijn (1975) have measured the frequency transfer function of the otolith VOR in the rabbit, and find it to fall off rather rapidly with increasing frequency (figure 2.17). The shape of this function is sufficiently different from that of the canal vestibulo-ocular transfer function that it is unlikely that their measurements were contaminated by a canal artefact. In man a similar low-pass response is seen, with a turn-over frequency of some 0.3–0.4 Hz, the gain corresponding roughly with what would be needed to aid the fixation of fairly distant objects (Buizza et al, 1981; Lichtenberg et al, 1982). The otolith contribution is clearly slow and tonic in nature.

One might expect that a subject with bilateral destruction of the labyrinths and with the head fixed with respect to the body would show no torsion at all on lateral tilt. This is not in fact the case, and such people often show a residual response that may be due to some remnant of labyrinthine activity, to visual cues (see section 3.2), or conceivably to differences in pressure on the two sides of the body (Miller and Graybiel, 1962; Krejcova et al, 1971).

The COR, eye movement in response to movements of the body while the head is kept still, seems first to have been described by Barany (1906b). Clinically, nystagmus and dizziness after damage to the cervical dorsal roots have been described by Philipszoon (1962). However, exactly what the proprioceptors of the neck—mainly joint receptors rather than muscle

receptors (McCouch et al, 1951)—contribute is not altogether clear. In most species, including man, the responses obtained are rather small and there appears to be very great variability between subjects and under different conditions (Young et al, 1966; Takemori and Suzuki, 1971; Peterson et al, 1981); in the rabbit the gain falls from 0.2 to less than 0.05 in the natural frequency range (Barmack et al, 1985), though in the frog, which uses its head for stabilizing its gaze much more extensively, the gain is some 0.8–0.9 in the same range (Dieringer and Precht, 1982). An *anticompensatory* rather than compensatory reflex, moving the eyes in the *same* direction when the head is moved, has been described in several species, when the vestibular apparatus is intact (Gresty, 1976; Fuller, 1980; Peterson and Goldberg, 1982; Berthoz, 1985). Would one actually *expect* the COR to be ‘compensatory’ in this sense? If the head is fixed in space, it is clearly not desirable for body movements to cause eye movements. If the body is still and the head moves, then one might suppose that the neck receptors ‘help’




Figure 2.17. Gain and phase of ocular counterrolling in the rabbit in response to sinusoidal variation of the effective direction of gravity by appropriate linear accelerations in the horizontal plane, as a function of frequency. Filled circles are for lateral movement and open circles for longitudinal; the vertical bar in each graph represents a typical value of the standard deviation of a single observation (after Baarsma and Collewijn, 1975).

the VOR; but again, if this help consists simply of adding an extra signal to the vestibular one, the result will be only to falsify the vestibular head-velocity signal. But in the case of voluntary movements of the head, when what is required is a shift of the gaze in the same direction, then an additional anticompensatory push from the COR might well be useful (see section 2.3.2 below). There are several observations that suggest that cervical effects are in some way dependent on the integrity of the vestibular system. In the rabbit the nystagmus normally provoked by unilateral block of cervical afferents by local injection of procaine is abolished by bilateral labyrinthectomy (Biemond and de Jong, 1961). Similarly, in the decerebrate cat it is very difficult to record any response at all to twist of the neck if the labyrinths on each side are destroyed (Carpenter, 1972a). Although the neck contribution to responses to head movement in the monkey is normally virtually nonexistent, after bilateral section of the vestibular nerves a certain degree of relearning occurs, so that after some seven weeks nearly all the response is recovered by an increase in gain of neck reflexes and by differences in the ‘programming’ of eye and head movements (Dichgans et al, 1973, 1974; Dichgans, 1975). Neck information is more useful when it is replacing the vestibular system, rather than supplementing it unnecessarily.

2.3 Dynamic vestibulo-ocular responses

2.3.1 Vestibular nystagmus

In man, movements of the eyes occur much more readily in response to rotation of the head than to linear acceleration, and indeed form the basis of the commonest clinical tests for vestibular function. We have seen that over a wide range of frequencies the canals signal the angular velocity of the head in different planes. The function of the VOR is to match the velocity of the eye to that of the head, in an attempt to keep the image of the outside world stationary on the retina. However, although with rotating chairs and other laboratory equipment we can rotate the head through as large an angle as we please, the eye obviously cannot go on turning indefinitely in the orbit. In fact, during prolonged rotations the compensatory movements of the eye in the opposite direction are periodically interrupted by fast flicks in the direction of rotation, which thus enable the eye to continue its task of matching velocity even though it has reset its *position* in relation to the outside world. Thus a graph of the instantaneous deviation of the eye as a function of time during such a rotation has a characteristic sawtooth appearance consisting of a *slow phase* in the compensatory direction and a *quick phase* in the anticompensatory direction: the whole pattern of response is what is meant by *nystagmus* (figure 2.18).

Vestibular nystagmus has been known for a considerable time: it seems first to have been described in 1794 by Erasmus Darwin (the remarkable physician, inventor and poet, and grandfather of Charles). Conventionally—and perhaps illogically—nystagmus is characterised not by the direction of

the slow phase but by that of the anticompany quick phase. Thus a rotation of the head to the right also leads to a nystagmus to the right.

The response of the eye to rotations of the head is similar to that of the vestibular fibres themselves, the transfer function of head movement to eye movement being of the first-order high-pass kind. During a prolonged rotation at constant velocity, for example, the slow phase velocity gradually declines, with a corresponding decline in the frequency of the quick phases (figure 2.19). If the motion is suddenly stopped, a nystagmus in the opposite direction (*post-rotatory* or *after-nystagmus*) is seen, which in its turn dies away with a similar time course.

By rotating a subject sinusoidally in the dark at different frequencies we can get a quantitative description of the dynamic properties of the VOR in




Figure 2.18. A record of human vestibular nystagmus, showing the slow and quick phases.




Figure 2.19. Vestibular nystagmus: plots of behaviour accompanying a constant-velocity turn: (a) the time course of the velocity of the head, (b) the resultant deviation of the cupula, and (c)—very diagrammatically—the ensuing vestibular nystagmus: the frequency of the quick phases has been drastically reduced, for clarity. The velocity of the slow phase is approximately proportional to cupular deviation, at every point.




Figure 2.20. An illustration of how the quick phases can be removed from a record of vestibular nystagmus to provide a curve showing the cumulative effect of the slow phase on its own.

terms of its frequency response (Niven and Hixson, 1961; Young et al, 1966; Carpenter, 1972a; Baarsma and Collewijn, 1974; Shinoda and Yoshida, 1974; Skavenski et al, 1979; Donaghy, 1980; Paige, 1983a; Godaux and Vanderkelen, 1984; van der Steen and Collewijn, 1984; Correia et al, 1985). For very small rotations, and in decerebrate animals, the resultant eye movements are sinusoidal; larger movements produce quick phase interruptions which can be removed from the raw records to isolate the slow phase response (figure 2.20). The result is a quasi-sinusoidal movement, showing asymmetric nonlinearities as the amplitude of the stimulus is increased (Carpenter, 1972a; Mathog, 1972; Shinoda and Yoshida, 1974; Landers and Taylor, 1975; Paige 1983a). These asymmetries are particularly marked in the vertical (pitch) plane, the response being stronger when the head is tilting up (Anderson, 1981; Darlot et al, 1981; Matsuo and Cohen, 1984). In the linear range, one can proceed to calculate the gain and phase of the response of the system at different frequencies: in all preparations studied, the resultant frequency response is virtually flat over a middle 'physiological' range of frequencies (figure 2.21), with a gain that is typically between about 0.8 and 0.9 in the dark. It has been suggested (Skavenski et al, 1979; see also chapter 6) that the gain is 'deliberately' not

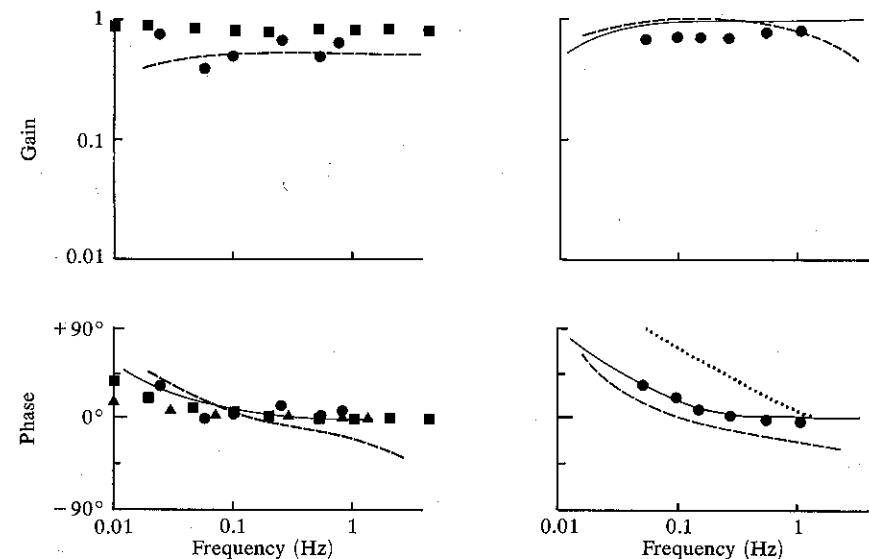


Figure 2.21. Transfer function of vestibulo-ocular reflex in different preparations. Left, primate: circles, man (Yasui and Young, 1984); continuous line, man (Benson, 1970); dashed line, man (Young et al, 1966); squares, monkey (Paige, 1983a); triangles, monkey (Skavenski and Robinson, 1973). Right, cat: circles, normal (Godaux and Vanderkelen, 1984); continuous line, normal (Landers and Taylor, 1975); dashed line, decerebrate (Carpenter, 1972a); dotted line, anaesthetised (Sugie and Melvill Jones, 1971).

unity, so that head movements may be permitted to contribute something to the retinal slip necessary to overcome the adaptational properties of the visual receptors. Another factor that must be borne in mind is that the axis of rotation of the skull is further from the visual target than that of the eye (by about 10 cm in man), so that for objects that are not at infinity the 'correct' gain is in any case less than unity (for example, about 0.5 for a target 10 cm away); there is some evidence that VOR gain is indeed affected by target distance (Blakemore and Donaghy, 1980; Biguer and Prablanc, 1981).

There is slightly less agreement about what happens at very high frequencies. One might expect to find an increase in gain corresponding to the high-frequency adaptional component of the canals mentioned earlier (section 2.1.4), and also perhaps on account of the inertia of the eyeball itself, which in this context is actually helpful. Although Keller (1978) reported a gain in the monkey that actually exceeded unity at around 4 Hz, others (Barr et al, 1976; Skavenski et al, 1979; Donaghy, 1980; Paige, 1983a) have not observed increased gains at high frequencies. Above some 2 Hz, the low-pass mechanics of the eye and the lags introduced by conduction delays would in any case be expected to dominate any phase lead originating from the canals.

Measurements of *phase* show some variation between different species (figure 2.21): for example, at just under 0.1 Hz, an etherised cat shows nearly 90° phase lead compared with an alert monkey (Sugie and Melvill Jones, 1971; Skavenski and Robinson, 1973). Some of this conflict, particularly at higher frequencies, is undoubtedly due to differences in procedure of phase measurements in the presence of nonlinearities. Another factor that may be relevant is that some of the preparations make quick phases during sinusoidal stimulation (which then have to be removed by the process shown in figure 2.20), whereas in others the quick phases are naturally absent. In the first case, we have to make the assumption that the quick phase is in effect only a brief displacement of the whole record: but it *might* have some disturbing effect on the stretch of record immediately following it, resulting in a distortion of the apparent phase. This is the explanation given by Sugie and Melvill Jones for the large phase advance of the slow phase seen when the depth of anaesthesia of their cats is increased to the point where the quick phases disappear: they suggest that each quick phase consists of a step of displacement followed by an exponential decline with a time constant of about 0.7 s, which is added to the slow phase component. Yet another factor likely to contribute to this variability is the remarkable degree of plasticity shown by the VOR under different adapting conditions, which in an extreme case can produce phase changes as large as 120° over the course of a week or two (see chapter 13). Another phenomenon which is perhaps best regarded as an example of plasticity is very slow adaptation to periods of constant angular acceleration, with time constants in the region of a minute or two (Malcolm and Melvill Jones,

1970; Stockwell et al, 1973; Landers and Taylor, 1975), so slow that they cannot contribute much under natural conditions of stimulation. Finally, it is perhaps worth emphasising that, although the great majority of work on the VOR has been done on horizontal movements of the eyes, nystagmus can equally be observed both in the vertical plane and [with some experimental sophistication (Davies and Merton, 1958; Melvill Jones, 1960; Petrov and Zenkin, 1973)] for torsional movements. The decay characteristics of vestibular nystagmus during constant accelerations are markedly different for rotations around different axes, no doubt because of the differences between the mechanical properties of the three canals mentioned above in section 2.1.4, and also because of differences in the degree of velocity storage (see below); the direction of gravity may also exert an independent effect (Benson and Bodin, 1966b; Benson, 1970; Darlot et al, 1981; Correia et al, 1985).

Although the time course of the eye in the VOR and of the firing of vestibular fibres are qualitatively similar, it turns out that at lower frequencies the eye actually does *better* than the signal from the canals would lead one to expect: in the monkey the time constant for the canals is some 5–6 s, but for the VOR, more than 20 s (Büttner and Waespe, 1981). This stretching of the time constant is often called *velocity storage*. A very similar phenomenon exists in the case of nystagmus evoked by linear acceleration, where the time constant is again greater than that of the responses from the afferent fibres, and also in optokinetic nystagmus (section 3.2.2), and it can be demonstrated already at the level of the vestibular nuclei for both kinds of stimulation (Waespe and Henn, 1977; Buizza et al, 1981; Büttner et al, 1981; Koenig and Dichgans, 1981). It seems overwhelmingly likely in fact that the same neural mechanism underlies all these phenomena, and that it may well consist of some kind of positive feedback loop whose output is an internal representation of head velocity in space, whether provided by the vestibular or the visual system (Robinson, 1977; Demer and Robinson, 1983; Paige, 1983a; Raphan and Cohen, 1985): the decay of its activity is predominantly linear rather than exponential (Collewijn et al, 1982a; Baloh et al, 1983; Collewijn, 1985). Possible mechanisms for the phenomenon are discussed in chapter 12.

It has often been suggested that the falloff in vestibular gain at lower frequencies is a 'deliberate' feature of the VOR (ter Braak, 1936; Melvill Jones, 1965; Baarsma and Collewijn, 1974). Optokinetic reflexes normally assist the vestibular system in maintaining fixation during head movements, but the visual contribution, though accurate, is very slow on account of both the sluggishness of the visual receptors themselves, and the many stages of neural processing that must intervene before visual signals can reach oculomotor pathways. As we shall see in chapter 3, the frequencies at which the visual component starts to fall off are similar to those at which the vestibular component begins to turn over onto its plateau, so that the two mechanisms working together can by simple addition of effects provide full

compensation of head movements right down to zero frequency. Equally, of course, one might argue that it is the low-pass filtering of the visual signals that has been matched to the given constraints in the semicircular canals (Robinson, 1977). Either way, the purpose of velocity storage is perhaps less to do with improving the frequency response of the whole thing, than with maintaining a consistent internal representation of head velocity from the diverse and often intermittent sources of information that are available to the brain, a view that is perhaps more compelling in connection with the optokinetic system than in the case of the VOR (section 3.2.2).

2.3.2 Natural head movements

Natural head movements are quite unlike what is experienced in a typical rotating chair experiment; nor, for that matter, are they much like sinusoidal rotations. Furthermore, most head movements are made with the express *purpose* of shifting the gaze, and indeed under free conditions the contribution of the head is about twice that of the eyes in looking from one object to another (Biguer et al, 1984). What then is the point of elaborate mechanisms for keeping the eyes still? Here again it is helpful to distinguish between the control of *velocity* and the control of *position*. Obviously a change in the position of fixation implies that at least for a limited period there must be relative motion between the retina and the visual world: the shorter this period is, the less visual information will be lost during the shift of gaze. But head movements must necessarily be much slower than eye movements because of the enormously greater moment of inertia of the head in comparison with that of the eye (see for example the data of Stark et al, 1980). The solution then is to use vestibular reflexes to keep the *velocity* of the eye at all times equal and opposite to that of the head, but at some point during the movement to superimpose on this basic response a *position* change in the same direction as the head movement that is as fast as the oculomotor system can manage. In this way the shift of gaze is achieved with the minimum possible period during which the visual system is out of action. Recordings of eye movements during natural movements of the head confirm that this is indeed what happens (Atkin, 1964; Morasso et al, 1973) and that, as a consequence, during the head movement the gaze moves smartly from one steady position to another. Movements of this kind can be represented graphically rather economically as a locus on the *gaze plane*, a plot of eye position against head position (Zangemeister and Stark, 1982b).

One might suppose that the fast anticomensatory movement was simply the same thing as the old quick phase of nystagmus, in a new guise. But a simple experiment makes it clear that the quick part of the response is, as it were, 'preprogrammed', and is not the result of vestibular stimulation at all (Morasso et al, 1973). If the subject's head is suddenly and unexpectedly restrained at the moment of beginning the head movement, the eye nevertheless moves smartly to the new position of gaze exactly as before; the slow vestibular component is of course abolished (figure 2.22).

In other words, the fast and slow components of the eye movement have quite different origins: the quick movement, like the head movement itself, is a *primary* movement, in the sense that it is the direct result of the desire to shift the gaze: but the slow component is *secondary* to the head movement and does not occur if the head is fixed. The quick component may indeed start before the head has moved (Bizzi et al, 1972), although the electrical response from the neck muscles generally precedes that from the eye muscles. All the same, one cannot rule out a vestibular contribution to the quick component, since the same anticomensatory movement is observed when the head is moved suddenly by the subject *without* the intention of shifting the gaze to a peripheral object (Henriksson et al, 1974). In the same way, Melvill Jones (1964) found that sudden externally applied turns of the head also produce an anticomensatory flick preceding the slow compensatory component (admittedly with rather large angular accelerations); comparable findings have been published by Barnes (1974, 1981). Similar responses have been described as a result of electrical stimulation of the ampullary nerve in the alert monkey, although in this case the quick component follows the slow, and may merely represent an isolated nystagmus quick phase, or the monkey's voluntary effort to refixate what he was looking at before the stimulus (Cohen et al, 1967a; Goto et al, 1968).

It is clearly desirable that the ordinary VOR should operate to stabilise the gaze before and after the saccadic component of this response, but what about *during* it? On the one hand, one might think that it would be easier for the system to deal with the saccade as if it were an extra signal superimposed on a continuing vestibular background; but on the other hand, by turning the VOR off during the saccade, the gaze will get to its final position more quickly, and vision will be less disrupted. Experiments show that the latter is what actually happens, and that as a consequence, the gaze is shifted more quickly than if the head is still. If the head movement is unexpectedly slowed during the saccade, gaze velocity drops by the same amount; and if sudden forced head movements are introduced during saccades—for example, by hitting the bite-bar—no compensatory VOR is seen (Roucoux et al, 1981b; Lauritis and Robinson, 1986). This implies a mechanism of




Figure 2.22. Head and eye movements made when looking at novel objects: (a) eye movement (saccade) made in fixating an eccentric target with the head fixed; (b) head (H) and eye (e) movements made under exactly the same conditions, but with the head free; E is gaze, the sum of H and e (after Morasso et al, 1973).

suppression in which the VOR is temporarily switched off, which is rather different from the *cancellation* found during smooth tracking with free use of the head and eyes. Here, what is found is that soon after bilateral destruction of the labyrinths the velocity of the gaze during tracking is actually greater than that of the target, implying that there is a signal derived from the intention to move the head that is trying to cancel a VOR that is not in fact present (Lanman et al, 1978; Tomlinson and Robinson, 1981; Robinson, 1982); the phenomenon cannot be as easily shown in patients with *chronic* labyrinthine deficiency, presumably because of long-term parametric adaptation (Leigh et al, 1987).

One should not forget that the movements of the head relative to the body are also under vestibular control, and at least in some species one can demonstrate nystagmus of the *head* as a result of this *vestibulocollic* reflex (see reviews by Peterson et al, 1981; Peterson and Goldberg, 1982; Berthoz, 1985). There are two reasons why these movements are generally not as prominent as those of the eyes. First, that most animals' heads are mechanically dominated by their inertia, and cannot easily execute high-frequency motion; and second, that the vestibulocollic reflexes have intrinsic negative feedback, in that any compensatory movement of the head necessarily reduces its own stimulus. The prominent head nystagmus of birds when walking was mentioned in section 1.1: but this is primarily visual rather than vestibular in origin (Dunlap and Mowrer, 1931). The *quick* phase of head nystagmus is, however, anticomensatory, and the extra vestibular stimulation it provides may cause a certain degree of synchrony between nystagmus of the head and of the eyes (Hinoki and Terayama, 1966a; 1966b; Outerbridge and Melvill Jones, 1971). Finally, to complicate things further, recordings from neck muscles suggest the existence of a kind of 'oculocollic' reflex, a tonic drive to the head that is strongly linked to horizontal eye position (Vidal et al, 1982; 1983; Lestienne et al, 1984).

2.3.3 The nature of the quick phase

How are the rhythmic movements of nystagmus produced? Two possible mechanisms come to mind. The first notion, expressed in some of the writings of Lorente de No (1933 for example) and held by implication if not explicitly by some more recent authors (for example McCabe, 1965), is that nystagmus is the product of a specifically organised neural oscillator—perhaps arranged in a similar manner to an astable multivibrator—that is set running by impulses from the semicircular canals. In Lorente de No's own words (1933, page 289): "Nystagmus is nothing else than an alternating reflex similar to other known rhythmic reflexes. The peripheral labyrinthine stimulation sets into activity a machinery which gives rise to nystagmus in the same way as the spinal cord sets up the scratch reflex as a response to a stimulus produced in the skin." The other view, advanced first by Barany (1907) and now more commonly held, is that the slow and quick phases are two quite separate phenomena, the slow phase being a continuous

background response to vestibular stimulation, on which the quick phases are superimposed by a different and distinct mechanism.

Perhaps one should not exaggerate the difference between these two concepts, which is largely one of emphasis: Lorente de No himself was fully aware, for example, that there are a large number of different circumstances under which the quick phase can be abolished, leaving the slow phase unaffected. These include brain lesions (for example Holmes, 1938; Wadia and Swami, 1971; Carpenter, 1972a) and sleep (di Giorgio, 1935; Nathanson and Bergmann, 1958); while many drugs have differential effects on fast and slow phases (Bender, 1955; Rashbass and Russell, 1961; Philipszoon, 1962; Janeke et al, 1969; Melvill Jones and Sugie, 1972; Haliska, 1973). Furthermore, although it is true that the frequency of nystagmus is roughly correlated with the speed of the slow phase, the relation is very far from being the close one that would be expected if both were the product of the same neuronal oscillator (Torok and Derbyshire, 1968; Tibbling, 1969; Sugie and Melvill Jones, 1971). It seems difficult to escape the conclusion that the slow and quick phases are independent entities, and not merely, as it were, the systole and diastole of an *essentially* rhythmic process.




Figure 2.23. Eye movement (upper trace) in response to sinusoidal head velocity (lower trace), showing the lack of regular correspondence between the stimulus and the moments at which the quick phases are initiated (Gonshor and Melvill Jones, 1971).

All the same, it is by no means clear exactly what the mechanism is by which the quick phases are repetitively generated during nystagmus. One might expect for example that the quick phase would be triggered when the eye reached a certain limiting deviation in the orbit, perhaps signalled by stretch receptors in the eye muscles. But complete removal of the eye and its muscles from the orbit has virtually no effect on the rhythmic pattern of discharge that can be recorded from the oculomotor neurons during vestibular nystagmus (McIntyre, 1939), and in any case, inspection of records of vestibular nystagmus in response to different patterns of rotation shows quite clearly that the quick phases are not in fact initiated at fixed deviations of the eye (figure 2.23). Further, Ohm showed in 1936 that if a small-amplitude high-frequency nystagmus is evoked at the same time as a large-amplitude low-frequency one—by two distinct modes of stimulation—the two sets of quick phases show no interaction (as would be expected if they were generated by the total resultant slow phase component) but

behave as if they were independently generated and added together in a linear manner. The result is thus a large sawtooth waveform with a smaller sawtooth on its edge.

During sinusoidal stimulation one finds that the point at which the quick phase is initiated varies systematically in each cycle, in such a way that the *mean* position of the eye itself moves sinusoidally, advanced in phase by some 90° relative to the sine wave traced out by the slow component (Mishkin and Melvill Jones, 1966). This suggests that the quick phase may be triggered when the velocity of the head (and thus the signal from the canals) reaches some fixed threshold value, and it seems also that the size of the quick phase is similarly related to the velocity of the head rather than to the position of the eye; an elaborate model of such an arrangement has been published by Chun and Robinson (1978). Possibly the mechanism of this phase advance is the same as the one producing the anticomensatory quick phase at the beginning of rapid head movements, discussed in the previous section; the matter is considered further in chapter 12.

A complicating factor is the observation that the mechanism generating the quick phases also appears to have an inherent *periodicity* (Cheng and Outerbridge, 1974b). When the stimulus is large (and steady) the distribution of intersaccadic intervals (that is, of the intervals of time between successive quick phases) is essentially unimodal. But as the degree of stimulation is increased, the distribution begins to break up into a series of peaks of uniform spacing, about half a second apart. This suggests that the quick phases are generated by a mechanism having an inherent rhythmic tendency with a frequency of about 2 Hz. It is difficult to see what the point of such an arrangement might be; conceivably it could be related to the intermittency associated with normal fixation saccades (see below, section 4.2.1)

Optokinesis and smooth pursuit

"It is not good to beholde things that move swiftilie, nor yet such as turne round."

This chapter is concerned with the mechanisms by which retinal slip causes movements of the eyes that tend to hold the image relatively stationary on the receptors. They are rather less reflex in character than the vestibular movements, for they are capable of being modified by an effort of will: but on the other hand, most subjects are quite unable to *generate* them voluntarily in the absence of an appropriate stimulus, and in that respect they are somewhat less 'voluntary' than saccades or vergence movements.

3.1 Optokinesis versus pursuit

It is customary to distinguish between *optokinesis*, evoked when a large part of the image moves uniformly over the retina, and *smooth pursuit*, in which the eye follows a small target. The main basis for the distinction is a functional one: optokinesis helps to stabilise the entire retinal image during head-turning or locomotion, whereas smooth pursuit is in addition concerned with the *foveation* of moving targets, which it tracks in cooperation with saccades. It is therefore not very surprising that whereas optokinesis is seen in all creatures with moveable eyes, smooth pursuit is only demonstrable in animals with a small central retinal area. Rabbits, with their extended visual streaks, show optokinesis but not smooth pursuit (Collewijn, 1969; Dubois and Collewijn, 1979a). Dogs will not follow a moving stripe with their eyes, but may do so if the stripe is replaced by a picture of a rabbit: even this may not work if the particular subject does not find rabbits sufficiently interesting (Rademaker and ter Braak, 1948). Optokinesis is on the whole less fussy, to the extent of tolerating as much as 20 diopters of misfocus (Dichgans, 1977); but in some animals, certain optokinetic patterns are more effective than others (Maioli and Precht, 1984; Sirkin et al, 1985a).

