2012

Human optokinetic nystagmus: a stochastic analysis

Waddington, Jonathan

Human optokinetic nystagmus: a stochastic analysis.

Jonathan Waddington.

Centre for Robotics and Neural Systems.

University of Plymouth

A thesis submitted to the University of Plymouth in partial fulfillment of the requirements for the degree of:

Doctor of Philosophy.

January 11th 2012.

Abstract

Human optokinetic nystagmus: a stochastic analysis.

Jonathan Waddington.

Optokinetic nystagmus (OKN) is a fundamental gaze-stabilising response in which eye movements attempt to compensate for the retinal slip caused by self-motion. The OKN response consists of a slow following movement made in the direction of stimulus motion interrupted by fast eye movements that are primarily made in the opposite direction. The timing and amplitude of these slow phases and quick phases are notably variable, but this variability is poorly understood.

In this study I performed principal component analysis on OKN parameters in order to investigate how the eigenvectors and eigenvalues of the underlying components contribute to the correlation between OKN parameters over time. I found three categories of principal components that could explain the variance within each cycle of OKN, and only parameters from within a single cycle contributed highly to any given component. Differences found in the correlation matrices of OKN parameters appear to reflect changes in the eigenvalues of components, while eigenvectors remain predominantly similar across participants, and trials.

I have developed a linear and stochastic model of OKN based on these results and demonstrated that OKN can be described as a 1^{st} order Markov process, with three sources of noise affecting SP velocity, QP triggering, and QP amplitude. I have used this model to make some important predictions about the optokinetic reflex: the transient response of SP velocity, the existence of signal dependent noise in the system, the target position of QPs, and the threshold at which QPs are generated. Finally, I investigate whether the significant variability within OKN may represent adaptive control of explicit and implicit parameters.

Contents

List of Figures

List of Tables

Acknowledgements

This work was supported by Plymouth University and funded by the Engineering and Physical Sciences Research Council. A number of people have helped me through my postgraduate career, and I wish to thank them all for their advice, inspiration and support.

My director of studies, Professor Chris Harris, has given me incredible encouragement during every stage of this project. His advice and feedback has been invaluable and I feel fortunate to have had him as a mentor during my PhD candidature. As my second supervisor, Professor Roman Borisyuk has likewise provided great assistance. His seminars on advanced mathematics at the very beginning of my candidature were truly inspiring and demonstrated to me what a powerful tool mathematical modelling can be. I would also like to thank my thesis committee for taking the time to read this thesis and provide feedback on the research I have conducted.

I am sincerely grateful to a number of other people from Plymouth University. Members of the Centre for Theoretical and Computational Neuroscience, particularly Lucy, Bassem, Salvador, Andy, Floortje, Masud, Ulli, and Beibei have all been influential in my postgraduate life and made the office an enjoyable environment to work in. In addition, Martin Coath saved my life by teaching me LaTeX.

My friends and family, both in and out of the Plymouth area, have supported me throughout and helped me maintain a healthy perspective on life. Lisa, Ashley, Murph, Jay, Deano, Bres, Becky, Natasha, Sophie, Jaysan, and Lacy have each helped fill my postgraduate years with fun and happy times. My mum and dad have been incredibly supportive, emotionally and financially, particularly as deadlines drew dauntingly close. My little sister and brothers, Leona, Colin and Ian, have also been inspirational in achieving so much in their own fields of interest.

Finally, I owe a very special debt of gratitude to Kitiara and Candice, for whose love and support I am eternally grateful.

Authors declaration

At no time during the registration for the degree of Doctor of Philosophy has the author been registered for any other University award.

Relevant scientific seminars and conferences were regularly attended at which work was often presented.

Signed: _____

Date: _____

Posters and conference presentations :

2011:

Perception and Action Group seminar series, *Cardiff University.* Invited talk: **Human optokinetic nystagmus: What, where, when and why?**

2009:

Centre for Robotics and Neural systems inaugural workshop, *Plymouth University.* Poster presentation: **Assessment of models describing the distribution of optokinetic nystagmus quick phase intervals**

British Oculomotor Group meeting, *University of Liverpool.* Poster presentation: **Distribution of optokinetic nystagmus quick phases revisited**

Word count for the main body of this thesis:43,360

Chapter 1

Introduction.

Vision and movement are inextricably linked. As we move through the world, the visual environment appears to stream by. This motion would cause our vision to become blurred, so it is essential that the visual system is able to compensate in order to maintain maximum visual acuity. Sensorimotor systems have evolved in order to keep gaze stable during head and body movements. The optokinetic response (OKR) is generated by the patterns of image motion that are presented to the visual system during locomotion, and optokinetic nystagmus (OKN) is the eye movement response generated. It is an alternating sequence of slow compensatory eye movements made in the direction of the image motion, and fast eye movements made predominantly in the opposite direction that keep gaze in a stable location. The timing and amplitude of these slow phases (SPs) and quick phases (QPs) are remarkably variable (fig. 1.1).

This thesis investigates the pattern of eye movements associated with human OKN and the complex interrelation of the parameters that characterise the movement. In the past, attempts to understand the details of OKN have recognised the remarkable variability in the parameters of this fundamental behaviour but have failed to find an explanation. Models used to explain this behaviour have often been based on the classical systems engineering approach of formulating a hypothesis using block diagrams, and testing these hypotheses using the techniques of

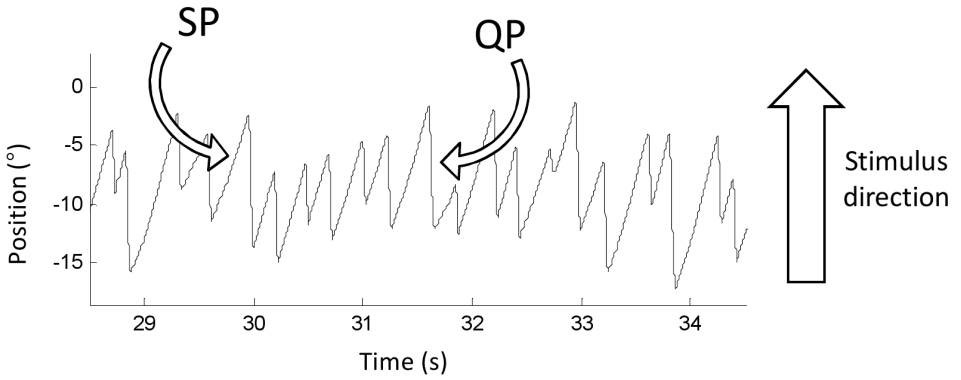

Figure 1.1: Example of an eye position trace during OKN stimulation, illustrating SPs made in the stimulus direction and QPs made opposite to the stimulus direction.

systems analysis. However, in this investigation I take a wholly different approach and apply methods of stochastic analysis to examine the statistical relationships between primary eye movement parameters such as the velocity, amplitude and start position of SPs and QPs. I will present results that illustrate that there is a structure underlying the variability of OKN eye movement parameters, and develop a linear stochastic model of the OKN system based on analysis of experimental results. Then I will use the model to illustrate the observed empirical behaviour, and make some important predictions about the specific goals of the response.

1.1 Aim of the research.

The aim of the present series of research was to demonstrate that despite the remarkable variability of individual OKN parameters, there is an underlying structure to the relationships between OKN parameters that can aid our understanding of how and why OKN behaves as it does. Both experimental and modelling approaches have been used to achieve this aim.

The primary aim of the experimental approach was to understand how different

OKN stimuli cause various OKN parameters to change, so the spatial and temporal characteristics of the OKN stimuli were varied in order to generate a range of responses in participants. A secondary aim of the experimental studies was to create a source of data that could be used to determine and validate a model of the OKN system, so each individual trial of the OKN responses was recorded over a long period in order to evaluate the random fluctuations of OKN parameters as they evolved over time.

The primary aim of the modelling approach was to develop a simple linear and stochastic model of how the OKN system behaves, which would be able to predict the distribution and interrelation of OKN parameters observed empirically. In order to achieve this, principal component analysis (PCA) was performed on the correlation matrices of OKN parameters in order to determine how the underlying components contribute to the correlation of OKN parameters. PCA is quite a well known method used in feature extraction and dimension reduction, and I shall give a detailed background to how it was implemented in the present series of research in chapters 3 and 5. Informally, it can be thought of as a method for determining the correlations between parameters that contain the most amount of information.

1.2 Organisation of the thesis.

The thesis is divided into nine chapters. After this introduction, a more in-depth background of the OKN literature is presented. The literature review gives an overview of the field and highlights those areas that are not currently well understood, giving details of those particular studies that have attempted to analyse the stochastic nature of OKN parameters. Details of the experimental study methods are presented in the third chapter, describing the range of experimental protocols that were used and the system for data acquisition. In chapters 4-8, the results

from the experimental and modelling studies are presented and discussed.

Chapter 4 presents descriptive statistics of the recorded OKN parameters and the results of repeated measures analysis of variance (RM-ANOVA) performed on the data from experiment 1. The RM-ANOVA results provide the main effects of changing stimulus speed on the mean value and standard deviation of OKN parameters.

Chapter 5 presents the results of PCA performed on OKN parameters recorded from the first experimental protocol. The results of this analysis highlight an underlying structure to the data that enables the development of a linear and stochastic model describing the OKN system. This system consists of three 1^{st} order Markov processes. The model is used to replicate empiricial behaviour, and we find four relationships between OKN parameters that remain constant across participants and despite changes in stimulus speed.

Chapter 6 presents the results of RM-ANOVA performed on the data from experiments 2 and 3. The main and interaction effects of changing stimulus parameters on the mean value and standard deviation of OKN parameters are reported. The within-subjects factors tested were the speed, spatial frequency, and pattern of the OKN stimulus. A number of significant main and interaction effects on OKN parameters were found. The results found with the Markov model are replicated with the data from experiments 2 and 3, and so we combine all data sets for a more thorough investigation of the behaviour of the model. We explore how the stimulus speed affects the model in order to generate the OKN response, and confirm the existance of signal dependent noise in the system. Using the model we examine the more poorly understood aspects of OKN, such as at what point (in time or space) QPs are generated and what location they are targeted towards.

Chapter 7 presents the results of an investigation into the particularly interesting

distribution of SP duration. It reports the results of a study designed to test the goodness of fit of proposed probability density functions (PDFs) to the histograms of SP duration from each trial. The distribution of SP duration is particularly interesting due to its similarity with the distribution of saccade latencies to visual targets. Stochastic models of saccade latency are often used in the eye movement literature to explain simple models of decision making. The histograms of both SP duration and saccade latency have a characteristic positive skewness and are significantly different to Gaussian, so the PDF that best fits the histograms of these data is often hotly debated. The results of this investigation demonstrate that the PDF predicted by the Markov model of OKN best fits the histograms of SP durations, when compared to a range of other proposed PDFs.

Chapter 8 proposes three possible cost functions to explain the variability in SP velocity, and presents the results of analysis on these models. Minimising one cost function in particular gives optimal values of SP velocity that match the mean values of SP velocity observed in the data, indicating that SP velocity may be being optimised. This cost function assumes the goal of the system is to minimise the average positional error in tracking a moving target over the course of each SP, and not only the retinal slip error.

Chapter 9 presents my conclusions from the major findings. Possible directions for future research based on the predictions that the models make are considered, as are possible improvements that could be made to the models based on their limitations.

Chapter 2

Background.

In this chapter the elements of the ocular motor system will be reviewed with regards to the essential characteristics of two important gaze stabilising mechanisms: the OKR, and the vestibulo-ocular reflex (VOR). In the second half of this chapter the literature that pertain directly to the methods and results that are presented in this thesis will be discussed.

2.1 The human eye and visual acuity.

In order to see the world around us light must travel from the field of view to the retina, at the back of the eye, through the visual axis. When light first enters the human eye it is refracted by the curvature of the cornea, the transparent component of the outer layer of the eye, accounting for a significant portion of the eye's total optical power. Once the light has passed through the cornea it is further refracted by the curvature of the lens, which is capable of being adjusted in order to focus on a specific depth plane via the accommodation system.

The light is focused on to the retina, which is comprised of hundreds of millions of neurons distributed into a number of layers. Many different types of neurons exist in the retina but those capable of phototransduction that contribute to sight directly are the rod cells and cone cells, of which the rod cells are most numerous. Rod cells can be triggered by a very small number of photons and detect dim

light signals under dark-adapted conditions. There are three types of cone cells in the retina that each have different peaks in their absorption spectra (Brown & Wald 1964) and respond differently depending on the wavelength and intensity of the light. Near the inner surface of the retina are other neurons that receive information from the rod cells and cone cells. Bipolar and amacrine cells project to retinal ganglion cells, which then transmit visual information from the retina to other areas of the central nervous system. Cone cells are capable of discriminating finer detail than rod cells as each cone cell sends information to only one retinal ganglion cell, whereas retinal ganglion cells recieve information from a number of rod cells.

Near the centre of the retina is the oval-shaped macula, approximately 5mm in diameter, it is defined as having two or more layers of retinal ganglion cells. The macula contains predominantly only cones, whereas in the rest of the retina the two different types of photoreceptor are interspersed. At the very centre of the macula is the fovea, a central pit that contains the highest concentration of cone cells in the retina arranged in the most efficient packing density of a hexagonal mosaic. The resolution of the retina is related to the density of the cone mosaic, and so the fovea is capable of resolving higher spatial frequency images than the peripheral retina.

Visual acuity is a measure of the spatial resolution of vision, and it can be limited by corneal or lens imperfections that cause the sharpness of retinal focus to fall. It can also be limited by the neural characteristics, such as the density of cones and rods in the retina, or pathological conditions of the central nervous system that affect the visual system. In the clinical environment it is usual to consider measurements of visual acuity at high spatial frequencies and high contrast. However, it is also useful to examine visual acuity at levels of low contrast. The threshold

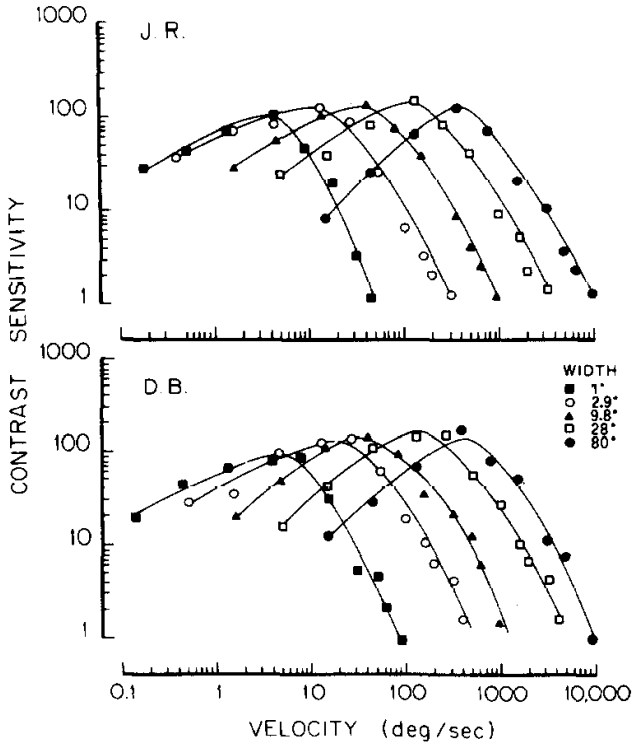

Figure 2.1: Contrast sensitivity as a function of image velocity. The contrast sensitivity function to single cycle, biphasic sinusoidal bars of various widths are shown as function of image velocity, for two participants. Figure taken from Burr & Ross (1982). Permission to reproduce this figure has been granted by Elsevier.

contrast necessary to perceive the details of a visual scene is a function of the spatial frequency content of that scene. Campbell & Robson (1968) found that the contrast sensitivity (1/threshold contrast) function peaked at around 4 cycles per degree (cyc/°) with sensitivity dropping off either side of the peak. This means that when viewing a stimulus at 4cyc/° we can detect smaller differences in contrast than at any other spatial frequency.

The contrast sensitivity for a moving visual scene is a function of the spatial frequency content, and the temporal frequency content, of that scene. This means that the contrast sensitivity is also a function of the speed of the stimulus. Burr & Ross (1982) found that for high spatial frequency stimuli the peak of the contrast

sensitivity function occurred at low stimulus speeds, but that for stimuli with very low spatial frequency the peak of the contrast sensitivity function occurred at much higher stimulus speeds. The study reported that participants were able to perceive a bar that was 80° in width moving at a remarkable 10,000°/s (fig. 2.1). However, for high spatial frequency stimuli the contrast sensitivity function decays logarithmically at speeds greater than 3°/s. It seems that an essential prerequisite for the accurate processing of visual information consists of a sharp retinal image that is kept relatively still on the fovea.

2.2 Eye movements and the oculomotor system.

At the beginning of the 20th century it was thought that eye movement strategies could be categorised into five different subsytems: the VOR, OKN, stationary eyes during fixation, very fast eye movements in between two periods of fixation (saccades), and the simultaneous movement of both eyes in opposite directions (vergence) (Dodge 1903). It is well known that different species have a predominance for different types of eye movement strategies, of particular note is the difference between foveate and afoveate animals.

In some afoveate animals there are specialised distributions of retinal cells, either circular (area centralis) or elongated (visual streak), where the cones become more densely packed. In afoveate animals that do not have these specialised distributions of retinal cells the position of the image on the retina is of lesser importance, as there is no particular area of the retina that has a higher degree of resolution. In these species, systems have evolved to stabilise images, and minimise the retinal slip caused by optic flow via the OKR and VOR.

In foveate animals, eye movements have evolved to bring objects of interest on to the fovea, where the object can be seen best (using saccades, fixations and

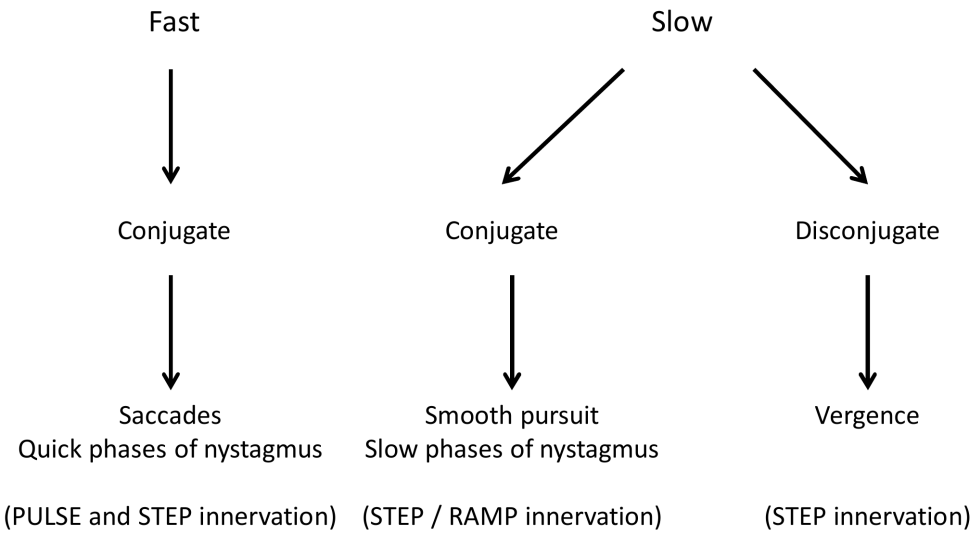

Figure 2.2: Subcategories of eye movements that are either fast (abrupt), or slow (smooth). Saccades are very fast eye movements made between two periods of fixation that generally bring images of interest on to the fovea. Smooth pursuit is a slow tracking movement that holds a small image of interest on the fovea. Vergence moves the eyes in opposite directions such that the images of an object are held on both foveas simultaneously.

vergence eye movements). Saccades are fast conjugate eye movements that bring objects of interest onto the fovea, while smooth pursuit holds a small moving target on the fovea. Vergence movements are disjunctive eye movements that are made so that images of a single object are placed on both foveas.

These two different eye movement strategies are often classified as either gaze-stabilising or gaze-shifting. However, there is a certin degree of overlap in these two categories, as eye movement strategies that shift gaze must also be employed during eye movement strategies that stabilise gaze. This has led to a more general classification of eye movement strategies as simply either fast or slow (Steinman et al. 1990).

Figure 2.2 illustrates the original five subcategories of eye movements as they correspond to either fast or slow eye movement strategies, and notes the type

of innervation required in order to generate each movement. A pulse (or burst) of innervation is required in order to overcome the viscous drag imposed by the tissue surrounding the eye when a fast eye movement is made, such as a saccade made to a visual target. A step, that is a new tonic level, of innervation is required in order to hold the eye in an eccentric position and overcome the elastic forces that tend to bring the eye back to its central position. It follows that a pulse-step command is required in order to produce a fast eye movement to an eccentric visual stimulus and maintain steady fixation of that stimulus (fig. 2.3). Slow eye movements require only a step (or ramp) command to follow a smoothly moving visual stimulus. Two pairs of muscles enable the eye to rotate around the horizontal and vertical axes (medial and lateral rectus, and superior and inferior rectus), and a third pair of muscles enables cyclotorsional movements (superior and inferior oblique). To a good approximation eye movements can be described as rotations around three axes.

The oculomotor system performs a number of strategies when moving the eyes in order to provide the central nervous system with adequate visual sensory information. Of particular importance in regards to the topic of this thesis is its role in compensating for optic flow. Optic flow is the apparent visual motion that is presented to the retina as we walk or move through the world. For example, when we move forward the retinal projections of objects in the visual field grow outward from a stationary central point. If we were sat in a train and looking out of the window the visual scene would appear to move in the opposite direction to the motion of the train and, while distant objects would appear to move very slowly, closer objects would appear to move much faster. We have discussed how high spatial frequency resolution decays when the visual image motion on the retina (retinal slip) exceeds speeds greater than $3^{\circ}/s$, and in order to maintain precise

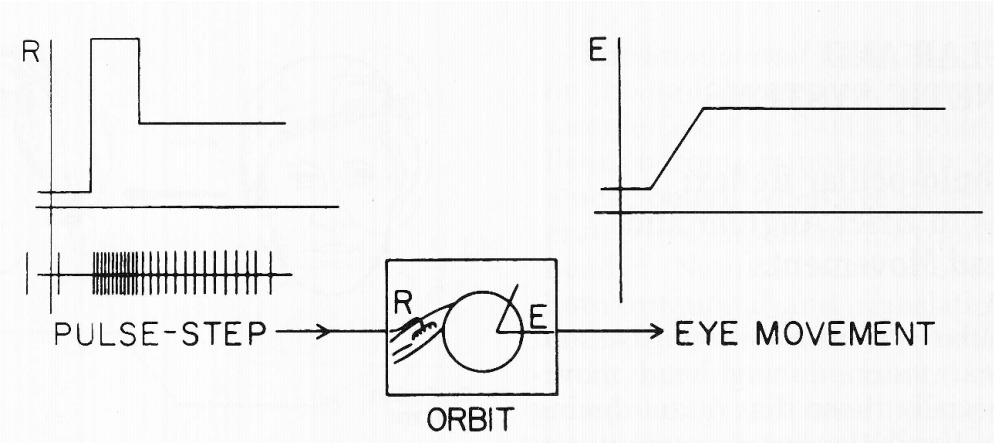

Figure 2.3: Pulse-step command for a saccade. The ocular motoneuron firing rate and eye position (in the orbit) are shown as a function of time, illustrating the pulse and step commands required to generate a saccade. The long single vertical line represents visual stimulus onset. The multiple vertical lines in the left plot represent action potentials of an ocular motoneuron. R, neuron discharge rate; E, eye position in the orbit. Figure taken from Leigh & Zee (1999). Permission to reproduce this figure has been granted by Oxford University Press.

vision it is essential that this motion is compensated for. We find that almost all vertebrates (Walls 1962, Huang & Neuhauss 2008), and some invertebrates with a mobile head or eyes (Land 1999), have evolved oculomotor subsystems that perform this function and stabilise the gaze of the animal on the visual scene. One such gaze-stabilising mechanism, which ensures high visual acuity is maintained during movements of the head, is the VOR.

2.3 The VOR.

The VOR stabilises images on the retina during head movement by producing an eye movement in the direction opposite to the head movement. The signals that drive the VOR originate in the vestibular apparatus in the inner ear, where the semicircular canals detect angular head acceleration and the otolithic organs (the utricle and saccule) detect linear (translational) head acceleration. As the eyes

are not at the center of rotation of the head, even pure head rotations produce a linear displacement of the eyes, and the size of the movements of each eye must be adjusted independently so that they remain pointed towards the same position in the depth plane. The system operates without any immediate feedback, so is an open loop (feedforward) system.

The VOR is able to respond to head movements with minimal delay, and eye movements can occur within 6-15ms from the initial head movement (Maas et al. 1989). The importance of the VOR is that it is the only eye movement that can be generated with a short enough latency to compensate for head movements that occur in the frequency range caused by footfall perturbation during walking (0.7-5.7Hz; Das et al. (1995)). It is primarily quantified by the ratio of the amplitude of eye rotation to the amplitude of the head rotation (VOR gain), and the temporal difference between eye velocity and head velocity (VOR phase). For sinusoidal stimuli the perfect VOR would have a gain of 1.0 and a phase of 180° indicating that eye and head velocity are exactly equal in magnitude and opposite in direction.

It is important to consider the differences between the VOR in light and in dark environments. Measured under light conditions, the VOR gain will approach 1.0 at most frequencies of head movements, but under dark conditions the amplitude of VOR gain decreases and changes in the variability of gain occur when the frequency of head movement drops below approximately 1Hz (Robinson 1976, Shelhamer et al. 1994, Fetter et al. 1995). Under light conditions the gain will approach 1.0 even at low frequencies of movement due to the additional contribution of the OKR generated by the motion of the visual image that the VOR does not compensate for. Not only does the VOR perform suboptimally at low frequencies, the gain of the VOR decays gradually over prolonged periods of rotation due to

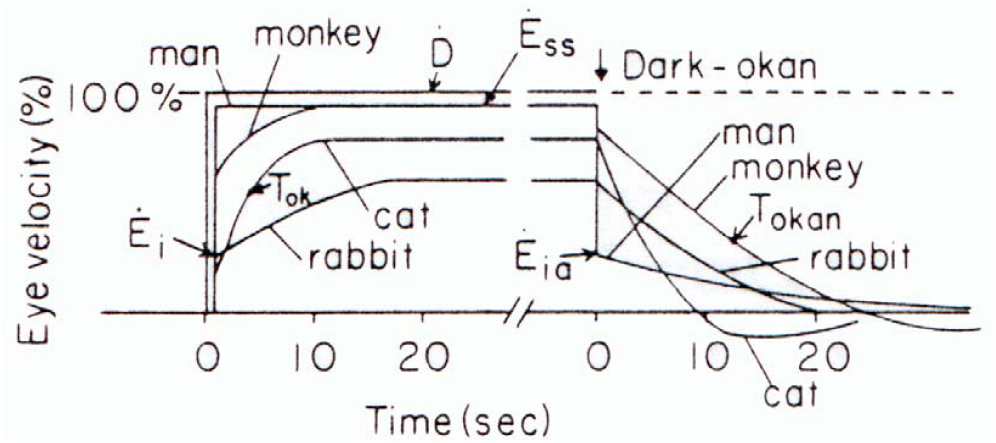

Figure 2.4: Time course of OKN and OKAN in man, monkey, cat and rabbit in response to a suddenly illuminated OKN stimulus, and later in time a sudden drop in illumination to complete darkness. $\dot{E}_i$, initial eye velocity; $\dot{E}_{ia}$, initial eye velocity after darkness; $\dot{E}_{SS}$, steady state eye velocity; T_{ok}, transient OKN; T_{okan}, transient optokinetic after-nystagmus. Figure taken from Robinson (1981*a*). Permission to reproduce this figure has been granted by Oxford University Press.

the mechanical properties of the semicircular canals. Prolonged rotation generates vestibular nystagmus, where initially SPs are compensatory and made in an equal and opposite direction to the head rotation and alternate with QPs made into the direction of head rotation. Over approximately 30s the VOR decays and the eyes become stationary. In the light, the OKR supplants the fading vestibular response over this period, and maintains a stable retinal image once there is no further contribution from the vestibular system.

2.4 The OKR.

In figure 2.4 the OKR to an OKN stimulus is illustrated for a number of different species including humans. The differences in the initial rise in SP velocity for the different species can be explained by the differences in the specialised distribution of retinal cells.

In afoveate lateral-eyed animals, such as the rabbit, a large portion of the retina

must be stimulated in order to elicit the OKR, but when the optokinetic system is tested using an artificial stimulus (such as a patterned drum rotating around the animal) the OKR slowly charges (SP velocity increases) over time (Collewijn 1981, Dubois & Hollewijn 1979). This is the compensatory behaviour discussed in section 2.3, where the OKR supplants the fading VOR in the light in order to maintain a stable retinal image. If the illumination in the room is suddenly turned off and the animal is left in complete darkness, the OKR will begin to discharge and will produce the optokinetic after-nystagmus. During the period of optokinetic after-nystagmus, the nystagmus will continue for some seconds with a decreasing SP velocity, counteracting the postrotational vestibular nystagmus that is made in the opposite direction to rotation as the mechanical properties of the semicircular canals return to normal (Barratt & Hood 1988). This charging and discharging behaviour is mediated by the velocity-storage mechanism, which also improves the ability of the VOR to transduce the low frequency components of head rotation. Raphan et al. (1979) modelled the velocity-storage mechanism as a non-ideal integrator coupled to the visual and peripheral vestibular systems, where the integrator stabilises eye velocity during whole field rotation and improves the performance of the VOR by extending the time over which the VOR can compensate for head movement. A direction bias for the OKR has also been found, predominantly in lateral-eyed animals but also in very young infants, that favours temporal-nasal directed stimuli. In lateral-eyed animals it is assumed that this bias reflects a need to avoid pinning the eyes onto the visual scene behind the animal as it moves forward, and so nasal-temporally directed stimuli.

In foveate animals, such as humans and other primates, there is a smooth pursuit system. The smooth pursuit system generates eye movements that are designed to keep a small visual target on the fovea, and maintain high visual acuity whilst

tracking. The latency of smooth pursuit eye movements are usually more than 200ms, but predictive mechanisms can adjust the eye movements when the motion of the target can be anticipated (Barnes 1993). Smooth pursuit is normally associated with tracking a foveal target whilst ignoring the motion of the background on the peripheral retina. However, there is some evidence that the pursuit system supplements the optokinetic system in order to boost SP velocity and match the speed of the stimulus motion. The ocular following response is a specific visual tracking mechanism similar to smooth pursuit, but with ultra-short latency (<85ms in humans), that is responsible for holding gaze on an object during self motion (Miles et al. 1986, Miles 1995, 1998). In this strategy the visual system is thought to compensate for movement of the head by generating eye movements appropriate for the depth plane within which the visual target exists, using extraretinal information. In effect, binocular images are compared, and where the images on both eyes match gives the position (and thus plane) of fixation. Whereas objects behind or in front of the plane of fixation are considered to have binocular disparity and are seen as double (fig. 2.5). The ocular following response is thought to follow only those parts of the visual scene that occupy the same corresponding positions on both retinas.

In a study by Cheng & Outerbridge (1974) selective optokinetic stimulation of different parts of the retina was achieved by deleting a variable portion of the stimulus that was controlled by an eye position signal from the participant, in such a way that the deleted area always remained in the central vision of the participant. They found that the gain of OKN was reduced as the size of the central vision deletion was increased. However, Hood (1967) had illustrated that patients with a scotomatous eye could still generate high SP gain, albeit with some differences that appeared to be characteristic of vestibular nystagmus and nystagmus under

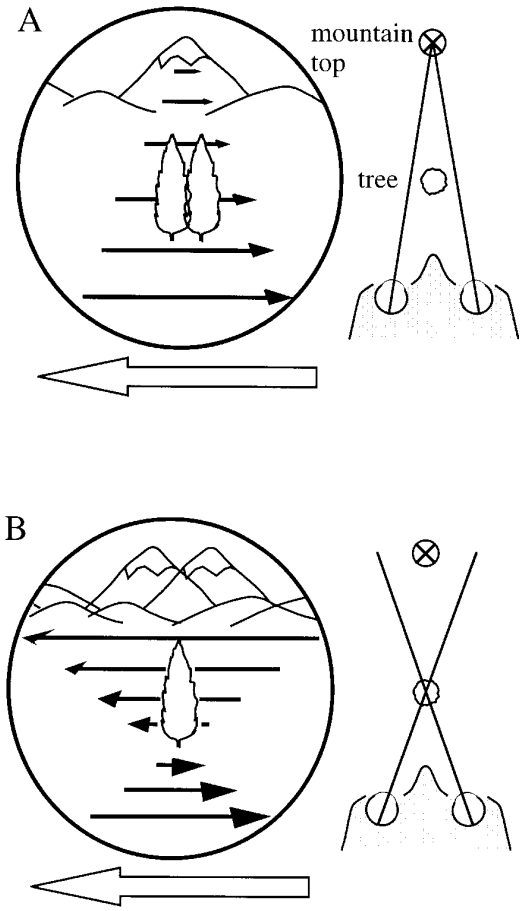

Figure 2.5: Illustration of optic flow experienced during translation, such as looking out of a train window towards the distance. (A) The observer is looking towards the mountain top and the tree in the foreground appears double due to binocular disparity on the retina. (B) The observer is looking towards the tree and the mountain top in the distance appears double due to binocular disparity. The large white arrow represents the motion of the observer, the smaller black arrows represent the magnitude of apparent motion of objects in that depth field. Figure taken from Busettini et al. (1996). Permission to reproduce this figure has been granted by Nature Publishing Group.

liminal illumination. In fact Hood (1967) claimed that SP gain was predominantly controlled by peripheral vision. In patients with unilateral scotoma, when the non-scotomatous eye was covered, the SP gain would exceed speeds of 90°/s as the stimulus speed was increased to 100°/s. Whereas when the scotomatous eye was covered the SP velocity would not exceed 50°/s, and even began to decrease as stimulus velocity was increased past 80°/s. van Die & Collewijn (1986) found a similar result in scotomatic patients, but also found that elimination of foveal stimulation was accompanied by a decrease in SP velocity when compared to whole field stimulation in normal participants. When participants are asked to fixate on a target in the centre of a moving visual field they are able to completely suppress the usual nystagmus response (Brandt et al. 1973). As nystagmus is completely suppressed during fixation, this has been cited as an indication that the pursuit system is able to dominate the OKR. It is clear that relative contribution of smooth pursuit mechanisms, such as the ocular following response, to the OKR are still not fully understood.

2.5 The neural substrate of OKN.

It appears that, in foveate species, both the smooth pursuit and optokinetic systems contribute to OKN during head translations and rotations. In humans, there is a very fast response within the first 1 to 2 seconds of OKR stimulation (the early OKN response; eOKN) that is followed by a slower buildup of stored neural activity that is revealed in optokinetic after-nystagmus when the subject is placed in darkness (the delayed OKN response; dOKN).

Neurophysiological studies in primates have revealed that vestibular nucleus neurons that respond to head rotation also respond to optokinetic stimuli (Henn et al. 1974), and that they can continue to fire even after the lights are turned

off (Waespe & Henn 1977). It appears that as the vestibular discharge declines during sustained rotation, the optokinetic system compensates in order to maintain a steady vestibular discharge, and this appears to be the neural correlate for dOKN.

Two important structures for the generation of OKN have been found in the pretectal complex, namely the nucleus of the optic tract (NOT), and the nuclei of the accessory optic system that consists of the dorsal, lateral, and medial terminal nucleus (DTN, LTN, MTN, respectively). The neural responses from NOT neurons are very similar to responses from DTN neurons, and since they cannot be differentiated anatomically are often referred to as the NOT-DTN. The importance of the NOT-DTN in generating OKN has been demonstrated by an ipsiversive reduction in SP velocity after damage to the NOT-DTN, and by electrical stimulation of the NOT-DTN producing ipsiversive SPs (Cohen et al. 1992). Neurophysiological studies have shown that NOT-DTN neurons respond with directional selectivity to visual motion (Hoffmann 1988), and the preferred directions of the DTN, LTN and MTN are roughly equivalent with the planes of action of the antagonistic pairs of extraocular muscles and the planes of the semicircular canals (Walley 1967, Simpson et al. 1979). In the primate LTN neurons respond with directional selectivity to upward motion (Mustari et al. 1988), but the primate MTN could not be found (Giolli 1963). Two different types of neurons were found in the primate NOT-DTN, in approximately half of the recorded neurons the firing stopped abruptly when the primate fixated a stationary target during OKN stimulation, and in the other half neuronal firing continued (Ilg & Hoffmann 1996). Further, when a smooth pursuit target was tracked across a textured background, the neurons that stopped responding when the primate had fixated a stationary target responded to the movement of the smooth pursuit target. Whereas the neurons

that continued to fire during fixation responded to the peripheral retinal image motion due to tracking the moving target over the textured background. These two different types of cells were termed "target velocity cells" and "background velocity cells". It appears that these two different cell types found in the NOT-DTN may be the neural correlates for the smooth pursuit and optokinetic components found in behavioural OKN studies.

2.6 The systems engineering approach.

Control theory deals with the behaviour of dynamical systems, where the external input of a system is manipulated by a controller to obtain the desired effect on the output of the system, and the input and output of a system are related to each other by the transfer function. The use of control systems analysis in oculomotor research was popularised by Robinson and revolutionised the way in which the ocular motor community continue to investigate eye movements. Fuchs et al. (1993) noted that Robinson had coined many phrases in common use today such as the "pulse and step", the "burst generator", and the "oculomotor integrator". Control systems analysis as applied to oculomotor research is essentially a method of modelling the sensorimotor control of the oculomotor systems using the block diagrams and algorithmic models usually employed in the control of systems. Using block diagrams as hypotheses for signal processing pathways in the brain we can pose questions about how the pathways may work, and simulate the output of the system given a certain input. If we define system output as some measureable quantity, such as the gain of the OKR, we can compare the simulated output with empirical data and determine the quality of model, based on the number of empirical observations it can simulate and the number of assumptions it requires.

An example of a simple negative feedback loop is illustrated in fig. 2.6. The OKR

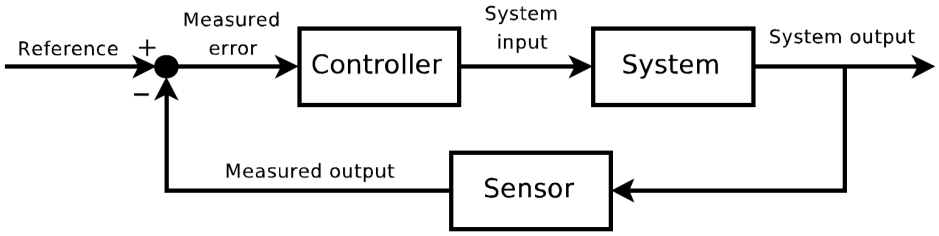

Figure 2.6: A simple model of a feedback loop. The dynamic behaviour of the system is controlled through negative feedback, rather than positive feedback, because the sensed value is subtracted from the desired value at the summing junction on the left to create the error signal, which is amplified by the controller.

is a closed-loop system, and can be considered as a simple negative feedback loop where the reference is the stimulus velocity, and the measured error is retinal slip, the controller updates SP velocity based on the measured error in such a way that retinal slip will be minimised. In such a closed-loop system we would predict that OKN gain would be very stable, and approximately 1.0. The finding that the visual system can be split into various component parts allows for models such as this to be developed. Robinson (1981*b*) proposed a more complicated model of the optokinetic system, which also included the system that generates the VOR (fig 2.7). In this box diagram it was assumed that the optokinetic system was only concerned with velocity, and not the position of the eye, head, and retinal images, and so the action of the neural integrator is not shown explicitly. The action of QPs were also ignored, in order to investigate only how the OKR responds to changes in head and eye velocity.

2.7 Relationships between parameters in human OKN.

While the control systems approach has allowed a greater understanding of the mechanisms that generate SP velocity and the gain of the system, and much is known about the physiology of the OKR, the processes that determine specific

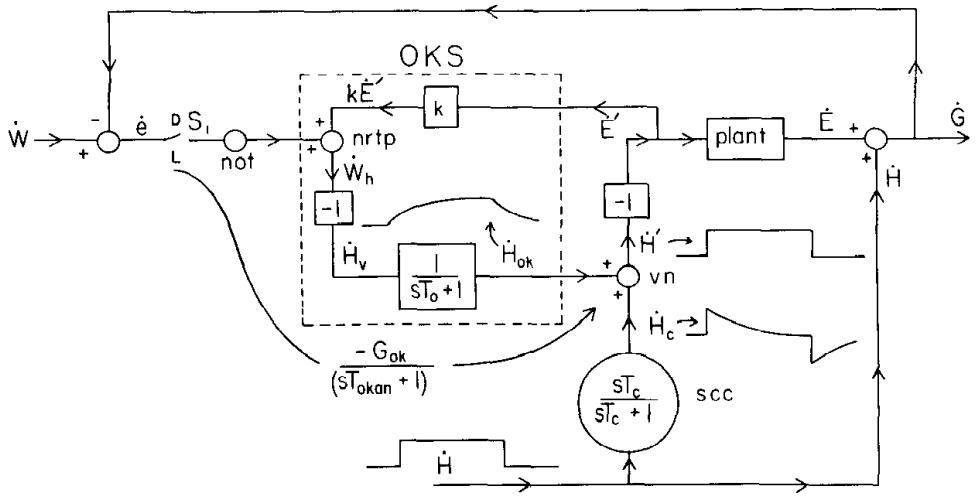

Figure 2.7: Block diagram illustrating a model of OKN, including its connection with the vestibulo-ocular reflex. $\dot{H}_{\rm ok}$, optokinetic signal; vn, vestibular nucleus; $\dot{H}$, head velocity; $\dot{H}_{\rm c}$, transient semicircular canal signal; $\dot{e}$, retinal slip velocity; not, NOT; $\rm G_{ok}$, optokinetic gain; $\rm T_{okan}$, optokinetic after-nystagmus time constant; $k\dot{E}'$, corollary discharge pathway; nrtp, nucleus reticularis tegmenti pontis; S_1, switch from light to dark. Figure taken from Robinson (1981*b*). Permission to reproduce this figure has been granted by Annual Reviews.

parameters such as when QPs are triggered or to where the end points of QPs are targetted remain elusive.

The earliest reported consideration on the threshold at which QPs are triggered is attributed to Ter Braak (1936) who proposed that QPs could be triggered by three possible mechanisms. First, QPs could be triggered by an extraocular proprioceptive signal indicating an eye position threshold. Second, the threshold may be determined by an internal clock or central pattern generator (attributed to Ohm (1928)). Finally, the threshold may be determined by the combination of both an eye position control system and an internal clock. Investigating the VOR, using the control systems approach, Chun & Robinson (1978) extended the original eye position hypothesis by modelling a local feedback loop in the pons which rapidly drives the eyes to a position specified by a vestibular signal, and proposed there were two internal signals that represented the start and end position of QPs. While the results appeared satisfactory they expressed concern at the ease to which mean OKN parameters can be predicted by models of the system, and concluded that measures of the distribution of parameters were necessary to determine the goodness of fit of a model. Lau et al. (1978) demonstrated that the threshold for triggering a (vestibular) QP in humans was dependent on instantaneous eye velocity, and not only on eye position.

An even more difficult problem to solve is understanding how the amplitude of QPs is determined, or where QPs are directed towards, particularly as there is no obvious visual target. QPs are usually made in the opposite direction to optic flow, and tend to reset eye position to maintain gaze around the desired visual region, but have a tendency to overshoot into the negative field (contraversion). The degree of contraversion in rotational OKN appears to be a function of the stimulus speed that saturates at speeds greater than 15°/s (Garbutt et al. 2002).

A well known relationship between the amplitude, duration and peak velocity of saccades is the main sequence. The trajectory determined by this relationship is an optimal compromise between the speed and accuracy of a saccade to a visual target, under the constraint that there is signal-dependent noise in the motor command (Harris & Wolpert 1998, 2006). Studies on the velocity profiles of QPs have demonstrated that they have a similar, if slightly slower, main sequence when compared with saccades made to visual targets (Garbutt et al. 2001, 2003, Kaminiarz et al. 2009).

There is a degree of variability intrinsic to the OKN mechanism that manifests itself in apparently random fluctuations of OKN parameters such as the SP duration, SP amplitude, QP amplitude and even SP velocity. The origin of this variability is not known. Significant nonlinear (and linear) predictability has been reported in the start and end positions of QPs, which indicates a deterministic component in the process. This has led to some speculation that while the apparently random fluctuations imply an embedded stochastic process, the variability may be a result of complicated deterministic behaviour manifesting as deterministic chaos (Shelhamer & Gross 1998, Shelhamer 1998, Trillenberg et al. 2001). A low fractional correlation dimension of the OKN waveform has been reported, ranging from 3.5 to 2.8 over a 120s period of OKN stimulation, implying that the system is chaotic (Shelhamer 1992, 1996). However, stochasticity affects the correlation dimension (Argyris et al. 1998), and a purely stochastic model of pathological nystagmus has demonstrated low fractional dimensions (Harris & Berry 2006). Finally, only a limited amount of predictability was found in SP velocity, and little or no predictability was found in SP duration and the amplitude of SPs and QPs, implying random behaviour.

Despite a highly variable gaze position from cycle to cycle, average position across

cycles remains steady, implying some adaptive mechanism or gross deterministic component that keeps gaze in the desired location. This indicates that SP and QP parameters must be dependent on each other. Dependency of parameters over time may indicate that working memory of previous OKN parameters is required to determine future OKN parameters, whereas short term dependencies may indicate that movements can be generated based solely on the current state of the system. Several relationships between OKN parameters within OKN cycles have been reported, such as a positive correlation between SP amplitude and SP velocity in humans viewing a random dot kinematogram (Watanabe et al. 1994), and a linear relationship between initial eye position and QP amplitude in turtles (Balaban & Ariel 1992). Differences in the parameters of human OKN can also be elicited by giving different instructions to a subject. Honrubia et al. (1968) found that instructing participants to maintain gaze on a region in space whilst viewing a large visual scene that moves produced a different type of OKN (stare OKN) than when the participant was asked to follow the moving bars (look OKN). Look OKN predominantly produces much larger amplitude SPs and QPs and generally has a higher gain than stare OKN.

There has yet to be a complete systematic dissection of the interelations between OKN parameters, but some headway has been made into understanding the threshold at which QPs are triggered through the testing of so-called "accumulator models".

2.8 Accumulator models.

Studies on the stochastic nature of OKN have often focused on the timing of QPs, and the development of models for QP generation that produce the characteristic positively skewed statistical distributions of SP durations (QP intervals). Several

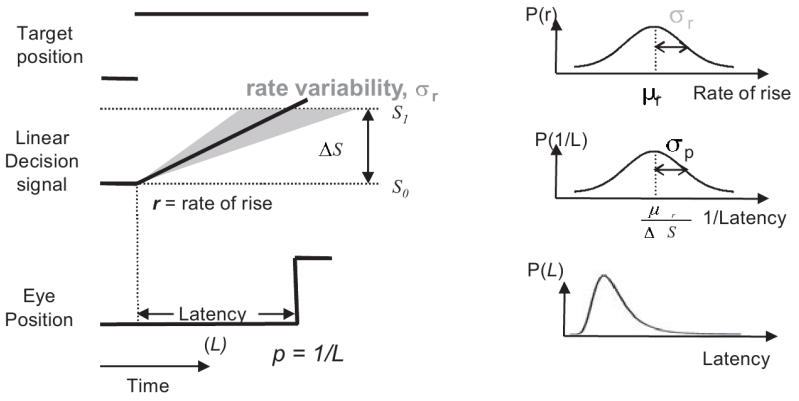

Figure 2.8: Illustration of the LATER model. In this model, when a visual target appears, a decision signal rises in the brain until it reaches threshold and triggers a response. The time it takes to reach threshold is the latency of the response. The rate of rise varies from trial to trial, and this variability is described by a normal distribution. Importantly, the rate remains constant within a trial. PDFs are illustrated on the right hand side of the figure for the rate of rise (top), reciprocal latency (middle), and latency (bottom). Figure taken from Harwood et al. (2008). Permission to reproduce this figure has been granted by The American Physiological Society.

processes have been suggested that fit the data with varying degrees of success. The most common are "accumulator models", that describe a decision signal or activation parameter in the brain that rises to threshold and triggers a QP once the threshold is reached. These models hypothesise different stochastic mechanisms by which the decision signal rises, and the variability in QP intervals is predicted to result from the random fluctuations in the model parameters. A series of distributions predicted by various statistical models of OKN were all tested to see which might best fit the histograms of QP intervals, with inconclusive results (Trillenberg et al. 2002) due to the amount of data required to discriminate between each distribution.

The linear approach to threshold with ergodic rate (LATER) model is one such accumulator model that has been used to predict a "recinormal" distribution of QP intervals (Carpenter 1993). The term recinormal is applied to mean that the

distribution of the reciprocal of the QP interval (QP rate) is normal. The LATER model predicts that a decision signal in the brain rises to threshold at a linear rate, and once it reaches the threshold a QP is triggered. The rate of rise of this signal varies from QP to QP with a normal distribution, but the rate remains constant during each QP, and the threshold also remains constant. As the rate at which the signal rises has a normal distribution, and the threshold is constant, the QP rate will also have a normal distribution. This leads to the latency of QPs having a recinormal distribution. This model is illustrated graphically in figure 2.8. However, while the LATER model does a surprisingly good job of modelling the distribution of QP intervals, an extension to the model was created in order to explain the particularly poor fit of very short duration QP intervals (Carpenter 1994). In this extension to the later model, a second "rogue" LATER unit is added in parallel with the first, and they both race to threshold. The first unit to reach threshold triggers a response and then the process resets and the race starts again. The mean rate of rise of the rogue unit is zero but as the standard deviation in the fluctuations between trials is so large it occasionally produces very high rates, hence very fast QPs. The LATER model has not only been used to model the latency of QPs, and is more often used to analyse the latency of saccades to visual target or manual reaction times (Carpenter & Williams 1995, Temel et al. 2009, Harwood et al. 2008).

Anastasio (1996) hypothesised that integration of noisy velocity signals by vestibular nucleus neurons would approximate a random walk with drift. Random walk models, or diffusion models, form one of the major classes of models analysed in reaction time experiments, particularly in two-choice decision tasks. The assumption is that a visual stimulus provides a certain amount of information that is accumulated over time towards one of two possible thresholds, that represent

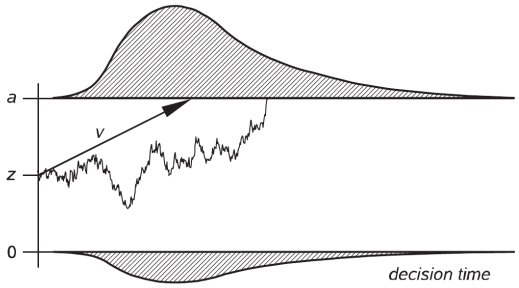

Figure 2.9: Illustration of the diffusion model. In this model, when a visual target appears, a decision signal in the brain rises and falls, in a random walk fashion, until it reaches a threshold and triggers a response. In the case presented here there are two different thresholds, one above and one below the inital level of the signal, representing a two-choice decision task where a different response is triggered depending on which threshold is reached first. The predicted distribution of the saccadic latency is depicted by the PDF illustrated with each threshold. Figure taken from Voss & Voss (2008). Permission to reproduce this figure has been granted by Elsevier.

two different responses, and when the signal reaches a threshold the respective response is activated (fig. 2.9). In the case where there is no variability in the drift rate between trials, and the drift rate is always positive (or zero), then the distribution of SP duration predicted by this model is an inverse Gaussian. Note that the major difference between this model and the LATER model is that the decision signal is allowed to vary during the course of each SP, although both assume a fixed threshold and an activation parameter rising to threshold. This is important if we consider that a normally distributed rate (as proposed by the LATER model) would eventually produce a rate with a negative value, leading to an infinitely long SP duration. As the activation parameter in the diffusion process is allowed to vary over the course of the SP in a random walk fashion, it will eventually cross the threshold, even if the drift rate is in the negative direction. Balaban & Ariel (1992) fitted experimental QP interval histograms with a lognormal distribution, that arises when a QP interval is the product of the previous QP interval and a normally distributed random variable. However, they

found that as the normally distributed random variable could take on negative values, the predicted QP intervals could also take on negative values, and it was necessary to truncate the variable at zero.

2.9 Concluding remark.

It is certainly evident that much is known about the physiology of how OKN is generated, and our understanding of the mechanisms that generate SP velocity have been greatly enhanced by the use of the control systems approach. However, the mechanisms that generate the QP component of OKN, or more specifically when the QP is triggered and to where they are targetted, is still poorly understood. The statistical approach to investigating the QP trigger threshold, and how it is reached, has had some limited success. However, differentiating between distributions created by these stochastic models requires a great deal of data. It seems that a systematic analysis of the relationships between OKN parameters, combined with an assessment of the statistical distributions of these parameters, would lead to a better understanding of these unresolved yet important mechanisms.

Chapter 3

Materials and Methods.

The results and conclusions presented in this thesis illustrate the results of statistical analysis applied to eye movement data recorded from normal human participants viewing optokinetic stimuli. Detailed information is given in each chapter on the specific analyses applied to the data. In order to avoid replicating the material and methods for collecting the eye movement data in each chapter, I present it here in full.

Three different experimental studies were performed in order to investigate how different spatiotemporal characteristics of the optokinetic stimuli affected eye movement parameters. The first study involved manipulating only the stimulus speed of the stimuli (Experiment 1: effect of stimulus speed). The second study involved manipulating both the spatial frequency and stimulus speed of square wave stimuli (Experiment 2: effect of stimulus spatiotemporal frequency). The third study involved manipulating the type of visual pattern used for the stimuli, comparing the use of a grid pattern with randomly allocated black and white squares, and a square wave pattern. The spatial frequency of the pattern was also manipulated in this experiment (Experiment 3: effect of stimulus pattern)

3.1 Participants.

28 participants were involved over the course of all experiments (18-35 years old, 14 female and 14 male). Participants were recruited through the departmental student participation scheme or known through the primary investigator. All participants were healthy, with no neurological, ophthalmological or vestibular impairments. Visual acuity for each eye was tested prior to each experimental session with a Snellen test (Keeler Snellen Test Type Panel (T O Y); Keeler Ltd., Windsor). All participants could read the 6/6 line with uncorrected vision. All protocols were approved by The University of Plymouth Faculty of Science Human Research Ethics Committee. A full description of the experimental procedures and an informed consent form were given to each participant before the investigation, and the option to withdraw from study at any time was explained prior to each procedure.

3.2 Eye tracking materials and procedure.

Each procedure required the participant to sit in a chair 1m from the middle of a flat white screen (viewable area of 1.57m by 1.17m landscape, subtending 76° by 60°) and the OKN stimulus was rear projected (EPSON EMP-500; Seiko Epson Corp., Japan) on to the screen with an SVGA resolution of 800 x 600 pixels. The participant's head was constrained using a chin rest. Eye movements were measured using a binocular head-mounted limbus tracker (Skalar IRIS Infrared Light Eye Tracker; Skalar Medical BV, Netherlands) that recorded horizontal movement with a resolution of 3 minarc at a sampling rate of 1kHz.

