

Computational Insights Into the Architecture and Operation of the Rat Head Direction Cell System

Hector J.I. Page
The Queen's College

Submitted for the Degree of DPhil in Experimental Psychology
The University of Oxford

Trinity Term 2015

Acknowledgements

No man is an island, or even a peninsula. This thesis would never have been possible without the love, support, and patience of so many who deserve my thanks.

Firstly, I am indebted to Sir Michael Hintze, of the Hintze Family Charitable Foundation, for their financial support via the Oxford Foundation for Theoretical Neuroscience and Artificial Intelligence (OFTNAI) in the last year of my studies.

I am also indebted to The Queen's College, who generously awarded me both a travel grant and the 650th anniversary trust fund grant.

I must principally thank my supervisor Dr. Simon Stringer. A keen theoretical mind, his enthusiasm and encouragement over the last 4 years has been unwavering, and his faith in my abilities has sustained me when my own has failed. His proofreading of multiple drafts of this thesis was invaluable.

For an exciting and fruitful collaboration, I thank Dr. Rebecca Knight and Professor Kate Jeffery. The latter in particular has set my research down a different path altogether, for which I am most grateful.

I must also give special thank Dr. Daniel Walters for his particular combination of technical expertise and warm friendship when teaching a psychologist how to bully a computer into doing his bidding. Without him I might well have done it all with pen and paper.

Thanks too, to the OFTNAI gang as a whole, and in particular to James Tromans for getting me started, to Ben Evans for his tough love, and to Juan Galeazzi-Gonzalez for his companionship.

For their personable and informative discussions, which helped clarify several aspects of this work, I thank Dr. Jeffrey Taube and Dr. Douglas Wallace.

To the friends I have made along the way in Oxford, over my entire time here, I say thank you for all the good times. I wish you all the best of luck in the uncertain future: learn by going where you have to go.

Thanks to the musicians past and present, whether we shared a stage or a classroom. For their many lessons, their company on late nights, and their friendship. This understanding extends to everything.

My eternal gratitude to the doctors and nurses of Oxford's John Radcliffe Hospital for teaching me how to walk again, and the physiotherapy department in particular for their skill with kintsugi.

The love of my family, in both emotional and financial form, deserves far more than this acknowledgement. They have taught, and continue to teach, me so much about how to face life's ups and downs with grace and charm.

Last, but not in any way least, I thank with all my heart my (significantly) better half, Y-Ling Chi. You delight and humble me in equal measure.

LIST OF ABBREVIATIONS

HD . . . . . Head Direction
COMB. . . . . Combination
CANN . . . . . Continuous Attractor Neural Network
AHV . . . . . Angular Head Velocity
MVN . . . . . Medial Vestibular Nuclei
SGN . . . . . Supragenual Nucleus
NPH . . . . . Nucleus Prepositus Hypoglossi
DTN . . . . . Dorsal Tegmental Nucleus
LMN . . . . . Lateral Mammillary Nucleus
ADN . . . . . Anterodorsal Thalamic Nucleus
PaS . . . . . Parasubiculum
PoS . . . . . Postsubiculum
RSC . . . . . Retrosplenial Cortex
MEC . . . . . Medial Entorhinal Cortex
HPC . . . . . Hippocampus (HPC)
ENT . . . . . Entorhinal Cortex
AVN . . . . . Anteroventral Thalamic Nucleus

Abstract

Computational Insights Into the Architecture and Operation of the Rat Head Direction Cell System
Hector J.I Page
The Queen's College
Submitted for the Degree of DPhil in Experimental Psychology
Trinity Term 2015

Head direction (HD) cells signal HD in the horizontal plane in a body-external (allothetic) frame of reference, and are a critical neural component of spatial cognition. This thesis uses computational modelling to investigate their properties. HD cell preferred firing directions are influenced depending on the discrepancy between sources of directional information, with allothetic information dominating up to a critical discrepancy and body-internal (idiothetic) information dominating after this. A computational mechanism is outlined which is capable of replicating these discrepancy-dependent changes in preferred firing direction. This mechanism depends on dynamic changes in HD cell responses to allothetic information and is shown to be affected by salience, defined as the degree to which a landmark has been previously experienced as a stable source of directional information. This model is then built on by providing a mechanism explaining how the preferred firing directions of any two given HD cells maintain their angular separation following idiothetic-allothetic discrepancies, a property known as isomapping. The accuracy of HD system updating solely via idiothetic information (path integration) in examined a model system. Two major limiting factors are found: the time neurons take to respond to stimulation, and the presence of direct reciprocal connectivity between HD cells. Finally a neural network model, based on known physiology, which does not incorporate direct HD to HD connectivity is proposed. This model is demonstrated to be capable of extremely accurate path integration.

Contents

Acknowledgements **i**

List of Abbreviations **iii**

Abstract **iv**

1 An introduction to Head Direction cell properties **1**

1.1 Head direction cell responses 4

1.1.1 Isomapping 6

1.1.2 Distribution and regional variation of HD cells 6

1.1.3 Angular head velocity by head direction cells 8

1.1.4 Angular head velocity cells 8

1.1.5 Variations in anticipatory firing 10

1.2 Structure of the HD cell system 10

1.2.1 Lesion studies indicate a roughly hierarchical organisation of HD cell areas . 12

1.2.2 The HD signal depends on vestibular input 13

1.2.3 Vestibular pathways to the HD system 15

1.2.4 The HD signal is most likely generated in the LMN-DTN complex 16

1.2.5 Other idiothetic information sources 19

1.3 Idiothetic vs Allothetic control of HD responses 22

1.3.1 Path integration 23

1.3.2 Visual landmark interaction with path integration mechanisms 25

1.3.3 Stability of landmark cues 31

1.3.4 Proximal vs distal landmarks . . . 32
1.3.5 Origins of allocentric input to the HD system . . . 33
1.3.6 Non-visual allothetic input to the HD system . . . 35
1.4 Relationship to other spatially-responsive cells . . . 36
1.5 Behavioural influence . . . 39

2 Modeling Head Direction cells **46**
2.1 Continuous Attractor Neural Networks (CANNs) . . . 49
2.2 Competitive Neural Networks . . . 52
2.3 Requirements of HD cell models . . . 53
2.4 HD cell models . . . 55
2.4.1 McNaughton, Chen and Markus (1991) . . . 56
2.4.2 Skaggs et al. (1995) . . . 57
2.4.3 Zhang (1996) . . . 58
2.4.4 Sharp et al. (1996) . . . 59
2.4.5 Redish et al. (1996) . . . 60
2.4.6 Goodridge and Touretzky (2000) . . . 62
2.4.7 Rubin et al. (2001) . . . 64
2.4.8 Arleo and Gerstner (2000; 2001) . . . 65
2.4.9 Xie et al (2002) . . . 66
2.4.10 Stringer et al. (2002) . . . 67
2.4.11 Song and Wang (2003; 2005) . . . 68
2.4.12 Hahnloser (2003) . . . 69
2.4.13 Degris et al. (2004) . . . 70
2.4.14 Stringer and Rolls (2006) . . . 71
2.4.15 van der Meer et al. (2007) . . . 72
2.4.16 Stratton et al. (2010, 2011) . . . 73
2.4.17 Massoud and Horiuchi (2011) . . . 75
2.4.18 Fortenberry et al. (2012) . . . 75

2.4.19 Walters and Stringer (2010)/Walters, Stringer and Rolls (2013) 76
2.5 General notes on the models in this thesis 78

3 Cue conflict resolution **80**
3.1 Model Description 81
3.2 Model Operation 84
3.2.1 Model Equations 84
3.2.2 Numerical Methods 89
3.2.3 Experimental Protocol 90
3.3 Resolving cue conflict: A hypothesis 92
3.4 Simulation Results 95
3.4.1 Main Result 98
3.5 Saliency of visual information 105
3.5.1 Introduction of homosynaptic LTD 107
3.5.2 Effect of learning rate manipulations 111
3.5.3 Varying spread of feed-forward and recurrent weights 113
3.5.4 Varying packet sizes 120
3.6 Discussion 124

4 Isomapping of Head Direction Cell preferred firing directions following environmental manipulations **133**
4.1 Model description 136
4.1.1 Experimental protocol 136
4.2 Isomapping: A hypothesis 138
4.2.1 The relationship between the change in feed-forward weight structure and head direction cell preferred firing direction 141
4.3 Simulation results 143
4.3.1 Main result 145
4.3.2 Effect of cue conflict direction 154
4.3.3 Accuracy of path integration 157

4.3.4 Salience of visual information 160
4.3.5 Effect of learning rate 162
4.3.6 Effect of path integration strength 166
4.3.7 Effect of rotation speed 172
4.4 Discussion 177

5 Factors affecting path integration speed **186**
5.1 Experiment 1: Pre-wired network 191
5.1.1 Model description 192
5.1.2 Model operation 193
5.1.3 Results 198
5.1.4 Estimating rise time 206
5.2 Experiment 2: Self-organising network 208
5.2.1 Model description 208
5.2.2 Model operation 208
5.2.3 Results 211
5.3 Discussion 220

6 Path integration in a two-layer system: Solving the speed problem **231**
6.1 Experiment 1: Pre-wired network 233
6.1.1 Model description 233
6.1.2 Results 248
6.2 Experiment 2: Self-organising network 262
6.2.1 Model description 262
6.2.2 Results 270
6.3 Discussion 283

7 Discussion **292**
7.1 Summary of findings 292
7.2 Experimental predictions 300

7.3 Future directions 308
7.3.1 Origins and nature of visual landmark input 308
7.3.2 Proximal-Distal landmark distinctions 312
7.3.3 Biological realism of AHV and COMB cells and more general path integration 319
7.3.4 HD attractor networks without visual input 324
7.3.5 Relationship to bursty HD cell firing following vestibular lesions 335
7.3.6 Preferred firing direction shifts and path integration accuracy 338
7.4 Closing statement 347

Bibliography **348**

Chapter 1

An introduction to Head Direction cell properties

Across a broad range of sentient organisms, the ability to navigate is indispensable. For ants, whales, rats, humans, birds, and a great number of other species, day to day survival can depend on effectively navigating the world around us. Indeed, some theories of artificial intelligence (Pfeifer and Scheier, 2001) require that any intelligent agent must be situated within an environment. By this, it is meant that an intelligent agent must exist in a changing environment within which they explore and with which they interact. As a result, intelligent behaviour arises from the interaction of an agent and the environment in which it exists. Logically it then follows that the ability to accurately navigate within a given environment is a key aspect of any agent of which intelligence is a characteristic. Whether an agent is avoiding prey, finding a nest, foraging for food, or even driving to work, spatial navigation is a fundamental cognition; regardless of whether one agrees that theories of situatedness are central to intelligence. Such navigation relies on knowledge of location within the world, coupled with a sense of directional heading.