Furthermore, in natural circumstances the two mechanisms are at odds with one another. Consider for example a cat intent on a mouse that is crossing a patch of undergrowth. The mouse is small, the background immense and detailed. If there were mere summation of effects between background and the mouse, it is clear that optokinesis would hold the cat's eyes firmly to the undergrowth, and the mouse would pass with impunity: any attempt at smooth pursuit by the cat would result in retinal slip of the background in the opposite direction, powerfully refixating the eye. Smooth pursuit is thus a 'higher', overriding, function, associated with more recent areas of the brain and more easily influenced by attention or the will. Indeed in man the same moving target may give rise to movements characteristic of either response, depending on the subject's attitude to the task (Pola and Wyatt, 1985).

3.2 Optokinesis

The easiest way to demonstrate optokinetic nystagmus in the laboratory is to observe the eye movements that are made whilst a subject is viewing a rotating striped drum that fills a substantial part of the visual field [better still, get the subject to sit inside it: see for example Lasker and Kummel (1985)]. As long as the drum is not moving too fast, the subject's eyes tend to follow the drum as it moves; eventually the gaze is carried too far from the primary position, and a fast anticompany flick is made that brings it back. Thus the time course of the movement is of the sawtooth form that we have already seen to result from head rotation (figure 3.1) and is called *optokinetic nystagmus* or OKN: Purkinje (1825) seems to have been the first to describe it, in the eyes of a crowd watching a procession of cavalry. Any detailed moving scene will do equally well—the view from a railway carriage window is excellent—and in fact moving patterns of randomly arranged dots have some advantages over stripes for experimental work, because the particular spacing of the stripes strongly influences the pattern of quick phases that is obtained (Cheng and Outerbridge, 1974a).

Optokinetic nystagmus occurs with equal facility in horizontal and vertical meridians, but around an anterior–posterior axis is very small in amplitude (Collewijn et al, 1985c): a more prominent response to a visual field rotating in a plane parallel to the frontal plane is a steady torsional deviation of the eye, reaching a maximum of some 2° at about 30° s^{-1} (figure 3.2) (Kertesz and Jones, 1969).

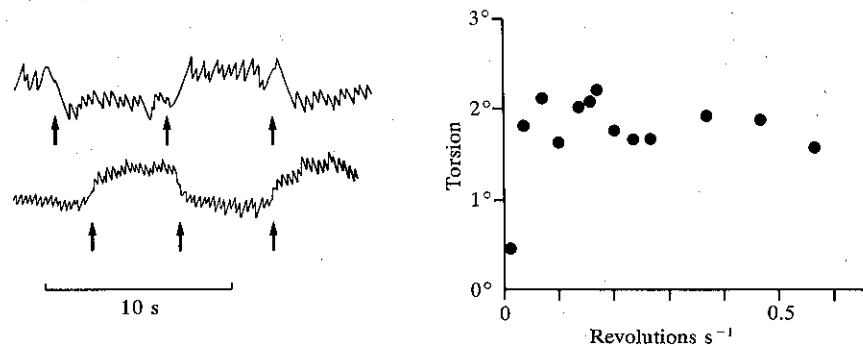


Figure 3.1. Human optokinetic nystagmus under two different conditions: in the upper trace, the subject was actively trying to follow the stripes of the drum; in the lower trace, the subject was gazing passively at it. The arrows indicate moments at which the movement of the drum was reversed: note the *average* deviation in the direction of drum movement in the first case, and in the opposite direction in the second (Hood and Leech, 1974).

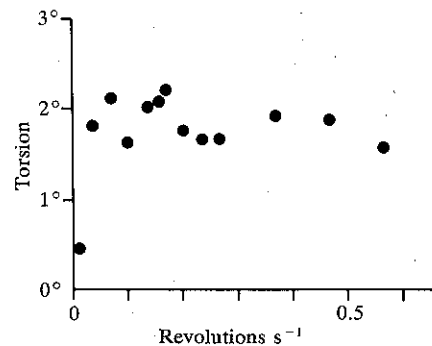


Figure 3.2. Steady torsional response of the human eye to a field rotating in the plane of fixation at different angular velocities (data from Kertesz and Jones, 1969).

The form and degree of OKN is strongly influenced by the subject's attitude to the task, and by the nature of the instructions given. Human subjects may be instructed either to *look* at the stimulus, to *stare* at it (that is, without making any effort to follow it with the eyes), or—a more difficult task—to match the velocity of their eyes to that of the target. This last command, paradoxically, produces the poorest correspondence between target velocity and eye velocity, leading to large cumulative positional errors (Puckett and Steinman, 1969). If the subject 'looks' at the drum, eye movements are large in extent and quick phases relatively infrequent; if instead the subject 'stares' at it, the nystagmus is small in amplitude and more frequent (figures 3.1 and 3.3) (Rademaker and ter Braak, 1948). What the subject *cannot* do is suppress the nystagmus altogether, except by fixating a stationary target in front of the drum; or given two superimposed patterns moving independently, he or she may choose which to follow (Murphy et al, 1975; Kowler et al, 1984b). However, a rabbit appears to be able to inhibit its optokinesis entirely when it is 'freezing' (Brecher, 1936).

At moderate speeds (less than some 30° s^{-1}) the velocity of the slow phase is close to the speed of the drum [see Honrubia et al (1968); similar figures are obtained for the cat and rabbit, much lower ones in the rat (Collewijn, 1972; Evinger and Fuchs, 1978; Fuller, 1985)]. With faster drum speeds, the slow phase lags more and more behind the movements of the drum, until at some 100° s^{-1} the system breaks down completely and only inappropriate movements—if any at all—are made (for example Dodge et al, 1930; Honrubia et al, 1967; and many others). These maximal velocities are not achieved immediately, but take time to build up, for reasons that will become apparent; the time course of the buildup often shows two components, one immediate (*early* or EOKN) and one delayed (*late* or LOKN), the relative size of each showing considerable interspecies variation (Collewijn, 1985) (figure 3.4). LOKN is normally evoked by optokinetic stimulation of large areas of the visual field and is typically more prominent

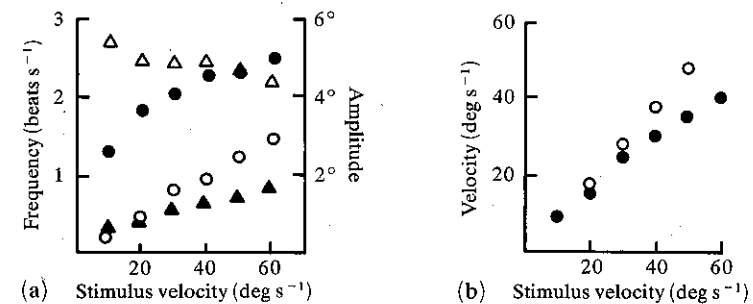
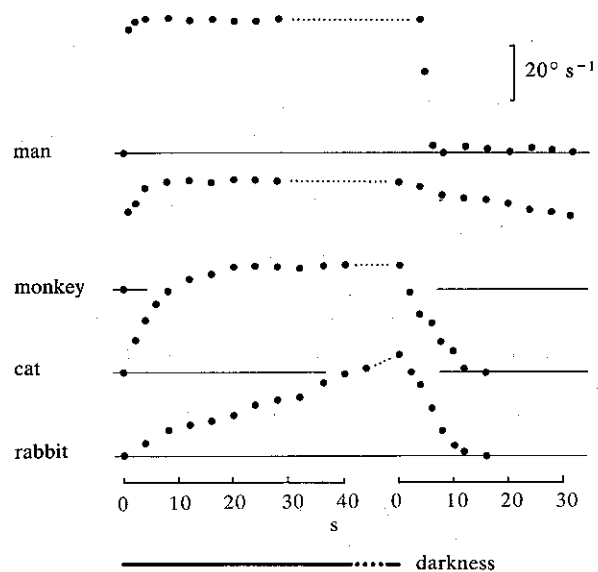



Figure 3.3. Optokinetic responses as a function of stimulus velocity. Filled symbols indicate conditions when the subject was told to *stare* at the drum; open symbols, when the subject was to *look* at it. In (a) the frequency (circles) and amplitude (triangles) are shown under the two conditions: in (b), the velocity of the slow phase (data from Honrubia et al, 1967).



(a)



(b)

Figure 3.4. LOKN, EOKN, and OKAN. (a) Slow-phase velocity of the eye during a 60° s^{-1} optokinetic stimulus (thick horizontal line), in man, monkey, cat, and rabbit (data from Collewijn, 1985). (b) Open-loop slow-phase velocity responses to pattern movement of constant velocity (\blacktriangle , 16° s^{-1} ; \bullet , 3.6° s^{-1}) in two opposite directions in the monkey. At the instant marked by the broken line, the pattern movement was stopped and all illumination extinguished. The eye velocity built up more or less linearly during stimulation, and gradually declined in the dark (optokinetic afternystagmus, OKAN) with a distinct overshoot (OKAN II) in which the direction of nystagmus was reversed (after Koerner and Schiller, 1972).

in animals with panoramic vision, whereas EOKN can be generated by small targets and is found in animals with a central retinal area: LOKN is suppressed by stationary contours in the visual field, EOKN is not (Dichgans, 1977). It is evident that there are strong similarities between EOKN, active smooth pursuit, and perhaps 'look' optokinesis on the one hand, and LOKN, 'stare' nystagmus, and passive smooth pursuit on the other, no doubt reflecting differences in the neural mechanisms serving each general class of movement (Pola and Wyatt, 1985).

On stopping the drum, the nystagmus quickly comes to a halt; but if the room is simultaneously plunged into darkness, so that there is no target to hold the eye in place, the OKN continues with decreasing velocity for some 20 s or more (figure 3.4) (Fox et al, 1931; Krieger and Bender, 1956; Mackensen and Wiegmann, 1959; Cohen et al, 1977; Maioli and Precht, 1984). In the rabbit the decline of this *optokinetic afternystagmus* or OKAN is apparently linear rather than exponential (Collewijn et al, 1980); one may also distinguish two or even three phases of OKAN, called I, II and III, between which the direction of the movement is reversed (Koerner and Schiller, 1972; Maioli and Precht, 1984; Collewijn, 1985). On the whole, the amplitude of OKAN is correlated with that of LOKN, and it is highly probable that both are due to the same mechanism of *velocity storage*, in which the velocity of the eye is partly determined by a signal that is a low-pass filtered version of the retinal slip velocity and thus builds up slowly and decays slowly. This slow component can be monitored in isolation by interrupting the optokinetic stimulation from time to time with periods of darkness; the buildup can then be seen directly, rather than having to be inferred (Segal and Liben, 1985). Velocity storage is discussed in more detail below.

3.2.1 Dynamic properties

Optokinesis provides a very nice example of a control system with built-in negative feedback. The visual receptors that provide the input to this system must of course move with the eye; so they signal not the absolute position or velocity of a stimulus target, but rather its position or velocity *relative* to the eye itself. Thus the signal from the visual receptors constitutes an *error signal* telling the control system how adequately it is compensating for the movements of the target.

If we assume for the moment that the system that converts this error signal into a corrective movement of the eye is linear, and has a forward gain (relating the slip velocity of the retinal image, \dot{a} , to the resultant eye velocity in the head, \dot{e}) of G , then we can represent the whole system in the closed-loop form of figure 3.5. It is a simple matter (see appendix 2) to calculate what the gain of the whole system (that is, from target velocity relative to the head, \dot{a} , to eye velocity \dot{e}) will be; it is

$$\frac{\dot{e}}{\dot{a}} = \frac{G}{1+G}$$

If G is very large, then it is clear from this expression that the overall gain will then approach unity, and the eye velocity will be close to that of the target. As we have seen, when looking at targets moving with a constant velocity of less than some 30° s^{-1} this is substantially true. Under these conditions, a small change in the value of G will make relatively little difference to the overall performance. To take a specific example, if G is simply a pure gain of 20, the gain of the whole thing can be seen to be $20/21$, or about 95.24%. If G is increased by a quarter of its value to 25, the overall gain then becomes $25/26$, or roughly 96.15%. Thus a 25% increase in the forward gain results here in less than 1% change in the closed-loop gain.

What this implies is that the *overall* performance—the relation between target and eye motion under natural, closed-loop, conditions—is likely to tell us rather little about the nature of the control system itself. For the behaviour of the whole thing is more determined by the fact that it *is* closed-loop than by the precise nature of G . The most useful studies of the control system have got round this difficulty by artificially ‘opening the loop’, converting the system into the form shown in figure 3.6. There are essentially two ways of doing this. One is to immobilise the eye that is viewing the target, and record the movements of the other eye (Ohm, 1926); but one cannot be sure that the lack of normal proprioceptive signals from the fixed eye is not affecting one’s results. Alternatively one can try to compensate for the eye movements by moving the target through the same




Figure 3.5. The closed-loop nature of optokinesis under natural conditions: on the left, an object moving with angular velocity \dot{a} is tracked with an eye velocity of \dot{e} . The resultant retinal image slip velocity $\dot{a} (= \dot{a} - \dot{e})$ is transformed by the operator G (which represents the central optokinetic mechanism) into the resultant eye velocity \dot{e} . The formal representation of the system is thus as shown at the top right: it is equivalent to a single operator $G/(1+G)$.

angle that the eye moves, in addition to any extra motion stimulus that one may wish to superimpose on this stabilised background (for example Collewijn and van der Mark, 1972; Pola and Wyatt, 1985). The two methods in fact lead to similar—and striking—conclusions. If the immobilised eye views a striped drum rotating at constant velocity, the velocity of the slow phase of the resultant nystagmus is found not to be constant in magnitude, but to increase steadily until it reaches a value typically very much greater than the stimulus velocity (ter Braak, 1936; Koerner and Schiller, 1972) (figures 3.4 and 3.7), when the angular velocity of the eye may be as great as 100° s^{-1} . In an animal such as the rabbit in which optokinesis appears to be almost entirely LOKN, this steady and uniform increase in velocity implies that the forward pathway is performing something like the time integral of the error signal to determine the eye velocity, or in other words that the open-loop gain is actually of the form $G_0 D^{-1}$: experimentally, it happens that the value of G_0 in the rabbit lies very close to unity (see figure 3.7).

Now the system cannot *really* be just a linear integrator: if it were, then with steady optokinetic stimulation under open-loop conditions the velocity of the eye would increase without limit. What actually happens is that the velocity of the eye levels off at a steady value that is normally many times greater than that of the target. This is partly because of *saturation* in the system (discussed in the next section), and partly because what we have is




Figure 3.6. Two ways of opening the intrinsic optokinetic feedback loop: (a) the viewing eye is immobilised, and the movements of the other are recorded; (b) the target is made to move with an added velocity that is arranged to be always equal to the eye’s own velocity, \dot{e} . In both cases the transfer function of the whole system is simply that of the open-loop operator G .

not a 'pure' integrator, but a *leaky* one, with a transfer function given not by $G_0 D^{-1}$ but by $G_0(\mu + D)^{-1}$: a low-pass filter. If we assume that what we are seeing during OKAN is simply the charge of the integrator leaking away, then the time constant of this decay should fix the value of μ . Although it is true that OKAN in different species lasts of the order of tens of seconds (Cohen et al, 1977; Collewijn, 1985; Segal and Liben, 1985), in the rabbit the time course of decline does not behave in the exponential way expected of a simple linear system and, as in the case of postrotatory vestibular nystagmus, is better described as a linear decay of some $2-5^\circ \text{ s}^{-2}$ (figure 3.8) (Collewijn et al, 1980): the integrator has more of a bilge pump than a leak!

Another way of determining the dynamic parameters of the system is to use sinusoidally-moving targets of different frequencies. Open-loop measurements of this kind have been made in the rabbit (Collewijn, 1969) and in man (Pola and Wyatt, 1985), and closed-loop measurements by many authors, including (in the rabbit) Baarsma and Collewijn (1974), in the cat, Godaux et al (1983a), and in the monkey, Paige (1983a). As has already been emphasised, the calculation of open-loop transfer functions from closed-loop data is liable to introduce large errors when the closed-loop gain is close to unity, as is the case at low frequencies. A further difficulty is the fact that the system responds only poorly to high stimulus velocities: since the peak velocity of a sinusoidal stimulus increases with frequency, such a velocity limitation is liable to make itself felt as an artificially steep high-frequency cutoff. Consequently it is best to ensure that at each frequency the peak velocity is kept constant, but this makes it difficult to interpret the transfer functions obtained: indeed Collewijn (1969) found that gain was much more correlated with stimulus velocity than with frequency, which rather makes a nonsense of the whole concept of frequency response. His open-loop measurements agreed qualitatively with




Figure 3.7. Open-loop responses to constant-velocity optokinetic stimulation in the rabbit: (a) timing marks and drum position; (b) resultant nystagmus; (c) the velocity of the slow phase of (b) is plotted as a function of time. It is clear that the velocity rises linearly during the first 15 s; the line is the expectation if the open-loop operator were simply an integrator of unity gain (data from ter Braak, 1936).




Figure 3.8. Decay of OKAN in the rabbit. OKAN velocity as a function of time after extinguishing an optokinetic target moving at 60° s^{-1} to right (above) or left (below): the line shows the least-squares best-fit linear interpolation (data from Collewijn et al, 1980).




Figure 3.9. Closed-loop transfer functions for optokinesis. ▲, ▼ two cats, amplitude $\pm 2.7^\circ$ (data from Evinger and Fuchs, 1978); △ cat, amplitude 3° (Godaux and Vanderkelen, 1984); ○ monkey, average of six subjects at a peak velocity of 40° s^{-1} (Paige, 1983a); ● man (Yasui and Young, 1984).

what would have been predicted from the closed-loop measurements, but not quantitatively. Comparisons with other species are difficult because of the variable additional contribution that may be made by EOKN. If we assume that LOKN and EOKN represent two systems operating in some sense in parallel, then it seems likely that EOKN also contains a low-pass filter. Its time constant would have to be considerably shorter than the 10 s or so of LOKN, yet its cutoff steep enough to produce no appreciable effect on the overall high-frequency response (see also Gillis et al, 1984). However, it is also possible to explain the apparent lack of EOKN contribution at high frequencies by assuming that it is even more limited in its velocity sensitivity than LOKN (Buizza and Schmid, 1985). Closed-loop transfer functions in various species are shown in figure 3.9.

Figure 3.10 shows the *simplest* kind of linear model that might be taken to represent the optokinetic system (cf Buizza and Schmid, 1985); however, many authors have preferred representations in which the LOKN integration is performed by means of a positive feedback loop (for example Demer and Robinson, 1983; Paige, 1983a; and Gillis et al, 1984; the idea perhaps originated with Fender and Nye, 1961). An equivalent scheme of that kind is shown in figure 3.11: it behaves indistinguishably from the other, with appropriate adjustment of the parameters. Nevertheless there are strong




Figure 3.10. A relatively simple representation of the optokinetic system.




Figure 3.11. A more complex representation of the optokinetic system, where the integration is performed by a positive feedback loop in which \dot{e} , an internal representation of the eye-in-head velocity, is added to retinal slip velocity \dot{a} to form an estimate \dot{a}' of target velocity \dot{a} relative to the head; this in turn drives the eye.

arguments in favour of the more complicated arrangement—presented most eloquently by Robinson, 1977—and these are discussed in chapter 12, particularly in relation to the way in which optokinesis and vestibular reflexes share out the job of limiting retinal slip velocity.

3.2.2 Nonlinearities

Two kinds of stimuli are capable of generating large signals at one or other part of the system, and thus of driving it into nonlinearity. These are, on the one hand, stimuli having a large instantaneous velocity at some point in their time course and, on the other hand, stimuli of smaller magnitude that are unidirectional and prolonged. The first kind of signal appears to overload the input to the system: in fact with input velocities greater than some 100° s^{-1} the optokinetic response breaks down completely, suggesting a nonlinearity at the input that is not merely a saturation but actually shows a *decline* in output with increasing stimulus velocity. This decline can be observed directly in the velocity-sensing neurons themselves (for example Oyster, 1968). Such behaviour might be positively useful in making the smooth pursuit system insensitive to saccadic movements of the eyes.

Saturation with prolonged stimulation can most readily be investigated under open-loop conditions (figure 3.12): at very low input velocities the output levels off at a value proportional to the input velocity, the final DC gain being of the order of one hundred or so in the rabbit (Collewijn, 1969) but apparently less in man (Hood and Leech, 1974; Leigh et al, 1982). Above about 0.1° s^{-1} the gain drops steadily so that in the region $10\text{--}100^\circ \text{ s}^{-1}$ the eye velocity is saturated and independent of target velocity. It is very difficult in experiments of this kind to disentangle effects due to saturation at the input from effects due to saturation at the output, or




Figure 3.12. Open-loop response to constant target velocity: the graph shows maximum eye velocity \dot{e} as a function of retinal slip velocity \dot{a} in the rabbit (squares: Collewijn, 1969) and in two monkeys (triangles: Koerner and Schiller, 1972). The gain is very high at low velocities, but saturates when retinal slip is more than some 0.5° s^{-1} .

either from the leakiness of the integrator. Another possibility is that the leak itself is nonlinear, a hypothesis that would simultaneously explain why the decay of OKAN is not exponential. It is perhaps worth emphasising that nonlinearities of this kind are in no sense 'faults' in the system, irksome though they can be to the experimenter trying to investigate it: optimum controllers are almost always nonlinear in behaviour, since the fastest responses are achieved by using the biggest possible control signals.

3.2.3 The effect of field size

It is not as easy as one might think to find out how much different parts of the retina contribute to optokinesis. At different eccentricities one may determine the minimum area for evoking nystagmus under constant conditions, or measure the optokinetic gain: such measurements agree in demonstrating that the smaller the part of the visual field that moves, or the more eccentric it is, the less it contributes (figure 3.13) (Easter, 1972; Koerner and Schiller, 1972; Cheng and Outerbridge, 1975; Dubois and Collewijn, 1979b; Barnes and Hill, 1984). One might therefore assume that masking the centre of the visual field would greatly reduce optokinesis, and many experiments have shown that at fairly high velocities this is the case (van Die and Collewijn, 1982; Howard and Ohmi, 1984). However, there are complications. Are the relationships between the central and peripheral parts of the visual field the same for EOKN and LOKN? Paradoxically it seems that in the monkey central retinal lesions less than 12° in extent have no effect on EOKN (Büttner et al, 1983). A patient who could only see in the periphery because of a central scotoma (Hood 1967; 1975; Hood and Leech, 1974), showed *enhanced* optokinetic nystagmus, matching the slow phase velocity to drum speed at much higher speeds than normal subjects (figure 3.13).




Figure 3.13. Eccentricity and optokinetic sensitivity: left, the graph shows the minimum area of optokinetic stimulus required to elicit nystagmus at different eccentricities in two monkeys (data from Koerner and Schiller, 1972). Right, velocity of slow phase of OKN as a function of stimulus velocity in a subject with a unilateral central scotoma (Hood, 1975).

Such findings do not really provide an answer to the problem of how the eye copes with pursuit against a fixed background, where one might have expected some kind of antagonism between the centre and the rest. In fact if the relative contributions of centre and periphery are measured not separately but when *both* are undergoing optokinetic stimulation, the results are quite different (Hood, 1975). A subject who views a field consisting of a central stripe pattern moving in one direction with a surround of stripes moving in the opposite direction shows an *enhancement* of the nystagmus response to the centre, as compared to the case when both sets of stripes move in the same direction.

However, there are considerable difficulties in trying to interpret even relatively simple observations like these. A central patch superimposed on a background moving in one direction will appear to the subject to be moving in the other, through a mechanism of *induction*: a well-known example is the moon appearing to sail through the clouds, when in fact it is the clouds that move and the moon that is stationary (Duncker, 1938). In experiments in which the central retina is masked by a natural or artificial scotoma, the mask itself may appear to be moving in the opposite direction to the periphery, and under such conditions one may even see some optokinesis in the *opposite* direction to the stimulation (Wyatt and Pola, 1984). In any case, it is likely that the presence of a stationary object in the visual field will tend to inhibit optokinesis, even if it is stabilised on the retina (Howard and Ohmi, 1984; Wyatt and Pola, 1984). Consequently, the *edge* between the mask and the optokinetic background is an important determinant of what is seen, and hence of the eye movements (Schor and Narayan, 1981): what matters seems to be not so much the area of stimulation but rather the distance of this edge from the fovea. If the boundary is deliberately blurred, then optokinesis is found to be very much less dependent on the width of the field, and an artificial central scotoma does indeed then cause a great reduction in gain (Howard and Ohmi, 1984). In addition, if we are content to regard smooth pursuit and optokinesis as entirely different phenomena, then experiments in which optokinesis (or its lack) in the centre is pitted against optokinesis in the periphery are not entirely relevant to the consideration of *pursuit* against backgrounds: this is considered in more detail in section 3.3 below.

At all events, the fact that OKN can be elicited despite the extensive loss of retinal or optic tract function, and the fact that it is relatively unaffected by acts of the will, make it a useful neurological tool. Some of these clinical aspects of optokinetic (and vestibular) nystagmus are helpfully discussed in Reinecke's (1961) and Jung and Kornhuber's (1964) extensive reviews. An interesting point in this context is that optokinesis can be evoked by contours seen stereoscopically in a pair of random-dot patterns (Fox et al, 1978), providing the possibility of an entirely *objective* test for fusion and stereopsis.

3.2.4 The quick phase

The quick phase of OKN shares many dynamic properties with the ordinary voluntary saccade, and these aspects are left to the next chapter. But a question specifically related to OKN that may be appropriately considered here is the one that we have already tried to answer in the case of vestibular nystagmus: what causes the quick phase to occur at one time rather than another?

There have been two schools of thought about the role of the quick phase in relation to the slow. The older and more immediately plausible view is that the primary movement is the *slow* phase, the quick phase shifting the eye to a new object only when the old must be relinquished because it is about to vanish off the edge of the field. A different view, proposed originally by Borries (1926), is that it is the *quick* phase that is primary, shifting the gaze to new objects appearing in the periphery as a result of the movement of the visual field ('haptation'), whereas the slow phase merely holds on to the object until a newer and more potent peripheral stimulus again appears. The difference between these two ideas is rather slight once the nystagmus has got going, but each makes different predictions about what will happen when the drum first starts to move. The first would lead us to expect that, on setting the drum in motion, the eyes would deviate slowly to follow it, and only subsequently flick back to a new point of fixation; whereas the second would predict that the eyes would first make an anticompany quick movement, and only then follow the drum (figure 3.14; cf figure 3.1). The latter is in fact what is observed in normal subjects (Hood, 1967; Easter, 1972), and results in a mean deviation of the eye during OKN in the direction *from* which the movement is coming. In dim illumination, however, and in patients with central scotomata, or if the subject is instructed to 'look' rather than 'stare' at the drum, the situation is reversed: it is the slow phase that is primary, and the mean deviation is in the compensatory direction. Perhaps it is more helpful not to think of the slow and quick phases as primary or secondary to each other at all, but




Figure 3.14. The two possible optokinetic responses at the beginning of drum movement (shown by the shaded area at the bottom): (a) *haptation*, the quick phase comes first; (b) *following*, the slow phase precedes the quick phase.

rather as two separate phenomena induced by different aspects of the stimulus and enjoying a measure of mutual independence: the background slow phase striving to maintain the *velocity* of the image over the retina within as small a range as possible, while the quick phases behave in precisely the same way as ordinary saccades, in acquiring new objects of interest in the field. There is an obvious parallel here with what has been said earlier about vestibular nystagmus: the relation between position and velocity control is further discussed in section 3.3.1.