The eye tracker measures the reflection of infrared light from the recorded eye, at the edge between the white sclera and the darker iris. The head-mounted measuring device is an arrangement of two infrared light transducers. Each trans-

ducer consists of a light emitting array of nine infrared light emitting diodes and a detector array of nine infrared light sensitive phototransistors. Infrared light is projected onto the eye from the light emitting array and the detector array converts the reflected infrared light into voltage. The voltage of the nasally located phototransistors is subtracted from the voltage of the temporally located phototransistors, and the voltage difference is demodulated and amplified. The resultant signal is proportional to the angular deviation of the eye. A feature of the IRIS eye tracker is that the infrared light emission is chopped in order to minimise interference from ambient light.

The transducers are attached to a lightweight padded headband that can be adjusted independently in three perpendicular directions. The transducers can be orientated in such a way that either horizontal or vertical eye movements can be recorded from both eyes, or horizontal eye movements can be recorded from one eye and vertical eye movements recorded from the other eye. During this investigation only horizontal eye movements were recorded, in both eyes.

The analogue voltage output for both eyes was recorded on computer (vsgEyetrace v.3.0.beta software for Windows; Cambridge Research Systems, UK) and stored on the hard disk for off-line analysis.

3.3 Calibration.

Calibration of eye movement recordings were performed in two steps. Green LED bar arrays on the front panel of the control module indicate momentary eye position as a function of the output voltage when the headset is worn. Zero position was determined by presenting a stimulus at the centre of the screen and adjusting the "zero" control on the front panel until the moving green LED bar lined up with the marker for zero. The "gain" control on the front panel was then

adjusted such that when the participant was viewing a stimulus presented 30° to the left or right of the centre the moving green LED bar lined up with the markers at the far left and right of the LED display. This rough calibration was performed before each experiment in order to make sure that the recorded eye movements would remain linear in the range of the recording equipment, while maintaining a gain high enough for good quality recordings.

After this initial calibration a total of forty fixation stimuli were pseudorandomly presented along the horizontal meridian in the range 20° left of the centre to 20° right of the centre, each for a random period between 3 and 5 seconds, and eye movements were recorded. These calibration recordings were made before the experimental session and, in experiments 2 and 3, after the break at the midway point in the session. A calibration scale factor and offset were calculated offline after each experimental session, by performing linear regression of the mean analogue voltage recorded during each fixation against the known stimulus position. Individual fixation periods were visually checked before calibration for signs of corrective or anticipatory saccades, and if either were observed a portion of the fixation period was selected manually for calculating the mean analogue voltage. The calibration scale factor and offset were then applied to the data collected from the experimental session following the calibration recording.

3.4 OKN stimuli.

During OKN presentation participants were asked to stare straight ahead at the centre of the screen, rather than follow the moving stripes, in order to evoke stare-OKN rather than look-OKN. Participants were given a break for one minute between each trial in order to alleviate discomfort, tiredness, and to minimise the effects of any optokinetic after-nystagmus. In experiments 2 and 3, where there

were a large number of trials, participants were also given a break for five minutes on completion of the first half of the trials.

3.4.1 Experiment 1: effect of stimulus velocity.

10 healthy adults (6 female and 4 male), mean age 25 (SD = 5) years participated in experiment 1. The OKN stimulus consisted of a flat vertical square wave grating, comprised of alternating black and white vertical stripes (0.1cycles/°), moving horizontally at a fixed speed (10°/s, 20°/s, 30°/s, or 40°/s). Recording sessions were composed of a sequence of four trials each lasting 160s. The direction and stimulus speed was pseudorandomly selected for each trial such that each participant was presented with all four stimulus speeds. The stimuli were created using the standard vsgEyetrace software protocol for optokinetic nystagmus stimuli.

3.4.2 Experiment 2: effect of stimulus spatiotemporal frequency.

8 healthy adults (2 female and 6 male), mean age 26 (SD = 4) years participated in experiment 3. The OKN stimulus consisted of a flat vertical square wave grating, comprised of alternating black and white vertical stripes, moving horizontally at a fixed speed. Recording sessions were composed of a pseudorandom sequence of 12 trials, each with a different spatial frequency (0.05cyc/°, 0.1cyc/° or 0.2cyc/°), stimulus speed (10°/s or 30°/s), or direction (leftward or rightward) lasting 100s. The stimuli were created using the standard vsgEyetrace software protocol for optokinetic nystagmus stimuli.

3.4.3 Experiment 3: effect of stimulus type.

10 (6 female and 4 male), mean age 19 (SD = 1) years participated in experiment 3. In this experiment two different types of OKN stimulus patterns were used. One pattern consisted of a flat vertical square wave grating (fig. 3.1a). The second type of pattern consisted of a grid of randomly allocated black and white squares

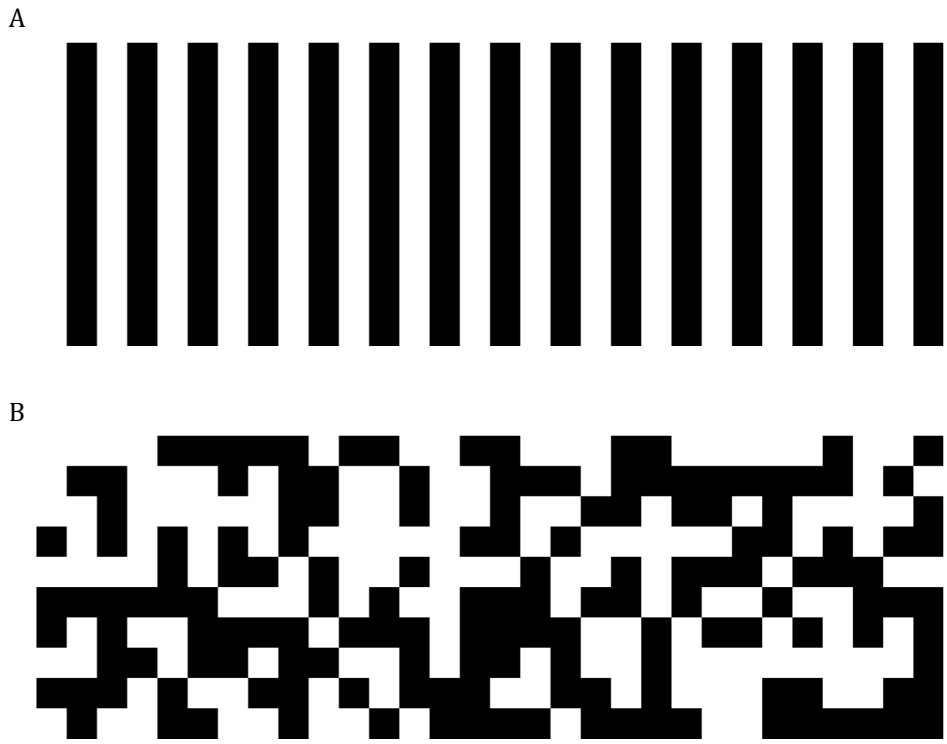

Figure 3.1: Examples of the two types of OKN stimulus patterns used: (A) a flat vertical square wave grating and (B) a "random grid" pattern. The random grid pattern in B is made up of squares that have a width and length equal to the width of the bars in A.

(fig. 3.1b). In order to compare the two different types of patterns the width (and thus the height) of squares was made to be exactly the same as the width of the vertical stripes in the respective square wave grating. Recording sessions were composed of a pseudorandom sequence of 12 trials, each with a different stimulus pattern (square wave or random grid), spatial frequency (0.083cyc/°, 0.165cyc/° or 0.248cyc/°), or direction (leftward or rightward) lasting 100s. The stimuli were created in Matlab (Matlab; Mathworks, USA) and exported as bitmap images. They were presented using the standard vsgEyetrace software protocol for scrolling picture stimuli at a stimulus speed of 31.5°/s.

The patterns were developed in Matlab by creating three matrixes with 600x1620 elements that were divided in to vertical subintervals containing m elements, where

$m = 20, 30$ or 60. The number of elements were specifically chosen in order that there would be an even number of subintervals in each matrix and that all the subintervals in a matrix could be the same width. For the square wave pattern, alternating vertical subintervals were allocated the numbers 0 then 1. Each matrix was then converted to a bitmap image using the built in Matlab function: *imwrite*, where the numbers 0 represent the colour black and the numbers 1 represent the colour white. In order to create the grid pattern each of the three 600x1620 element matrixes were further divided in to horizontal subintervals containing n elements, where $n = m$. A pseudorandom number generator: *rand* was used to generate a scalar value r from a uniform distribution on the unit scale for every square (mxn) subinterval. These numbers were then rounded up to 1 for $r \geq 0.5$ or down to 0 for $r < 0.5$. The process of assigning numbers to each square subinterval was repeated until an equal number of zeros and ones had been assigned to the matrix, and then the matrix was converted to a bitmap image using: *imwrite*.

3.5 Data acquisition.

All programs and algorithms for analysing data were developed and created in Matlab. Each eye was calibrated separately, and the average was computed to yield a "cyclopean" eye position in order to acquire a more accurate recording of eye position. Position traces of individual eyes were also compared for signs of vergence movements. At the higher stimulus speeds (30°/s and 40°/s) some participants would occasionally have divergent eye movements, but they were for brief periods (1 or 2s) and did not seem to occur in any systematic manner. Eye velocity was derived from the eye position using a central difference algorithm and then zero-phase digital filtering was performed, using an 8^{th}-order high pass Butterworth filter with a cut-off frequency of 80Hz. A Butterworth filter was used as it is maximally flat in the passband, and the cutoff frequency of 80Hz was chosen

based on previous studies investigating oculomotor tracking with the Skalar IRIS equipment that had used the same cutoff frequency (Barnes & Marsden 2002, Collins & Barnes 2006).

For the purposes of simplifying both analysis and presentation, data were normalised such that the stimulus direction was always positive. In this manner eye movement recordings for a rightward moving stimulus are unchanged but for a leftward moving stimulus are reversed. Thus SPs always had positive amplitude and velocity, whereas most QPs had negative amplitude and velocity. Each half of the stimulus was also described (relative to the stimulus centre) as positive or negative, such that a positive position is on the side which the stimulus is drifting towards and a negative position is on the side which the stimulus is drifting from. As an example, if the stimulus is drifting to the left, a position to the left of the stimulus centre would be positive and to the right would be negative.

Detecting QPs is a non-trivial process. SP Gain (stimulus speed/SP velocity) is often dependent on stimulus parameters, and varies in an apparently random fashion from trial to trial and also from SP to SP. A first pass of the data detected possible QPs when eye acceleration was greater than $1000^\circ/s^2$. Once a possible QP was detected the algorithm continued to search forwards in the data in order to find the peak velocity of the saccade. Peak velocity was calculated at the point where the eye velocity first began to decrease and then continuously decreased for a further period of 3ms. The possible QP was rejected if the value of eye velocity increased above $600^\circ/s$ before the peak velocity was determined as this was indicative of an anomaly caused by a blink rather than a real eye movement. The time at which peak velocity occured was indexed and the direction of the QP was determined from the sign of the peak velocity. Once the peak velocity had been calculated and indexed two more passes were made of the data, from the

time at which peak velocity occurred, forward and backward. The start and end of QPs made in the negative direction were determined when eye velocity returned to a value above 0°/s for more than 2ms. The start and end of QPs made in the positive direction were determined when eye velocity returned to a value under the stimulus velocity for a period longer than 2ms. In some trials when the SP gain was higher than 1, the eye velocity did not always return to a value under the stimulus velocity for a period longer than 2ms, and on visual inspection it was clear that some QPs were not detected. In these trials (30 of 256) a buffer zone of 5°/s was added to the threshold for detecting the start and end positions of QPs, and on visual inspection this buffer zone appeared to allow the QP detection algorithm to find all QPs. However, this criteria was not applied to all trials, as in typical OKN it would cause the detection algorithm to find the end and start points of QPs visibly earlier than they appeared to occur when examined by eye.

I defined a cycle of OKN as one SP and the following QP in series, and each cycle is numbered in chronological order. A matrix was created whose columns correspond to cycles and whose rows correspond to one of ten OKN variables recorded: SP start time, SP end time, SP duration, SP start position, SP end position, SP amplitude, SP velocity, QP amplitude, QP peak velocity, and QP duration. SP start times were determined to be the same as the end time of the previous QP, SP end times were determined to be the same as the start time of the following QP, and the SP duration was calculated from the SP start and end times. SP start position and end position were determined using the raw eye position data at the SP start and end times, and the SP amplitude was calculated from the SP start and end position. SP velocity was calculated by performing ordinary least squares linear regression on the raw positional data between the SP start and end times. QP amplitude was calculated from the end position of one SP to the start position

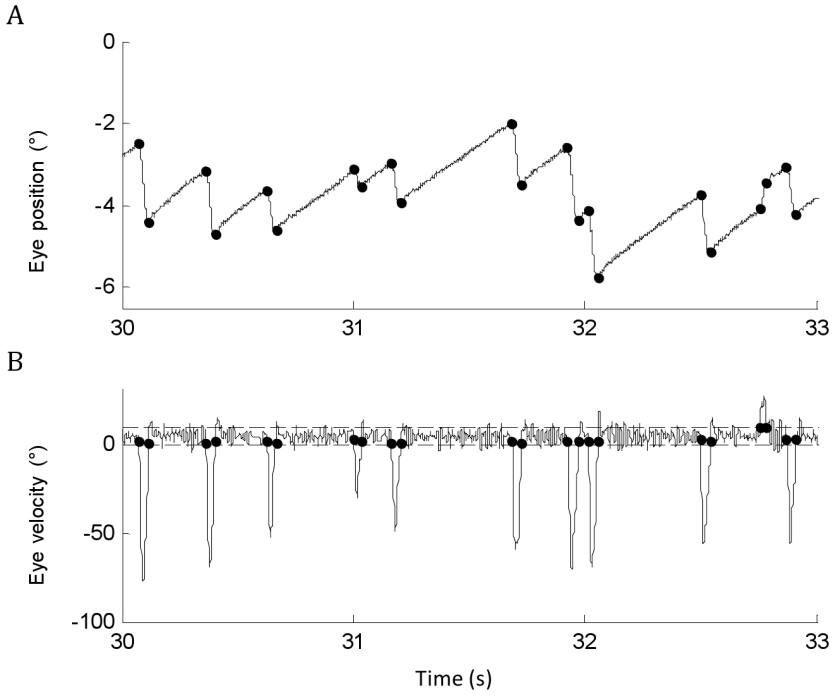

Figure 3.2: (A) Cyclopean eye position during 3 seconds of OKN stimulation. (B) Digitally filtered cyclopean eye velocity during the same period. Black dots represent the start and end position of QPs as determined by the QP detection algorithm without a buffer zone (i.e. for QPs made in the negative direction, when eye velocity returned to a value above 0°/s for more than 2ms, or for QPs made in the positive direction when eye velocity returned to a value under the stimulus velocity for a period longer than 2ms). Note the algorithm is capable of detecting very small QPs (at approximately 31s), small SPs (at approximately 32s), and QPs made in the direction opposite to stimulus motion (at approximately 32.7s). Upper dashed line in (B) represents stimulus speed (+10°/s), and lower dashed line in (B) represents 0°/s.

of the next SP. QP peak velocity was recorded during the process of determining the QP start and end positions explained above. QP duration was calculated from the end time of one SP to the start time of the next SP. All eye movements were then reviewed in a customised interactive graphical interface developed in Matlab. Blinks were detected manually and intervals containing blinks were marked and removed from analysis. After blinks were extracted, each trial contained m OKN cycles where m ranged from 49 to 536 (mean=227) for experiment 1, from 8 to 331 (mean=132) for experiment 2, and from 34 to 251 (mean=152) for experiment 3.

3.6 Data analysis.

RM-ANOVA was performed using the mean and standard deviation of OKN parameters as the dependent variables, and the stimulus parameters as the within-subjects factor, to determine the main and interaction effects of changing stimulus parameters on OKN parameters. These results are presented in chapter 4. The mean and standard deviation of SP velocity, SP gain, retinal slip (stimulus speed−SP velocity), SP amplitude, QP amplitude, QP start position, and SP start position were used as dependent variables. The within-subjects factors were stimulus speed, spatial frequency, direction and pattern type.

SP duration and QP rate (1/SP duration) were both skewed distributions with a heavy positive tail, so non-parametric tests were employed. A related-samples Friedman's two-way ANOVA by ranks test was performed, using either the median values of SP duration or the median QP rate as the dependent variable and stimulus speed as the within-subjects factor to test the main effect of stimulus speed in experiment 1. Performing F-tests on rank transformed data generally produces innacurate results for interaction effects. However, aligned rank trans-

formation can be performed on factorial nonparametric analyses and can handle RM-ANOVA to find both main and interaction effects (Higgins & Tashtoush 1994, Wobbrock et al. 2011). The median values of SP duration and QP rate were aligned and ranked using a software package *ARTool* (available by Wobbrock et al. at http://faculty.washington.edu/wobbrock/art/), and RM-ANOVA was carried out on each aligned data set in turn to find the significant main and interaction effects. The alignment process was used to estimate the main and interaction effects as marginal means and then all but one of these effects was stripped from the response variable to create an aligned variable. Ranks were assigned to the aligned variable and RM-ANOVA was performed on the ranked data, but only the effect for which the response was aligned was considered. This process was then systematically repeated until the response variable had been aligned and tested for all main and interaction effects.

In all RM-ANOVA tests where sphericity was violated, if the epsilon value was less than 0.75 the Greenhouse-Geisser correction was used to determine the corrected degrees of freedom (DOF), and if the epsilon value was greater than 0.75 the Huynh-Feldt correction was used. Bonferroni corrected confidence intervals were used to test significant differences in pairwise comparisons of values between stimulus conditions.

3.7 PCA.

PCA was performed on the OKN data in order to investigate the underlying components that explain the correlation of OKN parameters within each cycle of OKN, and the possible correlation of OKN parameters between cycles of OKN. These results are presented in chapter 5.

The method of PCA essentially passes a line through an n-dimensional cloud of

data points (through the centroid) and rotates the line to minimise the square of the orthogonal distance of each point to that line. This line is said to go through the maximum variation of the data and is termed the first "principal component" or "axis of variability". Successive components can be found by repeating this procedure to explain the maximum proportion of remaining variance, with the added constraint that it must be orthogonal to all preceding components. The total number of components that can be discovered in the cloud of data points is less than or equal to n. Each component is made up of a weighted linear sum of the original n parameters (the eigenvector) and explains a certain percentage of the variance in the data that is proportional to the eigenvalue of each eigenvector. The individual numbers in an eigenvector are called loadings (or component coefficients) and these represent the contribution of each of the original n dimensions to the principal components, scaled by the amount of variance explained by the component. An example illustrating the 2 principal components of a 2-dimensional data set are shown in figure 3.3, where the two dimensions are clearly correlated. The first component passes through the direction of maximum variance of the data (80%), and the second component is placed orthogonal to the first to explain the remaining variance in the data (20%).

Mathematically PCA is performed by eigenvalue decomposition, where the nxn input correlation (or covariance) matrix is diagonalised to yield n eigenvalues and n eigenvectors, which are linear sums of the n original dimensions. The eigenvectors of a real symmetric square matrix are the (non-zero) vectors that, after being multiplied by the matrix, remain parallel to the original vector. The corresponding eigenvalue is the factor by which the eigenvector is scaled. This concept is most easily represented graphically (fig. 3.4). The vector $\mathbf{x}$ is illustrated by an arrow with a length and direction in Cartesian coordinates. The vector $\mathbf{x}$

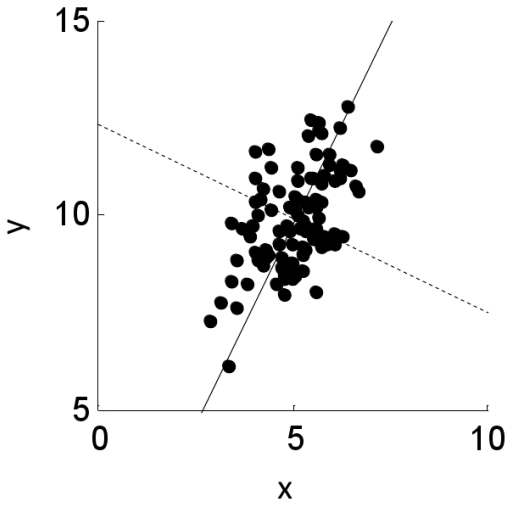

Figure 3.3: Scatter plot of 2-dimensional simulated data, with the two principal component overlayed on top. The first principal component passes through the direction of maximum variance of the data and is given by the formula PC1=0.44x+0.90y (solid line). The second principal component lies orthogonal to the first principal component and is given by the formula PC2=-0.90x+0.44y (dotted line).

is an eigenvector of the square (nxn) matrix $\mathbf{R}$ if multiplication by $\mathbf{R}$ causes the arrow to either remain the same, stretch or shrink, but to otherwise point in the exact same direction or in the exact opposite direction. The eigenvalue λ is the factor by which the arrow is stretched. Thus, the vector $\mathbf{x}$ is an eigenvector of the matrix $\mathbf{R}$ with the scalar eigenvalue λ, if and only if it satisfies the following equation:

$$\mathbf{Rx} = \lambda\mathbf{x}. \tag{3.1}$$

PCA was performed using the correlation matrices (rather than the covariance matrices) as variables measured on a different scale could influence the results of the analysis disproportionately (Blunch 2008). The correlation matrix $\mathbf{R}$ is

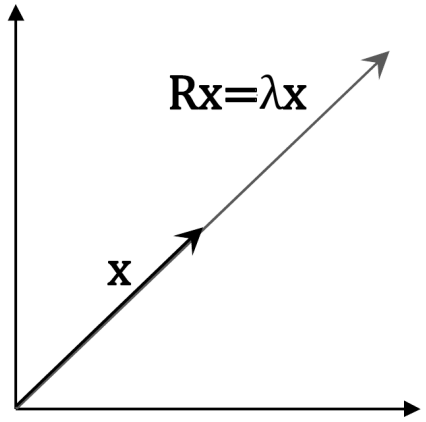

Figure 3.4: An example of an eigenvector illustrated as an arrow in Cartesian coordinates. The matrix $\mathbf{R}$ stretches the vector $\mathbf{x}$ but does not change its direction, therefore the vector $\mathbf{x}$ is an eigenvector of the matrix $\mathbf{R}$. The scalar λ is the eigenvalue of the eigenvector, and represents the factor by which the arrow is stretched.

diagonalised with the equation:

$$\mathbf{R} = \mathbf{V}^{\mathrm{T}}\mathbf{D}\mathbf{V} \tag{3.2}$$

where $\mathbf{D}$ is the diagonal matrix of eigenvalues of $\mathbf{R}$, $\mathbf{V}$ is the eigenvector matrix, and $\mathbf{V}^{\mathrm{T}}$ is the transpose of $\mathbf{V}$. Rearranging equation 3.2 and taking the square root of the eigenvalue matrix gives:

$$\mathbf{R} = (\mathbf{V}\sqrt{\mathbf{D}})(\mathbf{V}^{\mathrm{T}}\sqrt{\mathbf{D}}). \tag{3.3}$$

We let $\mathbf{V}\sqrt{\mathbf{D}} = \mathbf{A}$ and $\mathbf{V}^{\mathrm{T}}\sqrt{\mathbf{D}} = \mathbf{A}^{\mathrm{T}}$ to give:

$$\mathbf{R} = \mathbf{A}\mathbf{A}^{\mathrm{T}} \tag{3.4}$$

where $\mathbf{A}$ is the loading matrix and $\mathbf{A}^{\mathrm{T}}$ is its transpose, and each is a combination of the eigenvectors and the square roots of eigenvalues. The correlation matrix is the product of the loading matrix and its transpose, and the loading matrix represents the correlation between each component and each of the original dimensions. PCA was performed on the correlation matrix using the built in Matlab function: *pcacov*.

After extraction of the principal components it is sometimes preferable to discard components that only explain a small amount of variance before interpreting the results. There are several methods for determining how many components should be retained, but methods of best practice indicate that the use of a scree test is the best choice in order to obtain interpretable results (Costello & Osborne 2005). Scree tests involve plotting the eigenvalues of the principal components in descending order to create a scree plot. This is then examined for an inflection point, or "elbow". An elbow exists where the slope of the plot reaches some value close to zero, and indicates the point at which discarding the component associated with that eigenvalue (and all other smaller eigenvalues) would not result in discarding significant variance. An example scree plot is illustrated in fig. 3.5. The slope of the scree plot flattens off after the fourth component and the eigenvalues reach values less than 0.1, consequently explaining less than 2% of the variance in the system. If components 5-7 were discarded, the remaining four components would still explain 98% of the variance in the system. Details on the number of components retained, and the number discarded before further analysis, are

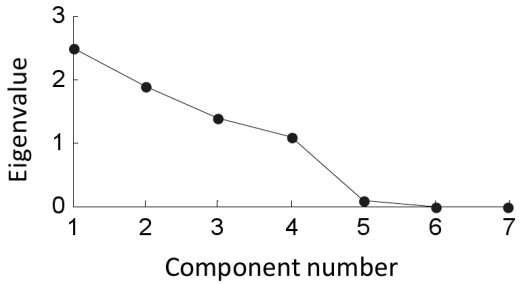

Figure 3.5: An example scree plot. An "elbow" in the plot after component 4 indicates there are 4 significant dimensions in the data.

explained in chapter 5 in relation to the maximum DOF in the system.

In my application factor rotation was performed on the retained components using the "varimax" strategy. Varimax rotation is a rotation of the component axes that maximises the sum of the variance of the squared loadings of a component. This results in components that have high loadings for a few variables and low loadings for the other variables, and aids in making the underlying structure of the loading patterns more clear. An example of rotation using the varimax strategy is illustrated in figure 3.6. In this example, rotation of the coordinate system brings the axes of the principal components closer to two clusters of loading values representing two components in the data. Oblique rotation strategies relax the restriction of orthogonal components and allow axes to pass through the centre of the clusters. Oblique rotations were tested as well as orthogonal rotations, as suggested by Costello & Osborne (2005), using the "promax" strategy, but there were no qualitative differences in the results. After factor rotation, similar loading patterns were observed among trials, and a heuristic algorithm was applied to sort the patterns into categories. The details and results of this sorting are explained in chapter 5.

There are a number of other possible methods that can be used to investigate the underlying components of correlated variables, or to investigate time series data

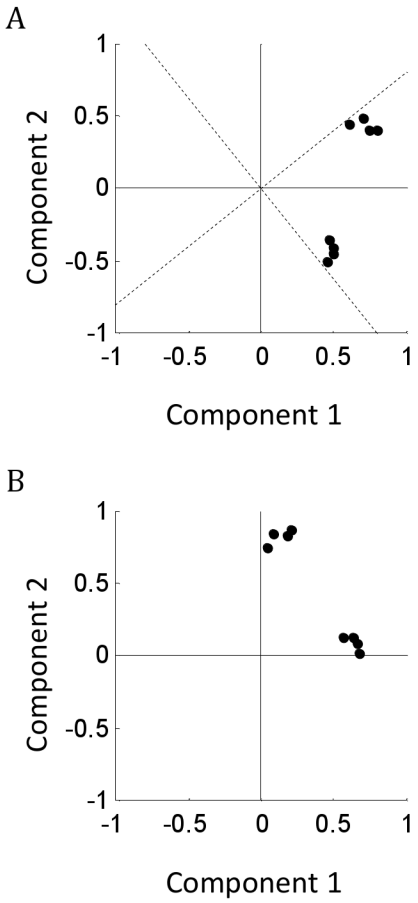

Figure 3.6: (A) Scatter plot of loading values from 8 original dimensions on the principal component axes (solid lines), with the new varimax rotated axes overlayed on top (dotted lines). (B) Scatter plot of loading values after varimax rotation. Loading values of the original dimensions are relatively higher on the new axis for one component and relatively lower on the other axis, allowing clusters of loadings to be more easily detected.

specifically, that were considered inappropriate for use in this analysis. These methods are Independent components analysis, factor analysis, and ARIMA modelling. I give a brief explanation below as to why these methods were not employed.

Independent component analysis is a similar method to PCA, popular in signal processing. However, performing independent components analysis finds the underlying components by maximising the statistical independence of the estimated components rather than maximising the proportion of explained variance. While this appears to make independent components analysis a more powerful tool than PCA, in order to identify the underlying components using this method no more than one of them may be Gaussian (Comon 1994). In fact the algorithms used to identify the underlying independent components rotate the n-dimensional data cloud such that it finds the least Gaussian component. The classical central limit theorem states that the mean of a sufficiently large number of independent random variables, each with a finite mean and variance, will be approximately normally distributed. The least Gaussian component extracted by independent components analysis is considered to be the most statistically independent, as a consequence of the central limit theorem. Independent component analysis is a particularly useful tool in blind source separation, such as in separating mixed EEG signals where neural networks are infrequently "on", so signals are often logistically distributed. However, it is inapplicable in the case where the underlying components actually represent multiple sources of Gaussian noise. Initial examination of the data revealed that many OKN parameters had an approximately Gaussian distribution. This is corroborated in chapter 6 where we illustrate that after detrending OKN parameters, using the model derived in chapter 5, the residuals are approximately Gaussian.

Factor analysis is another technique similar to PCA. The aim of factor analysis is

to reveal those factors that cause the observed parameters to covary. The analysis estimates the "specific variance" and "error variance" that is unique to each parameter, and includes only the estimated shared variance (the communalities) in the calculation of the factors. PCA does not assume an underlying causal structure, it is a mathematical transformation typically used for feature extraction or as an objective method of dimension reduction. It is not a statistical model and does not estimate any unique variance of the observed parameters, instead it includes the total variance of parameters in the calculation of the principal components (Blunch 2008). The PCA method of extraction was chosen in order to conserve as much information as possible without making any statistical assumptions about the data.

The Box-Jenkins (ARIMA - autoregressive integrated moving average models) approach is a common method used to analyse time series data and evaluate models based on their forecasting ability. However, these techniques assume that values in the time series occur at equally spaced intervals, do not contain missing data, and the variance in the fluctuations over time are constant (Box & Jenkins 1976). Whereas there is a remarkable amount of variability in the timing of both SPs and QPs, blinks during eye movement recordings create variable gaps in the time series, and there is known to be some signal dependent noise in the ocular motor system (e.g. in the amplitude of saccades). Although ARIMA methods allow for statistical testing of the validity of time series models, the statistical significance can be interpreted incorrectly, and sometimes cannot be computed at all, when the assumptions of the model are not met. However, methods specified for identifying the order of an autoregressive or moving average model, such as examining the autocorrelation function and partial autocorrelation function of parameters for significant values were used in this analysis.

The autocorrelation function describes the correlation between the same parameter at different points in time, as a function of the time difference. Given a time series x_t, the partial autocorrelation of lag k, is the autocorrelation between x_t and x_{t+k} that is not already accounted for by lags 1 up to and including $k-1$. The sample autocorrelation function was calculated in this analysis using the built in Matlab function: *corrcoeff*. The sample partial auto-correlation function was calculated in SPSS 19.0 from the menu Analyze > Forecasting > Autocorrelations.

When identifying the order of an autoregressive or moving average model, it is recommended that the auto-correlation function and the partial auto-correlation function should be examined for certain patterns. Exponential decay of the auto-correlation function to zero is one indication that the data can be modelled with an autoregressive model rather than a moving average model. The partial auto-correlation function should then be examined to determine whether there is a lag at which there is a sharp cutoff in the function, or the function becomes non-significant. A sharp cutoff at a lag of one in the partial autocorrelation function is an indication that the time series data follow a 1^{st} order autoregressive model. A 1^{st} order autoregressive model explains the change in x_t at each cycle as a function of x_{t-1} plus some unexplained noise. Whereas, a 1^{st} order moving average model explains the change in x_t at each cycle as a function of only the error term (or unexplained noise) in the previous cycle, rather than as a function of x_{t-1}. These techniques were used in conjunction with the results of PCA in order to interpret the most likely update dynamics of OKN parameters between cycles.

Chapter 4

Results of experiment 1: description of OKN parameters and stimulus speed effects.

In this chapter the general characteristics of the distribution of OKN parameters is presented alongside the results of RM-ANOVA performed on OKN parameters recorded from experiment 1. The effect of stimulus speed, and the remarkable amount of variability in OKN parameters between and within trials is illustrated, and important differences in the distributions of OKN parameters are highlighted.

4.1 SP velocity, retinal slip, and gain.

The most notable characteristic of SP velocity recorded in experiments 1-3 is the remarkable amount of apparently random fluctuations, not just between trials, but also between SPs. The distribution of SP velocity within a trial is unimodal and has a mean value, but the SP velocity during any given SP can be much less than, or much more than, the mean value (fig. 4.1a and 4.1b). The coefficient of variation (CV; standard deviation/mean) of SP velocity during each trial ranged from 0.08 to 1.34 (mean=0.32). To put this in context, if we assume a normal distribution, a CV of SP velocity greater than 0.5 would indicate that there is approximately

a 2% chance for SP velocity to drop below 0°/s and actually drive SPs in the opposite direction to stimulus motion. If the CV of SP velocity was greater than 1 the chance of SP velocity dropping below 0°/s would be approximately 14%. Despite the relatively high values for the CV of SP velocity, no SPs were observed in the direction opposite to stimulus motion, although some SPs did reach very low values close to zero. However, the SP velocity did occasionally exceed the stimulus speed, producing negative values of retinal slip (fig. 4.1c) and values of gain higher than 1 (fig. 4.1e). A representative example of an eye movement trace, with a corresponding plot of the SP velocity is illustrated in figure 4.2 containing 10 seconds of OKN eye movement data from one trial. In this example, the mean SP velocity is much lower than the stimulus speed (mean SP gain=0.38) and SP velocity appears to change in a haphazard fashion from one SP to the next. In general the mean SP gain was less than 1 during each trial (mean=0.74), and tended to decrease with stimulus speed as presented in the results of RM-ANOVA below.

4.1.1 Effect of stimulus speed.

The mean SP velocity, retinal slip, and gain were calculated for each presentation and figure 4.3 contains three boxplots demonstrating how these parameters were affected by stimulus speed. A one-way RM-ANOVA was performed using each parameter as the dependant variable and stimulus speed as the within-subjects factor to determine the main effects of changing stimulus speed.

There was a significant main effect of stimulus speed on the mean SP velocity ($F(1.37,12.3)=18.2$, $p=0.001$), retinal slip ($F(1.37,12.3)=39.4$, $p<0.001$) and SP gain ($F(3,27)=15.4$, $p<0.001$). Bonferroni corrected confidence intervals for multiple comparisons indicated there was a significant difference in the mean values of SP velocity between presentations with a stimulus speed of 10°/s and 20°/s

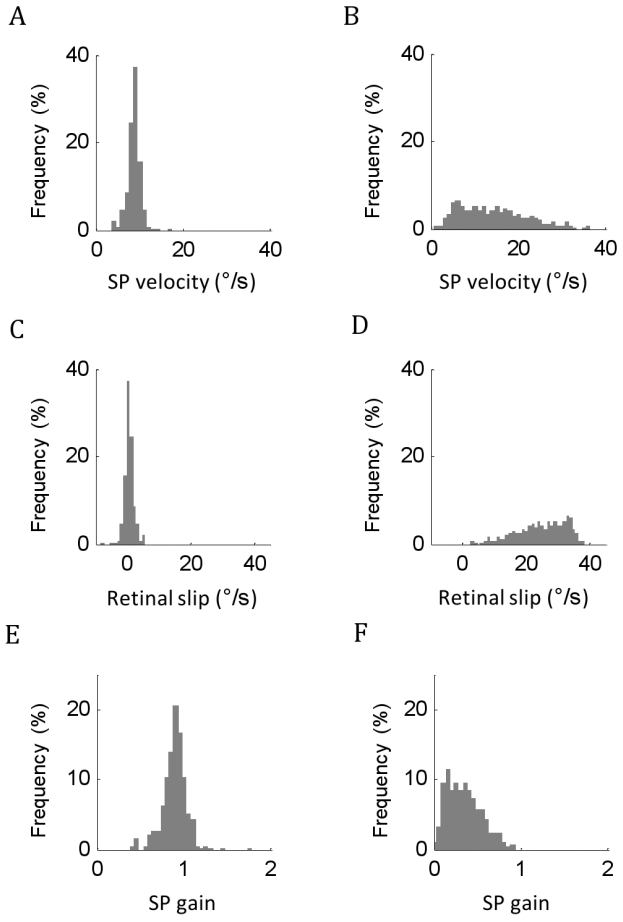

Figure 4.1: (A and B) Distribution of SP velocity. (C and D) Distribution of retinal slip. (E and F) Distribution of SP gain. Data taken from participant 1. Left panel, stimulus speed 10°/s; Right panel, stimulus speed 40°/s.

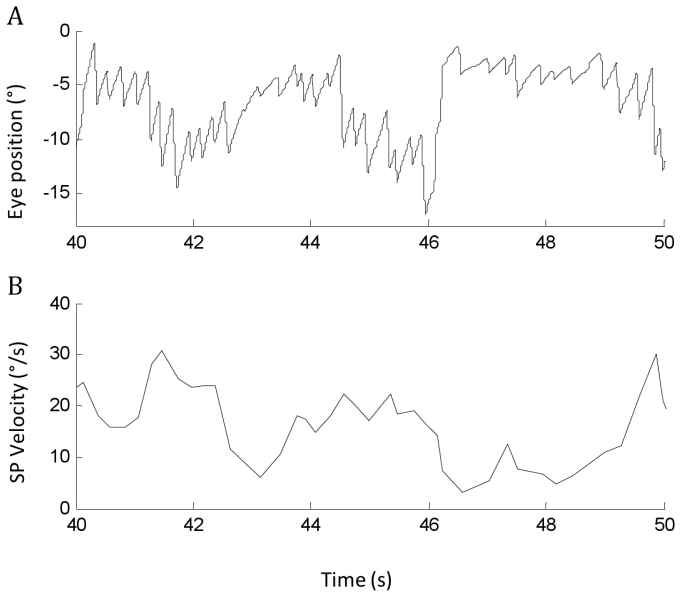

Figure 4.2: (A) Example eye movement trace and (B) a plot of the corresponding average eye velocity during each SP. Data taken from participant 1, stimulus speed 40°/s. Mean SP velocity=15°/s. SP velocity CV= 0.56.

(p<0.001), 10°/s and 30°/s (p<0.001), and 10°and 40°/s (p=0.006). Significant differences in mean retinal slip were found between all presentations, between 10°/s and 20°/s (p=0.005), 10°/s and 30°/s (p<0.001), 10°/s and 40°/s (p<0.001), 20°/s and 30°/s (p=0.007), 20°/s and 40°/s (p=0.001), and 30°/s and 40°/s (p=0.003). Significant differences in the mean values of SP gain were found between presentations with a stimulus speed of 10°/s and 30°/s (p=0.007), and 10°/s and 40°/s (p=0.005).

One-way repeated measures ANOVA performed on the standard deviation of these parameters found a main effect of stimulus speed on the standard deviation of SP velocity and retinal slip (F(1.61,14.5)=41.3, p<0.001), but not SP gain (F(1.75,15.7)=1.0, p=0.385). Bonferroni corrected confidence intervals for multiple comparisons indicated there was a significant difference in the standard deviation of SP velocity and retinal slip, between presentations with a stimulus

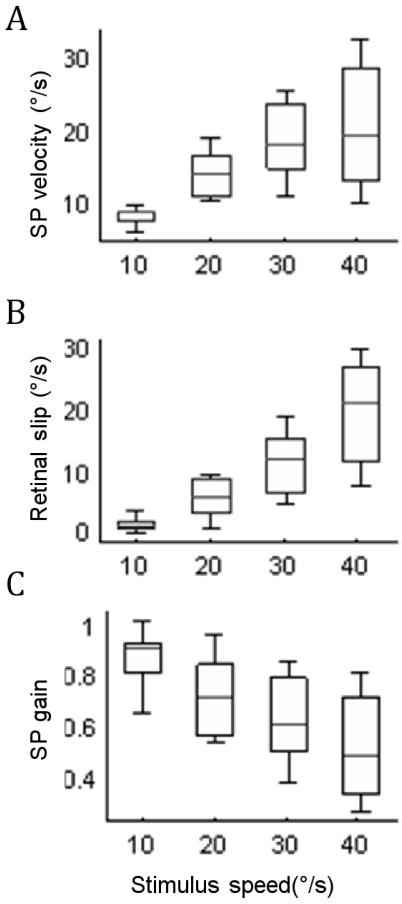

Figure 4.3: Boxplot illustrating the effect of stimulus speed on (A) SP velocity, (B) retinal slip, and (C) SP gain. Each box represents the parameter grouped across all participants for each stimulus condition. Stimulus conditions tested were: stimulus speed, 10°/s, 20°/s, 30°/s or 40°/s. Box represents interquartile range, whiskers represent maximum and minimum data values (excluding outliers), and dots represent outliers found more than 1.5 times the interquartile range from the ends of the box.

speed of 10°/s and 20°/s (p=0.001), 10°/s and 30°/ (p<0.001), 10°/s and 40°/s (p<0.001), 20°/s and 30°/s (p=0.014), 20°/s and 40°/s (p=0.006).

In summary, increasing stimulus speed resulted in an increase in the mean SP velocity and retinal slip that saturated at higher stimulus speeds, and a decrease in SP gain. Increasing stimulus speed also resulted in an increase in the standard deviation of SP velocity and retinal slip, but had no effect on the standard deviation of SP gain. The mean value of the standard deviation of SP gain across all trials was 0.2.

4.2 SP amplitude, QP amplitude, and the start and end points of QPs.

QPs were predominantly made in the negative direction (opposite to stimulus motion), although there were a small population of QPs that were made in the direction of stimulus motion, and these positive amplitude QPs made up 5% of the total population of QPs across all trials. The shape of the QP distribution sometimes appeared bimodal with a dip around 0° that extended approximately 0.5-0.75° in both directions. Manual examination of the eye movement records verified that these dips were real, and not an artefact of the QP detection algorithm, although we cannot rule out the possibility that very small microsaccades were made that were not possible to detect. When the CV of the QP amplitude was smaller than 0.35 (167 of 256 trials) there were rarely any QPs made in the positive direction at all, so this "dead zone" often could not be observed and the distribution appeared unimodal. Also, when the dead zone could be observed it was clear that the cutoff was not absolute as some very small amplitude QPs would occasionally occur. Figures 4.4a and 4.4b illustrate representative examples of the distribution of QP amplitude from two participants viewing an OKN stimulus at

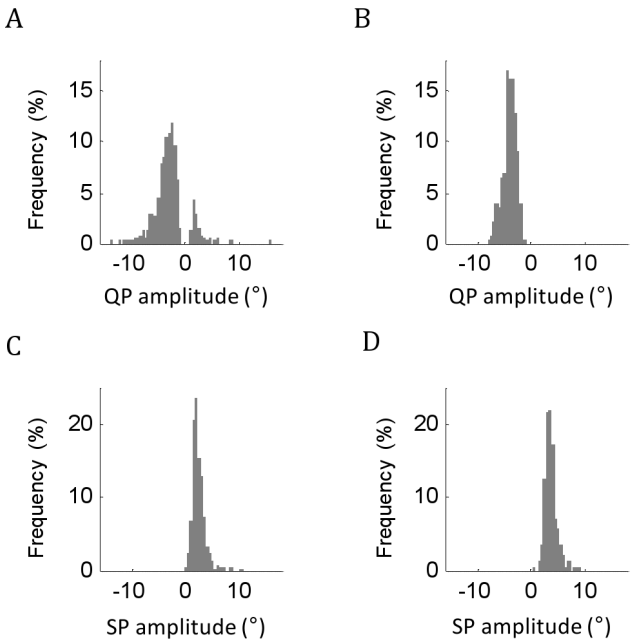

Figure 4.4: (A and B) Distribution of QP amplitude. (C and D) Distribution of SP amplitude. Left panel: data collected from the same trial, participant 1, stimulus speed 10°/s; Right panel: data collected from the same trial, participant 6, stimulus speed 10°/s. Note the "dead zone" around 0° in (A).

the same stimulus speed (10°/s), that had very different values for the CV of QP amplitude, and in figures 4.4c and 4.4d the distribution of SP amplitude from the same trials are shown. It was clear that mean QP amplitude and mean SP amplitude had a strong tendency to compensate for each other (i.e. mean values of magnitude were equal, but were made in the opposite direction), but that the distributions of SP and QP amplitudes could be quite different. The SP amplitude distribution was always unimodal, and often slightly positively skewed (mean skewness=0.88), particularly at low stimulus speeds. SPs were always made in the direction of stimulus motion and no negative SPs were found in the data. The standard deviation of SP amplitude was usually smaller than the standard deviation of QP amplitude (237 of 256 trials).

Over the course of a trial, QPs tend to compensate for the eye displacement due

to SPs, in order to stop the average eye position from drifting towards the limit of gaze. However, from figures 4.4c and 4.4d it appears that QPs are generally triggered when the eye position is close to the centre of the stimulus, and they are driven into the negative direction of the stimulus field. We define contraversion as the mean eye position over the course of a trial and found that it was on average 3° in the negative direction, with no significant stimulus velocity effect (F=0.6, p=0.50). An increase in contraversion has been observed with increasing stimulus speed when elicited with a rotating full-field patterned curtain (Garbutt et al. 2002). An increase in contraversion has also been observed during periods of perceived self-motion (circular vection) that occur during rotational OKN (Thilo et al. 2000). The lack of a significant stimulus speed effect on contraversion in the data could be the result of using a translational rather than rotational stimulus and the lack of any strong linear vection.

4.2.1 Effect of stimulus speed.

The mean SP amplitude, QP amplitude, and the start and end points of QPs were calculated for each presentation and boxplots were plotted for each of these parameters (fig. 4.5) illustrating how these parameters were affected by stimulus speed. A one-way RM-ANOVA was performed using each parameter as the dependent variable and stimulus speed as the within-subjects factor.

There was a significant main effect of stimulus speed on the mean SP amplitude ($F(3,27)=13.1$, $p<0.001$) and QP amplitude ($F(3,27)=14.9$, $p<0.001$), but not on QP start position ($F(2.27,20.4)=1.1$, $p=0.356$) nor QP end position ($F(1.69,15.2)=1.1$, $p=0.345$). Bonferroni corrected confidence intervals for multiple comparisons indicated there was a significant difference in the mean values of SP amplitude between presentations with a stimulus speed of 10°/s and 20°/s ($p=0.041$), 10°/s and 30°/s ($p=0.005$), and 10°/s and 40°/s ($p=0.007$). There were signifi-

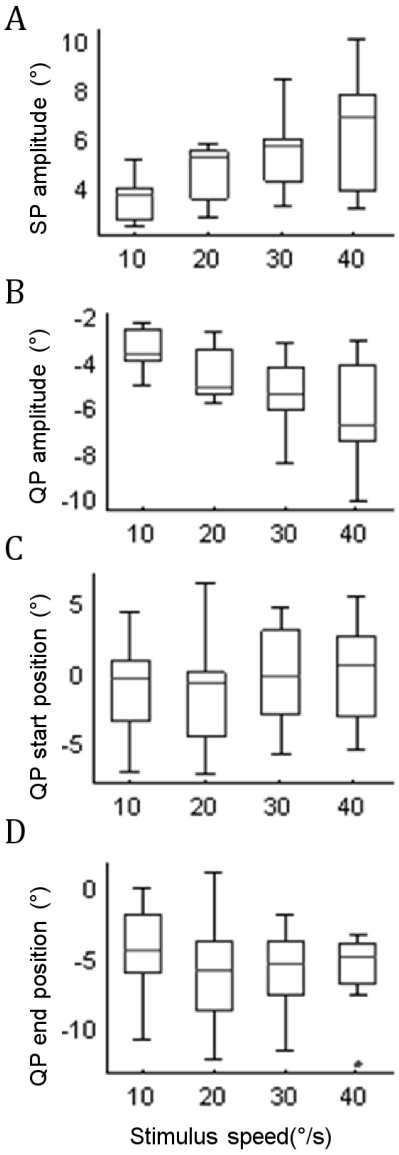

Figure 4.5: . Boxplot illustrating the effect of stimulus speed on (A) SP amplitude, (B) QP amplitude, (C) QP start position, and (D) QP end position. Each box represents the parameter grouped across all participants for each stimulus condition. Stimulus conditions tested were: stimulus speed, 10°/s, 20°/s, 30°/s or 40°/s. Box represents interquartile range, whiskers represent maximum and minimum data values (excluding outliers), and dots represent outliers found more than 1.5 times the interquartile range from the ends of the box.

cant differences in the mean value of QP amplitude between presentations with a stimulus speed of 10°/s and 30°/s (p=0.007), and 10°/s and 40°/s (p=0.004).

One-way RM-ANOVA on the standard deviations of these parameters found a main effect of stimulus speed on the standard deviation of SP amplitude (F(3,27) =12.0, p<0.001), QP amplitude (F(3,27)=5.7, p=0.004), QP start position (F(3, 27)=4.9, p=0.007), and QP end position (F(3,27)=6.7, p=0.002). Bonferroni correctd confidence intervals for multiple comparisons indicated there were significant differences in the standard deviation of SP amplitude between presentations with a stimulus speed of 10°/s and 30°/s (p=0.049), 10°/s and 40°/s (p=0.011), and 20°/s and 40°/s (p=0.028). Significant differences in the standard deviation of QP amplitude (p=0.031), QP start position (p=0.018), and QP end position (p=0.009) were only found between presentations with a stimulus speed of 10°/s and 40°/s.

In summary, increasing stimulus speed resulted in an increase in the mean SP and QP magnitude, but had no effect on the mean start and end position of QPs (and SPs). Increasing stimulus speed also resulted in an increase in the standard deviation of the mean SP and QP magnitude, and the start and end position of QPs (and SPs).

4.3 SP duration and QP rate.

The distribution of SP duration was almost always positively skewed (mean skewness=1.7). Lilliefors tests conducted on all 256 trials (corrected using the Holm-Bonferroni method) showed that 207 of 256 histograms were significantly different to Gaussian. A reciprocal transformation has been cited as a method for transforming the SP duration to a Gaussian distribution (Carpenter 1993). However, the distribution of the reciprocal transformation of SP duration (which we will

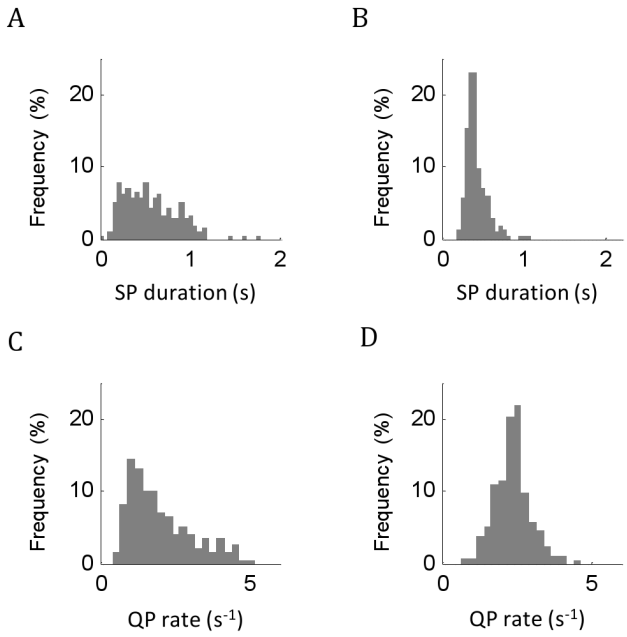

Figure 4.6: (A and B) Distribution of SP duration. (C and D) Distribution of QP rate. Left panel, data collected from the same trial: participant 2, stimulus speed 10°/s; Right panel, data collected from the same trial: participant 6, stimulus speed 10°/s.

term QP rate) is often more positively skewed than the distribution of SP duration itself (mean skewness=3.7). Lilliefors tests conducted on all 256 trials (corrected using the Holm-Bonferroni method) showed that 212 of 256 histograms were significantly different to Gaussian. The distribution of SP duration and QP rate for two trials, from two different participants at the same stimulus speed, are illustrated in figure 4.6. One participant had a number of particularly short duration SPs and a highly skewed distribution of QP rate, the other participant did not demonstrate any SPs under 200ms and the corresponding distribution of QP rate was approximately symmetrical.

Care must be taken when dealing with positively skewed distributions that have particularly heavy tails. One consideration is that if the tails of the PDF decay slowly it may not have moments (mean, variance, skewness, kurtosis) that converge to a finite value. Although sample moments can be calculated from data

sampled from such a distribution, they are really erroneous values if the distribution of the parent population does not have finite moments. Measures of the sample moments, such as the measures of skewness presented above, should be considered with a deal of caution, and it is more useful to describe the population using non-parametric measures such as the median and the interquartile range. Also, it is usually best to use non-parametric methods when performing statistical tests on data from such distributions, as non-parametric methods do not make assumptions about the asymptotic tail behaviour of the model to be tested in order to produce accurate results.

4.3.1 Effect of stimulus speed.

The median SP duration and QP rate were calculated for each presentation and figure 4.7 contains two boxplots demonstrating how these parameters are affected by stimulus speed. A related-samples Friedman's two-way ANOVA by ranks was performed using the rank-transformed median values of each parameter as the dependent variable and stimulus speed as the within-subjects factor.

There was a significant main effect of stimulus speed on SP duration and QP rate ($\chi^2(3)=14.5,\ p=0.002$). Bonferroni corrected confidence intervals for multiple comparisons indicated there was a significant difference in SP duration and QP rate between presentations with a stimulus speed of 10°/s and 30°/s ($p=0.006$), and 10°/s and 40°/s ($p=0.006$). Increasing stimulus speed resulted in a decrease in SP duration with the mean asymptote at approximately 250ms for speeds greater than 20°/s, and an increase in QP rate with the mean asymptote at approximately 4Hz.

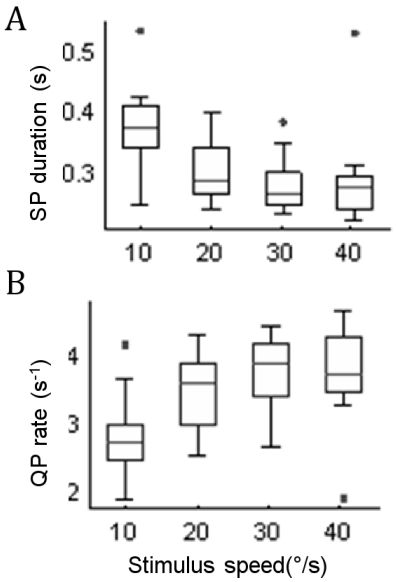

Figure 4.7: Boxplot illustrating the effect of stimulus speed on (A) SP duration and (B) QP rate. Each box represents the parameter grouped across all participants for each stimulus condition. Stimulus conditions tested were: stimulus speed, 10°/s, 20°/s, 30°/s or 40°/s. Box represents interquartile range, whiskers represent maximum and minimum data values (excluding outliers), and dots represent outliers found more than 1.5 times the interquartile range from the ends of the box.

4.4 OKN cycles

In the next section we will examine the correlations between OKN parameters, but before we do so it useful to introduce the concept of an OKN cycle. We will define an OKN cycle to be one SP and the following QP in series. The concept of investigating OKN behaviour in terms of cycles is not common in the eye movement literature. As discussed in chapter 2 it is much more common to investigate OKN using a classical systems engineering approach. However, investigating how OKN parameters change and interrelate within a single cycle of OKN makes sense when we consider that within each and every cycle of OKN there are three decisions that must always be made. These three decisions require determining three values:

SP velocity. After each QP the system must settle at the new SP velocity. The common assumption is that SP velocity minimises retinal slip, so we might assume that SP velocity will always be equal to stimulus speed. However, we have found that the mean SP velocity of a given trial is often lower than stimulus speed, and in individual cycles the SP velocity often reaches values much lower or higher than the mean SP velocity. This indicates that the system is either limited, remarkably variable, or that the system is purposefully generating a SP velocity below stimulus speed.