An understanding of spatial cognition may provide insight into various cognitive disorders involving spatial processing. These include topographical disorientation (Berthoz and Viaud-Delmon, 1999; Iaria et al., 2009), as well a variety of episodic memory deficits, which necessarily have a spatial aspect, such as Korsakoffs amnesia (Holdstock et al., 1999). More broadly, spatial cognitive deficits are an important aspect of a variety of neurodegenerative diseases (Possin, 2010), such as Alzheimers (Monacelli et al., 2003) and Parkinsons (Owen et al., 1997). In many cases, subtle spatial deficits may be early predictors of future cognitive decline, or provide differential diagnoses of similar conditions (Johnson et al., 2005). Deficits in spatial cognition appears to be a component of learning difficulties, with patients displaying an impaired ability to process spatial information in nonverbal learning disability (Mammarella et al., 2010) and spatial processing being related to poor problem solving (Passolunghi and Mammarella, 2010).

What underpins the brains considerable ability to make sense of spatial information? Several cell types contribute. Place cells (O'Keefe and Dostrovsky, 1971; O'Keefe, 1979) fire when the animal is located at a specific location within the environment. Grid cells (Fyhn et al., 2004; Hafting et al., 2005) fire at multiple locations spaced within an environment, in a regular grid-like pattern. Boundary vector cells (Barry et al., 2006; Lever et al., 2009) fire when the animal is at a particular distance from environmental boundaries, including no distance (these latter cells are a subtype of boundary vector cells, known as border cells). Head Direction (HD) cells (Ranck, 1985; Taube et al., 1990a) firing signals head direction in the horizontal plane. It is likely that these kinds of cells provide estimates of direction, distance travelled, and location within an environment, therefore providing a neurophysiological basis for spatial processing. This thesis concerns the last of these cell types: Head Direction

cells. Recently, these cells have been theorised to be involved not only in spatial cognition, but also in providing a spatial reference to episodic memory (Hasselmo et al., 2010; Byrne et al., 2007). Indeed, this may help to explain the fact that episodic memory can be impaired following lesions of HD cell areas (Aggleton and Brown, 1999).

This chapter presents a selected review of the experimental literature on Head Direction (HD) cells, with the intent of placing in context the simulation work carried out in this thesis. Several essential facts about the HD system are established. HD cells are introduced as signalling head direction in an allocentric (body-external) spatial reference frame. They are shown to exist in multiple interconnected brain areas, and are hierarchically organised, with the lateral mammillary nucleus and dorsal tegmental nucleus given as the source of the HD signal. The dependence of the HD signal on vestibular input is established. The convergence of idiothetic (body-internal) and allothetic (body external) information sources is discussed, as well as the physiological correlates of these sources. It is established that the HD cell system can use idiothetic information alone, but such a process accumulates error, and for this reason allothetic landmarks can recalibrate the HD cell system. However, a caveat is introduced that landmarks must be learned as stable indicators of directional information before they have any influence over HD cell firing. A brief discussion is provided of the role of HD cells in a broader context; both in relation to other spatially-responsive cells and in terms of their influence on behaviour.

1.1 Head direction cell responses

Head direction (HD) cells, first discovered in the rat postsubiculum (PoS) (Ranck, 1985), respond to head direction in the horizontal (also known as azimuth, or yaw) plane. Crucially, HD cell firing appears not to be affected by body position, location, current behaviour, or geomagnetic fields. Taube and colleagues (Taube et al., 1990a) showed core postsubicular HD cell characteristics to be as follows: individual HD cells show maximal firing rate at a single preferred direction, with firing rate dropping off symmetrically as true head direction deviates from preferred direction on either side. Beyond approximately $90°$ cells showed little to no firing. Cell tuning curves can therefore be described as roughly Gaussian. A typical tuning curve is pictured in Figure 1.1. As a population, HD cell preferred firing directions are evenly distributed from $0 - 360°$. True head direction is therefore represented by the pattern of overlapping firing of a population of HD cells. It should be noted that no underlying topographical organisation has been discovered among HD cell populations.

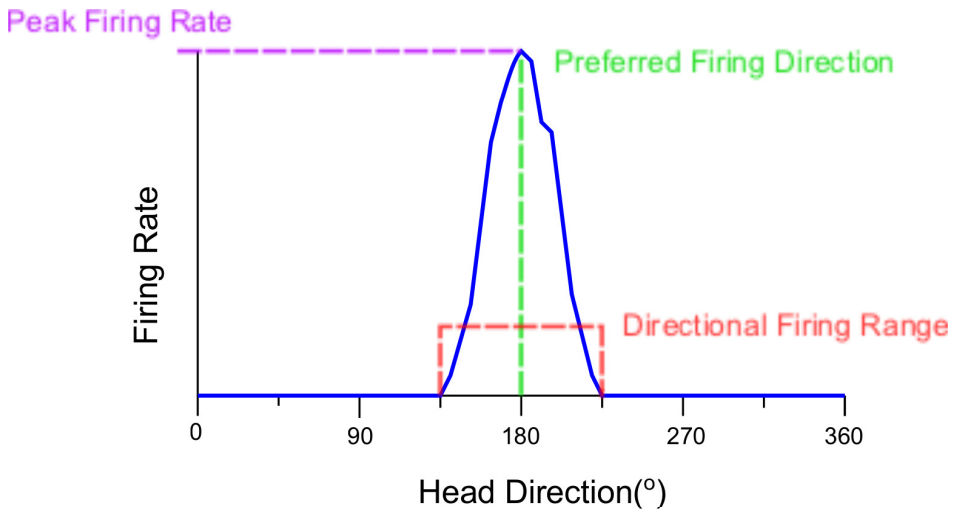


Figure 1.1: Stylised HD cell response curve, show preferred firing direction, peak firing rate, and directional firing range.

The recording method used by Taube and colleagues is based on pre-existing methodology (Muller et al., 1987; Muller and Kubie, 1987) and is paradigmatic for HD cells - it reoccurs, with some modifications, repeatedly in the literature. A hungry rat is placed in a grey cylindrical apparatus (Figure 1.2) and searches for food pellets, which are periodically dropped. The floor material can be changed between recording sessions, and the apparatus as a whole is isolated from the rest of the laboratory by a featureless curtain. A single white cue card provides the only source of orienting information to the rat. Head direction, position within the apparatus, and cell firing are all recorded as the rat explores the apparatus. Two LEDs (one red, one green) are affixed to the rat along with recording electrodes, in order to help with video tracking of head direction and location.

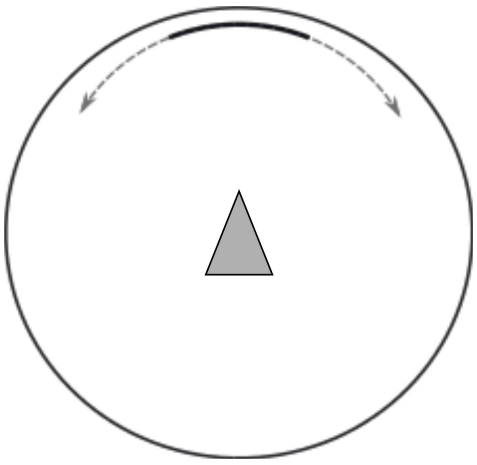


Figure 1.2: Paradigmatic cylinderical apparatus used for recording HD cells. The rat (triangle) explores a cylindrical area containing a single cue card, which can be rotated in order to generate a mismatch between vision and other sources of directional information.

The HD cell response is very stable. Provided that head direction remains the same, cells will show little to no firing rate adaptation (Taube and Muller, 1998). Neither the preferred direction, the peak firing rate, nor directional firing range are seen to change over a period of weeks or months (Taube,

2007). It is also supramodal, in that it is influenced by a number of sensory modalities (Blair et al., 1998; Wiener et al., 2002; Knight et al., 2014). This latter point will be unpacked in some depth in this thesis. Briefly, HD cells represent a convergence of interally-derived (idiothetic) and externally-derived (allothetic) information. Together these facts indicate a role of the HD system in generating a reliable signal representing current head direction in a body-external (allocentric) frame of reference. Indeed, the HD cell system can be seen as a convergence of two streams; an ascending idiothetic stream and a descending allothetic stream (Bassett and Taube, 2005).

1.1.1 Isomapping

Although HD cells may have different preferred firing directions depending on the environment, the angular separation of preferred firing directions between individual HD cells remains the same in all cases. For example, in room 1 HD cell 1 has a preferred direction of 60° and HD cell 2 has a preferred direction of 90°. In room 2, if HD cell 1 has a new preferred direction of 100°, HD cell 2 will have a new preferred firing direction of 130°. Thus, the difference in preferred firing directions of these two cells is always 30°, regardless of their absolute values. This property is known as isomapping, and is observed experimentally in a number of cases (Taube et al., 1990b; Taube and Burton, 1995; Zugaro et al., 2004; Yoganarasimha et al., 2006; Taube, 2007).

1.1.2 Distribution and regional variation of HD cells

As mentioned in the previous section, HD cells were initially discovered in PoS. However, they are found in a wide variety of interconnected limbic system areas roughly corresponding to the Papez circuit (Papez, 1937). These areas include the lateral mammillary nucleus (LMN) (Stackman and Taube, 1998), the anterodorsal thalamic nucleus (ADN) (Taube, 1995), the retrosplenial cortex (RSP) (Chen

et al., 1994; Cho and Sharp, 2001), the dorsal tegmental nucleus of Gudden (DTN) (Sharp et al., 2001), the entorhinal cortex (ENT) (Sargolini et al., 2006), anteroventral thalamic nucleus (AVN) (Tsanov et al., 2011), and the parasubiculum (PaS) (Boccara et al., 2010).

HD cells have also been recorded from a number of regions outside of this limbic circuit, including the lateral dorsal thalamus (Mizumori and Williams, 1993; Yoganarasimha et al., 2006), the dorsal striatum (Wiener, 1993; Mizumori et al., 2000), the hippocampal CA1 region (Leutgeb et al., 2000), the medial prestriate cortex (Chen et al., 1994), and the medial precentral cortex (Mizumori et al., 2005).

It is not clear why the directional signal is replicated across so many different brains regions. None of these areas can be considered strictly HD cell areas, given that the proportion of cells displaying HD cell firing characteristics varies between areas (Taube, 2007). For example, in ADN approximately 56% of neurons can be classified as HD cells (Taube, 1995; Taube and Bassett, 2003), whilst this figure is closer to 25% in PoS (Taube et al., 1990a) and LMN (Stackman and Taube, 1998). Across limbic regions, cell firing properties are broadly similar. Peak firing rate tends to vary within regions as much as between regions. There are variations in average tuning curve width, with cells showing firing over more or less than the 90° originally observed.

Many areas containing HD cells also contain cells which signal angular head velocity (AHV), as well as cells which respond to both HD and AHV. This is significant, given that AHV is likely a useful idiothetic information source for the HD cell system.