Another parallel between the two kinds of nystagmus is the existence of an inherent rhythmicity (Miyoshi et al, 1970; Cheng and Outerbridge, 1974a). Analysis of the intervals between successive quick phases of OKN shows a similar change from unimodal to multimodal form as the amplitude of the stimulus (in this case drum velocity) is decreased. Cheng and Outerbridge (1974a) find a distinct difference in the distributions if subjects are instructed to 'look' rather than 'stare' at the stimulus. In the second case the similarity to vestibular nystagmus is striking (with a periodicity of around 3 Hz), whereas in the first there is a natural tendency for the subject to follow a particular aspect of the stimulus until it disappears off the edge, which leads to a bimodal distribution in which the longer peak simply corresponds to the time taken for any part of the stimulus to traverse the field, and this underlines once again the close relation between 'look' nystagmus and smooth pursuit.

3.2.5 Asymmetries

So far we have only considered a situation in which both eyes are stimulated by a moving pattern over a field that is circularly symmetric with respect to the retina. By using one eye one might hope to reveal lateral asymmetries of response that are normally hidden. It turns out that there are indeed asymmetries of this kind, but that they show marked species differences related to the laterality of vision and to the relative importance of the cortical areas of the visual system. These asymmetries are seen in their purest form in rats, rabbits, and guinea pigs (Fukuda and Tokita, 1957; Precht, 1981). Monocular stimulation in these animals only produces optokinesis if it is in the *nasal* direction (for example, from right to left, for the right eye: figure 3.15). Cats and pigeons show stronger responses to nasal than to temporal stimulation (Gioanni et al, 1981; Precht, 1981), and in the pigeon there is a tendency for the stimulated eye to move more than the unstimulated.

In man such asymmetries are rather slight. Nasal optokinetic stimuli are a little more effective than temporal, as is stimulation of the right eye, but there are no left-right or nasal-temporal differences within each visual field: in the periphery, centrifugal stimulation is more effective than centripetal (Dubois and Collewijn, 1979b; van Die and Collewijn, 1982). An asymmetry in the *output* has also been described (Williams and Fender, 1979), consisting of small velocity differences between the eyes, nasal

movement being faster than temporal. By presenting subjects with separate optokinetic targets moving in opposite directions for each eye, Erkelens and Collewijn (1985) were able to show very substantial optokinesis in *opposite* directions in the two eyes, with a gain of up to 20%: these were not disparity vergence responses in the usual sense since the two targets were unrelated to each other. A curious phenomenon that might be mentioned in this context is that of *latent nystagmus* (van Vliet, 1973). The majority of normal subjects show nystagmus as the result of shutting one eye (but not both), or when covering one eye, with the slow phase in the temporal direction; it is not the unilateral loss of vision that causes it, but rather the *act* of covering the eye: secret covering has no effect. Certain clinical conditions may be associated with optokinetic asymmetry. Rod monochromats show marked asymmetry as well as a predominance of LOKN: optokinetically they behave like rabbits (Baloh et al, 1980; Yee et al, 1981; 1985). Nasal preference may also be found in amblyopes (Westall and Schor, 1985), in some cases of disturbance of binocular vision (Hine, 1985), and in very young children,




Figure 3.15. Optokinetic asymmetry. Above, top view of an animal with asymmetric OKN; with one eye covered, nasal optokinetic stimulation is much more effective than temporal, which may be absent altogether. Below, a general scheme that can account for such asymmetries in (left) an idealised asymmetrical animal and (right) a symmetrical one. In each case, the solid lines show the flow of excitation on monocular stimulation.

where in addition the odd phenomenon of nasal OKAN after nasal or temporal stimulation may be observed (Naegele and Held, 1982; Schor et al, 1983).

Finally, an asymmetry of a different kind may be observed in albino rats and rabbits (Collewijn et al, 1978; Batini et al, 1979; Precht and Cazin, 1979; Sirkin et al, 1985a). What is found is *inversion* of the direction of nystagmus on stimulating the anterior field, perhaps because of aberrant projections of the optic nerve: as a consequence, whole-field stimulation is rather ineffective, but the animal may follow single slowly-moving stripes or projected patterns of dots if they are sufficiently bright.

The simplest explanation for the optokinetic asymmetries is that there are two levels of integration of the response: a lower one consisting of two unidirectional subunits, each of which is driven monocularly, and a higher level that is driven by both eyes and has access to both of the lower subunits (figure 3.15). The anatomical basis for such an arrangement is discussed in chapter 9.

3.3 Smooth pursuit

When the eyes track a small moving target, they do it partly by making saccades that bring the fovea on to the target, and partly by slower, continuous movements whose velocity is closely related to that of the target. This second component is what is meant by *smooth pursuit*, and can be studied on its own by subtracting the saccades from the full tracking movement (figure 3.16). As we have seen, it is a mechanism that is distinct from optokinesis (though it may well share some of its neural mechanisms, particularly in the case of EOKN), for many animals that show optokinesis with large moving fields do not respond to single moving targets. Furthermore, a target may be pursued against a stationary detailed background, when optokinesis ought to hold the eye still. Under such circumstances the




Figure 3.16. The separation of smooth and saccadic components of foveal tracking. *a*, target position; *e*, eye position; *e_{sm}*, cumulative eye position obtained by removal of saccades; *e_{sa}*, the isolated saccadic component (after Collewijn and Tamminga, 1984).

velocity of smooth pursuit is only slightly reduced, the increased position error being corrected by making more saccades; conversely, if the background moves with the target, gain increases: but the interaction is not just one of simple summation (Evinger and Fuchs, 1978; Kowler et al, 1978; Yee et al, 1983b; Collewijn and Tamminga, 1984; 1986). Closer investigation, looking at the first 100 ms of pursuit of a target that suddenly starts to move (which implies an effectively open-loop response, since there is not time for visual feedback) shows that there are in fact two components of the response (Lisberger and Westbrook, 1985). In the first 20 ms or so, the acceleration of the eye is constant, and unrelated to any aspect of the stimulus. Thereafter, it is strongly dependent on the position and velocity of the target, and the background. This suggests that there are two components to smooth pursuit: an early one that is either invariable or at least easily saturated, and a later one influenced more obviously by the nature of the stimulus. In the monkey, Keller and Khan (1986) have recently shown that the presence of a textured background reduces the acceleration of the eye at first, but that after several hundred milliseconds have elapsed it has little effect.

A target does not have to lie on the fovea for smooth pursuit to occur, and parafoveal tracking of this kind can apparently be performed without preliminary practice (Winterson and Steinman, 1978; Yee et al, 1983a; Pola and Wyatt, 1985; Collewijn and Tamminga, 1986). Pursuit can occur in any meridian, but is on the whole smoother and more precise in the horizontal direction (Collewijn and Tamminga, 1984). The maximal velocities achieved depend on the species, and the kind of target used. In man and monkeys, estimates range from some 80 to 160° s^{-1} (Lisberger et al, 1981b; Buizza et al, 1984; Meyer et al, 1985); in the cat, peak velocities are much smaller, perhaps some 10° s^{-1} at most (Evinger and Fuchs, 1978).

3.3.1 What is being pursued?

On the face of it, smooth pursuit *sounds* quite simple. The image of a target moves across the retina, generating a slip velocity that makes the eyes move in such a way as to track it. However, there are many kinds of observation that suggest an entirely different process. What is envisaged is that the brain forms an internal estimate or *prediction* of what the movement of the target is in space, retinal slip being only one of a number of sources of information that may be used as evidence about target movement; this estimate of target velocity is then used to drive the eye. Thus pursuit can be evoked by smoothly moving auditory or even tactile stimulation (Lackner and Evanoff, 1977; Gottschalk et al, 1978; Zambardi et al, 1981). One can track the movement of one's own arm equally well in the dark; and if a vibrator is applied to the biceps, it is the *illusory* position that is tracked, unless there is also a direct visual indication of limb position as well (Jordan, 1970; Levine and Lackner, 1979; Lackner and Mather, 1981). In the same way, a target can be tracked better when moved by the subject than when moved by

someone else (Steinbach, 1969; Collewijn et al, 1985b; but see also Mather and Lackner, 1981), using fewer saccades and with higher velocity and lower latency. Furthermore, in the presence of a stabilised retinal image such as an afterimage (Heywood and Churcher, 1971), a subject in the dark can track an entirely *imaginary* moving target [though not very realistically (Cushman et al, 1984)]! The idea that it is target velocity in *space* rather than relative to the head that is being tracked is also supported by an experiment of Lanman et al (1978). If a monkey is trained to track a target by using its head as well as its eyes, and the head is suddenly stopped, the eye continues to track with a latency of only 15 ms, far too short to be the result of a perceived change in retinal slip.

Even with purely visual stimulation it is found that smooth pursuit is much more closely related to the *perceived* target velocity than its actual velocity, when these two are at odds. Thus a moving frame makes a stationary target inside it appear to move in the opposite direction, and this is found to increase the velocity of tracking when the target does actually move in counterphase (Wyatt and Pola, 1979; Pola and Wyatt, 1980; Mack et al, 1982). Another example comes from the use of *anorthoscopic* stimuli (figure 3.17). If a simple shape such as a circle is moved back and forth behind a narrow vertical slit, all one can see at any one moment is a single cross section of the stimulus, the edges of which are moving up and down. Nevertheless, the stimulus is perceived as a horizontally moving object of approximately the right shape, and can be tracked [though with reduced amplitude (Steinbach, 1976; Mack et al, 1982)].

Finally, perhaps the most familiar of these kinds of illusion is that of continuous motion of an intermittently moving target (Wertheimer's *β -movement*, 1912): without it, cinema and television could not happen. So long as the discrete jumps are less than some 150 ms apart, such a target is tracked smoothly (Morgan and Turnbull, 1978; van der Steen et al, 1983); the perception of smoothness actually breaks down a little sooner than




Figure 3.17. The anorthoscope. A circle (or some other shape) moves backwards and forwards behind a narrow vertical slit so that all that can be seen is vertical motion of its edges. Yet its shape and motion are perceived, and it can be tracked—approximately—with the eyes.

smooth pursuit as this interval is increased. A phenomenon related to β -movement is σ -movement (Stoper, 1973; see also ter Braak, 1972): here a regularly structured pattern such as a row of lights or a grating pattern, with a spatial period p , is stroboscopically illuminated at a frequency f . The subject then tends to perceive a movement of the pattern of velocity pf , and his eyes track this velocity smoothly (Adler and Grüsser, 1979; 1981; van der Steen et al, 1983). The choice of the direction of the apparent motion is under voluntary control, as is the path taken when two such rows of dots intersect, whereas the perceived speed is most closely related to gaze velocity (Collewijn et al, 1981; 1982b; Curio and Grüsser, 1985). This illusion accords very nicely with the idea of a self-maintained centrally generated estimate of target velocity which is topped up from time to time by visual information, since the stroboscopic illumination appears to confirm the correctness of the perceived velocity and eye movement, *whichever* direction the subject happens to choose! In the same way, an afterimage provides a similarly self-validating stimulus that may account for the increased gain in the VOR in the dark that is then observed (Yasui and Young, 1975; 1984).

3.3.2 Responses to repetitive stimuli

In order to avoid both the Scylla of transient input saturation, and the Charybdis of prolonged output saturation, so evident in optokinesis to unidirectional targets, many experimenters have used sinusoidally moving targets to investigate the pursuit system. One may show, for instance, that when the relation between target velocity and eye velocity is measured by using ramps, the system saturates in a very nonlinear manner, at around 40° s^{-1} ; during sinusoidal stimulation, however, the relation between peak eye and target velocity is linear up to some 80° s^{-1} (Buizza et al, 1984). As in the case of the vestibular system, sinusoidal inputs are likely both to be more 'physiological' from this point of view, and also to provide more detailed information about the dynamic properties of the system.

Unfortunately, the use of repetitive and therefore predictable stimuli brings problems of its own. It turns out that the pursuit system in man [but not in monkeys (Fuchs, 1967b)] is capable of learning to improve its response to such stimuli with practice, so that, although a subject's performance at the task is at first poor, with many corrective saccades, in quite a short time, eye movements smoothly match the movements of the target. One might say that the subject is learning to *predict* where the target is going to be by the time the movements he or she initiates take their effect. St Cyr and Fender (1969c) suggest that this is an unjustifiable description: it is not so much that the system predicts the position of the target, but rather that it can learn to reduce the delay with which it responds. Whichever form of words one prefers, the fact is that the phase lag between target and eye is greatly reduced—particularly at higher frequencies—as a result of practice.

It is the unpractised response that is the more interesting as far as our understanding of the mechanism of pursuit is concerned, at least if we want a nice simple model unencumbered with parametric adaptation.

Does this rule out the possibility of making measurements of the frequency response of the system? If we take a number of sine waves that differ only slightly in frequency and add them together, the resultant waveform will in general only repeat itself at a frequency very much lower than that of the sine components of which it is made, and, if the components are properly chosen, such a waveform is to all intents unpredictable. But since the component frequencies are close together on the frequency axis, measurement of the response of the eye to a number of such clusters, centred on different frequency bands, will enable the experimenter to plot a frequency response for the system, though with rather less resolution than is possible with single sinusoids. This does not result in much loss of information in practice, since the frequency responses show little by way of fine structure. Another way of achieving the same result is to start with a source of random Gaussian 'noise', and pass it through broadly tuned filters centred on various frequencies: again, a signal is achieved that, though unpredictable, is concentrated in a narrow frequency range. The predictability of these signals can be varied by altering the width of the filter, or in the case of multiple sine waves, by using a larger or smaller number of individual components (Stark et al, 1962; Michael and Melvill Jones, 1966; St Cyr and Fender, 1969b; 1969c). A third approach is to generate a 'random walk of sinusoids' (Lisberger et al, 1981b) in which at the end of each cycle of a sine wave a random choice is made as to whether or not to change direction (figure 3.18). Finally, almost any kind of random stimulus will do, if one is prepared to carry out the calculations necessary to derive a transfer function from the response: one that is simple to generate consists of ramps of random velocity and duration (Bahill et al, 1980).




Figure 3.18. A 'random walk of sinusoids' as used by Lisberger et al (1981b). At each point of zero velocity there is a given probability of the motion changing direction.

An important point to bear in mind in this section is that until recently (Bahill et al 1980; Collewijn and Tamminga, 1984; Yasui and Young, 1984) experimenters in this area have tended to consider the *entire* tracking response in relation to the stimulus, rather than separating out that part of tracking that is performed by smooth pursuit, and the part that is done by saccades. Figure 3.19 shows transfer functions for pursuit of a bandwidth-

limited pseudorandomly moving target, with 'dual-mode' tracking (smooth pursuit plus saccades) and 'single-mode' (with the saccades removed): the contribution of the saccades at high frequencies is obvious (Bahill et al, 1980).

With all of these methods of generating band-limited random stimuli it is found that the phase lag of the eye behind the target is directly related to the predictability of the stimulus in the region around 0.3–2 Hz: in the very limited range above 2 Hz, the response appears to be constrained by other factors (which must include, for example, the mechanical properties of the eye), whereas at lower frequencies the performance is equally good whether or not the stimulus is predictable (figure 3.20). St Cyr and Fender (1969c) find that the relationship between phase lag and predictability is roughly compatible with the idea that there is a delay in the system that increases in proportion to the quantity of information carried by the signal: similar relations between reaction time and information transfer have long been known to experimental psychologists (see for example Edwards, 1969).

Investigation of the dynamics of two-dimensional tracking has led to interesting findings (Goodwin and Fender, 1973a; 1973b). If a target is moved sinusoidally in the horizontal direction, and randomly in the vertical direction, tracking is found to be essentially independent in the two channels (in the sense that very little extra power at the sinusoidal frequency is observed to leak into the vertical eye movement component): similarly, little interaction is observed between saccades and slow pursuit in perpendicular meridians (Feinstein and Williams, 1972). Furthermore, it is found that the usual reduction in phase lag for the sinusoidal stimulus occurs (relative to that for the random stimulus) even if the two axes are tilted away from the horizontal and vertical. Thus it is clear that this




Figure 3.19. Dual versus single mode tracking of a pseudorandom stimulus in man. Gain of pursuit measured as a function of frequency under dual-mode (triangles) and single-mode (circles) conditions—that is, with and without saccades (data from Bahill et al, 1980).

independence of processing of the two components is not the result of a fixed division into horizontal and vertical vectors, but rather the result of the ability of the pursuit system to seek out for itself, and adapt itself to, a predictable *component* of the two-dimensional motion, whatever its orientation in space.

Perhaps the most dramatic demonstration of the ability of the system to tailor its performance to the type of stimulus it is given is the fact that, if one looks at the response to each of the separate components within a small cluster of sinusoids, the gain of the system is found actually to *increase* with frequency within the group; if one group is compared with another, the relation is the expected one of falling gain with rising frequency (figure 3.21) (St Cyr and Fender, 1969b). A more recent report (Barnes and Donnelly, 1986) suggests an even more complex state of affairs when small sets of sinusoids are used. Increasing the frequency of only the highest component appears to reduce the gain for the lower-frequency components, that for the highest component itself remaining what it would have been if the others were not there. In other words, it seems that in such circumstances the mechanism that is doing the predicting deals only with the highest frequency present, and ignores the others. Perhaps it is only able to envisage a relatively small number of types of pattern of movement.

Adaptive behaviour of this kind is no doubt advantageous to the system, but it does make life difficult for the experimenter, who would prefer to have to deal with a system with fixed properties. Although the results of using sinusoidal inputs lead to definite conclusions that can be expressed




Figure 3.20. Closed-loop phase response for smooth pursuit of targets moving sinusoidally (solid symbols) and unpredictably (open symbols) by human subjects. The sources of the data are: ●○ Yasui and Young (1984); ■□ Dallos and Jones (1963); ▲ Fender and Nye (1961); ◇ Michael and Melvill Jones (1966).

quantitatively as elaborate models of the control system, we cannot assume that these models have any predictive power if stimuli other than simple sinusoids are presented. The likelihood is that there is a repertoire of potential 'models' embodied in the oculomotor system, any one of which can be selected by the system to provide the best possible response to a particular input (Bahill and McDonald, 1983). Such an idea is a logical extension of the notion of a centrally generated estimate of target velocity: such an estimate may be based not only on current information, but on past experience, and will continue to assume things are the same until it receives new information that contradicts it. A good demonstration of this is what




Figure 3.21. Closed-loop frequency response of smooth pursuit of a target moving with four sinusoidal components of similar frequency: it can be seen that for each of the two subjects, gain tends to rise with frequency when comparing the responses to each individual component *within* the frequency band, but that the overall response *comparing one band with another* is similar to that of figure 3.20 (data from St Cyr and Fender, 1969b).




Figure 3.22. Perseveration of smooth pursuit: (a) perseveration in tracking a 'random walk of sinusoids' (Lisberger et al, 1981b); (b) eye movement *e* in response to a target *a* that moves at constant velocity for the time indicated and then vanishes; the eye continues for a period inversely related to the duration of the target motion (after Mitrani and Dimitrov, 1978).

happens when a target that has previously been moving in a predictable manner is extinguished, for instance a pendulum swinging behind a narrow screen. After the moment of disappearance the eye continues at the same velocity for some 200 ms, then decelerates rapidly to about 60% of its previous velocity (Becker and Fuchs, 1985; see also Mitrani and Dimitrov, 1978; Muratore and Zee, 1979). In the same way, when in a 'random walk of sinusoids' the target direction unexpectedly changes, for a short period the eye continues along the trajectory it would have followed if the motion had continued (figure 3.22) (Lisberger et al, 1981b). A related observation is that the contrast of a target influences smooth pursuit when it moves unpredictably, but not when it moves predictably (Haegerstrom-Portnoy and Brown, 1979). In the latter case, the visual input is rather uncritical, since it is merely required as a confirmation that all is well: the complacency of such a self-satisfying system perhaps requires pretty definite proof that it is wrong.

3.3.3 Open-loop dynamics

The general desirability of measuring open-loop rather than closed-loop responses has already been emphasised. But open-loop studies introduce certain complications when the system we are dealing with is one that is trying to generate a response that matches the stimulus, rather than one that is driven directly by an error signal. Such a system is likely to use error signals from the retina mostly as *parametric* feedback, continually monitoring the match between prediction and stimulus and using the information either to improve the model, or perhaps to select a different one. In any case, if it is indeed parametric feedback that governs the programmes of behaviour, then under open-loop conditions, when the system will see errors that go on occurring *whatever* it does to improve its performance, it would be understandable if it went into a sulk and refused to behave very sensibly at all. At all events, one need not expect any very clear relationship between open-loop and closed-loop responses, most particularly when dealing with very predictable stimuli.

Rather few experimenters have made open-loop investigations of the tracking response to sinusoidal target motion (Fender and Nye, 1961; Wyatt and Pola, 1983; Pola and Wyatt, 1985) (figure 3.23). Under open-loop conditions nonlinearities are necessarily particularly prominent, which may contribute to the large phase lags seen at lower frequencies. Nevertheless, despite the pessimism voiced in the preceding paragraph, Pola and Wyatt (1985) conclude that closed-loop responses *can* be adequately predicted from open-loop measurements, and (Wyatt and Pola, 1983) that prediction occurs in much the same way under both conditions. This is not at all easy to reconcile with the view that prediction is the result of selecting some kind of stereotyped response from a repertoire, and in fact tends to lead back to the idea that smooth pursuit is after all primarily a function of retinal slip velocity.

A different kind of observation that can also be taken as open-loop, mentioned earlier, is that of eye velocity within 100 ms or so of the start of pursuit to a new stimulus or change of stimulus (Lisberger et al, 1981b;




Figure 3.23. Open-loop transfer functions for smooth pursuit in human subjects (■, Fender and Nye, 1961; ▲, ○, active and passive in one subject; ▼, active in another, Pola and Wyatt, 1985).




Figure 3.24. Pursuit afternystagmus. On the left of the record, a small target moving in sawtooth fashion was tracked in the dark; at the arrow it was extinguished but nystagmus continued for a while (modified from Muratore and Zee, 1979).

Behrens et al, 1985; Lisberger and Westbrook, 1985; Keller and Khan, 1986; Tychsen and Lisberger, 1986). The basic finding is that eye acceleration is consistently related to retinal error velocity, which implies a process of integration borne out—as far as they go—by the general form of the open-loop frequency responses, and similar to what has already been noted in the case of OKN. The nature of these observations of course precludes the measurement of long-term (low-frequency) behaviour, and thus of determining whether here too we should think in terms of a *leaky* integrator; the existence of pursuit afternystagmus (figure 3.24) (Muratore and Zee, 1979), having a similar time constant to OKAN, suggests that we should. It is perhaps not necessary to point out that integration in this context amounts to much the same thing as the kind of perseveration introduced by self-sustaining positive feedback, and thus lends a little support to the idea of a centrally sustained internal model of target velocity. Details of how such a system might operate are considered in chapter 12.

3.3.4 Artificial feedback

In most open-loop experiments, we create a second, artificial, loop whose gain is adjusted to be equal and opposite to the -1 provided by the natural relation between eye movement and image movement. The same apparatus can be modified to create any overall feedback relation that we wish. Now all real feedback systems are at best only conditionally stable (see appendix 2); if the loop gain is steadily increased, there must eventually come a point when the gain associated with the frequency for which the phase lag round the loop is 180° exceeds unity, and this would result in spontaneous oscillation at that frequency. This is in fact precisely what is observed, although one must be careful to separate saccadic corrections from those due to smooth pursuit, for the nature of their potential instability is different. Robinson (1965) arranged to filter any saccadic movements out of the electrical signal representing the eye movements: this signal was then fed back to move the target, after a stage of amplification of variable gain. With an initially stationary target, the system is found to be stable until the feedback gain is increased from -1 (that is, natural conditions) to -8 , when spontaneous oscillations begin at a frequency near 3.3 Hz. However, it is still possible to trigger transient oscillations in response to the onset of target motion (figure 3.25), even when the feedback gain is only -2 , a reminder again that we are not dealing with a simple linear system.

A striking feature of the oscillations with these intermediate values of feedback is that their amplitudes wax and wane during the stimulus, finally dying out not in the exponential manner that might be expected, but considerably more abruptly. A possible reason for this phenomenon comes from a consideration of the differences in phase lags already noted as between unpredictable (Gaussian) and predictable (sinusoidal) stimuli. Presumably, after a few cycles of the oscillation the system parameters adjust themselves to correspond with the predictability of the stimulus: the

phase lag is then reduced, and consequently the stability of the system increased (Robinson, 1965). These results suggest that the gain of the system is just about as large as it can be without the danger of instability: the difference between a feedback gain of -1 and one of -2 is not very great. One might expect to find transient oscillations in response to sudden accelerations of the visual field even under natural conditions, although the saturation of the input in response to high angular velocities would tend to reduce such an effect. It may be that the intermittent discontinuities of velocity that have sometimes been described in the oculomotor response to stimuli of changing velocity (for example Westheimer, 1954b; Puckett and Steinman, 1969) really represent an oscillational instability superimposed on the basic smooth response. Certainly the interval between these periods of apparent change of velocity (some 200 ms) is of the right order of magnitude to be consistent with such oscillations.

Another kind of instability can be induced if the feedback is arranged to be positive instead of negative, so that any corrective movement made by the eye makes the error larger instead of smaller. As might be expected, only a small amount of feedback of this kind (for example a feedback gain of around $+0.3$) is needed to induce this type of instability (Fender and Nye, 1961); it does not tell one much about the system.




Figure 3.25. Human pursuit of a target that suddenly moves off with a velocity of 10° s^{-1} , under artificial feedback conditions: the feedback gain was arranged to be -4 for smooth pursuit, but -1 for saccades. The thick trace is the open-loop stimulus, the thin trace is the response of the eye. After a pair of initial saccades, irregular oscillations of the eye record are evident, with a frequency of some 2.9 Hz and peak velocity deviation of $1.25^{\circ}\text{ s}^{-1}$ (after Robinson, 1965).

3.3.5 Smooth control of position

In a simple linear system we would expect to find the closest possible relation between its responses to velocity and its responses to position, since one is simply the differential of the other. If we find that the immediate response of the system to a step of retinal slip velocity is a ramp of steadily increasing eye velocity, then we would certainly expect its response to a step of *position* to be a steady increase in deviation, of constant velocity. But in making this prediction we are assuming that both position and velocity are processed by the same mechanism, and that each is simply a different aspect of one and the same signal. Is this actually justified?