SP end position. The system must determine at what point in time or space to end the SP and trigger a QP. At some point the system must generate a QP or the eye will reach the limit of the orbit and remain there. However, we have found that QPs are made when the eye position is directed approximately towards the centre of the OKN stimulus. This indicates that there must be some trigger that causes the eyes to shift gaze long before it

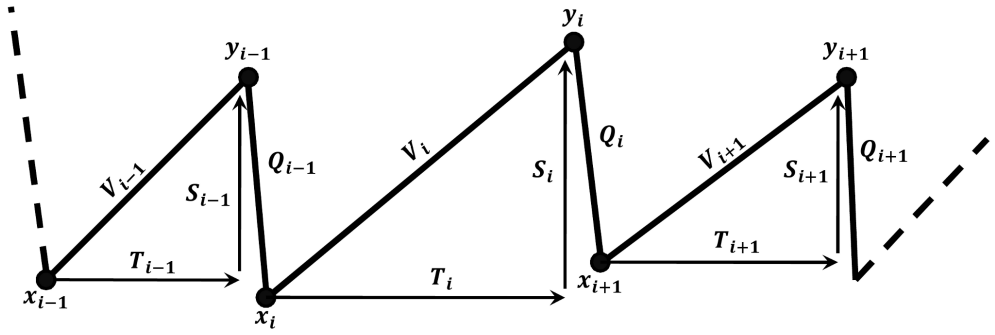

Figure 4.8: Representation of OKN and its parameters. The i^{th} cycle in a series contains one SP followed by one QP and is defined by six characteristic parameters; SP start position (x_i), SP amplitude (S_i), average SP velocity (V_i), SP duration (T_i), QP start position (y_i), and QP amplitude (Q_i).

becomes necessary due to reaching the limit of gaze.

QP end position. The system must determine where to target the QP towards in order to maintain a stable gaze. When participants are asked to stare directly at the centre of the screen we might expect the end position of QPs to be directed straight ahead to the centre of the stimulus. However, we have found that QPs are generally targetted into the negative field of view. This indicates that there is some target position, or area, within the negative field that the eyes are directed towards during OKN.

Figure 4.8 illustrates three cycles of OKN in series where the central cycle is the ith cycle, the cycle before it is the i-1 cycle, and the cycle after is the i+1 cycle. In this illustration each cycle is characterised by the same 6 OKN parameters (x_i, S_i, V_i, T_i, y_i and Q_i), so we can define each of these cycles with a vector containing just 6 parameters and examine the relationships between and within these vectors. It is worth considering how OKN parameters may be sequentially dependent on other OKN parameters, and we note that direct causal relationships between three or more OKN parameters will manifest as many "indirect" relationships. As

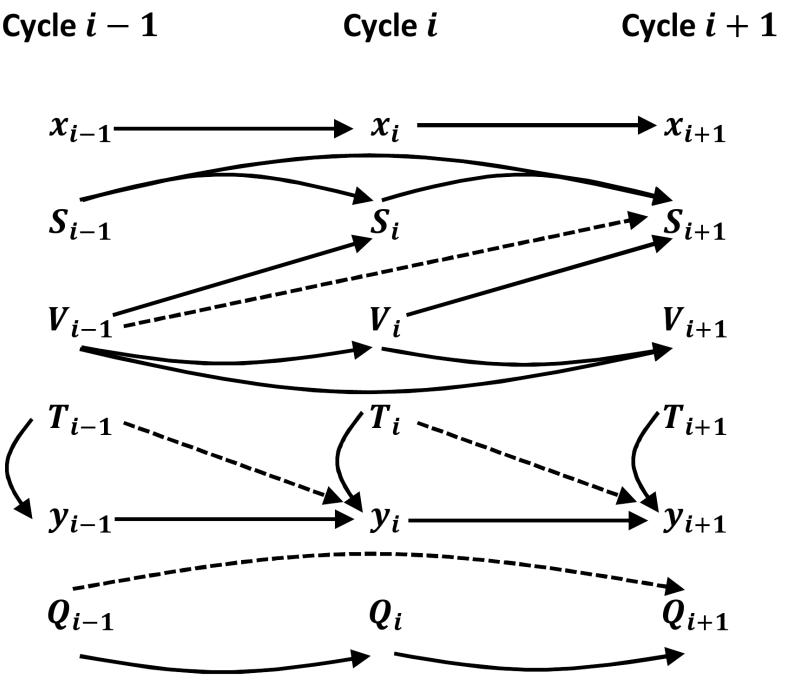

Figure 4.9: Representation of possible direct and indirect relationships between parameters and cycles of OKN. Direct correlations are represented by solid arrows, and the resultant indirect correlations are represented by dotted arrows. Direct correlations can occur within cycles (e.g. T_i with y_i), across one cycle (e.g. x_i with x_{i+1}), or across multiple cycles (e.g. S_{i-1} with S_{i+1}). Indirect correlations are caused by a number of direct correlations in series and can also occur within cycles, across one cycle, or across multiple cycles. For example, Q_{i-1} directly correlates with Q_i, and Q_i directly correlates with Q_{i+1}, causing an indirect correlation between Q_{i-1} and Q_{i+1}.

more cycles and more direct relationships are considered the number of indirect relationships grows such that we will find correlations between many (if not all) of the OKN parameters (fig. 4.9). We applied PCA to the correlation matrices of OKN parameters in order to determine the underlying structure of the correlations between OKN parameters within and between cycles of OKN.

4.5 Results of investigating the correlation of OKN parameters between and within cycles.

Multiple pair-wise scatter plots illustrated that most relationships were linear, except for the hyperbolic relationship between SP velocity and SP duration and a number of parameters that did not seem to be related at all. The QP "dead zone" around zero degrees was also clearly visible when the coefficient of variation of the QP amplitude was greater than 0.5. Plotting the same pair-wise correlations for all 40 trials revealed some general similarities but often some clear differences as well. This implied that there was some underlying structure to the relationships between parameters, but that there might be specific variables or relationships that change between participants or stimulus conditions. An example of the pair-wise scatter plots is shown in figure 4.10, where a variety of strong, weak, positive and negative correlations are evident. Note that in this data set the coefficient of variation of the QP amplitude is only 0.34, so the distribution of QP amplitude does not cross zero and the QP "dead zone" cannot be observed.

The correlation matrices of OKN parameters were examined in order to determine if there were likely to be significant correlations between parameters. Three examples are shown in fig. 4.11 that illustrate clear differences in the correlation matrices between trials with different stimulus speeds, and between participants.

The sample auto-correlation function of each parameter was examined in order to determine if there were likely to be relationships that persisted over a number of cycles. A typical example of the auto-correlograms of OKN parameters is shown in fig. 4.12, where the sample auto-correlation is illustrated as a function of the integer cycle iterations (or time "lags"). SP velocity, SP start position, and QP start position have some correlation with previous cycles, sometimes extending back as

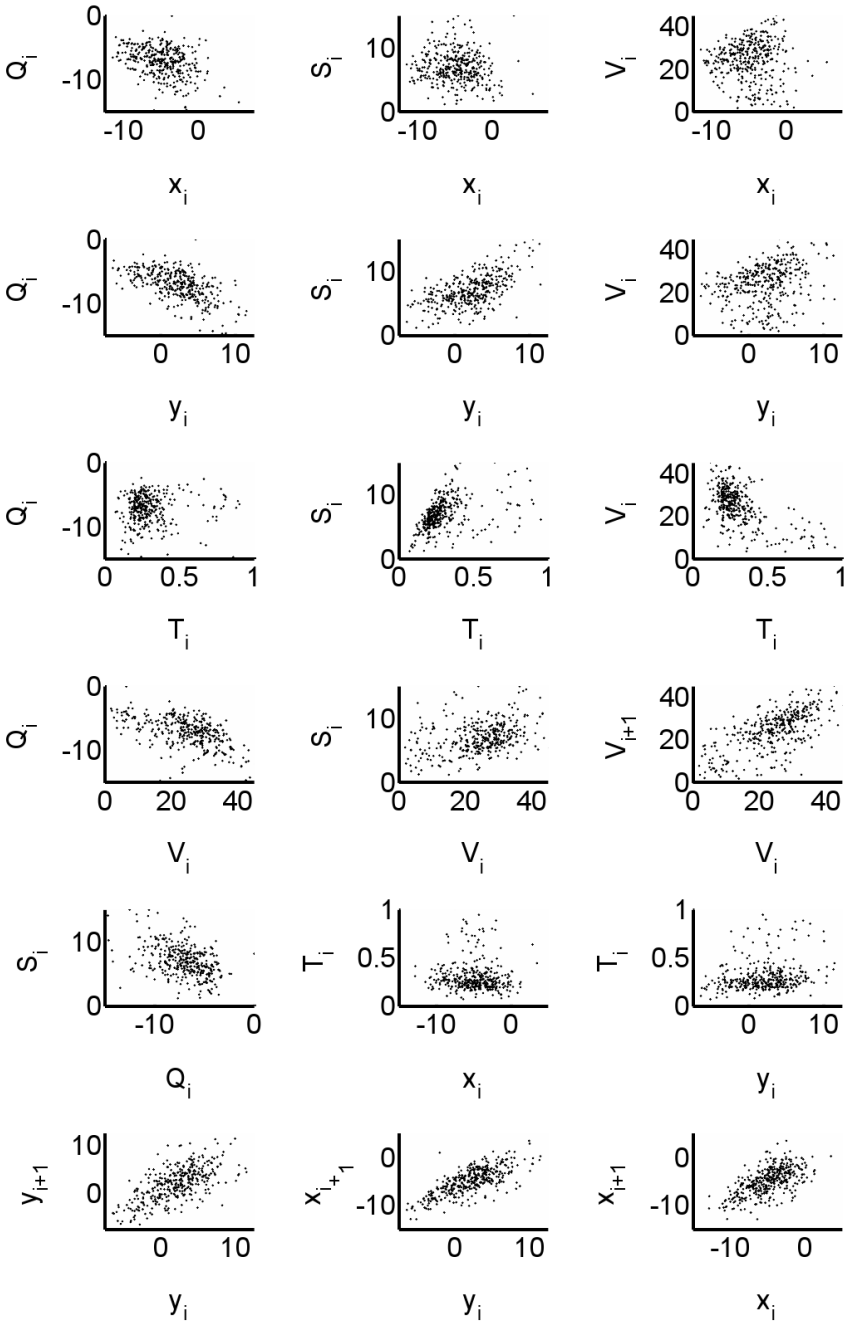

Figure 4.10: A representative sample of pair-wise scatter plots, illustrating the direct and indirect relationships between OKN variables. A number of strong and weak linear relationships, as well as a potentially non-linear relationship between V_i and T_i, can be seen. The degree of correlation between participants does vary from trial to trial. Illustrated data obtained from participant 5, stimulus speed 40°/s.

A

	Xi	Si	Vi	Ti	Yi	Qi	Xi+1
Xi	1						
Si	-0.175	1					
Vi	-0.253	-0.056	1				
Ti	-0.059	0.916	-0.351	1			
Yi	0.917	0.232	-0.273	0.312	1		
Qi	0.502	0.077	-0.085	0.121	0.528	1	
Xi+1	0.415	0.156	-0.190	0.193	0.473	-0.499	1

B

	Xi	Si	Vi	Ti	Yi	Qi	Xi+1
Xi	1						
Si	-0.579	1					
Vi	-0.566	0.526	1				
Ti	0.076	0.293	-0.414	1			
Yi	0.816	-0.001	-0.321	0.301	1		
Qi	0.049	0.3366	0.335	-0.042	0.298	1	
Xi+1	0.617	-0.296	-0.553	0.280	0.546	-0.634	1

C

	Xi	Si	Vi	Ti	Yi	Qi	Xi+1
Xi	1						
Si	-0.582	1					
Vi	-0.189	0.153	1				
Ti	-0.390	0.817	-0.151	1			
Yi	0.458	0.457	-0.040	0.467	1		
Qi	0.246	0.444	0.085	0.333	0.754	1	
Xi+1	0.059	-0.239	-0.163	-0.063	-0.196	-0.792	1

Figure 4.11: Three correlation matrices, illustrating the correlations between OKN parameters. (A) Data from participant 1, stimulus speed 10°/s. (B) Data from participants 2, stimulus speed 40°/s. (C) Data from participant 2, stimulus speed 10°/s.

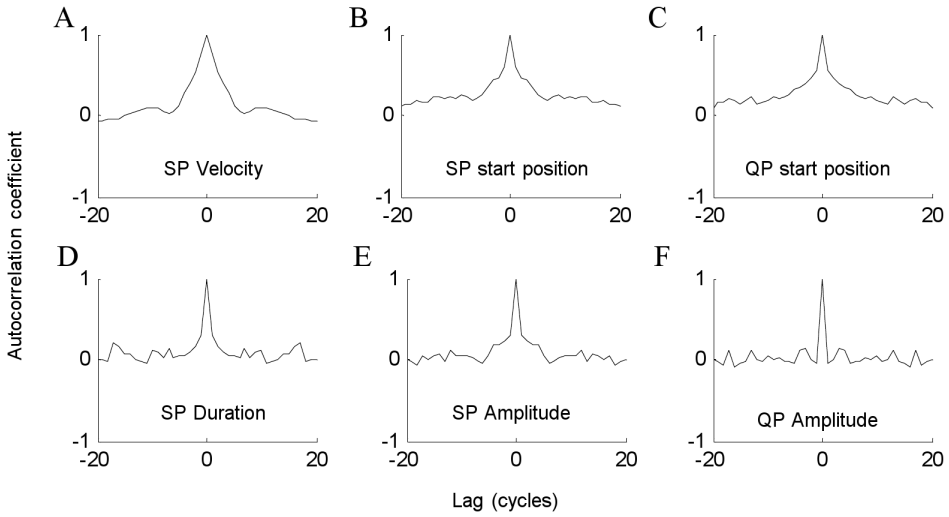

Figure 4.12: Auto-correlograms of (A) SP velocity, (B) SP start position, (C) SP end position, (D) SP duration, (E) SP amplitude and (F) QP amplitude. Illustrated data obtained from participant 1, stimulus speed 40°/s.

far as 5 cycles, although the decay of the auto-correlation was fast. Plotting the same autocorrelograms for all 40 trials revealed a similar pattern. SP velocity, SP start position, and QP start position always illustrated some auto-correlation but the strength of the autocorrelation and the number of cycles over which the correlation extends varied from trial to trial. An auto-correlation function with an exponential shape decaying to zero indicates an autoregressive model in ARIMA model identification, and examination of the partial autocorrelation function was warranted, in order to identify the order of the autoregressive process. The sample auto-correlation function of SP duration, SP amplitude and QP amplitude is approximately zero for all lags other than zero, which appeared to indicate that these data are essentially random.

The partial auto-correlation functions of SP velocity, the start position of SPs, and the start position of QPs were examined. Significant partial auto-correlations were only found at a lag of 1 cycle for SP velocity, indicating that the auto-

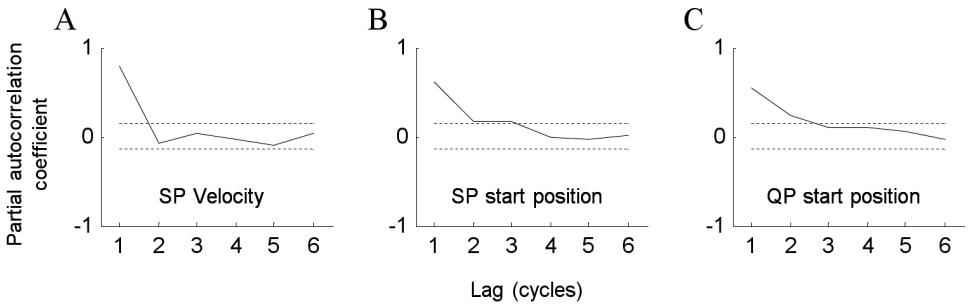

Figure 4.13: Partial auto-correlograms of (A) SP velocity, (B) SP start position and (C) QP start position. Illustrated data obtained from participant 1, stimulus speed 40°/s. Dotted lines, upper and lower confidence limits.

correlation at a lag of 1 cycle could explain all the higher-order auto-correlations in SP velocity. Significant partial auto-correlations were predominantly found at a lag of 1 cycle for the start position of SPs and QPs, although for some trials the partial auto-correlation at a lag of 2 could be significantly different to zero (fig. 4.13). This indicated that the auto-correlation at a lag of 1 cycle might be enough to explain all the higher-order auto-correlations in the start position of SPs and QPs. Although a second-order term might appear to be required for some trials, considering how close these "significant" partial auto-correlations were to the edge of the confidence interval it seemed likely that they may have only appeared significant by chance, due to comparing multiple trials. Taken with the results of the auto-correlograms, the results of the partial auto-correlograms indicated that SP velocity, and the start and end position of SPs could be described by a 1^{st} order autoregressive model.

The scatter plots of pair-wise correlations and the auto-correlograms provide evidence of many correlations among OKN variables both within and between cycles. Other than the hyperbolic relationship between SP velocity and SP duration, when relationships did exist between parameters they appeared to be linear or non-existant.

4.6 Discussion of results.

The most striking observation from these experiments was the remarkable degree of variability in SP velocity. The OKR is generally understood as a servo mechanism that minimises retinal slip, as the contrast sensitivity for high spatial frequency images drops at a logarithmic rate when images move faster than 4°/s across the retina. We have found that the mean value of SP velocity can be much lower than stimulus speed, allowing the average retinal slip to reach values much higher than 4°/s. SP velocity also appears to vary in a random fashion from SP to SP and we have found that during individual SPs, the SP velocity can actually drop to values close to 0°/s. We also found that as stimulus speed was increased the mean SP velocity recorded during each trial increased, but this increase saturated at higher stimulus speeds resulting in more retinal slip, and lower values of SP gain. The variability of SP velocity and retinal slip also increased with stimulus speed, but the variability of gain did not. If the standard deviation of SP gain was approximately constant ($\approx$0.2) during each trial then we would expect that as stimulus speed was increased, the standard deviation of SP velocity (and correspondingly retinal slip) would increase proportionally. This result may point to a specific role for SP gain in the OKN system.

Increasing stimulus speed also resulted in an increase in the magnitude of SPs and QPs, a decrease in SP duration and an increase in QP rate. Correlation between SP amplitude and SP velocity has been previously reported (Watanabe et al. 1994), so we could not be certain if the increase in SP amplitude was due to an increase in stimulus speed or SP velocity, or both. Similarly, as QPs must compensate for SPs in order that the mean eye position does not wander to the limit of gaze, the mean QP amplitude must be correlated with mean SP amplitude and the increase in QP amplitude with stimulus speed could be caused indirectly

via SP amplitude or SP velocity.

We have examined the statistical relationships among OKN parameters (the start position, end position and amplitude of the SP and QP, the SP duration, and the SP velocity). We found a number of correlations exist between parameters but the correlation matrices vary markedly between trials, and clearly the stochasticity of OKN cannot be summarised by a single correlation matrix. The results also indicate that SP velocity, and the start and end position of SPs are autocorrelated, and might be described by a 1^{st} order autoregressive model. In chapter 5 we examine the correlations between OKN parameters within and between adjacent OKN cycles with PCA, in order to determine the underlying structure of the relationships between parameters.

Chapter 5

Results of PCA performed on OKN data: a stochastic model of OKN.

In this chapter the results of performing PCA on OKN parameters within and across adjacent OKN cycles is reported. This analysis was performed in order to investigate how the underlying components contribute to the correlation of OKN parameters within and across adjacent OKN cycles. The results indicate that the differences in the correlation matrices of OKN parameters, between participants and stimulus conditions, appear to reflect changes in the eigenvalues of components while the eigenvectors remain predominantly similar across trials and across cycles. A model is developed from these results that describes OKN as a purely stochastic process that has dependencies only among adjacent cycles, but with three sources of noise affecting SP velocity, QP triggering (SP amplitude), and QP amplitude. Changes in the amplitude of QPs and SPs are explained by a dependence on the position of the eye at the beginning of the respective QP or SP, as well as the SP velocity during that cycle of OKN. Future SP velocity is explained by a dependence on the current SP velocity, and the dynamic behaviour of SP velocity is modelled as a 1^{st} order Markov chain that updates iteratively during each cyclce of OKN.

5.1 Method of principal components analysis.

For any nxn correlation matrix the maximum number of components is n, where n is the number of parameters (or original dimensions) to be considered. However, when applying PCA to the OKN parameters, n is an overestimate of the maximum DOF because of the constraints imposed by the geometric relationships among the parameters illustrated in figure 4.8. The start position of a SP must equal the start position of the QP plus the QP amplitude. Likewise, the start position of a QP must equal the start position of the SP plus the SP amplitude. These constraints, and the assumption that SP velocity is constant within each cycle, are defined mathematically as:

$$x_{i+1} = y_i + Q_i \tag{5.1}$$

$$y_i = x_i + S_i \tag{5.2}$$

$$V_i = (y_i + x_i)/T_i \tag{5.3}$$

Thus there is effectively only a maximum of d dimensions in the data, and we can assume that the smallest n-d components will not contain any meaningful information. The exact number of d depends on how many cycles and how many of the OKN parameters we consider. If we consider all of the characteristic parameters

(x_i, S_i, V_i, T_i, y_i, and Q_i) then d is three times the number of cycles included plus one, so $d = 4$ when we consider only one cycle with all of the characteristic parameters. The values of d for each analysis performed is given in section 5.2. The d dimensions do not all pass through the same axes as the n dimensions, and so it is important to include all n dimensions in PCA extraction, but only the d largest components are retained for further analysis. During this investigation the discarded eigenvectors had eigenvalues that were either negligible ($< 10^{-5}$) or had values less than 1 (0.3 ± 0.2) due to the residual error from the linear approximation of eq. 5.3 implicit in the PCA. This error can be partially overcome by performing a reciprocal transformation of T_i, although this introduces a new nonlinear relationship between S_i and $1/T_i$ so cannot solve the problem entirely. For completeness PCA was also performed using $1/T_i$ as a parameter in place of T_i with qualitatively similar results, but with the sign of the loadings from $1/T_i$ reversed when compared with the loadings from T_i on each component. After discarding $n - d$ components, factor rotation was performed using the "varimax" strategy in order to obtain orthogonal (uncorrelated) rotated components. After factor rotation, similar loading patterns were observed among the trials, but with different eigenvalues.

The d components from each trial were sorted into categories manually by eye, based on the whether they shared the same high (or low) component loadings. Later, when it appeared that the similarities in loading patterns were consistent across individuals this procedure was automated using a numerical algorithm. These heuristics can be found in Appendix A. It should be noted that principal components can be rotated, such that they face in the opposite direction but remain in the same dimension, causing all the loadings on the component to have the opposite sign. During sorting it was necessary to flip the sign of all

loadings on these "mirrored" components so that they could be sorted in to the correct categories. In order to demonstrate the stability of components across participants and stimulus speeds, each category was expressed graphically as a line plot of component loading against the original parameters, and components placed in the same category were plotted on the same axis (fig. 5.1 - 5.6).

There are limits to categorising components using this type of heuristic approach. As the criteria for creating categories were initially derived by examining the components by eye, they are quite arbitrary and applying these criteria to a different set of data would likely generate some different results. The sorting is limited to the complexity of the heuristic algorithm. Any categories not accounted for by the heuristic criteria could result in components being sorted into very different categories, and subcategories of components could also be missed. However, I decided to use the heuristic approach based on the observation that components did fit quite clearly into similar loading patterns. Had there been striking differences between components sorted into the same category a more sophisticated form of cluster analysis would have been warranted, but this appeared to be unnecesary.

5.2 Sequences of OKN parameters included in the analysed correlation matrices.

Six measurements were taken from each OKN cycle: x_i, S_i, y_i, Q_i, V_i and T_i ($i = 1...m$) according to the scheme shown in fig. 4.8, and we defined these as the six characteristic parameters of an OKN cycle. Eye velocity changed between cycles but could also vary during a single SP, however the variability introduced by these non-linear SPs was not analysed in this study. For the purposes of PCA the SP velocity is measured as the difference in eye position at the beginning and end of the SP divided by the SP duration, such that $V_i = (y_i - x_i)/T_i$.

Measurements from adjacent OKN cycles were grouped to create a vector that was used to generate a correlation matrix for each trial (40 trials, 10 subjects x 4 speeds). Several different vectors were considered in order to generate the correlation matrixes. The first analysis was performed on the 1x5 vector consisting of the parameters: x_i, S_i, y_i, Q_i and x_{i+1}. This vector is simple as it encompasses only one SP and one QP in series and does not include the non-linear relationship between SP velocity and duration. The Pearson correlation coefficients were calculated using the built-in Matlab function: *corrcoef*, for all possible pairs among the 5 variables using m-1 vectors (the m^{th} is dropped as there is no x_{m+1} parameter). This generated a symmetric 5x5 correlation matrix with 10 unique off-diagonal coefficients. PCA was then performed on the correlation matrixes.

The second vector that was considered consisted of the parameters: x_i, S_i, V_i, T_i y_i, Q_i and x_{i+1}. This vector still encompasses only one SP and one QP in series but introduces the non-linear relationship between SP velocity and duration. This generated a symmetric 7x7 correlation matrix with 21 unique off-diagonal coefficients. PCA was then performed on these correlation matrixes.

The third vector considered consisted of the parameters: x_i, S_i, V_i, T_i y_i, Q_i, x_{i+1}, S_{i+1}, V_{i+1}, T_{i+1} y_{i+1}, Q_{i+1} and x_{i+2}. This vector includes the non-linear relationship between SP velocity and duration, and introduces a second cycle of OKN. There was no overlap between vectors and in this and subsequent analyses a vector was not included if it contained one or more blinks. This generated a symmetric 13x13 correlation matrix with 78 unique off-diagonal coefficients.

The fourth vector considered consisted of the parameters from three adjacent cycles and the fifth vector consisted of the parameters from four adjacent cycles. These vectors generated a symmetric 19x19 correlation matrix with 171 unique off-diagonal coefficients and a symmetric 25x25 correlation matrix with 300 unique

off-diagonal coefficients respectively. There was no qualitative difference in the results of PCA performed on the third, fourth and fifth vectors and so exploration was terminated at 4 adjacent cycles.

Due to the constraints imposed by equations 5.1 and 5.2 the DOF for our first vector (x_i, S_i, y_i, Q_i and x_{i+1}) were reduced by two. For the second vector (x_i, S_i, V_i, T_i, y_i, Q_i and x_{i+1}) the DOF were reduced by three (eq. 5.1, 5.2 and 5.3). For the third vector, including two cycles of OKN, the DOF were reduced by six. For all five vectors the maximum possible DOF were therefore d=3, 4, 7, 10 and 13 respectively. PCA is often used to reduce the number of dimensions in a data set. However, the goal of this analysis was to examine all of the d DOF in order to explore the linear weights of the components and look for any underlying patterns. In this case the number of eigenvalues to be retained is already known and the components with the $n - d$ lowest eigenvalues can be discarded before rotation. Eigenvalues of the discarded components were checked and found to be either negligible ($< 10^{-5}$) or had values less than 1 (0.3 ± 0.2) due to the residual error from the linear approximation of eq. 5.3 implicit in PCA.

5.3 Results of PCA performed on a single cycle of OKN parameters.

The first analysis was performed on the simplest non-trivial vector of a single OKN cycle consisting of the 5 variables (x_i, S_i, y_i, Q_i and x_{i+1}) with a maximum of 3 DOF (section 5.2). Figure 5.1 illustrates, step by step, the analysis of OKN data from two different participants with qualitatively different traces of eye position. In the left panel (fig. 5.1a, 5.1c and 5.1e) are the results from participant 2 (stimulus speed 10°/s, m=201 cycles) and in the right panel (fig. 5.1b, 5.1d and 5.1f) are the results from participant 1 (stimulus speed 10°/s, m=388 cycles).

The correlation matrix in the left panel (fig. 5.1b) contains a mixture of high and

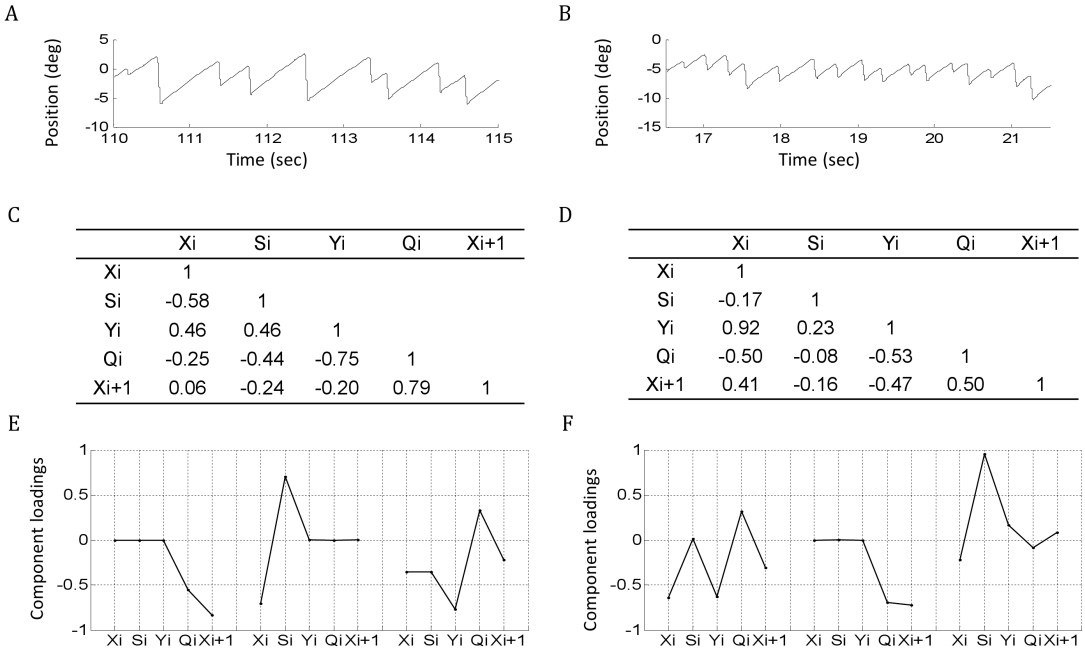

C

	Xi	Si	Yi	Qi	Xi+1
Xi	1				
Si	-0.58	1			
Yi	0.46	0.46	1		
Qi	-0.25	-0.44	-0.75	1	
Xi+1	0.06	-0.24	-0.20	0.79	1

D

	Xi	Si	Yi	Qi	Xi+1
Xi	1				
Si	-0.17	1			
Yi	0.92	0.23	1		
Qi	-0.50	-0.08	-0.53	1	
Xi+1	0.41	-0.16	-0.47	0.50	1

Figure 5.1: Step-by-step example of PCA analysis. Top panels (A and B) illustrate eye position during OKN presentation. Middle panels (C and D) illustrate correlation matrixes of OKN parameters from the data set corresponding to the eye position trace above them. Bottom panels (E and F) illustrate the component loadings for the three primary components extracted from the data and are plotted, in the order of highest to lowest eigenvalues, from left to right. Left panels (A, C and E) illustrates data from participant 2, stimulus speed $10°/s$. Right panels (B, D and F) illustrates data from participant 1, stimulus speed $10°/s$. First component in panel E and second component in panel F represents QPs with a more negative amplitude have a more negative end position. Second component in panel E and third component in panel F represents SPs with a more negative start position have a more positive amplitude. Third component in panel E and first component in panel F represents QPs with a more negative start position have a more positive amplitude.

low correlation coefficients between the OKN parameters. PCA performed on this matrix produced 5 components with the eigenvalues 2.50, 1.59, 0.90, $3.09 \cdot 10^{-6}$ and $3.63 \cdot 10^{-7}$. The 4^{th} and 5^{th} eigenvalues were negligible as expected because the constraints in eq. 5.1 and 5.2 reduce the DOF to a maximum of 3. The first three components were retained and rotated using the "varimax" strategy. After rotation each component is illustrated as a line plot of the rotated component loadings against the original OKN parameters, and the components are placed in order of descending eigenvalue (fig. 5.1e).

In order to interpret the loading patterns it was necessary to consider that they represent the relative contribution of the original parameters to the components. For example the first component expressed in figure 5.1e has negative loadings from Q_i and x_{i+1}, thus represents QPs with more negative amplitude (bigger magnitude) having a more negative end position. In other words bigger QPs end further in the direction that the stimulus is drifting from. Note that this interpretation is equally valid if the sign of all loadings were reversed, as more positive amplitude (smaller magnitude) QPs would have more positive end position. The second component expressed in figure 5.1e has negative loading from x_i and positive loading from S_i, and represents SPs that start further in the direction that the stimulus is drifting from having a more positive amplitude (bigger magnitude). Finally, the third component represents QPs that start further in the direction that the stimulus is drifting from having a more positive amplitude (smaller magnitude).

The correlation matrix in the right panel (fig. 5.1d) was very different from the first example. PCA performed on this matrix produced 5 components with the eigenvalues 2.41, 1.51, 1.08, $5.58 \cdot 10^{-6}$ and $2.37 \cdot 10^{-6}$. The first three components were retained and rotated using the "varimax" strategy, and the rotated com-

ponent loadings were plotted against the original OKN parameters in descending eigenvalue order as before (fig. 5.1f). In spite of the qualitative differences between the two OKN traces and the quantitative differences in the correlation matrixes, there is a similar pattern of component loadings but in a different eigenvalue order. The first component expressed in figure 5.1f is similar to the third component in figure 5.1e, whereas the second component expressed in figure 5.1f is similar to the first component in figure 5.1e, and the third component expressed in figure 5.1f is similar to the second component in figure 5.1e.

Performing PCA on the correlation matrixes obtained from all 40 trials it became clear that, despite the differences across individuals and stimulus speeds, the component loadings always fell into 3 similar patterns. In every trial, 3 significant eigenvalues were found corresponding to the maximum number of DOF. Sometimes these components might appear in a different order of eigenvalues, and so 3 simple heuristics were developed in order to sort components in to categories based on their component loadings rather than their eigenvalues (fig. 5.2). At first these categories were simply labelled category "X", "Q" or "S". Category "X" contained components with high positive loading from QP amplitude and high negative loading from QP start position. We will return to category "X" in section 5.5 as it appears this category is closely linked to the "Q" category. Category "Q" contained components associated with high negative loading from QP amplitude and high negative loading from QP end position. Category "S" contained components with high positive loading from SP amplitude.

These results implied that there was some underlying structure to the correlations between OKN parameters despite the clear differences in correlation matrices between participants, and that there were three uncorrelated sources of variance in the first analysis, including x_i, S_i, y_i, Q_i and x_{i+1}.

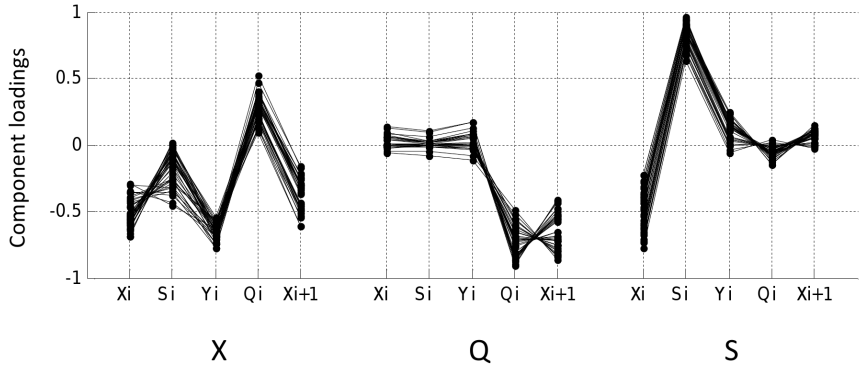

Figure 5.2: Results of single cycle PCA. Component loadings of 3 principal components extracted from OKN data for all 40 trials. For each trial the 3 principal components extracted are sorted heuristically in to 3 categories. From left to right, categories "X", "Q" and "S".

5.4 Results of PCA performed on a single cycle of OKN parameters, including SP velocity and SP duration.

The 2nd analysis was performed on the correlation matrix created from a vector including SP velocity and duration, so that there were now 7 variables (x_i, S_i, V_i, T_i, y_i, Q_i and x_{i+1}) and a maximum of 4 DOF. Performing PCA on the correlation matrixes from all 40 trials revealed 4 significant eigenvalues corresponding to the maximum number of DOF. The 6th and 7th eigenvalues were always near zero and the 5th was always very small ($0.3\pm$ 0.2). This small variance appeared to arise due to the non-linear constraint (eq. 5.3). Simulating a simple non-linear relationship between three parameters, where one parameter was subject to random fluctuations, it was possible to demonstrate that PCA cannot completely re-express the non-linear combination of variables as a linear weighted sum. This leaves some small fraction of variance seemingly "unexplained".

As before, PCA performed on every trial generated these 4 significant eigenvalues and 4 simple heuristics were developed in order to sort the components in to categories based on their component loadings rather than their eigenvalues (fig.

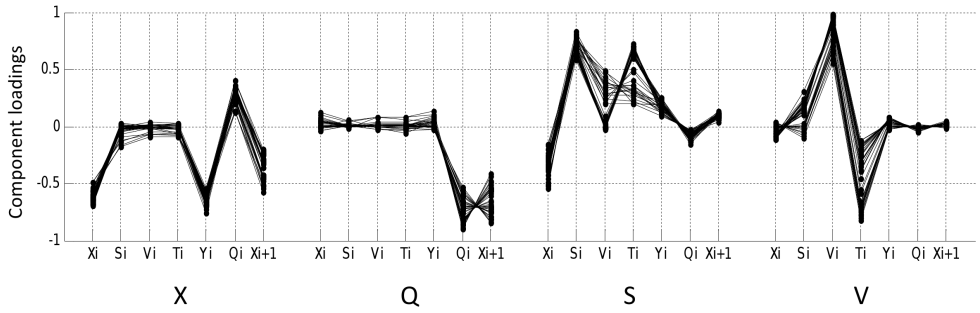

Figure 5.3: Results of PCA performed on a single cycle of OKN, including SP velocity and duration. Component loadings of 4 principal components extracted from OKN data for all 40 trials. For each trial the 4 principal components associated with the 4 highest eigenvalues are sorted heuristically in to 4 categories. From left to right, categories "X", "Q", "S" and "V".

5.3). The most clear effect was to introduce the new category "V" that contained components with high positive loading from SP velocity and high negative loading from SP duration. This seemed to imply that variability in these two variables tends to cancel to keep SP amplitude roughly constant. Category "S" was also affected. For most trials a high positive loading from SP duration was introduced, and a very small positive loading from SP velocity. There were several trials, however, where SP velocity had high positive loadings on the component and there were correspondingly lower positive loadings from SP duration. Categories "X" and "Q" were largely unaffected.

These results implied that the introduction of the parameters SP velocity and duration generated only one more significant source of variance in a single cycle of OKN, the "V" component.

5.5 Results of PCA performed on multiple cycles of OKN parameters.

The 3rd, 4th and 5th analyses were performed on correlation matrices created by vectors including parameters from across multiple cycles of OKN. Vectors of OKN parameters from 2, 3 and 4 cycles were considered with similar results. Here, the results of analysis on the 4 cycle vector are presented.

For n cycles (including V_i and T_i) there are $6n+1$ parameters and $3n+1$ maximum DOF. Vectors of 4 cycles contained 25 parameters, 300 unique off-diagonal correlation coefficients and a maximum of 13 DOF. Performing PCA on the correlation matrices from all 40 trials revealed 13 significant eigenvalues corresponding to the maximum number of DOF.

Further, we found that there were three broad groups of components and that these groups represented the categories with similar loading patterns but displaced by one, two or three cycles (fig. 5.4 - 5.6), and OKN parameters from only one cycle (or just over one cycle) had high loadings in any given component. At this point it was also possible to see that the categories that were initially labelled "X" and "Q" belonged to essentially the same group of components but separated by one cycle. Component 1 and component 5 (fig. 5.4) represent the original "X" and "Q" components respectively (cf.fig. 5.2).

While eigenvectors demonstrated remarkably similar loading patterns their eigenvalues were much more variable. The eigenvalue order of the sorted components often changed between trials, as illustrated in figure 5.7 where no particular component maintained a high or low eigenvalue across all of the trials, except for component 5. Component 5 was characterised with high loading from just two parameters (Q4 and X5, in fig. 5.4) and maintained a relatively low eigenvalue of

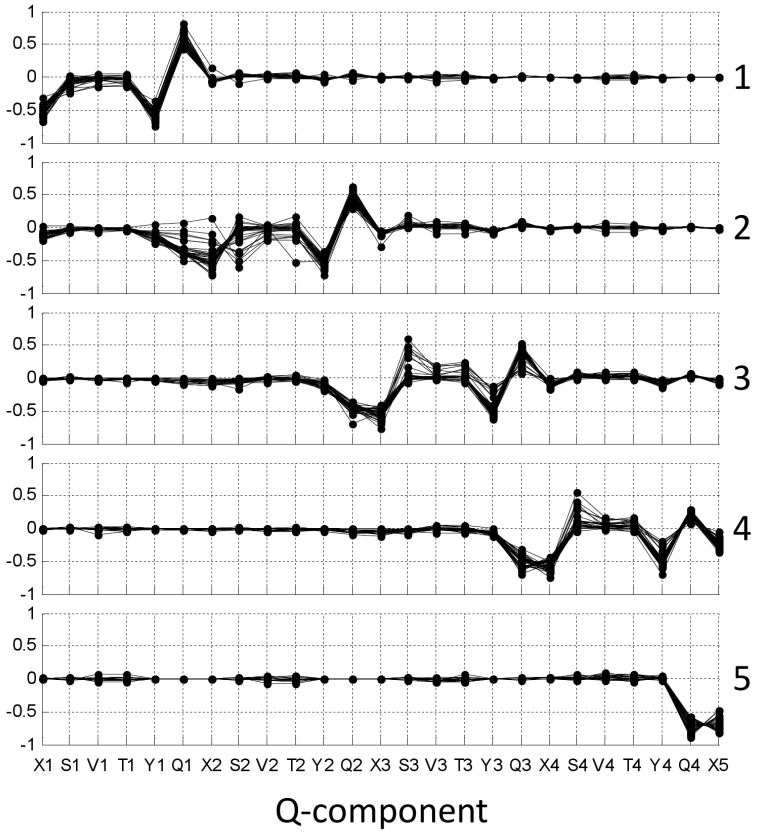

Figure 5.4: Results of PCA performed on four cycles of OKN parameters in series, including SP velocity and duration. Component loadings of 5 principal components extracted from OKN data from all 40 trials. A total of 13 principal components were extracted during analysis. Five components were similar to the categories "X" and "Q" in the single cycle analysis and are illustrated in this figure as one broad category of components, the Q-component category. Numbers are used only to label components and do not represent the position of the corresponding eigenvalue on the scree plot as these varied between trials.

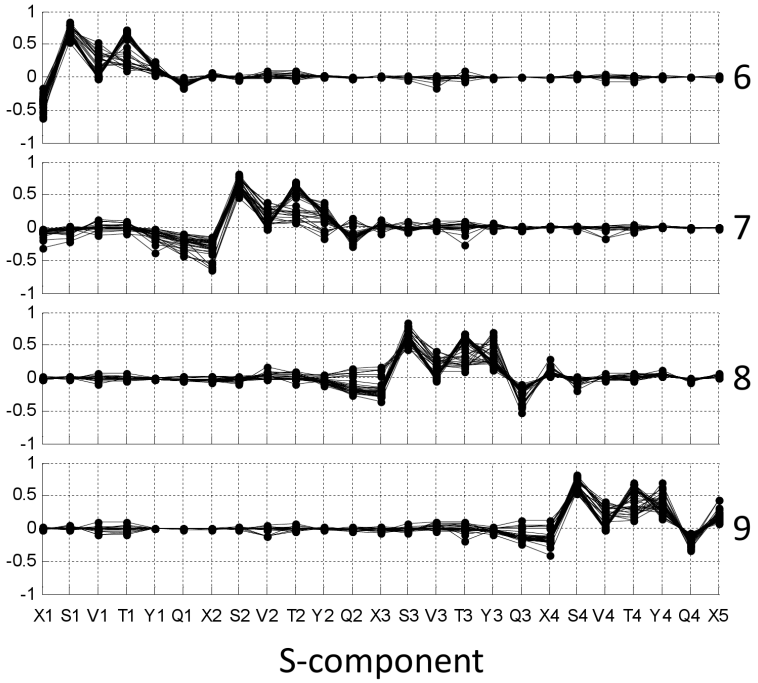

Figure 5.5: Results of PCA performed on four cycles of OKN parameters in series, including SP velocity and duration. Component loadings of 4 principal components extracted from OKN data from all 40 trials. A total of 13 principal components were extracted during analysis. 4 components were similar to the category "S" in the single cycle analysis and are illustrated in this figure as one broad category of components, the S-component category. Numbers are used only to label components and do not represent the position of the corresponding eigenvalue on the scree plot as these varied between trials.

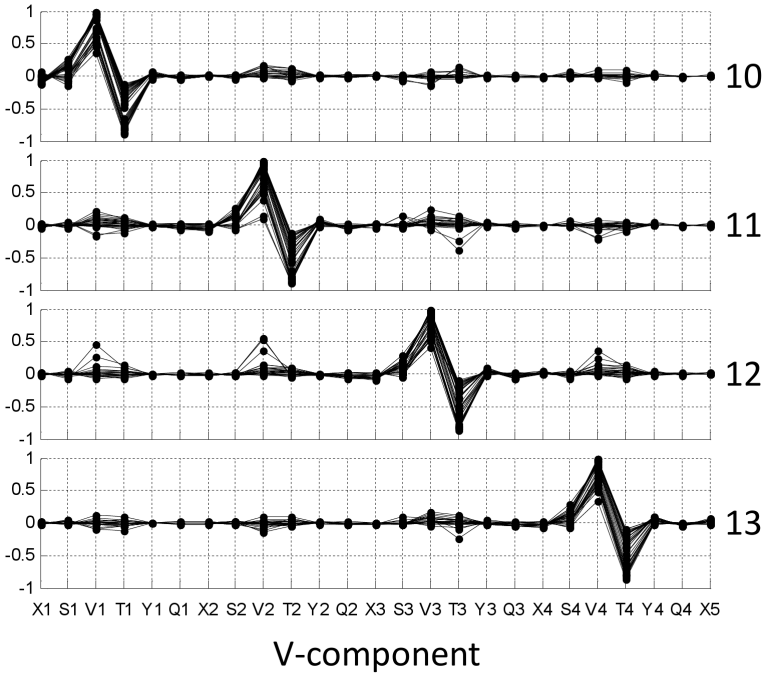

Figure 5.6: Results of PCA performed on four cycles of OKN parameters in series, including SP velocity and duration. Component loadings of 4 principal components extracted from OKN data from all 40 presentations. A total of 13 principal components were extracted during analysis, 4 components were similar to the category "V" in the single cycle analysis and are illustrated in this figure as one broad category of components, the V-component category. Numbers are used only to label components and do not represent the position of the corresponding eigenvalue on the scree plot as these varied between trials.

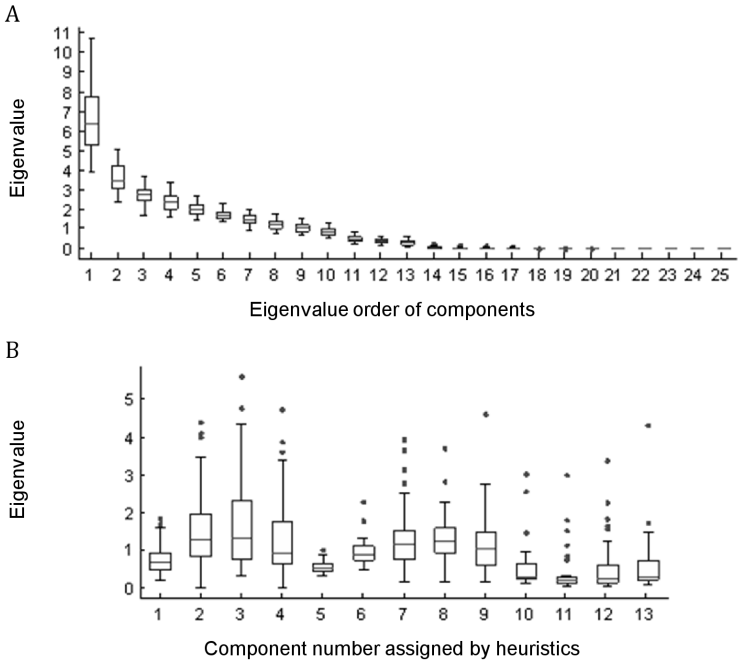

Figure 5.7: (A) Boxplot illustrating the eigenvalues of principal components found in four cycles of OKN parameters in series before being sorted by heuristics. Number on principal component axis represents the eigenvalue order of components from largest to smallest. (B) Boxplot illustrating the eigenvalues of principal components extracted from four cycles of OKN parameters in series after being sorted by heuristics. Sorted component numbers 1-5, "Q-component"; 6-9, "S-component"; 10-13, "V-component". Boxes represents interquartile range, whiskers represent maximum and minimum data values (excluding outliers), and dots represent outliers found more than 1.5 times the interquartile range from the ends of the box.

an order just under 1 across all trials.

The finding that there are essentially only three broad groups of components, each associated with a given cycle of OKN, greatly simplifies our results. The three groups of components indicate that each cycle of OKN has three orthogonal (uncorrelated) noise processes, which we call the "Q-component" (originally the "X" and "Q" categories), the "S-component", and the "V-component".

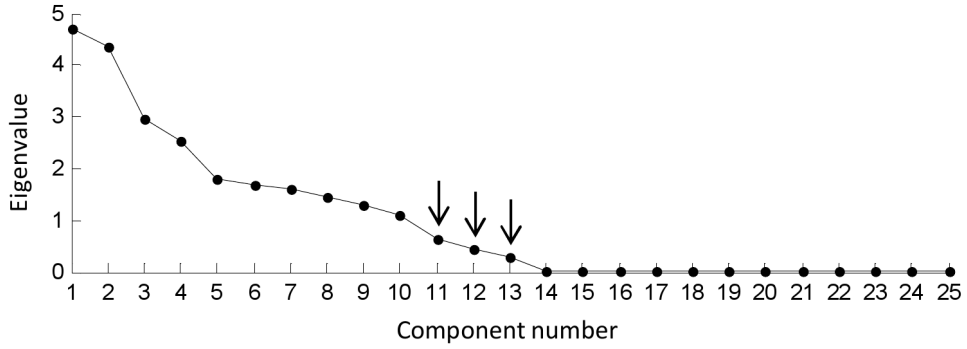

Figure 5.8: A representative example of a scree plot, illustrating the eigenvalues of all 25 components in four cycles of OKN parameters. The eigenvalues of the 11^{th}-13^{tm} largest eigenvectors sum to 1.45, indicating that they explained only 5.8% of the variance in the data. Data taken from participant 10, stimulus speed 40°/s

5.6 The results of retaining only ten eigenvectors.

PCA is often used as a dimension-reducing scheme for data sets with numerous observed dimensions of data. During extraction a number of eigenvectors can be discarded before factor rotation, leaving only a certain percentage of the variance explained by the components retained. The purpose of our analysis was not to deliberately reduce the dimensions of the data other than the reduction already imposed by the geometric relationships between parameters, which would still leave 100% of the variance explained. However, for completeness, and in order to check the validity and possible interpretation of the results, scree plots were created for each trial when four cycles of OKN parameters were considered and dimension reduction was conducted in order to determine if it was at all possible to reduce the number of components.

A representitive scree plot is illustrated in figure 5.8. The 10 largest components could explain 94% of the variance in the 25 parameters, and when only 10 eigenvectors were retained for rotation another interesting pattern of categories emerged. All of the components from the "V" category appeared to merge their loading

patterns into a single component that illustrated high loading from SP velocity across all cycles (fig. 5.9). The loading patterns of the "S" and "Q" category of components remained relatively similar, but with increased loading from SP velocity across all cycles. This appeared to indicate that the 'S" and "Q" categories were somewhat dependent on SP velocity.

Other than when all thirteen components were retained, these clear loading patterns of components only occurred when the ten largest components were retained for factor rotation. When only eleven components were retained, the nine components representing the "S" and "Q" category of components remained relatively similar but with an increased loading from SP velocity across all cycles, as when only ten components were retained. The remaining two components were characterised by high loading from SP velocity from two separate cycles, but these cycles appeared arbitrary. From the parameters in a given trial, one of these components might have high loading from SP velocity in cycles 1 and 4 whereas the other component would have high loading from SP velocity in cycles 2 and 3. However, in a separate trial one component might have high loading from SP velocity in cycles 1 and 2 whereas the other component would have high loading from SP velocity in cycles 3 and 4. Only when ten components were retained for factor rotation did all four of the "V" category of components merge into one component. These results indicated that when considered together the three smallest components from the d DOF represented uncorrelated sources of noise in SP velocity that occurred during each cycle of OKN. However, when considered individually these components occured in different eigenvalue orders.

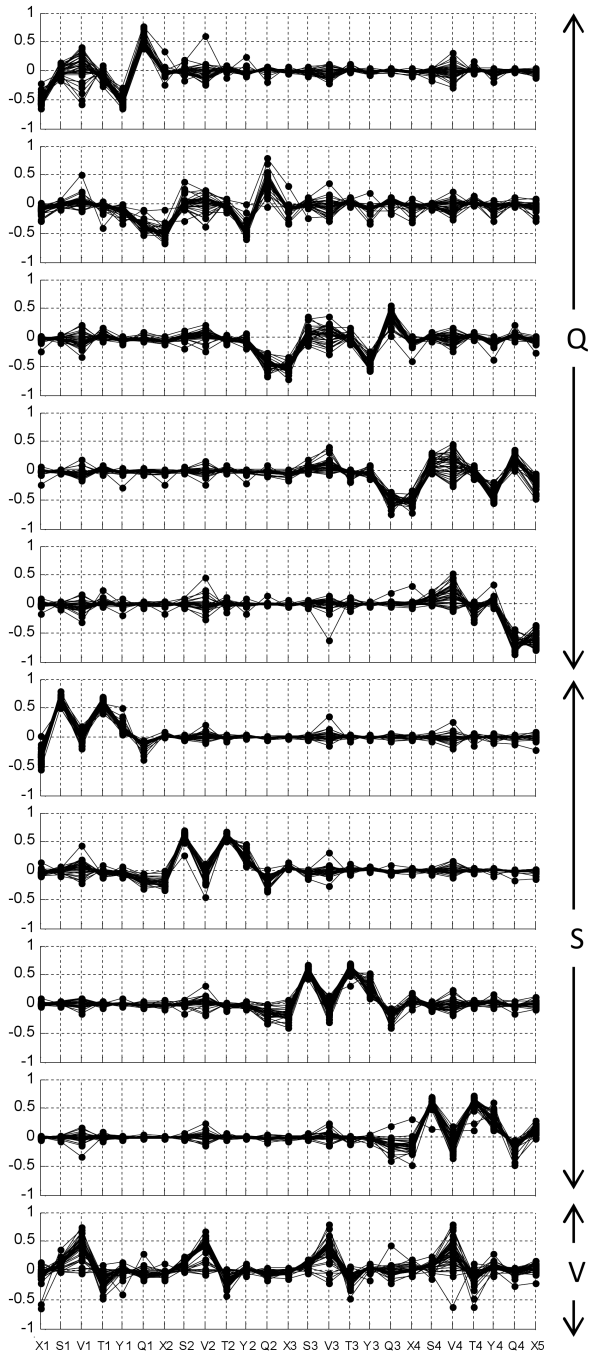

Figure 5.9: Results of PCA performed on four cycles of OKN in series, including SP velocity and duration. Component loadings of the 10 largest principal components extracted from OKN data from all 40 presentations. The "V" category of components found in the previous results have all merged in to one very different component. SP velocity now loads on to both the "S" and "Q" category of components from across all cycles.