1.1.3 Angular head velocity by head direction cells

Some HD cells have angular head velocity as a secondary firing correlate. Stackman and Taube (1998) reported that LMN HD cell firing rates were modulated by angular head velocity, with faster head turns being accompanied by higher firing rates. Specifically, cells showed a lateralised effect of rotational direction, with right LMN cells having higher peak firing rates during clockwise turns, and left LMN during counter clockwise turns. In the DTN a slightly different type of modulation is found (Bassett and Taube, 2001), with a small proportion of recorded cells responding principally to angular head velocity yet showing increased firing over a range of particular head directions. However, these cells did show broader tuning and lower peak rates than true HD cells. Further studies have also found DTN HD cell firing rates to be modulated by AHV (Sharp et al., 2001). Interestingly, this latter study found an additional effect of AHV on HD cell tuning curve width. HD cell tuning functions were found to be vary with head turns which were contra or ipsiversive to the cells hemisphere. This lateralised effect of AHV on tuning curve width has also been seen in LMN, where ispiversive head turns cause a narrowing of HD cell tuning curves. Right hemisphere HD cells showed narrower tuning curves during clockwise turns and left hemisphere HD cells showed narrower tuning curves during counterclockwise turns (Blair et al., 1998). Although the correlation is not as strong as in other areas, ADN HD cell firing is also modulated by AHV, increasing in rate with faster head turns (Taube, 1995).

1.1.4 Angular head velocity cells

It should be noted that many HD areas also contain pure angular head velocity (AHV) cells. In the DTN, 75% of cells recorded by Bassett and Taube (2001) fired principally as a function of AHV. However, these cells varied in their response. Some cells were symmetric, with firing rates correlated

to AHV in both clockwise and counter clockwise directions. Other cells were asymmetrical, having firing rates which signalled AHV in a single direction. In the opposite direction these asymmetrical cells showed either no firing, or firing negatively correlated with AHV. Other studies have found AHV in the DTN, although only of the asymmetric variety (Sharp et al., 2001). Both symmetric and asymmetric AHV cells have also been found in the LMN (Stackman and Taube, 1998). LMN AHV cells were classified into two subcategories: fast cells showed a positive correlation between firing rate and AHV, thus resembling AHV cells found in DTN, whilst slow cells showed a negative correlation. Figure 1.3 shows stylised tuning curves for various types of AHV cells.

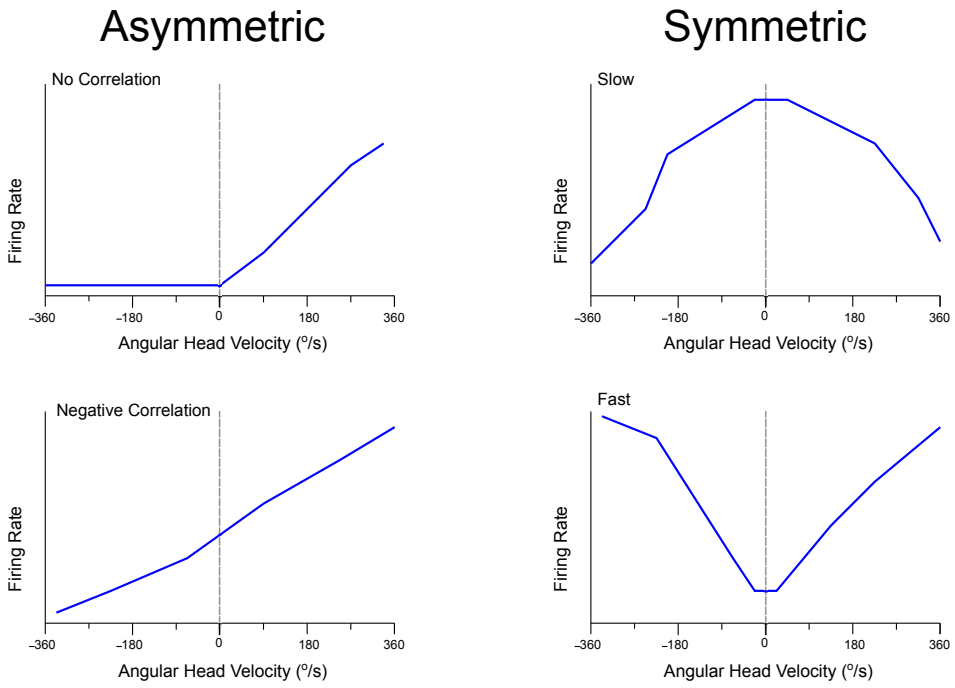


Figure 1.3: Stylised AHV cell responses, showing asymmetrical cells (left) with both no correlation (top) and negative correlations (bottom). Also shown are symmetrical AHV cells (right), with slow (top) and fast (bottom) responses.

1.1.5 Variations in anticipatory firing

Interestingly, not all HD areas signal current HD (Stackman et al., 2003). It seems that HD cells in ADN anticipate future heading by approximately 25ms (Blair and Sharp, 1995; Taube, 1998), as do HD cells in RSP (Cho and Sharp, 2001). Greater anticipation has been observed in LMN, although the amount varies from 40ms (Blair et al., 1998) to 75ms (Stackman and Taube, 1998). In contrast, PoS cells show no anticipation, thus signalling current HD (Sharp, 2005; Taube, 2007). Whilst the origins of this anticipation are not well understood, it is possible that this is related to vestibular input dynamics (van der Meer et al., 2007), or from interactions of multiple differently-timed idiothetic inputs (Fortenberry et al., 2012).

1.2 Structure of the HD cell system

Figure 1.4 shows a diagram of the known principal HD cell circuitry. Regions containing HD cells are densely interconnected, and Figure 1.4 should be taken as somewhat of a simplification for the sake of clarity. A general, but by no means strict, organisational principle for the HD circuit is a dependence of the cortical HD signal on subcortical HD cell directional firing. This is supported by a large body of combined lesion and electrophysiological studies demonstrating an abolition of cortical HD firing following subcortical lesions, but not vice versa (Taube, 2007; Clark and Taube, 2012).

There exists between LMN and DTN reciprocal connectivity (Allen and Hopkins, 1988, 1989), and this LMN-DTN complex is thought to be the locus of HD signal generation. As well as back to DTN, LMN projects to ADN (Hayakawa and Zyo, 1984). ADN itself projects reciprocally to RSC (Van Groen and Wyss, 2003) as well as to PoS and PaS (Shibata, 1993), both directly and via RSC

(Van Groen and Wyss, 1990). PoS projects downstream to EC, but also back upstream to earlier HD areas such as the dorsal thalamic nuclei (ADN, AVN) and LMN (Van Groen and Wyss, 1990). RSP itself projects strongly to, although does not receive strongly from, PoS (van Groen and Wyss, 1992) as well as back to LDN and AVN (Shibata, 1998).

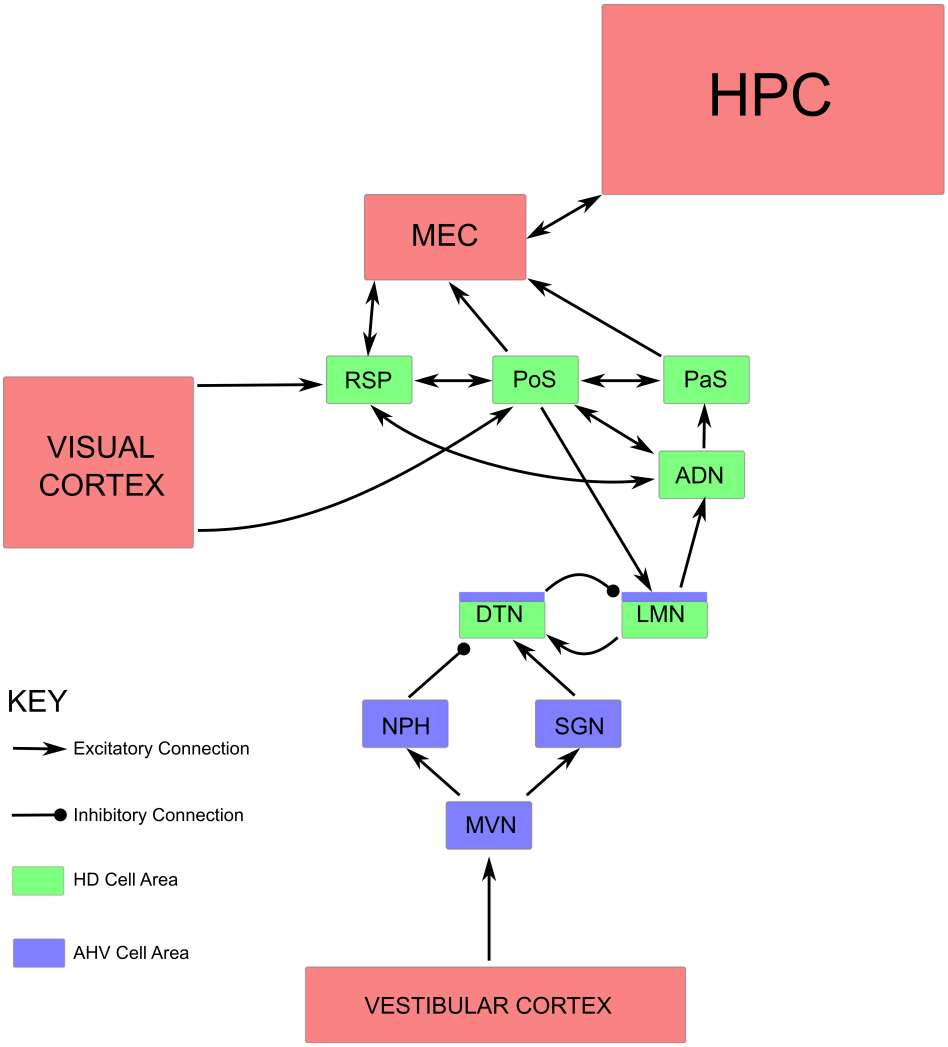


Figure 1.4: Simplified connectivity of HD cell (green), AHV cell (blue), and input/output regions (red). Abbreviated areas are: Medial Vestibular Nuclei (MVN), Supragenual Nucleus (SGN), Nucleus Prepositus Hypoglossi (NPH), Dorsal Tegmental Nucleus (DTN), Lateral Mammillary Nucleus (LMN), Anterodorsal Thalamic Nucleus (ADN), Parasubiculum (PaS), Postsubiculum (PoS), Retrosplenial Cortex (RSC), Medial Entorhinal Cortex (MEC), and Hippocampus (HPC).

1.2.1 Lesion studies indicate a roughly hierarchical organisation of HD cell areas

Goodridge and Taube (1997) found that ADN lesions eliminated the HD signal in PoS. ADN lesions have also been found to eliminate HD cell firing in the medial entorhinal cortex and PaS (Clark et al., 2011).

However, Goodridge and Taube (1997) did not find any dependence of ADN HD cell firing on intact PoS. Other studies have a found a lack of ADN HD cell disruption following lesions of RSP (Clark et al., 2010), medial entorhinal cortex (Clark and Taube, 2011), and parietal cortex (Calton et al., 2008). There has been similar lack of ADN or PoS HD signal disruption following lesions of the hippocampus (Golob and Taube, 1997). PoS lesions do, however, increase the directional range of ADN HD cells, demonstrating a reduction in accuracy of the ADN HD signal (Goodridge and Taube, 1997).

ADN HD cell firing is, however, abolished by either DTN (Bassett et al., 2007) or LMN (Blair et al., 1997; Bassett et al., 2007). Indeed, LMN lesions also seem to abolish HD firing in PoS and PaS, as well as medial entorhinal cortex (Sharp and Koester, 2008). This is hardly surprising given the established dependence of these areas on an intact ADN HD signal, which is itself abolished by LMN lesions.