The truth is that the visual system processes these two quantities in entirely different ways. Position in the visual system is signalled qualitatively, by *which* of the whole array of visual receptors is active at any moment. Velocity is treated quite differently, being extracted at a very early stage by local 'microprocessing' of firing patterns in small subgroups of receptors, and thereafter coded *quantitatively*: knowledge of how fast a target is moving comes not from which fibres are firing but from *how much* they fire. These two kinds of information make very different computational demands on the oculomotor system. The final output is of course of the 'how much' kind—at least, as far as any one eye muscle is concerned. Visual velocity information is already in this form, and it is not hard to devise neuronal circuits that would convert one into the other. On the other hand, positional 'which' information must undergo radical recoding before it can be used to control muscles, and the relative complexity of this task no doubt underlies the slowness of positional oculomotor responses compared with those driven by velocity. To some extent, in fact, the control of velocity and position are somewhat incompatible, as Rashbass's striking analogy makes clear:

"Imagine two drivers of the same car, one trying to keep the speedometer at 30 m.p.h. and one trying to keep the car alongside another also travelling at 30 m.p.h. Unless there is communication between the two drivers it is easy to see that the velocity man is going to frustrate the efforts of the position man. Changes in velocity due to the saccades must not be corrected by smooth tracking. Saccadic suppression of vision—the position man putting his hand over the speedometer every time he does something—could account for this ..." (Rashbass, 1971).

Equally, the fact that optokinetic responses break down completely at the kind of velocity levels that are found during saccades would serve the same purpose, without having to suppose a special mechanism of interaction (saccadic suppression is discussed in chapter 11). In practice, the systems are indeed nearly independent, in the sense that positional errors are essentially generators of saccades, while velocity errors produce smooth movements, but not entirely. Saccades can be initiated by pure velocity errors (section 4.2.3), and conversely, pure errors of position can generate smooth movements.

This can be demonstrated most easily by arranging to displace a stabilised image through a small angle (for example, Robinson, 1965; Steinbach and Pearce, 1972), or even by viewing a stationary stabilised image—an afterimage will do very well—and concentrating the attention on one or other side of it (Kommerell and Taumer, 1972; the importance of attention suggests a rather high level of development of the response). Now we saw earlier that, if the system were linear and essentially unitary, we would expect the response to a step of position to be a movement of constant velocity. In fact this is not at all what is observed: step displacements give rise to movements whose velocities increase steadily (figure 3.26) (Robinson, 1965) in a way that depends on the size of the step (Wyatt and Pola, 1981; but see also Steinbach and Pearce, 1972).

A rather clear-cut demonstration of the relative strength of errors of retinal velocity and position is given by the open-loop response to a target that steps to the left at the same time as starting to move at constant velocity to the right (Pola and Wyatt, 1980): the initial target movement is to the left, even though this actually *increases* retinal slip velocity. Smooth responses to target position are considered by Pola and Wyatt (1985) to be a property of 'active' foveal pursuit, and not of the 'passive' optokinetic kind. A phenomenon that may well be related to smooth position control is the involuntary drift towards a target that is about to be fixated, described by Kowler et al (1984a); there is probably also a connection with *slow correction* ("corrective drift": section 6.3.2). Torsion in response to a tilted visual stimulus, once denied but undoubtedly real though rather small in magnitude (1° at most: Howard and Templeton, 1964; Crone, 1975; Goodenough et al, 1979; Merker and Held, 1981), is superficially similar but rather different in its dynamic properties.




Figure 3.26. Human open-loop responses to a small step displacement of the target. The thick line represents target position, the thin line eye position: the accelerating nature of the response is evident (after Robinson, 1965).

Saccades

"To be short, they be wholly given to follow the motions of the minde, they doe change themselves in a moment, they doe alter and conforme themselves unto it in such maner, as that *Blemor* the Arabian, and *Syreneus* the Phisition of Cypres, thought it no absurditie to affirme that the soule dwelt in the eyes ..."

In the strictest sense, saccades are the fast movements of the eyes that are used to bring a new part of the visual field to the foveal region. They are essentially voluntary, and indeed are the only voluntary eye movements that one can make without special training. However, some other fast eye movements, less voluntary in character, share many properties with voluntary saccades and are almost certainly generated by the same mechanism. These include the quick phase of vestibular or optokinetic nystagmus and the microsaccades which can be observed during fixation, to be described in chapter 6. The voluntary torsional movements of up to 30° that subjects can be trained to execute are also partly saccadic in nature (Balliet and Nakayama, 1978). The word 'saccade' appears first to have been used in the oculomotor sense by Javal (1879): Westheimer (1973) has briefly reviewed the historical development of ideas about this kind of eye movement.

4.1 The time course of saccades

Saccades are remarkably stereotyped: for a particular subject, the time course of a saccade of a given amplitude is largely independent of the means by which it is evoked, whether voluntarily to an existing visual target, or involuntarily in response to the sudden appearance of a target; its *latency* on the other hand is subject to considerable random variation. Saccades are so fast that there is normally no time for visual feedback to guide the eye to its final position: except in the case of the very largest saccades, the delay in the visual feedback loop is longer than the duration of the movement itself. The saccade control system must therefore calculate in advance a pattern of muscle activation that will throw the eye exactly to the desired position. The resultant movement is thus preprogrammed or *ballistic* (the same property that distinguishes ballistic missiles from guided missiles), in the sense that alterations in the target occurring during a saccade cannot modify its trajectory. (This should not be taken to imply that a saccade, once under way, cannot be modified. Although the time course of a saccade is normally quite rigidly determined by the stimulus that evoked it, under special conditions a subsequent visual stimulus, provided that it occurs soon enough, may modify the movement in midflight: section 4.1.3 below.) The way in which the time course of this movement 'package' varies for saccades of different sizes can tell us something about how the control system performs what is, on the face of it, a complex calculation in which distances across the retina have to be converted into temporal patterns of muscle activity.

4.1.1. Amplitude-velocity-duration relationships

Figure 4.1 shows the time courses of a number of human saccades of different amplitudes in the horizontal plane: saccades in other meridians, including torsional saccades, do not differ in their essential characteristics (Gurevich, 1961; Balliet and Nakayama, 1978), nor is much interspecies variation observed [thus very similar results are found in the goldfish: see Hermann and Constantine (1971) and Easter (1975)]. A notable feature of such records is how fast the eye moves, often reaching more than 700° s^{-1} for large amplitudes. Not all recording techniques are suitable for measuring such high velocities: slippage of contact lenses may lead to underestimates (Byford, 1962), as does electro-oculography (Byford, 1963; Stryker and Blakemore, 1972; Boghen et al, 1974). It can also be seen that the duration of the complete movement is not constant, but increases with increasing amplitude (figure 4.2); the duration of saccades larger than some 5° in amplitude is roughly 20–30 ms plus about 2 ms for every degree of amplitude (Dodge and Cline, 1901; Hyde, 1959; Robinson, 1964).

This dependence of duration on amplitude has sometimes—misleadingly—been taken to imply that the system is in a sense nonlinear. But we cannot tell whether a system is linear or not simply by looking at a sample of its output: we must also know what the corresponding *input* was that produced it. The implicit assumption here is that the input that produces saccades of different sizes is a step function of varying height: if this were true, then differences in response duration would indeed imply some kind of nonlinearity. But records of the electrical activity of the motor nerves show that the fast rising part of the saccade is generated not by a step of activity of variable height, but by a *pulse* of essentially constant height whose *duration*




Figure 4.1. Human saccades of different sizes. The traces have been superimposed so that the beginning of each movement is at time zero. The dependence of the duration of the saccade on its amplitude can be seen, as can second correction saccades for larger amplitudes, at around $t = 200$ ms (Robinson, 1964).

determines the amplitude of the saccade (see section 7.5.5). Such a pulse, acting on a linear model of the mechanical properties of the eye, produces amplitude-duration relationships very similar to those observed in actual voluntary saccades (Robinson, 1964). For a small saccade the pulse is very short and the response is dominated by the mechanical properties of the eye, so that the movement has a nearly constant duration; under these circumstances the peak velocity varies in proportion to the amplitude. The situation is rather like that of a man falling off a cliff: at first, acceleration dominates, and his peak velocity depends on how far he falls. But if his drop is a long one, most of the way he will be falling at his terminal velocity, so the duration of his fall will be in proportion to the height of the cliff. Either way, his initial trajectory will be the same. In the case of the eye, one can observe almost identical patterns of acceleration at the *beginning* of the movement, whatever its amplitude (Hyde, 1959). Larger saccades show a levelling-off of peak velocity as a function of amplitude, though it never becomes completely flat (figure 4.3).

This relatively fixed relation between amplitude, duration, and peak velocity for different saccades leads naturally to the idea of a sort of saccadic norm by which one can judge whether a particular saccade is aberrant, whether indeed it *really* is a saccade. [By analogy with the astronomers' classification of stars that is based on a similarly normative relation between brightness and temperature, this set of saccadic data has been called the *main sequence* (Bahill et al, 1975c; 1981).] Thus very similar relationships between amplitude, duration, and velocity have been reported for microsaccades and for the quick phase of nystagmus (Mackensen and Schumacher, 1960; Ron et al, 1972; Guitton and Mandl, 1980a), which lends weight to the notion that these movements are essentially no different




Figure 4.2. Saccade duration as a function of amplitude: \blacktriangle nasal saccades and \blacktriangledown temporal saccades (Robinson, 1964); \circ Yarus (1956); \bullet calculated from Hyde's (1959) records. The line represents the function $(2.2A + 21)$ ms, where A is the saccade amplitude in degrees.

from ordinary saccades. Under natural conditions, with the subject moving freely in his normal surroundings, a histogram of the frequency of occurrence of fast eye movements (taking saccades and quick phases together) as a function of their amplitudes appears to show a single population, following a simple exponential function with a characteristic amplitude of around 7° (Bahill et al, 1975a) (figure 4.4). Thus more than 85% of natural saccades have amplitudes of less than 15° [as also noted much earlier by Lancaster (1941); a similar relationship is found in the rabbit and in the cat (Collewijn, 1970b; 1977a; 1977b)], and again, saccades and quick phases seem to fall into a single statistical population.




Figure 4.3. The main sequence, a graph of the peak velocity of human saccades as a function of their duration (Bahill et al, 1975c).




Figure 4.4. Distributions of saccadic amplitude under natural conditions. The graph shows the proportion of saccades that are less than or equal to the amplitudes shown on the abscissa, for: ● average of three human subjects (Bahill et al, 1975a); ○ cat (Collewijn, 1977a; 1977b).

Saccades that are slower than the main sequence would predict are found in infants (Hainline, 1984; Hainline et al, 1984), with auditory rather than visual targets, or in the dark (Becker and Fuchs, 1969; Riggs et al, 1974; Koerner, 1975; Zambardi et al, 1981; 1982), as a result of alcohol or barbiturates (Aschoff, 1968; Franck and Kuhlo, 1970; Gentles and Llewellyn Thomas, 1971; Wilkinson et al, 1974; Bittencourt et al, 1981), or as a result of voluntary control (Crawford, 1984), in certain clinical conditions (Starkman et al, 1972; see Leigh and Zee, 1983), and perhaps with fatigue (Bahill and Stark, 1975b; Fuchs and Binder, 1983). Age appears to have little effect (Abel et al, 1983), nor blindness (Leigh and Zee, 1980). Differences are also found in more extreme lateral eye positions, which may be explained by the mechanical properties of the eye (Abel et al, 1979). In cats, peak velocities are smaller, and the saccadic parameters tend to be more variable (Evinger and Fuchs, 1978). One might wonder whether the normal duration-velocity relation would still be seen when a saccade is made on a continuing background of smooth pursuit, or whether the smooth pursuit velocity would not add on to the saccadic velocity. Summation of this kind is not in fact seen (Jürgens and Becker, 1975) suggesting that here, as noted in section 2.3.2 in the case of the vestibulo-ocular reflex, other inputs to the oculomotor system are in effect switched off during saccades. Another situation where some deviation from the main sequence might perhaps have been expected is after damage to the oculomotor nerves; it turns out, however, that subjects learn to overcome the resultant weakness by generating longer pulse durations for a given amplitude, the velocity-duration relationship remaining normal (Abel et al, 1978).

Differences in the velocity of saccades made in different directions have sometimes been noted, though not always consistently (for a summary see Fuchs, 1971). In the case of oblique saccades—those involving more than one pair of muscles—one might wonder whether each pair of muscles receives a command of equal duration, but of a size scaled according to its contribution, or whether the horizontal and vertical components of the movement are of different durations. Experiments show that it is the first of these possibilities that actually happens, and that the smaller component is




Figure 4.5. Horizontal (H) and vertical (V) components of an oblique saccade made by a cat: T is target displacement. The durations of each component (marked by arrows) are approximately equal (Evinger et al, 1981a).

stretched and slowed in such a way that both components finish more or less simultaneously (Viviani and Berthoz, 1977; Blakemore and Donaghy, 1980; Guitton and Mandl, 1980b; Evinger et al, 1981a; van Gisbergen et al, 1985) (figure 4.5). Possible ways in which the oculomotor system may do this are considered in chapter 12.

Comparisons with other kinds of movement can be instructive. Quite similar relationships between velocity, amplitude, and duration are found for voluntary movements of the the wrist (Fortuyn, 1982), suggesting that here too the size of a movement is essentially encoded in the duration of a motor command. Head movements, however, appear to be programmed differently: velocity increases more markedly with amplitude, and duration is virtually constant (Stark et al, 1980; Dieringer et al, 1982).

4.1.2 The termination of saccades

The pulse of activity that drives the eye to its new position must be followed by a steady (and lower) level of stimulation that will hold it there. If these two components, the *pulse* and the *step*, are not matched correctly to one another, at the end of the saccade the eye will move from the position to which the pulse has thrown it, to the position corresponding to the size of the step component. Movements of this kind are rather slow, taking as much as a second to reach completion, and are called *glissades*: their time course corresponds with the step response of the eye mechanics (Weber and Daroff, 1972; Easter, 1973; Bahill et al, 1975d; 1976). If the pulse is too long in relation to the size of the step, an overshoot with glissadic return will be seen; if too small, glissadic undershoot (figure 4.6). In addition to mismatch between the step and pulse components of activation of any one muscle, there may be differences of timing or amplitude between its activity and that of its antagonist. In particular, there is some evidence that saccades




Figure 4.6. Modes of saccade termination. Left, glissadic undershoot and overshoot compared with a normal saccade, together with the corresponding idealised neural commands. Right, binocular recording of human 5° saccades showing dynamic overshoot in the abducting eye (*abd*) and postsaccadic drift in the adducting eye (*add*) (after Kapoula et al, 1986).

CONFUSING

may be braked by a pulse of activity in the antagonist at the end of the movement (chapter 7), and this may also provide a source of error. Mismatch of this kind is believed to account for a different kind of behaviour seen in a substantial proportion of saccades called *dynamic overshoot*. Here the eye returns to its final position within 15–20 ms at a velocity that is much higher than in the case of a glissade (figure 4.6) (Westheimer, 1954a; Thomas, 1961; Fuchs, 1967a; Bahill et al, 1975b; 1976; Lehman and Stark, 1983; Kapoula et al, 1986). It is more commonly seen in the abducting eye, and is often associated with dynamic undershoot in the adducting one; it is characteristic of certain diseases of the periphery such as myasthenia gravis (Feldon et al, 1982). A third category of behaviour after a saccade is *postsaccadic drift*, more common in the adducting eye, and more like a glissade: the eye moves very slowly and at a more constant velocity (Kapoula et al, 1986). Whereas the glissade is the result of the eye settling to a new steady level of excitation, postsaccadic drift is thought to be the result of slow changes in the excitation itself, very likely as the result of drift in the saccadic integrator, moderated by the mechanism of slow visual oculomotor control (described in section 6.3.2: see for example Kowler and Steinman, 1979a; 1979b), and perhaps also by vergence. Thus postsaccadic drift is a prominent feature of the saccades of patients who have become blind well after birth (Leigh and Zee, 1980), demonstrating that in this case at least, postsaccadic drift is not primarily driven by vergence. A curious feature of the way in which the oculomotor

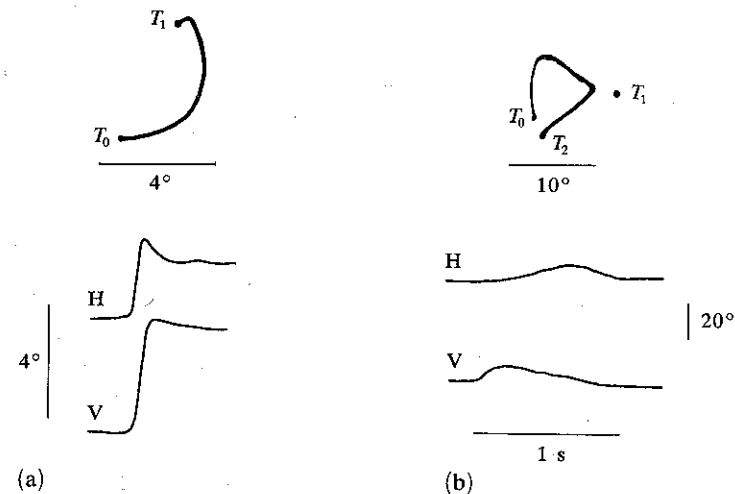


Figure 4.7. Two-dimensional trajectories of oblique saccades: (a) in man (after Bahill and Stark, 1979); and (b) in the cat (Evinger and Fuchs, 1978). The records underneath show the corresponding horizontal (H) and vertical (V) components as a function of time. The cat saccade is an aberrant one, the saccade being modified in midflight. T_0 is the initial target position, T_1 and T_2 are successive ones.

system copes with peripheral defects of the mechanics is that modification of the pulse-step pattern can apparently only be made for *both* eyes equally: with unilateral weakness, one eye will make accurate saccades at the expense of the other, which shows overshoot and a prominent postsaccadic drift (Kommerell et al, 1976; Optican, 1985).

Finally, it should perhaps be borne in mind that some of the postsaccadic behaviour that has been described may possibly be caused by movements of the eye other than simple deviations, in particular the retraction and torsion that has been described by Enright (1987). An observation that reinforces this view is that neurons in the oculomotor nuclei sometimes show slow postsaccadic changes in firing rate that are *not* accompanied by corresponding deviations of the eye (Mays et al, 1985).

One consequence of all this irregular behaviour is that the two-dimensional trajectory of an oblique saccade is generally far from straight, despite the pulse-stretching mechanism described in the previous section; this is exacerbated by the lack of synchrony between the onset of activity in the two pairs of muscles that is sometimes observed (figure 4.7).

4.1.3 Latency

The complexity of the calculation necessary to transform retinal distances into eye movements is reflected in the rather long reaction times associated with saccadic movements. A common experimental arrangement is for the subject to fixate a stimulus light that is suddenly switched off while a nearby light is simultaneously switched on: the saccade the subject then makes to the new target typically does not begin until about a fifth of a second later. A puzzling feature of saccadic latencies is their random *variability*: latency histograms in such an experiment typically stretch from some 120 ms up to 350 ms or more. The distribution is skewed, with a long tail towards longer latencies: by plotting it not as a function of latency but of its reciprocal (*promptness*), a curve that is very close to a normal distribution is obtained (figure 4.8) (Carpenter, 1981), generating a straight line on probit paper. A very similar relationship is found for many other kinds of reaction time data, and leads one to ask what type of mechanism might give rise to such a distribution.

The simple notion of transport delays along a chain of neurons seems out of the question, since the number of neurons in such a chain would have to be considerably greater than any plausible estimate of the neural length from retina to eye muscles; and in any case, the very large degree of *variability* would be hard to explain in such a scheme. Nor does it seem likely to be the result of the time it takes to calculate the neural signal necessary to move the eye to the required position. We shall see later that visual information arriving as little as 50 ms before the saccade can be used to help compute its size (Barmack, 1970b), and short-latency saccades on the whole show no more scatter than long (Becker and Jürgens, 1979). Rather, this delay seems to have to do with the task of deciding whether in fact there

is a target at all, and if so, whether to look at it, something altogether more high-level and complex. The complexity of such a process may not seem so obvious in the very artificial laboratory situation of a single light that may either be illuminated or not, but is immediately evident when we consider saccades under natural conditions. Here we are surrounded by a rich variety

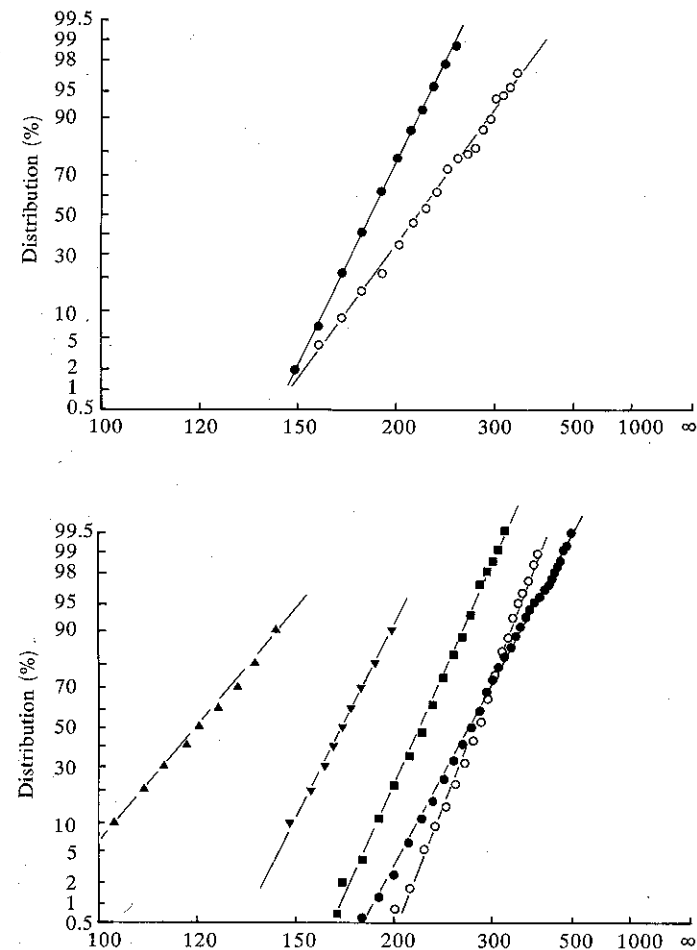


Figure 4.8. Top, cumulative distributions of latency for human (●) and simian (○) saccades, plotted on scales such that a normal distribution of reciprocal latency (promptness) would result in a straight line (data from Fuchs, 1967a; Carpenter, 1981). Bottom, reaction times for some very different tasks plotted in the same way, to show the generality of the relationship. ● pressing one of two keys in response to lights (Welford, 1959); responding to a ▲ 100 dB or ▼ 50 db auditory stimulus (McGill, 1963); ■ responding to a tactile stimulus (Kiesow, 1904); and ○ in a visual discrimination task (Johnson, 1918).

of possible objects, which do not obligingly signal themselves by lighting up in the manner of laboratory targets, and whose relative importance depends on their meaning: are they *worth* looking at? That it is this kind of decision process that ultimately determines saccadic latencies is evident from experiments where the subject is required to make an oculomotor response to a visual target that is not simply that of looking at it. The subject may for example be required to make a saccade that is deliberately directed in the opposite direction to the stimulus (an 'antisaccade': Hallett, 1978; Hallett and Adams, 1980), or a colour change in the fixation light may tell the subject which of two alternative targets to fixate: in all such cases, as in those where the target is difficult to detect because of reduced contrast or higher noise-levels, latencies are greatly increased. How might such decisions be made?

One plausible notion is that for each possible action that might be taken, there is a corresponding *decision signal* that increases in response to stimuli that make that action appear more desirable and decreases with those that seem to make it less so. When this signal reaches some predetermined threshold, the command to carry out the action is given, and the decision signals of all the possible actions are reset to zero (figure 4.9). In effect, the system runs a race between the different possibilities, the first one to reach threshold being the one that determines the response. The rate of rise of such a signal will be expected to depend on the rate at which information is received (itself depending on the signal-to-noise ratio), but may also be expected to have a deliberate *random* element, to prevent responses from becoming too stereotyped. If this random component is taken to be Gaussian, the distribution of latencies expected from such a model in response to a novel stimulus is precisely of the rather odd kind previously noted, namely normal with respect to reciprocal latency rather than to




Figure 4.9. A model of the saccadic latency process. Left, the activity of a decision unit rises linearly at a rate R on presentation of the stimulus, and initiates a saccade when its activity reaches a predetermined level, L ; it then resets all other units. Thus the first unit to reach its threshold determines the choice of target and the latency. Right, if R is subject to random Gaussian variability, the resulting distribution of latencies will be Gaussian with respect to the *reciprocal* of the reaction time, rather than to the reaction time itself.

latency itself (Carpenter, 1981). The variability is on a rather short time scale: there is no detectable correlation between the latencies to two target jumps 500 ms apart. Its origin is certainly central rather than peripheral, at a rather high hierarchical level: latencies are generally identical in the two eyes (Williams and Fender, 1977), and during many kinds of complex natural movements, such as reaching out for an object, there is a very strong correlation between saccadic latency and that of head and arm movements (Herman et al, 1981; Biguer et al, 1982; Zangemeister and Stark, 1982a). In the same way, if a subject is asked both to look at, and also to reach out for, one of a pair of targets simultaneously presented, both eye and hand make the same choice (Gielen et al, 1984); the decision mechanism is evidently not merely a part of the oculomotor system. Possible neural mechanisms that might correspond to such a decision process are considered in chapter 12.

Cumulative plots of the distribution of saccadic latency as a function of promptness greatly facilitate the comparison of the effects of various stimulus conditions on saccadic latency. With reduced illumination, latencies are increased (figure 4.10) (Wheeless et al, 1967), although the effects here are not very marked until the luminance of the target is actually below foveal threshold, introducing an obvious difficulty of interpretation! Latency increases to some extent with increased saccade amplitude, for example by some 40 ms for a 40° movement (figure 4.11, left) (White et al, 1962; Bartz, 1962). Paradoxically, it also increases if the required amplitude is




Figure 4.10. Saccade latency and luminance: cumulative latency probabilities are plotted as in figure 4.8 for five different luminance conditions: \blacktriangle target luminance equal to the foveal threshold; \bullet 1 log unit, \blacksquare 2 log units above foveal threshold; \circ 1 log unit, \square 1.5 log units below foveal threshold (data from Wheeless et al, 1967).

reduced (figure 4.11, right) (Wyman and Steinman, 1973b); it seems that the probability (per unit time) of making the saccade falls sharply if the size of the positional error is very small. It is probably true to say that no target displacement is so small that it will not eventually evoke a saccade: certainly saccades can be repeatedly elicited by target movements rather less than $10'$ of arc in amplitude (Timberlake et al, 1972; Wyman and Steinman, 1973a; Haddad and Steinman, 1973; see section 6.4).

But the greatest changes in latency are caused by providing the subject with prior information about the saccadic target. If the subject knows in advance *where* the target is going to appear, he or she tends to anticipate and shows shorter latencies than if the target can appear in one of two possible positions. [On the other hand, the latency to a target that has just *been* looked at is found to be increased slightly (Vaughan, 1984).] However, if the number of possible target positions is increased beyond two, the reaction time is apparently not correspondingly increased (Saslow, 1967b): saccadic reaction times thus differ in this respect from what is found in other reaction time experiments as the number of possible responses is increased (see for example Edwards, 1969). If the subject knows both where and *when* the stimulus is to appear—as for example in tracking a spot that is jumping back and forth at regular intervals—performance rapidly improves, and after a few cycles the subject can produce saccades that are virtually in synchrony with the stimulus movement (Westheimer, 1954b; Stark et al, 1962; Dallos and Jones, 1963; Fuchs, 1967b) (figure 4.12). As was previously noted in the case of sinusoidal tracking, monkeys are apparently unable to make use of this kind of redundancy in the input, and the saccades they make to repetitive stimuli show no improvement with time (Fuchs, 1967b).