5.7 Development of the 1st order Markov model.

Some correlations between OKN parameters would be forced due to the constraints imposed by eq. 5.1 - 5.3 and so it was not possible to simply take the rotated component loadings from each category. Instead, it was necessary to interpret the loading patterns and constraints together to yield a stochastic model. A number of possible linear statistical relationships were considered between the original OKN parameters that could give the observed categories of principal components.

In order to arrive at a model based on PCA results, that could reproduce realistic OKN with the same component loadings seen in fig. 5.4 - 5.6 it was necessary to make a few assumptions. The first assumption was that there were 3 sources of variability or "noise" within each individual cycle, corresponding to the S-, Q- and V-component groups, and that this noise was normally distributed. PCA performed with "varimax" rotation identifies orthogonal (uncorrelated) components after discarding components that represent a negligible amount of variance, and so these sources of variability would also be uncorrelated and independent of each other if they were jointly normal. Second, as the results of PCA indicated a lack of long range correlations across cycles, it seemed that the process could be explained by a 1st order Markov process. For a 1st order Markov process, the state of the system is conditional only on the immediate previous state and not explicitly on earlier states. This assumption also agrees with the finding that methods for determining the order of an ARIMA model would lead to a 1st order autoregressive model being identified (section 4.5). Third, the value of parameters from the model may differ across individuals and stimulus speeds as evidenced by the differences in correlation matrices between trials, so it would be necessary to derive parameters from individual trial data using multiple linear regression.

In order to test the model, I performed Monte-Carlo simulations to generate 40 artificial sequences of OKN parameters (x_i, S_i, V_i, T_i, y_i and Q_i). The length of each sequence was matched to the corresponding number of cycles from each trial, and the coefficients of the linear model were based on the robust regression analysis (weighted least squares method) of the relevant parameters from each trial using the Matlab function: *robustfit*. The same PCA procedure was then performed on these sequences, in order to determine if they would produce the same categories of components as the original data. Other possible linear relationships were also explored, but the simulated sequences produced from these other models generated poorer examples of the component categories and pairwise scatter plots.

5.7.1 S-component.

In the S-component there was high positive loading from S_i, positive loading from V_i, and negative loading from x_i (fig. 5.5). This loading pattern indicated that SP amplitude increases (becomes more positive) as SP velocity increases, and as the eye position at the start of the SP is directed further in to the negative direction. These relationships can be seen in the regression of S_i against x_i and V_i (fig. 5.10a and 5.10b), both illustrating a linear trend, and can be expressed in the form:

$$S_i = ax_i + bV_i + \widehat{s} + \epsilon_s(i) \tag{5.4}$$

where a, b and $\widehat{s}$ are constants, and $\epsilon_s(i)$ is a normal random process with standard deviation σ_s.

In the S-component there are also positive loadings from T_i and some from y_i. It seems most likely that the positive loading from y_i is forced by the constraint in eq. 5.2, considering that y_i cannot have a direct causal effect on the parameters

that precede it. Similarly the positive loadings from T_i will be forced by the constraint in eq. 5.3. It could be argued that the loading pattern indicates that SP duration increases with SP velocity, and that it is the positive loading from S_i that is forced by the constraint in eq. 5.3. I decided to take the consistently higher loading from S_i over T_i in the S-component to mean that S_i is the dependent variable here.

5.7.2 Q-component.

There was a similar loading pattern in the Q-component, with positive loading from Q_i and negative loading from y_i, indicating that the QP magnitude increases (becomes more negative) as the eye position at the start of the QP is directed further in to the positive direction. Evidence of V_i loading on the Q-component was not as clear in fig. 5.4, but is apparent in the results of retaining only 10 components fig. 5.9, and is sometimes positive and sometimes negative. Regression of Q_i against V_i did reveal a linear trend, as did the regression of Q_i against y_i (fig. 5.10c and 5.10d), and can be expressed in the form:

$$Q_i = cx_i + dV_i + \widehat{q} + \epsilon_q(i) \tag{5.5}$$

where c, d and $\widehat{q}$ are constants, and $\epsilon_q(i)$ is a normal random process with standard deviation σ_q. We interpret the other loadings in the Q-component to be forced by the constraints expressed in eq. 5.1, 5.2 and 5.3.

5.7.3 V-component.

In the V-component of each cycle there was a high positive loading from V_i and corresponding negative laoding from T_i with a small or absent loading from S_i

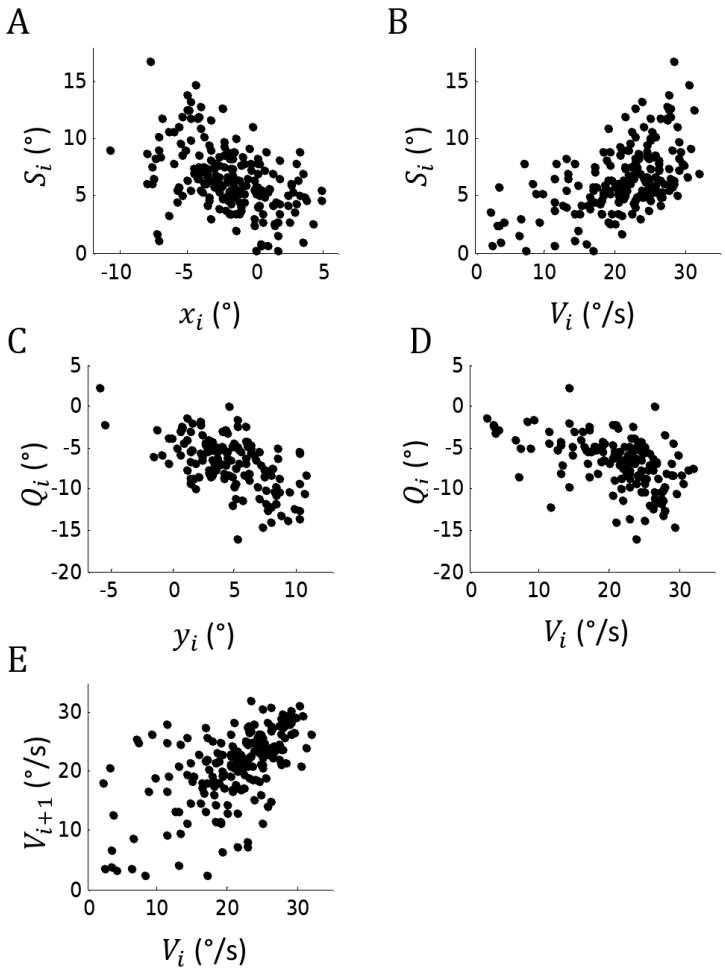

Figure 5.10: Scatter plots illustrating linear relationships between (A) SP start position and SP amplitude, (B) SP velocity and SP amplitude, (C) QP start position and QP amplitude, (D) SP velocity and QP amplitude, and (E) SP velocity during one cycle and SP velocity during the next cycle. Data taken from participant 3, stimulus speed $30°/s$.

(fig. 5.6). This indicates that the V-component describes whether SP velocity is fast or slow independently of its amplitude. This is expected from the constraint in 5.3 and implies that the V-component is self-contained, and does not depend directly on the other variables. However, from figure 5.9 we can see that there is a clear cross cycle dependency if we ignore the uncorrelated sources of noise in SP velocity that occur during each OKN cycle. This auto-correlation is also visible in the autocorrelogram of SP velocity (fig. 4.12a), and the regression of V_{i+1} against V_i (5.10e). Results from the autocorrelograms and partial autocorrelograms of SP velocity indicated that this relationship should be modelled in the form of an autoregressive, rather than a moving average model:

$$V_{i+1} = eV_i + \widehat{v} + \epsilon_v(i) \tag{5.6}$$

where e and $\widehat{v}$ are constants, and $\epsilon_v(i)$ is a normal random process with a standard deviation σ_v.

5.8 Results of Monte Carlo simulations.

Eq. 5.1 - 5.6 describe the model of the OKN system, containing three constraints and three discrete uncorrelated stochastic processes that are each expressed as a simple linear combination of OKN variables. Eq. 5.1 and 5.6 describe the update dynamics as a 1^{st} order Markov process where the subsequent value of the parameters is dependent only on the current value of the parameters.

For the purposes of estimating the model variables, they were assumed to be constant during each individual OKN presentation. Robust multiple linear regression was performed on the parameters recorded from the real OKN data in order to

estimate the values of a, b, c, d, e, $\widehat{s}$, $\widehat{q}$ and $\widehat{v}$ from each trial using a weighted least squares method. The values of σ_s, σ_q and σ_v were also estimated from each trial using the standard deviation of the residuals of linear regression. Using all of these values, 40 Monte Carlo simulations were performed in order to create simulated data to mirror the original data sets. The simulations were created iteratively using eq. 5.1 - 5.6 in a simple Matlab function, with the additional constraints:

$$S_i > 0 \tag{5.7}$$

$$V_i > 0 \tag{5.8}$$

$$|Q_i| > z \tag{5.9}$$

where z is a small positive value to mimic the QP dead zone, which was set to 1°. If a value of S_i or V_i was generated equal to or below zero the algorithm would generate a new value, and this process would continue until a value above zero was generated. The same system was put in place for generating a value of Q_i but the process continued until a value above $|z|$. PCA was then performed on the simulated data in exactly the same way as with the empirical data, and a similar pattern emerged. Thirteen significant components were found in each simulated data set when four cycles of OKN parameters were considered, and the eigenvalue order of the sorted components changed between trials (fig. 5.11) as

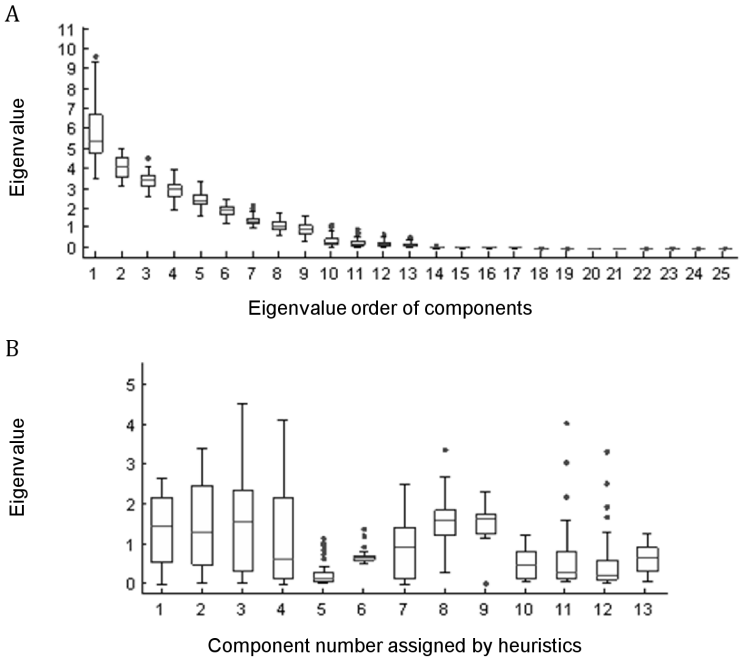

Figure 5.11: (A) Boxplot illustrating the eigenvalues of principal components found in four cycles of simulated OKN parameters in series before being sorted by heuristics. Number on principal component axis represents the eigenvalue order of components from largest to smallest. (B) Boxplot illustrating the eigenvalues of principal components extracted from four cycles of simulated OKN parameters in series after being sorted by heuristics. Sorted component numbers 1-5, "Q-component"; 6-9, "S-component"; 10-13, "V-component". Boxes represents interquartile range, whiskers represent maximum and minimum data values (excluding outliers), and dots represent outliers found more than 1.5 times the interquartile range from the ends of the box.

before. The components were sorted using the same numerical heuristics, and the results are illustrated in Fig. 5.12 - 5.14. The number of components extracted, and the patterns and values of component loadings were identical to those found in the empirical data (cf. fig. 5.4 - 5.6).

The mean (and median) values of simulated parameters accurately predict the mean (and median) values observed in OKN parameters (fig. 5.15). However, it is worth noting that while this is a nice quantitative result, it is not surprising considering the number of parameters that are fit to the data. It is therefore also

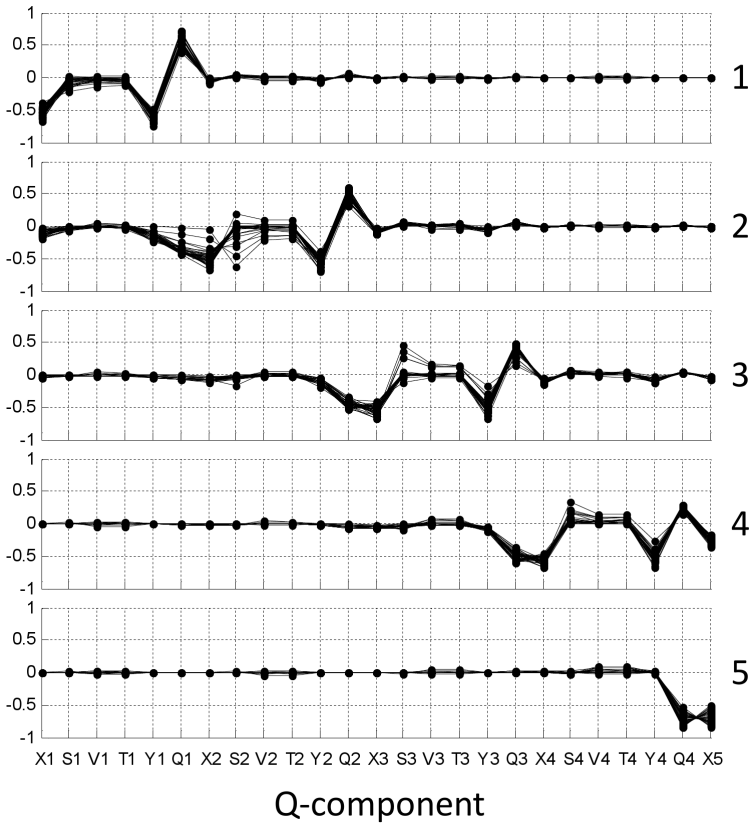

Figure 5.12: Results of PCA performed on four cycles of simulated OKN in series, including SP velocity and duration. Component loadings of five principal components extracted from 40 sets of OKN data simulated from eq. 5.1 - 5.9. A total of 13 principal components were extracted during analysis, 5 components are illustrated in this figure as the Q-component category.

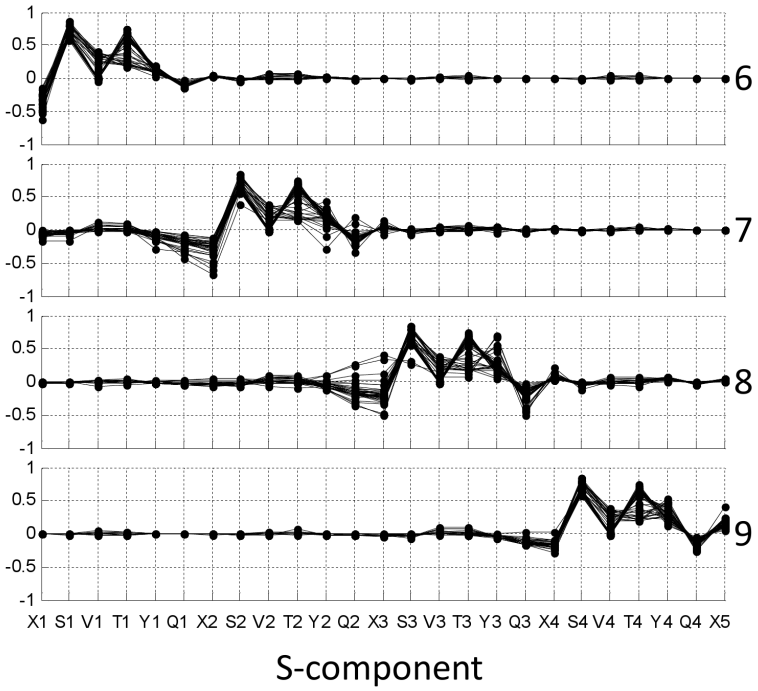

Figure 5.13: Results of PCA performed on four cycles of simulated OKN in series. Component loadings of four principal components extracted from 40 sets of OKN data simulated from eq. 5.1 - 5.9. A total of 13 principal components were extracted during analysis, 4 components are illustrated in this figure as the S-component category.

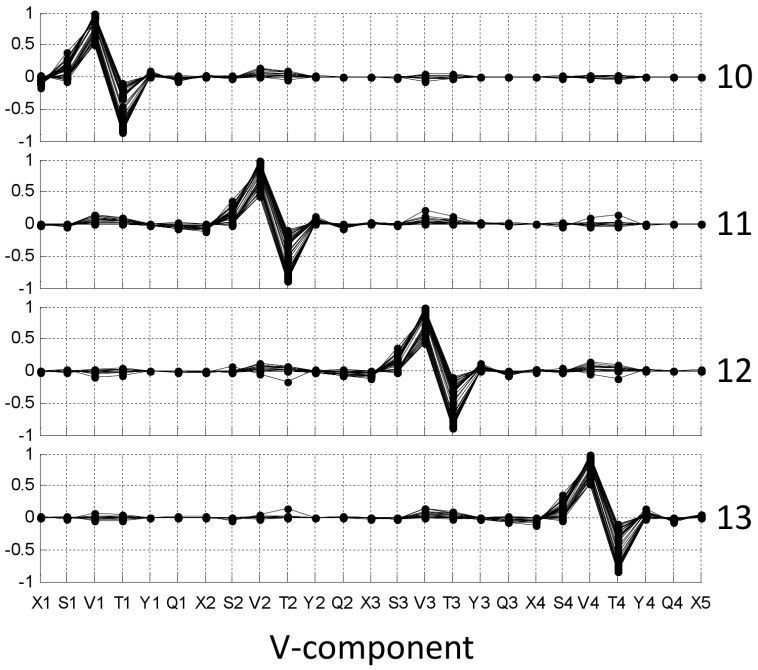

Figure 5.14: Results of PCA performed on four cycles of simulated OKN in series. Component loadings of five principal components extracted from 40 sets of OKN data simulated from eq. 5.1 - 5.9. A total of 13 principal components were extracted during analysis, 4 components are illustrated in this figure as the V-component category.

necessary to consider a qualitative comparison of the simulated data with observed OKN parameters. Pair-wise scatter plots between the simulated OKN parameters were created and assessed for how well the model could describe the indirect as well as the direct relationships in the OKN system (fig. 5.16). The simulated data capture the subtle shape of each scatter plot very accurately (cf. fig. 4.10), such as the non-linear relationship between V_i and T_i, the distributions of OKN parameters, the variability of SP velocity, and certain indirect relationships (such as between S_i and Q_i, and between V_i and x_i).

Auto-correlograms were also created from the simulated data (fig. 5.17) and illustrate the cross-cycle correlations within SP velocity and the start and end position of SPs found empirically (cf.fig. 4.12), despite the update dynamics being only 1^{st} order Markov. The partial auto-correlograms of the simulated data also illustrated the sharp cutoff in the partial autocorrelation function at a lag of 1 (fig. 5.18) seen in the data as expected for a 1^{st} order autoregressive model.

An estimate of the main sequence was used to determine values of QP duration from simulated values of QP amplitude:

$$T_Q = 18.5\,|Q|^{0.54} \tag{5.10}$$

where T_Q is the QP duration (Kaminiarz et al. 2009). Simulated eye position traces were created using the simulated OKN parameters and the estimated values of QP duration. The simulated eye traces were compared to real data and the simulation could capture both the random nature of OKN and other gross traits such as the degree of contraversion and the percentage of QPs made in the direction of stimulus motion (fig. 5.19).

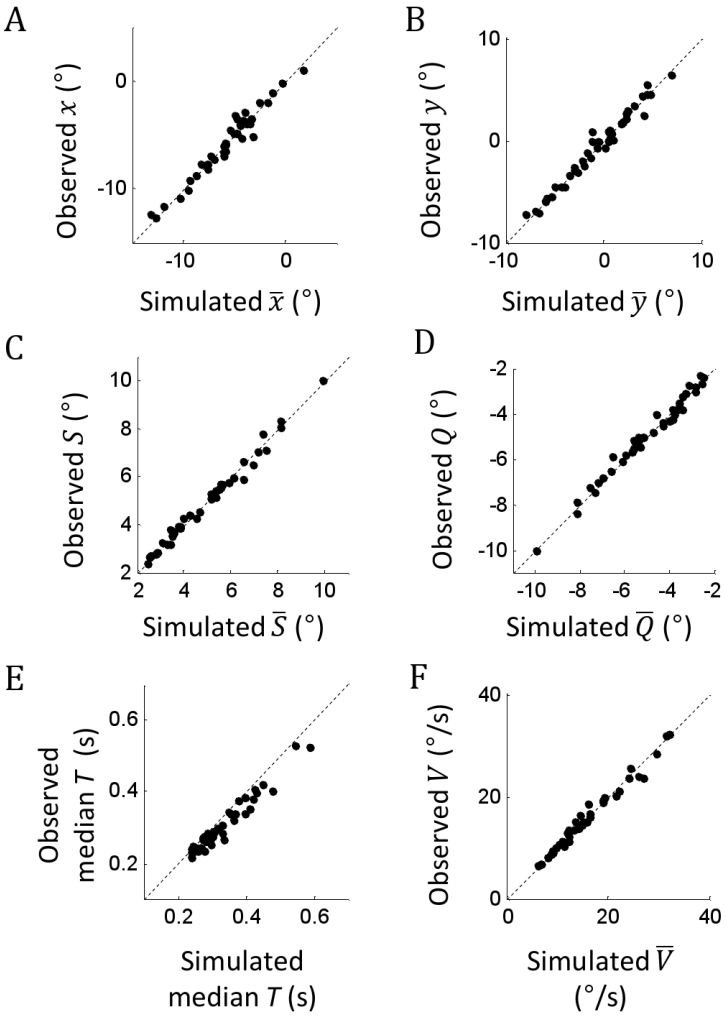

Figure 5.15: Scatter plots of the values of central tendency from the observed OKN parameters for each trial in experiment 1, against the values of central tendency from the simulated OKN parameters. Simulated data were created using the iterative 1st order Markov model. The number of simulated samples of each parameter was set to match the number of observed samples for each parameter. Dotted lines represent identity.

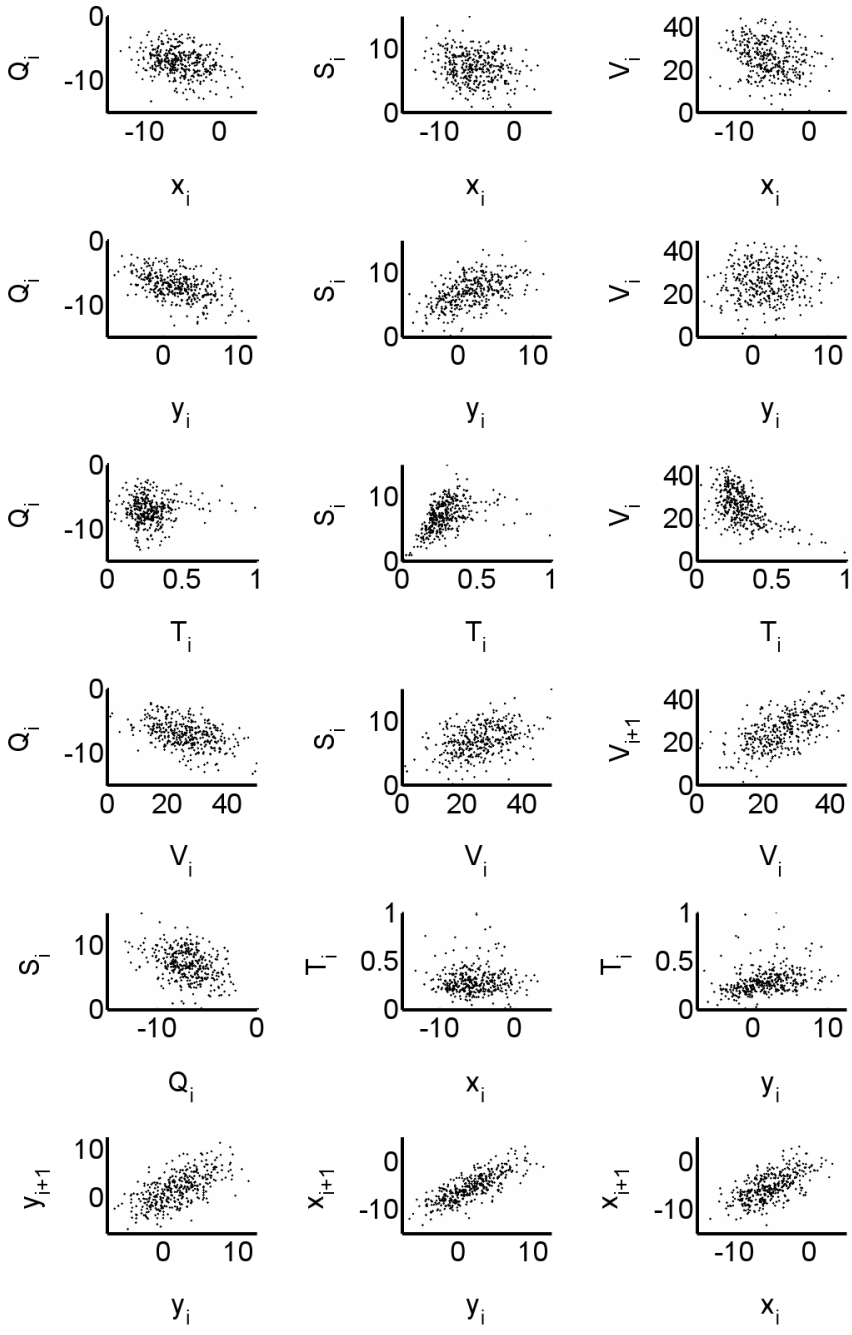

Figure 5.16: Pair-wise scatter plots of simulated data, illustrating the direct and indirect relationships between OKN parameters in the model. Relationships in the simulated data match relationships seen in the empirical data set from which regression coefficients were obtained. Simulated data were created using the iterative 1st order Markov model, with the regression coefficients obtained from participant 5, stimulus speed 40°/s.

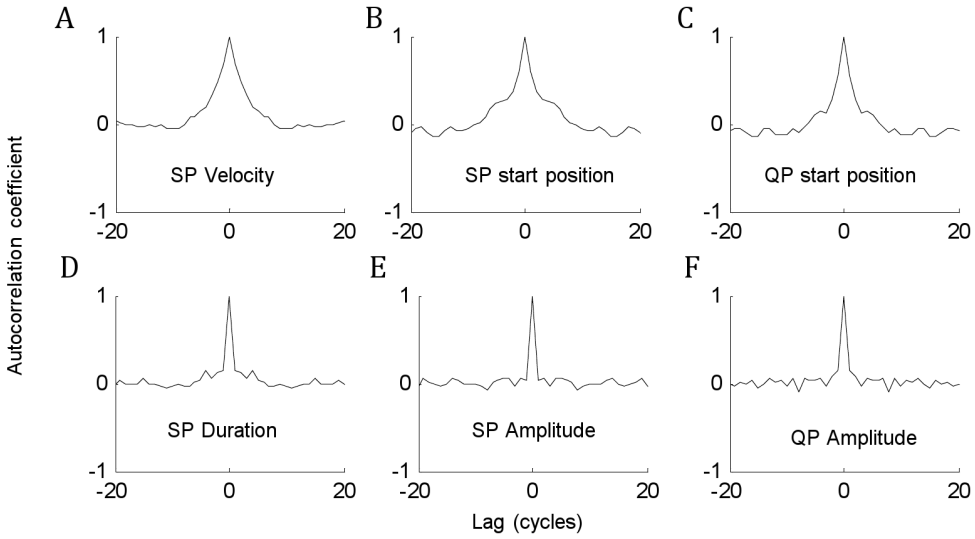

Figure 5.17: Auto-correlograms of simulated (A) SP start position, (B) SP amplitude, (C) SP velocity, (D) SP duration, (E) QP amplitude and (F) QP start position. Data simulated using regression coefficients obtained from participant 1, stimulus speed 40°/s.

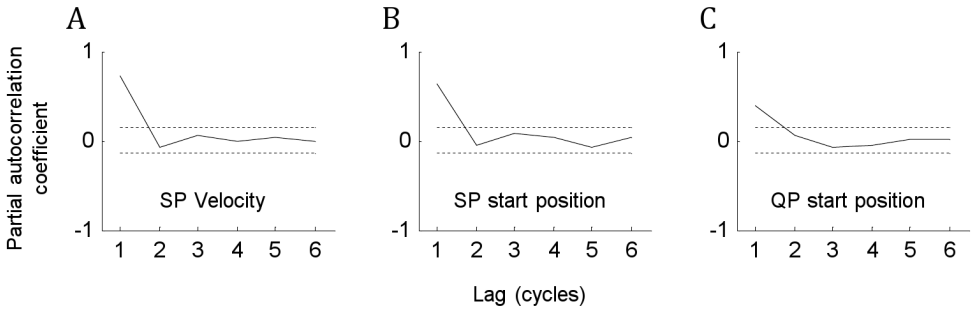

Figure 5.18: Partial auto-correlograms of simulated (A) SP velocity, (B) SP start position and (C) QP start position. Data simulated using regression coefficients obtained from participant 1, stimulus speed 40°/s. Dotted lines, upper and lower confidence limits.

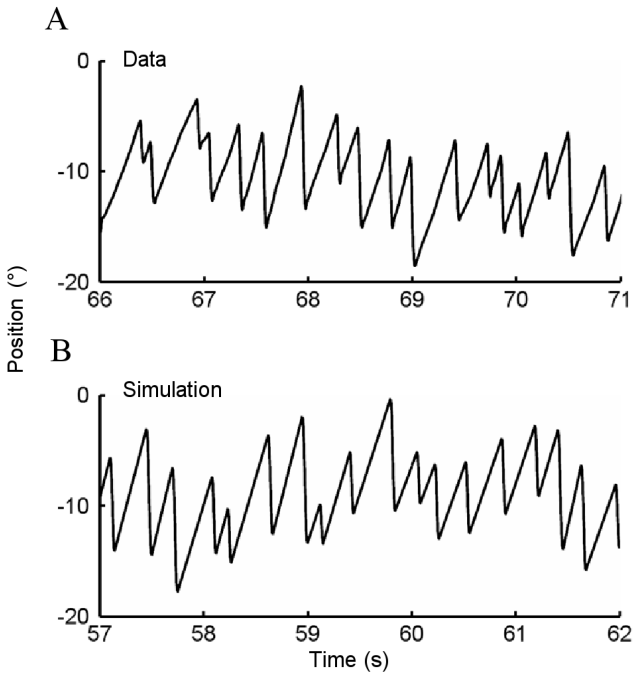

Figure 5.19: Eye position traces from (A) empirical data and (B) a computer simulation of the model. Data is taken from participant 10, stimulus speed 40°/s, and simulation is computed using the regression coefficients estimated from the same data set.

5.9 Model parameters.

Considering the similar loading patterns of the eigenvectors it seemed that relationships between parameters that were within a single OKN cycle could be predominantly the same across trials and participants, and might be considered constants rather than free variables in the model. Taking the mean values of the estimates for a, b, c and d as constant values the values of $\widehat{s}$ and $\widehat{q}$ were recalculated by rearranging the eqs. 5.4 and 5.5 respectively to give:

$$\widehat{s} = S_i - aX_i - bV_i \tag{5.11}$$

and

$$\widehat{q} = Q_i - cY_i - dV_i \tag{5.12}$$

The values of $\widehat{s}$ and $\widehat{q}$ were then calculated separately for each trial, by substituting the mean values of the respective OKN parameters (e.g. $\overline{S}$, $\overline{X}$ and $\overline{V}$ in to S_i, X_i and V_i) for that trial and the constants a, b, c and d into eq. 5.11 and 5.12. The value of e was kept as a free variable because it represented correlations across multiple cycles of OKN, and not within a single cycle. The constants a, b, c and d were calculated to be -0.250, 0.158, -0.478 and -0.166 respectively. The parameters e, $\widehat{s}$, $\widehat{q}$, $\widehat{v}$, ϵ_s, ϵ_q and ϵ_v remained free variables in the model. A table of these free variables can be found in Appendix B.

Mean values of the new simulated parameters were compared with the observed

mean values of OKN parameters, and still demonstrated a very close fit (fig. 5.20) despite the reduced number of free variables in the model. Pair-wise scatter plots between OKN parameters were also created from the new simulated data, and there was no qualitative loss in the description of the OKN parameters (fig. 5.21). Scatter plots of the mean (and median) values of the new simulated parameters were compared to the mean (and median) values of observed parameters, in order to determine if using constant values for a, b, c and d had an effect on predicting the values of observed OKN parameters, and no difference was found.

It was clear that the relationships expressed by the terms ax_i, bV_i, cy_i and dV_i were relatively constant across stimulus speeds and across participants. This was particularly useful as it reduced the number of free parameters in our model from 8 to just 4 (or from 11 to 7 if the standard deviation of the error terms are to be included).

5.10 Autonomous equations for the update dynamics of OKN parameters.

A clearer picture of the system can be obtained by solving eq. 5.1 - 5.6 to find the autonomous update dynamics of x_{i+1}, y_{i+1}, S_{i+1} and Q_{i+1}, given by:

$$\begin{aligned} x_{i+1} =&(1+a)(1+c)x_i + ((1+c)b+d)V_i + (1+c)(\widehat{s}+\epsilon_s(i)) \\ &+ \widehat{q} + \epsilon_q(i), \end{aligned} \tag{5.13}$$

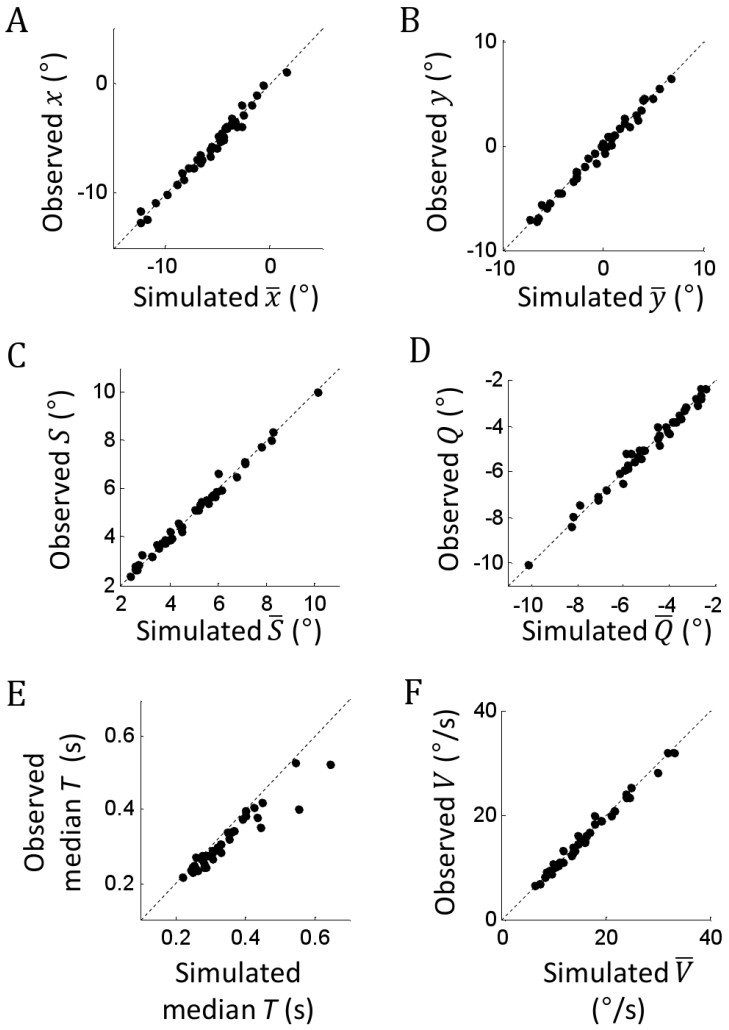

Figure 5.20: Scatter plots of the values of central tendency from the observed OKN parameters for each trial in experiment 1, against the values of central tendency from the simulated OKN parameters. These simulated data were created using Monte Carlo simulations of equations 5.1 - 5.9, where the parameters a, b, c and d were considered constants and not free parameters. The number of simulated samples of each parameter was set to match the number of observed samples for each parameter. Dotted lines represent identity.

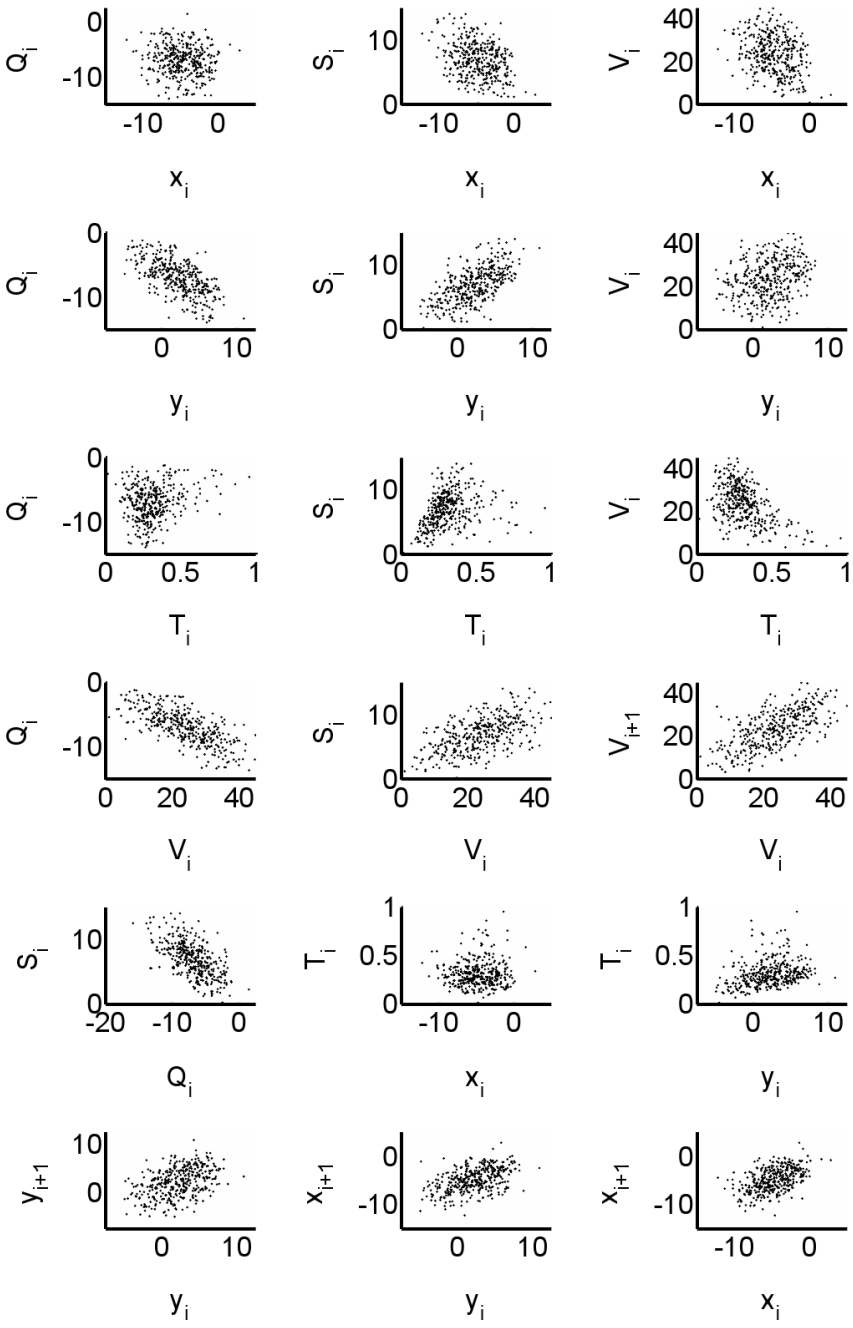

Figure 5.21: Pair-wise scatter plots of simulated data, illustrating the direct and indirect relationships between OKN parameters in the model. These simulated data were created using Monte Carlo simulations of equations 5.1 - 5.9, where the parameters a, b, c and d were considered constants and not free parameters. Relationships in the simulated data match relationships seen in the empirical data set from which regression coefficients were obtained. Data simulated using regression coefficients obtained from participant 5, stimulus speed 40°/s.

$$\begin{aligned} y_{i+1} =&(1+a)(1+c)y_i + ((1+a)d+be)V_i + \widehat{s} + \epsilon_s(i+1) \\ &+ (1+a)(\widehat{q} + \epsilon_q(i)) + b(\widehat{v} + \epsilon_v(i)), \end{aligned} \tag{5.14}$$

$$\begin{aligned} S_{i+1} =&(1+a)(1+c)S_i + (-b(1+c-e)+ad)V_i - c(\widehat{s} + \epsilon_s(i)) \\ &+ a(\widehat{q} + \epsilon_q(i)) + b(\widehat{v} + \epsilon_v(i)), \end{aligned} \tag{5.15}$$

and

$$\begin{aligned} Q_{i+1} =&(1+a)(1+c)Q_i + (-d(1+a-e)+cbe)V_i + c(\widehat{s} + \epsilon_s(i+1)) \\ &- a(\widehat{q} + \epsilon_q(i)) + (cb+d)(\widehat{v} + \epsilon_v(i)), \end{aligned} \tag{5.16}$$

where

$$V_{i+1} = eV_i + \widehat{v} + \epsilon_v(i).$$

From these equations it is clear that each of these parameters depends on its value only in the previous cycle, and also the SP velocity in the previous cycle. In other words, the dynamics are 1^{st} order Markov and driven by SP velocity, which is itself another 1^{st} order Markov process. The variables x_{i+1}, y_{i+1}, S_{i+1} and Q_{i+1} are all correlated with each other through a common dependence on SP velocity

and the three noise sources $\epsilon_s(i)$, $\epsilon_q(i)$ and $\epsilon_v(i)$.

5.11 Discussion of results.

The underlying structure of the correlation matrices was examined using PCA and remarkably similar eigenvectors across subjects and conditions were found that were obviously clustered into three categories. The eigenvalues were much more variable than the eigenvectors between trials, often affecting the order in which eigenvectors were expressed (in terms of their explained variance) on a scree plot. It appears that the variation in the correlation matrices between trials is primarily caused by the variation in the eigenvalues associated with each eigenvector rather than the eigenvectors themselves. In other words the differences found in the correlation matrixes are caused by the differences in noise associated with SP velocity, SP amplitude and QP amplitude between participants and between stimulus speeds. In the terms of our model, it is the varying values of $\epsilon_s(i)$, $\epsilon_q(i)$ and $\epsilon_v(i)$ that cause the differences in the correlation matrices, not the values of a, b, c and d.

Only three categories of components were detected and retained in this analysis. PCA is often used as a dimension-reducing scheme, in which only a certain percentage of the variance is explained by the components retained. In this analysis the low number of categories of eigenvectors found emerged from the data due to the DOF imposed by the geometric relationships between parameters, and was not due to any deliberate dimension-reducing scheme. Almost 100% of the variance was explained in the components that were retained, with only a slight loss due to the non-linear relationship between SP velocity and SP duration. We conclude that OKN is a stochastic process with three uncorrelated noise sources, which we label as the "S", "Q" and "V" components (eq. 5.4 - 5.6).

Parameters such as the start and end position of SPs and QPs are correlated over several OKN cycles, and originally direct correlations among parameters across many OKN cycles had been anticipated, as a long-term adaptive mechanism could explain why average eye position does not randomly wander off to the limit of gaze. Surprisingly, only correlations across adjacent cycles were apparent indicating no memory of parameters across cycles (1^{st} order Markov process). It seems that the Q and S components control mean eye position in a rather simple fashion. QP amplitude is negatively correlated with the start position of the QP determined by c in eq. 5.5. QP amplitude will therefore tend to correct for the end position of the previous SP. Similarly, SP amplitude is negatively correlated with SP start position as determined by a in eq. 5.4 such that it tends to correct for end position of the previous QP. It appears that the Q and S components cooperate to maintain mean eye position, via a rapid start position "feedback".

The start of a QP is determined by the saccadic trigger, which could be determined by SP duration, eye position, or a combination of processes. In this model the threshold for the generating a QP is the SP amplitude, which is determined by the SP velocity and SP start position. It is tempting to consider the dependence of SP amplitude on SP velocity as a refractory period during which QPs are not triggered. The mean value of b estimated from the data was 0.16, which could indicate an average minimum duration for SPs of 160ms. However, this conclusion does require careful consideration of the other terms expressed in eq. 5.4 particularly as SPs clearly do occur under 160ms in duration. The proposed model also predicts that QP amplitude is dependent on SP velocity and the QP start position. Remarkably, it seems that the values of all these correlations are relatively constant, even between participants and stimulus speeds. The threshold for generating a vestibular QP in humans has been reported to depend on instantaneous

eye velocity as well as eye position (Lau et al. 1978). The dependence of both SP and QP amplitude on eye position and eye velocity was also observed in the vestibular nystagmus of rabbits (Lau & Honrubia 1986), but with a much greater dependence of QP amplitude than SP amplitude on SP velocity.

A remarkable finding was that SP velocity can be described by an autonomous 1st order Markov process. V_i does not depend on any of the other OKN parameters (although they all depend on V_i). V_i depends on V_{i+1}, but not explicitly on the state of earlier OKN cycles. It is difficult to reconcile the Markovian behaviour of SP velocity with the conventional view of OKN as a deterministic control system driven by retinal slip. This type of system requires internal feed-forward and/or feedback open loop gains, together with loop delays, to place limits on the steady-state closed loop gain of the system. However, when we examine OKN cycle-by-cycle, we observe gain changing haphazardly as SP velocity wanders in Markov fashion often far above or far below the mean gain. Clearly, high gain can occur (i.e. SP velocity is not limited within the range of speeds we have investigated), so then why is a high gain not always maintained? In the next chapter we investigate further the interaction of SP velocity with the other OKN parameters, and consider the contribution of all three sources of noise to the system. SP velocity can be described by an autonomous 1st order Markov process, but all other OKN parameters that we have investigated are dependent on SP velocity and this has important consequences on how the system behaves.

Chapter 6

Results of experiments 2 and 3: the description and predictions of OKN behaviour.

In this chapter the results of RM-ANOVA performed on OKN parameters recorded from experiments 2 and 3 are presented. The main and interaction effects of changing stimulus parameters on the mean and standard deviation of OKN parameters are reported. The within-subjects factors tested were the speed, spatial frequency, and pattern of the OKN stimulus. The previous chapter reported the results of performing PCA on OKN data from experiment 1, and from these results we developed a Markov model of the OKN system that could replicate the empirical behaviour. Here, we apply the Markov model to the OKN parameters recorded from experiment 2 and 3 separately, in order to demonstrate that the model is capable of fitting all of the data. We then use the combined data from all experiments, and the model, to make a number of observations and predictions about the behaviour of the OKN system.

6.1 RM-ANOVA results for experiment 2: effects of spatial frequency

6.1.1 SP velocity, retinal slip, and gain.

Figure 6.1 contains three boxplots demonstrating how the mean SP velocity, retinal slip and SP gain were affected by spatial frequency and speed of the stimulus. A three-way RM-ANOVA was performed using each parameter in turn as the dependant variable, and spatial frequency, speed, and direction of the stimulus as the within-subjects factors.

There was a significant main effect of stimulus speed on the mean SP velocity ($F(1,7)$=11.5, p=0.012), retinal slip ($F(1,7)$=26.8, p=0.001), and SP gain ($F(1,7)$=18.7, p=0.003), verifying the results of experiment 1.

There was also a significant main effect of spatial frequency on the mean SP velocity ($F(2,14)$=6.0, p=0.013), retinal slip ($F(2,14)$=6.1, p=0.013), and SP gain ($F(2,14)$=6.9, p=0.008). However, Bonferroni corrected confidence intervals could not indicate a significant difference in pairwise comparisons of mean SP velocity, retinal slip, or SP gain, between presentations of different spatial frequency stimuli.

There was no significant interaction effect between spatial frequency and stimulus velocity on the mean SP velocity ($F(2,14)$=13.5, p=0.152), retinal slip ($F(2,14)$= 16.3, p=0.138), or SP gain ($F(2,14)$=0.0, p=0.931).

There was no significant main effect of direction on the mean SP velocity ($F(1,7)$= 0.0, p=0.938), retinal slip ($F(1,7)$=0.0, p=0.897), or SP gain ($F(1,7)$=0.0, p= 0.812). There were also no significant interaction effects between direction and any other factor on the mean SP velocity, retinal slip or gain.

One-way RM-ANOVA performed on the standard deviation of these parameters

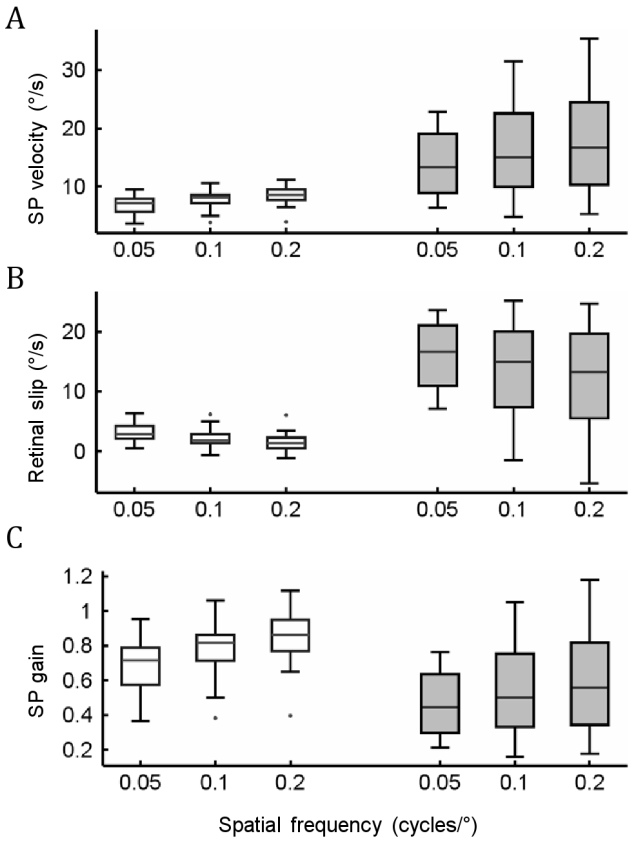

Figure 6.1: Boxplot illustrating the effect of spatial frequency and stimulus speed on (A) SP velocity, (B) retinal slip, and (C) SP gain). Each box represents the parameter grouped across all participants for each stimulus condition. Stimulus conditions tested were: stimulus speed, 10°/s or 30°/s; spatial frequency, 0.05cyc/°, 0.1cyc/°, or 0.2cyc/°. Grey boxes represent data from trials with a stimulus speed of 30°/s, white boxes represent data from trials with a stimulus speed of 10°/s. Box represents interquartile range, whiskers represent maximum and minimum data values (excluding outliers), and dots represent outliers found more than 1.5 times the interquartile range from the ends of the box.

found a main effect of stimulus speed on the standard deviation of SP velocity and retinal slip ($F(1,7)$=41.3, p<0.001), but not SP gain ($F(1,7)$=1.0, p=0.239), as before.

In summary, the results of varying stimulus speed verified the findings in experiment 1 that there was a main effect of stimulus speed on mean values of SP velocity, retinal slip and SP gain, as well as the standard deviation of SP velocity and retinal slip. Increasing the spatial frequency of the stimulus resulted in an increase in mean SP velocity and gain, and a decrease in retinal slip.

6.1.2 SP amplitude, QP amplitude, and the start and end points of QPs.

Figure 6.2 contains four boxplots demonstrating how the mean SP amplitude, QP amplitude, QP start position and QP end position were affected by spatial frequency and speed of the stimulus. A three-way RM-ANOVA was performed using each parameter in turn as the dependant variable, and spatial frequency, speed, and direction of the stimulus as the within-subjects factors.

There was a significiant main effect of stimulus speed on the mean SP amplitude ($F(1,7)$=17.2, p=0.004) and QP amplitude ($F(1,7)$=18.8, p=0.003), but not on QP start position ($F(1,7)$=0.2, p=0.664) or QP end position ($F(1,7)$=1.2, p=0.309), verifying the results of experiment 1.

There was a significant main effect of spatial frequency on the mean SP amplitude ($F(2,14)$=9.1, p=0.003) and QP amplitude ($F(2,14)$=10.2, p=0.002), but not on QP start position ($F(2,14)$=1.0, p=0.374) or QP end position ($F(2,14)$=0.8, p=0.475). Corrected confidence intervals indicated a significant difference in pairwise comparisons of mean SP amplitude (p=0.016) and QP amplitude (p=0.012) between presentations of stimuli with a spatial frequency of 0.05cyc/° and 0.2cyc/°.

There was a significant interaction effect between spatial frequency and stimulus

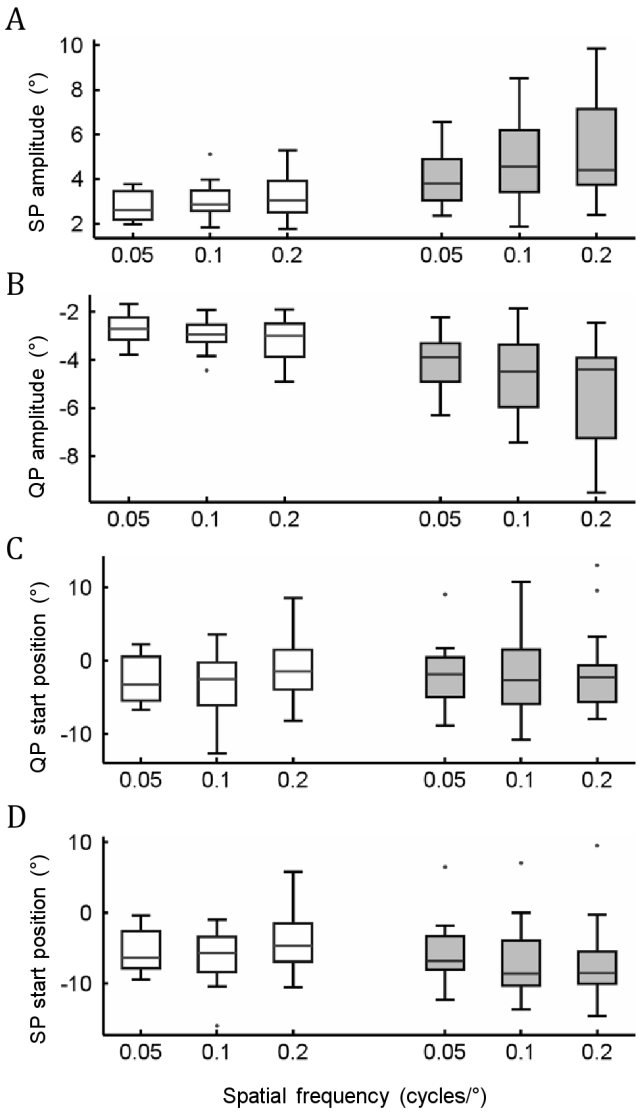

Figure 6.2: Boxplot illustrating the effect of spatial frequency and stimulus speed on (A) SP amplitude, (B) QP amplitude, (C) QP start position, and (D) QP end position. Each box represents the parameter grouped across all participants for each stimulus condition. Stimulus conditions tested were: stimulus speed, 10°/s or 30°/s; spatial frequency, 0.05cyc/°, 0.1cyc/° or 0.2cyc/°. Grey boxes represent data from trials with a stimulus speed of 30°/s, white boxes represent data from trials with a stimulus speed of 10°/s. Box represents interquartile range, whiskers represent maximum and minimum data values (excluding outliers), and dots represent outliers found more than 1.5 times the interquartile range from the ends of the box.

speed on the mean SP amplitude ($F(2,14)=4.7$, $p=0.028$) and QP amplitude ($F(2,14)=5.6$, $p=0.016$), but not on QP start position ($F(2,14)=1.2$, $p=0.322$) nor on QP end position ($F(2,14)=2.9$, $p=0.087$).