Combined, these studies reflect a roughly hierarchical flow of the HD signal in a central circuit from DTN /LMN to ADN to PoS and other parahippocampal structures, with DTN/LMN appearing the be central to the HD signal. This information can then be conveyed onwards to other brain areas such as the hippocampus.

1.2.2 The HD signal depends on vestibular input

It is likely that an accurate HD signal depends on ascending vestibular input (Clark et al., 2012). Indeed, the HD signal can be derived from angular head velocity information, which would originate in vestibular areas (Taube and Bassett, 2003). In addition to being important for HD signal generation, it is likely that idiothetic (e.g., vestibular) information is important for updating the HD representation in the absence of allothetic directional cues (McNaughton et al., 1991) (see this chapter Section 1.3 and Chapter 2 for a full discussion of path integration theories).

It seems that general vestibular lesions impair navigation on a behavioural level, especially when navigation relies more on idiothetic information (see Yoder and Taube 2014 for review). This effect is potentially due to the vestibular contribution to the HD signal, itself a major component of spatial navigation. However, the link between the HD system and spatial navigation is not clearly understood. More convincing evidence of vestibular importance for the HD signal comes from examining the effects of vestibular lesions on HD cell firing. Damage to vestibular hair cells by neurotoxic lesions via administration sodium arsanilate has been shown to eliminate HD cell activity in ADN (Stackman and Taube, 1997). In fact, this damage has an effect over a long duration, several months, suggesting that vestibular input plays a vital and irreplaceable role in generation of the HD signal. This has been verified by temporary inactivation of vestibular hair cells with tetrodotoxin, which also abolishes HD cell activity in ADN (Stackman and Herbert, 2002).

More precise damage reveals specific roles for different components of the vestibular system. One major part of the vestibular apparatus is the semicircular canals, which detect acceleration in three

perpendicular axes (pitch, roll and yaw). Blocking the semicircular canals, part of the inner ear, affects the functioning of the vestibular system (Plotnik et al., 2002). This blocking has also been found to disrupt HD firing in the ADN (Muir et al., 2009). The importance of the semicircular canals is also indicated by a lack of HD cell firing in the ADN of mice with a genetic mutation causing a lack of horizontal canals (Yoder and Taube, 2014).

Another important part of the vestibular apparatus is the otolith organs, which detect linear accelerations and displacements of the head, particularly in reference to gravity. Studies involving mice with genetically deficient otoliths show that while ADN HD cells are present, they are significantly less stable than expected (Yoder and Taube, 2009).

Otolith-based perception of gravity may play an important role in updating the HD signal when locomoting in planes other than the earth horizontal. The HD cell system tends to treat vertical surfaces as a new horizontal (Stackman et al., 2000), but loses directionality during upside down (i.e., on the ceiling) locomotion (Calton and Taube, 2005; Gibson et al., 2013). Contrastingly, HD cell firing was disrupted when locomoting on any plane other than the horizontal in situations where the otolithic contribution is minimised due to microgravity during parabolic flights (Taube et al., 2004). It is possible then, that an otolithic sense of gravity is involved both in updating the HD system to treat vertical surfaces as a new horizontal and in the disruption of the HD signal during inverted locomotion.

Together, these studies indicate that whilst otolith organs are important contributors to the stability of HD cell firing and may play an important role during locomotion in planes other than the earth

horizontal, it is the semicircular canals that are truly necessary for the HD signal.

1.2.3 Vestibular pathways to the HD system

Vestibular information is likely to be conveyed to the HD system via the medial vestibular nuclei (MVN) (Winter and Taube, 2014). MVN responds to stimulation of semicircular canals (Uchino et al., 2005), and contains AHV cells (Leigh and Zee, 2015). One study from Liu and colleagues (Liu et al., 1984) identified a direct MVN-DTN projection. However, this finding has not been replicated (Biazoli et al., 2006) and it is generally taken that there is no evidence for direct MVN-DTN pathways (Hitier et al., 2013). However there are indirect MVN-DTN projections via the supragenual nucleus (SGN), to which MVN projects both directly and via the dorsal paragigantocellular reticular nucleus dorsal(PGRNd) and via the nucleus prepositus hypoglossi (NPH) (Biazoli et al., 2006). NPH and SGN themselves both project to DTN (Yoder and Taube, 2014). Studies involving transynaptic tracers tentatively support the idea of SGN and NPH as important vestibular inputs to the DTN (Brown et al., 2005). Crucially, the AHV signal seen in MVN is also found in both NPH and SGN, as well as in PGRNd (Winter and Taube, 2014).

The contribution of the SGN to the HD signal was tested with both bilateral electrolytic lesions, as well as unilateral neurotoxic lesions. It was found that whilst bilateral lesions abolished HD cell firing in ADN, unilateral lesions merely disrupted the stability of HD cell preferred directions in situations that tested idiothetic HD cell updating (Clark et al., 2012). This demonstrates a role for SGN in both maintaining and updating the HD signal based on vestibular information. The contribution of the NPH has also been tested, with NPH lesions resulting in a loss of the HD signal in ADN (Butler and Taube, 2015). Tellingly, partial lesions of NPH caused, much as in cases with unilateral SGN le-

sions, instability of HD cell preferred firing directions in situations requiring idiothetic update. Taken together, these studies suggest a role for both NPH and SGN in conveying vestibular information of a dual purpose: generating the HD signal, and updating HD cell firing based on idiothetic information.

1.2.4 The HD signal is most likely generated in the LMN-DTN complex

Contemporary opinion suggests that the LMN-DTN complex is a likely candidate substrate for the generation of the HD signal (e.g., Song and Wang 2005 and Taube 2007). Certainly, the combined lesion and electrophysiological studies cited above show that lesions to this complex tend to eliminate the HD signal in later areas. They also show that lesions to other areas within the HD system do not affect LMN/DTN HD cell firing.

The DTN in particular receives prominent connections from the vestibular system (see Section 1.2.3), which is thought to be critical to both HD signal generation and the updating of HD cell firing based on idiothetic movement information (see Section 1.3). Indeed, the DTN/LMN contains cells that are sensitive to both angular head velocity, as well as combinations of angular head velocity and head direction; information that is used to update HD cell firing to accurately track true HD.

The LMN-DTN complex shows reciprocal connectivity (Allen and Hopkins, 1988, 1989), which is particularly important given that the vast majority of computational HD cell models require this type of connectivity, as they are based on continuous attractor neural networks (CANNS, see Chapter 2 for further details). In brief, HD cells are thought to be able to sustain firing solely through HD to HD connectivity. The strength of this connectivity depends on HD cell preferred firing directions, such that HD cells are most strongly connected to other cells with similar preferred firing directions. At a

given HD, cells drive one another via this connectivity and, when this excitatory drive is balanced by mutual inhibition between cells, current HD is represented by a pattern of firing across a population of HD cells which are active to various degrees.

Part of the role of idiothetic information could be to provide external input to this network of reciprocally connected HD cells to alter the firing pattern and thus the representation of current HD (a process known as path integration, see Section 1.3) in the absence of visual input to the system. Without this idiothetic information, the packet of HD activity would be decoupled from idiothetic motion, and instead shift based on a variety of allothetic cues. Depending on other available inputs, the system might lose the ability to track true HD. Individual HD cells would still show preferred firing directions but would fire in a manner decoupled from true HD, with burst of activity as if the head were moving through a given cells preferred firing direction.

Several of the lesion studies cited above reported this bursty activity, as if the HD packet were moving around a CANN network that was no longer influenced by idiothetic inputs (Muir et al., 2009; Yoder and Taube, 2009). Several key observations from these studies suggest that bursty cells are indeed HD cells decoupled from vital inputs (Clark and Taube, 2012). HD and bursty cells were never observed at the same time, but did show similar spike train characteristics. The temporal order of bursty cell firing varied depending on direction of rotation, as would be expected from a CANN model. Bursting activity was depending on movement velocity, with faster rotations yielding shorter bursts of firing. Additionally, the percentage of spikes during bursts and samples containing bursts were approximately equal between HD cells and bursty cells. Further support for idea of an intact CANN decoupled from

idiothetic information comes from the observation that lesions of the LMN-DTN circuit do not result in bursty activity in ADN (Bassett et al., 2007). This suggests that an intact CANN, which would account for bursting after vestibular lesions, is no longer present following LMN-DTN damage. Despite this collection of evidence, several other lesion studies have reported no such bursty HD cell firing following vestibular lesions (Stackman and Taube, 1997; Stackman et al., 2000). It is not yet clear why this might be the case.

Further evidence of a CANN network comes from ensemble recordings of HD cells in awake and sleeping rats (Peyrache et al., 2015). In concordance with a CANN hypothesis of the HD cell system, the temporal relationship of ADN HD cells was preserved across brain states. This strongly suggests an internally-generated HD signal. Indeed, examining the population firing revealed a Gaussian packet of activity as theorised in CANN networks. This suggests that either ADN, or an area upstream of ADN such as DTN/LMN, provides attractor dynamics to generate the HD signal (Schultheiss and Redish, 2015).

If bursty cells do represent HD cells decoupled from their inputs, and HD cells are influenced by both idiothetic and allothetic sources of information, why are vestibular lesions enough to produce bursty cells? One possibility may be that allothetic cues to HD need to be learnt as stable with respect to idiothetic information (for more information, see Section 1.3.3). Thus, loss of idiothetic information prevents establishment of a consistent allothetic influence on the HD system. This would explain how bursty activity can shift between cells, with allothetic information updating the current HD representation but in an inconsistent manner. It may well be that other sources of idiothetic information

remain to contribute to the HD signal, as will be outlined below in Section 1.2.5. A full treatment of bursty firing and its possible root causes is given in the discussion of this thesis.

The LMN-DTN complex shows both the internal and the efferent/afferent connectivity necessary to fulfil a role as a generative circuit for the HD signal. It contains the necessary cell firing properties to both represent current HD and update this representation on the basis of idiothetic information. Damage to this complex eliminates the HD signal in downstream areas. The specific results of lesion studies in downstream areas, with bursty activity not seen after LMN-DTN lesions but rather after vestibular lesions, possibly indicate the presence of a CANN as the HD signal generator. Therefore it is likely that the LMN-DTN complex is the source of the HD signal, acting as a CANN network, whilst vestibular input acts to couple this generative circuit to idiothetic information.

1.2.5 Other idiothetic information sources

Although the primary idiothetic information source, as discussed so far, can be thought of as vestibular, a number of other sources of idiothetic information are available to the HD cell system. When motor commands are executed, copies of these commands, known as motor efference copy, convey the intended movement to a number of brain systems. Additional information comes from the translation of the visual scene during motion, optic flow, which can provide self-motion information. Although neither of these sources of information appear to be as crucial to the HD system, or as well studied, as vestibular inputs, they have a modest effect on HD cell firing.