Figure 4.11. Left, cumulative plots of promptness, as in figure 4.8, of saccades of different amplitudes: ● 10°; ■ 20°; ▲ 40° (data from White et al, 1962). Right, average latencies (for two subjects) of saccades in response to rather small target displacements, as a function of their size (data from Wyman and Steinman, 1973b).

In the same way, any kind of *warning* signal before a saccadic stimulus, that helps the subject to predict when the target will appear, reduces latency (Ross and Ross, 1980). The warning may simply consist of extinction of the fixation light at a fixed interval before illumination of the target; in this situation saccades with unusually short latencies ('express saccades') are observed both in human subjects and monkeys, particularly if the target location is also known in advance (figure 4.13) (Saslow, 1967a; Becker, 1972; Findlay, 1981a; Boch et al, 1984; Fischer and Ramsperger, 1984; Fischer et al, 1984). Histograms recorded under such conditions tend to be bimodal, the longer mode being nearly normal and the shorter one corresponding to anticipations. Indeed, even in the classical saccadic latency




Figure 4.12. Progressive reduction in saccade latency with square-wave targets: left, human eye movement response (thick line) to target jumping from side to side at a frequency of 0.8 Hz; right, progressive reduction of latency as a function of the number of cycles of such a stimulus of 0.5 Hz (●) and 0.8 Hz (○) (after Fuchs, 1967b).




Figure 4.13. Mean saccadic latency for different relationships between the time of extinguishing the fixation light and of illuminating the target. On the abscissa, negative times correspond to the case where the fixation light is extinguished before presentation of the target ('gap'), and positive times to when it is after ('overlap') (data from Saslow, 1967a).

experiment in which the fixation light is extinguished at the same moment as the appearance of the target, this extinction still serves as an extra cue that reduces saccadic latency. If the fixation light is left on, so that the stimulus is simply the extra appearance of the target, latencies are increased by 100 ms or more (figure 4.13) (Saslow, 1967a; Ross and Ross, 1980; Reulen, 1984a). It is arguable in fact that this form of saccadic stimulus is both simpler than the traditional one, and also more natural: it is not often in everyday life that something we are looking at happens to vanish at the very same instant that another object suddenly appears elsewhere in our visual field! The classical stimulus is really one of sudden *movement* rather than of simple stimulation of a particular retinal locus, and from the point of view of latency the extinction of the fixation light seems to be more important than the appearance of the target.

Another factor influencing the latency of saccadic movements is the apparent *refractoriness* that is observed after a saccade; as a rule, one saccade cannot follow another with an interval of less than some 180 ms. This can most easily be demonstrated by arranging for the visual target to switch back to the centre again after the initial displacement, providing a brief pulse of displacement rather than the step used in conventional saccade measurements (figure 4.14). If the duration of this displacement pulse is less than some 150 ms, one of two responses may occur: either the eye simply ignores the brief excursion, and does not move, or it makes a full saccade to the eccentric target, followed by a return movement. This return movement is not separated from the excursion by a time equal to the width of the pulse, as would be expected from a simple delay, but by a relatively fixed interval of some 150–200 ms (Westheimer, 1954b). It makes little difference if the second saccade is in a different meridian from the first (Feinstein and Williams, 1972). However, examination of large numbers of responses to short pulses of this kind, or to pairs of target displacements in the same direction (*double steps*) shows that refractoriness is not always strictly observed, and that a second saccade may be initiated before the first




Figure 4.14. Saccadic responses to long and short pulses of displacement. In the second case it is clear that the initial saccadic response is unaffected by the fact that it is completely inappropriate—the target has already returned to its starting position—and the latency of the return saccade is very greatly increased.

would normally have finished: this modification of the first may occur within 80 msec of the second step, a time that is considerably less than the normal latency (Bahill and Stark, 1975b; Becker and Jürgens, 1975; Evinger and Fuchs, 1978; Findlay and Harris, 1984) [see figure 4.7(b)].

A similar refractoriness is observed in the microsaccades of fixation (Nachmias, 1959), a fact which raises a doubt as to whether conventional measurements of saccade latencies really measure reaction times at all. For if, in the course of fixating the central target, the subject happens to have executed a microsaccade very shortly before the target movement occurs, the saccade will be delayed by an interval representing the refractory period, rather than the true reaction time. Indeed, one could imagine a microsaccade occurring up to one reaction time *after* the moment of target motion (correcting for where the target *was*), resulting in a maximum latency of one refractory period plus one true reaction time; the minimum possible latency is of course the reaction time. Finally, one can demonstrate refractoriness of the quick phase of nystagmus after voluntary saccades (Judge, 1973), underlining once again the probable identity of these two movements. However, a quick phase does *not* apparently delay a subsequent saccade in the expected manner, perhaps because a novel visual stimulus in some way suppresses the more primitive nystagmus mechanism. The possibility that refractoriness is the result of some kind of intermittent gating process is discussed in section 4.2.2 below. But it is also a direct consequence of the decision mechanism proposed above, since it will take a certain period of time after all the decision units have been reset at the initiation of a saccade for any of them once more to reach threshold.

General factors that increase latency include age (Abel et al, 1983), mental state (Kapoula, 1984), and certain kinds of psychopathology such as schizophrenia; in this respect saccadic reaction times behave very similarly to other kinds (Nettelbeck, 1980; Done and Frith, 1984). Increased latency can also be characteristic of amblyopia (Ciuffreda et al, 1978a; 1978b).

4.2 The control of saccades in relation to the visual stimulus

4.2.1 Correction saccades

A rather puzzling feature of large saccades is that they almost invariably fall short of their targets, and are followed—after the expected refractory period—by a secondary, or *correction*, saccade (figure 4.1). Making a saccade to a target is thus not unlike taking a hole at golf. For target amplitudes greater than some 20° the size of the correction saccade is typically about 10% of that of the primary component, though there is a tendency for correction saccades to the midline to be smaller than when the movement is to the periphery (Becker, 1972; Kapoula and Robinson, 1986). As the target amplitude is reduced, the percentage is also reduced: at around 5–10° the primary saccade is accurate, and below that amplitude *overshoot* rather than *undershoot* is seen (Clark, 1936; Weber and Daroff, 1971; 1972; Henson, 1978; 1979; Fischer and Kunz, 1981; Kapoula, 1985). However, there is an

oddity about these observations, which is that the crossover point on the graph of primary saccadic amplitude against target amplitude depends on the range of target amplitudes presented in any particular session, undershoot and overshoot depending to some extent on whether an amplitude is *relatively* bigger or smaller than others in the range recently presented [the *range effect*, commonly observed in manual tasks as well (Kapoula, 1985; Kapoula and Robinson, 1986)].

Are such saccades genuinely corrective, the result of seeing at the end of the primary saccade that there is still some way to go? If so, then in the case of the largest saccades, which can last 100 ms or more, it would imply a very short reaction time (rather less than 100 ms) for the generation of the correction saccade, though one that is perfectly compatible with latencies seen when, as here, the *time* of the saccadic stimulus is predictable. Perhaps indeed this is the explanation for having a refractory period at all: the system must obviously ensure that the eye has stopped moving at the time the estimate of the error is made, and this could perhaps be arranged most easily by holding up the calculation until some fixed time has elapsed after the beginning of the last one, long enough to allow for even the largest eye movements. One might suppose that it would be more efficient to allow more settling time for larger saccades than for smaller: but this does not seem to occur. The latency of the correction as measured from the *end* of the main saccade is greater, the smaller the size of the original movement (Becker, 1972); but the latency as measured from the *beginning* of the main saccade is roughly constant whatever its size.

However, a number of experiments indicate that the correction saccades seen in large movements are not in fact the result of sensing a visual error. They are still observed in total darkness, when of course no visual monitoring is possible (Becker and Fuchs, 1969), and also in 'antisaccades', where again there can be no direct error information from the retina (Hallett, 1978). Extinguishing the target just before the primary saccade does not abolish correction saccades, although they may then be less accurate: although visual feedback is not necessary, it can be used to increase accuracy provided it is available some 60 ms before the time of the secondary saccade (Pernier et al, 1969; Barnes and Gresty, 1973; Prablanc et al, 1978; Deubel et al, 1982). But in any case, the characteristics of the error in the first saccade do not suggest that the system is doing its best to get it right first time. If it were, we would surely expect the endpoint of the first saccade to be randomly distributed about the target, whereas in fact the first saccade shows rather little random variation and appears to be deliberately and quite accurately aimed short of the target (Becker, 1972). Nor is it simply that the eye is incapable of executing a single saccade of more than a certain amplitude: a saccade of 30° is executed in two stages, even though the first component of a 40° saccade is itself a saccade of more than 30° (figure 4.1). Further evidence that the eye is very deliberately

undershooting comes from experiments in which the target is moved back a little way at the moment of each primary saccade, so that what was previously an undershoot now lands on target. Far from being contented with the resultant hole-in-one, the saccadic control system quite rapidly adapts to the new state of affairs by reducing the amplitude of the primary component so as to require secondary saccades exactly as before (Henson, 1978).

Why are large saccades executed in this peculiar way? One might argue that undershoot simplifies the calculation of the secondary saccade, since at least its *direction* is known; perhaps also it is helpful for the same side of the brain to be involved as for the primary saccade (Becker 1972; Henson, 1978). Or perhaps it is that by providing a comparison of expected error with actual error, undershoot can help the system to optimise its parameters in relation to the oculomotor plant; deliberate errors are sometimes used for a similar purpose in man-made control systems (Elgerd, 1967; Optican, 1982). But while it is quite plausible that the system might deliberately miss the target so as to improve its *theoretical* performance in the long term, it is perhaps less plausible that the system should choose to play educational games with itself at the very moment when it is required to perform as well as possible. Another suggestion is that undershoot is the result of the head being constrained: under natural conditions a head movement is initiated in parallel with the saccade, and it may very well be that the oculomotor system is trying to make allowance for a part of this expected contribution by the head. However, the sequence of saccades often seen with large shifts of gaze when the head is free suggests that this is not the case (see for example Guitton et al, 1984). A final possibility comes from a consideration of natural saccadic targets. As was emphasised earlier, targets do not commonly jump instantaneously from one fixed position to another in the real world. Usually a new target appears in the periphery, and it appears there because it is *moving into* the visual field. Perhaps then the undershoot in fixating sudden peripheral targets arises through anticipation of this movement, the saccadic system assuming that by the time the saccade occurs, the target will be that much nearer the centre.

But in addition to these rather pathological correction saccades following large eye movements, there are also those associated with much smaller movements, that are certainly simply the result of sensing an error of fixation. One way of demonstrating this is by opening the intrinsic visual feedback loop (in the manner described in the previous chapter) so that every time the eye makes a saccade toward the target, the target simultaneously jumps away from the eye by the same amount. One then finds that the response of the eye to an initial step displacement consists of a regular series of saccades of equal amplitude, forming a staircase in which the treads are some 200 ms long [figure 4.15(f)] (Young and Stark, 1963; Robinson, 1965; Fuchs, 1967a; Fleming et al, 1969; Evinger and Fuchs, 1978). If the negative feedback is increased beyond its natural value of -1 , a step displacement of the target produces oscillations (again, as in the case of the

smooth pursuit mechanism) which either die away or increase in amplitude, depending on the value of the feedback gain [figure 4.15(a)–(c)]. For a gain of -2 they neither grow nor shrink, and the result is a stable oscillation; with larger negative gains, around -2.3 to -5 , spontaneous oscillation is seen. All this is of course only what would be expected if the saccade size is matched accurately to the perceived target error, and tells us little new about the system beyond confirming that correction saccades can be generated by visual feedback, and really are corrective. The correctiveness of *microsaccades* is considered in section 6.4.




Figure 4.15. Saccadic responses to a step of target position (top trace) under different artificial feedback conditions: -1 is the natural condition (after Young and Stark, 1963).

4.2.2 Saccadic intermittency

Given that saccades are discrete, prepackaged, events, one might wonder whether the visual information used to generate them is also obtained in a similarly discontinuous manner. Is the visual scene *sampled*? The fact that every eye movement causes a shift of the visual world puts certain bounds on the periods of time preceding a saccade during which visual information about the position of a target is actually useful. Information obtained before a previous saccade is clearly out of date; equally, information acquired within one reaction time of the new saccade is too late. Thus there is a definite time window outside which visual information about target position is of little use. Note, however, that this does not necessarily apply to *velocity* information, which is not rendered out-of-date by an intervening saccade: this point is considered in the next section. It is also true that for *flashed* targets, information acquired during or even before a prior saccade might be better than none at all (see for example Hallett and Lightstone, 1976a; 1976b; and section 11.1); but the argument presented here concerns 'natural' objects that are continuously visible. Furthermore, not all the

information gathered during this window is equally useful: clearly more weight ought to be attached to the most recent information. In fact, in a noise-free situation (where the target is not subject to relatively fast random motion, and where the monitoring device has perfect certainty of the target's position at any instant) the best strategy is to make a rather short sample of the visual information, at the last possible moment. The greater the noise in the system, the longer this sampling interval will have to be for accurate performance. A simple way of meeting all these requirements would be to sample the visual scene briefly at regular intervals of time, the intervals being chosen to be greater than the sum of the reaction time and the duration of the longest saccades, that is, of the order of 200 ms. More than forty years ago, the desirability of this kind of intermittency in motor control in general had been identified by Craik (1966): "It is possible here that a system inevitably containing time-lags (for example, a human being) acts best if it acts discontinuously, so that its actions are always based on the last 'static' situation (that is, the instant when the last response, with its time-lag, has taken effect and the comparison of that with the situation to be met is observable) rather than on a shifting and misleading state of affairs". The suggestion that intermittent sampling might be used by the oculomotor system seems first to have been made by Vossius (1960) and by Young (1962; Young and Stark, 1963). Does it in fact occur?

The experiments that best support the notion of intermittent sampling are those that involve pulses of target displacement, of the kind described above in connection with the demonstration of the refractory period (for example Westheimer, 1954b). If the sampling pulse happens to coincide with the brief displacement of the target, one would expect the system to respond with a saccade; if not, no movement would be expected. This is exactly what is observed: on some trials a saccade is made to the pulsed target, whereas on others nothing happens at all. On the whole the responses are 'all-or-none', implying a rather short sampling period, although sometimes one may observe saccades of intermediate size.

Similar results have been obtained by using a slight modification of this arrangement, namely the *pulse-step* stimulus (Wheless et al, 1966; extended by Komoda et al, 1973; Carlow et al, 1975; Becker and Jürgens, 1979; and others). Here the target first jumps to one side of a central fixation spot, and then after a brief interval W jumps to the other side of it and stays there (figure 4.16): to prevent adaptational changes, this stimulus is randomly interleaved with suitable controls. The resultant response is again found to be probabilistic, the eye making a saccade either to one or to the other side of the centre. As the duration of the pulse is increased, the probability of the initial saccade being in the same direction as the pulse—a *type A* response—increases linearly (figure 4.16). [The exact relationship varies considerably from subject to subject, but this variability can be greatly reduced if allowance is made for the mean reaction times of different individuals (Lisberger et al, 1975).] Such a relationship is exactly what

would be expected if brief samples of the visual input were being made at regular intervals: the longer the pulse, the greater in proportion will be the probability that the sample will happen to fall within it. Even with a pulse lasting 200 ms, a substantial percentage of the responses do not follow it, suggesting that the sampling interval is somewhat longer than this. The fact that zero probability is reached with a pulse duration of some 40 ms might be taken to suggest that the sampling period is of about this size. This view is supported by the observation that both here and also in the case of double-step stimuli a few saccades of intermediate size can be observed in the transition region between the full-blooded type A and type B and responses (figure 4.17) (Becker and Jürgens, 1979; van Gisbergen et al, 1981a; Ottes et al, 1984), suggesting low-pass filtering of perceived target position of the kind that a finite sampling window would be expected to produce.




Figure 4.16. The pulse-step experiment. On the left, the upper trace shows the target movement: it steps to one side for a duration W , then steps back to the same distance on the other side of the midline. The eye (lower trace) either follows it back and forth (type A response), or moves only in the final direction (type B response). The graph shows the probability of a type A response as a function of W (● Wheelless et al, 1966; ■ Carlow et al, 1975; ○ Komoda et al, 1973; □ Becker and Jürgens, 1979).




Figure 4.17. Transitional behaviour with step, double-step, and pulse-step targets. In each case, the endpoints of saccades that finished at particular times are shown as dots, superimposed on the time course of the stimulus itself (data from van Gisbergen et al, 1981a: subject JG).

These results clearly provide *prima facie* evidence for the idea of intermittent sampling. But if we start to look at what the model implies about the distribution of *latencies* in different saccadic tasks, it becomes obvious that there is something terribly wrong. Consider first an ordinary single target step: on different occasions a stimulus sample will occur with responses at different times relative to the step itself, so that we would expect to get a random variation of latencies with a spread of some 200 ms, which is perfectly true. But the expected *distribution* of these latencies will be a rectangular function, flat in the region λ_0 to $\lambda_0 + s$, where λ_0 is the reaction time of the system, and s is the intersample interval [figure 4.18(a)]. Actual latency distributions are not of this form, as we have seen, but it is likely that other sources of variation exist which will tend to round off the corners. (Latour claimed in 1967 that there was a *periodic* component to distributions of saccadic latency, linked to the EEG; but its frequency was an order of magnitude greater than the kind of clock being considered here.) Consider next a pulse-step stimulus in which the width of the pulse is W . The expectation then is that all responses whose latencies lie in the region λ_0 to $\lambda_0 + W$ will follow the pulse, whereas those in the region $\lambda_0 + W$ to $\lambda_0 + s$ will follow the step, resulting in a pair of distributions for the two responses, that do not overlap [figure 4.18(b)].

The actual distributions found in such experiments (figure 4.19) do not look like this. Although the number and distribution of pulse-following responses agree well with the model (as can be seen by comparing their cumulative distributions with those for the simple step stimulus: see also Lisberger et al, 1975, and Becker and Jürgens, 1979), the distribution of latencies for the responses that follow the step is markedly shifted to the right, by some 130 ms. Now if the *only* information available to the system about the position of the target was obtained by a series of samples more than 200 ms apart, the system could not possibly know that the step had




Figure 4.18. Distributions of reaction times expected in the Wheelless experiment, on the very simple hypothesis of regular sampling at an interval s (uncorrelated with the stimulus) and a simple reaction time: expected distribution (a) in response to a step of position, and (b) for each of the two kinds of response to the Wheelless stimulus; W is the duration of the stimulus pulse.

been preceded by a pulse in the opposite direction. Yet the system *must* in fact know that these are not just normal steps, for it responds to them with latencies that are enormously increased. One is therefore forced to conclude that the system does *not* rely solely on intermittent brief samples in calculating its saccades. In the pulse-step experiments the presence of the step can in some way cancel the pulse response that would otherwise have occurred, and conversely the presence of the pulse can profoundly modify the response to the ensuing step. Rather similar effects were shown by Becker and Fuchs (1969), using stimuli consisting of a large step followed by a smaller step. If the interval between the two is large enough for the second step to take place after the saccadic responses to the first, extremely large latencies—again of the order of 400 ms—are observed in the second saccade.

Observations of this kind do not rule out the notion of intermittent sampling altogether: but they imply that visual input is *continuously* available for the purpose of cancelling saccades that are in the course of elaboration and that are going to be inappropriate. A hypothetical system of this type is

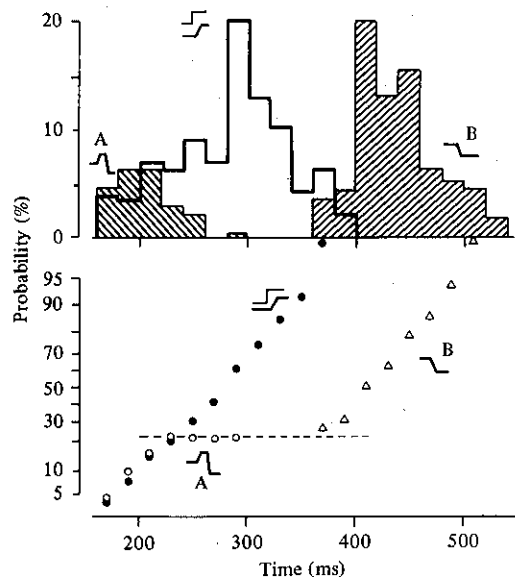


Figure 4.19. Statistical properties of the latencies in the experiment shown in figure 4.20. Above, the shaded areas show the observed latency histograms for the two kinds of response, type A and type B: the unshaded histogram is for a simple single step of the same amplitude under identical conditions. Below, the same data are plotted in the form of cumulative probabilities, showing that the type A responses lie on the curve for ordinary single-step saccades, whereas the latencies for type B responses are as if the single-step distribution were cut in two at around $t = 230$ ms, and the later responses delayed as a whole by some 130 ms (data from Wheelless et al, 1966).

represented in figure 4.20. After experiencing a delay λ_v , visual information is used in two ways: it is intermittently allowed access to the process that calculates the next saccade, and it is also able at any time to cancel a calculation that is in progress, if some change in target position has occurred which means that the saccade is going to be wrong. Figure 4.21 shows the expected distributions from this model in the Wheelless experiment: with a suitable choice of parameters ($s = 200$ ms, $\lambda_v = 130$ ms, $\lambda_p = 50$ ms) it accounts well for the time relations observed in the actual distributions, but not their shapes (because of the assumption of an essentially flat stochastic process). It is also necessary to assume that the cancel signal lasts at least 50 ms, as otherwise the earliest type B responses, due to samples occurring just after the second target displacement, would not be cancelled: the corresponding 'hole' in the distribution is obvious in the data of figure 4.17, but it is evidently not absolute. A similar mechanism of cancellation can also be used to explain the delayed responses observed when a subject is




Figure 4.20. A hypothetical model that shows some of the features observed in the saccadic responses to pulse-step target movement.




Figure 4.21. Latency distributions that would be predicted by the model of figure 4.20: above, expected distribution of latencies to a single-step movement; below, expected latencies of the two kinds of response to pulse-step stimuli, for $W = 100$ ms.

required to make an 'antisaccadic' response to a target (section 4.1.3) (Hallett and Adams, 1980).

However, there is another respect in which the idea of intermittent sampling must be modified. If the sampling clock is ticking away at a constant rate, regardless of any visual stimuli or saccadic responses, then there should be a strong correlation between the latencies to two consecutive target steps—in fact the two saccades should always be separated by an exact multiple of the clock interval. Täumer (1975) found some evidence for this tendency in certain of his subjects, though its contribution was small in relation to the overall variability of the latency of the second saccade. With a pulse-step stimulus, Becker and Jürgens (1975) found that the smaller the interval between the second step and the first saccade the longer, on average, the interval between the two saccades—as would be expected if the second response latency were effectively uncorrelated with the first—whereas longer intervals tend to lead to single saccades that jump immediately to the final position of the target. The most critical test of the idea of a regular sampling clock is to arrange for the second step of a double-step stimulus to be triggered at a fixed time after the first saccade, rather than after the first target step (a DT stimulus). This ought in effect to synchronise the second step to the supposed clock, and should therefore entirely abolish (or at least greatly reduce) the variability of the second latency. This simply does not happen: the scatter of the second latencies is about the same as that of the first (figure 4.22). Sampling does *not* occur at regular intervals of time.

So does it occur at all? In particular, is there some other way of explaining the probabilistic nature of the responses to pulse-step targets?




Figure 4.22. Saccadic latencies for one subject in a DT experiment, where the second target step is triggered by the first saccade. The latency of the second saccade is plotted as a function of that of the first; the interval, D , between the beginning of the first saccade and the appearance of the target was 250 ms. The scatter of the second response is evidently no less than that of the first, and there is no obvious correlation between the two.

Robinson (1973a) has presented a model of saccadic timing that postulates no clocks or sampling, but is still consistent with the kinds of observations described above. In essence, he proposes an array of independent processors acting in parallel and corresponding to different points on the retina. Each takes a certain time to reach its decision about the command to be issued, and when the decision is reached, it both initiates a saccade, and cancels any decisionmaking activity that may be going on in the other units. The time needed for any particular unit to make its decision is subject to random variation: thus if—as in the Wheelless experiments—one unit is activated only a short period of time after another, it may nevertheless complete its operations sooner and thus produce a saccade that jumps immediately to the final target position. By choosing plausible probability functions for the decision times of the individual units, it is possible to make quantitative predictions of many kinds of double-step responses, without having to postulate the existence of a regular sampling process. In particular, if each decision unit works in the way suggested in section 4.1.3 above (Carpenter, 1981), one may obtain quantitative descriptions of the latency distributions that are actually observed.

One factor that complicates the consideration of the responses to double steps is that the relative *directions* of the two components interact with one another in a complex way. Thus the observed amplitude-transition function (figure 4.17) is much wider when the two steps are in the same direction rather than in opposite directions (Becker and Jürgens, 1979), and the correlation of the latencies of the first and second response is different in the two cases (figure 4.23). When the target movement has both vertical and horizontal components the results are more complex still (Feinstein and




Figure 4.23. Correlation between first and second saccadic latencies plotted for one subject as in figure 4.22, but for an ordinary double-step experiment in which the interval between the steps was 150 ms, and the second step was either in the same direction as the first (\bullet), or in the opposite direction (\circ).

Williams, 1972; Lisberger et al, 1975; Hou and Fender, 1979; Findlay and Harris, 1984; Ottes et al, 1984), suggesting that there are special subsystems that deal with direction and magnitude.

4.2.3 Saccades to moving targets

We saw in the previous chapter that smooth pursuit is distinctly sluggish, exhibiting large phase lags even at quite moderate frequencies. In the case of a step change in target velocity, as when a target that is initially stationary moves off with constant velocity, this sluggishness implies that in the absence of any other mechanism, although the eye will quite soon reach the same *velocity* as the target, it should do so with a permanent error in *position* (figure 4.24). In practice, however, what happens is that the saccadic system intervenes, generating a saccade that is of exactly the right size to correct the positional error. The astonishing thing is that this correction saccade is typically initiated long before the positional error is fully manifest (Westheimer, 1954b; Rashbass, 1961). One is forced to conclude that the saccadic control system must estimate from the velocity of the target how large a correction is going to be required: in doing so it must also of course allow for the dynamic properties of the slow pursuit mechanism. It appears that this estimate is added to any other step displacement that may be required. If at the moment of moving off, the target is also simultaneously displaced, the resultant saccade is still of the right size to correct the total positional error (figure 4.25) (Fuchs, 1967a; Evinger and Fuchs, 1978). In particular, if the target is displaced backwards by a suitable amount at the



Figure 4.24. Response to a target moving off at constant velocity: left, response of smooth pursuit system alone leads to a permanent positional lag because of time required for the eye to approach the target velocity; right, diagrammatic representation of how a saccade of appropriate size can eliminate positional error of this kind.