There was no significant main effect of direction on the mean SP amplitude ($F(1,7)=0.0$, $p=0.955$), QP amplitude ($F(1,7)=0.0$, $p=0.912$), QP start position ($F(1,7)=2.0$, $p=0.200$), or QP end position ($F(1,7)=2.0$, $p=0.199$). There were also no significant interaction effects between direction and any other factor on the mean SP amplitude, QP amplitude, QP start position, or QP end position.

Three-way RM-ANOVA on the standard deviation of these parameters found a main effect of stimulus speed on the standard deviation of SP amplitude ($F(1,7)=46.1$, $p<0.001$), QP amplitude ($F(1,7)=27.0$, $p=0.001$), QP start position ($F(1,7)=18.4$, $p=0.004$) and QP end position ($F(1,7)=29.9$, $p=0.001$). There was also a significant interaction effect between spatial frequency and stimulus velocity on the standard deviation of QP amplitude ($F(2,14)=4.7$, $p=0.027$) but not on SP amplitude ($F(2,14)=2.5$, $p=0.114$).

In summary, the results of varying stimulus speed verified the findings in experiment 1 that there was a main effect of stimulus speed on mean SP and QP amplitude, as well as the standard deviation of SP and QP amplitude, and the start and end position of QPs (and SPs). Increasing the spatial frequency of the stimulus resulted in an increase in the mean SP and QP magnitude, but had no effect on the mean start and end position of QPs (and SPs). Increasing the spatial frequency also resulted in an increase in the standard deviation of SP and QP magnitude, and an increase in the standard deviation of the start and end position of QPs (and SPs). There was also a significant interaction effect between spatial frequency and stimulus velocity on the mean SP and QP amplitude, indicating a possible temporal frequency effect of the stimulus.

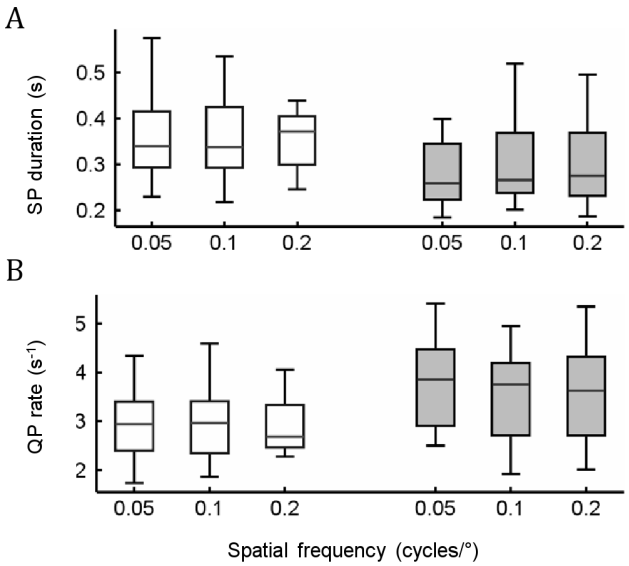

Figure 6.3: Boxplot illustrating the effect of spatial frequency and stimulus speed on (A) SP duration and (B) QP rate. Each box represents the parameter grouped across all participants for each stimulus condition. Stimulus conditions tested were: stimulus speed, 10°/s or 30°/s; spatial frequency, 0.05cyc/°, 0.1cyc/° or 0.2cyc°. Grey boxes represent data from trials with a stimulus speed of 30°/s, white boxes represent data from trials with a stimulus speed of 10°/s. Box represents interquartile range, whiskers represent maximum and minimum data values (excluding outliers), and dots represent outliers found more than 1.5 times the interquartile range from the ends of the box.

6.1.3 SP duration and QP rate.

Figure 6.3 contains two boxplots demonstrating how the median SP duration and QP rate are affected by spatial frequency and speed of the stimulus. RM-ANOVA was performed using the aligned rank-transformed median values of each parameter as the dependent variables, and spatial frequency, speed, and direction of the stimulus as the within-subjects factors.

There was a significant main effect of stimulus speed on SP duration ($F(1,7)=15.1$, $p=0.006$) and on QP rate ($F(1,7)=14.0$, $p=0.007$), verifying the results of experiment 1.

There was no significant main effect of spatial frequency on SP duration ($F(2,14)=0.3$, $p=0.717$) or QP rate ($F(2,14)=1.0$, $p=0.388$), and there was no significant main effect of direction on SP duration ($F(1,7)=0.0$, $p=0.867$) or QP rate ($F(1,7)=0.1$, $p=0.735$).

There was no significant interaction effect of spatial frequency and stimulus velocity on SP duration ($F(2,14)=0.8$, $p=0.459$) or QP rate ($F(2,14)=0.0$, $p=0.933$). There were also no significant interaction effects between direction and any other factor on the SP duration or QP rate.

6.2 Investigating the principal components found in experiment 2.

The same PCA procedure applied to the data from experiment 1 was applied to the data from experiment 2. The analysis was performed on the correlation matrixes of parameters from across four adjacent cycles. However, in one trial there were such a large number of blinks that only 8 cycles of data could be retrieved, and it was not possible to find four consecutive cycles of OKN at all, so this trial was ignored for the purposes of PCA analysis. PCA performed on the correlation matrixes of the other 95 trials did reveal 13 significant eigenvalues, corresponding to the maximum number of DOF (fig. 6.4a). The heuristic algorithm that had been applied to the components found in the data from experiment 1 performed well in sorting the components found in experiment 2 into individual categories, but not perfectly as 3 components (out of 1248) did not fit the sorting criteria for any of the categories. The heuristic criteria for sorting were then changed such that all the components found in the data from experiment 2 (and experiment 1) could be placed into individual categories. These results indicated that there were 13 qualitatively different components within the data from each trial in experiment 2, but they could be sorted into similar categories as found in experiment 1 (i.e.

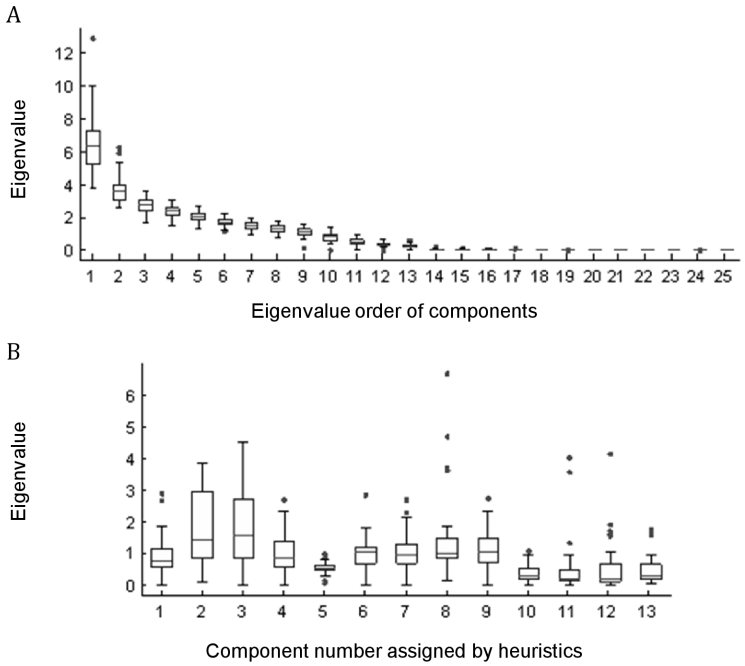

Figure 6.4: (A) Boxplot illustrating the eigenvalues of principal components found in four cycles of OKN parameters in series before being sorted by heuristics. Number on principal component axis represents the eigenvalue order of components from largest to smallest. (B) Boxplot illustrating the eigenvalues of principal components extracted from four cycles of OKN parameters in series after being sorted by heuristics. Sorted component numbers 1-5, "Q-component"; 6-9, "S-component"; 10-13, "V-component". Boxes represents interquartile range, whiskers represent maximum and minimum data values (excluding outliers), and dots represent outliers found more than 1.5 times the interquartile range from the ends of the box. OKN data taken from experiment 2 only.

they had similar loading patterns). As in the results from experiment 1, each of these sorted components had variable eigenvalues between trials (fig. 6.4b).

6.3 Results of applying the model to OKN parameters recorded from experiment 2.

In the previous chapter we demonstrated that the parameters a, b, c and d could be considered constants rather than free variables in the model, and did not ap-

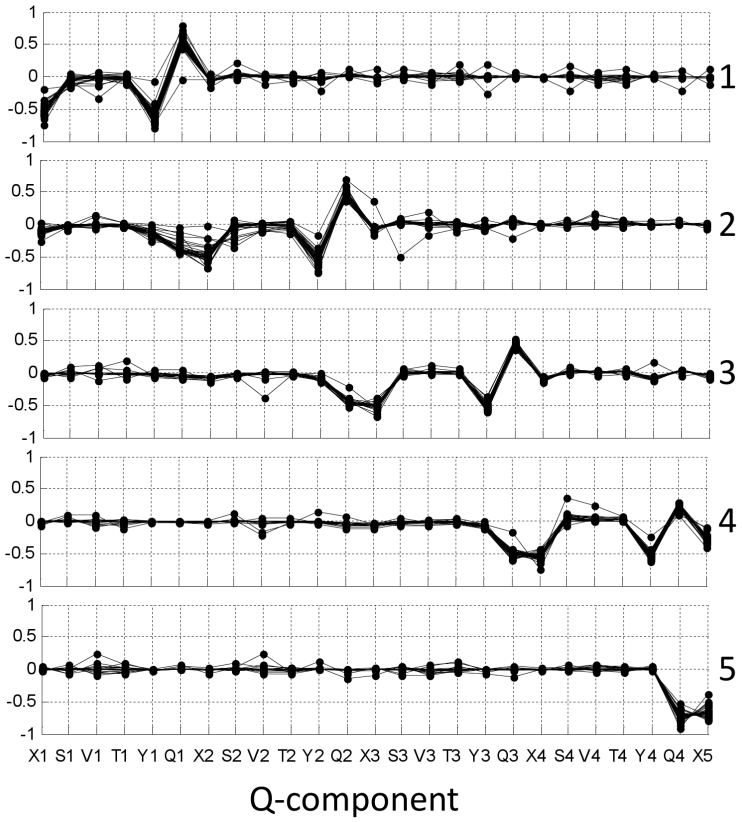

Figure 6.5: Results of PCA performed on four cycles of OKN in series, including SP velocity and duration. Component loadings of 5 principal components extracted from OKN data from experiment 2 (95 trials). A total of 13 principal components were extracted during analysis. These 5 components verify the existence of the Q-component category in the data recorded from experiment 2. Numbers are used only to label components and do not represent the position of the corresponding eigenvalue on the scree plot as these varied between trials.

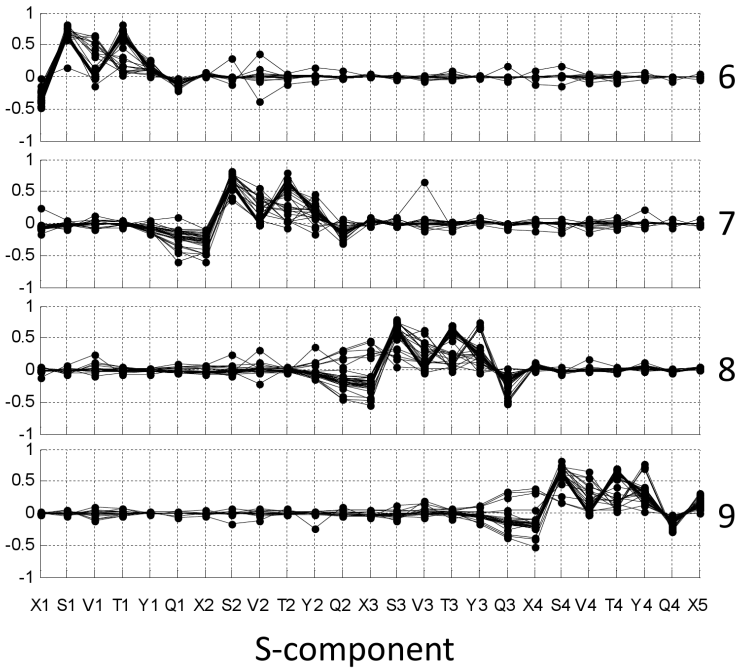

Figure 6.6: Results of PCA performed on four cycles of OKN in series, including SP velocity and duration. Component loadings of 4 principal components extracted from OKN data from experiment 2 (95 trials). A total of 13 principal components were extracted during analysis. These 4 components verify the existence of the S-component category in the data recorded from experiment 2. Numbers are used only to label components and do not represent the position of the corresponding eigenvalue on the scree plot as these varied between trials.

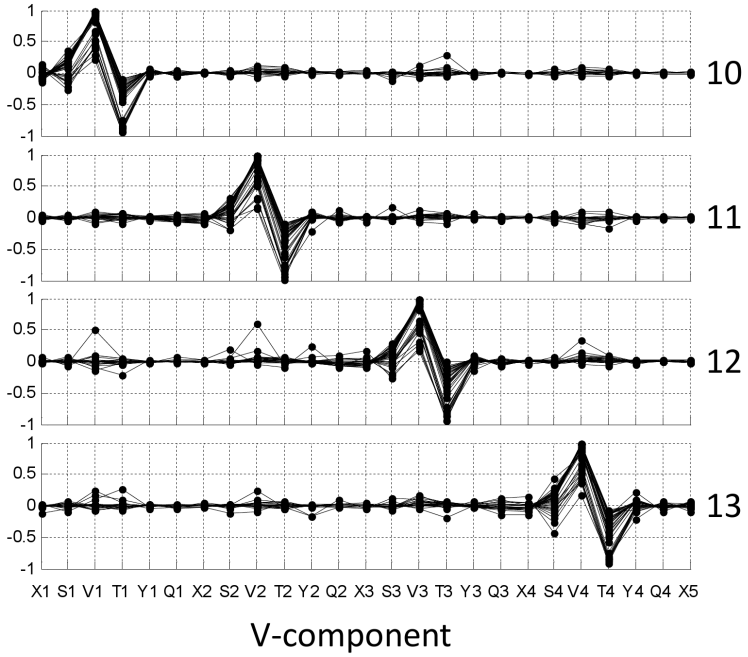

Figure 6.7: Results of PCA performed on four cycles of OKN in series, including SP velocity and duration. Component loadings of 4 principal components extracted from OKN data from experiments 2 (95 trials). A total of 13 principal components were extracted during analysis. These 4 components verify the existence of the V-component category in the data recorded from experiment 2. Numbers are used only to label components and do not represent the position of the corresponding eigenvalue on the scree plot as these varied between trials.

pear to be affected by stimulus velocity. In order to determine if these parameters changed with spatial frequency, multiple linear regression (weighted least squares) was performed on the OKN parameters recorded from experiment 2. Regression was performed on SP amplitude, SP start position and SP velocity in order to estimate the coefficients a and b ($S_i = ax_i + bV_i + \widehat{s} + \epsilon_s(i)$). Regression was performed on QP amplitude, QP start position and SP velocity in order to estimate the coefficients c and d ($Q_i = cy_i + dV_i + \widehat{q} + \epsilon_q(i)$). The mean values of a, b, c and d found from regression of the parameters in experiment 2 were -0.251 (SD=0.20), 0.124 (SD=0.20), -0.397 (SD=0.21) and -0.197 (SD=0.16) respectively. Note that the mean values of a, b, c and d found in experiment 1 were -0.250 (SD=0.15), 0.158 (SD=0.16), -0.478 (SD=0.27) and -0.166 (SD=0.10) respectively, so there appeared to be no significant difference (p>0.089) in these values between experiments 1 and 2.

In order to determine if these non-significant differences would have an effect on the predictive power of the model the values of the constants from experiment 1 were substituted in to the model, and simulated parameters were created using the free variables estimated from the trials in experiment 2. The simulated parameters were then compared to the observed OKN parameters from those trials. The values of e, $\widehat{v}$, σ_s, σ_q and σ_v from each trial were estimated with robust linear regression performed on the OKN parameters recorded from experiment 2. The values of $\widehat{s}$ and $\widehat{q}$ were calculated by substituting the constants and mean OKN parameters into eq. 5.4 and 5.5. For example, $\overline{S}$, $\overline{x}$ and $\overline{V}$ were substituted into $\widehat{s} = S_i - ax_i - bV_i + \epsilon_s(i)$, with the values of a and b from experiment 1, to find $\widehat{s}$.

Simulated data sets were created iteratively in a Matlab function with eqs. 5.1 - 5.9 and a number of simulated parameters were created to match the number of observed samples of each parameter. Then the mean (and median) values of

simulated parameters were compared with the mean (and median) values of the observed data (fig. 6.8). Also, the standard deviation (and interquartile range) of simulated parameters were compared with the standard deviation (and interquartile range) of observed data (fig. 6.9). R^2 for mean (and median) values of OKN parameters range from 0.78 to 1.00, and the R^2 for standard deviation (and interquartile range) values range from 0.55 to 0.80. This indicated that in both cases the simulated values accurately reflected the observed values, and that the model could capture the distribution of OKN parameters well with only four free variables in the model (or 7, including the estimated values of standard deviation in the underlying components).

6.4 RM-ANOVA results for experiment 3: effects of stimulus pattern type.

6.4.1 SP velocity, retinal slip, and gain.

Figure 6.10 contains three boxplots demonstrating how the mean SP velocity, retinal slip and SP gain were affected by different stimulus patterns and by changing the spatial frequency of those stimulus patterns. A three-way RM-ANOVA was performed using each parameter in turn as the dependant variable, and stimulus pattern, spatial frequency, and direction of the stimulus as the within-subjects factors.

There was a significant main effect of spatial frequency on the mean SP velocity, retinal slip, and SP gain ($F(2,18)=5.7$, $p=0.012$). However, corrected confidence intervals could not indicate a significant difference in pairwise comparisons of presentations with different spatial frequency stimuli.

There was also a significant main effect of stimulus pattern on mean SP velocity, retinal slip, and SP gain ($F(1,9)=19.6$, $p=0.002$).

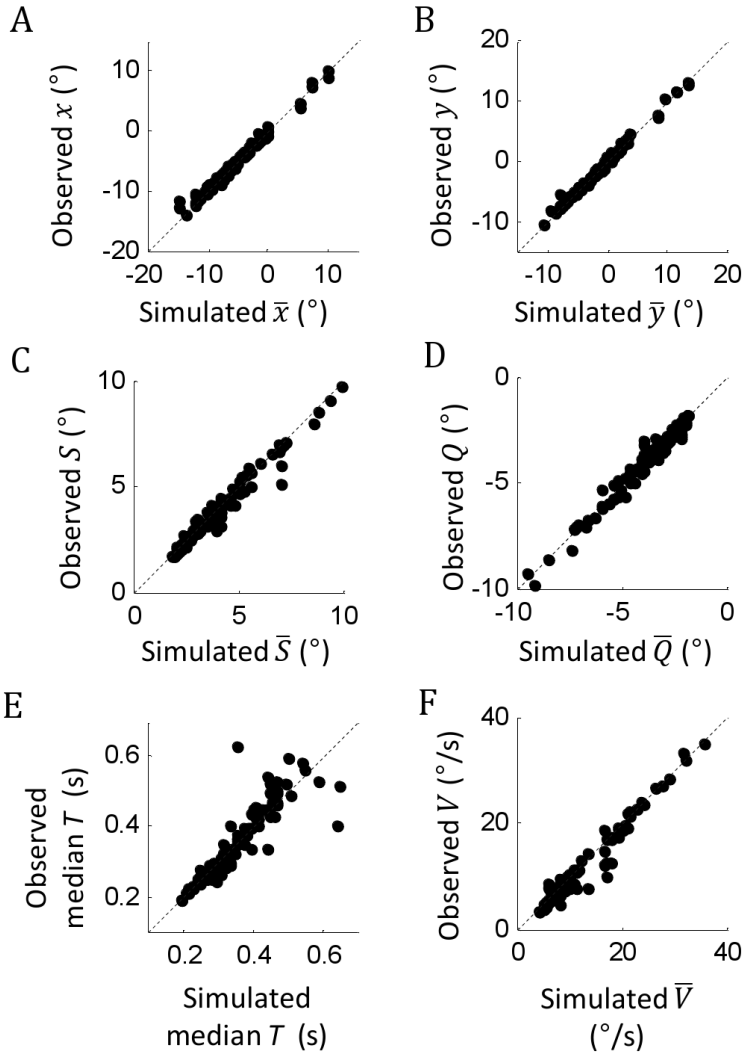

Figure 6.8: Scatter plots of the values of central tendency from the observed OKN parameters for each trial in experiment 2, against the values of central tendency from the simulated OKN parameters. These simulated data were created using Monte Carlo simulations of equations 5.4 - 5.6, where the parameters a, b, c and d were estimated from experiment 1. The number of simulated samples of each parameter was set to match the number of observed samples for each parameter. Dotted lines represent identity.

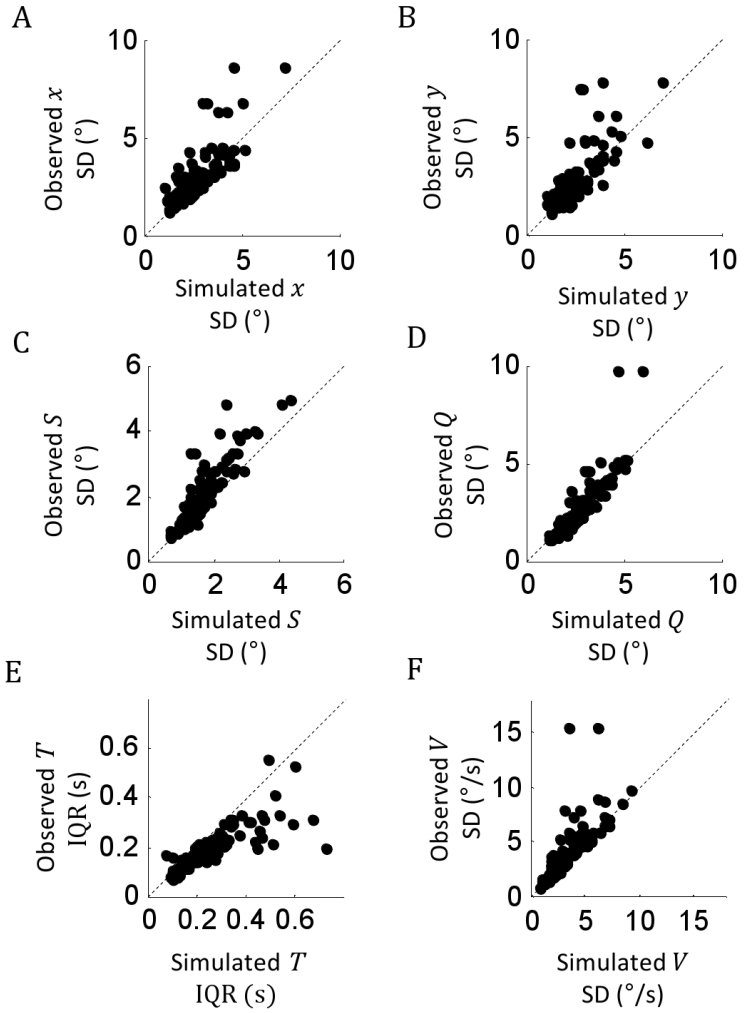

Figure 6.9: Scatter plots of the values of the measure of dispersion from the observed OKN parameters for each trial in experiment 2, against the values of the measure of dispersion from the simulated OKN parameters. These simulated data were created using Monte Carlo simulations of equations 5.4 - 5.6, where the parameters a, b, c and d were estimated from experiment 1. The number of simulated samples of each parameter was set to match the number of observed samples for each parameter. SD, standard deviation; IQR, Inter-quartile range. Dotted lines represent identity.

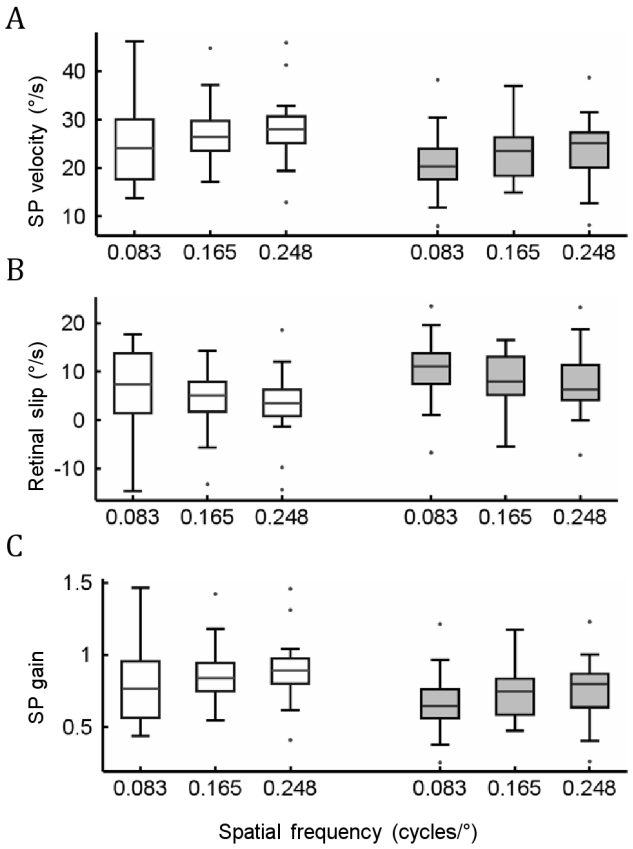

Figure 6.10: Boxplot illustrating the effect of spatial frequency and stimulus pattern on (A) SP velocity, (B) retinal slip, and (C) SP gain. Each box represents the parameter grouped across all participants for each stimulus condition. Stimulus conditions tested were: stimulus pattern, square wave or random grid; spatial frequency, 0.083cyc/°, 0.165cyc/°or 0.248cyc/°. Grey boxes represent data from trials with a square wave stimulus pattern, white boxes represent data from trials with a random grid stimulus pattern. Box represents interquartile range, whiskers represent maximum and minimum data values (excluding outliers), and dots represent outliers found more than 1.5 times the interquartile range from the ends of the box.

As before, there was no significant effect of direction on mean SP velocity, retinal slip, or SP gain ($F(1,9)$=1.0, p=0.282). There were also no significant interaction effects between any of the parameters on SP velocity, retinal slip, or SP gain.

Three-way RM-ANOVA performed on the standard deviation of these parameters found a main effect of spatial frequency on the standard deviation of SP velocity, retinal slip, and SP gain ($F(2,18)$=5.4, p=0.015).

In summary, the results of varying spatial frequency verified the findings in experiment 2 that there was a main effect of spatial frequency on mean values of SP velocity, retinal slip and SP gain. Using a random grid stimulus pattern resulted in an increase in mean SP velocity and gain, and a decrease in retinal slip, when compared with using a square wave stimulus pattern.

6.4.2 SP amplitude, QP amplitude, and the start and end points of QPs.

Figure 6.11 contains four boxplots demonstrating how the mean SP amplitude, QP amplitude, QP start position and QP end position are affected by different stimulus patterns and by changing the spatial frequency of those stimulus patterns. A three-way RM-ANOVA was performed using each parameter in turn as the dependant variable, and stimulus pattern, spatial frequency, and direction of the stimulus as the within-subjects factors.

There was a significant main effect of spatial frequency on the mean SP amplitude ($F(2,18)$=7.6, p=0.004) and QP amplitude ($F(2,18)$=7.0, p=0.006), but not for QP start position ($F(2,18)$=0.7, p=0.507) or QP end position ($F(2,18)$=0.1, p=0.906). Corrected confidence intervals indicated a significant difference in pairwise comparisons of SP amplitude between presentations with a spatial frequency of 0.083cyc/° and 0.165cyc/° (p=0.041), and 0.083cyc/° and 0.248cyc/° (p=0.022). Corrected confidence intervals also indicated a significant difference in

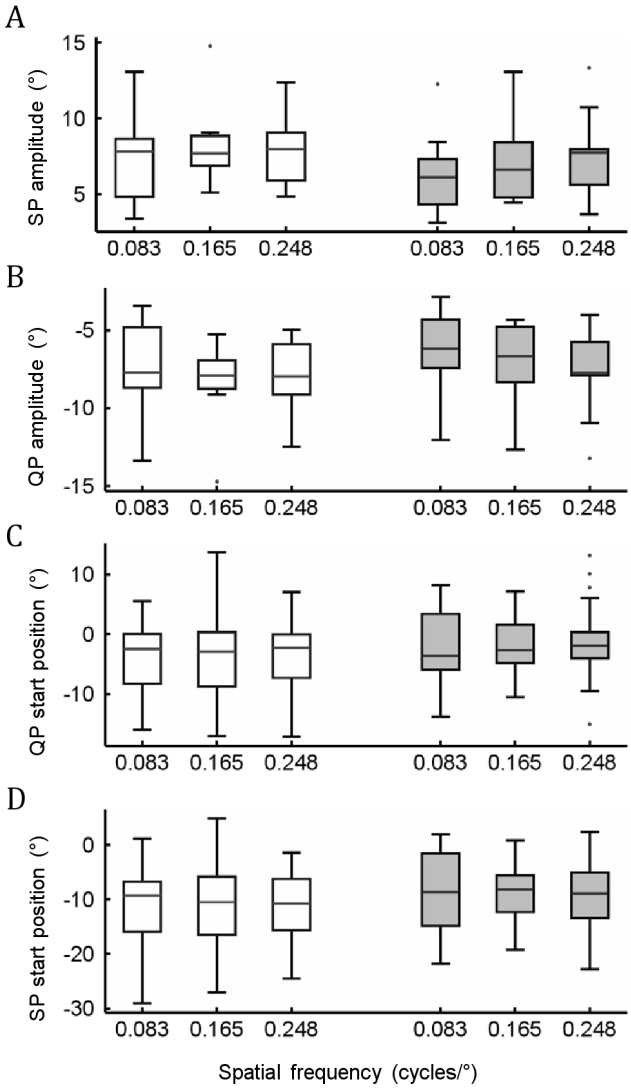

Figure 6.11: Boxplot illustrating the effect of spatial frequency and stimulus pattern on (A) SP amplitude, (B) QP amplitude, (C) QP start position, and (D) QP end position. Each box represents the parameter grouped across all participants for each stimulus condition. Stimulus conditions tested were: stimulus pattern, square wave or random grid; spatial frequency, 0.083cyc/°, 0.165cyc/°or 0.248cyc/°. Grey boxes represent data from trials with a square wave stimulus pattern, white boxes represent data from trials with a random grid stimulus pattern. Box represents interquartile range, whiskers represent maximum and minimum data values (excluding outliers), and dots represent outliers found more than 1.5 times the interquartile range from the ends of the box.

QP amplitude between presentations with a spatial frequency of 0.083cyc/° and 0.165cyc/° (p=0.041), and between 0.083cyc/° and 0.248cyc/° (p=0.022).

There was a significant main effect of stimulus pattern on mean SP amplitude (F(1,9)=16.8, p=0.003), QP amplitude (F(1,9)=15.1, p=0.004), QP start position (F(1,9)=6.1, p=0.035), and QP end position (F(1,9)=10.5, p=0.010).

As before, there was no significant main effect of direction on mean SP amplitude (F(1,9)=14.0, p=0.165), QP amplitude (F(1,9)=2.5, p=0.150) or QP start position (F(1,9)=4.8, p=0.056), although there was a significant main effect on QP end position (F(1,9)=6.0, p=0.036). There were no significant interaction effects between any of the stimulus conditions on SP amplitude, QP amplitude, QP start position or QP end position.

Three-way RM-ANOVA on the standard deviation of these parameters found a main effect of spatial frequency on the standard deviation of SP amplitude (F(1.28,11.5)=1.2, p=0.323), QP amplitude (F(2,18)=0.5, p=0.624), QP start position (F(2,18)=0.9, p=0.427) and QP end position (F(2,18)=2.9, p=0.084). There was also no main effect of stimulus pattern on the standard deviation of SP amplitude (F(1,9)=1.2, p=0.301) or QP amplitude (F(1,9)=2.0, p=0.192), but there was an effect on QP start position (F(1,9)=10.2, p=0.011) and QP end position (F(1,9), p=0.026.)

In summary, these results verified the findings in experiment 2 that there was a main effect of spatial frequency on mean values of SP amplitude and QP amplitude, as well as on the standard deviation of QP and SP amplitude, and the standard deviation of the start and end position of QPs and SPs. Using a random grid stimulus pattern resulted in an increase in the mean SP and QP amplitude, and the standard deviation of the start and end positions of QPs and SPs, when compared with using a square wave stimulus pattern. Using a random grid stim-

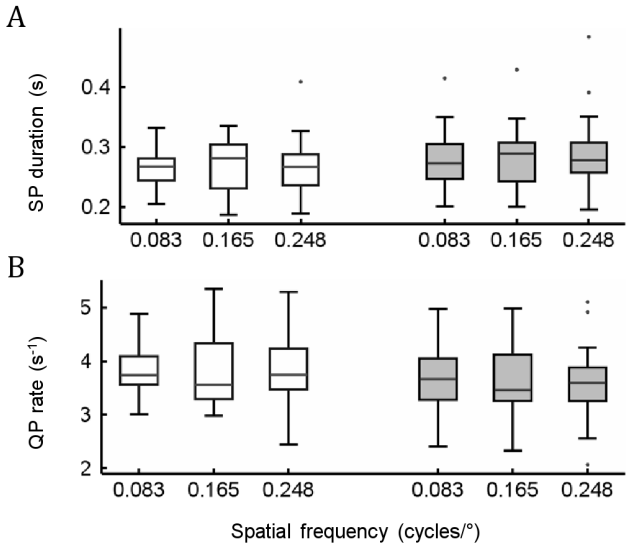

Figure 6.12: Boxplot illustrating the effect of spatial frequency and stimulus pattern on (A) SP duration and (B) QP rate. Each box represents the parameter grouped across all participants for each stimulus condition. Stimulus conditions tested were: stimulus pattern, square wave or random grid; spatial frequency, 0.083cyc/°, 0.165cyc/°or 0.248cyc/°. Grey boxes represent data from trials with a square wave stimulus pattern, white boxes represent data from trials with a random grid stimulus pattern. Box represents interquartile range, whiskers represent maximum and minimum data values (excluding outliers), and dots represent outliers found more than 1.5 times the interquartile range from the ends of the box.

ulus pattern also resulted in more negative start and end positions of QPs (and SPs) when compared with using a square wave stimulus pattern.

6.4.3 SP duration and QP rate.

Figure 6.12 contains two boxplots demonstrating how the median SP duration and QP rate are affected by different stimulus patterns and by changing the relevant spatial frequency of those stimulus patterns. A three-way RM-ANOVA was performed on the aligned rank-transformed median values of each parameter as the dependent variables, and the stimulus pattern, spatial frequency, and direction of the stimulus as the within-subjects factors.

There was no significant main effect of spatial frequency on SP duration ($F(2,18)=0.4$, $p=0.691$) or QP rate ($F(2,18)=0.5$, $p=0.616$), as expected from the results of experiment 2.

There was a significant main effect of stimulus pattern on SP duration ($F(1,9)=5.8$, $p=0.039$) and QP rate ($F(1,9)=5.6$, $p=0.042$).

There was no significant main effect of direction on SP duration ($F(1,9)=0.0$, $p=0.874$) or QP rate ($F(1,9)=0.0$, $p=0.935$). There were also no significant interaction effects between any of the stimulus conditions on SP duration or QP rate.

6.5 Investigating the principal components found in experiment 3.

As with experiment 2, the same PCA procedure that had been applied to the data from experiment 1 was applied to the data from experiment 3. The analysis was performed on the correlation matrices of parameters from across four adjacent cycles. PCA performed on the correlation matrices of the 120 trials did reveal 13 significant eigenvalues, corresponding to the maximum number of DOF (fig. 6.13a). The adjusted heuristics that had been created in order to sort the components from experiment 2 into individual categories were capable of sorting all of these components into individual categories. These results illustrated that there were 13 qualitatively different components within the data from each trial in experiment 3, and they could be sorted into the same categories as found in experiment 1 and 2 (i.e. they had similar loading patterns). As in the results from experiment 1 and 2, each of these sorted components occurred in a different eigenvalue order (fig. 6.13b).

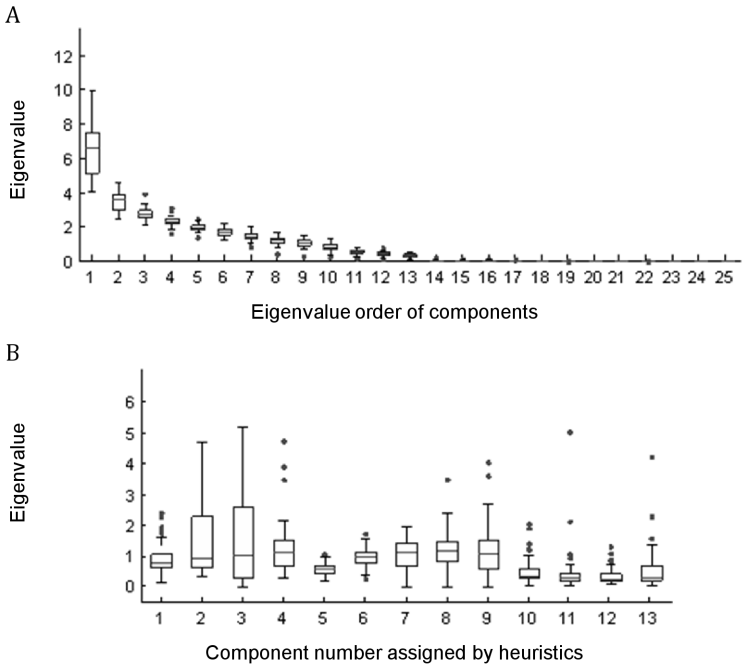

Figure 6.13: (A) Boxplot illustrating the eigenvalues of principal components found in four cycles of OKN parameters in series before being sorted by heuristics. Number on principal component axis represents the eigenvalue order of components from largest to smallest. (B) Boxplot illustrating the eigenvalues of principal components extracted from four cycles of OKN parameters in series after being sorted by heuristics. Sorted component numbers 1-5, "Q-component"; 6-9, "S-component"; 10-13, "V-component". Boxes represents interquartile range, whiskers represent maximum and minimum data values (excluding outliers), and dots represent outliers found more than 1.5 times the interquartile range from the ends of the box. OKN data taken from experiment 3 only.

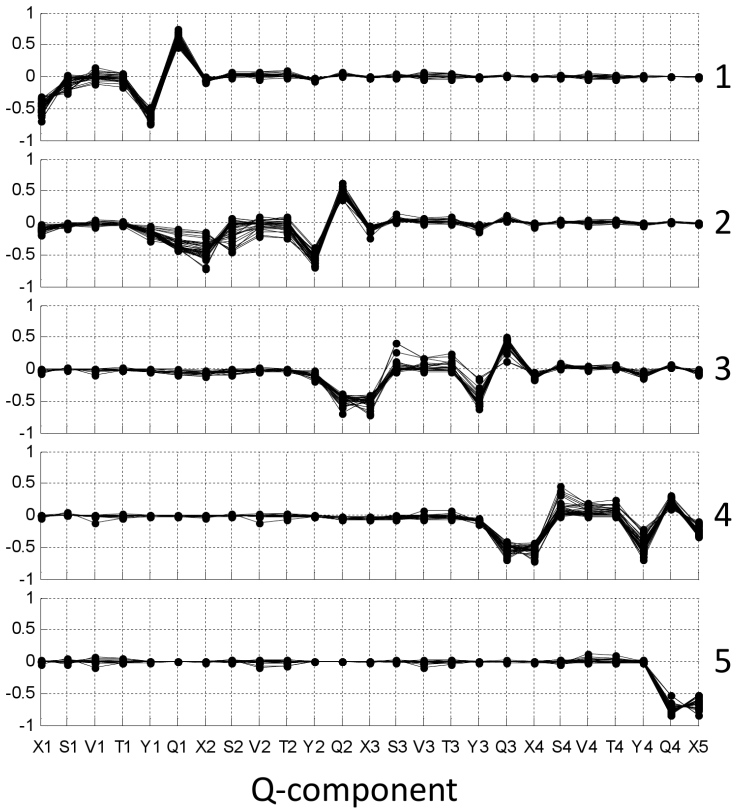

Figure 6.14: Results of PCA performed on four cycles of OKN in series, including SP velocity and duration. Component loadings of 5 principal components extracted from OKN data from experiment 3 (120 trials). A total of 13 principal components were extracted during analysis. These 5 components verify the existence of the Q-component category in the data recorded from experiment 3. Numbers are used only to label components and do not represent the position of the corresponding eigenvalue on the scree plot as these varied between trials.

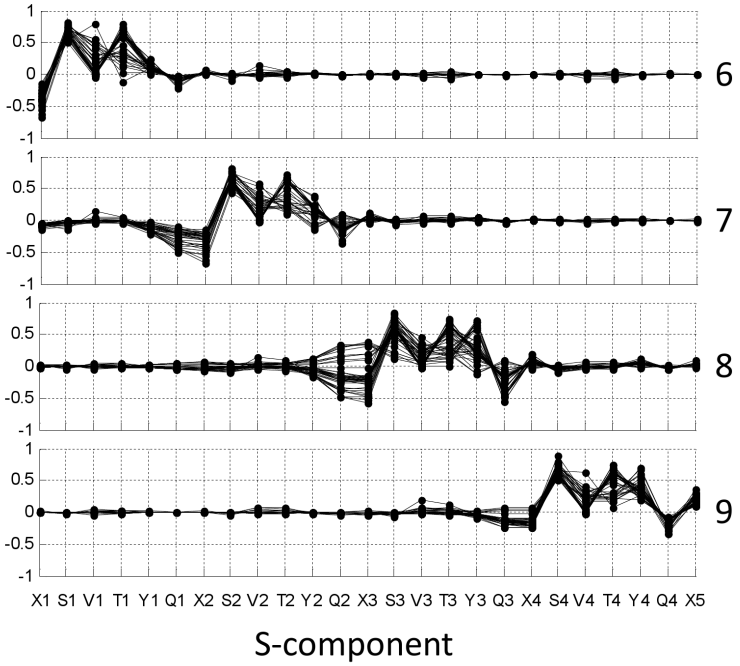

Figure 6.15: Results of PCA performed on four cycles of OKN in series, including SP velocity and duration. Component loadings of 4 principal components extracted from OKN data from experiment 3 (120 trials). A total of 13 principal components were extracted during analysis. These 4 components verify the existence of the S-component category in the data recorded from experiment 3. Numbers are used only to label components and do not represent the position of the corresponding eigenvalue on the scree plot as these varied between trials.

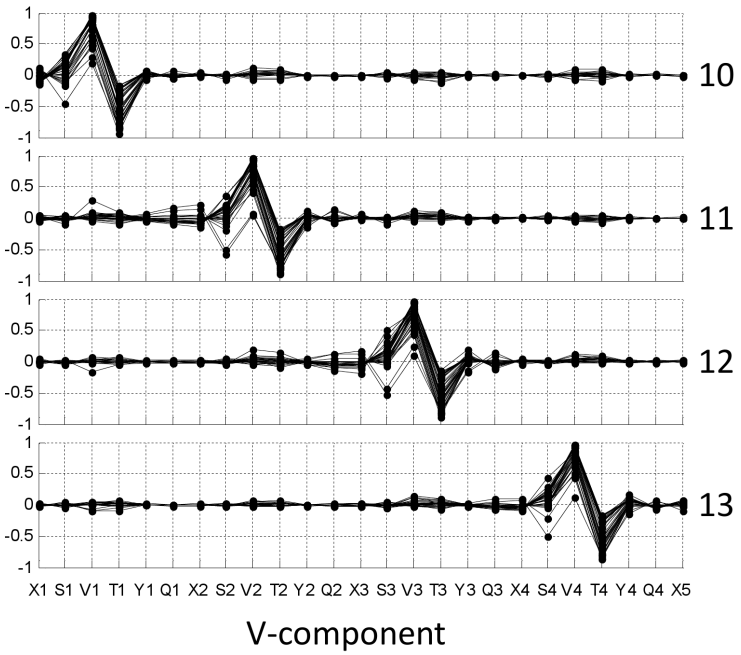

Figure 6.16: Results of PCA performed on four cycles of OKN in series, including SP velocity and duration. Component loadings of 4 principal components extracted from OKN data from experiments 3 (120 trials). A total of 13 principal components were extracted during analysis. These 4 components verify the existence of the V-component category in the data recorded from experiment 3. Numbers are used only to label components and do not represent the position of the corresponding eigenvalue on the scree plot as these varied between trials.

6.6 Results of applying the model to OKN parameters recorded from experiment 3.

Multiple linear regression was performed on the OKN parameters recorded from experiment 3 in order to determine the coefficients a, b, c and d in exactly the same way as with experiment 2, and they were found to be -0.318 (SD=0.19), 0.120 (SD=0.14), -0.352 (SD=0.25) and -0.137 (SD=0.10) respectively. Bonferroni corrected t-tests indicated a possible significant difference in c between experiments 1 and 3 (p=0.010).

We applied the model to the recorded parameters of experiment 3 in order to determine if the differences in parameters would have a significant impact on the predictive power of the model. Robust linear regression was performed on the parameters recorded from the real OKN data in order to estimate the values of e, $\widehat{v}$, σ_s, σ_q and σ_v from each trial. The mean values of the estimates for a, b, c and d that had been found in the data from experiment 1 (-0.250, 0.158, -0.478 and -0.166), and the mean values of the respective OKN parameters recorded from experiment 3, were substituted into eq. 5.4 and 5.5 in order to find the values of $\widehat{s}$ and $\widehat{q}$ from each trial.

Simulated data sets were created iteratively in a Matlab function with eqs. 5.1 - 5.9 and the number of simulated samples was set to match the number of observed samples for each parameter. The mean (and median) values of simulated parameters were compared with the mean (and median) values of observed data (fig. 6.17). Also, the standard deviation (and interquartile range) of simulated parameters were compared with the standard deviation (and interquartile range) of observed data (fig. 6.18). In both cases the simulated values accurately reflected the observed values. R^2 for mean (and median) values of OKN parameters range

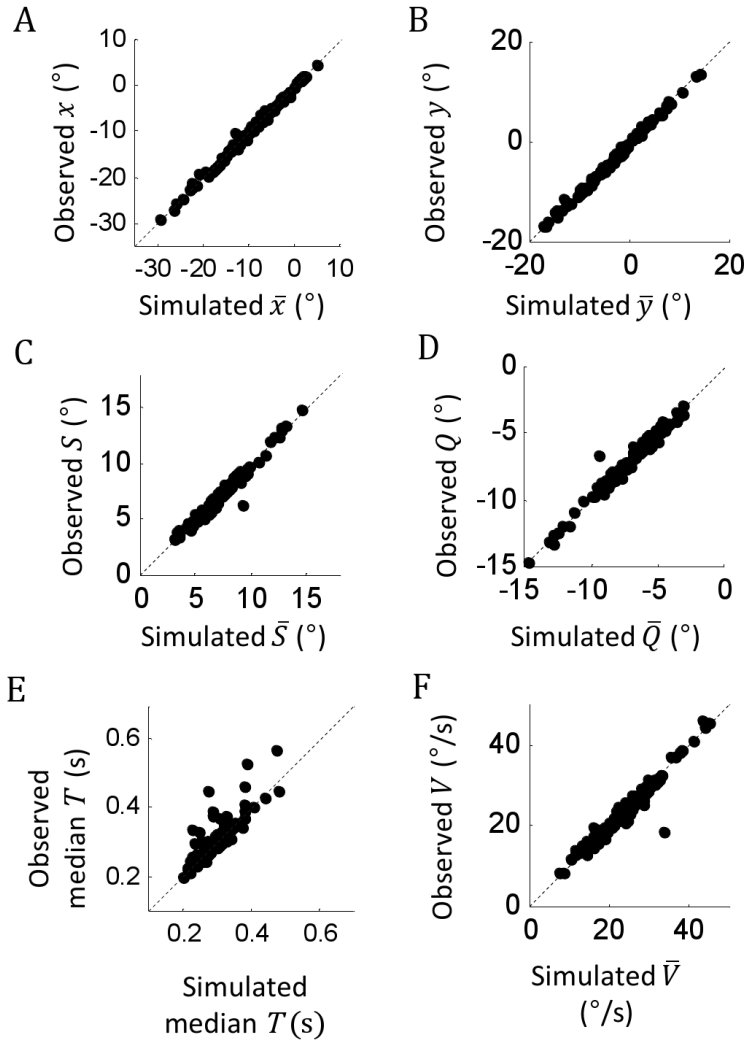

Figure 6.17: Scatter plots of the values of central tendency from the observed OKN parameters for each trial in experiment 3, against the values of central tendency from the simulated OKN parameters. These simulated data were created using Monte Carlo simulations of equations 5.4 - 5.6, where the parameters a, b, c and d were estimated from results in experiment 1. The number of simulated samples of each parameter was set to match the number of observed samples for each parameter. Dotted lines represent identity.

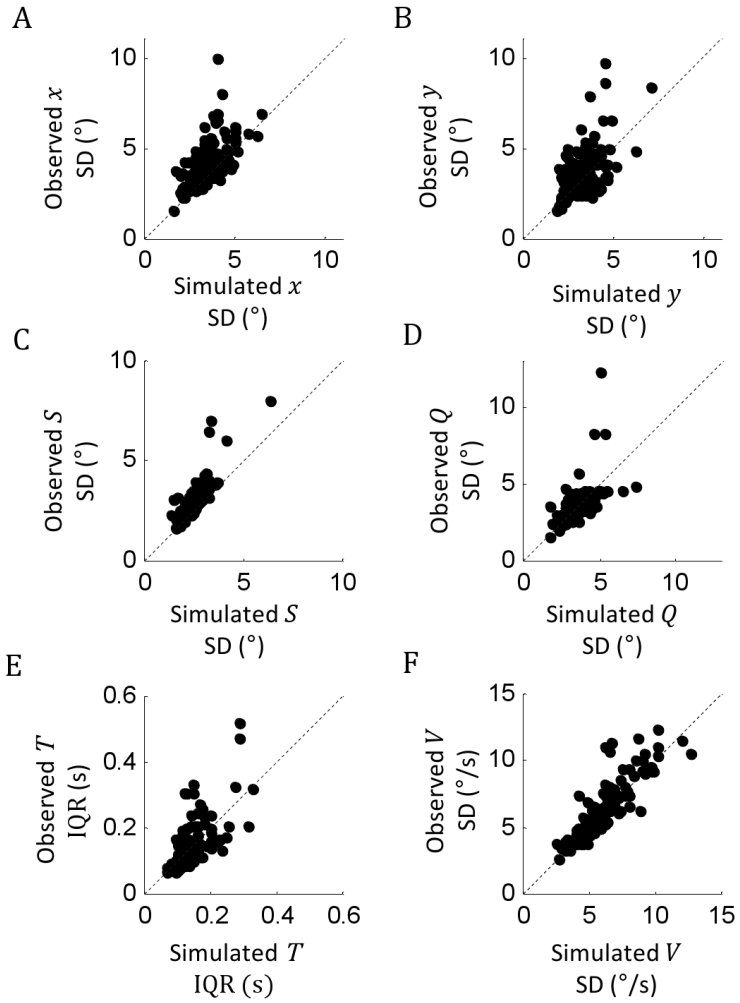

Figure 6.18: Scatter plots of the values of the measure of dispersion from the observed OKN parameters for each trial in experiment 3, against the values of the measure of dispersion from the simulated OKN parameters. These simulated data were created using Monte Carlo simulations of equations 5.4 - 5.6, where the parameters a, b, c and d were estimated from results in experiment 1. The number of simulated samples of each parameter was set to match the number of observed samples for each parameter. SD, standard deviation; IQR, Inter-quartile range. Dotted lines represent identity.

Experiment	a	b	c	d
Experiment 1	-0.250 (0.15)	0.158 (0.16)	-0.478 (0.27)	-0.166 (0.10)
Experiment 2	-0.251 (0.20)	0.124 (0.20)	-0.397 (0.21)	-0.197 (0.16)
Experiment 3	-0.318 (0.19)	0.120 (0.14)	-0.352 (0.25)	-0.137 (0.10)
Combined	-0.282 (0.19)	0.127 (0.17)	-0.389 (0.24)	-0.164 (0.13)

Table 6.1: Table of constants found for the Markov model from each experiment. The mean values of a, b, c and d are given from across all trials for each experiment, and finally combined across all trials. The standard deviation of these parameters are given in brackets.

from 0.74 to 1.00, and the R^2 for standard deviation (and interquartile range) values range from 0.32 to 0.77. This illustrated that the model could capture the distribution of OKN parameters well with only four free variables in the model (or 7, including the estimated values of standard deviation in the underlying components).

When taken across all 256 trials, the mean values of a, b, c and d were -0.282 (SD=0.19), 0.127 (SD=0.17), -0.389 (SD=0.24) and -0.164 (SD=0.13) respectively. For the purpose of future simulations using the Markov model, the values of the constants a, b, c and d were set to -0.282, 0.127, -0.389 and -0.164. These values had been determined from 256 trials as opposed to just the first 40 trials in experiment 1, so were likely to be more representative of the actual values of these parameters. A table containing the estimated values of a, b, c and d from each experiment, and combined across all experiments, is illustrated in table 6.1.