Motor efference copy might well be a useful source of idiothetic information for the HD system. This was initially thought to be the case on the basis that PoS HD cell firing rates were noticeably

reduced during passive rotation (Taube et al., 1990b). This can be interpreted as a dampening of HD cell firing in the absence of motor input generated by active movements, and has been confirmed in several other studies, which show a large or even total reduction in ADN HD cell firing rate of rats which are restrained by being wrapped in a towel and passively rotated (Knierim et al., 1995; Taube, 1995).

There are a number of other factors that might be responsible for the reduced HD cell firing observed in these studies, the stressful experience of being restrained for example. Zugaro and colleagues (Zugaro et al., 2001) trained rats to remain stationary while they were passively rotated. They showed that firing rates of ADN HD cells were reduced during such passive rotations. This finding confirms a role for motor efference copy as an idiothetic input to the HD cell system. However, the reduction in firing observed was on average only 27%. This suggests that previous findings were perhaps confounded by the effect of tighter restraint, and that the role of motor efference copy is more modest than previously thought.

When rats transition from a familiar environment to a novel one, preferred firing directions tend to remain stable (Taube and Burton, 1995), which is indicative of idiothetic maintenance of HD cell preferred firing directions in the absence of familiar landmarks. The role of motor efference copy in this process was investigated by an extension of this original study (Stackman et al., 2003). In the extended paradigm, rats could actively walk between familiar and novel environments, or be passively moved inside a clear plexiglass container on a small wheeled cart. Both active and passive movement could occur with the lights on or in the dark. When moving from novel to familiar locations, ADN and

PoS HD cell preferred directions were found to change by larger and more variable amounts during passive movement with no significant effect of lighting. This instability during passive transport is considered a result of a mismatch of motor efference signals with the rats changing orientation, and highlights the importance of motor input for HD cell stability in conditions where emphasis is placed on idiothetic processing.

Motor efference information might enter the HD system via NPH, being conveyed from the entopeduncular nucleus (a major motor system output in the basal ganglia) via the habenula and interpeduncular nucleus (IPN) (Winter and Taube, 2014). IPN projects to, and receives from, DTN (Groenewegen et al., 1986), and lesions of this area reduce HD cell stability (Clark et al., 2009).

The fact that lighting does not affect the results described above of Stackman and colleagues suggests that optic flow is not a significant idiothetic HD cell input. This lack of optic flow influence on HD firing has been replicated in later studies (Blair and Sharp, 1996). However, this has been contradicted slightly by later studies finding a small effect of optic flow. Rotation of a projected array of irregular luminous points, in order to generate optic flow information, was found to cause a corresponding drift of ADN HD cell preferred directions (Arleo et al., 2013). This difference might possibly be related to the stimuli used to generate optic flow (stripes in the former study, whilst the point array used in the latter is more typical of experiments on optic flow). It is possible that optic flow contributes more to HD cell stability than core firing properties. Anatomically, optic flow information enters the HD system via integration into the NPH; a structure containing AHV cells which can be driven by optic flow information (Lannou et al., 1984).

These studies reveal that both motor efference copy and optic flow are important idiothetic HD cell inputs. However, the precise nature and role of these inputs is not clear, and has not been investigated in much depth.

1.3 Idiothetic vs Allothetic control of HD responses

Effective navigation can be performed on the basis of two different types of spatial information. Allothetic environmental landmarks, of a variety of modalities, can be used to provide orienting information. In this case, landmarks would provide directional information provided that they have previously been learnt (i.e., following prior experience of a given environment). Such information might be used intermittently, for example placing a target location in relation to a known landmark. Alternatively, navigation can be based on idiothetic self-motion cues; again of a variety of modalities. This process, known as path integration, happens continuously. It is likely that these two sources of information are utilized together by navigation systems. External landmarks can be learned of in the context of an internally-derived sense of direction, but once learned can help to correct errors which might accumulate in path integration.

Whilst HD cells are influenced by idiothetic and allothetic information sources, these information streams are utilised differently. The HD system is capable of sustaining and updating activity based on idiothetic information alone. However, this process can accumulate error. Allothetic information, whilst not essential for the HD cell firing, acts to recalibrate the HD system and correct such errors. Thus, a dual process exists to ensure HD system accuracy: a continuous iterative usage of idiothetic

information regarding self-motion (path integration) and a discrete usage of allothetic information (landmark navigation). However, it takes time for the HD system to become anchored to a given allothetic information source, as external landmarks need to be established over time as a reliable source of spatial information (Knierim et al., 1995; Jeffery, 1997).

1.3.1 Path integration

Path integration can be defined as the use of internally-derived information to estimate current location and direction in order to effectively navigate (Mittelstaedt and Mittelstaedt, 1980). Specific to HD cells, path integration is the ability of HD cells to sustain and update their firing on the basis of idiothetic information alone (McNaughton et al., 2006). Experimental observations show that HD cells are able to continue to update in the dark (Taube et al., 1990b; Goodridge et al., 1998). It seems that path integration is an important property of HD cells. Indeed, this should be apparent from the studies cited above regarding the essential nature of vestibular input to the HD signal.

Taube and Burton (1995) recorded HD cells from rats navigating between a familiar arena and a novel one. They found that cell preferred firing directions remained relatively stable (randomly distributed preferred direction shifts of no more than 18°) when transitioning from the familiar arena to the novel one. Given that no familiar visual landmarks were available; this result is suggestive of a HD representation which can be sustained by idiothetic information sources alone. There also exist a number of studies demonstrating that HD cells continue to fire in the absence of visual information, such as when blindfolded (Goodridge and Taube, 1995; Goodridge et al., 1998).

Interestingly, rat pups of 2 weeks are shown to have HD cells with adult-like properties, just hours

after the eyelids have unsealed (Langston et al., 2010; Wills et al., 2010). Indeed, HD cell activity has been recorded in rat pups as young as 12 days old; three days before eye opening (Bjerknes et al., 2015; Tan et al., 2015). It should be noted however, that in the latter study HD cell stability improved dramatically following eye opening. This indicates that whilst a functional HD cell system exists without visual input, there is a clear role for this type of input in calibrating cell responses. It may be that an initial HD cell network is established on the basis of idiothetic input alone, although rat pups of this age do have access to allocentric olfactory information (Polan and Hofer, 1998; Goodridge et al., 1998).

Vestibular input, in addition for being important for HD signal generation, has been empirically demonstrated to update the HD representation. When rats explore a paradigmatic circular apparatus, HD cells maintain the correct preferred firing direction during apparatus rotations that are fast enough to be above vestibular detection threshold, but not with those that slow enough not to be detected by the vestibular system (Blair and Sharp, 1996). Passive rotation of rats in darkness results in maintenance of preferred firing directions, provided that rotations were not at a high constant velocity thereby optimally stimulating the vestibular system, for a period of some minutes (Knierim et al., 1998). Beyond this, HD preferred directions started to shift, as path integration accumulated errors. HD cell firing is observed to be heightened in a manner related to passive unrestrained rotation speed (Zugaro et al., 2002), demonstrating that HD firing is modulated by angular head velocity signals of vestibular origin.

1.3.2 Visual landmark interaction with path integration mechanisms

A number of studies suggest that visual landmarks can dominate head direction cell responses. Taube and colleagues (Taube et al., 1990b) studied the effects of cue card rotation in the paradigmatic HD cell recording session outlined in Section 1.1. They recorded from PoS HD cells of rats that freely explored a grey cylindrical apparatus. On the cylinder wall was a white card that represented the only orienting visual landmark. No other visual cues were available with the cylinder (or beyond, due to a plain black floor to ceiling curtain surrounding the cylinder). Once rats had become familiar with exploring the apparatus, they were removed. The floor was replaced (to prevent any possible olfactory orientation cues) and the cue card was rotated by a set amount. When rats were reintroduced to the apparatus, it was found that HD cell preferred firing directions rotated in a manner corresponding to cue card rotation, indicating the strong influence of visual landmarks on HD cell firing. This effect of cue card rotation has been replicated in brain areas other than PoS, including ADN (Taube, 1995), LMN (Stackman and Taube, 1998), and retrosplenial cortex (Chen et al., 1994). However, it seems that idiothetic information still plays a role. In situations where rotation of a visual landmark cue is done with the rat removed from the apparatus and disorientated (to disrupt path integration), landmark control of HD firing was complete (Taube et al., 1990b), but in situations where the rat was present during cue rotation, and thus a conflict is introduced between path integration and landmark navigation, cue control is weakened (Taube, 2007).

Although the studies cited above show a strong influence of allothetic visual landmark cues over HD cell preferred firing directions, this control was not always complete. Typically, preferred directions tend to under-rotate relative to the cue card, indicating that other information sources are conflicting

with visual landmark information. Mizumori and Williams (1993) recorded head direction cells in the dorsolateral thalamic nuclei of rats exploring a radial maze. These recordings showed preferred directions drifting in the dark, and being re-established once the lights were turned back on. However, it was also shown that directional specificity could be maintained without visual cues for a short period. This again suggests that whilst visual cues have a dominant influence, idiothetic information is an important aspect of the head direction signal.

Goodridge and Taube (1995) recorded head direction cells from the PoS of rats exploring a paradigmatic apparatus, with a single white cue card, subtending 100° of visual angle, attached to the wall. Both rat and card were removed, before the rat was replaced and explored the apparatus alone. If preferred directions were found to have rotated further than 30°, the cue card was replaced in its original position; but if preferred direction drift was below this figure, the card was reintroduced 90° away from its original position. In the majority of cases (8 of 11) where preferred directions had drifted more than 30°, returning the cue card to its original position caused an immediate shift of preferred directions back to their original values. In all cases where shifts of less than 30° occurred, reintroduction of the cue card at 90° rotation to its original position caused preferred directions to shift correspondingly. However, this shift was never exactly 90°, averaging 54° (with a range of 6 to 80°). Taken together, these results imply that although visual landmarks do seem to have a dominant role in determining preferred directions; idiothetic inputs do also play some role.

The idea of partial visual landmark control over preferred directions is supported by further experimental work (Blair and Sharp, 1996). In this study, a similar apparatus was used to the one described

above, but visual landmark information originated from vertical black and white stripes on the inside cylinder wall. Rats explored the apparatus, then the lights were turned out and the cylinder was rotated $45°$. This would change the configuration of stripes on the wall, and thus landmark information would suggest a rotation of $45°$. Recordings, made from anterior thalamic head direction cells, showed that preferred direction shifts varied across cells from 0 to $45°$. This is in agreement with the idea that both idiothetic information and visual landmark cues play a role in determining preferred directions. However, this study does not further the case for landmark cues as the dominant influence.

Another study, from Zugaro and colleagues (2000), also arrived at a similar conclusion, with the influence of visual landmarks being relatively stronger than path integration. Rats were placed in a paradigmatic cylindrical apparatus whose walls, upon which was placed a single polarising visual cue card, could be rotated separately to the floor. Whilst a rat foraged for pellets inside the apparatus, the walls and floor were rotated in various combinations by 45, 90, or $180°$. Recordings from anterodorsal thalamic and postsubicular head direction cells showed a rotation of preferred directions by approximately 90% of the rotation of either the walls alone or the floor and walls together. This suggests that visual cues dominate idiothetic cues. However, this control is incomplete, as the consistent undershoot of 10% shows a small but significant influence of conflicting idiothetic cues. This is confirmed by the result that, when the floor is rotated so that visual and idiothetic cues both indicated rotation by the same amount, this 10% undershoot is eliminated.