Figure 4.25. Monkey eye movements in response to ramp and ramp-step target motion: (a)–(c), ramps of 10° s^{-1} , showing clear saccadic components; (d), (e), ramps of 13° s^{-1} combined with backwards steps: the saccades are reduced or eliminated (Fuchs, 1967a).

start, no saccade at all is necessary—and none occurs. If the target is extinguished shortly before the first response, a saccade still takes place to the point where it would otherwise have arrived, and the eye continues for some 200 ms with the previously appropriate velocity (Westheimer, 1954b).

This seemingly intelligent behaviour could in fact be produced by quite a simple mechanism. All the system needs to do is to monitor the velocity of the target, and generate a saccade whose size is proportional to this velocity, since (if linear) the permanent positional lag of the pursuit system will vary directly with the velocity of the stimulus. One could conceive of a more sophisticated system that took into account not only the velocity of the target but also any acceleration it might have, to form a better estimate of its future position. But experiments with targets moving off with constant acceleration show that this does not in fact occur, and that in these circumstances the extra error due to the acceleration remains uncorrected (figure 4.26) (Fleming et al, 1969). If the target velocity is very high, so that there is still a residual *velocity* error by the time the saccade has taken effect, one or more additional saccades may be made to correct for the positional lag that will result from it (Fleming et al, 1969). One might perhaps expect that in these circumstances the system would simply have arranged for the whole ultimate error to be cancelled at a stroke; but this would result in a larger mean deviation of the eye over the whole movement than piecemeal correction by a number of smaller saccades. It may in any case be the result of the saturation nonlinearity of the smooth pursuit mechanism.

This particular mechanism also shows a certain degree of adaptive plasticity; if monkeys are repeatedly trained at one ramp velocity, on switching to a different ramp velocity they tend to persevere with the (inappropriate) previous saccade size (Barmack, 1970a). A similar observation in man is that if the target velocity is altered at the moment of the saccade, so that the eye movement has an inappropriate velocity, with training the eye learns to move with a velocity matched to the second target velocity rather than to the first. These adaptational aspects of saccades are discussed in chapter 13. *Attention* can also be used to select one moving target rather than another. If while tracking one moving spot across an oscilloscope screen, it becomes necessary to shift the gaze to another spot moving with a different velocity in a different part of the field, a saccade of appropriate magnitude is generated to reach it despite the powerful and




Figure 4.26. Human responses to targets moving off with constant acceleration, showing gradually increasing error despite a sequence of corrective saccades (after Fleming et al, 1969).

misleading visual slip signal being generated by the objectively stationary surroundings (Atkin, 1969).

It was argued in the last section that intermittent sampling of the velocity of visual targets was unnecessary, since velocity information is not rendered out of date by intervening saccades. A simple experiment (Barmack, 1970b) shows that even though it is used to control saccades, velocity information is indeed continuously monitored [at least, by *some* subjects; others apparently do not (Heywood and Churcher, 1981)]. If we arrange for a target to move off at constant velocity from an initial stationary position (as in the previous experiments), suddenly reducing its velocity shortly after it has set off, the resultant saccade is found to be correspondingly reduced, by more or less the right amount to bring the eye finally on to the trajectory of the target, so long as the change in velocity occurs more than some 50 ms before the saccade (figure 4.27). It is difficult to avoid the conclusion that all the velocity information gathered for a considerable period of time before the saccade is available for the purpose of computing its size. It also seems that the latency for using it is very much shorter than that for positional information, a conclusion that is perhaps not so very surprising, bearing in mind the difference in complexity between velocity and position control that has repeatedly been urged in previous chapters.




Figure 4.27. Latency of responses to altered target velocity. Thin lines show target movement, thick lines show the eye response. In the control experiment the target moves off at constant velocity and the eye responds in the usual way with a saccade of appropriate size. In test runs, the target velocity is altered soon after it has started to move: nevertheless the saccade size is correctly modified, with considerably reduced latency (after Barmack, 1970b).

4.2.4 Selection of the saccadic goal

The question of how the saccadic system decides *what* to look at is obviously one that requires a consideration of topics like perception and attention that are well beyond the scope of this book. The subject has considerable applied interest, and much work has been done by psychologists on patterns of eye movements made when for example a jet pilot views his instrument panel, or a shopper scans the supermarket shelves. Tasks discussed in a recent symposium (Gale and Johnson, 1984) range from searching for lung tumours on radiographs and for blemishes on

apples to judging shot-putters and spotting tanks on battlefields. Although many experimenters have been interested in the *principles* that determine what features of a scene attract the gaze (for example Mackworth and Mackworth, 1958; Yarbus, 1967; Stark, 1971), their interests have tended to lie at a psychological rather than physiological level. Although Gould and Peeples (1970) claim that a subject's *interpretation* of a simple figure has no effect on the eye movements he makes when viewing it, and that it is only its 'physical attributes' that determine the movements, it is plain that this cannot be true of more complex targets such as photographs of faces, where the eye movements are obviously strongly influenced by the high-level perception of such 'attractive' features as the eyes and mouth (figure A1.11).

One physiological question that is of considerable interest is how, out of all the possible points on a particular visual object, the oculomotor system selects *one* point as the target for a saccade. Suppose for example that a subject is looking at a fixation spot with instructions to look at any visual target we may flash nearby on the screen. If the target we present consists simply of a small spot, then we know that the subject's gaze will jump quite accurately onto it. But if we give an extended object such as a triangle, it is not clear *a priori* whether the subject will fixate its centre, or an edge, or even one of the points. If we believe that the system behaves mechanistically, it must presumably have built into it some algorithm for converting two-dimensional retinal images into a single point which is the desired endpoint of saccades used for looking at it. An obvious approach to finding out how this algorithm works is to consider the simplest of all stimuli where the problem arises at all, namely a pair of simple targets presented simultaneously (Findlay, 1981b; 1982; Deubel et al, 1984). Under some conditions, an eye movement is made to a point lying between the two targets, generally nearer to the less peripheral target, to the bigger one, or to the brighter. Such results can be explained relatively simply by assuming that there are target-detectors with large receptive fields that are a function of retinal eccentricity; but such a model cannot easily be extrapolated to deal with edges and vertices and other salient aspects of more complex stimuli.

The model presented earlier, in which a large number of decision units compete in parallel to decide the next saccadic goal, can explain some aspects of the selection of saccadic targets. But these units cannot simply correspond to single points in the outside world: most of them must code for whole objects that cover an extended area of the retina. So how are their commands translated back into a single saccadic goal? For example, a subject can make saccades to fixate the vertices of a figure created by pairs of Julesz random-dot stereograms (Dimitrov et al, 1976), where the figure exists only at a high level in the brain and there is no corresponding retinal pattern that might be referred to. Similarly, the amplitudes of the eye movements made in scanning the pairs of intersections of the Müller-Lyer illusion figures show the same metrical changes as the perceived illusion

itself (figure 4.28) (Yarbus, 1967), and similar results are found for other illusions (see Coren's general review, 1981); again, there cannot here be reference back to the raw retinal image to determine the saccadic goal. On the other hand, if an illusory movement of a saccadic target is induced, by shifting a frame round it, the subsequent saccade is made to a position corresponding with the *retinal* rather than the perceived position: though if the retinal information is no longer present, so that the target position has to be remembered, the eye goes to the perceived location (Wong and Mack, 1981; Mack et al, 1985). It is not easy to formulate a plausible mechanism that will deal both with the global and with the local aspects of such a process.

Another experimental approach is to measure not the initial eye movements made in response to a figure suddenly presented, but rather the relative popularity of various parts of figures as resting places for the eye, as measured by the relative amounts of time spent there. For example, in free scanning of gratings the eye tends to spend more time at one part of each cycle than another, resulting in a pattern of average illumination that is itself grating-like (Arend and Skavenski, 1979). Earlier work by Kaufman and Richards (1969; see also Richards and Kaufman, 1969) showed that in simple line drawings the centre of gravity is attractive, as are also the edges and corners (figure 4.29): comparable findings have been presented by Lévy-Schoen (1973). Similarly, in scanning polygons (Bozkov et al, 1982)




Figure 4.28 (a) Two-dimensional records of eye movements made when viewing Müller-Lyer illusion figures, showing that the amplitude of the eye movements corresponds on the whole to the illusory size. The comparison can be made more easily in (b), which shows the time course of the horizontal eye movements made by a subject instructed to scan his eyes between the intersections of each figure (Yarbus, 1967).

the eye tends to linger at the vertices, even where these have been erased by curtailing the sides of the figure. In general it seems to be 'informative detail' that attracts the gaze (see Mackworth and Morandi, 1967; and also Findlay's intelligent discussion, 1981b).




Figure 4.29. Fixational tendencies in prolonged viewing of extended discs: the shaded areas include 86% of recorded fixations by two observers, and the dots show mean fixation positions (after Kaufman and Richards, 1969).

A further aspect of scanning by the eye relates to the question of the extent to which the *recognition* of objects depends on recognising the pattern of eye movements that it evokes, the notion that 'you know what you see when you see what you do'. The idea seems first to have been proposed by Bain (1855), who considered that the perception of "... naked outlines, as in the diagrams of Euclid and the alphabetical characters, are to say the least of it, three parts muscular and one part optical". Subjects repeatedly viewing a set of pictures generate relatively stereotyped patterns of scan that are specific for each picture, and one may choose to think of such patterns as the eye looking in sequence at individual features of the picture which are then ticked off on a mental checklist (Noton and Stark, 1971). For very complex recognitions (*is this flower a meadowsweet?*) this kind of serial recognition process is evidently used (*are the leaves interruptedly pinnate? the terminal leaflets three-lobed, and downy white on the underside? the calyces five-lobed?*); but equally, it is clear that even the most complex recognitions can ultimately be performed far more quickly than the time that would be required to make the requisite number of saccades. Faces presented tachistoscopically can still be recognised (Groner et al, 1984), and so of course can objects presented in mirror-reversal or at a different angle of view. One specific case where it might naturally be assumed that eye movements are used for recognition is *counting*: Landolt (1891) observed that objects closer than some 5' of arc apart could not be counted, and he supposed this to be due to an inability to make the necessary small voluntary saccades, which are of about the same size (Haddad and Steinman, 1973). However, direct experiment shows that accuracy is actually *improved* if saccades are suppressed (Kowler and Steinman, 1977), partly perhaps because of the danger of losing one's place if an inaccurate movement is made. But at a higher, more conscious level of thought, there is no doubt of the very strong connections between gaze and serial logical

analysis: they are particularly evident when recording the eye movements of someone playing chess (Tikhomirov and Poznyanskaya, 1967), when the contemplation of possible moves and their consequences is reflected in corresponding shifts of the gaze around the board.

One special and important case of visual scanning occurs of course in *reading*. Normal eye movements made while reading consist of a sequence of saccades along the text, at intervals of some 250 ms, with an amplitude and a frequency that depend on the skill of the reader, the difficulty of the reading matter, and details of the typography (for example Hering, 1879a; Erdmann and Dodge, 1898; Vernon, 1931; Gruber, 1962; McConkie and Zola, 1984) (figure 4.30). A curious feature of the saccades is that they are closely similar in skilled readers both of Chinese and English (Sun et al, 1985), despite the very different nature of the symbols in each case, and direction of movement: it seems that typography has evolved in each case in such a way as to present information at a linear rate that suits the eye movement control system. Saccades do not have to be particularly closely tailored to the visual stimulus for comprehension to be possible: reading can be performed quite adequately if the subject fixates a stationary point, while the whole text is stepped sideways in a saccade-like fashion at an appropriate regular rate (Bouma and de Voogd, 1974).




Figure 4.30. Typical eye movements made while reading, showing fixational pauses (F), interfixational saccades (S), and a regressional movement (R).

The control of eye movements in relation to the text being read presents many very interesting questions, particularly concerning *when* the eye chooses to move from one point to the next, and *where* it decides to move to. The dwell time for a fixation is not preset, but depends on the meaning of the word: but it would be unreasonably inefficient for the command for the eye to move to be issued only when all processing is complete, as the saccadic latent period would then be spent in idleness (Vonk, 1984; McConkie et al, 1984). Furthermore, while the word one is looking at is being processed, there must be simultaneous preliminary analysis, using parafoveal vision, of what is to come (Carpenter and Just, 1983; Jennings and Underwood, 1984). Saccades do not simply jump from one word to the next but are closely related to the information content of what is read, implying control at a very high level by the large-scale meaning of the text.

For example, the time spent fixating a pronoun that refers back to a previous noun depends on whether or not the reference is ambiguous (as in 'he put his hand on the stove because it was cold'), a distinction that clearly requires some knowledge of the properties of the real world (Kerr and Underwood, 1984). In situations of this kind, eye movements provide a subconscious running commentary on higher cognitive processes, opening up exciting possibilities for research.

Vergence

"These two eyes, although they be farre enough separated the one from the other, have such a fellow-feeling, and doe so well agree the one with the other in their actions, as that the one of them cannot move without, or otherwise than the other: for it is not in our abilitie, to looke up with the one and downe with the other, or els to stir the one and hold the other still."

So far, we have considered only *conjugate* movements of the eyes, that is movements which are closely similar in amplitude and direction in the two eyes and thus obey Hering's principle of 'equal innervation' (Hering, 1868). The requirements for binocular visual fusion as the gaze is transferred from one distant object to another clearly demand that this should be so. But if the image of an object that is near the eyes is to be brought to the fovea of each retina, it is equally clear that the lines of sight of the two eyes must be brought to converge from the parallelity they assume for very distant objects. Horizontal *vergence* movements are thus a consequence of overlap between the visual fields of the two eyes, and are most highly developed in those species that have forward-facing eyes. Although at one time vergence was commonly said to be restricted to the primates, careful observation indicates that not only do other front-eyed animals such as cats show a considerable, though infrequently exercised, power of convergence—up to 14° (Stryker and Blakemore, 1972; Hughes, 1972; Olson and Freeman, 1978)—but lateral-eyed animals such as the rabbit and certain birds can do the same (Collewijn and Zuidam, 1977; Zuidam and Collewijn, 1979; Bloch et al, 1981; Wallman and Pettigrew, 1985). Perhaps because of their relatively recent evolution, vergence movements are the slowest to reach full development in children (Fowler et al, 1985), and disturbances of vergence—resulting in double or unilaterally suppressed vision—are amongst the most frequent clinical ophthalmological symptoms; it is a common experience that it is these movements that are first affected by




Figure 5.1. Defining angles of convergence: on the left, I is the interocular distance, d is the distance of the point of fixation from the baseline, and θ is the angle of convergence. On the right, a prism whose power is one prism diopter (1Δ) deviating a light ray by 1 cm at a distance of 1 m.

fatigue, and alcohol and other drugs, sometimes long before any other behavioural manifestations.

The simplest way to specify a particular degree of horizontal vergence is by the angle between the two lines of sight, which is the angle subtended by the interocular distance at the point of binocular fixation (neglecting the fact that the line of sight may not pass through the centre of rotation). In clinical practice, it is usual to express this angle in *prism diopters* (Δ): a prism of 1Δ deviates a light ray by 1 cm at 1 m (figure 5.1). Thus for a point of fixation at a distance d m from the baseline, in the sagittal plane, the angle of convergence in prism diopters is given by I/d , where I is the interocular distance in centimetres: but this definition is not applicable to very close objects. A similar objection applies to the use of the *metre angle* (given by the reciprocal of d : see the discussion in Alpern, 1969b, page 110), and physiological studies of vergence are made simpler if the angle of vergence is specified directly in degrees or radians.

5.1 Sensory stimuli for vergence movements

People vary enormously in the degree to which their vergence movements are susceptible to *voluntary* control. Many find great difficulty in converging on command onto the rather artificial targets—for example, points of light in a darkened room—that are commonly used in the laboratory, while others can converge and diverge at will, whether or not there is anything to look at. Sir James Barrett (1921) describes a patient who could voluntarily diverge either eye whilst keeping the other still: he "had been able to do this all his life and ... it was just the same thing to him as moving the hand". Yet this is by no means a difficult feat, if—like him—one is allowed the use of a fixation point. I myself can hold my left eye fixed, while moving my right eye at will through a range of about 25° inward to 8° outward from parallelity. Again, it is commonly said that one cannot voluntarily diverge the eyes in the absence of a suitable visual target: yet I can do it with ease, and so can quite a number of my students. The voluntary control of vergence is probably not a profitable field of enquiry.

Vergence movements are of course normally performed unconsciously in response to visual objects that are closer or farther away. Fixation of a near object induces not only convergence, but also accommodation of the lens and constriction of the pupil: these three together constitute the *near response* (triple response or near reaction). Now there are many ways in which the nearness of an object can be recognised: these include binocular *retinal disparity*, the degree of *accommodation* required for optimum focusing, and a number of factors such as perspective, size of known objects, overlap, etc, that can be classed together as '*psychic*' or high-level cues (Maddox, 1907). There is ample evidence that the first two sources of information are capable of generating large vergence movements on their own. The higher-level cues have not received as much attention, although it is clear that they can contribute both to vergence and to accommodation

[for example see Hofstetter (1942), Ittelson and Ames (1950), Alpern (1958), and Erkelens and Regan (1986); Landolt suggested in 1886 that the eye strain frequently resulting from the use of optical instruments like microscopes, especially by inexperienced users, might be due to inappropriate vergence and accommodation responses caused by the subject's awareness of the proximity of the eyepiece]. The next two sections consider the relative contributions of disparity and accommodative information in evoking horizontal vergence.

5.1.1 Disparity vergence

Suppose that a subject uses both eyes to fixate a point target (A, figure 5.2) at a certain distance, so that the image of A falls on the fovea of each eye. Then the two images of a target such as B, lying at a different distance, will in general fall outside the fovea, at locations on the two retinæ that do not correspond with one another (b and b'). Such points are called *disparate*, and the angular distance between them (imagining one retina laid exactly upon the other) is the *retinal disparity*. It is a simple matter to show that disparity information alone can generate compensatory vergence movements, such as to reduce the degree of disparity. If, while a subject is binocularly fixating a visual target, a weak horizontal prism is suddenly introduced before one eye, the disparity which results from the shift of the image of the target in that eye leads to a vergence movement of which the subject is generally unaware (figure 5.3). Clearly there is no stimulus to accommodation in this situation. Similarly, in the stereoscope—where different parts of the binocular scene have different disparities, to give an illusion of depth—the observer performs vergence movements while shifting attention from one part of the field to another. Again, the actual distance of the visual plane remains unchanged, and so accommodation can play no part in evoking the movements. On the other hand, the degree of accommodation is itself influenced by pure disparity information of this sort (Fry, 1937; 1939), as may also be pupil size (Knoll, 1949; Marg and Morgan, 1949; 1950); these interactions are considered in more detail in section 5.1.2.




Figure 5.2. Retinal disparity as a consequence of different target distances.

An important point that is not always fully appreciated is that the retinal disparity of a target that is seen as fused is not necessarily zero, since in addition to the work done by vergence in bringing disparate images together, there is a *central* mechanism of fusion that is not well understood but is undoubtedly related to the existence in the visual cortex of neurons that are wired up to noncorresponding areas of the two retinæ, and are thus preferentially excited by stimuli having particular disparities (Blakemore, 1970; see also chapter 9). The resulting residual disparity when viewing a target is called the *fixational disparity*, and in man may be as large as 2° (Martens and Ogle, 1959; Erkelens and Collewijn, 1985). Consequently there is an area, called *Panum's fusional area* (Panum, 1858) surrounding the point of binocular fixation, inside which point targets are seen as single even though their retinal images are disparate; it may be as large as 20° in extent if measured with binocularly stabilised images (Fender and Julesz, 1967). It follows that the range of disparities over which a stimulus can be fused may be considerably greater than the corresponding vergence movements, and in primates may extend from some 6Δ base in (that is, divergent) to 22Δ base out: bigger stimuli generate larger fusional ranges, both through increased vergence and also increased central fusion (Boltz and Harwerth, 1979; Kertesz, 1981).

In natural visual surroundings, just as in the stereoscope, different parts of the visual field will have different degrees of binocular disparity: and we would certainly expect to find—as in the case of smooth pursuit—some mechanism for selecting one area rather than another as an input to the vergence system. A simple way to do this would be if only the fovea were capable of eliciting disparity vergence, so that one would automatically converge on what one was looking at. Ludvigh and McKinnon (1966; see also Ludvigh et al, 1965) showed that disparity in the fovea was indeed more effective than disparity in the rest of the visual field, although extensive objects in the periphery could override the fovea if of sufficient contrast and detail. There is certainly also a purely voluntary attention component, as the reader can demonstrate for himself. Fixate a point on a distant wall, and hold a pencil up near the face a few degrees to the left of




Figure 5.3. Prism vergence: introducing a weak prism in front of one eye when both are binocularly fixating necessitates unilateral vergence to bring the two retinal images back into correspondence.

the line of sight of the left eye: with a little practice the binocular images of the pencil can be fused *without* having to fixate it, and with the gaze still in the original direction.

One question that is of some interest from the point of view of elucidating the sensory mechanisms of disparity vergence is how similar the retinal images in the two eyes must be to provoke compensatory fusional movements. It is clear for example that slight optical blur of one of the images does not matter much, since subjects with an uncorrected refractive error on one side can nevertheless perform such movements quite adequately. Westheimer and Mitchell (1969) made a systematic investigation of the tolerance of the fusional mechanism to such differences between the two retinal images, presenting the images in brief flashes so that movements would not result in actual fusion: what was measured was the initial vergence response of the system. They showed that the mechanism was startlingly indifferent to a lack of similarity between the two images (figure 5.4). Such grossly dissimilar targets as a vertical and a horizontal line, or a circle and a cross, could still initiate vergence, as could short disparate lines separated in the *vertical* direction by as much as 4° . Neither dimming of one of the figures by a factor of thirty, nor even reversal of its light and dark areas, prevented the correct vergence responses from being elicited. The conclusion must be that disparity information is tapped off the visual pathway at quite an early stage, before any coding of information into channels coding for shape. A distinction is sometimes made between *disparity* vergence, in which dissimilar targets, briefly exposed, may initiate vergence movements but not complete them, and *fusional* vergence, which requires the long exposure of very similar targets, and completes and maintains the correspondence of the retinal images (Milés, 1985), though whether this is a very clear-cut distinction remains to be shown.

In any case, there are limits to the system's tolerance, and in clinical practice a Maddox rod (Maddox, 1907) can be introduced in front of one eye to disable the disparity vergence system. This consists of one or more




Figure 5.4. Examples of pairs of dissimilar binocular targets which when presented one to each eye can nevertheless initiate vergence movements (after Westheimer and Mitchell, 1969).

small red horizontal transparent cylinders arranged side-by-side in a frame: a distant point source of light viewed through it is drawn out into a vertical red streak which is sufficiently unlike the source itself as not to evoke vergence. Under these circumstances, or indeed if one eye is covered or if the binocular scene is completely blank, the eyes assume the 'physiological position of rest', in which their visual axes may or may not be parallel. If they are not, the condition is described as *heterophoria*, a state of tonic convergence or divergence in the absence of fusible stimuli (in the presence of such stimuli it is called *tropia* or *squint*). Clinical aspects of phoria lie beyond the scope of the present account, and the interested reader should consult a standard textbook of ophthalmology (for example Duke-Elder and Wybar, 1973; Leigh and Zee, 1983; Reading, 1983; Pickwell, 1984; Walsh, 1985) for further information. To the physiologist, phoria is of interest because it may reveal long-term adaptational changes in the vergence control systems: as such it is considered in chapter 13.

Although horizontal disparity vergence movements are by far the most important, since the disparities that generate them result naturally from the relative positions of the two eyes in the head, under artificial conditions other kinds of vergence may also be demonstrated. If a weak prism is put before the eye in a suitable position to provoke ordinary horizontal vergence, and then slowly turned through a right angle so that its base is horizontal, fusion can still be maintained: the vertical disparity has produced a state of *vertical* disjunction (Helmholtz, 1867). Vertical fusional ranges as large as 6° have been reported in man (Duwaer, 1982b), most of which can be accounted for by vertical vergence (Perlmutter and Kertesz, 1978; Kertesz et al, 1983); in rhesus monkeys the fusional range for vertical disparity is some $2-4^\circ$ (Boltz et al, 1980). *Diagonal* disparities demonstrate nonlinear interactions between the vertical and horizontal components (Boman and Kertesz, 1983).

One might hope in the same sort of way to be able to generate *torsional* vergence movements, by allowing one eye to view a scene that is twisted relative to the other about an anterior-posterior axis. Such a situation is not as unnatural as it sounds, since the retinal images of any line that forms an angle with the frontal plane will show rotational disparity of this kind. In addition, a consequence of the torsional movements made in correspondence with Listing's law when viewing objects in extreme tertiary positions (although it is doubtful whether the law is actually strictly obeyed in these circumstances: see chapter 7) is that a state of torsional disparity must exist between the two retinal images. Although it is certainly true that an observer can learn to fuse pairs of images having rotational disparity (by *cyclofusion*: see Ogle and Ellerbrock, 1946; Crone and Everhard-Halm, 1975), it is only recently that it has been possible to make sufficiently accurate objective measurements of the corresponding eye movements. Although an earlier report had suggested that fusion was entirely central (Kertesz and Jones, 1970), it is now clear that torsional vergence does occur in response to such

stimuli, but its amplitude depends markedly on the size of the target: 5° disparity of a 50° stimulus can be fused, about half of the fusion being the result of cyclotorsion (Kertesz and Sullivan, 1978). Cyclotorsion is also seen as an accompaniment of the near reflex, together with displacement of the eye (A W Volkmann, in Helmholtz, 1867; Allen and Carter, 1967; Enright, 1980): these second-order effects are discussed in section 7.3, in the context of ocular kinematics.

5.1.2 Accommodation vergence

In a classical experiment, Johannes Müller showed that an accommodative stimulus—retinal blur—can produce vergence movements in the absence of disparity (Müller, 1826). If one eye is covered while the other fixates a target, and then a negative lens is suddenly introduced before the seeing eye, the other responds by converging (figure 5.5). The only possible origin of the vergence in this experiment is the accommodative effort made in refocusing the target. Many experimenters have investigated the relationship between the magnitude of the accommodative stimulus and that of the resultant accommodation convergence: the ratio between these last two quantities is called the AC/A ratio, and may be measured in prism diopters per diopter; it actually has the dimensions of length, and might be better expressed in centimetres. One cannot of course assume that the actual degree of accommodation is necessarily exactly what is required to cancel the negative lens, and strictly speaking one should also distinguish between stimulus AC/A and response AC/A; but in practice, the differences are usually small, of the order of 8% (Alpern, 1969b). For fairly small accommodation stimuli (less than some 5 diopters), vergence in most




Figure 5.5. Asymmetrical convergence: above, under normal binocular viewing conditions, movement of the fixation point along the line of sight of one eye results in a mixed conjunct and disjunct movement of both. Below, if the eye which has to move is covered, then it alone moves in a simple monotonic manner: this is pure accommodation convergence (data from Alpern and Ellen, 1956; Alpern, 1957).

subjects is effectively a linear function of accommodation (figure 5.6) (Flom, 1960, part 1) and the AC/A ratio is constant at around 3–4 Δ per diopter, with considerable variation between subjects (figure 5.7) (Morgan, 1944; Emmes, 1949; Morgan and Peters, 1951; Flom, 1960, part 2).