6.7 Verifying the distribution of the noise processes.

In the previous chapter we assumed the noise process in each component was Gaussian. We can test this hypothesis by finding the residuals of the observed

OKN data, i.e. the deviation of the observed OKN parameters from the estimated value of the OKN parameters, and their distribution in the empirical data. The residuals of the S-, Q-, and V-component are given by a simple rearrangement of the eqs. 5.4 - 5.5 respectively:

$$\epsilon_s(i) = S_i - ax_i - bV_i - \widehat{s}, \tag{6.1}$$

$$\epsilon_q(i) = Q_i - cy_i - dV_i - \widehat{q}, \tag{6.2}$$

and

$$\epsilon_v(i) = V_{i+1} - eV_i - \widehat{v}. \tag{6.3}$$

Figures 6.19 - 6.21 illustrate these residuals taken from one participant at different stimulus speeds. These histograms are representative of those created from all trials, and by eye might appear Gaussian. However, Lilliefors tests conducted on all 256 trials (corrected using the Holm-Bonferroni method) show that 110 histograms of the S-component residuals, 63 histograms of the Q-component residuals, and 100 histograms of the V-component residuals are significantly different to the Gaussian distribution. This appeared to be due to the tendency for residuals to be leptokurtic, with a more acute peak and fatter tails than normal. The kurtosis of a Gaussian distribution is 3.0, whereas the average kurtosis of the S-component

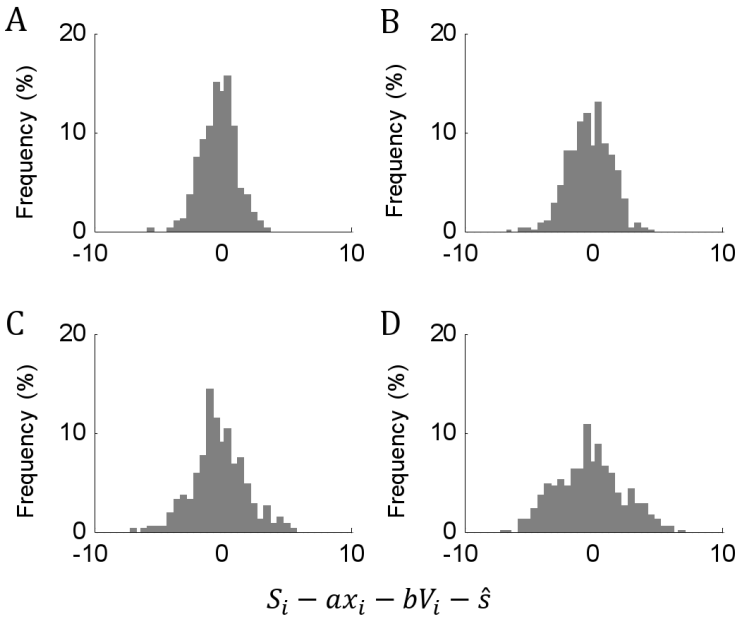

Figure 6.19: Frequency histograms of the noise process from the S-component as calculated by finding the residuals of $S_i - ax_i - bV_i - \widehat{s}$. Data from participant 10, stimulus speed (A) 10°/s, (B) 20°/s, (C) 30°/s and (D) 40°/s.

residuals is 5.2, of the Q-component residuals is 5.2, and of the V-component residuals is 8.8. The Q-component and V-component residuals are approximately symmetric (mean skewness = 0.2 and 0.1 respectively). However, the S-component residuals sometimes demonstrated a moderate amount of positive skewness (mean skewness = 0.7). Another interesting pattern found in the V-component residuals was that the standard deviation of the residuals appeared to increase with stimulus speed. This indicated that there may be signal dependent noise in the V-component.

6.8 Signal-dependent noise.

In order to determine if any of the three stochastic processes demonstrated noise proportional to the magnitude of the process, the values of σ_s and σ_q were plotted against the values $\overline{S}$ and $\overline{Q}$ respectively. The relationship between σ_s and $\overline{S}$, and σ_q

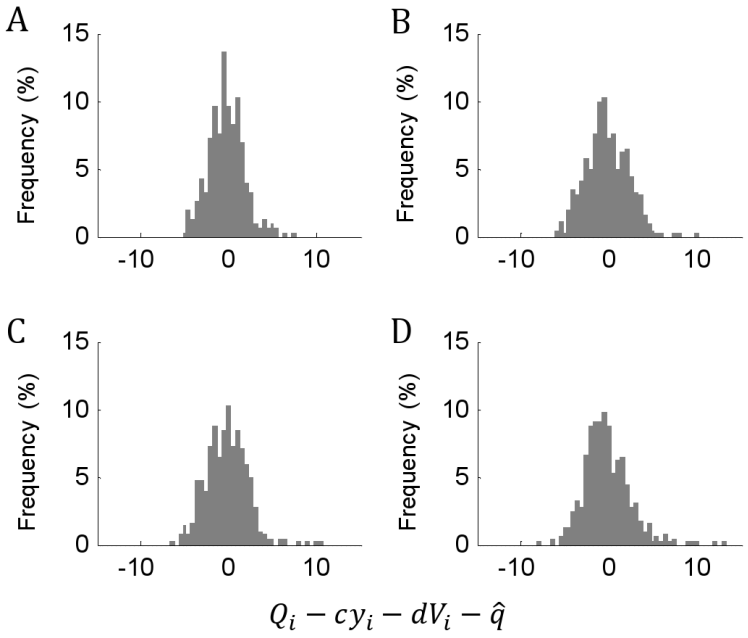

Figure 6.20: Frequency histograms of the noise process from the Q-component as calculated by finding the residuals of $Q_i - cy_i - dV_i - \widehat{q}$. Data from participant 10, stimulus speed (A) 10°/s, (B) 20°/s, (C) 30°/s and (D) 40°/s.

and $\overline{Q}$ was determined using the bivariate total least squares method (orthogonal regression). Both the S- and Q-components demonstrated proportional noise, as illustrated by a linear increase in the standard deviation of the error term with the magnitude of $\overline{S}$ and $\overline{Q}$ respectively, and a possible constant source of noise illustrated by the significant positive intercept of the line of regression on the y-axis (fig. 6.22a and 6.22b). Ordinary least squares regression indicated that the 95% confidence interval of the intercept did not contain zero, also indicating a constant source of noise.

The values of σ_v were also plotted against $\overline{V}$, stimulus speed, and retinal slip. There was a relatively weak relationship between σ_v and $\overline{V}$ ($r^2 = 0.21$), and a very weak relationship between σ_v and the absolute mean retinal slip ($r^2 = 0.05$). However, there was a moderate relationship between σ_v and stimulus speed ($r^2 =$

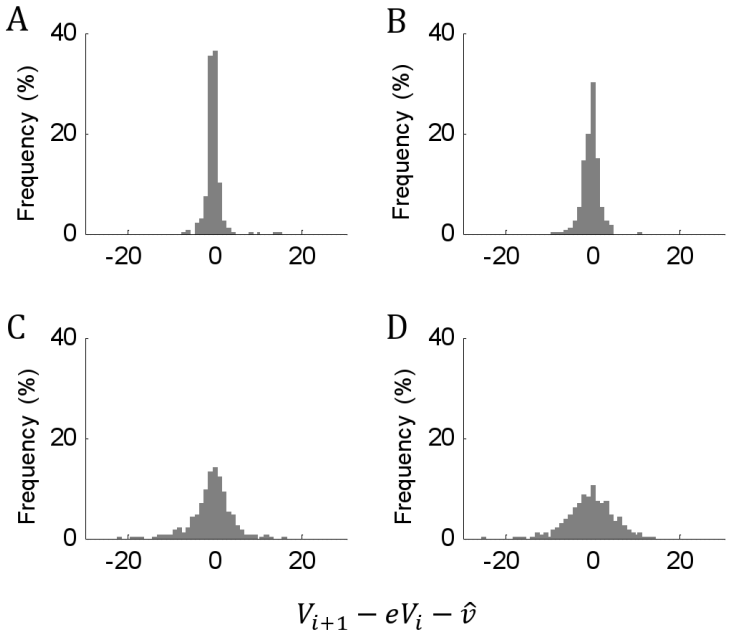

Figure 6.21: Frequency histograms of the noise process from the V-component as calculated by finding the residuals of $V_{i+1} - eV_i - \widehat{v}$. Data from participant 10, stimulus speed (A) 10°/s, (B) 20°/s, (C) 30°/s and (D) 40°/s.

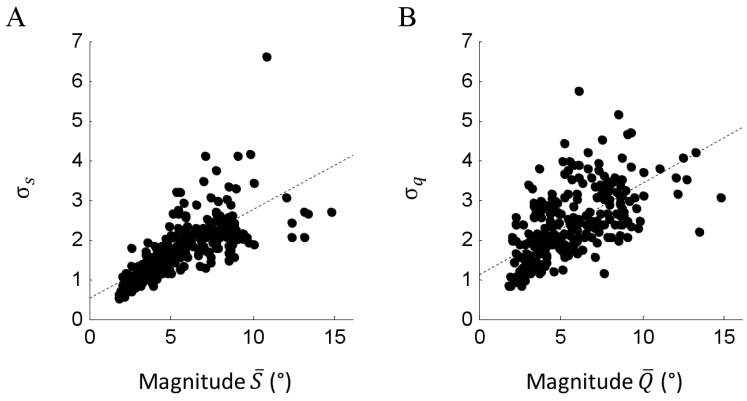

Figure 6.22: Scatter plots illustrating proportional noise in all three stochastic processes. Dashed lines represent lines of orthogonal regression. In (A) S-component error increases with SP amplitude ($\sigma_s = 0.22\overline{S} + 0.52$) and in (B) Q-component error increases with QP amplitude ($\sigma_q = 0.23\overline{Q} + 1.1$). Error values ($\sigma_s$ and σ_q) are the standard deviations of the error term estimated from linear regression performed on the data.

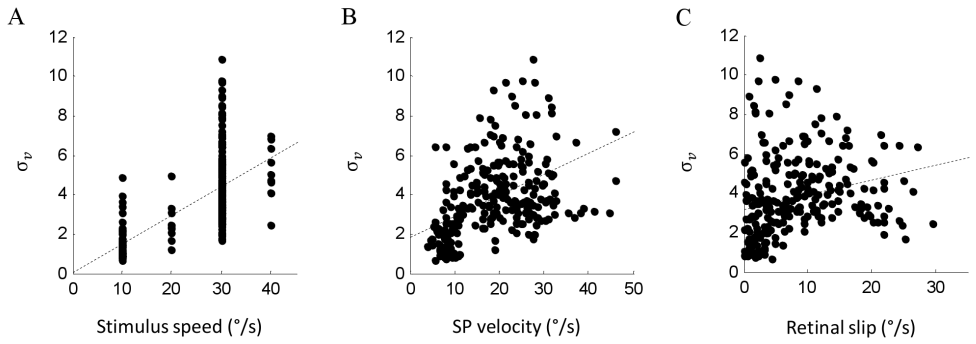

Figure 6.23: Scatter plots illustrating proportional noise in V components. Dashed lines represent lines of regression. In (A) V-component error increases with stimulus speed ($\sigma_v = 0.15$(stimulus speed) + 0.05), in (B) V-component error increases with SP velocity ($\sigma_v = 0.11\overline{V} + 1.8$), and in (C) V-component error increases with the absolute mean retinal slip ($\sigma_v = 0.05$(retinal slip) + 3.5). The error value (σ_v) is the standard deviation of the error term estimated from linear regression performed on the data.

0.37). The relationships between the σ_v and $\overline{V}$, σ_v and stimulus speed, and σ_v and retinal slip, using orthogonal regression (fig. 6.23). The relationship between σ_v and stimulus speed has no constant source of noise (fig. 6.23a).

The proportional relationships of the signal dependent noise in each component were:

$$\sigma_s = 0.22\overline{S} + 0.52, \tag{6.4}$$

$$\sigma_q = 0.23\overline{Q} + 1.1, \tag{6.5}$$

and

$$\sigma_v = 0.15V_S + 0.05 \tag{6.6}$$

where V_S is the stimulus speed. This means that during trials where the mean QP amplitude is 5.6° (this is the mean value across all trials) the constant component of noise contributes approximately 46% towards the total amount of noise in the Q-component. Whereas when the mean SP amplitude is 5.6° the constant component contributes approximately 30% towards the total amount of noise in the S-component. Of course, as the QP or SP amplitude increases the contribution from the level of constant noise will decrease.

The possibility of using these proportional relationships to define the standard deviation of the error terms in the Markov model, rather than estimating the standard deviation of the error term from every trial was considered. The mean values of the estimates for a, b, c and d that had been found in the data from experiments 1-3 (-0.282, 0.127, -0.389 and -0.164), and the respective mean values of the OKN parameters recorded from experiments 1-3, were substituted into eq. 5.4 and 5.5 to find the values of $\widehat{s}$ and $\widehat{q}$ from each trial. The values of σ_s, σ_q and σ_v from each trial were estimated using eqs. 6.4 - 6.6. The free variables e and $\widehat{v}$ were estimated with robust linear regression as before. Simulated data sets were created iteratively in a Matlab function with eqs. 5.1 - 5.9 and the number of simulated samples was set to match the number of observed samples for each parameter. The mean (and median) values of simulated parameters were compared with the mean (and median) values of observed data (fig. 6.24). Also, the standard deviation (and interquartile range) of simulated parameters were compared with the standard deviation (and interquartile range) of observed data (fig. 6.25).

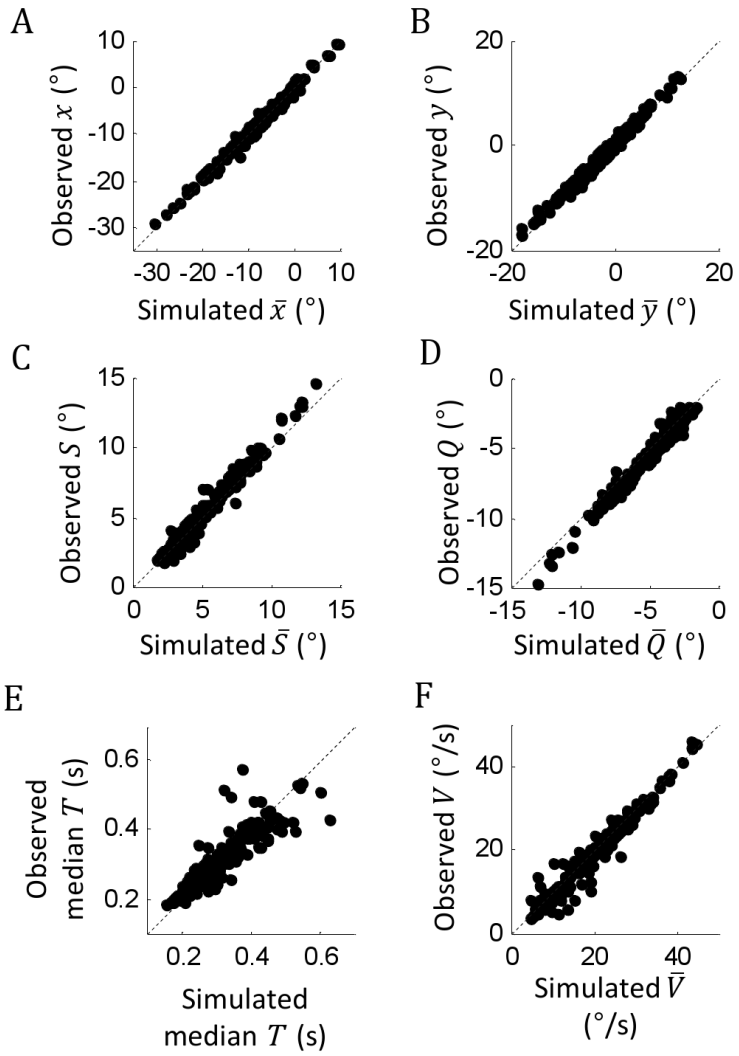

Figure 6.24: Scatter plots of the values of central tendency from the observed OKN parameters for each trial in experiments 2 and 3, against the values of central tendency from the simulated OKN parameters. These simulated data were created using Monte Carlo simulations of equations 5.4 - 5.6. The parameters a, b, c and d were considered constants and not free parameters, and the standard deviation of the noise processes $\epsilon_s(i)$, $\epsilon_q(i)$ and $\epsilon_v(i)$ were calculated from the constants of proportionality found in the signal dependent noise results. The number of simulated samples of each parameter was set to match the number of observed samples for each parameter. Dotted lines represent identity.

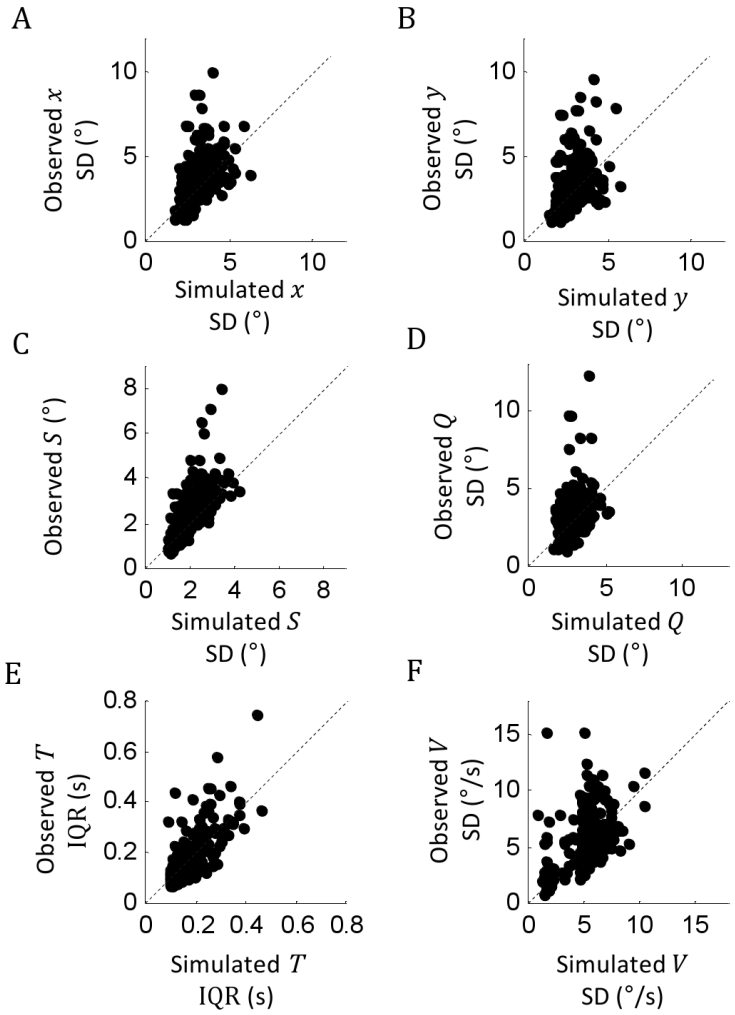

Figure 6.25: Scatter plots of the values of the measure of dispersion from the observed OKN parameters for each trial in experiments 2 and 3, against the values of the measure of dispersion from the simulated OKN parameters. These simulated data were created using Monte Carlo simulations of equations 5.4 - 5.6. The parameters a, b, c and d were considered constants and not free parameters, and the standard deviation of the noise processes $\epsilon_s(i)$, $\epsilon_q(i)$ and $\epsilon_v(i)$ were calculated from the constants of proportionality found in the signal dependent noise results. The number of simulated samples of each parameter was set to match the number of observed samples for each parameter. SD, standard deviation; IQR, Inter-quartile range. Dotted lines represent identity.

The model could still accurately predict the observed mean (and median) values when the standard deviation of the error terms in each component was estimated using eqs. 6.4 - 6.6. However, the R^2 for the values of standard deviation (and interquartile range) of OKN parameters dropped to a range between 0.11 and 0.60 when the standard deviation of the error terms was estimated from eqs. 6.4 - 6.6, illustrating that the power of the model in predicting the distribution of parameters was somewhat reduced from when the standard deviation was estimated from the residuals of linear regression.

6.9 How does stimulus speed affect SP velocity in the Markov model?

In section 6.8 we illustrated how the stimulus speed appears to affect the standard deviation of the error term in the V-component. It is important to consider how the stimulus speed affects the model, and in particular how stimulus speed affects SP velocity. From figure 4.3a it is obvious that SP velocity depends on stimulus speed and so we might assume that the stimulus dependence is captured somehow in eq. 5.6 either by e or by $\widehat{v}$. In order to investigate how stimulus speed affects the parameters of the model some assumptions were made about the system that generates OKN. A closed loop system driven by a retinal slip signal in the forward loop, and efference copy of the SP velocity in a feedback loop is illustrated in figure 6.26a. The values g_1, g_2, g_3 and g_4 are the gain of retinal slip, efference copy, combined retinal slip and efference copy, and SP velocity respectively. Assuming that g_1, g_2, g_3 and g_4 are linear and time-invariant, this system gives the relationship:

$$V_{i+1} = g_4(g_3(g_1(V_S - V_i) + g_2(V_{EC}))) \tag{6.7}$$

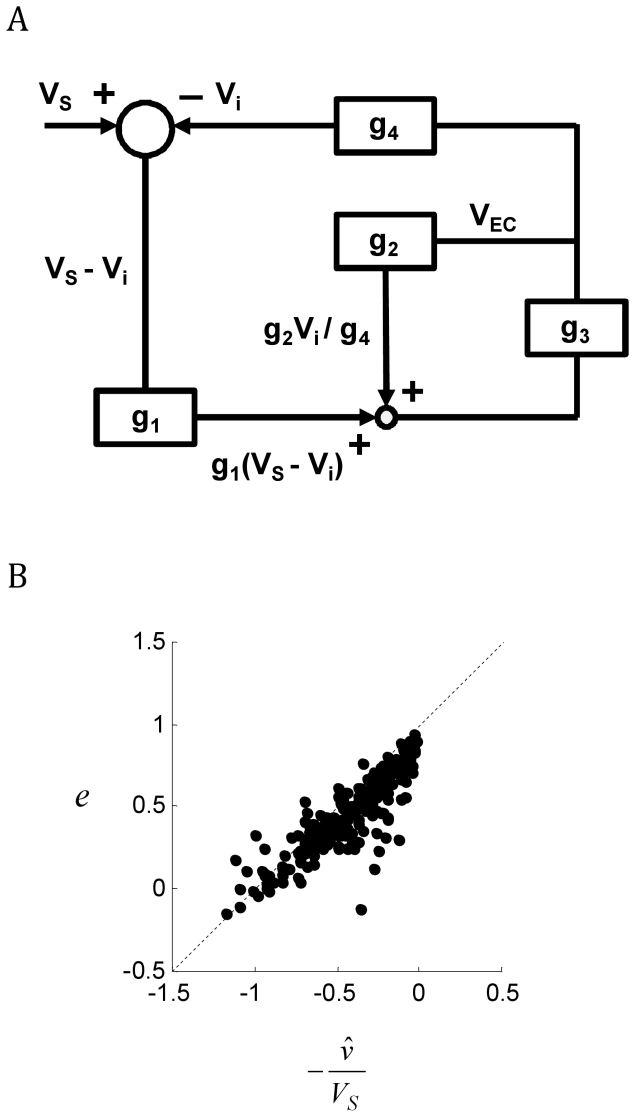

Figure 6.26: (A) Block diagram illustrating a simple system that could generate OKN, where a retinal slip signal and an efference copy signal are summed to generate an eye velocity signal. The values g_1, g_2 and g_3 are the respective gain elements of retinal slip in the forward loop, efference copy in the feedback loop, and SP velocity in the forward loop. The value g_4 represents a linear plant. V_S, stimulus velocity; $V_S - V_i$, retinal slip; V_i, SP velocity; V_{EC}, efference copy. (B) Scatter plot relating the values of $\widehat{v}$ and e. Data points represent values estimated from each trial by linear regression. Note that under the assumptions of the model $\widehat{v} = g_1 g_3 g_4 V_S$ (the product of the gain elements in the forward loop and stimulus velocity) and $e = g_2 g_3 - g_1 g_3 g_4$ (the product of the gain elements in the feedback loop, minus the product of the gain elements in the forward loop). Dotted line, $e = 1 - \frac{\widehat{v}}{V_S}$.

where $V_S - V_i$ and V_{EC} are retinal slip and efference copy respectively. We can relate SP velocity and efference copy with:

$$V_{EC} = \frac{V_{i+1}}{g_4} \tag{6.8}$$

Substituting eq. 6.8 into eq. 6.7 gives:

$$V_{i+1} = g_1 g_3 g_4 (V_S - V_i) + g_2 g_3 V_{i+1} \tag{6.9}$$

and collecting terms in eq. 6.9 gives:

$$V_{i+1} = g_3(g_2 - g_1 g_4) V_i + g_1 g_3 g_4 V_S \tag{6.10}$$

Substituting eq. 6.10 into eq. 5.6 gives:

$$e = g_3(g_2 - g_1 g_4) \tag{6.11}$$

and

$$\widehat{v} = g_1 g_3 g_4 V_S \tag{6.12}$$

It appears that the term $\widehat{v}$ captures the dependence of the Markov model on the OKN stimulus speed, when considered as part of the system described in figure 6.26a. However, if this is the case then the terms e and $\widehat{v}$ must be related as they share the term $g_1 g_3 g_4$. Rearranging eq. 6.12 and substituting into eq. 6.11 illustrates that:

$$e = g_2 g_3 - \frac{\widehat{v}}{V_S} \tag{6.13}$$

This relationship can be observed in the scatter plot of the values of e against $-\widehat{v}/V_S$ estimated from each trial (fig. 6.26b). Data points tend to fall below the line defined by $e = 1 - \widehat{v}/V_S$, and bivariate linear regression performed on these data give a slope of 0.882, and a y-intercept of 0.842. There is some variability in this realtionship as not all points fall exactly on the line of regression so we might suppose that the error term in the V component is due to one (or more) of the gain values (g_1, g_2, g_3 or g_4) being variable. An interesting possibility is given if g_1, g_3 or g_4 varies. As the product of g_1, g_3 and g_4 is multiplied by the stimulus speed to arrive at $\widehat{v}$ (eq. 6.12) this would explain our finding that the error term in the V-component is proportional to stimulus speed.

6.10 Steady state of OKN parameters.

Any ergodic stable 1st order Markov process will tend towards a statistical steady state, where the probability distribution approaches some fixed function, independent of starting conditions (Kijima 1997). Therefore, the mean and variance will tend to become constant as $i \to \infty$. Once the system reaches the steady state the mean values for each of these OKN parameters can be calculated in relation to

$\overline{V}$, $\widehat{s}$ and $\widehat{q}$, from eq. 5.3, 5.6 and 5.13 - 5.16 using the values estimated for the constant parameters (a, b, c, and d), and are given by:

$$\begin{aligned}\overline{x} &= \frac{((1+c)b+d)}{1-(1+a)(1+c)}\overline{V} + \frac{(1+c)\widehat{s}+\widehat{q}}{1-(1+a)(1+c)} \\ &\approx -0.154\overline{V} + 1.09\widehat{s} + 1.78\widehat{q}\end{aligned} \tag{6.14}$$

$$\begin{aligned}\overline{y} &= \frac{((1+a)d+b)}{1-(1+a)(1+c)}\overline{V} + \frac{\widehat{s}+(1+a)\widehat{q}}{1-(1+a)(1+c)} \\ &\approx 0.017\overline{V} + 1.78\widehat{s} + 1.28\widehat{q}\end{aligned} \tag{6.15}$$

$$\begin{aligned}\overline{S} &= \frac{ad-bc}{1-(1+a)(1+c)}\overline{V} + \frac{a\widehat{q}-c\widehat{s}}{1-(1+a)(1+c)} \\ &\approx 0.17\overline{V} + 0.69\widehat{s} - 0.50\widehat{q}\end{aligned} \tag{6.16}$$

$$\begin{aligned}\overline{Q} &= \frac{bc-ad}{1-(1+a)(1+c)}\overline{V} + \frac{c\widehat{s}-a\widehat{q}}{1-(1+a)(1+c)} \\ &\approx -0.17\overline{V} - 0.69\widehat{s} + 0.50\widehat{q}\end{aligned} \tag{6.17}$$

$$\begin{aligned}\overline{T} &= \frac{ad-bc}{1-(1+a)(1+c)} + \frac{a\widehat{q}-c\widehat{s}}{1-(1+a)(1+c)}\frac{1}{\overline{V}} \\ &\approx 0.171 + \frac{0.38\widehat{s}-0.28\widehat{q}}{0.56\overline{V}}\end{aligned} \tag{6.18}$$

and

$$\overline{V} = \frac{\widehat{v}}{1-e} \tag{6.19}$$

The model clearly predicts that mean SP and QP amplitudes depend linearly with mean SP velocity, whereas mean SP duration is inversely related, as empirically observed. As the mean SP and QP amplitude both depend on the mean SP velocity and the values of $\widehat{s}$ and $\widehat{q}$, the standard deviation of the error terms in the S-component and Q-component will also depend on the mean SP velocity and the values of $\widehat{s}$ and $\widehat{q}$, due to the signal dependent noise processes.

6.11 Transient state of OKN parameters.

In the inital period before a Markov process reaches steady state it is said to be in a transient state that is dependent on the initial conditions of the system. The behaviour of transients may give additional interesting information about a system, and here we investigate the behaviour of the transient state in the simple case where the initial eye velocity and position are both 0°/s. In the Markov model SP velocity is only dependent on SP velocity in the previous cycle and will converge to a stable steady state if the value of the parameter e is less than ± 1,

otherwise the process will diverge and continue to increase in magnitude. The value of e ranged from -0.15 to 0.94, consequently SP velocity should reach a stable steady state. Although we have not measured transient responses directly, they can be easily computed.

The update dynamics of SP velocity were simulated with eq. 5.6 using the values of the parameters $\widehat{v}$ and e obtained through linear regression from all 256 data sets, then the number of cycles it took to reach 63% of the steady state SP velocity (eq. 6.19) from an initial eye velocity of 0°/s was calculated. When the point at which SP velocity reached 63% of the steady state value was between one cycle and the next, linear interpolation was used to determine the fractional value at which SP velocity reached 63% of the steady state value.

From an initial eye velocity of 0°/s, the transient state of SP velocity lasted for just a few cycles. To reach 63% of the steady state SP velocity took an average 2.1 cycles (standard deviation=1.8) across all 256 simulated data sets. This is in good agreement with the study by Abadi et al. (2005) who showed that steady state was usually reached in two OKN cycles. We compared the transient response of SP velocity at different stimulus speeds. When the simulated transient responses from trials with a stimulus speed of 10°/s were averaged, it took 1.3 cycles (standard deviation=0.9) to reach 63% of SP velocity. The average simulated transient response from trials with a stimulus speed of 20°/s was 2.7 (standard deviation=2.7), at 30°/s it was 2.3 (standard deviation=1.9), and at 40°/s it was 2.4 (standard deviation=1.9). Figure 6.27a illustrates the transient response of SP velocity, when all the simulations created using parameters from trials with the same stimulus speed were averaged. There was a tendency for SP velocity to reach steady state more quickly when the stimulus speed was 10°/s than at other stimulus speeds, but the transient response times were quite variable.

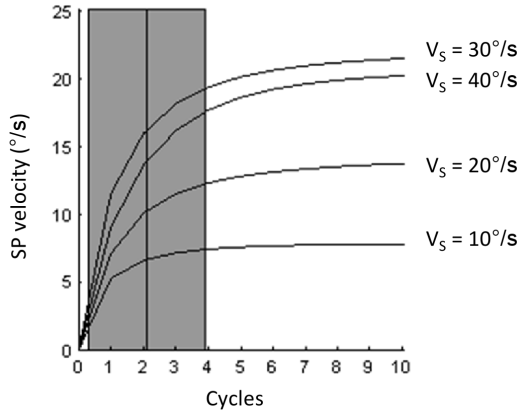

Figure 6.27: Simulated transient response of SP velocity starting from 0°/s. Simulations that were created using the parameters from trials with the same stimulus speed were averaged. Each line represents one averaged transient response of SP velocity with a matching stimulus speed. V_S notes stimulus speed. Solid line at 2.1 cycles represents mean of the cycles required to reach 63% of SP velocity, shaded area represents ± 1 standard deviation either side of the mean value.

6.12 QP targetting.

As the participants were asked to stare towards the centre of the screen we might assume that the desired target for returning QPs was straight ahead. However, the end position of QPs tends to overshoot the centre and drive eye position in to the negative field of view. The data illustrates that the amplitude of QPs depend on their start position and our model predicts that there is a start position where, if a QP is triggered, the system would generate a QP with mean amplitude of zero. Assuming that there is some desired location towards which QPs are targetted, we can make the hypothesis that the position where a zero amplitude QP would be generated is the desired target location of QPs. In order to estimate what this position is, we solve eq. 5.5 for y_i when $Q_i = 0$ to give:

$$y_T = \frac{-(dV_i + \widehat{q})}{c} \tag{6.20}$$

where y_T is the desired target location. Using the values of c, d and $\widehat{q}$ from our data, and estimating V_i with $\overline{V}$ the average y_T for each trial was found. Mean values of y_T were -9.8°, -13.1°, -18.9° and -15.6° for trials with a stimulus speed of 10°/s, 20°/s, 30°/s and 40°/s respectively (mean SP velocity was 8.3°/s, 14.4°/s, 22.0°/s and 20.7°/s respectively). Figure 6.28a illustrates how the proposed QP target position changes with SP velocity and with different values of $\widehat{q}$. Increasing SP velocity causes a linear increase in the QP target position, in the negative direction. However, making the value of $\widehat{q}$ more negative also causes the QP target position to become more negative. These results indicate that QPs were being targeted into the negative field and not straight ahead. Calculating values of $\overline{y}$ and $\overline{x}$ from eq. 6.14 and 6.15 we find that the average amplitude of QPs and the average distance to the target location. The model predicts QPs on average undershoot the target location with mean values of -4.7°, -6.8°, -9.4° and -9.0°for trials with a stimulus speed of 10°/s, 20°/s, 30°/s and 40°/s respectively. When this undershoot bias was considered as saccadic gain (QP amplitude ÷ QP target amplitude) the mean value of the saccadic gain was 0.39. Figure 6.28b demonstrates how the mean QP start position, QP end position and QP target position relate to the mean SP velocity in the model, and illustrates the undershoot bias. Thus the model predicts that QPs have a tendency to undershoot their desired target location, possibly due to adaptive changes in the value of saccadic gain or as a result of not fully compensating for the movement away from the target location during the SP.

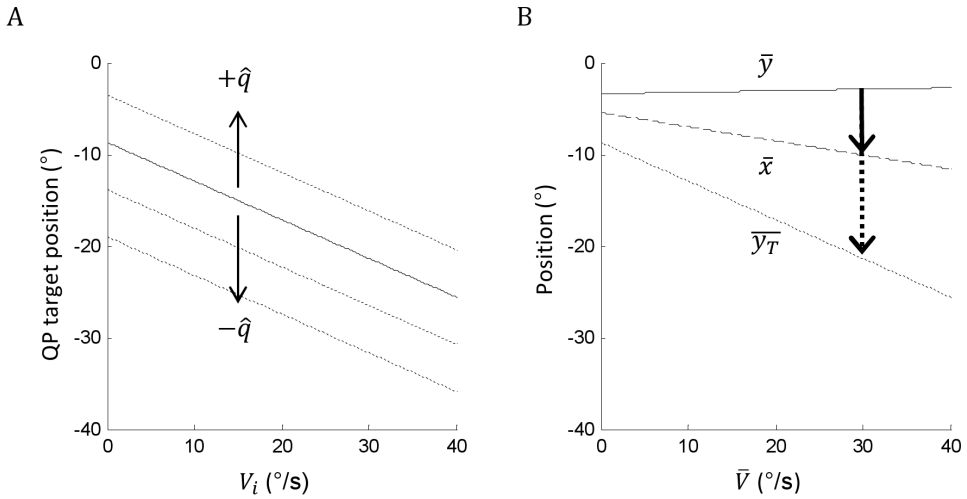

Figure 6.28: (A) Diagram illustrating the relationship between SP velocity and the proposed QP target position. Solid line indicates the proposed QP target position, using the mean value of $\widehat{q}$. Dotted lines indicate QP target position after increasing or decreasing $\widehat{q}$ by two each time. (B) Diagram illustrating the relationship between SP velocity and the proposed QP target position ($\overline{y_T}$), mean QP start position ($\overline{y}$), and mean QP end position ($\overline{x}$). Dotted arrow illustrates the amplitude of the target QP. Solid arrow illustrates the actual (mean) amplitude of QPs.

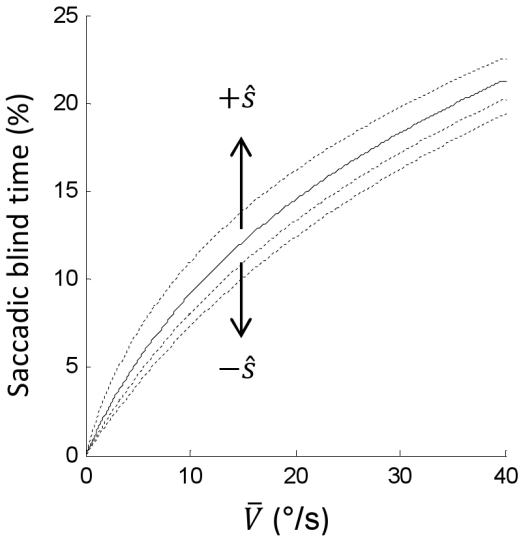

Figure 6.29: Relationship between SP velocity and the percentage of the total time spent in saccadic flight (saccadic blind time). Solid line indicate saccadic blind time, using the mean values of $\widehat{s}$ and $\widehat{q}$. Dotted lines indicate saccadic blind time after increasing or decreasing $\widehat{s}$ by two each time.

6.13 The effect of SP velocity on QP duration.

The average QP duration is related to the average SP duration through the dependence of $\overline{Q}$ on $\overline{V}$ and the dependence of $\overline{T}$ on $1/\overline{V}$. Substituting the average values of $\widehat{s}$ and $\widehat{q}$ from the data ($\widehat{s} = 0.51$ and $\widehat{q} = -3.38$) into eq. 6.18 and 6.17 the values of $\overline{T}$ and $\overline{Q}$ were found for a variety of different $\overline{V}$. I then used eq. 5.10 to find the values $\overline{T_Q}$, where $\overline{T_Q}$ is the average QP duration.

When $\overline{V} = 40°/s$, $\overline{T} = 222$ms and $\overline{T_Q} = 60$ms, and when $\overline{V} = 10°/s$, $\overline{T} = 374$ms and $\overline{T_Q} = 38$ms. This result illustrates that if SP velocity increases both the SP duration decreases and the QP duration increases regardless of stimulus speed. In this case increasing SP velocity from 10°/s to 40°/s lead to an increase in the average saccadic flight time from 9% to 21% of the total time. Figure 6.29 illustrates how the percentage of the total time spent in saccadic flight (saccadic blind time) changes with SP velocity and with different values of $\widehat{s}$. Increasing SP

velocity causes an approximately logarithmic increase in the saccadic blind time. Changing $\widehat{s}$ and $\widehat{q}$ has a relatively minor effect on the saccadic blind time, but in general increasing $\widehat{s}$ increases the saccadic blind time and decreasing $\widehat{s}$ decreases the saccadic blind time.

6.14 Discussion of results.

6.14.1 Discussion of RM-ANOVA results.

An interesting finding in the RM-ANOVA results was that increasing the spatial frequency of the stimulus caused an increase in the mean SP velocity, such that SP gain was increased and retinal slip was reduced. In chapter 2 we discussed that the contrast sensitivity drops rapidly when images move faster than 4°/s across the retina, but the relationship between contrast sensitivity and the speed at which the visual stimulus moves across the retina depends on the spatial frequency of the stimulus, and the contrast sensitivity to low spatial frequency gratings can peak at high levels of retinal slip. The spatial frequency of the square wave grating used in these experiments ranged from 0.05cyc/° to 0.248cyc/°, which equates to bar widths of 2° to 10°. The peak contrast sensitivity for detecting the motion of a bar that is 1° in width is approximately 4°/s, but at 2.9° in width it is between 10 and 20°/s, and at 9.8°it is between 30 and 40°/s (Burr & Ross 1982). The increase in retinal slip that is observed when the spatial frequency of the OKN stimulus is decreased may reflect an attempt to maximise visual contrast, by matching retinal slip to the velocity at which peak contrast sensitivity is achieved for the spatial frequency of that stimulus.

Changing the stimulus pattern also caused a change in SP velocity. Using a random grid pattern rather than a square wave grating caused the SP velocity to increase such that SP gain was increased and retinal slip was decreased, in a

similar manner to increasing the spatial frequency of the stimulus. If we assume that retinal slip does decrease with increasing spatial frequency in an attempt to maintain optimal visual contrast, we might consider that the constantly changing spatial frequency of the random grid stimulus pattern also causes retinal slip to change. However, as the spatial frequency of the stimulus is aperiodic, the degree of retinal slip that would be required to maximise visual contrast is not clear. If retinal slip changed as the spatial frequency of the stimulus currently being viewed changed, we might expect this effect to average out over the period of stimulation as the number of black and white squares across the width of the random grid pattern were matched to the number of black and white stripes in the square wave grating. However, this effect did not appear to average out, indicating that a more complex compromise might be being made.

6.14.2 Discussion of model results.

In this chapter we have also investigated the Markov model developed in chapter 5 further, in order to verify that the model can be fit to all of the OKN trials tested, and to examine the description and predictions the model makes about OKN behaviour. For example, we have used the model to demonstrate that OKN has a transient response of just a few cycles for SP velocity, which is in agreement with Abadi et al. (2005) who demonstrated that the steady state of SP velocity was reached in two cycles. The model also predicts a fast transient response for SP start position of just two or three cycles.

We found that the standard deviation for the Q-component (σ_q) is linearly related to the mean QP amplitude indicating that QPs are subject to proportional noise. This result is not entirely unexpected, as it is well known that larger saccades to static visual targets have larger errors. What is surprising is that the constant of proportionality is about 23%, which is much more than the typical values found

for saccades to static target found by van Beers (2008) ($\approx$ 10%). The reason for this is unknown but it may reflect targetting error for a moving target (note that retinal slip is usually not zero in OKN). The standard deviation for the S-component (σ_s) was also linearly related to the mean SP amplitude, indicating the presence of proportional noise in determining the threshold at which to trigger a QP, and this result was rather more unexpected.

We have also found that the noise associated with the V-component is related to stimulus velocity, in accordance with Kolarik et al. (2009) who found that variation in velocity between SPs increased linearly with stimulus speed. Stimulus speed, retinal slip, and $\overline{V}$ are each dependent on the others (retinal slip = stimulus speed - eye velocity), however the correlation between the error term in eq. 5.6 and stimulus speed was much stronger than the correlation of the error term with either retinal slip or $\overline{V}$. It is important to consider that only retinal slip, and eye velocity by way of efference copy (Sperry 1950), are readily available to the system whereas stimulus speed must be reconstructed internally. The error term expressed in eq. 5.6 then appears to reflect an internally generated noise, possibly in reconstructing stimulus velocity as part of a forward model whose output allows prediction of the motor output (a dynamic predictor). I propose that this is most likely explained by a variable gain in the OKN forward loop.

As we have noted, QP amplitude is influenced by proportional noise that increases linearly with the magnitude of the QP. The constant of proportionality for these movements is approximately 23%, meaning that QPs with an average magnitude of 2° would only have a standard deviation of 0.5° (plus 1.1°) whereas QPs with an average magnitude of 10° would have a standard deviation of 2.3° (plus 1.1°). The error introduced by the proportional noise in QPs would direct QPs of a large magnitude into a correspondingly larger target area, consistently overshooting and

undershooting the desired location. Making QPs smaller by reducing SP velocity would allow participants to make QPs to a more accurately predicted location. We have determined a possible location towards which QPs are targeted that appears to be dependent on SP velocity. However, we predict that QPs to this target constantly undershoot with a saccadic gain as low as 0.39, possibly as a result of not fully compensating for the movement away from the target location during the SP. Daye et al. (2010) demonstrated that memory guided saccades made during a period of head-unrestrained visual tracking could compensate for approximately 62% of the smooth gaze displacement. This compensation was a function of saccade latency that was maximal (62%) at latencies greater than 400ms, and reduced to no compensation for saccades made with a latency under 200ms. As the median SP duration is typically between 200 and 400ms we might expect that compensation of the movement away from the target during the SP would be somewhere between 0 and 62%.

It seems that correspondingly very small saccades would drastically lower the fraction of time spent making a SP during a given OKN cycle. The visual system integrates contrast of both stationary and moving stimuli over approximately 120ms (Burr 1981). Visual information during a SP of 12ms will have only one tenth the effective contrast obtained from a SP of 120ms or longer. Also, small saccades have only a small effect on improving the contrast of a tracked target (Harris & Berry 2006), such that their cost may exceed their benefit. It seems likely that it would not pay to make very small QPs at all, and this could explain the QP dead-zone that has been consistently observed in the data.

The data (and the model) demonstrate that, as SP velocity increases, QP amplitude increases and SP duration decreases regardless of stimulus speed. The increase in QP amplitude would cause a corresponding increase in QP duration

due to the main sequence of saccadic eye movements (Garbutt et al. 2001, 2003, Kaminiarz et al. 2009). This would result in a longer period of time making a QP during each cycle and, combined with the shorter period of time making a SP, the ratio of QP duration to the total duration of the OKN cycle would increase. Since high spatial frequency vision is lost during a QP, the average visual contrast per cycle would decrease. Closely matching stimulus velocity would maintain visual acuity during the SP but this could be offset by the contrast loss during the QP, and an optimal compromise might lead to an allowance for SP velocity to decrease in order to keep the QP amplitude low. However, visual acuity was not tested during this investigation and it is not possible to determine whether participants were attempting to maintain maximal visual contrast.

In Chapters 5 and 6 we have examined the parameters of human OKN cycle by cycle, and the results of PCA performed on this data have led us to consider OKN not as an archaic primitive oculomotor sub-sytem, but as three stochastic processes acting simultaneously in concert or in competition. It is possible that the goal of OKN is not to simply minimise retinal slip, but to find some optimal value of SP velocity that maximises the benefit (or minimises the cost) provided by each of these three processes. This hypothesis will be further explored in Chapter 8.

Chapter 7

Results of fitting SP duration histograms: the ratio distribution of two truncated normal variables.

In this chapter we find the maximum likelihood estimates (MLEs) of parameters for a number of PDFs to histograms of SP duration, and present the goodness of fit of these PDFs. The measures of goodness of fit from the PDFs predicted by the Markov model are then compared to the goodness of fit of PDFs predicted by other stochastic models of OKN that have been proposed in the literature, predominantly those that result from internal accumulator models.

7.1 Ratio distributions.

In our Markov model of OKN the SP duration is determined by the ratio of SP amplitude and SP velocity. SP velocity is described by a 1^{st} order Markov process. The update dynamics of SP amplitude are slightly more complicated, as SP amplitude during one cycle of OKN depends on the SP amplitude during the previous cycle (and not on the SP amplitude during earlier cycles), but also on the SP velocity during the previous cycle (eq. 5.15). However, a stable 1^{st} order Markov process tends towards a steady state where the probability density

approaches a fixed function, independent of the initial starting conditions. We have demonstrated that in the Markov model steady state SP velocity is reached in approximately 2 cycles (section 6.11).

For steady-state OKN, the model predicts that the probability distribution of SP duration would be a ratio of two random variables (S_i/V_i). We approximate the error terms of components with a Gaussian distribution although this is not strictly observed, as in chapter 6 we illustrated that the error terms had a tendency to be leptokurtic for approximately 1/3 of the trials. However, if normality of the error terms is assumed the resultant ratio distribution would be a positively skewed distribution that is at least qualitatively similar to the observed SP duration distribution. This distribution has been examined before in the context of saccadic latency (Nakahara et al. 2006), but as the result of an internal accumulator model and not as the result of the geometric relationships between primary eye movement parameters. However, the saccadic latencies investigated in the Nakahara et al. (2006) study were not recorded from OKN data, and there was no correlation between the threshold at which the saccades were triggered and the rate of the accumulating decision signal. Whereas the proposed correlation between SP amplitude and SP velocity is included in our analysis, and slightly complicates the PDFs of the proposed distributions. Two ratio distributions were tested for the goodness of fit of their PDF to the histograms of SP durations: the FRD and the DRTN. The explanation of how the PDF of these distributions is determined is given below.

7.1.1 Fieller's ratio distribution.

If X_1 and X_2 are two normally distributed random variables $X_1 \sim N(\mu_1, \sigma_1^2)$ and $X_2 \sim N(\mu_2, \sigma_2^2)$ with Pearson correlation coefficient ρ then let T be the random variable resulting from the ratio of X_1 and X_2 ($T = X_1/X_2$). Fieller (1932) derived

the PDF $f(\tau)$ for the distribution of this random variable, which we call Fieller's ratio distribution (FRD). We denote the variable $T \sim FRD(\mu_1, \sigma_1^2, \mu_2, \sigma_2^2, \rho)$, and give the PDF:

$$\begin{aligned} f(\tau) =& \frac{b(\tau)d(\tau)}{\sigma_1\sigma_2 a^3(\tau)\sqrt{2\pi}} \left[\Phi\left(\frac{b(\tau)}{a(\tau)\sqrt{1-\rho^2}} \right) - \Phi\left(-\frac{b(\tau)}{a(\tau)\sqrt{1-\rho^2}} \right) \right] \\ &+ \frac{\sqrt{1-\rho^2}}{\pi\sigma_1\sigma_2 a^2(\tau)} e^{-\frac{c}{2(1-\rho^2)}}, \end{aligned} \tag{7.1}$$

where

$$a(\tau) = \sqrt{\frac{\tau^2}{\sigma_1^2} - \frac{2\rho\tau}{\sigma_1\sigma_2} + \frac{1}{\sigma_2^2}}, \tag{7.2}$$

$$b(\tau) = \frac{\mu_1\tau}{\sigma_1^2} - \frac{\rho\left(\mu_1 + \mu_2\tau\right)}{\mu_1\mu_2} + \frac{\mu_2}{\sigma_2^2}, \tag{7.3}$$

$$c = \frac{\mu_1^2}{\sigma_1^2} - \frac{2\rho\mu_1\mu_2}{\sigma_1\sigma_2} + \frac{\mu_2}{\sigma_2^2}, \tag{7.4}$$

$$d(\tau) = e^{\frac{b^2(\tau) - ca^2(\tau)}{2(1-\rho^2)a^2(\tau)}}, \tag{7.5}$$

and Φ is the CDF of the standard normal distribution.

For the purposes of this study it is important to note that due to a simple scaling factor this distribution can be described by four free parameters rather than the five suggested by the formulation above. To characterise the distribution it is only necessary to know the ratio between the two means ($\kappa = \mu_1/\mu_2$), the CV of both original variables ($\lambda_1 = \sigma_1/\mu_1$ and $\lambda_2 = \sigma_2/\mu_2$), and the correlation coefficient (ρ).

The PDF described in eq. 7.1 can predict a small proportion of negative SP durations even when μ_1 and μ_2 are positive. In the context of our model this would equate to a SP with negative velocity and positive amplitude or vice versa, something that cannot happen in real terms without producing a SP of infinite duration. In order to overcome this problem I have assumed that these negative SP durations cannot happen and investigated the truncated PDF, with a lower bound of 0. The zero-truncated PDF $Tr(\tau)$ is given by renormalising the probability density in the interval $(0, \infty)$ such that it sums to 1, using the value of the CDF at $\tau = 0$. Hinkley (1969) derived the CDF $F(\tau)$ for this distribution:

$$\begin{aligned} F(\tau) =& L\left\{\frac{\mu_1 - \mu_2\tau}{\sigma_1\sigma_2 a(\tau)}, -\frac{\mu_2}{\sigma_2}; \frac{\sigma_2\tau - \rho\sigma_1}{\sigma_1\sigma_2 a(\tau)}\right\} \\ &+ L\left\{\frac{\mu_2\tau - \mu_1}{\sigma_1\sigma_2 a(\tau)}, -\frac{\mu_2}{\sigma_2}; \frac{\sigma_2\tau - \rho\sigma_1}{\sigma_1\sigma_2 a(\tau)}\right\}, \end{aligned} \tag{7.6}$$

where $L\{h, k; \gamma\}$ is the value at (h, k) of the CDF of a standard bivariate normal distribution with correlation coefficient γ. There is no analytical expression of the multivariate normal CDF but it can be approximated with numerical integration, and we use the Matlab function: *mvncdf* to do so. Thus $Tr(\tau)$ for $T \sim FRD(\kappa, \lambda_1, \lambda_2, \rho)$ is:

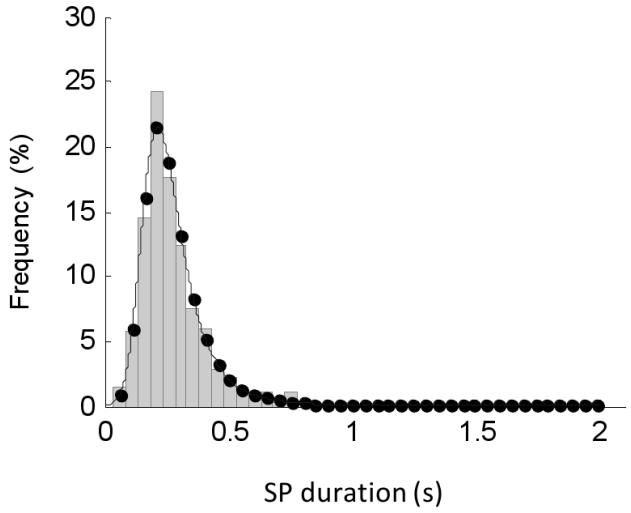

Figure 7.1: Histogram of SP duration from one trial (grey bars). PDF of the FRD fitted to the data (using MLEs of PDF parameters) (black line). Values of the FRD PDF at the central value of each bin (black dots). Data from participant 1, stimulus speed 10°/s.

$$Tr(\tau) = f(\tau)(1 - F(0)), \tau > 0 \tag{7.7}$$

where $f(\tau)$ is eq. 7.1 and $F(\tau)$ is eq. 7.6. An example of this function fitted to a histogram of SP duration is illustrated in figure 7.1. The results of formal goodness of fit tests are given later in this chapter (section 7.4).