Although these studies suggest a mixture of landmark information and path integration determines preferred directions, and, for the most part, that visual landmarks are dominant, other studies present

a contradictory view. Wiener (1993) recorded dorsal striatal head direction cells from rats in a 60cm cubic arena containing a single polarising cue card. As the rats visited each corner in turn, via the centre of the arena, all lights were extinguished and the apparatus was rotated by a multiple of 90°. Recordings showed that preferred directions rotated with the box, suggesting that visual cues do not principally govern head direction cell firing. A further study showed that, in response to rotations of visual landmark cues, head direction cells in the posterior cortex were found to rotate their preferred directions erratically, or not at all (Chen et al., 1994). This implies that visual landmark cues exert a relatively weak influence on head direction cells. However, this study has been criticised for its use of albino rats, who are known to have poorer vision than other strains, and which would therefore artificially decrease the influence of visual landmark cues (Stackman and Zugaro, 2005). Such mixed results have also been reported in rat place cells, with some cells seemingly dominated by visual cues, and others by idiothetic cues (Wiener et al., 1995).

One explanation for mixed results as to whether visual landmarks are a dominant influence on head direction cell firing is the effect of cue familiarity. It may be the case that landmark cues experienced by rats as unstable lose their prevailing influence, and path integration mechanisms begin to dominate. Certainly this is seen in place cells (Jeffery and O'Keefe, 1999), where place fields of rats were recorded before and after conflicting motion of a cue card and a circular platform upon which the rat was standing. Initially, all place fields were controlled by the cue card. Rats could view all rotations, partially view rotations, or not view rotations at all. The cue card controlled place cell firing in all rats, but lost this control in rats that viewed cue conflicts and partially lost control in rats that only viewed some conflicts. This is interpreted as a demonstration of unstable cues losing their influence,

as results show a direct link between this loss and degree of perceived cue instability.

Although cue stability is a possible factor in cue control, other studies have challenged this theory; reporting no effect of number of cue rotations observed ((Zugaro et al., 2000). Another explanation is the effect of conflict size, that is, the distance between the directions indicated by idiothetic cues and by visual landmarks. This was explored in a study from Knierim and colleagues (Knierim et al., 1998), who recorded anterior thalamic nucleus head direction cells in rats gathering food pellets in a high-walled cylindrical apparatus with a single visual cue card. The entire apparatus, floor included, was then rotated by a small amount, $45°$, or by large amounts, 135, 150, or $180°$. This resulted in a conflict between idiothetic cues, which indicated a rotation, and visual cues, which retained their relationship to the rat and thus did not indicate rotation. For rotations of $45°$, the cue card exhibited a strong degree of control over preferred directions, but for larger rotations this control was significantly weakened. This led the authors to suggest that idiothetic cues are a primary source of information for head direction cell preferred directions, but that visual cues have an overriding corrective influence which depends on the degree of conflict between cue types.

Although the previous study could be criticised for a use of albino rats with poorer vision, which would reduce the influence of visual landmark cues, it has been replicated. Recent experimental work (Knight et al., 2014) has explored the effect of distance on cue integration in more detail, by systematically varying the amount of conflict between a visual cue and path integration mechanisms. This would allow a more nuanced finding than that obtained with only two different conflict sizes. In this study, a rat explored a circular dark environment with a single orienting distal visual cue, in the form

of a light on the perimeter of the apparatus. After 4 minutes of exploration, the light was extinguished and the rat was left in total darkness for 10 seconds. This ensured that the representation of head direction was maintained through path integration. A second light was then turned on in a conflicting location relative to that of the original light, and thus to the current internal head direction representation. This procedure was repeated, with the conflict systematically varied in 20° increments from 20° conflict up to a maximum conflict of 180°. This protocol is pictured in Figure 1.5.

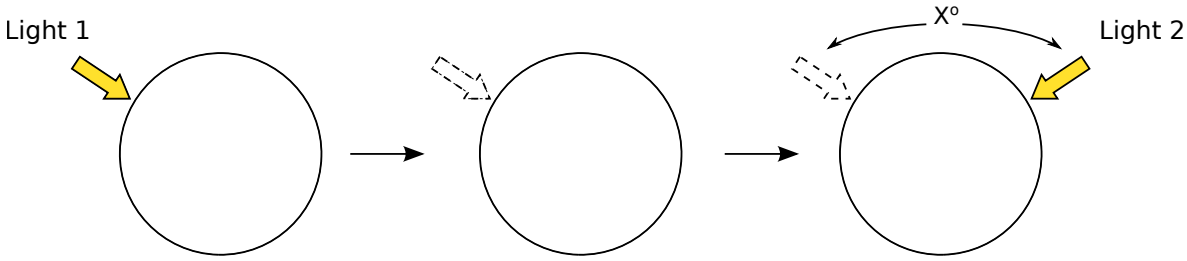


Figure 1.5: Experimental protocol used by Knight et al (2014).

Single cell recording revealed that, in response to cue conflict, cells shift their preferred firing direction in a manner dependent on conflict size. For conflict sizes up to 120°, cells shifted their firing directions by $>50\%$ of the distance between the first and second light cues. Crucially however, cells did not shift their preferred firing as if it were tied exclusively to the second visual cue, but rather undershot this location. Beyond 120° separation, cells shifted by $<50\%$ of the distance between the two cues. As separation increased, the cell shift lessened. These results suggest that the head direction cell system always uses a combination of visual and path integration cues, and the degree of conflict between visual and path integration cues determines how this conflict is resolved. This leads to head direction cell preferred direction responses being largely determined by a visual cue up to conflicts of 120°, and by path integration beyond this. These results are pictured in Figure 1.6.

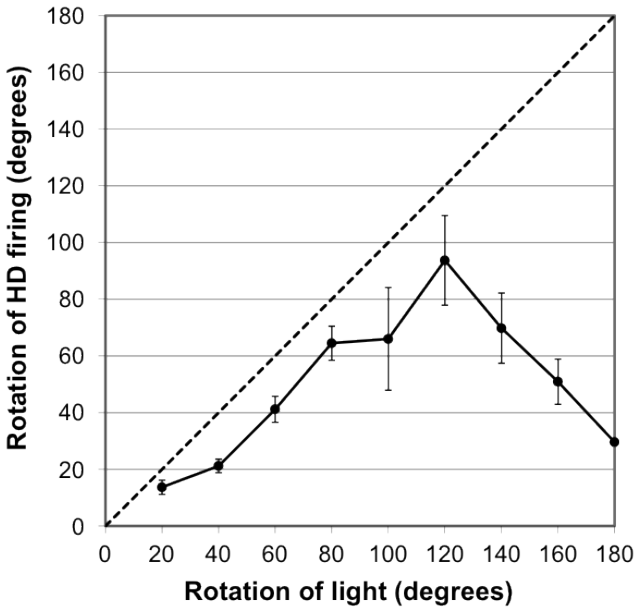


Figure 1.6: Plot detailing the relationship of expected shifts of head direction cell preferred firing directions based on conflict size given total visual cue capture (dashed line) and actual shift in cell preferred directions (solid line). Reproduced, with permission, from Knight and colleagues (2014).

These studies point to a role for visual information as a source of calibrating input for the HD cell system. Once a landmark is established, it serves to reset the system and avoid any buildup of errors in path integration. In this way, landmark navigation and path integration complement one another to ensure accurate representations of current HD.

1.3.3 Stability of landmark cues

Secondary results from the Jeffery Lab (Knight et al., 2014) suggest that, in addition to conflict size, saliency of visual cues may have an effect. Experimentally, this seems to be the case where visual landmark cues have been learned as unstable, that is, lack correspondence with ideothetic information. In cases where visual landmarks are less salient, it seems that shifts in cell preferred directions are governed increasingly by ideothetic information. Conversely, in cases where visual landmarks are very salient (taken by Knight et al. 2014 to mean visual landmarks which have previously not been

experienced as unstable with relation to ideothetic information), a visual landmark may completely dominate shifts in cell preferred directions even when at large conflict sizes.

This idea of establishment of cue control (McNaughton et al., 1991) is evidenced in the literature, with allothetic landmarks becoming, over time, stably associated with idiothetic information. If this relationship is inconsistent, landmark control is weakened. Knierim and colleagues (Knierim et al., 1995) placed rats in a paradigmatic apparatus and performed cue card rotations. Before each rotation, rats were removed from the apparatus and half were disorientated. It was found that cue control was weaker in rats that experienced disorientation, as there was no consistent relationship between idiothetic representations and the landmark cue. It seems this process takes time, and cue control has been observed to be proportional to time spent experiencing a stable landmark (Goodridge et al., 1998).

1.3.4 Proximal vs distal landmarks

Related to the idea of landmark stability is the concept of distinction between different types of landmarks. It seems that HD cells consistently prefer distal visual landmarks to proximal ones. Zugaro and colleagues (Zugaro et al., 2001) recorded HD cell responses from rats navigating within a paradigmatic cylindrical apparatus. This apparatus contained three objects, serving as visual landmarks. When the rats were removed from the apparatus and the configuration of objects rotated as an ensemble, HD cells were found to rotate by a similar amount upon reintroduction of the rat. However, if this entire experiment is repeated with the cylinder removed, such that the entire room is visible and the landmarks are now proximal rather than distal, rotation of the landmarks failed to induce a similar rotation of preferred firing directions. The authors concluded that it is only distal landmarks which are used by the HD cell system, as simple removal of the cylinder, and a subsequent shift in perspective

changing the objects from distal to proximal, was enough to abolish landmark control. Interestingly, this finding has been replicated even in situations in which place cells lose coherence (Yoganarasimha et al., 2006).

This proximal-distal distinction is intimately related to the idea of stability. As a rat navigates through the environment, proximal objects will move around within the field of vision, and therefore do not have a stable relationship with idiothetic information about current heading. The breakdown of this relationship, as described above, will prevent the establishment of cue control.

A follow up study (Zugaro et al., 2004) supports this hypothesis, demonstrating a role for visual motion signals in proximal-distal distinctions. The experimenters used an apparatus with two cue cards. One presented proximally to a rat on a central platform, and one presented distally. The cards were identical, and sized proportionally to subtend the same visual angle from the rats perspective. The large majority (30/53) of HD cell preferred responses were found to shift with rotation of the distal cue. When this experiment was repeated under stroboscopic lighting, this distal cue preference was not observed. This suggests that disruption of visual motion cues under stroboscopic lighting conditions prevents a distal-proximal distinction.

1.3.5 Origins of allocentric input to the HD system

Where might visual landmark information enter the HD cell system? Initial findings indicate that, although upstream HD area lesions do not tend to overly affect downstream areas, lesions of PoS affect visual control of HD cell responses in ADN (Goodridge and Taube, 1997) and LMN (Yoder and Taube, 2008). Indeed, contemporary theorists suggest PoS as a key entry point to the HD system for

landmark processing (Yoder and Taube, 2014). Recent work has demonstrated clearly that bilateral lesions of PoS disrupt HD cell use of visual landmarks (Yoder et al., 2015).