Now it is easy to verify by simple geometry that a perfect accommodation vergence response (that is, one that generated exactly the amount of vergence corresponding to the apparent movement of the fixation point) would show an AC/A ratio equal to the interocular distance in centimetres for relatively distant targets. Since this distance in adults is around 6 cm (Hofstetter, 1972), it is clear that the accommodation convergence component in fixating near objects is substantial, accounting for rather more than half the total response. Indeed, it was once held (Maddox, 1886) that vergence was primarily controlled by accommodation, and merely supplemented by disparity. If this were so, one would expect little divergence of an eye suddenly covered during binocular viewing of a near stimulus: this is not what is observed (Semmlow and Heerema, 1979). However, it is in any case misleading to think of vergence as being simply the linear sum of two independent stimuli, because of the complex feedback relationships that exist between vergence and accommodation on the one hand, and disparity and blur on the other. These relationships appear to be perfectly symmetrical, in that changes in disparity elicit not only vergence but also accommodation: one may measure CA/C ratios as well as AC/A. In the words of Porterfield (1759):

“... it is evident, that there is a necessary Connection and Dependence established between those Motions, whereby the Conformation of our Eyes is changed, and certain corresponding Motions in our Axes of Vision, which make it impossible for us to direct our Eyes to any Object within the limits of distinct Vision, without at the same time giving them that Disposition that is necessary for seeing distinctly at that Distance ...”

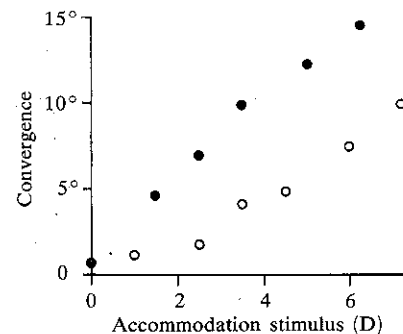


Figure 5.6. The relationship between accommodation and convergence in two subjects, showing the effectively linear relationship in each case, with individual differences in slope (data from Flom, 1960, part 1).




Figure 5.7. A measured distribution of AC/A ratios amongst 104 subjects (after Ogle and Martens, 1957).

Thus the most general form of the vergence control system must be something like figure 5.8, in which the parameters α , β , γ , and δ may be taken to represent the relative strengths of the pathways respectively linking blur to accommodation, disparity to vergence, blur to vergence, and disparity to accommodation. [Similar representations have been published by Krishnan et al (1977), Judge (1985), and Schor and Kotulak (1986), amongst others. The question of the *dynamic* properties of these pathways is left until the next section.]

If we represent the accommodative stimulus (in diopters) by S , target disparity (in prism diopters) by U , resultant accommodation by R , and vergence by V , then the system can be written as:

$$R = \alpha(S - R) + \delta(U - V),$$

$$V = \beta(U - V) + \gamma(S - R).$$

Then

$$\mu R = (\mu - 1 - \beta)S + \delta U,$$

$$\mu V = (\mu - 1 - \alpha)U + \gamma S,$$

where $\mu = (1 + \alpha)(1 + \beta) - \gamma\delta$.

With one eye covered—and a few assumptions!—one can take $\beta = \delta = 0$, so $\mu = 1 + \alpha$. Then

$$R(1 + \alpha) = S\alpha, \text{ and } V(1 + \alpha) = S\gamma,$$

so that the stimulus AC/A ratio (V/S) is $\gamma/(1 + \alpha)$, and the response AC/A ratio (V/R) is γ/α . One may similarly calculate fixation disparity ($U - V$) as:

$$\frac{1}{\mu} [(1 + \alpha)U - \gamma S].$$

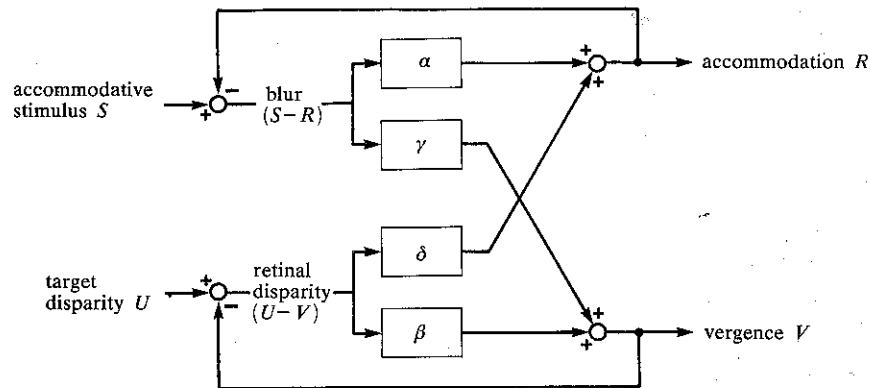


Figure 5.8. The general relationship between accommodative stimulus (S), target disparity (U), accommodation (R), and vergence (V): $R = \alpha(S - R) + \delta(U - V)$; $V = \beta(U - V) + \gamma(S - R)$.

It follows that the ratio of the powers of the prism and lens that produce the same degree of fixational disparity is also a measure of stimulus AC/A (Judge, 1985). One would expect from the symmetry of all these relationships that steady-state errors in accommodation would be found, equivalent to fixational disparity: in practice, errors of the order of 0.5 diopters or more are indeed found (Alpern and David, 1958).

5.2 Dynamics of vergence movements

5.2.1 Natural viewing

Vergence movements are necessarily rather small (in fixating an object that moves from one metre away to two, each eye moves through less than a degree), so that quantitative studies of the control of vergence movements have had to wait for the development of sensitive and accurate methods of measuring small eye movements (see for example Riggs and Niehl, 1960; Rashbass and Westheimer, 1961a).

Disparity vergence control can best be investigated by arranging for each eye to view separate and independently moveable visual targets: in this way, visual stimuli can be presented whose disparity can be varied at will. In response to a step change in disparity, after a reaction time of some 160 ms the eyes move smoothly and comparatively slowly to their final positions, in a roughly exponential manner (figure 5.9) (Westheimer and Mitchell, 1956; Riggs and Niehl, 1960; Rashbass and Westheimer, 1961a): the whole movement may take nearly a second to complete. The latency of accommodation vergence is similar (some 150–200 ms): the time course is also comparable but possibly somewhat slower, being typically only 90% complete after 1.2 s (Alpern and Ellen, 1956; Robinson, 1966; Cumming and Judge, 1983). By analysis of covariance between random fluctuations of accommodation and of vergence, it is possible to show that most of this delay arises in pathways that are common to both phenomena, that is, essentially in the visual system itself (Wilson, 1973a); by contrast, the latency of blur accommodation is considerably longer—some 370 ms—and would suggest an additional peripheral delay (Campbell and Westheimer, 1960) were it not that disparity accommodation appears to have a latency that is similar to vergence (Cumming and Judge, 1983).




Figure 5.9. Vergence response (recorded from one eye) to a sudden step of disparity applied at the moment indicated by the arrow: closed-loop conditions (Rashbass and Westheimer, 1961a).

As is evident in figure 5.8, the disparity vergence control system is essentially a feedback one. The relation between vergence V and disparity d ($= U - V$) is:

$$V = d \left(\beta - \frac{\gamma}{1 + \alpha} \delta \right) + S \frac{\gamma}{1 + \alpha}.$$

It is not difficult to show that the control of vergence by disparity is indeed guided and not ballistic, by using pulses rather than steps of disparity. In chapter 4 we saw that the saccade control system responded to pulses of displacement by generating either a full-sized saccade or none at all: the response of the vergence system is quite different (figure 5.10). One reaction time after the beginning of the pulse, the eyes begin their vergence movement; and one reaction time after the end of the pulse, the eyes start to return to their original positions. Thus it is clear that these movements are not ballistic, and also do not exhibit the refractory behaviour characteristic of saccades. If the disparity is made to vary sinusoidally, the resulting vergence is found to match it quite accurately so long as the peak velocity of the stimulus is not too great. Adequate matching continues until the stimulus frequency increases to around 1 Hz, when the response begins to break down: the gain is 3dB down at 1.2 Hz, and is reduced to less than 1% before 4 Hz is reached (Richards, 1972; Cumming and Judge, 1983). Recently, Erkelens and Collewijn (1985) have made accurate measurements of phase and gain not with small fusional targets but with large-field random-dot stereograms, moved in or out of phase with one another, or with one stationary and the other moving: some of their results are shown in figure 5.11. Another point of similarity is in their reaction to barbiturates: with even quite moderate doses of sodium amytal, both accommodation and disparity vergence responses are considerably slowed (Westheimer and Rashbass, 1961; Westheimer, 1963); smooth movements behave similarly (Rashbass, 1959).




Figure 5.10. (a) Demonstration of the ballistic behaviour of saccades when a subject is presented with a short pulse of target movement; (b) simple additivity of vergence responses to a short pulse of disparity. The total response is in each case approximately the sum of the responses to the rising and falling edges of the disparity pulse, and there is no refractory period.




Figure 5.11. Eye movements made while viewing a random-dot stereo pair, the two halves moving laterally in counterphase with a sinusoidal time course (0.25 Hz). From top to bottom: left eye, right eye, vergence (L-R), version ($\frac{1}{2}L + \frac{1}{2}R$); the continuous line shows eye movement, the dotted line, target motion (after Erkelens and Collewijn, 1985).

5.2.2 Open-loop studies

Although disparity vergence embodies intrinsic feedback, in the sense that any vergence movement made immediately alters the disparity stimulus, accommodation vergence on the other hand is intrinsically open-loop in nature; but the interaction between accommodation and disparity vergence discussed in section 5.1 implies that each of the two systems must partake of some of the properties of the other. For example, a purely accommodative binocular stimulus—we might suddenly introduce a negative lens in front of each eye—induces a vergence movement: but this is immediately opposed by the disparity vergence system, which in turn, as we have seen, may generate an associated accommodation response. Even in the classical demonstration of accommodation vergence, in which one eye is covered, one cannot be sure to what extent the vergence response itself modifies the accommodation effort and thus induces second-order effects. Indeed it is not altogether clear whether the accommodation vergence is more closely related to the actual accommodation of the eye, or only to that *part* of the accommodation that is due to defocusing of the retinal image rather than other potential sources such as disparity. At all events, it is clearly desirable to remove at least some of this obscurity by 'opening the loop' and arranging for the disparity stimulus, at least, to be independent of any vergence movements that may be made: under these circumstances the gain, V/U , will be given by $1 - (1 + \alpha)/\mu$. More information about the crossed pathways linking accommodation to vergence and vice-versa can be obtained by measuring the transfer functions of response AC/A and CA/C , either under open-loop or under closed-loop conditions (figure 5.12) (Schor and Kotulak, 1986).

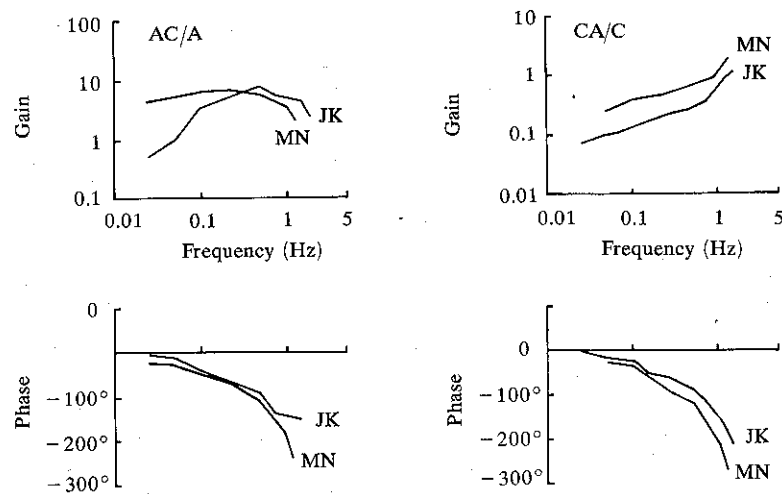


Figure 5.12. Crossed dynamics in the vergence/accommodation system. Gain and phase for two subjects under closed-loop conditions: left, of response AC/A; right, of response CA/C (data from Schor and Kotulak, 1986).




Figure 5.13. A method of opening the feedback loop in the control of disparity vergence. The movement of each eye is measured separately: the difference between these signals (proportional to the vergence) is added to the applied disparity stimulus and used to generate reciprocal movements of the spots on two cathode-ray tube screens such that the subject's vergence responses are cancelled out. Disparities may thus be applied and 'clamped' at any desired value.

If the disparity stimuli are generated by means of moveable spots on the faces of a pair of cathode-ray tubes, with suitable arrangements to ensure that each eye sees only its own spot, one can arrange for an electrical signal proportional to the difference in angular position between the two eyes—that is, proportional to the vergence—to be continuously added to whatever input disparity signal is used (figure 5.13) (Rashbass and Westheimer, 1961a; Zuber and Stark, 1968). If the gain of this artificial external feedback loop is correctly adjusted, it cancels out the effect of the intrinsic loop, and the response of the system then corresponds only to the forward part of the control pathway. If a step of disparity is now applied, it is found that after the usual reaction time of about 160 ms the vergence of the eyes starts to increase at a constant rate. Thus the response to a square-wave disparity input is a triangular-wave vergence output (figure 5.14) in which the slope is directly proportional to the amplitude of the input, up to some 4° of disparity (figure 5.15). This suggests, as in the case of the smooth pursuit system, that the input signal is being integrated with respect to time to generate the vergence response. If this were true, we would expect that the amplitude-frequency response for sinusoidal disparity inputs under open-loop conditions would also be characteristic of an integrator, namely a straight line with a slope of -1 in logarithmic coordinates. This is found to be the case, at least for a rather moderate range of frequencies (figure 5.16), but the phase response is more puzzling. A pure integrator would be expected to show a phase lag of 90° at all frequencies, while the observed reaction time of 160 ms would add to this a component increasing with frequency, amounting to an extra 90° of lag at about 1.5 Hz.

The actual phase lags that are observed do not correspond with this expectation at all (figure 5.16), and are consistently smaller than would be anticipated from the amplitude response, assuming the system to be minimum-phase apart from the constant 160 ms delay. Now we saw in the case of the smooth pursuit system that such unusually small phase lags can be associated with adaptive behaviour to repetitive stimuli, whereby the




Figure 5.14. Open-loop vergence response (upper trace) to applied square-wave modulation of disparity (lower trace). The integrator-like response of the eye is clearly seen (Rashbass and Westheimer, 1961a).




Figure 5.15. Open-loop disparity vergence: the relation between disparity and rate of change of vergence for different applied disparities: the relation is evidently a linear one until disparities of around 5° are reached (data from Rashbass and Westheimer, 1961a).

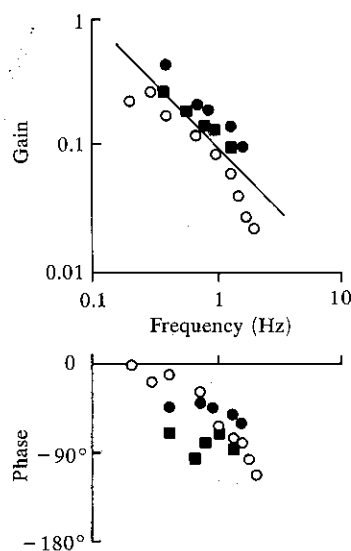


Figure 5.16. Frequency response of open-loop vergence responses to sinusoidal variation in applied disparity: \circ Zuber and Stark (1968); \bullet and \blacksquare , two subjects from Rashbass and Westheimer (1961a). The line in the gain plot shows what would be the expected response from a simple integrator.

system effectively learns to predict where the stimulus is likely to be. Under closed-loop (natural) conditions, reduced phase lags to repetitive stimuli are in fact observed, and system gain is greatly reduced if less predictable stimuli are used (figure 5.17) (Zuber and Stark, 1968; Cumming and Judge, 1983; Schor and Kotulak, 1986).

Even under open-loop conditions, prediction of a rather different kind can occur: the system often exhibits an unexpectedly short response time to some brief stimuli that are presented only once, as for example the ramp-step combination of figure 5.18 (Rashbass and Westheimer, 1961a). The vergence movement is of the parabolic form expected as a result of integrating the ramp: but the turning point of the parabola, which would not be expected until some 160 ms after the moment when the stimulus passes through zero disparity, actually *anticipates* this moment by at least 100 ms. One possibility is that information about *rate* of change of disparity, as well as disparity itself, is used to predict the future degree of vergence required. Ludvigh and McKinnon (1968) did indeed find that the rate of change of disparity under natural (closed-loop) conditions influenced the size of vergence movements that could be made: but it is difficult to see why if such a component exists it does not show up in the observed amplitude frequency responses.




Figure 5.17. Frequency response of closed-loop vergence responses to sinusoidal variation in applied disparity [data from: \bullet Yoshida and Watanabe, 1969; \blacksquare Erkelens and Collewyn, 1985 (5° amplitude); $+$ Zuber and Stark, 1968]. The open circles show Zuber and Stark's measurements, using summed sinusoidal components, for unpredictable target motion.




Figure 5.18. Anticipatory vergence response to ramp-step of disparity: see text.

Another problem with trying to represent the disparity vergence system in linear form is that the closed-loop responses to steps of disparity are markedly nonlinear: convergence is considerably more rapid than divergence (figure 5.19) (Zuber and Stark, 1968). Again, there appears to be a discrepancy here between closed-loop and open-loop behaviour, for the latter shows no nonlinearities of this magnitude (see figure 5.16). Conceivably the more complicated behaviour of the system under closed-loop conditions reflects the intrusion of linked accommodation responses in the manner suggested earlier. Another complicating factor is the influence of higher visual cognitive functions on the fusional response: it is a common experience when viewing difficult stereograms that fusion may be hard to achieve at first, but gets progressively easier over the course of a minute or so, as the various depth relationships in the figure are recognised and understood (see for example Helmholtz, 1867). More generally, it is clear that the distinction between disparity and fusional vergence mentioned earlier extends also to the time course of vergence responses: a short period of convergence results in fast relaxation on covering one eye, whereas a longer period—more than half a minute—induces incomplete relaxation (Schor, 1979b). Though it is certainly possible to encompass such observations by stuffing one's model of the system with extra integrators (figure 5.20) (Schor and Kotulak, 1986), they may also be regarded, like other kinds of phoria, as long-term adaptational phenomena: as such, they are considered in chapter 13.




Figure 5.19. Differences in the time course of convergence and divergence responses of the same amplitude (after Zuber and Stark, 1968).

One curious aspect of natural disparity vergence that invites speculation is its *slowness*: not the latency of 160 ms or so, which is after all slightly better than what is achieved by the saccadic system, but the leisurely fashion in which the movement gradually grinds to a halt over a period of a second or more. A possible explanation comes to mind from consideration of the

stability of the closed-loop system. The slowness of the response is simply a function of the gain of the open-loop integrator, whose transfer function we may write as kD^{-1} . If k is large, the system will respond to small errors of disparity with a large corrective vergence velocity that will smartly restore the error to zero: if speed were the only consideration, k should be made as large as possible. But if there is also some intrinsic delay round the loop, then if k is too large the system will overreact, and there will be no time for feedback to moderate its response: the result will be an overshoot, and a subsequent 'hunting' oscillation as the eyes swing back and forth about the correct degree of vergence.

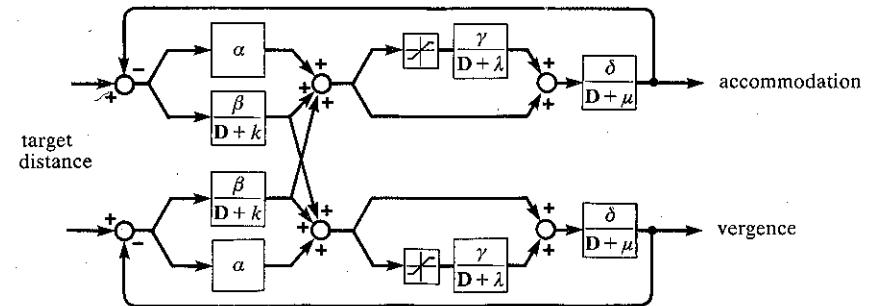


Figure 5.20. Dynamic components of the complete accommodation/convergence system, according to Schor and Kotulak (1986).

It is a relatively straightforward matter to calculate, for a particular loop delay, how large k may safely be (see section 2.8 of appendix 2): on the assumption of pure integration in the forward pathway, k must have a value of less than about 10 s^{-1} to prevent oscillation of this kind, and nearer half this figure if transient instabilities are to be avoided. The value of k in man can be deduced from the results shown in figure 5.15: it is somewhere between 4 and 5 s^{-1} . In other words, the system is about as fast as it can be if it is not to be unstable. However, the fact that the measured phase delays are considerably shorter than would be expected from a simple integrator with a delay in series suggests that the calculated maximum value for k is on the conservative side, and it would be interesting to determine by direct experiment (by appropriate manipulation of the external feedback loop in an arrangement such as Rashbass and Westheimer's) how far the open-loop gain may in fact be increased with safety. Of course, one is still entitled to ask why the vergence system chooses to operate as a feedback, guided, system at all, and not go for ballistic control as the saccadic system has done. The answer may be that the vergence system is still young, and has not yet had time to evolve the rather more sophisticated mechanisms needed to achieve ballistic control: possibly its rather puzzling predictive properties are the fruits of its first tentative steps in that direction.

The dynamics of accommodation vergence do not appear to be markedly different from those of disparity vergence. The vergence response to sinusoidal modulation of an accommodative stimulus is broadly similar to that of sinusoidal disparity vergence (figure 5.21; cf figure 5.12), although there is some disagreement as to the magnitude of the system gain. Like disparity vergence, accommodation vergence also demonstrates a certain degree of predictive behaviour (Krishnan et al, 1973; Cumming and Judge, 1983).




Figure 5.21. Frequency responses of vergence response to sinusoidal variations of accommodation (data from: ■ Troelstra et al, 1963; ● Yoshida and Watanabe, 1969).

5.3 Mixed vergence and version

The natural movements made in looking successively at different objects in space normally of course involve components both of vergence and of version, and it is reasonable to ask whether the eye control system makes any attempt to coordinate their separate requirements, or whether on the other hand the two systems issue independent commands that merely summate in the final common path. A critical experiment that helps to answer this question is that of unilateral prism vergence (Alpern, 1957), performed without the other eye being covered (figure 5.3). Only the eye that receives the prism needs to move: and one might expect that the other eye would remain motionless throughout the operation, maintaining steady fixation on the target. This is not in fact what happens at all. After the usual 160 ms or so, both eyes participate in a symmetrical vergence movement: but shortly after it has begun, a saccade occurs that brings their mean position of gaze in line with the target, the vergence movement meanwhile

proceeding to completion. The same pattern of movements—a slow vergence movement with a saccade superimposed in the middle—is seen under more natural conditions when a subject's gaze is shifted between objects situated at different distances and in different directions (figure 5.22) (Yarbus, 1957a; 1967).




Figure 5.22. The sequence of events in a mixed vergence and version movement. The thick line on the left in each diagram shows the locus traced out by the point of fixation; on the right of each diagram, the time course of the movement of each eye is plotted separately.

If Hering's Law were being obeyed in such cases, then during the saccade the two eyes should move equally in the same direction; and in the vergence component, by equal amounts in *opposite* directions. However, the contribution of each eye to vergence movements is, in general, not *exactly* equal, but depends on the relative dominance of the two eyes; the conjunctive component in asymmetric vergence of this kind tends to bring the eye to a point lying on the line joining the target to not the midpoint of the baseline, but to the *binoculus*, or subjective sighting centre, which tends to lie to the side of the more dominant eye (figure 5.23) (Pickwell, 1972a; 1972b). Indeed, with small displacements of the target, most movements are carried out without using saccades at all (Ono and Tam, 1981), and the probability




Figure 5.23. Mixed vergence and version with unequal contributions of the two eyes. On the left, the fixation locus (as in figure 5.22) when the eyes contribute equally to vergence: on the right, when the left eye is dominant and contributes less to vergence. The asterisk shows the position of the binoculus in each case.

of unequal vergence movements depends in general on the geometrical relations between the target lights and the eyes (Ono and Nakamizo, 1977; 1978). These inequalities need not in themselves imply a violation of Hering's principle during the *saccadic* component, but careful observations nevertheless show that the saccades made by each eye are not in general equal in size, and that a substantial part of the entire change in vergence required by the target is actually achieved during the saccade, very much more than can be accounted for on the assumption of simple additivity between the saccade and vergence (Enright, 1984a; 1986).

It appears from these observations that there is a certain degree of cooperation between the two systems, but not to the extent of avoiding unnecessary motion altogether. Of course, the most blatant violation of Hering's Law is Müller's classical demonstration of accommodation vergence, with one eye covered: under these circumstances the seeing eye remains stationary and only the occluded eye moves (Alpern and Ellen, 1956), suggesting that the normal relationship depends in some way on the presence of natural visual feedback. Indeed, Kenyon et al (1978) describe small movements of the viewing eye under these conditions, and suggest that they represent attempts at vergence and saccades which are thwarted by optokinesis acting only on the viewing eye. Evidence that retinal slip in one eye can be used to generate asymmetrical movements of this kind comes from an extension of the experiments of Erkelens and Collewijn (1985) mentioned earlier, in which *nonfusible* pairs of random-dot images were presented to the two eyes: at low frequencies, retinal slip acting in opposite directions in each eye could give rise to vergence, with a gain of up to 20%. Conceivably this is a mechanism driven by retinal errors of position rather than velocity, and thus the same as slow correction ('Corrective drift'—see section 6.3.2); it would also be expected to contribute to the violations of Hering's Law in asymmetrical vergence.

A strict division of eye movements into those that are purely conjunct and those that are purely disjunct leads to the notion of an 'oculomotor map' of visual space, divided by lines of equal version (in the sense that any point on such a line may be reached from any other point on it by means of a pure vergence movement) and lines of equal vergence. The latter—which could be called *isophores*—correspond to the Vieth-Müller circles (see Shipley and Rawlings, 1970a; 1970b) that link lines of equal disparity when the visual axes are parallel and horizontal: in a horizontal plane including the baseline, they consist of a series of circles passing through the centres of rotation of each eye (figure 5.24). The lines of equal version—which one might call *isotropes*—form a series of rectangular hyperbolae whose centre is the midpoint of the baseline and which pass through the two centres of rotation. Similar hyperbolae were introduced by Hillebrand (1893) to describe the perceived geometry of visual space. Any movement of the eyes from one point to another on such a map can be resolved into its versional and vergence components along these coordinates. That a similar system of

coordinates is also appropriate for perceptual space has been convincingly argued by several authors (for example Luneburg, 1948; 1950; Leibovic et al, 1971), and no doubt stems from its utility in the calculation of three-dimensional eye movements.




Figure 5.24. A horizontal section through a system of space coordinates based on lines of equal vergence and lines of equal version. The two dots are the notional centres of rotation; the circular arcs are *isophores* (lines of equal vergence), and the rectangular hyperbolae are *isotropes* (lines of equal mean version) (after Luneburg, 1948).

Miniature movements

"... but I assure you... the eye standeth not still but moveth incessantly."