7.1.2 Distribution of the ratio of two truncated normal variables.

Empirically SPs are never made in the opposite direction to the stimulus motion, a constraint that is also included in the Markov model of OKN. A more strict description of the distribution of SP duration then would be given by the ratio of two random variables that are defined only in the positive domain. Under this assumption the Markov model predicts that the SP duration distribution will be the distribution of the ratio of two truncated normal variables (DRTN). It is

necessary to consider the joint distribution of two normally distributed variables. Suppose X_1 and X_2 are two normally distributed random variables $X_1 \sim N(\mu_1, \sigma_1^2)$ and $X_2 \sim N(\mu_2, \sigma_2^2)$ with Pearson correlation coefficient ρ. The multivariate distribution $f(x_1, x_2)$ for $X \sim N_2\left(\begin{bmatrix} \mu_1 \\ \mu_2 \end{bmatrix}, \begin{bmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{bmatrix}\right)$ is:

$$G(x_1, x_2) = \frac{1}{2\pi\sigma_1\sigma_2\sqrt{1-\rho^2}} e^{\left(-\frac{1}{2(1-\rho^2)}\left[\frac{(x_1-\mu_1)^2}{\sigma_1^2} + \frac{(x_2-\mu_2)^2}{\sigma_2^2} - \frac{2\rho(x_1-\mu_1)(x_2-\mu_2)}{\sigma_1\sigma_2}\right]\right)} \quad (7.8)$$

The PDF of the ratio of these two variables $g(\tau)$ can then be derived through integration (Curtiss 1941), and is given by:

$$g(\tau) = \int_{-\infty}^{+\infty} |x_2|\, G(\tau x_2, x_2)\, dx_2 \quad (7.9)$$

However, if we consider both X_1 and X_2 with a lower bound of 0 the situation is more complicated. Harris (2011, personal communication) has derived a formula to calculate this PDF that can be found in appendix C. The PDF $f(\tau)$ for $T \sim DRTN(\kappa, \lambda_1, \lambda_2, \rho)$ is given by:

$$f(\tau) = -\frac{\exp(c)}{2ak_1}\left[1 + b\sqrt{\frac{\pi}{-4a}}\exp(-b^2/4a)\mathrm{erfc}(-b/2\sqrt{-a})\right], \tau > 0 \quad (7.10)$$

where

$$a = \frac{-1}{2(1-\rho^2)} \frac{\tau^2\sigma_{x_2}^2 - 2\tau\rho\sigma_{x_1}\sigma_{x_2} + \sigma_{x_1}^2}{\sigma_{x_2}^2\sigma_{x_1}^2} < 0 \tag{7.11}$$

$$b = \frac{1}{2(1-\rho^2)} \frac{2\tau\mu_{x_1}\sigma_{x_2}^2 + 2\mu_{x_2}\sigma_{x_1}^2 - 2\tau\rho\mu_{x_1}\sigma_{x_1}\sigma_{x_2}}{\sigma_{x_2}^2\sigma_{x_1}^2} \tag{7.12}$$

$$c = \frac{-1}{2(1-\rho^2)} \frac{\mu_{x_1}^2\sigma_{x_2}^2 - 2\rho\mu_{x_1}\mu_{x_2}\sigma_{x_1}\sigma_{x_2} + \mu_{x_2}^2\sigma_{x_1}^2}{\sigma_{x_2}^2\sigma_{x_1}^2} \tag{7.13}$$

and

$$k_1 = \int_0^\infty \int_0^\infty G(x_1, x_2) dx_1 dx_2 \tag{7.14}$$

where k_1 is found numerically. An example of this function fitted to a histogram of SP duration is illustrated in figure 7.2.

7.2 Distributions predicted by saccadic initiation and decision-making models.

In chapter 2 we discussed a number of stochastic models of OKN that have been used to describe the distribution of SP duration or, in a more general sense, inter-saccadic intervals and saccade latency. Here we formally introduce the equations that describe the PDF of SP duration, that are predicted by these models.

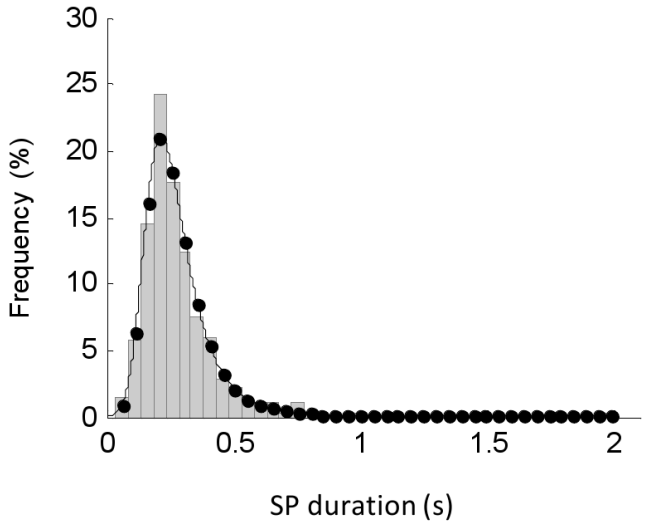

Figure 7.2: Histogram of SP duration from one trial (grey bars). PDF of the DRTN fitted to the data (using MLEs of PDF parameters) (black line). Values of the DRTN PDF at the central value of each bin (black dots). Data from participant 1, stimulus speed 10°/s.

7.2.1 Recinormal distribution.

The LATER model is a model of a decision process that has been used to describe the PDF of saccade latencies (Carpenter & Williams 1995) and OKN SP durations (Carpenter 1993). It can be simply described as a decision signal ($S(t)$) that rises at a linear rate (r) from an initial level (S_0) to a threshold (S_T), and generates a response (i.e. the QP) when the signal reaches this threshold. The time taken to make the decision to generate a QP is then given by $T = \Delta S/r$ where ΔS is the change in the level of the decision signal ($\Delta S = S_T - S_0$) required to generate the QP.

Importantly r varies randomly from decision to decision and is assumed to be a normally distributed random variable $r \sim N(\mu_r, \sigma_r^2)$. If r is normally distributed then the reciprocal decision time (also called "promptness") will also be a normally distributed random variable $\frac{1}{\tau} \sim N(\mu_p, \sigma_p^2)$ and the distribution of decision times is thus called recinormal. The mean rate of the decision signal is directly related

to the mean promptness by the distance to threshold ($\mu_r = \mu_p \Delta S$). Assuming there is no stochastic variation in ΔS we can see that the standard deviation of the rate is also directly related to the standard deviation of the promptness ($\sigma_r = \sigma_p \Delta S$).

Developing the normal PDF it is possible to obtain the PDF $f(\tau)$ of the recinormal distribution (RND) for the variable $T \sim RND(\mu_p, \sigma_p^2)$:

$$f(\tau) = \begin{cases} \frac{1}{\tau^2\sqrt{2\pi\sigma_p^2}} e^{-\frac{(1-\tau\mu_p)^2}{2\tau^2\sigma_p^2}} & \text{if } \tau \neq 0 \\ 0 & \text{if } \tau = 0 \end{cases} \tag{7.15}$$

This PDF is always bimodal and can predict a small proportion of negative SP durations even when μ_p is positive. As in the case of the FRD, we wish to define the variable T only in the positive domain. Thus the zero-truncated PDF $Tr(\tau)$ for $T \sim RND(\mu_p, \sigma_p^2)$ supported only in the positive domain is:

$$Tr(\tau) = \frac{1}{k_2\tau^2\sqrt{2\pi\sigma_p^2}} e^{-\frac{(1-\tau\mu_p)^2}{2\tau^2\sigma_p^2}}, \tau > 0 \tag{7.16}$$

where k_2 is a normalising constant given by:

$$k_2 = 1 - \Phi(-\mu/\sigma) \tag{7.17}$$

and Φ is the standard Normal CDF.

An example of this function fitted to a histogram of SP duration is illustrated in

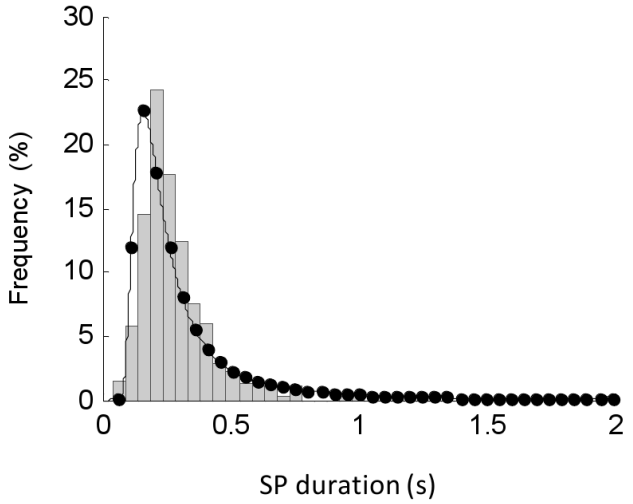

Figure 7.3: Histogram of SP duration from one trial (grey bars). PDF of the RND fitted to the data (using MLEs of PDF parameters) (black line). Values of the RND PDF at the central value of each bin (black dots). Data from participant 1, stimulus speed 10°/s.

figure 7.3.

7.2.2 Recinormal mixture distribution.

An extension to the LATER model has been provided to account for the existence of extremely short latency (express) saccades, and has also been used to model what appear to be SPs of extremely short duration (Carpenter 1994). This extension to the LATER model has two LATER units in parallel. The promptness of both units are normally distributed random variables $p_1 \sim N(\mu_1, \sigma_1^2)$ and $p_0 \sim N(\mu_0, \sigma_0^2)$. We consider the case where the decision signal with $p = p_1$ is described as a standard LATER unit and the decision signal with $p = p_0$ is described as a rogue unit with $\mu_0 = 0$ and $\sigma_0^2 > \sigma_1^2$. These two units then compete in a race to threshold and the response (i.e. the QP) is triggered when either signal reaches threshold first. This results in a recinormal mixture distribution (RNMD) where the probability of the standard unit reaching threshold first is given the probability p, and the probability of the rogue unit reaching threshold first is given by the

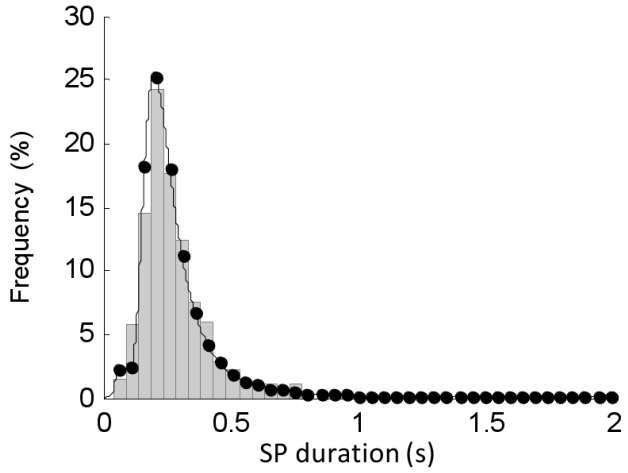

Figure 7.4: Histogram of SP duration from one trial (grey bars). PDF of the RNMD fitted to the data (using MLEs of PDF parameters) (black line). Values of the RNMD PDF at the central value of each bin (black dots). Data from participant 1, stimulus speed 10°/s.

probability $1-p$. The PDF $f(\tau)$ for $T \sim RNMD(\mu_1, \sigma_1^2, \sigma_0^2, p)$ supported only in the positive domain is:

$$f(\tau) = p\left(\frac{1}{k_3\tau^2\sqrt{2\pi\sigma_1^2}}e^{-\frac{(\tau-1-\tau\mu_1)^2}{2\tau^2\sigma_1^2}}\right) + (1-p)\left(\frac{1}{0.5\tau^2\sqrt{2\pi\sigma_0^2}}e^{-\frac{(\tau-1)^2}{2\tau^2\sigma_0^2}}\right) \quad (7.18)$$

where k_3 is a normalising constant given by:

$$k_3 = 1 - \Phi(-\mu_1/\sigma_1) \quad (7.19)$$

and Φ is the standard Normal CDF.

An example of this function fitted to a histogram of SP duration is illustrated in figure 7.4.

7.2.3 Inverse Gaussian distribution.

Another common accumulator model that randomly rises to some fixed threshold before activating a response can be described by a decision signal subject to a random walk, with a drift rate that tends to cause the decision signal to rise. This is the well-known diffusion process, the PDF of which can be described by the inverse Gaussian distribution (IGD). The PDF of the IGD $f(\tau)$ for $T \sim IGD(\mu, \lambda)$ supported only in the positive domain is:

$$f(\tau) = \sqrt{\frac{\lambda}{2\pi\tau^3}} e^{\frac{-\lambda(\tau-\mu)^2}{2\mu^2\tau}}, \tau > 0 \tag{7.20}$$

where μ is the mean and λ is the shape parameter. An example of this function fitted to a histogram of SP duration is illustrated in figure 7.5. This model was previously proposed by Anastasio (1996) who suggested that the the integration of noisy vestibular nucleus neurons could be described in the same manner.

7.2.4 Gamma distribution.

Another simple model of a decision process that rises to threshold before activating a response is described by a neuron (or network of neurons) that integrates incoming spikes that appear at randomly distributed time intervals until a certain number of spikes have been received (Tuckwell 1988). When the random intervals of spikes is exponentially distributed, the distribution of the total decision time (T) is given by a gamma distribution (GD). The gamma distributed variable is then given by $T \sim GD(a, b)$ where a is the threshold number of spikes and $1/b$ is the probability per unit time of a spike occurring. A GD of the first order where $a = 1$ is simply an exponential distribution with a rate parameter $\lambda = 1/b$.

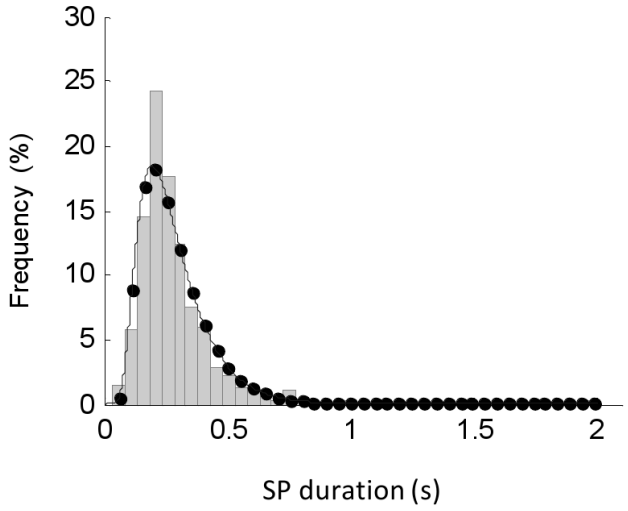

Figure 7.5: Histogram of SP duration from one trial (grey bars). PDF of the IGD fitted to the data (using MLEs of PDF parameters) (black line). Values of the IGD PDF at the central value of each bin (black dots). Data from participant 1, stimulus speed 10°/s.

If Γ is the gamma function defined by the integral:

$$\Gamma(a) = \int_0^\infty e^{-t} t^{a-1}\, dt \tag{7.21}$$

then the PDF of the GD $f(\tau)$ for $T \sim GD(a, b)$ supported only in the positive domain is:

$$f(\tau) = \frac{1}{b^a \Gamma(a)} \tau^{a-1} e^{-\frac{\tau}{b}}, \tau > 0 \tag{7.22}$$

where a is the shape parameter and b is the rate parameter of the distribution. An example of this function fitted to a histogram of SP duration is illustrated in figure 7.6.

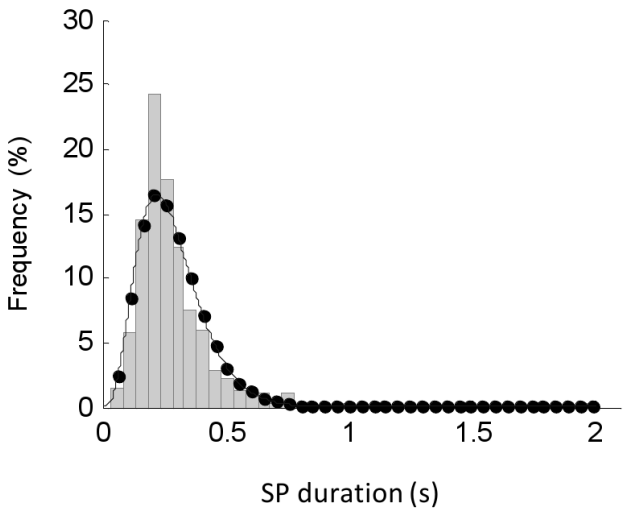

Figure 7.6: Histogram of SP duration from one trial (grey bars). PDF of the GD fitted to the data (using MLEs of PDF parameters) (black line). Values of the GD PDF at the central value of each bin (black dots). Data from participant 1, stimulus speed 10°/s.

7.2.5 Lognormal distribution.

In a study on OKN in turtles, Balaban & Ariel (1992) described the duration of SPs with the lognormal distribution (LND), and proposed that this was based on an integrate-and-fire neuron model. The distribution they proposed would arise if SP duration was determined by the product of the duration of the previous SP and a zero-truncated normal variable with a mean of 1. The PDF of the LND $f(\tau)$ for $T \sim LND(\mu, \sigma)$ supported only in the positive domain is:

$$f(\tau) = \frac{1}{\tau\sigma\sqrt{2\pi}} e^{-\frac{(ln\tau-\mu)^2}{2\sigma^2}}, \tau > 0 \tag{7.23}$$

where μ is the location parameter, and σ is the scale parameter of the distribution. An example of this function fitted to a histogram of SP duration is illustrated in figure 7.7.

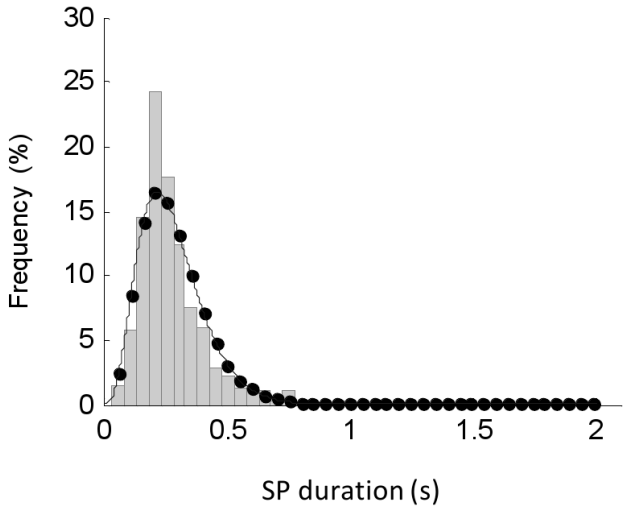

Figure 7.7: Histogram of SP duration from one trial (grey bars). PDF of the LND fitted to the data (using MLEs of PDF parameters) (black line). Values of the LND PDF at the central value of each bin (black dots). Data from participant 1, stimulus speed 10°/s.

7.3 Distribution fitting.

7 PDFs were tested for their goodness of fit to the histogram of SP duration from each trial. The distributions tested were the Fieller's ratio (FRD), the distribution of the ratio of two truncated normal variables (DRTN), the recinormal (RND), the recinormal mixture (RNMD), the inverse Gaussian (IGD), the gamma (GD), and the lognormal (LND) distributions. Further details of the tested distributions, including the equations for calculating the PDF, and the number and interpretation of parameters required to define each PDF, is given in Chapter 7. Likewise, the details for calculating the cumulative distribution function for the FRD and the DRTN are also given in Chapter 7.

To fit the distributions to each SP duration histogram maximum likelihood estimates (MLEs) of the parameters for each distribution were obtained using the function from the Matlab statistics toolbox: *mle*. This routine searches the parameter space for a minimum value of the negative log likelihood criterion. The

search is conducted using a simplex search algorithm: *fminsearch* that is both robust and can find the minimum of a multiparameter function effectively (Lacouture & Cousineau 2008, Myung 2003). It is particularly useful as it allows the user to define a custom PDF to fit to the data, enabling the fitting procedure for PDFs that are not supported by the statistics toolbox such as truncated distributions and the DRTN. A simplex is a d-dimensional polygon with $d + 1$ vertices, where d is the number of parameters to be estimated. A simplex search works by calculating the values of the objective function at each corner of the simplex and then reorienting the simplex around the corner with the lowest value of the function, and either contracting or expanding. This search continues for a number of iterations until the simplex is smaller than some minimum volume (Nelder & Mead 1965).

The maximum number of iterations of the simplex search was set to $2{\cdot}10^3$, and the maximum number of function evaluations allowed was set to $5{\cdot}10^3$. Termination tolerance on the values of parameters was set to $1{\cdot}10^{-6}$, and the termination tolerance on the values of the function value was set to $1{\cdot}10^{-6}$. These values were chosen in order to give sufficient time for the simplex search to converge. Estimates of the correlation coefficient could not be reliably retrieved for FRD and DRTN using the simplex search to minimise maximum likelihood, as the maximum likelihood function became "lumpy" when this parameter was included in the search. The estimates of the correlation coefficient for these PDFs was instead fixed to the value of the sample correlation coefficient between the SP amplitude and SP velocity for the trial from which the SP duration histogram had been calculated. A lower bound of zero was set for the estimation of parameters only defined in the positive domain, such as the shape and scale parameters of GD, the mean and shape parameters of IGD, and values of standard deviation.

An upper bound of ∞ was set for the estimation of all parameters.

After finding the MLEs for each distribution the goodness of fit was also estimated with the chi-square criterion for the purpose of comparing the goodness of fit between the different PDFs tested. This test is sensitive to the choice of bin sizes and so they were matched to those bin sizes from a similar study (Trillenberg et al. 2002), where each SP duration histogram was covered by 40 bins of equal size. The observed frequency was taken from the histogram, and the expected frequency was calculated using the PDF of the distribution under study with the parameters obtained from MLE. The chi-square criterion is given by: $\chi^2 = \int_{i=1}^{k}(O_i - E_i)^2/E_i$, where O_i is the observed frequency for bin i and E_i is the expectd frequency for bin i. The actual χ^2 statistic was calculated using the Matlab function: *chi2gof* by specifying the observed and expected bin counts. This test performs poorly if the expected frequency in any given bin is small and so *chi2gof* combines bins in the tails of each distribution until there is a minimum expected frequency of 5 in each bin.

The reduced chi-squared test statistic calculated follows the χ^2 distribution with $\upsilon = N - 1 - n$ DOF, where N is the number of non-empty bins in the histogram and n is the number of parameters estimated in the distribution function. The goodness of fit of a test distribution to a single histogram can be tested by considering the null hypothesis that the data are from the tested distribution. The null hypothesis must be rejected if $\chi^2 > \chi^2_{\alpha,\upsilon}$, where $\chi^2_{\alpha,\upsilon}$ is the chi-square inverse CDF and α is the level of significance. In some trials with small data sets and highly peaked distributions there were actually more parameters to estimate than bins, after bins in the tails have been combined. In these instances the total DOF were reduced to zero (or less) and the null hypothesis was automatically rejected. The number of trials for which this occured are given in the results in chapter 7.

Testing the null hypothesis on multiple histograms simultaneously requires a method to counter the problem of multiple testing. In a situation where it is more interesting to retain the null hypothesis, simple Bonferroni correction is actually non-conservative. Instead, the Holm-Bonferroni method was applied in order to retain the power of the test procedure. The Holm-Bonferroni method first requires that the p-values of all the null hypotheses tested are sorted in ascending order. The smallest p-value was compared to the value α/k, where α was the overall significance level (type 1 error rate) and k was the number of null hypotheses to be tested. If that p-value was less than α/k, then that hypothesis was rejected and the process was repeated with $k-1$ hypotheses, and so the next p-value was compared to the value $\alpha/k-1$. This process was continued until the smallest p-value could not be rejected. The remaining null hypotheses, that were not rejected at previous steps, were all accepted.

The results of fitting one histogram with two different PDFs were compared by taking the ratio of the two test statistics scaled appropriately by their DOF. $F = (\chi_1^2/\upsilon_1)(\chi_2^2/\upsilon_2)$ follows the F distribution with υ_1, υ_2 DOF, where χ_1^2 and χ_2^2 are the test statistics derived from fitting two different distributions to a single histogram and υ_1 and υ_2 are their respective DOF. The cumulative results of fitting all histograms with two different PDFs was also compared by summing all the values of χ_1^2, χ_2^2, υ_1 and υ_2 from each histogram tested, such that $F = (\sum \chi_1^2/\sum \upsilon_1)(\sum \chi_2^2/\sum \upsilon_2)$ follows the F distribution with $\sum \upsilon_1$ and $\sum \upsilon_2$ DOF.

7.4 Results of testing goodness of fit.

On testing the goodness of fit of PDFs from trials with small data sets and highly peaked distributions there were some cases where there were more parameters to estimate than there were bins, after bins in the tails had been combined. This

occurred when testing the FRD for 5 histograms, the DRTN for 5 histograms, the RND for 7 histograms, the RNMD for 7 histograms, the IGD for 6 histograms, the GD for 6 histograms, and the LND for 4 histograms. These histograms were generally from the same group of trials and the low number of bins was due to a large number of empty bins between the mode of the distribution and extreme outliers in the positive tail of these specific trials. In these cases, where we were unable to determine the reduced chi-squared statistic we assumed that the model was significantly different to the observed values. The reduced chi-squared statistics were almost always greater than 1, and there was no evidence of "over-fitting" the data.

When testing the goodness of fit to individual histograms, out of 256 trials the FRD demonstrated a significantly good fit (Holm-Bonferroni corrected) for 244 trials, the DRTN for 245, the RNMD for 229, the LND for 207, the GD for 197, the IGD for 181, and the RND for 114. This indicated three different groups of PDFs, one group that gave a good fit to over 95% of the histograms of SP duration (FRD and DRTN), one that gave a good fit to between 71% and 89% of the histograms (IGD, GD, LND and RNMD), and one that gave a good fit to only 45% of the histograms (RND).

When comparing the goodness of fit from different PDFs to individual histograms only the RND produced significantly worse fits than the other PDFs. The RND was significantly worse than the FRD in 21 cases, the DRTN in 20 cases, the GD in 20 cases, the RNMD in 7 cases, and the LND in 7 cases.

Although there were relatively few significant differences in the goodness of fit of different PDFs for single histograms, there were a number of significant differences in the goodness of fit when the chi-squared values and the DOF for the individual trials were added, and combined for testing. The table 7.1 lists the values for the

PDF	$\sum \chi^2$	$\sum f$	$\sum \chi^2 / \sum f$
FRD	3597	2083	1.73
DRTN	3650	2088	1.75
LND	5731	2654	2.16
GD	6538	2914	2.24
RNMD	3843	1673	2.30
IGD	6842	2615	2.62
RND	17297	2107	8.21

Table 7.1: Table of the reduced chi-squared statistics for goodness of fit tests between PDFs for multiple histograms. $\sum \chi^2$ is the sum of the reduced chi-squared statistics from all histograms tested, for each PDF. $\sum f$ is the sum of the DOF from all histograms tested, for each PDF.

$\sum \chi^2$, $\sum f$, and $\sum \chi^2 / \sum f$.

The goodness of fit of the FRD and the DRTN were compared to the goodness of fit of the other PDFs. There was no significant difference in goodness of fit between FRD and DRTN ($F_{2083,2088} = 0.988$, $p = 0.39$), but FRD was a significantly better fit than the LND ($F_{2083,2654} = 0.799$, $p = 3.41 \cdot 10^{-8}$), GD ($F_{2083,2914} = 0.767$, $p = 4.43 \cdot 10^{-11}$), RNMD ($F_{2083,1673} = 0.755$, $p = 5.88 \cdot 10^{-10}$), IGD ($F_{2083,2615} = 0.657$, $p = 8.60 \cdot 10^{-24}$), and RND ($F_{2083,2107} = 0.211$, $p = 6.09 \cdot 10^{-252}$). The DRTN was a significantly better fit than the LND ($F_{2088,2654} = 0.809$, $p = 1.62 \cdot 10^{-7}$), GD ($F_{2088,2914} = 0.776$, $p = 2.86 \cdot 10^{-10}$), RNMD ($F_{2088,1673} = 0.765$, $p = 3.28 \cdot 10^{-9}$), IGD ($F_{2088,2615} = 0.665$, $p = 1.34 \cdot 10^{-22}$), and RND ($F_{2088,2107} = 0.214$, $p = 1.55 \cdot 10^{-248}$).

The goodness of fit of the remaining PDFs were also compared. The LND was a significantly better fit than the IGD ($F_{2654,2615} = 0.825$, $p = 4.21 \cdot 10^{-7}$) and the RND ($F_{2654,2107} = 0.266$, $p = 1.13 \cdot 10^{-221}$), but not the GD ($F_{2654,2914} = 0.963$, $p = 0.16$) nor the RNMD ($F_{2654,1673} = 0.947$), $p = 0.11$). The GD was a

significantly better fit than the IGD ($F_{2914,2615} = 0.857$, $p = 2.54 \cdot 10^{-5}$) and the RND ($F_{2914,2107} = 0.276$, $p = 1.68 \cdot 10^{-222}$), but not the RNMD ($F_{2914,1673} = 0.985$, $p = 0.37$). The RNMD was a significantly better fit than the IGD ($F_{1673,2615} = 0.869$, $p = 8.18 \cdot 10^{-4}$) and the RND ($F_{1673,2107} = 0.322$, $p = 2.80 \cdot 10^{-149}$). The IGD was a significantly better fit than the RND ($F_{2615,2107} = 0.322$, $p = 2.65 \cdot 10^{-163}$).

7.5 Discussion of results.

The Markov model predicts that the distribution of SP duration is a ratio distribution of two random variables S_i and V_i. The results of this analysis give compelling evidence to support this hypothesis.

The results appear to indicate three groups of PDFs that each have a different level of goodness of fit for the histograms of SP duration. The group with the poorest fit consisted solely of the RND, which was significantly different to between 55% of the histograms tested. The second group consisted of the IGD, GD and LND, and represented PDFs that were significantly different to approximately 10% to 30% of histograms. The third group consisted of the DRTN and the FRD which were significantly different to less than 5% of histograms.

Comparing the goodness of fit from different PDFs to individual histograms, revealed only significantly worse fits from the RND than with other distributions. However, comparing the goodness of fit from different PDFs using the combined values for the reduced chi-squared statistic and DOF revealed a number of statistically significant differences. While the goodness of fit of the FRD and DRTN were not significantly different to each other, they were both significantly better fits than the distributions proposed by the accumulator models.

It is often reported that the distribution of saccade latency and more generally reaction times, appear to have the same characteristic shape, that of a positively

skewed distribution. It is not the intention of this analysis to declare that the FRD or the DRTN are the best distributions for describing all intersaccadic interval distributions, but rather that they are representative of the distribution of the SP duration specifically. There is no *a priori* reason to believe that because the distribution of saccadic latency to a visual target appears to be best described by one distribution (such as the RND described by the LATER process), that the same distribution necessarily describes all intersaccadic intervals.

The relationship proposed in the Markov model between SP velocity and SP duration, as defined by the constraint in eq. 5.3, is the only nonlinear component that we have allowed for during the development of the model. This is because PCA is only capable of analysing the linear relationships (correlations) between parameters. The analysis performed here is important in determining the validity of this imposed constraint, and indicates that not only do the distributions predicted by the Markov model provide a good fit to the SP duration distribution, they are also a significanlty better fit than a number of other PDFs that have been proposed to fit this distribution.

Chapter 8

Proposal for a cost model of OKN: does OKN minimise the error in tracking a moving target?

This chapter proposes that mean SP velocity is optimal for some function given the constraints of the Markov model, and despite the apparently low gains observed in the data. I investigate three possible cost functions that the system might attempt to minimise: the cost of retinal slip, the cost of saccadic blind time, and the cost of error in tracking a moving target. We solve the problem of finding the optimal mean SP velocity that minimises the proposed functions, and illustrate that the observed mean SP velocity is near optimal for minimising the cost of error in tracking a moving target.

8.1 The cost of retinal slip.

The OKR is generally considered a reflex whose primary aim is to minimise retinal slip. If this is the case then we would expect the mean SP velocity to remain close to the OKN stimulus speed, in order to keep retinal slip close to zero. We can propose a cost function to describe this, where cost is zero when retinal slip is zero and the cost increases linearly as retinal slip increases in magnitude. We define

this cost function:

$$J_{RS1} = \alpha \left| V_S - \overline{V} \right| \tag{8.1}$$

where V_S is the stimulus speed and $\overline{V}$ is the mean SP velocity.

However, we might suppose that the reason that OKN is attempting to minimise retinal slip is to maximise the contrast of the visual image. As discussed in chapter 2, the peak contrast sensitivity for a stimulus of a given spatial frequency is non-zero, and the contrast sensitivity function decays approximately logarithmically at speeds greater than the optimal value. We assume that the system is either attempting to maximise the contrast sensitivity for the spatial frequency content of the visual scene, or that it is attempting to maximise the contrast of high spatial frequency content.

The width of the bars used for the OKN stimulus in experiment 1 were 5°. On a log-log plot a function of the form $y = ax^b$ will appear as a straight line in which b will be the slope of the line and a will be the y value corresponding to $x = 1$. Approximating from fig. 2.1 (pg. 8) we find that the slope of the decaying contrast function as image motion increases is ≈ -1.5, and the $x = 1$ intercept is ≈ 9000. The slope of the decay as image motion decreases is ≈ 0.5, and the $x = 1$ intercept is ≈ 30. The intercept of these two lines is at $[\sqrt{300}, 30(300^{0.25})]$, thus the velocity at the peak of the contrast sensitivity curve is $\approx 17°/s$. This seems quite a good match to the peak on the contrast sensitivity curve for a bar of 2.9° so is a slightly conservative estimate for bars with a width of 5°(fig. 2.1). As with our cost function for retinal slip, we set zero cost to be at the retinal slip speed that gave the optimal contrast. We define this cost function as:

$$J_{RS2} = \alpha \begin{cases} 30(300^{0.25}) - (30\left|V_S - \overline{V}\right|^{0.5}) & \text{if } \left|V_S - \overline{V}\right| < \sqrt{300} \\ 0 & \text{if } \left|V_S - \overline{V}\right| = \sqrt{300} \\ 30(300^{0.25}) - (9000\left|V_S - \overline{V}\right|^{-1.5}) & \text{if } \left|V_S - \overline{V}\right| > \sqrt{300} \end{cases} \tag{8.2}$$

If the system is attempting to maximise the contrast of high spatial frequency content we would anticipate that the optimal retinal slip speed would be much closer to zero. We extrapolate from fig. 2.1 a second contrast sensitivity function that has an intercept at $[\sqrt{10}, 50(10^{0.25})]$, thus the velocity at the peak of the contrast sensitivity curve is $\approx 3^\circ$/s. We define this cost function as:

$$J_{RS2} = \alpha \begin{cases} 50(10^{0.25}) - (50\left|V_S - \overline{V}\right|^{0.5}) & \text{if } \left|V_S - \overline{V}\right| < \sqrt{10} \\ 0 & \text{if } \left|V_S - \overline{V}\right| = \sqrt{10} \\ 50(10^{0.25}) - (500\left|V_S - \overline{V}\right|^{-1.5}) & \text{if } \left|V_S - \overline{V}\right| > \sqrt{300} \end{cases} \tag{8.3}$$

8.2 The cost of saccadic blind time.

In chapter 6 we demonstrated the effect that SP velocity has on QP duration and SP duration, and illustrated how increasing mean SP velocity resulted in a logarithmic increase in the mean saccadic blind time. Harris (1995) investigated the possibility that saccadic undershoot is an economical strategy for maximising the period of clear vision, by minimising saccadic flight time. Here, we consider the possibility that the OKN system is also attempting to maximise the period of clear vision by minimising the duration of QPs. We assume a mean QP duration of zero has zero cost, and define the cost function by substituting eq. 6.17 into

eq. 5.10:

$$J_{MS} = \beta 18.5 \left| -0.17\overline{V} - 0.79\bar{s} + 0.41\bar{q} \right|^{0.54} \tag{8.4}$$

We take $\bar{s}$ and $\bar{q}$ to be the mean values of $\widehat{s}$ and $\widehat{q}$ estimated from all forty trials in experiment 1. We do not include estimates of $\widehat{s}$ and $\widehat{q}$ from experiments 2 and 3 for simplicity, as the spatial frequency content of the stimulus in these trials was variable, and we intend to investigate only the mean SP velocity that is optimal for a given stimulus speed and cost function. We found the values of $\bar{s}$ and $\bar{q}$ to be 0.925 and -2.635 respectively.

8.3 The cost of target error.

In chapter 2 we discussed the possibility that foveal tracking plays an important part in maintaining OKN gain. Here, we consider the case where a visual target is moving at a constant speed over a period of time T_i, from a position p to q in space (fig. 8.1a), and an eye movement attempts to follow the target with an initial eye position error and a constant (linear) eye velocity. We define the eye movement start position as x and the end position as y, this is illustrated in figure 8.1b, where $x \neq p$ and $y \neq q$. It is possible to consider the relative position of the eyes to the target by subtracting the value of p from the start position of both the target movement and the eye movement, and the value of q from the end position of both movements, such that we consider the relative position of the target as zero (fig. 8.1c). We then consider the total positional error (TPE) over the period of the movement to be the absolute relative position of the two movements integrated over the period of the movement (fig. 8.1d).

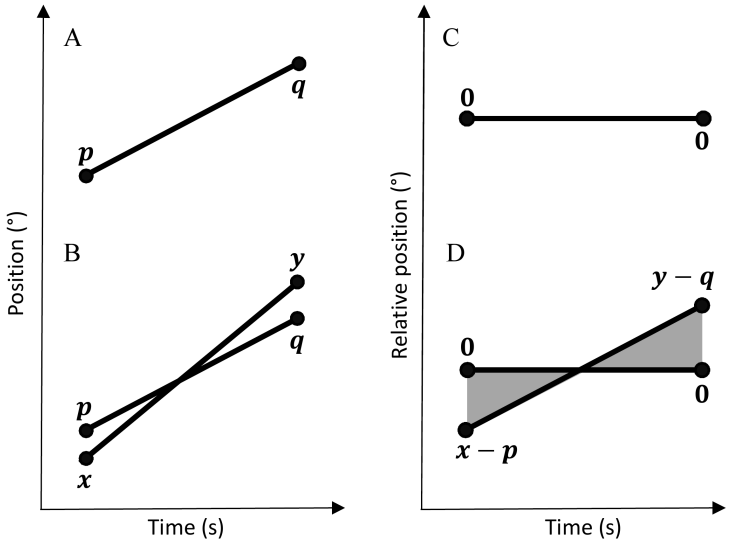

Figure 8.1: An illustration of (A) a target moving at constant speed from position p to position q, and (B) an eye movement attempting to follow the moving target at a constant speed, but starting from a different position to the target. (C) Relative position measured over the course of the movement is defined as zero for the target movement, and so (D) the relative start position of the eye movement is $x - p$ and the relative end position of the eye movement is $y - q$. The shaded area between these two lines represents the TPE (total positional error) over the period of the movement.

There are six fundamentally different patterns of eye movements that can occur in tracking a moving target when we consider this arrangement. The first, is perfect tracking, where the saccade to the target is exact and the velocity of the following movement is also exact (fig. 8.2a). The second, is velocity error only, where the saccade to the target is exact but the velocity of the following movement is not exact (fig. 8.2b). The third, is positional error only, where the saccade to the target is not exact but the velocity of the following movement is (fig. 8.2c). The fourth, is positional error with non-compensating velocity, where the saccade to the target is not exact and the velocity error brings the eyes even further from the target over the duration of the movement (fig. 8.2d). The fifth, is positional error with compensating velocity, where the saccade to the target is not exact and the velocity error brings the eyes closer to the target over the duration of the movement (fig. 8.2e). The sixth, is positional error with over-compensating velocity, where the saccade to the target is not exact and the velocity error brings the eyes back to the target during the movement but then continues past the target for the remainder of the movement (fig. 8.2f).

We can consider this movement to be a QP made to a target position and the following SP. We define the cost due to not remaining on target during the SP to be proportional to the positional error between the eye and the target, integrated over the course of the SP. We assume that the velocity of the target is always equal to the OKN stimulus speed, and that eye velocity is always constant during a SP. In the case where the trajectory of the eye and the trajectory of the target do not cross this is given by:

$$J_{TPE} = \gamma \left| \int_{t_0}^{t_0+T_i} (V_S t + p_i) - (\overline{V} t + x_i) dt \right| \tag{8.5}$$

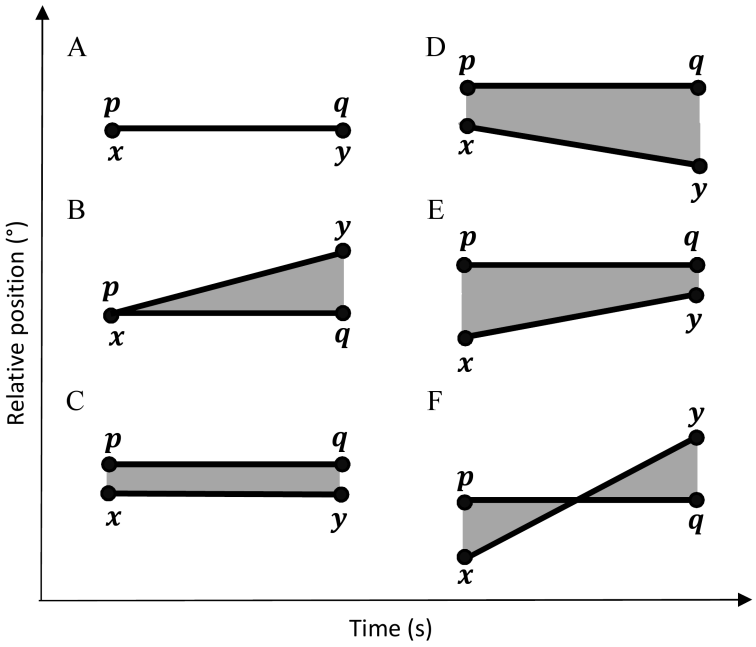

Figure 8.2: Illustration of six different patterns of eye movements when tracking a moving target, which generate different patterns of positional error. (A) The eyes follow the target exactly. (B) The eyes start in the correct position but over the course of the movement drift away from the relative position of the target. (C) Eye position starts with an initial error, and over the course of the movement they move at the same speed as the target. (D) Eye position starts with an initial error, and over the course of the movement they drift further away from the relative position of the target. (E) Eye position starts with an initial error, and over the course of the movement drift closer to the relative position of the target. (F) Eye position starts with an initial error, and over the course of the movement drifts passed the relative position of the target.The shaded area between these two lines represents the TPE over the period of the movement.

where γ is the proportional cost of the TPE, t_0 is the time when the movement begins, T_i is the duration of the movement, V_S is the stimulus speed, p_i is the start position of the target during that movement, $\overline{V}$ is the eye velocity during that movement, and x_i is the start position of the eye during that movement.

In the case where the trajectory of the eye and the trajectory of the target do cross this is given by:

$$J_{TPE} = \gamma \left(\left| \int_{t_0}^{(x-p)/(V_S-\overline{V})} (V_S t + p_i) - (\overline{V} t + x_i) dt \right| + \left| \int_{(x-p)(V_S-\overline{V})}^{t_0+T_i} (V_S t + p_i) - (\overline{V} t + x_i) dt \right| \right) \tag{8.6}$$

8.3.1 Constraints of the Markov model.

We examine the cost due to positional error numerically by simulating $2.5{\cdot}10^5$ OKN cycles, consisting of a QP followed by a SP, and we make some assumptions about each OKN cycle here, based on the Markov model that we have developed. We will consider a mean position from which all QPs are made:

$$\overline{y_0} = 0.017\overline{V} + 1.78\bar{s} + 1.28\bar{q} \tag{8.7}$$

where $\overline{y_0}$ is the average QP start position. We assume that a target, with velocity V_S has a start position defined by the difference between $\overline{y_0}$ and the mean QP amplitude defined by the Markov model:

$$p_i = \overline{y_0} - (-0.164\overline{V} - 0.389\overline{y_0} + \bar{q}) \tag{8.8}$$

The QP amplitude during each cycle is given by:

$$Q_i = p_i - \overline{y_0} + \epsilon_q(i) \tag{8.9}$$

where ϵ_q is a normal random variable with zero mean, and a standard deviation σ_q that varies from cycle to cycle (generated with the Matlab function: *randn*). This sequence of variables is not strictly independent and identically distributed due to signal dependent noise causing the standard deviation to change during each cycle such that $\sigma_q(i)$ is proportional to Q_i:

$$\sigma_q = 0.23(p_i - \overline{y_0}) \tag{8.10}$$

The start position of each SP is defined as:

$$x_i = \overline{y_0} - Q_i \tag{8.11}$$

We constrain the SP amplitude to be always positive, and define it as:

$$\begin{cases} S_i = -0.282x_i + 0.127\overline{V} + \bar{s} + \epsilon_s(i) & \text{if} - 0.282x_i + 0.127\overline{V} + \bar{s} + \epsilon_s(i) \geq 0 \\ S_i = 0 & \text{if} - 0.282x_i + 0.127\overline{V} + \bar{s} + \epsilon_s(i) < 0 \end{cases} \tag{8.12}$$

where ϵ_s is a normal random variable with zero mean, and a standard deviation σ_s that varies from cycle to cycle (generated with the Matlab function: *randn*). This sequence of variables is not strictly independent and identically distributed due to signal dependent noise, as σ_s is proportional to the S_i such that:

$$\sigma_s = 0.23S_i \tag{8.13}$$

We define the SP duration as:

$$T_i = S_i/\overline{V} \tag{8.14}$$

In this definition of the model we have not added a constant component of noise to either σ_q or σ_s. However, we shall investigate the effect of changing the level of constant noise in the results.

We then simulate the $2.5{\cdot}10^5$ OKN cycles and calculate the cost J_{TPE} for each cycle using equations 8.5 and 8.6, and find the mean cost across all SPs in order to calculate our cost function for J_{TPE}

8.4 Cost model results.

In order to determine if SP velocity is optimised for a given cost function we first consider the observed values of SP velocity in experiment 1, where the spatial frequency of the stimulus was not changed. The mean values of SP velocity were 8.6 (SD=1.2), 14.4 (SD=3.2), 18.5 (SD=4.9) and 20.7°/s (SD=8.1) when stimulus speeds were 10, 20, 30 and 40°/s respectively. Note that retinal slip increases with stimulus speed such that mean SP gain was 0.86 (SD=0.12), 0.72 (SD=0.16), 0.62 (SD=0.16), and 0.52 (SD=0.24) at each speed respectively. We now investigate whether the mean SP velocity that minimises each cost function is equivalent to the observed values of SP velocity.

The cost functions J_{RS1}, J_{RS2} and J_{RS3} are illustrated in fig. 8.3 for stimulus speeds of 10, 20, 30 and 40°/s when $\alpha = 1$. It is quite clear that for each individual cost function, the optimal value of mean retinal slip will remain the same regardless of stimululs speed. For J_{RS1} this value is zero, and for J_{RS2} and J_{RS3} it is an absolute value of retinal slip that is equivalent to the image motion that produces peak contrast sensitivity: $\approx \pm 17$°/s for J_{RS2}, and $\approx \pm 3$°/s for J_{RS3}. We also note that the optimal value of SP velocity for minimising J_{RS2} at a stimulus speed of 10°/s is much higher than the stimulus speed. As the optimal value of retinal slip is $\approx \pm 17$°/s for minimising this cost function, the lower value of SP velocity that would minimise cost would actually be negative, and made in the opposite direction to stimulus motion. In the data we find that retinal slip increases with stimulus speed and does not remain the same, and certainly does not remain at zero. Therefore, the observed mean SP velocity is not optimal for minimising J_{RS1}, J_{RS2} or J_{RS3}.

The cost function J_{MS} is illustrated in fig. 8.4. As stimulus speed does not effect

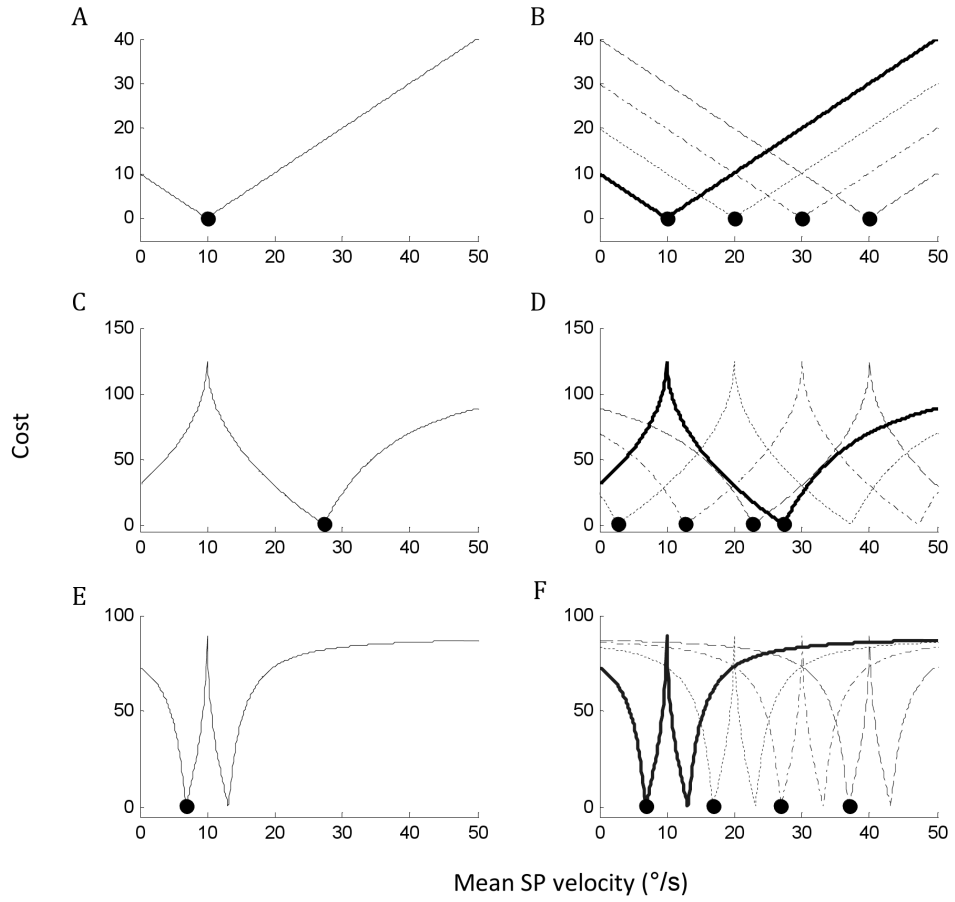

Figure 8.3: Illustration of the cost functions (A and B) J_{RS1} (retinal slip), (C and D) J_{RS2} (negative contrast sensitivity for low spatial frequency bars), and (E and F) J_{RS3} (negative contrast sensitivity for high spatial frequency content) as a function of mean SP velocity, for $\alpha = 1$. For clarity, left panels illustrate the cost function at stimulus speed 10°/s, and right panels illustrate the cost function at all stimulus speeds. Bold line, 10°/s; dotted line, 20°/s; dash-dot line, 30°/s; dashed line, 40°/s. Circles illustrate minimum cost. Where the minimum cost occurs at two different values of mean SP velocity we mark the function at the lower value of SP velocity for clarity.

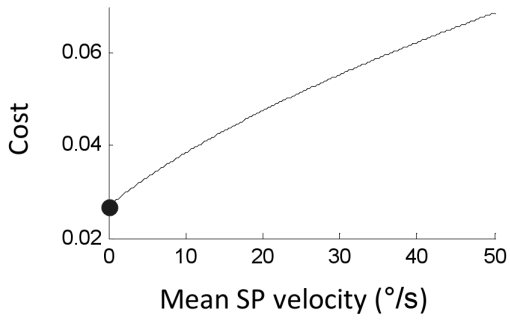

Figure 8.4: Illustration of the cost function J_{MS} as a function of mean SP velocity for $\beta = 1$. Stimulus speed has no effect on QP duration independent of changes in SP velocity, so all stimulus speeds produce the same cost function and the same minimum cost. Circle illustrates minimum cost.

QP duration independently of changes in SP velocity, all stimulus speeds produce the same cost as a function of SP velocity. The mean SP velocity that minimises this cost function is zero. In the data we find that mean SP velocity is always above zero, so the observed mean SP velocity is not optimal for minimising J_{MS}.

The results of these analyses are not surprising, but are useful in illustrating that the observed mean SP velocity is not optimal for minimising these plausible cost functions. We also investigate the possibility that SP velocity is optimal for minimising a combination of these cost functions. Figure 8.5 illustrates the combination of the cost functions $J_{MS} + J_{RS1}$ for stimulus speeds of $10°/s$ and $40°/s$, and for different values of β. Across the range of values tested we find that for each stimulus speed there are two possible optimal values of SP velocity. Both of these values represent the optimal value of SP velocity for minimising each individual cost function without considering the other. It appears that the optimal value of SP velocity in this case will be determined by the relative weight of each cost function, such that if β is above a critical value then the optimal value will be determined solely by J_{MS} and if β is below that critical value then the optimal value will be determined solely by J_{RS1}. As we have already noted

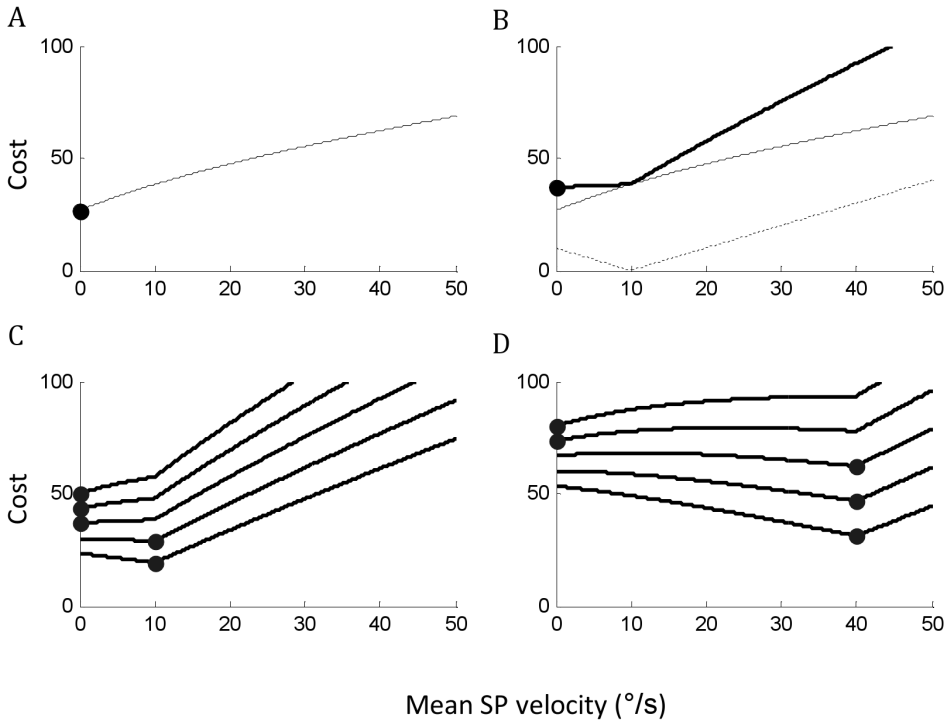

Figure 8.5: Illustration of the cost functions (A) J_{MS} for $\alpha = 1000$ and (B) $J_{MS} + J_{RS1}$, J_{MS}, and J_{RS1} (from top to bottom) for $\alpha = 1$ and $\beta = 1000$. (C) $J_{MS} + J_{RS1}$ for $\alpha = 1$ and β =500, 750, 1000, 1250 and 1500 (from bottom to top) for stimulus speed 10°/s. (D) $J_{MS} + J_{RS1}$ for $\alpha = 1$ and β =500, 750, 1000, 1250 and 1500 (from bottom to top) for stimulus speed 40°/s. Bold lines, $J_{MS} + J_{RS1}$; solid lines, J_{MS}; dotted line, J_{RS1}; Circles, minimum cost.