Any landmark utilised for directional information must be identified, assessed for usefulness, and processed in relation to the spatial configuration of the environment. Therefore, knowledge of visual landmarks potentially involves processing aspects of both the dorsal (where) and ventral (what) visual streams. PoS certainly has access to both processing streams both directly and via RSP (Van Groen and Wyss, 1990). However, lesion studies attempting to assess various sources of visual input to PoS are equivocal in their results, and thus the exact structure of the visual input pathways has not been clearly identified. It is likely that RSP is involved however, given that RSP lesions cause a moderate impairment of landmark control over HD firing (Clark et al., 2010) and disrupts learning of distal visual cues (Vann and Aggleton, 2005). Furthermore, RSP and PoS have strong reciprocal connectivity (Wyss and Van Groen, 1992).

A role for landmark processing is supported by studies of RSP in humans (Vass and Epstein, 2013). Of particular significance is the finding that RSP is sensitive to only those objects which appear to be permanent and relatively unchanging within a given environment (Auger et al., 2012). Indeed, one prominent theory of human RSP function is as a locus of translation between egocentric and allocentric reference frames (Knight and Hayman, 2014), a viewpoint which is consistent with the combinations of egocentric and allocentric information observed in rat RSP (Cho and Sharp, 2001). Given the above results that PoS is critical to landmark control of HD cell firing, it seems likely that the influence of RSP is via PoS.

1.3.6 Non-visual allothetic input to the HD system

Discussion of allothetic influences on the HD system has thus far been limited to visual information. This is because vision has been demonstrated as the most significant allothetic HD input, and the literature has been focused accordingly. However, other allothetic directional cues can influence HD firing, in particular olfactory and auditory cues. The modalities have a distinct, although modest, effect on HD cell firing. Certainly this is likely to be the case given that place cells can use a variety of sensory modalities. Place cells still develop in blind rats and are still influenced by allothetic landmarks (Save et al., 1998). Indeed, it has been found that olfactory cues in particular have a strong influence on the place cell system (Zhang and Manahan-Vaughan, 2015). These findings indicate that olfactory and tactile information provides viable information to spatial cognition systems.

One major study examining non-visual allothetic cues (Goodridge et al., 1998) looked at the effects of shifting olfactory and auditory cues on HD cell firing. Olfactory cues took the form of a series of four cotton buds spaced equally at 90° around the exterior of a paradigmatic cylindrical apparatus, with one bud dipped in peppermint extract. The existence of multiple cotton buds was to eliminate visual information from the presence of a single bud. Auditory cues were in the form of auditory clicks emitted by four speakers also spaced equally at 90° around the apparatus. It was found that, in the absence of visual information, HD cells tended to shift in the same direction as rotations of olfactory cues. However, HD cells did not shift to the same extent as in situations with visual cue rotation, suggesting that olfactory cues are secondary to visual cues. Auditory cue rotation was also performed, with the result that no consistent relationship between cue rotation and HD cell preferred direction shifts occurred. That is, HD cell preferred directions did not always shift in the same direction as

the auditory cue. However, there was a statistically significant change in cell preferred directions, suggesting that auditory cues have some influence. Given that visual landmarks are by the far the best-investigated source of allothetic spatial information and there has been relatively little investigation of other allothetic information sources, the work in this thesis focuses specifically on the visual modality. However, there is the possibility to extend the ideas contained herein to other modalities.

1.4 Relationship to other spatially-responsive cells

The directional information conveyed by HD cells is very likely a key component of spatial cognition. Successful navigation within a given environment relies on knowledge of both location and direction. Whilst place and boundary vector cells can provide information about location, and grid cells can provide a metric of distance travelled, a knowledge of heading direction is also required for navigation. Given the specificity and robustness of their directional response, HD cells could complete the picture by reliably providing a complementary neural compass to a map-like representation of the current environment provided by other cell types. Given that these cell types function in a complimentary manner, it is interesting to consider what their interrelationship might be, particularly in terms of functional interdependence.

Although the relationship between various spatially-responsive cells is not yet clearly-defined, HD cells are theorised in a variety of models to contribute to the formation of grid cells, either by oscillatory interference (Burgess, 2008; Burgess et al., 2007; Hasselmo, 2008a) or through more traditional attractor dynamics (Burak and Fiete, 2009; Fuhs and Touretzky, 2006; McNaughton et al., 2006). Grid cells in turn are theorised to contribute to the generation of place cell responses (de Almeida et al.,

2009; Hasselmo, 2008b). It has also been theorised that place cell firing is generated by HD cell firing (McNaughton et al., 1996; Touretzky and Redish, 1996; Sharp, 1999). These theories therefore present somewhat of a hierarchical system for spatial cognition, with each layer somewhat dependent on its predecessor.

In many situations, these cell types appear to be coupled. When distal landmarks are rotated, both HD and place cells shift (Taube et al., 1990b; Knierim et al., 1995; Yoganarasimha and Knierim, 2005). This has been supported by experiments showing that HD, grid, and place cells act in a co-ordinated manner following environmental manipulations (Knierim et al., 1995, 1998; Hafting et al., 2005; Yoganarasimha et al., 2006; Hargreaves et al., 2007). However, the idea of hierarchical spatial processing is not always supported experimentally.

Coupling between HD and place cells is not always consistent. In situations with conflicting proximal and distal cue rotations, HD cells follow distal cues, whilst place cells split their representation (Yoganarasimha et al., 2006). In general, it seems that whilst HD cells might follow distal cues, place cells show some variation in their proximal-distal cue distinctions (Renaudineau et al., 2007). The relationship between HD and place cells is also complicated by differences in coherence of place cell representations in CA3 and CA1 (Lee et al., 2004).

It is possible, though, that HD cell firing has some influence on place cell firing, most notably in terms of stability. Mammillary body lesions eliminate HD cell firing, but merely disrupts place cell firing(Sharp and Koester, 2008). Other lesion studies find that intact HD firing contributes to place

cell stability but is not necessary for place cell firing in general (Calton et al., 2003). Interestingly, early findings on the coherence of HD, grid, and place cell firing did show that unstable HD cells following disorientation led to an unstable place cell representation (Knierim et al., 1998). Thus it seems that, whilst important for place cell firing, HD cells are not involved in all aspects of place cell activity. This could be thought to be supported by the observation that HD cells, unlike place cells, do not show replay of sequences during sleep (Brandon et al., 2012). However, later studies recording from a great number of cells (Peyrache et al., 2015) have found HD cell replay in both ADN and PoS. Thus, we cannot conclude that place cell replay is independent of HD cell input.

Grid cells also have a complex relationship with HD cell firing. Lesions of ADN HD cells abolish grid cell firing in MEC (Clark et al., 2011). Grid cells are found to shift with HD cells during rotations of a visual landmark cue (Gupta et al., 2014). These suggest a strong HD influence on grid cell firing. However, environmental manipulations can, whilst not affecting the alignment of HD and grid cells, change the spatial phase of grid cells. This suggests that some aspects of grid cells are independent of HD cell firing. It has also been found that HD and grid cells can lose concordance in changing environments (Whitlock and Derdikman, 2012). Specifically, grid cell fragmentation was accompanied by a stable and consistent HD cell response.

Further insight into HD-grid cell interactions comes from observations of cells in MEC which are responsive to both grid and HD (Sargolini et al., 2006; Moser and Moser, 2008). It has been found that the HD and grid aspects of such conjunctive cells can be independently altered. In the above study from Whitlock and Derdikman (2012), it was found that conjunctive cells showed fragmentation of

their grid responses, but not of their HD responses. Pharmacological inactivation of the medial septum disrupts grid cell firing, and following such lesions, HD by grid conjunctive cells lose their grid characteristics, but retain their HD firing properties (Brandon et al., 2011). Abolishing ADN HD cell firing with lidocaine also abolishes grid cell firing, but not theta rhythmicity in MEC (Winter et al., 2015). Certainly then, it seems that HD cells are not influenced by grid cells, but it is likely that there is an influence of HD cells on grid cell firing. Further evidence against a hierarchical view of spatial processing comes from the fact that place cells retain, even re-establishing in novel environments, their firing properties during grid cell disruption by septal inactivation (Koenig et al., 2011; Brandon et al., 2014a).

In general, it seems that the concept of hierarchical processing, as seen so clearly in HD cell areas, does not apply to the relationship between other spatially-responsive cells. Whilst HD cells appear to an extent to be upstream of place and grid cells, especially as HD cells appear to contribute to place cell stability and grid cell firing, there seems to be a certain degree of independence of cell responses. This is certainly in keeping with contemporary views of place and grid cell interactions being complimentary and non-hierarchical (Bush et al., 2014), and of spatial information processing in the entorhinal cortex and hippocampus as parallel in nature (Brandon et al., 2014b). The precise of nature of the interaction and interrelationship of spatially-responsive cells is still being uncovered.

1.5 Behavioural influence

How important is HD cell firing to spatial processing? In many ways, this is an unexpectedly difficult question to answer. Since HD cells exist across a set of densely interconnected areas, it is hard to

make definitive statements about the effects of HD area lesions. Adding to this difficulty is the fact that many HD cell areas are not exclusively dedicated to processing the HD signal, and thus the results of lesions to these areas cannot always be ascribed to a loss of HD cell functionality. Indeed, lesions to one area might well result in damage to others. Added to this issue is the parallel nature of spatial processing (see Section 1.4 above), which implies that some spatial processing is preserved without HD cell firing. At present, the HD system is known to be important for navigation, but discovering the exact contribution is an active area of research.

Although the results of lesion studies should be interpreted cautiously, it seems that damage to HD cell areas results in spatial processing deficits. In general, results of lesion studies can be taken to show that an intact HD system certainly contributes to performance on spatial tasks, but is not the only contributor. Lesions to the mammillary body, containing the LMN, result in mild deficits in spatial memory tasks (Vann and Aggleton, 2004). However, it is likely that these deficits reflect more than damage to LMN HD cells, as most lesions have been to the mammillary nuclei as a collective unit. This intuition, that non-HD aspects of HD cell areas might contribute to observed spatial deficits, is supported by the observation that lesions of anterior thalamic nuclei (Warburton et al., 1999) result in spatial memory deficits of a more severe nature than those following lesions of ADN alone (Aggleton et al., 1996). The behavioural significance of PoS HD cells has also been tested via lesion studies. It is known that PoS lesions impair performance on some spatial tasks (Bett et al., 2012). However, other spatial tasks have not been affected by PoS lesions; potentially because PoS serves as only one output for the HD system, and PoS lesions do not affect HD signal generation (Yoder and Taube, 2014).

Although these studies show that the HD system contributes to spatial processing, they do not cover the hierarchical organisation of HD system areas. For example, ADN lesions are known not to disrupt HD firing in LMN and DTN, and thus an intact HD signal remains in these areas, and could contribute to spatial processing. It is potentially more informative to block HD cell firing in all areas. It has, however, been shown that bilateral damage to the vestibular system, known to eliminate the HD signal in all downstream areas, produces deficient spatial processing in situations where visual information is not available (Smith et al., 2010). This supports previous results as to the importance of the HD cell system to navigation, suggesting a role in processing self-movement information. This result is with the caveat that vestibular damage might affect navigation via processing in areas other than the HD system.