Each of the kinds of movement considered so far confers some obvious benefit on the visual system, either serving to shift the gaze from one point to another, or helping to keep the retinal image relatively still. But normal subjects make a number of other, less prominent, types of eye movement whose usefulness is not so obvious. These movements are an order of magnitude smaller than anything that has been considered so far, and can be observed best when a subject is attempting to hold his fixation on a stationary object (figure 6.1): they can be described as *micro-* or *miniature* movements, or as movements of fixation. Because of their small amplitude, exact knowledge of their characteristics had to wait for the development of recording techniques exquisitely sensitive to eye movement, but unresponsive to head movement; Barlow (1952) gives a good idea of the heroic methods that may be necessary to achieve this aim.




Figure 6.1. Verheijen's demonstration of the micromovements of fixation (Verheijen, 1961). The reader should first fixate the centre of the large white dot, and then after a time sufficient to get a good afterimage, some 20 s, fixate the small black dot. Small displacements of the afterimage relative to the figure itself are then made strikingly obvious, and it should be possible to see the slow drifting movements of the eye as well as the microsaccades that tend to refixate it.

6.1 General characteristics of miniature eye movements

General accounts of miniature eye movements have been given by, among others, Adler and Fliegelman (1934), Lord and Wright (1948), Ratliff and Riggs (1950), Barlow (1952), Ditchburn and Ginsborg (1953), Nachmias

(1959), Yarbus (1967), Ditchburn (1973) and Steinman et al (1982). It is generally agreed that in man there are essentially three different components of the miniature eye movements: *tremor*, *drift*, and *microsaccades* (figure 6.2). Ratliff and Riggs (1950) have also described slow random movements in a frequency range of about 2–5 Hz and with amplitudes of the order 1–5', which they consider to be separable from slow drift, but this distinction is not generally recognised.

The combined effect of tremor, microsaccades, and drift is to move the retinal image about, irregularly and incessantly, so that it traces out a path on the retina whose dimensions are considerably larger than that of the foveal receptors (figure 6.3). This can hardly pass unnoticed by the visual system; the somewhat controversial question as to whether these movements do or do not help us to see is deferred until chapter 11.

There is a considerable degree of variation in these movements between species, no doubt reflecting differences in the importance of foveal vision. Rhesus monkey is not very different from man (Skavenski et al, 1975; Motter and Poggio, 1984). In the cat, ocular stability is probably about as good as ours, but is maintained solely by slow control (corrective drift, see below) rather than by microsaccades; drift velocities are on average perhaps twice as large (Winterson and Robinson, 1975). [Earlier reports (Hebbard and Marg, 1960; Pritchard and Heron, 1960) of drift movements of much larger amplitudes were probably caused by periods in which the animal




Figure 6.2. Simultaneous records of the miniature movements of the two eyes. The small-amplitude high-frequency component is the *tremor*, the large and relatively slow excursions are *drift*, and at the arrow both eyes execute a *microsaccade*. It can be seen that drift movements are essentially dissociated in the two eyes, whereas the microsaccade is virtually exactly conjugate (after Yarbus, 1967).

Figure 6.3. Schematic representation of the kind of two-dimensional motion of the fixation point consequent on the miniature movements. The large circle has a diameter of 10' (after Pritchard, 1964).

dozed off (figure 6.4.)] A similar absence of fixational microsaccades is shown by rabbits, and in darkness a large and slow drift develops; ocular stability is thus normally maintained by slow correction using visual feedback (Collewijn and van der Mark, 1972; van der Steen and Collewijn, 1984). The miniature eye movements of birds show some striking differences. Tremor is undetectable in the pigeon, but both here and in the owl one can record short bursts of oscillation, with frequencies in the region of 20–30 Hz, and with comparatively large amplitudes, in the region of 2° or more (Nye, 1969; Steinbach and Money, 1973); it is not known what initiates these bursts, or what function they may have.




Figure 6.4. Left, record of horizontal and vertical eye position in a cat. Initially, the animal was drowsy, and showed large drift movements. At the arrow, it was aroused by shouting: subsequently, the stability of eye position was maintained despite the absence of microsaccades. (The bar marks a region in which the amplifier was adjusted.) Right, record of horizontal eye movement in a cat trying to maintain its direction of gaze in the dark (after Winterson and Robinson, 1975).

6.2 Tremor

It can be seen in figure 6.2 that the amplitude of tremor is substantially smaller even than that of the other miniature eye movements, and is in fact of the order of the diameter of the smallest cones (some 24"). Estimates include Ratliff and Riggs's (1950) median of 17.5", Ditchburn's (1955) range of 5–15", and Yarbus's (1967) of 20–40". Smaller estimates are probably nearer the truth, since experimental difficulties are likely to increase rather than decrease the estimated amplitude. Tremor generates a component of retinal image *velocity* that is of the order of $10' s^{-1}$, and thus not very different from what is introduced by drift and microsaccades (Ditchburn 1980). The bandwidth of ocular tremor is some 90 Hz: measurements of the power spectrum show a function that falls essentially monotonically with increasing frequency above 10 Hz, until the level of background noise is reached, typically in the 150–200 Hz region (figure 6.5) (Fender and Nye, 1961; Bengi and Thomas, 1968b; Findlay, 1971; St Cyr, 1973). It may be that there are significant resonance bumps on the curve, particularly in the regions around 40 Hz and 80 Hz (Bengi and Thomas, 1968b), but these are not consistently reported: the possibility that they may be due to resonance in the contact lens attachments used for the measurements, or even in the globe itself, cannot be completely ruled out (Boyce and West, 1968).

St Cyr (1973), discussing the rate of decline of the power spectrum with increasing frequency, distinguishes two possibilities: the falloff may either be due to the intrinsic frequency characteristics of the noise source itself, or it may be the result of filtering a flat (white) noise with a peripheral low-pass filter. He suggests that the observed rate of decline is what would be expected if the noise spectrum were essentially determined by the mechanical properties of the eye (assumed to be dominated by *inertia* at high frequencies). But it is equally compatible with white noise passed through a central integrator and then further filtered by the more commonly accepted *viscosity*-dominated eye mechanics, a notion that is in some ways more attractive (see chapter 12). Wilson (1973b) has similarly suggested that the decreased high-frequency noise in accommodation vergence in comparison with that of accommodation itself may also be the result of differential peripheral filtering.

Binocular measurements of tremor indicate that it is uncorrelated between the two eyes (Riggs and Ratliff, 1951), suggesting a rather peripheral origin. There is no evidence that the amplitude of tremor can be influenced by the visual conditions, or by efforts of the will.




Figure 6.5. Amplitude frequency spectrum of the tremor component of the miniature movements. The two lines show amplitudes that fall off with the reciprocal of the frequency, and with the square of the reciprocal: their slopes are -1 and -2 respectively (data from Findlay, 1971).

6.3 Drift

Drift movements are comparatively large and slow, having velocities in man in the region of $4' s^{-1}$, and median amplitudes of around 2–5' [Ditchburn (1973) has an exhaustive list of previous measurements]. Boyce (1967) has published statistical analyses of the frequency of occurrence of drifts of various extents, showing a roughly exponential decline in frequency with increasing amplitude. Since under normal conditions of fixation in man,

each drift movement is necessarily terminated by a microsaccade, such distributions really reflect the statistical behaviour of the timing of the microsaccades, rather than any property associated with the mechanism underlying drift; for this reason, and also because of its more obvious relevance to vision, it is more helpful to consider the velocities introduced by different kinds of micromovements, rather than their amplitudes.

6.3.1 The randomness of drift in the dark

There are two kinds of random processes that could give rise to drift, that one might call *position* noise and *velocity* noise. In the first case, one can imagine some source of (low-frequency) random disturbance that causes the position of the eye to wander from moment to moment within a probability distribution fixed with respect to the desired point of fixation. In the second, what is envisaged is a process in which it is the velocity rather than the position of the eye that is controlled, and where the amplitude and direction of this velocity are similarly subject to random variation: such a process is technically known as a *random walk*. The fundamental distinction between these two models lies in how the power spectrum of the eye position behaves at very low frequencies: in the first case it tends to a finite limit; in the second, it increases indefinitely with decreasing frequency.

However, the difference between these two models is blurred to some extent by the likelihood that if the eye wanders too far from where it ought to be looking, some other *corrective* process will bring it back. Such a process might be visual or proprioceptive, and it might operate either saccadically or through a slower mechanism akin to smooth pursuit or optokinesis. We shall see later that the microsaccades of fixation do indeed come into this category, and that there is also clear evidence, in the absence of microsaccades, for a slow mechanism of correction. The latter mechanism is often called *corrective drift*, perhaps an unfortunate expression since it confuses correction of the error that the drift has caused with the drift itself: it might be better called *slow correction*. The existence of such a mechanism makes it impossible, simply by examining records of drift, to determine whether it is due essentially to position or velocity noise. However, if it is assumed that such slow correction is visually driven, then in the dark we should be able to record drift in its pure and uncontaminated form; and one consequence of such conditions is that the mean rate of drift should increase. This is indeed seen in the rabbit (Collewyn and van der Mark, 1972), and in man was observed by Ditchburn and Ginsborg (1953) and by Nachmias (1961), but not by Cornsweet (1956). [Fixation movements in the dark seem to be especially idiosyncratic, in respect of both drift and microsaccades, and a subject may even learn to use quite unrelated sensations such as artificial auditory feedback to control his eye position (McLaughlin et al, 1968).] Becker and Klein (1973) report that drift in the dark is on the whole directed to the primary position, even when eccentric fixation is attempted, which suggests a second mechanism of slow correction using

information about eye position obtained from muscle stretch receptors or by monitoring motor commands (these possibilities are further discussed in chapter 10). Lackner and Zabkar (1977) reported that another kind of proprioception—grasping the target with the hand—could limit drift, although Winterson and Steinman (1979) were unable to confirm this.

Whatever the source of the information used to generate slow correction of this kind, observation of long stretches of drift (possible in the dark because of the reduced frequency of microsaccades) should be able to tell us whether in fact such correction is taking place at all. For if the direction and amplitude of drift is wholly random with respect to the position of the desired direction of gaze in the field, the trajectory of the direction of gaze will be of the 'random walk' form: in particular, the mean square of the deviation from the initial position will grow linearly with time. This is not in general true if the gaze is also subject to a correcting influence tending to pull it back to the origin. Actual experiments indicate that in fact the mean square rule is followed quite accurately, and that, although significant deviations may occur, they are relatively small (figure 6.6) (Riggs et al, 1954; Matin et al, 1970; Findlay, 1974). Skavenski and Steinman (1970) similarly found that any corrections in the dark were made by microsaccades and not by drift. The fact that drift in the dark can be well described by a random walk model suggests—as in the case of tremor—that the source of the noise that produces drift lies central to an integrator of some kind.




Figure 6.6. Drift in the dark. The open circles show measurements of median deviation of the eye from the original fixation point for different periods of time in the dark; the triangles show measurements of the accuracy of perceiving the position of a flashed object under similar circumstances, and indicate a distance that could be discriminated correctly 75% of the time. The time axis is such that, if the data followed a 'random walk' model, it would lie on a straight line (data from Findlay, 1974).

6.3.2 Slow visual correction of drift

Investigation of slow visual correction in the light is complicated to some extent by the existence of fixational microsaccades. Many species do not exhibit microsaccades, and ocular stability is almost entirely maintained by visual control of drift that prevents the image of the target from moving too far from the fovea. Under such conditions, and in man when microsaccades are voluntarily suppressed, there is no question but that drift is constrained by slow correction, maintaining ocular stability nearly as well as microsaccades do (see for example Kowler and Steinman, 1980). Recently, de Bie and van den Brink (1984) have shown that this correction drift is driven partly by retinal positional error and partly by its rate of change, retinal slip velocity.

But under conditions in which fixation is maintained by microsaccades, different considerations apply. Each period of drift is necessarily terminated by a microsaccade, and, as we shall see, the microsaccades are on the whole such as to bring the visual target nearer the fovea. It must follow that even if the drift movements were entirely independent of the visual input, we would *still* expect to find a correlation between drift amplitude and direction and the position of the fixation target on the retina. For if on average the endpoint of a microsaccade is nearer the target than was its point of initiation, the converse must necessarily follow for drifts; and on average, this is—as it must be—precisely what is observed (Nachmias, 1959; 1961). But the possibility still remains that particular subclasses of drift movement (for example, in particular directions) may be corrective under such conditions, and others anticorrective: and in fact many experiments (Nachmias, 1959; 1961; Boyce, 1967) have shown that, under particular circumstances, drifts in the compensatory direction may be observed. Thus one of Nachmias's subjects showed compensatory drift near the vertical meridian, but none elsewhere (see for example figure 6.11).

Drifts in the two eyes are apparently uncorrelated (Krauskopf et al, 1960), although it has been suggested that over shorter time spans periods of significant *negative* correlation—amounting to fluctuating waves of vergence—may be observed (Ditchburn and Ginsborg, 1953). However, such periods are only to be expected from time to time on a purely random basis, balanced by a comparable number of periods of positive correlation, and arguments based on such 'runs of luck' do not seem very convincing. A further complicating factor is introduced if random errors in vergence are considered. Such errors are the inevitable consequence of independence of drift movements in the two eyes: and we saw in the previous chapter that mechanisms exist that generate slow corrective vergence movements in response to small changes in disparity, and which would therefore act to oppose drifts of opposite sense in the two eyes. However, if the vergence errors to be corrected are themselves the result of drift, the presence of such a mechanism, though effectively reducing the size of anticorrective drifts, *cannot* make the average direction of drift revert to the compensatory

direction. Yet St Cyr and Fender (1969a) found that on average, in binocular fixation, the position of the eyes was nearer to the target at the ends of horizontal (but not vertical) drift movements than at the beginning. One must conclude from this, contrary to the suggestion of Krauskopf et al (1960), and to Easter's (1971) finding in the goldfish, that from the point of view of the control of *vergence*, microsaccades are not on average corrective at all, but anticorrective.

6.4 Microsaccades

Since microsaccades have broadly the same function as ordinary gross saccades—that of bringing a visual target to the centre of the fovea—one might perhaps hope that they would be found to share the same control mechanisms. Some evidence that this is the case comes from a comparison of the time course of ordinary saccades and microsaccades of different amplitudes. Figure 6.7 (Zuber et al, 1965) shows the peak-velocity-duration relationship for microsaccades, plotted together with the same relationship for voluntary saccades: it is quite clear that the microsaccades fall convincingly on an extrapolation of the main curve. Similarly, Ginsborg (1953) was able to show that the *mean* velocities of microsaccades also lie on an extrapolation of the corresponding curve for ordinary saccades; such mean velocities in man are of the order of $200' s^{-1}$ (Ditchburn, 1980). The implication is that the same temporal pattern of activation is generated in the ocular muscles to produce the two kinds of saccade, from the smallest microsaccades of around 1–2' (Nachmias, 1961) to the largest voluntary saccades having amplitudes some six thousand times larger.




Figure 6.7. Maximum velocities of saccades and microsaccades as a function of their amplitudes, showing the essential continuity in the properties of the two kinds of movement (after Zuber et al, 1965).

One difference between microsaccades and gross saccades apart from size is that the former show a much greater degree of random variation both in direction and amplitude relative to the desired target, and in the time of onset of the movement. The statistical properties of these parameters have been investigated by many workers, and are summarised by Ditchburn (1973, pages 84–85) in tabular form. Median saccade amplitudes show great intersubject variability, ranging from about $1'$ (Nachmias, 1961) to nearly $23'$, while mean intersaccadic intervals vary from about 300 ms to 5 s or more. Microsaccades are predominantly corrective in nature, and the probability of their occurrence increases on average with the distance of the visual target from the centre of the fovea (figure 6.8) (Cornsweet, 1956; Boyce, 1967; de Bie and van den Brink, 1984): the same mechanism no doubt underlies the greatly increased saccade latency associated with very small target displacements (cf figure 4.5).

But since these movements take place on a background of continuous slow drift, it is not easy to disentangle a process in which the probability of a




Figure 6.8. Corrective properties of microsaccades. Below, probability of a microsaccade occurring in a period of 1 s, when the eye is at various deviations from the long-term mean fixation position. Above, the arrows show weighted mean amplitudes of microsaccades from various initial deviations. The length and direction of each is given by (mean absolute microsaccade amplitude) $\times (p_R - p_L)$, where p_R and p_L are respectively the probabilities of a microsaccade being directed to the right and to the left, at any particular initial deviation. It can be seen that in each case the effect of microsaccades is, on average, to bring the point of fixation close to its long-term mean (data from Cornsweet, 1956).

saccade depends essentially on the distance of the visual target, and one in which it merely depends on the time lapsed since the last saccade. Some evidence for the latter view is provided by the fact that the frequency of microsaccades does not increase in the expected manner when the drift rate is artificially increased by having the subject accommodate (Nachmias, 1959); though more recently de Bie and van den Brink (1984) have found, in contradiction, that reduction of retinal drift velocity does reduce the microsaccade rate. The effect of drift is on average to cause a steady increase in the fixation error as time passes after a microsaccade, so that one would expect the frequency distribution of intersaccadic times to show a rather steeper falloff than is the case for a strictly Poissonian process in which the probability is stationary with respect to time: such published data as there are (Ginsborg, 1953) contain too few observations for exact comparisons to be made, but indicate that such a model is too simple, and that at least in some subjects there are other factors that strongly influence the probability of microsaccades. It is now well established, for example, that microsaccades can be greatly influenced by *volition*: an increase in attention (as for example in fixating a fresh object) is associated with an increased frequency of microsaccades (Barlow, 1952). Subjects can also reduce the frequency of their microsaccades at will, or even abolish them altogether for limited periods (Fiorentini and Ercoles, 1966; Steinman et al, 1967); this is not achieved merely by relaxing the accommodation to degrade the quality of the retinal image (Steinman et al, 1969). Suppression of microsaccades is not something that requires special training; naive subjects can produce it immediately with suitable verbal instruction (Kowler and Steinman, 1980), and indeed it is observed spontaneously when subjects carry out tasks requiring particularly high visual acuity, such as threading a needle or firing a rifle (Winterson and Collewijn, 1976; Bridgeman and Palca, 1980). We have already seen that in many species fixational microsaccades are not observed at all (Collewijn and van der Mark, 1972; Winterson and Robinson, 1975), and in monkeys, microsaccade suppression can be elicited voluntarily under suitable conditions (Skavenski et al, 1975).

The level of illumination and colour of the target have little influence on the mean intersaccadic interval, within the photopic range (Steinman, 1965; Boyce, 1967): since the fovea is relatively ineffective at scotopic levels, one would expect a considerable alteration in the pattern of microsaccades, and this is indeed observed (Steinman and Cunitz, 1968). Under these circumstances, the saccades are found on average to move the target not *into* the fovea, but *out* of it. In complete darkness, as we have seen, the eye gradually drifts away from its initial position—as much as $1-2^\circ$ away after two minutes (Skavenski and Steinman, 1970)—and deviations of this magnitude are tolerated without correction by the saccadic system (Ditchburn and Ginsborg, 1953). Under *intermittent* illumination, particularly at frequencies in the range 3–4 Hz, the intersaccadic interval is significantly

reduced (West and Boyce, 1968): the explanation for this phenomenon is obscure. The shape and size of a target have surprisingly little effect on fixational stability: one is nearly as good at fixating the unmarked centre of a circle as one is at fixating a dot (Rattle, 1969; Bedell et al, 1984). Presumably the visual mechanisms involved are ones that decline relatively slowly from fovea to peripheral retina: localisation is not a function that demands particularly high visual acuity.

One might well wonder what on earth microsaccades are *for*, since so many animals seem to manage perfectly well without them, and indeed we ourselves suppress them when visually demanding work has to be done. As Bridgeman and Palca (1980) point out, because of saccadic suppression (discussed in chapter 11), microsaccades oblige us to spend a tenth of our fixational time in a state of partial blindness. Why do we do it? Steinman and his colleagues (1973) have suggested that microsaccades may be the result of the artificial circumstances of the laboratory, in which deliberate and unnatural fixation of objects is demanded for periods of time that far exceed their intrinsic interest. Microsaccades are apparently not observed in the fixational pauses of reading [Cunitz and Steinman (1969); but it is perhaps debatable whether reading can really be considered a 'natural' pattern of activity]. These authors suggest that microsaccades represent tiny searching or scanning movements, although this view conflicts to a certain extent with the fact that these movements are so strongly corrective in




Figure 6.9. Correlation between fixation errors and microsaccades. The graph shows correlation coefficients (r) between directed magnitudes of saccade components and projected retinal image positions at different times before the beginning of the saccade: the directions used are indicated by the orientation of the bisections of the circles. The broken horizontal line shows the smallest value of r which is significant at the 1% probability level (data from Nachmias, 1959).

nature; as may be seen from figure 6.8, or from the evident correlation between the directed magnitude of a microsaccade and the visual error immediately preceding it (figure 6.9) (Nachmias, 1959). The notion of deliberate small-scale scanning is considered further in section 6.5.1. In summary, it seems clear that microsaccades contribute little to maintaining fixation; but there remains the possibility that they help vision by introducing high rates of visual slip, a matter pursued in chapter 11.

6.5 Statistical variation in the point of fixation

If we ignore the distinction between the three components of fixational noise, we can describe their combined effects simply as if the point of intersection of the visual axis and the plane of fixation were subject to random perturbations around the fixation spot. We can then proceed to make an analysis of the two-dimensional statistics of these disturbances using as data the coordinates of this point at successive instants of time, and paying no particular attention to the velocity of the eye movement, or indeed to any other aspect of their actual time course. If the plane of fixation is divided up into a number of small equal areas, one can reduce this data to a two-dimensional histogram showing the total period spent by the point of regard in each of the sampling areas over some specified time interval, and display the results as contour maps (Bennet-Clark, 1964) (figure 6.10). (The whole process may be automated by reflecting a beam of light off the eye and on to a photographic plate, when the density in any area will be proportional to the fraction of the time that the eye has spent in the equivalent direction.) Histograms of this sort, if derived from data collected over a period of around half a minute, are typically close to a bivariate normal distribution, that is, a distribution of which any cross section is a normal one-dimensional distribution (Nachmias, 1959; Steinman, 1965). Thus one may make a further reduction of the data by describing such a distribution in terms of its mean, maximum and minimum standard




Figure 6.10. Two typical spatial distributions of fixation dwell times. The contours define areas within which the point of fixation was to be found 25%, 50%, 75%, and 100% of the time, in order from darkest to lightest (after Bennet-Clark, 1964).

deviations, and orientation to the horizontal. Since the maximum and minimum standard deviations are generally significantly different, the observed contour maps are distinctly elliptical in shape: the same directional preponderance is also observed in the angular distribution of microsaccades and drift (figure 6.11) (Nachmias, 1959). The reader can observe the directionality of his own miniature eye movements by steadily fixating the centre of a figure like that of figure 6.12: moiré fringes generated by the relative movement of the afterimage of the figure against the figure itself make directional preponderance particularly obvious.

A further simplification of the statistical description can be made by calculating the standard deviation of the distance of the point of fixation from the mean point of fixation, without regard to direction: if we call this quantity, σ , then

$$\sigma^2 = \sigma_1^2 + \sigma_2^2,$$

where σ_1 and σ_2 are the maximum and minimum standard deviations. The first measurement of σ was made by Barlow (1952), although he confined




Figure 6.11. Relative frequency of drift and microsaccades in different directions. The graph shows the percentage of drift and microsaccades made in particular 10° sectors by a single subject (subject JA, Nachmias, 1959). The upper, dotted area is for microsaccades, and the lower for drift.




Figure 6.12. A subjective demonstration of directional preponderance in the micromovements. If the reader fixates the centre of the pattern, micromovements will become apparent as moiré fringe patterns between the afterimage of the pattern and the pattern itself: the number of lobes in the pattern can be used to estimate the amplitude of the miniature movement.

himself to disturbances due to drift and tremor: his value of $15''$ is considerably smaller than those of later authors (extensively reviewed in Ditchburn, 1973), who considered all three components. Typical values of σ lie between about $1.5'$ and $4'$, and depend to some extent on the nature of the fixation target (Steinman, 1965; Boyce, 1967; Sansbury et al, 1973; but see also Murphy et al, 1974), and its position in the visual field. This implies that for over 60% of the time, the point of regard lies within an area of about one hundredth of a square degree, surrounding the visual target.

6.5.1 The stationarity of fixation

All these methods of data reduction make the implicit assumption that the process giving rise to the disturbance is a stationary one; that is, that its statistical properties do not vary with time. One such possible time-dependant factor that has sometimes been suggested (for example Barlow, 1952) is that the retinal point of fixation (that is, the point on the retina to which visual objects are brought) might be shifted around within a small area from time to time, and that this might in some way be beneficial to the visual system, perhaps by allowing different sets of receptors to examine the image, and possibly akin to the mechanism by which the attention can be directed to different parts of stabilised images (see Pritchard, 1958). The unimodal character of distributions obtained over periods of half a minute or so suggests that such adjustments are either rather small in comparison with the other disturbances, or that they occur over much longer time scales. Boyce (1967) used a 'cumulative sum' technique (calculating, in effect, the time integral in two dimensions of the eye position) and claimed to be able to show that slow changes of this type do in fact occur. However, the method of cumulative sums is somewhat tendentious, in that it amplifies low-frequency spectral components in proportion to their period and thus gives them a misleading and unnatural prominence. Strict statistical analysis has so far failed to provide evidence for slow adjustments in the retinal point of fixation, or for slow small-scale scanning of extended visual targets (to which it is equivalent). However, motivation is certainly a factor that upsets stationarity: the offer of a milk shake for good fixation was sufficient to reduce one subject's standard deviation by nearly a factor of two (Snodderly and Kurtz, 1985)!

6.5.2 Stability of gaze under natural conditions

Measurements of the kind just described are made under conditions in which the head is clamped so as to be as nearly motionless as possible. If what we are interested in are the visual consequences of movement of the retinal image across the retina, we might well wonder whether in practice the micromovements of fixation are really of any significance in comparison with the disturbances introduced by natural movements of the head. It turns out in fact that visual stability in all species examined depends much more critically on stability of the head than on any contribution from the fixational micromovements (Kowler and Steinman, 1980; Steinman and

Collewijn, 1980; Dieringer and Precht, 1982; Duwaer, 1982a; Steinman et al, 1982; van der Steen and Collewijn, 1984). In man, for example, average image velocities of some $15' s^{-1}$ are observed when fixating with the head clamped; with the head voluntarily held still this rises to some $30' s^{-1}$; and under conditions of natural head movement, to about $1-3^\circ s^{-1}$, a figure that would be expected to lead to a substantial loss of high-frequency acuity [Skavenski et al (1979); Steinman and Collewijn (1980); Steinman et al (1982); it has, however, been suggested that the last figure may be an overestimate, through instrumental error (Duwaer, 1982a)]. Similar relationships are found in the rabbit (van der Steen and Collewijn, 1984), though the range of velocities is rather higher. A curious point noted by Winterson et al (1979; see also Skavenski et al, 1979) is that the gain of the vestibulo-ocular reflex is smaller for very small head movement than for large; they suggest that this might be a deliberate feature serving to maintain the average velocity of random retinal slip, for the benefit of the visual system!

Structure