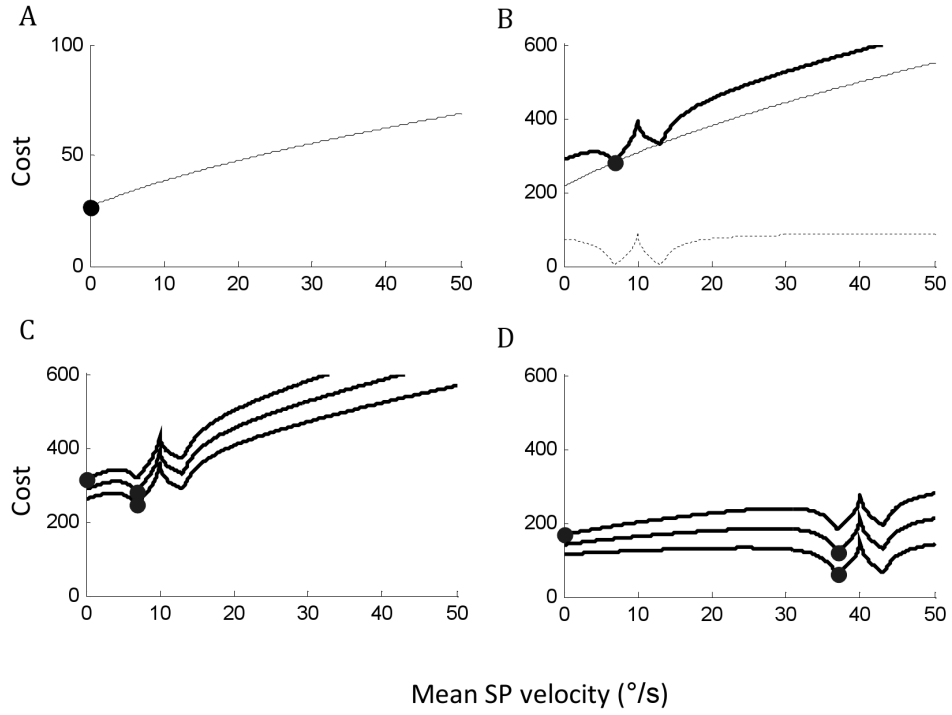

Figure 8.6: Illustration of the cost functions (A) J_{MS} for $\alpha = 1000$ and (B) $J_{MS} + J_{RS3}$, J_{MS}, and J_{RS1} (from top to bottom) for $\alpha = 1$ and $\beta = 2000$. (C) $J_{MS} + J_{RS1}$ for $\alpha = 1$ and β =1000, 2000 and 3000 (from bottom to top) for stimulus speed 10°/s. (D) $J_{MS} + J_{RS1}$ for $\alpha = 1$ and β =7000, 8000 and 9000 (from bottom to top) for stimulus speed 40°/s. Bold lines, $J_{MS} + J_{RS1}$; solid lines, J_{MS}; dotted line, J_{RS1}; Circles, minimum cost.

that the observed mean SP velocity is not optimal for minimising either of these cost functions then it cannot be optimal for minimising the sum. We find the same result when we attempt to minimise the combination of the cost functions $J_{MS} + J_{RS3}$ (fig. 8.6)

The cost functions for J_{TPE} are illustrated in fig. 8.7 for stimulus speeds of 10, 20, 30 and 40°/s. Here, we find some interesting results. The optimal values of mean SP velocity for minimising this function are 8.9, 15.4, 20.9 and 25.7°/s, which not only indicate an increasing retinal slip, but a decreasing gain, with stimulus speed. The values of gain are 0.89, 0.77, 0.70 and 0.64 for stimulus speeds of 10,

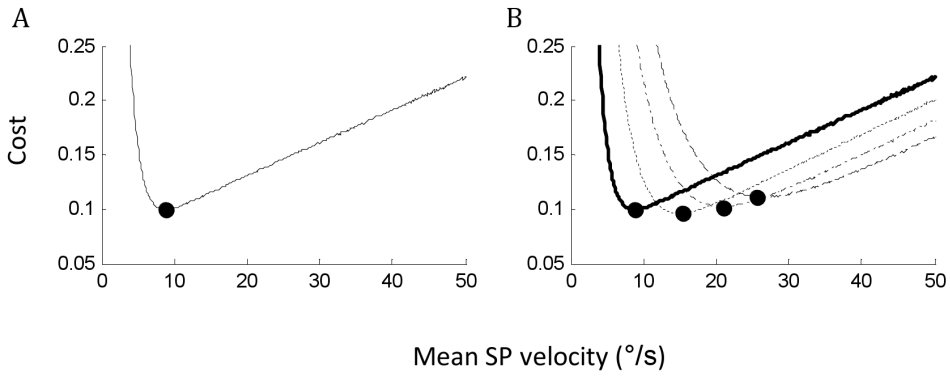

Figure 8.7: Cost function (A and B) J_{TPE} (tracking a moving target) for $\gamma = 1$. For clarity (A) illustrates the cost function at stimulus speed of 10°/s, and (B) illustrates the cost function at all stimulus speeds. Bold line, 10°/s; dotted line, 20°/s; dash-dot line, 30°/s; dashed line, 40°/s. Circles illustrate minimum cost.

20, 30 and 40°/s respectively. This pattern is very similar to that observed in the data (0.86, 0.72, 0.62 and 0.52). Although the values of gain are slightly higher than observed, this cost model does much better at predicting the observed SP velocity, and the corresponding increase in retinal slip with increasing stimulus speed.

If we examine the shape of the cost functions we can see that when SP velocity is low they are dominated by a hyperbolic component which appears to result from an increased SP duration, that will tend to ∞ as SP velocity $\rightarrow 0$. As SP velocity increases the cost function is dominated by a gradual increase which appears to result from an increase in targetting error. This increase in targetting error is due to a combination of SP velocity becoming faster than stimulus velocity and the targetting error on QPs increasing due to SP velocity increasing the magnitude of QPs. The minimum cost falls between these two extremes in a basin that becomes more shallow with increasing stimulus speed. This is an interesting result in itself, as if the OKN system was attempting to "search" for an unreferenced optimum SP velocity it may wander further along a more shallow basin than a more steep

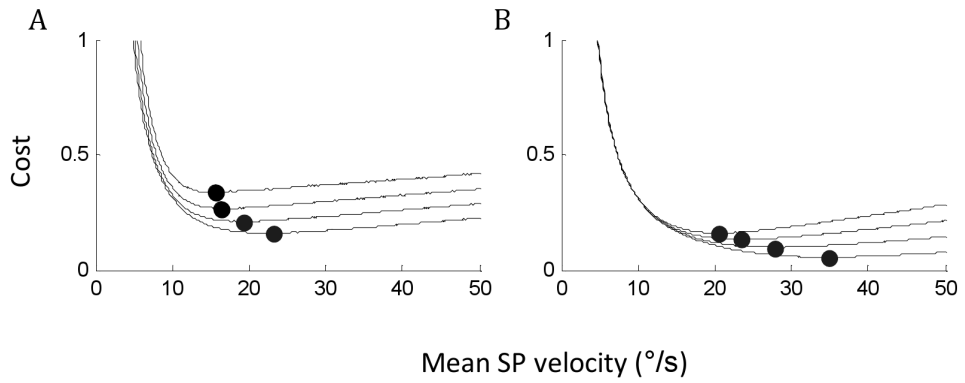

Figure 8.8: Cost function J_{TPE} (tracking a moving target) for $\gamma = 1$ and stimulus speed 40°/s. (A) Adding contant noise on QP amplitude with a standard deviation of 1, 2, 3 and 4°/s (represented by lines from bottom to top) and multiplicative noise with a standard deviation of $0.23|Q_i|$. (B) Adding multiplicative noise on QP amplitude with a standard deviation of 0.1, 0.2, 0.3 and $0.4|Q_i|$ (represented by lines from bottom to top), and no constant noise. Circles illustrate minimum cost.

basin. This would result in a more variable SP velocity for higher stimulus speeds.

We have also examined the effect that adding and increasing constant noise to the QP amplitude has on the cost function (fig. 8.8). It appears that increasing the constant noise component changes the shape of the cost function such that the optimum value of SP velocity decreases, as does increasing the multiplicative noise component, but with some subtle differences as seen in the figure.

8.5 Discussion of results.

These results indicate that the observed SP velocity is not optimal for minimising retinal slip or maximising visual contrast. The observed retinal slip increases with stimulus speed but the optimal retinal slip for each of these cost functions always remains the same. The results also demonstrate that SP velocity is not optimal for minimising the duration of QPs, as this cost model always predicts a SP velocity of zero in order to make QPs as small as possible (and thus shorter in duration

via the main sequence). Even when these cost functions are combined they are unable to explain the low values of SP gain observed in terms of a minimising retinal slip or maximising visual contrast.

The most interesting results are from the cost function that models OKN as optimising SP velocity to minimise the TPE in tracking a moving target over the course of a SP. It is capable of predicting optimum values of SP velocity that are very near the observed mean SP velocity found over a range of stimulus speeds, from experiment 1. We have also found that increasing the constant and multiplicative noise components in the QP amplitude causes the shape of the cost function to change, such that the optimum SP velocity decreases. This raises the interesting possibility that some participants may have a lower gain than others due to an increased amount of noise in the QP amplitude.

Chapter 9

Summary and conclusion.

The primary aim of the series of research presented in this thesis was to demonstrate the remarkable variability of individual OKN parameters, and to find an underlying structure in the relationships between OKN parameters that could be used to develop a mathematical model to describe how and why OKN behaves as it does.

9.1 Summary of findings.

In chapters 4 to 8 I have presented a range of results, both from analysis of eye movement data recorded from participants during OKN stimulation, and from extensive investigation of proposed models of the OKN system.

The results of RM-ANOVA performed on eye movement data during OKN stimulation illustrate that the most obvious effects on OKN parameters were from stimulus speed, which had a significant main effect on almost all parameters. The only parameters tested that were not affected were the standard deviation of SP gain, and the mean values of the SP start and end positions.

An interesting finding was that the spatial frequency and the type of stimulus pattern used also had a main effect on the mean SP velocity. Increasing the spatial frequency of the stimulus, or using a random grid pattern rather than a square wave pattern, both resulted in an increase in SP velocity such that gain

was increased and retinal slip was reduced. A related result was found in a study by Wester et al. (2007) where the maximum stimulus velocity at which an OKN response could be elicited in visually impaired patients was negatively related to the log spatial frequency of the stimulus. Also, Watanabe et al. (1994) observed differences in the OKN response using two different types of stimulus patterns (a wide interval stripe and a random dot pattern), but in their study no attempt was made to control patterns for their spatial frequency content.

One possibility for the dependence of SP velocity on the spatial frequency of the stimulus is that the peak contrast sensitivity for low spatial frequency gratings occurs at increased levels of retinal slip (Burr & Ross 1982). Therefore, when viewing an OKN stimulus with a low spatial frequency grating, increased levels of retinal slip would be allowed without adversely effecting the contrast of the image. The effect of changing the stimulus pattern is more difficult to explain, as the spatial frequency content of the image constantly changes but the average spatial frequency content remains the same. However, it is an important consideration for OKN in the natural environment, where spatial frequency content is likely to change during locomotion. I believe that in the case that low spatial frequency content of the visual stimulus is interspersed with high spatial frequency content, it is better to maintain minimal retinal slip in order to preferentially maximise the resolution of the high spatial frequency content.

The most unexpected result found in the data was that SP velocity randomly fluctuated between SPs to such a remarkable degree that during any given trial the SP gain might reach values much higher or much lower than the mean gain. The degree of the variability in SP velocity over long periods is rarely reported, and the study by Kolarik et al. (2009) is the only study that I am aware of that has investigated the variability of SP velocity over such a timescale. In our

investigation we have found a significant main effect of stimulus speed on the standard deviation of SP velocity and retinal slip, in agreement with the results of Kolarik et al. (2009) who illustrated that variation in SP velocity between SPs was a linear function of stimulus speed. One possible explanation for this is that a variable gain in the neural pathway that generates OKN (e.g. as part of a forward or feedback loop) could lead to multiplicative noise that is dependent on the stimulus speed. A variable gain element shared by both the VOR and the OKR has been proposed to explain changes in steady state OKN gain after long term VOR adaptation in cats (Demer 1981), and VOR pathology in humans (Baloh & Demer 1993), although this variable gain element is thought to adapt over long periods of time rather than fluctuate from cycle to cycle. Another possibility results from our proposed cost model that the OKN system is searching for an unreferenced optimum value of SP velocity from cycle to cycle, and at higher stimulus speeds there is a wider basin to search resulting in a wider range of SP velocity.

The results of PCA performed on OKN parameters illustrates how the eigenvectors and eigenvalues of the underlying principal components contribute to the correlation of OKN parameters within and across adjacent OKN cycles, and indicate that the remarkable variability in OKN parameters observed in the different correlation matrixes of OKN parameters between trials is the result of varying eigenvalues and not varying eigenvectors. Eigenvectors, and thus the coefficients of the linear relationships between parameters, remain predominantly similar across cycles and even across participants and stimulus conditions.

With these results and the results of linear regression I have developed a linear and stochastic model of OKN with three uncorrelated sources of noise affecting the SP velocity, the threshold for triggering a QP, and the amplitude of the returning

QP during each OKN cycle. The SP velocity was determined by the SP velocity in the cycle before, and not on the SP velocity during any previous cycles, i.e. it was a 1st order Markov process. This explained the variability in SP velocity from SP to SP, as SP velocity wanders in a Markov fashion often far above or below the mean gain. However, this behaviour is difficult to reconcile with the common model of OKN as a deep closed-loop feedback servomechanism. The threshold for triggering a QP was determined by the start position of the SP and the SP velocity during that cycle. These results indicate that a position threshold is not enough to explain when QPs are triggered and that there is also a velocity component to the trigger, as similarly observed in vestibular nystagmus (Lau et al. 1978) and as indicated by the correlation between SP amplitude and SP velocity observed by Watanabe et al. (1994). The amplitude of the returning QP was determined by a linear relationship with the start position of the QP, as also observed in turtles (Balaban & Ariel 1992), and the SP velocity during that cycle.

Solving the equations of the system in order to determine the autonomous update dynamics for each parameter demonstrated that all the OKN parameters analysed depended on their value only in the previous cycle, and also on the SP velocity in the previous cycle. The OKN parameters were all correlated with each other through a common dependence on SP velocity and the three noise sources of the S-, Q-, and V-components. This has important consequences for the behaviour of OKN, as it indicates that SP velocity does not only contribute to the compensatory SPs of OKN, but also the QP component.

The investigation into how QPs were affected by SP velocity illustrated that saccadic blind time was a function of SP velocity, and that there was a position towards which QPs were targetted that was also dependent on SP velocity. QPs to this target appear to consistently undershoot with a saccadic gain as low as

0.38. While there is a known undershoot bias for saccades to visual targets, the undershoot is considered to be approximately 10% of the target distance and does not explain why QPs would undershoot their target to such a degree. One possible explanation is that the visual system is not able to fully compensate for the movement made away from the target location during short latency SPs, as has also been observed when memory guided saccades were made during a period of vestibular rotation (Daye et al. 2010). However, it has also been proposed that the saccadic undershoot bias minimises saccadic flight time, as smaller magnitude saccades would have a smaller end point variance due to signal dependent noise in the motor command and would thus require fewer corrective saccades in total, and a greater amount of signal dependent noise would lead to this optimal undershoot bias being even greater (Harris 1994). In this case low saccadic gain occurs due to a speed-accuracy tradeoff, and the undershoot is even larger for QPs than visually guided saccades possibly due to an even greater amount of signal dependent noise in the motor command for QPs. Neither one of these possibilities excludes the other, so it is possible that both a lack of compensation for the SP movement and a saccadic undershoot bias are in effect.

Comparing the magnitude of SP and QP amplitudes with the standard deviation of the S and Q-components revealed that the noise in both components was linearly proportional to the magnitude of both movements i.e. signal dependent noise was present in QPs and SPs. The presence of signal dependent noise in the Q-component was not unexpected as it is well-known that larger amplitude saccade to static visual targets have larger errors (van Beers 2007). However, it was surprising that the constant of proportionality was so high. Examining the residuals of the S, Q, and V-components revealed that the noise processes were not all strictly Gaussian, but had a tendency to be slightly leptokurtic. It is not

entirely clear what mechanism generates this noise but the excess kurtosis found in the residuals may be a product of the signal-dependent noise process.

An investigation of the transient response of OKN parameters predicted by the model illustrated that both SP velocity and the start position of SPs have a fast transient response, taking just over 2 cycles to reach 63% of the simulated steady state value, a result that has been illustrated empirically by Abadi et al. (2005).

The results of fitting SP duration histograms with a range of PDFs found that two possible PDFs predicted by the Markov model (the FRD and DRTN) gave very good fits to the data. Not only that, but they could give significantly better fits to the data than any other distribution tested when all histograms were combined, and were not significantly different to each other. The RND gave significantly worse fits to the data than all other PDFs tested, even to individual histograms in agreement with a similar study performed by Trillenberg et al. (2002), and I do not recommend using the RND to fit the histograms of OKN SP duration in any future models.

Crucially, the significantly good fit of the two ratio distributions to the histograms of SP duration is in good agreement with the hypothesis that SP duration is determined by the ratio of a variable amplitude threshold and variable SP velocity. This in stark contrast to most studies on OKN that have modelled or assumed that SP duration is a fundamental variable of OKN and that it is determined by some interval generator process such as the accumulator models tested in this thesis (Balaban & Ariel 1992, Carpenter 1993, 1994, Anastasio 1996). While these accumulator models have been successful in describing the distribution of voluntary saccade latency to stationary targets (Carpenter 1999, Smith & Ratcliff 2004) during OKN the eyes are primarily in a state of continuous motion in the direction of optic flow and as QPs apparently act to redirect the eyes to a more

central location it seems likely that the threshold for triggering a QP is dependent on the position of the eyes in the orbit, as described by the Markov model.

The results of our proposed cost model of OKN have been very interesting, although it is a preliminary model and still requires more thorough testing. SP velocity does not appear to be optimal for minimising retinal slip or QP duration, nor does it appear to be optimal for maximising contrast sensitivity. The results of the cost analysis indicate that SP velocity is optimal for minimising the positional error in tracking a moving target integrated over the course of each SP, indicating that foveal tracking plays a vital role in OKN and that positional error may play a more important part than previously considered in determining SP velocity.

Miles (1998) has noted a functional distinction between the rotational and translational VOR, as a single rotational eye movement cannot stabilise the entire visual field on the retina during translational motion whereas a single rotation of the eyes can stabilise the entire visual field during rotational motion. As stabilisation of the entire visual field is not possible, it has been proposed that the translational VOR preferentially stabilises images on the fovea and the amplitude of translational VOR eye movements depends on the gaze direction (Tomko & Paige 1992, Angelaki & Hess 2001), which appears to agree with this hypothesis. It has also been previously observed that the 3D kinematic properties of the OFR are similar to those of smooth pursuit and the translational VOR, which exhibit an eye-position dependent torsional component as predicted by Listing's law (Adeyemo & Angelaki 2005). The kinematics of rotational and translational OKN also appear to differ in the same manner as the kinematics of rotational and translational VOR (Tian et al. 2007) so it seems plausible that translational OKN, much like the translational VOR, preferentially stabilises images on the fovea as

the entire visual field cannot be stabilised during pure translational motion.

It is important to consider that positional error is of significant importance in primates due to the existence of the fovea, which is capable of resolving much higher spatial frequency images than the peripheral retina. In animals with less retinal specialisation, the significance of positional error becomes much less profound as all areas of the retina possess similar visual acuity. Indeed, it is possible that the differences in the distribution of retinal cells may account for some of the differences in the gain and time course of OKN in different species.

While the classical systems engineering approach to investigating OKN often ignores the QP component of OKN and considers only the neural correlates of the velocity signal to be important in determining SP gain, the results of this series of investigations indicates that accounting for positional error may be crucial in effectively modelling SP gain in humans.

9.2 Future research.

One concern of the Markov model is that the PCA methods used to develop it were only designed to find components that are linear combinations of OKN parameters. However, nonlinear generalisations of PCA are available (e.g. Gorban et al. (2007)), and could be applied in a similar manner to the methods used in this thesis in order to determine if the "principal curves" are the same as the principal components. Nonlinear extensions to the model might allow the relationships between OKN parameters to be investigated in more detail, such as the proposed hyperbolic relationship between SP velocity and duration, or to explain complications of the model.

One complication does occur with the model when SP velocity reaches (or begins) at a value of $0°/s$, and causes the SP duration to become infinite. In the Markov

model, SP velocity updates only after each QP, but if the value of SP velocity reaches 0°/s an infinitely long duration SP would occur and SP velocity would be unable to update. There are three possible solutions to this complication. One possibility is that another system, outside the scope of this model, generates a new QP in order to allow the SP velocity to update. This is not entirely without evidence, as Cheng & Outerbridge (1974) found that histograms of SP duration can range in shape from unimodal to multimodal form, although we did not observe this in our data. A second possibility comes from considering possible nonlinear SPs, where SP velocity is allowed to update in time and not just between SPs. This is certainly very likely at the very start of OKN stimulation where it appears the ocular following response begins within 80ms, without first generating a QP. A third possibility is that SPs never reach a SP velocity of 0°/s. This is the constraint that is currently implemented in the model (eq. 5.8), but it is unsatisfactory as extremely long SPs can still occur. An investigation of the nonlinearity of SPs would clarify whether SP velocity updates in real time, between cycles, rather than simply settling to a new SP velocity after every QP. While SPs are classically considered linear, variation in eye velocity during SPs has been reported (Kolarik et al. 2009).

It is also necessary to consider the type of OKN stimulus used. In the natural environment visual scenes are complicated, and do not have constant spatial frequency content, unlike the alternating black and white bars of the typical OKN stimuli used in these experiments. A necessary step in verifying the Markov model would be to compare how well it predicts the behaviour of OKN when different patterns of OKN stimuli are used and different instructions given to participants. A comparison of the values of the constants a, b, c and d from the model, when applied to OKN parameters recorded while participants are viewing different types

of stimulus patterns would further verify the results presented in this thesis. Of particular interest to this author is whether a pure retinal slip signal, using limited lifetime dot stimuli, has an effect on the Markov model parameters or predictive power of the cost model.

The approach that I have taken in this thesis could be applied to other forms of physiological nystagmus such as the vestibular nystagmus of the VOR, or even pathological forms of nystagmus such as congenital or acquired nystagmus. The cause of congenital nystagmus is still not currently known, and an understanding of the threshold at which QPs are triggered, and to where they are targetted, could feasibly help us understand more about the mechanism that causes this condition. However, the SPs of pathological nystagmus are notably non-linear, specifically they accelerate in congenital nystamgus and decelerate in manifest latent nystagmus (Abadi & Bierre 2002). It would be necessary to use non-linear methods to analyse these SPs adequately, or at the very least to transform the data into parameters where linear methods might be effective.

9.3 Conclusion.

In this thesis I have examined the loading patterns of components extracted from OKN parameters and found that only parameters from within a single cycle of OKN load highly on to any given component, with no evidence for long term correlations over multiple cycles. I have developed a model that describes OKN as a purely stochastic process with three sources of noise affecting SP velocity, QP triggering, and QP amplitude: a triple 1^{st} order Markov process. The random fluctuations in SP velocity as it wanders in Markov fashion from SP to SP is difficult to reconcile with the common assumption of OKN as a simple servomechanism. I propose that SP velocity may minimise the positional error over the course of the

SP in order to stabilise the visual field on the retina preferentially at the fovea, as the entire visual field cannot be fully stabilised during pure translational motion. As such, future models of OKN should consider the positional components of the system as well as the usual velocity components, in order to adequately predict SP gain.

Appendix A

Heuristics for sorting components.

Here I give the heuristics for sorting the principal components in to their respective categories after extraction using PCA, based on the loading patterns of the components rather than their eigenvalue. I have included the heuristics for all categories that have been illustrated in this thesis. Heuristics are given in pseudocode format and parameters are labelled as they were illustrated in figures 5.2 - 5.6, and 5.9.

A.1 Heuristics for sorting the three principal components from one cycle of OKN parameters, not including SP velocity and SP duration.

FOR components 1 to 3

IF the absolute value of Y_i is greater than 0.4 THEN

Category = "X"

IF Yi is greater than 0 THEN

Loadings = -Loadings

END IF

ELSE IF the absolute value of Q_i+X_{i+1} is greater than 0.4 THEN

Category = "Q"

IF Q_i is greater than 0 THEN

Loadings = -Loadings

END IF

ELSE IF the absolute value of S_i is greater than 0.4 THEN

Category = "S"

IF S_i is less than 0 THEN

Loadings = -Loadings

END IF

END IF

END FOR

A.2 Heuristics for sorting the four principal components from one cycle of OKN parameters, including SP velocity and SP duration.

FOR components 1 to 4

IF the absolute value of Y_i is greater than 0.4 THEN

Category = "X"

IF Yi is greater than 0 THEN

Loadings = -Loadings

END IF

ELSE IF the absolute value of Q_i+X_{i+1} is greater than 0.4 THEN

Category = "Q"

IF Q_i is greater than 0 THEN

Loadings = -Loadings

END IF

ELSE IF the absolute value of S_i is greater than 0.4 THEN

Category = "S"

IF S_i is less than 0 THEN

Loadings = -Loadings

END IF

ELSE IF the absolute value of V_i is greater than 0.4 THEN

Category = "V"

IF V_i is less than 0 THEN

Loadings = -Loadings

END IF

END IF

END FOR

A.3 Heuristics for sorting the thirteen principal components from four cycles of OKN parameters, including SP velocity and SP duration.

FOR components 1 to 13

IF the absolute value of Y1-Q1 is greater than 0.9 THEN

Category = "1"

IF Q1 is less than 0 THEN

Loadings = -Loadings

END IF

ELSE IF the absolute value of Q1+X2+Y2-Q2 is greater than 1 THEN

Category = "2"

IF Q2 is less than 0 THEN

Loadings = -Loadings

END IF

ELSE IF the absolute value of Q2+X3 is greater than 0.8 AND

the absolute value of Q2-Y2 is less than 1.2 THEN

Category = “3”

IF Q2 is greater than 0 THEN

Loadings = -Loadings

END IF

ELSE IF the absolute value of Q3+X4 is greater than 0.7 THEN

Category = “4”

IF Q3 is greater than 0 THEN

Loadings = -Loadings

END IF

ELSE IF the absolute value of Q4+X5 is greater than 0.8 THEN

Category = “5”

IF Q4 is greater than 0 THEN

Loadings = -Loadings

END IF

ELSE IF the absolute value of S1+V1+T1> is greater than 0.8 THEN

Category = “6”

IF S1 is less than 0 THEN

Loadings = -Loadings

END IF

ELSE IF the absolute value of S2+V2+T2+Y2-X2 is greater than 1 AND

the absolute value of Q2-Y2 is less than 0.8 THEN

Category = “7”

IF S2 is less than 0 THEN

Loadings = -Loadings

END IF

ELSE IF the absolute value of S3+V3+T3+Y3 is greater than 0.8 THEN

Category = "8"

IF S3 is less than 0 THEN

Loadings = -Loadings

END IF

ELSE IF the absolute value of S4+V4+T4+Y4 is greater than 0.8 THEN

Category = "9"

IF S4 is less than 0 THEN

Loadings = -Loadings

END IF

ELSE IF the absolute value of V1-T1 is greater than 1 THEN

Category = "10"

IF V1 is less than 0 THEN

Loadings = -Loadings

END IF

ELSE IF the absolute value of V2-T2 is greater than 0.8 THEN

Category = "11"

IF V2 is less than 0 THEN

Loadings = -Loadings

END IF

ELSE IF the absolute value of V3-T3 is greater than 0.8 THEN

Category = "12"

IF V3 is less than 0 THEN

Loadings = -Loadings

END IF

ELSE IF the absolute value of V4-T4 is greater than 0.8 THEN

Category = "13"

IF V4 is less than 0 THEN

Loadings = -Loadings

END IF

END IF

END FOR

A.4 Heuristics for sorting the ten principal components with the largest eigenvalues from four cycles of OKN parameters, including SP velocity and SP duration.

FOR components 1 to 10

IF the absolute value of Y1-Q1 is greater than 0.7 THEN

Category = "1"

IF Q1 is less than 0 THEN

Loadings = -Loadings

END IF

ELSE IF the absolute value of Q1+X2+Y2-Q2 is greater than 0.7 THEN

Category = "2"

IF X2 is greater than 0 THEN

Loadings = -Loadings

END IF

ELSE IF the absolute value of Q2+X3 is greater than 0.6 AND

the absolute value of Q2-Y2 is less than 0.9 THEN

Category = “3”

IF Q2 is greater than 0 THEN

Loadings = -Loadings

END IF

ELSE IF the absolute value of Q3+X4 is greater than 0.6 THEN

Category = “4”

IF Q3 is greater than 0 THEN

Loadings = -Loadings

END IF

ELSE IF the absolute value of Q4+X5 is greater than 0.6 THEN

Category = “5”

IF Q4 is greater than 0 THEN

Loadings = -Loadings

END IF

ELSE IF the absolute value of S1+T1> is greater than 0.8 THEN

Category = “6”

IF S1 is less than 0 THEN

Loadings = -Loadings

END IF

ELSE IF the absolute value of S2+T2+Y2-X2 is greater than 1 AND

the absolute value of Q2-Y2 is less than 0.8 THEN

Category = “7”

IF S2 is less than 0 THEN

Loadings = -Loadings

END IF

ELSE IF the absolute value of S3+T3+Y3 is greater than 0.8 THEN

Category = “8”

IF S3 is less than 0 THEN

Loadings = -Loadings

END IF

ELSE IF the absolute value of S4+T4+Y4 is greater than 0.8 THEN

Category = “9”

IF S4 is less than 0 THEN

Loadings = -Loadings

END IF

ELSE Category = “10”

IF V1+V2+V3+V4 is less than 0 THEN

Loadings = -Loadings

END IF

END IF

END FOR

Appendix B

Table of free variables from experiment one.

In Chapter 5, I presented the results of analysis performed on Monte Carlo simulated data that had been created using the constants $a = -0.25$, $b = 0.158$, $c = -0.478$, and $d = -0.166$, and the free variables e, $\widehat{s}$, $\widehat{q}$, $\widehat{v}$, σ_s, σ_q and σ_v. Here I present the values of the free variables used to create the simulated data. The variables e, $\widehat{v}$, σ_s, σ_q and σ_v were found using a weighted least squares linear regression model. The variables $\widehat{s}$ and $\widehat{q}$ were calculated by substituting the mean values of the relevant OKN parameters for that trial and the constants a, b, c and d into eq. 5.4 and 5.5 respectively.

Par	V_S	e	$\widehat{s}$	$\widehat{q}$	$\widehat{v}$	σ_s	σ_q	σ_v
1	10	0.41	-0.32	-2.84	5.35	0.95	2.39	1.10
1	20	0.61	0.03	-2.88	6.60	1.29	2.86	2.30
1	30	0.85	-0.11	-3.05	2.17	1.58	2.33	3.48
1	40	0.81	-0.34	-2.80	2.58	1.31	2.55	4.61
2	10	0.27	2.52	-3.40	6.81	2.08	2.64	1.28
2	20	0.51	2.12	-3.11	6.15	2.65	3.65	4.95
2	30	0.47	2.61	-2.17	6.86	2.95	4.45	5.31
2	40	0.32	4.33	-4.05	8.47	3.76	4.51	4.10

3	10	0.33	0.81	-2.46	4.25	1.21	2.14	1.97
3	20	0.65	3.38	-0.24	4.63	2.10	2.12	3.32
3	30	0.66	2.72	-1.08	7.25	2.29	2.32	4.91
3	40	0.53	2.69	-0.95	13.72	2.47	3.04	4.70
4	10	0.08	1.79	-0.76	9.22	1.42	1.15	0.81
4	20	0.27	1.50	-1.30	14.08	1.81	1.74	1.23
4	30	0.75	1.43	-0.58	6.32	2.05	2.03	3.73
4	40	0.55	0.73	-1.24	7.79	1.50	1.80	7.01
5	10	0.74	2.95	-0.97	2.30	1.76	1.96	2.07
5	20	0.86	0.89	-1.67	1.53	1.25	1.29	3.10
5	30	0.72	0.91	-1.18	5.97	1.55	1.26	2.85
5	40	0.69	1.99	-1.98	7.94	2.11	1.57	6.80
6	10	0.12	1.01	-3.30	7.88	0.82	1.22	0.86
6	20	0.43	1.56	-3.52	9.32	1.32	1.92	2.29
6	30	0.71	1.28	-4.33	4.59	1.70	1.76	3.59
6	40	0.59	1.82	-3.21	7.91	1.75	2.17	5.64
7	10	0.44	1.71	-2.21	5.18	1.59	2.91	1.13
7	20	0.55	-0.25	-6.26	6.31	1.53	2.50	2.46
7	30	0.65	3.10	-2.41	8.82	3.04	3.20	3.59
7	40	0.58	3.79	-2.04	13.66	3.45	3.70	4.11
8	10	0.42	-1.18	-4.34	4.75	0.71	0.93	1.01
8	20	0.90	-1.08	-3.91	1.00	0.78	1.13	2.07
8	30	0.78	-0.48	-3.60	2.67	0.93	1.28	2.78
8	40	0.86	-0.39	-3.52	1.46	0.93	1.21	2.47
9	10	0.37	1.11	-0.91	4.37	1.03	0.92	0.93
9	20	0.36	0.58	-1.43	6.95	1.37	1.86	2.49

9	30	0.29	0.41	-2.38	10.98	1.78	2.12	4.39
9	40	0.13	0.87	-1.82	11.02	1.44	2.05	6.33
10	10	0.09	-0.45	-5.61	8.38	1.22	1.92	1.00
10	20	0.24	0.00	-4.54	14.56	1.68	2.38	1.64
10	30	0.69	-0.66	-4.88	7.52	1.87	2.08	3.54
10	40	0.50	-1.07	-4.53	15.99	2.51	2.29	5.01

Appendix C

The PDF of the ratio of truncated (and untruncated) normal random variables.

The PDF of the ratio of two Normal random variables, x and y, is given by:

$$z = x/y \tag{C.1}$$

From basic probability theory, the PDF of z is given by:

$$f_z(z) = \int_{-\infty}^{\infty} |y| \, f_{xy}(yz, y) dy \tag{C.2}$$

where $f(x, y)$ is the bivariate Normal PDF. If x and y are left truncated at zero, then the PDF of z becomes:

$$f_z(z) = \int_{0}^{\infty} y f_{xy}(yz, y) dy \tag{C.3}$$

The bivariate Normal PDF $f(x, y)$ is given by the normalised two-dimensional

Gaussian, $G(x, y)$:

$$f(x, y) = \frac{1}{k} G(x, y) \tag{C.4}$$

where

$$G(x, y) = \exp\left(-\frac{1}{2(1-\rho^2)}\left[\frac{(x-\mu_x)^2}{\sigma_x^2} + \frac{(y-\mu_y)^2}{\sigma_y^2} - \frac{2\rho(x-\mu_x)(y-\mu_y)}{\sigma_x \sigma_y}\right]\right) \tag{C.5}$$

and k is a normalising constant such that the area under $G(x, y)$ is unity for the quadrants considered. For all quadrants:

$$k = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} G(x, y) dx dy = 2\pi \sigma_x \sigma_y \sqrt{1-\rho^2}. \tag{C.6}$$

For the ratio of truncated Normals

$$k_1 = \int_{0}^{\infty} \int_{0}^{\infty} G(x, y) dx dy \tag{C.7}$$

and is found numerically.

We can expand the Gaussian in eq. C.5 in the form:

$$G(zy, y) = \exp(ay^2 + by + c) \tag{C.8}$$

where

$$a = \frac{-1}{2(1-\rho^2)} \frac{z^2\sigma_y^2 - 2z\rho\sigma_x\sigma_y + \sigma_x^2}{\sigma_y^2\sigma_x^2} < 0 \tag{C.9}$$

$$b = \frac{1}{2(1-\rho^2)} \frac{2z\mu_x\sigma_y^2 + 2\mu_y\sigma_x^2 - 2z\rho\mu_x\sigma_x\sigma_y}{\sigma_y^2\sigma_x^2} \tag{C.10}$$

$$c = \frac{-1}{2(1-\rho^2)} \frac{\mu_x^2\sigma_y^2 - 2\rho\mu_x\mu_y\sigma_x\sigma_y + \mu_y^2\sigma_x^2}{\sigma_y^2\sigma_x^2} \tag{C.11}$$

Now consider the integral

$$f_z(z) = \frac{1}{k_1}\int_0^\infty yG(zy, y)dy = \frac{1}{k_1}\int_0^\infty y\exp(ay^2 + by + c)dy \tag{C.12}$$

With integration by parts:

$$f_z(z) = -\frac{\exp(c)}{2ak_1}\left(1 + b\int_0^\infty \exp(ay^2 + by)dy\right) \tag{C.13}$$

and substituting $u = y\sqrt{-a} - b/2\sqrt{-a}$, we have

$$f_z(z) = -\frac{\exp(c)}{2ak_1}\left[1 + b\sqrt{\frac{\pi}{-4a}}\exp(-b^2/4a)\mathrm{erfc}(-b/2\sqrt{-a})\right] \tag{C.14}$$

For all quadrants, we have

$$f_z(z) = \frac{1}{k}\int_{-\infty}^{\infty} |y| \exp(ay^2 + by + c)dy \tag{C.15}$$

$$= -\frac{\exp(c)}{2ak}\left[1 + b\sqrt{\pi} - 4a\exp(-b^2/4a)\mathrm{erfc}(b/2\sqrt{-a})\right] \tag{C.16}$$

Glossary.

CDF Cumulative distribution function

CV Coefficient of variation

DOF Degrees of freedom

DRTN Distribution of the ratio of two truncated normal variables

DTN Dorsal terminal nucleus of the accessory optic system

FRD Fieller's ratio distribution

GD Gamma distribution

IGD Inverse Gaussian distribution

LND Lognormal distribution

LTN lateral terminal nucleus of the accessory optic system

MLE Maximum likelihood estimate

MTN medial terminal nucleus of the accessory optic system

NOT Nucleus of the optic tract

OKN Optokinetic nystagmus

OKR Optokinetic response

PCA Principal component analysis

PDF Probability density function

QP Quick phase

RM-ANOVA Repeated measures analysis of variance

RND Recinormal distribution

RNMD Recinormal mixture distribution

SP Slow phase

TPE Total positional error

VOR Vestibulo-ocular reflex

List of references.

Abadi, R. & Bierre, A. (2002), 'Motor and sensory characteristics of infantile nystagmus', *Br J Ophthalmol* **86**, 1152–1160.

Abadi, R., Howard, I., Ohmi, M. & Lee, E. (2005), 'The effect of central and peripheral field stimulation on the rise time and gain of human optokinetic nystagmus', *Perception* **34**, 1015–1024.

Adeyemo, B. & Angelaki, D. E. (2005), 'Similar kinematic properties for ocular following and smooth pursuit eye movements', *J Neurophysiol* **93**, 1710–1717.

Anastasio, T. (1996), 'A random walk model of fast-phase timing during optokinetic nystagmus', *Biol Cybern* **75**, 1–9.

Angelaki, D. E. & Hess, B. J. M. (2001), 'Direction of heading and vestibular control of binocular eye movements', *Vision Res* **41**, 3215–3228.

Argyris, J., Andreadis, I., Pavlos, G. & Athanasiou, M. (1998), 'The influence of noise on the correlation dimension of chaotic attractors', *Chaos Solition Fract* **9**, 343–361.

Balaban, C. & Ariel, M. (1992), 'A "beat-to-beat" interval generator for optokinetic nystagmus', *Biol Cybern* **66**, 203–216.

Baloh, R. & Demer, J. (1993), 'Optokinetic-vestibular interaction in patients with increased gain of the vestibulo-ocular reflex', *Exp Brain Res* **97**, 334–342.

Barnes, G. (1993), 'Visual-vestibular interaction in the control of head and eye movement: the role of visual feedback and predictive mechanisms', *Prog Neurobiol* **41**, 435–472.

Barnes, G. & Marsden, J. (2002), 'Anticipatory control of hand and eye movements in humans during oculo-manual tracking', *J Physiol* **539**, 317–330.

Barratt, H. & Hood, J. (1988), 'Transfer of optokinetic activity to vestibular nystagmus', *Acta Oto-Laryngol* **105**, 318–327.

Blunch, N. (2008), *Introduction to structural equation modelling using SPSS and AMOS*, Sage Publications Ltd.

Box, G. & Jenkins, G. (1976), *Time series analysis: forecasting and control*, Holden-Day.

Brandt, T., Dichigans, J. & Koenig, E. (1973), 'Differential effects of central versus peripheral vision on egocentric and exocentric motion perception', *Exp Brain Res* **16**, 476–491.

Brown, P. & Wald, G. (1964), 'Visual pigments in single rods and cones of human retina', *Science* **144**, 45–52.

Burr, D. (1981), 'Temporal summation of moving images by the human visual system', *Proc R Soc Lond B Biol Sci* **211**, 321–339.

Burr, D. & Ross, J. (1982), 'Contrast sensitivity at high velocities', *Vision Res* **22**, 479–484.

Busettini, C., Masson, G. & Miles, F. (1996), 'A role for stereoscopic depth cues in the rapid visual stabilization of the eyes', *Nature* **380**, 342–345.

Campbell, F. & Robson, J. (1968), 'Application of Fourier analysis to the visibility of gratings', *J Physiol* **197**, 551–566.

Carpenter, R. (1993), 'Distribution of quick-phase intervals in optokinetic nystagmus', *Ophthalmic Res* **25**, 91–93.

Carpenter, R. (1994), Express optokinetic nystagmus, *in* 'Contemporary ocular motor and vestibular research', Georg Thieme.

Carpenter, R. (1999), 'Visual selection: neurons that make up their minds', *Curr Biol* **9**, 595–598.

Carpenter, R. & Williams, M. (1995), 'Neural computation of log likelihood in control of saccadic eye movements', *Nature* **377**, 59–62.

Cheng, M. & Outerbridge, J. (1974), 'Inter-saccadic interval analysis of optokinetic nystagmus', *Vision Res* **14**, 1053–1058.

Chun, K. & Robinson, D. (1978), 'A model of quick phase generation in the vestibuloocular reflex', *Biol Cybern* **28**, 209–221.

Cohen, B., Reisine, H., Yokota, J. & Raphan, T. (1992), 'The nucleus of the optic tract', *Ann NY Acad Sci* **656**, 277–296.

Collewijn, H. (1981), *The oculomotor system of the rabbit and its plasticity*, Springer-Verlag.

Collins, C. & Barnes, G. (2006), 'The occluded onset pursuit paradigm: prolonging anticipatory smooth pursuit in the absence of visual feedback', *Exp Brain Res* **175**, 11–20.

Comon, P. (1994), 'Independent component analysis, a new concept?', *Signal Processing* **36**, 287–314.

Costello, A. & Osborne, J. (2005), 'Best practices in exploratory factor analysis: Four recommendations for getting the most from your analysis', *Practical Assessment Research & Evaluation* **10**, 1–9.

Curtiss, J. (1941), 'On the distribution of the quotient of two chance variables', *Ann Math Stat* **12**, 409–421.

Das, V., Zivotofsky, A., DiScenna, A. & Leigh, R. (1995), 'Head perturbations during walking while viewing a head-fixed target', *Aviat Space Environ Med* **66**, 728–732.

Daye, P., Blohm, G. & Lefevre, P. (2010), 'Saccadic compensation for smooth eye and head movements during head-unrestrained two-dimensional tracking', *J Neurophysiol* **103**, 543–556.

Demer, J. (1981), 'The variable gain element of the vestbulo-ocular reflex is common to the optokinetic system of the cat', *Brain Res* **229**, 1–13.

Dodge, R. (1903), 'Five types of eye movements in the horizontal meridian plane of the field of regard', *Am J Physiol Lond* **8**, 307–329.

Dubois, M. & Hollewijn, H. (1979), 'Optokinetic reactions in man elicited by localised retinal motion stimuli', *Vision Res* **19**, 1105–1115.

Fetter, M., Misslisch, H., Sievering, D. & Tweed, D. (1995), 'Effects of full-field visual input on the three-dimensional properties of the human vestibuloocular reflex', *J Vestib Res* **5**, 201–209.

Fieller, E. (1932), 'The distribution of the index in a normal bivariate population', *Biometrika* **24**, 428–440.

Fuchs, A., Brandt, T., Buttner, U. & Zee, D., eds (1993), *Contemporary ocular motor and vestibular research: a tribute to David A. Robinson*, Thieme.

Garbutt, S., Han, Y., Kumar, A., Harwood, M., Harris, C. & Leigh, R. (2003), 'Vertical optokinetic nystagmus and saccades in normal human subjects', *Invest Ophth Vis Sci* **44**, 3833–3841.

Garbutt, S., Harwood, M. & Harris, C. (2001), 'Comparison of the main sequence of reflexive saccades and the quick phases of optokinetic nystagmus', *Brit J Ophthalmol* **85**, 1477–1483.

Garbutt, S., Harwood, M. & Harris, C. (2002), 'Anticompensatory eye position ("contraversion") in optokinetic nystagmus', *Ann NY Acad Sci* **956**, 445–448.

Giolli, R. (1963), 'An experimental study of the accessory optic system in the cynomologous monkey', *J Comp Neurol* **121**, 89–107.

Gorban, A., Kegi, B., Wunsch, D. & Zinovyev, A., eds (2007), *Principal manifolds for data visualisation and dimension reduction*, Springer.

Harris, C. (1995), 'Does saccadic undershoot minimize saccadic flight-time? a Monte-Carlo study', *Vision Res* **35**, 691–701.

Harris, C. & Berry, D. (2006), 'A distal model of congenital nystagmus as nonlinear adaptive oscillations', *Nonlinear Dynam* **44**, 367–380.

Harris, C. M. (1994), 'Does saccadic undershoot minimize saccadic flight-time? a monte-carlo study', *Vision Res* **35**, 691–701.

Harris, C. & Wolpert, D. (1998), 'Signal-dependent noise determines motor planning', *Nature* **394**, 780–784.

Harris, C. & Wolpert, D. (2006), 'The main sequence of saccades optimizes speed-accuracy trade-off', *Biol Cybern* **95**, 21–29.

Harwood, M., Madelain, L., Krauzlis, R. & Wallman, J. (2008), 'The spatial scale of attention strongly modulates saccade latencies', *J Neurophysiol* **99**, 1743–1757.

Henn, V., Young, L. & Finley, C. (1974), 'Vestibular nucleus units in alert monkeys are also influenced by moving visual fields', *Brain Res* **71**, 144–149.

Higgins, J. & Tashtoush, S. (1994), 'An aligned rank transform test for interaction', *Nonlinear World* **1**, 201–211.

Hinkley, D. (1969), 'On the ratio of two correlated normal random variables', *Biometrika* **56**, 635–639.

Hoffmann, K. (1988), 'Responses of single neurons in the pretectum of monkeys to visual stimuli in three-dimensional space', *Ann NY Acad Sci* **545**, 180–186.

Honrubia, V., Downey, W., Mitchell, D. & Ward, P. (1968), 'Experimental studies on optokinetic nystagmus', *Acta Oto-Laryngol* **65**, 441–448.

Hood, J. (1967), 'Observations upon the neurological mechanism of optokinetic nystagmus with especial reference to the contribution of peripheral vision', *Acta Oto-Laryngol* **65**, 208–215.

Huang, Y. & Neuhauss, S. (2008), 'The optokinetic response in zebrafish and its applications', *Front Biosci* **13**, 1899–1916.

Ilg, U. & Hoffmann, K. (1996), 'Responses of neurons of the nucleus of the optic tract and the dorsal terminal nucleus of the accessory optic tract in the awake monkey', *Eur J Neurosci* **8**, 92–105.

Kaminiarz, A., Konigs, K. & Bremmer, F. (2009), 'Task influences on the dynamic properties of fast eye movements', *J Vision* **9**, 1–11.

Kijima, M. (1997), *Markov processes for stochastic modeling*, Chapman & Hall.

Kolarik, A., Margrain, T. & Freeman, T. (2009), 'Precision and accuracy of ocular following: influence of age and type of eye movement', *Exp Brain Res* **201**, 271–282.

Lacouture, Y. & Cousineau, D. (2008), 'How to use MATLAB to fit the ex-Gaussian and other probability functions to a distribution of response times', *Tutorial Quant Meth Psychol* **4**, 35–45.

Land, M. (1999), 'Motion and vision: why animals move their eyes', *J Comp Physiol* **185**, 341–352.

Lau, C. & Honrubia, V. (1986), 'Fast component threshold for vestibular nystagmus in the rabbit', *J Comp Physiol A* **160**, 585–592.

Lau, C., Honrubia, V. & Baloh, R. (1978), *Vestibular mechanisms in health and disease*, Academic, chapter The pattern of eye movement trajectories during physiological nystagmus in humans, pp. 37–44.

Leigh, R. & Zee, D., eds (1999), *The neurology of eye movements*, Oxford University Press.

Maas, E., Huebner, W., Seidman, S. & Leigh, R. (1989), 'Behavior of human horizontal vestibulo-ocular reflex in response to high-acceleration stimuli', *Brain Res* **499**, 153–156.

Miles, F. (1995), The sensing of optic flow by the primate optokinetic system, *in* 'Eye movement research: mechanisms, processes and applications', Elsevier.

Miles, F. (1998), 'The neural processing of 3-d visual information: evidence from eye movements', *Eur J Neurosci* **10**, 811–822.

Miles, F., Kawano, K. & Optican, L. (1986), 'Short-latency ocular following responses of monkey. i. dependence on temporospatial properties of the visual input', *J Neurophysiol* **10**, 153–158.

Mustari, M., Fuchs, A. & Wallman, J. (1988), 'Response properties of dorsolateral pontine units during smooth pursuit in the rhesus macaque', *J Neurophysiol* **60**, 664–686.

Myung, I. (2003), 'Tutorial on maximum likelihood estimation', *J Math Psychol* **47**, 90–100.

Nakahara, H., Nakamura, K. & Hikosaka, O. (2006), 'Extended LATER model can account for trial-by-trial variability of both pre- and post- processes', *Neural Networks* **19**, 1027–1046.

Nelder, J. & Mead, R. (1965), 'A simplex method for function minimization', *Comput J* **7**, 308–313.

Ohm, J. (1928), *Zur Tätigkeit des Augenmuskelsenders*, Selbstverlag des Verfassers, in Kommission bei W. Postberg.

Raphan, T., Matsuo, V. & Cohen, B. (1979), 'Velocity storage in the vestibulo-ocular reflex arc (VOR)', *Exp Brain Res* **35**, 229–248.

Robinson, D. (1976), 'Adaptive gain control of the vestibulo-ocular reflex by the cerebellum', *J Neurophysiol* **39**, 954–969.

Robinson, D. (1981*a*), *The Handbook of Physiology, The nervous system, Vol 2.*, Oxford University Press, USA, chapter Control of eye movements, pp. 1275–1320.

Robinson, D. (1981*b*), 'The use of control systems analysis in the neurophysiology of eye movements', *Annu Rev Neurosci* **4**, 463–503.

Shelhamer, M. (1992), 'Correlation dimension of optokinetic nystagmus as evidence of chaos in the oculomotor system', *IEEE T Bio-Med Eng* **39**, 1319–1321.

Shelhamer, M. (1996), 'On the correlation dimension of optokinetic nystagmus eye movements: computational parameters, filtering, nonstationarity, and surrogate data', *Biol Cybern* **76**, 237–250.

Shelhamer, M. (1998), 'Nonlinear dynamic systems evaluation of 'rhythmic' eye movements (optokinetic nystagmus)', *J Neurosci Methods* **83**, 45–56.

Shelhamer, M. & Gross, C. (1998), 'Prediction of the sequence of optokinetic nystagmus eye movements reveals deterministic structure in reflexive oculomotor behaviour', *IEEE T Bio-Med Eng* **45**, 668–670.

Shelhamer, M., Tiliket, C., Roberts, D., Kramer, P. & Zee, D. (1994), 'Short-term vestibulo-ocular reflex adaptation in humans', *Exp Brain Res* **100**, 328–336.

Simpson, J., Soodak, R. & Hess, R. (1979), *Reflex control of posture and movement. Prog. Brain. Res., Vol.50*, Elsevier, chapter The accessory optic system and its relation to the vestibulocerebellum, pp. 715–724.

Smith, P. & Ratcliff, R. (2004), 'Psychology and neurobiology of simple decisions', *Trends Neurosci* **27**, 161–168.

Sperry, R. (1950), 'Neural basis of the spontaneous optokinetic response produced by visual inversion', *J Comp Physiol Psychol* **43**, 482–489.

Steinman, R., Kowler, E. & Collewijn, H. (1990), 'New directions for oculomotor reserach', *Vision Res* **30**, 1845–1864.

Temel, Y., Visser-Vandewalle, V. & Carpenter, R. (2009), 'Saccadometry: a novel clinical tool for quantification of the motor effects of subthalamic nucleus stimulation in parkinson's disease', *Exp Neurol* **216**, 481–489.

Ter Braak, J. (1936), 'Untersuchungen über den optokinetischen nystagmus', *Arch Neerl Physiol* **21**, 309–376. Translated in: Collewijn H (ed). The oculomotor system of the rabbit and its plasticity. Studies of brain function. Springer–Verlag, Berlin 1981.

Thilo, K., Guerraz, M., Bronstein, A. & Gresty, M. (2000), 'Changes in horizontal oculomotor behaviour coincide with a shift in visual motion perception', *Neuroreport* **11**, 1987–1990.

Tian, J., Zee, D. S. & Walker, M. F. (2007), 'Rotational and translational optokinetic nystagmus have different kinematics', *Vision Res* **47**, 1003–1010.

Tomko, D. L. & Paige, G. D. (1992), 'Linear vestibuloocular reflex during motion along axes between nasoocipital and interaural', *Ann NY Acad Sci* **656**, 233–241.

Trillenberg, P., Gross, C. & Shelhamer, M. (2001), 'Random walks, random sequences, and nonlinear dynamics in human optokinetic nystagmus', *J Appl Physiol* **91**, 1750–1759.

Trillenberg, P., Zee, D. & Shelhamer, M. (2002), 'On the distribution of fast-phase intervals in optokinetic and vestibular nystagmus', *Biol Cybern* **87**, 68–78.

Tuckwell, H. (1988), *Introduction to theoretical neurobiology: Nonlinear and Stochastic Theories*, Cambridge University Press.

van Beers, R. (2007), 'The sources of variability in saccadic eye movements', *J Neurosci* **27**, 8757–8770.

van Beers, R. (2008), 'Saccadic eye movements minimize the consequences of motor noise', *PLoS ONE* **3**, e2070.

van Die, G. & Collewijn, H. (1986), 'Control of human optokinetic nystagmus by the central and peripheral retina: effects of partial field masking, scotopic vision and central retinal scotomata', *Brain Res* **383**, 185–194.

Voss, A. & Voss, J. (2008), 'A fast numerical algorithm for the estimation of diffusion model parameters', *J Math Psychol* **52**, 1–9.

Waespe, W. & Henn, V. (1977), 'Vestibular nuclei activity during optokinetic after-nystagmus (OKAN) in the alert monkey', *Exp Brain Res* **30**, 323–330.

Walley, R. (1967), 'Receptive fields in the accessory optic system of the rabbit', *Exp Neurol* **17**, 27–43.

Walls, G. (1962), 'The evolutionary history of eye movements', *Vision Res* **2**, 69–80.

Watanabe, Y., Ohmura, A., Shojaku, H. & Mizukoshi, K. (1994), 'Optokinetic nystagmus elicited by a random dot pattern and a wide interval stripe pattern in normal subjects', *Acta Oto-Laryngol* **511**, 104–108.

Wester, S. T., Rizzo, J. F., Balkwill, M. D. & Wall, C. (2007), 'Optokinetic nystagmus as a measure of visual function in severely visually impaired patients', *Invest Ophth Vis Sci* **48**, 4542–4548.

Wobbrock, J., Findlater, L., Gergle, D. & Higgins, J. (2011), The aligned rank transform for nonparametric factorial analyses using only anova procedures, *in* 'CHI 2011 Proceedings of the 2011 annual conference on human factors in computing systems'.