Further studies have attempted to uncover precisely the contribution of the HD cell system itself to navigation, with the general result that the HD system plays a role in processing self-motion information for navigation. One study that addresses hierarchical HD system information flow (Frohardt et al., 2006) tested spatial navigation in rats with either ADN (downstream) or DTN (upstream) lesions. Following lesions of each area both navigation based on self-motion and landmark navigation was impaired, indicating a role for the HD system. Additionally, impairments in navigation were more severe following DTN lesions, implying that preservation of HD firing in upstream areas following ADN lesions is of some behavioural significance.

Lesions of the mammillothalamic tract, disrupting LMN HD cells and therefore HD signal generation, have been found to disrupt performance on both food-hoarding and watermaze spatial tasks

(Winter et al., 2011). In the food-hoarding task, rats search for a food pellet and then return home. This can be done in the light with both self-motion and landmark information available, or in the dark, with this latter scenario being a test of self-motion cue usage. Lesioned rats were impaired in their use of self-motion cues to return home, but were able to use landmark cues as well as controls. In the watermaze task, rats swim to locate a hidden goal platform. Lesioned rats were impaired at learning the platform location. This can be interpreted as in impairment in the usage of self-motion cues, and presumably reliance instead of environmental landmarks to find a platform. Once a location is learnt, lesions did not affect performance. This is a slightly confusing finding, as self-motion cues would be expected to still have an effect. One explanation might be that the HD system does not affect usage of pre-learned trajectories. Usage of landmark navigation following damage to the HD cell system is evidenced by the results of another watermaze study (Stackman et al., 2012), which found that muscimol inactivation of ADN forced rats to utilise distal landmark cues, rather than proximal ones, for navigation. However, this has been contradicted by findings that DTN lesions impair the use of landmarks during watermaze navigation (Clark et al., 2013). This slightly contradictory result might stem from the different strategies rats have for making use of landmark information it is possible to find a goal location by moving in a specific direction from a landmark, or by using landmarks to locate goal at a particular place on a cognitive map. Thus Clark et al. (2013) argue it is place landmark navigation that is impaired, but not directional landmark usage. These differences might also arise from the fact that ADN HD cell inactivation spares upstream HD cell firing.

HD cell firing is also absent when rats are navigating in the inverted plane (Calton and Taube, 2005). In these situations, it has been found that rats are unable to successfully find an escape hole from

a novel location(Gibson et al., 2013). This deficit did not extend to using a pre-learned route to an escape hole from a familiar location, an ability which was spared. This, it is proposed, is because the HD signal is used during navigation requiring flexible usage of a cognitive map, acting as a continuous signal to help guide such behaviour (Winter and Taube, 2014), as opposed to associative spatial learning of specific trajectories towards targets (Valerio et al., 2010). This is supported by studies that show HD cell preferred firing direction shifts are consistent with errors in flexible navigation (Valerio and Taube, 2012) and could explain the contradictory findings as to the effects of HD system lesions on landmark navigation. It is likely that several different strategies are available to the rat, and it has been shown that, following vestibular lesions, rats are able to perform spatial navigation only when provided with visual landmarks (Stackman and Herbert, 2002). It is possible that such a landmark-based strategy can be utilised once cue control is established (Dudchenko and Bruce, 2005). Given that landmarks must be learnt with respect to idiothetic information, the HD cell system could plausibly be involved in the establishment of cue control. However, it is also possible that familiar landmarks can guide navigation independently of the HD cell system, possibly relating back to the directional vs place landmark strategies discussed by Clark and colleagues (2013). A test of this hypothesis would be to disorient rats during training, to prevent cue control, and see whether landmark navigation is intact following vestibular lesions or HD system inactivation.

One interesting theory that might make explain unclear results for an HD system role in navigation, and that accounts for potential mnemonic roles of the HD system comes from Knierim and Hamilton (2011). They hypothesise that the HD system guides navigation indirectly, by setting the orientation of the grid and place cell systems and disambiguating boundaries based on direction. The authors sug-

gest that navigation is therefore preferentially via mapping strategies rather than directional strategies.

Although the studies discussed thus far have concerned the HD systems role in navigation, it is also theorised to play a role in some mnemonic processing. This is consistent with episodic memory impairments following lesions to human HD cell areas such as LMN and ADN (Aggleton, 2005). The exact theorised role differs. Some researchers maintain that the HD system contributes information on orientation and velocity during memory encoding and retrieval (Hasselmo et al., 2010), whilst others suggest it is important for associating egocentric information with an allocentric representation of space (Byrne et al., 2007). Indeed, it may even be that the HD system role in memory extends to a general role for hippocampal memory systems in mental imagery (Bird and Burgess, 2008). The role of the HD system in mnemonic tasks is logical given theorised HD system role in formation of the cognitive map, and the usefulness of such a spatial map for episodic memory. However, a full review of the literature on the mnemonic functions of various spatially-responsive cells is beyond the scope of this thesis.

In summary, it is clear that the HD cell system is an important component of spatial cognition. However, spatial processing can be performed using a number of different strategies (e.g., spatial association, landmark navigation, path integration), and can involve a distributed set of brain areas processing spatial information at least partly in parallel. The contribution of HD cells to navigation is principally centred on using a cognitive mapping strategy, with the HD cell system providing a neural compass to guide spatial behaviour and allow flexible navigation. Specific lesions studies have indicated that the HD system is a vital component of a system for path integration, but also for some types of landmark

navigation. However, navigation based on pre-learned routes and trajectories may not be dependent on HD cell firing, and it could be that landmark usage with reference to a cognitive map is particularly affected by HD cell area damage. Precisely how the HD system can guide navigation is still under investigation. Additionally, it has been theorised that the HD system plays a critical role in spatial memory, as well as more immediate spatial processing. This HD contribution to spatial memory is still being clarified. It is hoped that this summary has provided an overview of the range of factors and strategies to take into account when considering HD cell contributions to behaviour via broader cognitive systems.

Chapter 2

Modeling Head Direction cells

This chapter summarises the computational modelling literature to date, as it applies to understanding the HD cell system. However, before this is done it seems appropriate to explain what exactly, in the eyes of this author, modelling does and does not do. A computational model should principally be thought of as a tool for hypothesis refinement. Indeed, it can be considered an instantiation of a current hypothesis regarding the relationship between cell firing properties, inter-cell connectivity, and function. Assuming that a simulation achieves brain-like function in the same manner, we would of course have a greater ability, in comparison to a real brain, to pick it apart and understand it. This ability to examine in detail how the physical aspects of a model (e.g., neuronal responses and connectivity) relate to function is a crucial aspect of computational modelling. With these facts in mind, several statements are presented below. It is hoped that these statements and the following discussion make clear the perceived role of the modelling work contained in this thesis.

1. A model can only be designed to be as close to the real brain as our understanding of the brain permits.

2. Unexpected results may occasionally arise from a model, despite our understanding of its construction.

3. We can use our ability to pick apart simulated systems in detail to gain insight into model behaviour, regardless of whether simulations results are expected or not.

4. The better we understand a model, the better we can attempt to explain its behaviour.

5. The process of interating a model towards brain-like function and architecture moves our understanding of the brain itself forward.

6. Computational models provide a valuable context in which to make sense of experimental findings from different neuroscience sub-disciplines.

We can apply these principles to modelling the HD system. First, we decide which aspect of HD cell firing properties we wish to explore, and create a hypothesis based on what we already know. We then utilise the existing experimental literature to understand how to structure our model to replicate the known neurophysiology, responses, and inputs of the HD cell system. Simulation will then provide us with some insight into this particular aspect of the HD system, and in turn will generate ideas for future empirical investigation. An abstract representation of this process is given in Figure 2.1.

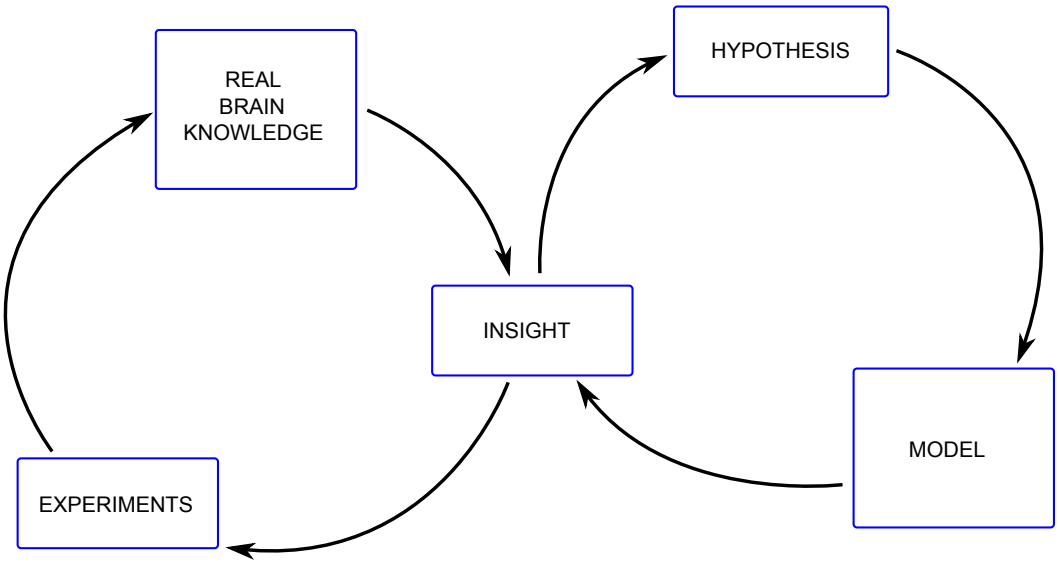


Figure 2.1: A graphical representation of the relationship between empirical and theoretical investigation of the brain. We start with some knowledge about the brain. This represents the sum of the factual statements we are able to make about the brain. This collection of factual knowledge is tied together, allowing us to have some insight into how the brain works. This can be thought of as our understanding of the brain. Using this insight/understanding, we can perform further experimental investigations, which will hopefully grace us with further factual information. In this way, insight guides experiments, which generate real brain knowledge. It should be noted that experiments are of course hypothesis-driven, but this detail is omitted for the sake of clarity. Insight into the brain also drives the formation of hypotheses at the level of computation, directed towards the formation of models. Simulation results from these hypothesis-driven models will provide us with some insight into the brain. This in turn will allow us to perform further experiments, or to generate further models. In this manner, it is the interrelationship between models and experiments that progresses our insight into the brain. It should be noted that models themselves cannot lead directly to real knowledge about the brain, other than indirectly by guiding experimental work.

When applying the modelling process, it is useful to limit exploration to a single, often quite basic, question. This can be related to points 1, 2, 3, and 4 listed above: it is easier to accurately construct and understand the behaviour of a model whose focus is tighter, as this avoids being overwhelmed by complexity. In some cases, it is necessary to simplify a model in order to be able to understand it enough to gain insight. Once we have a firmer theoretical footing, we can introduce more complexity and refine our hypothesis. This is often a balancing act, as the model must be close enough to the real HD system to provide insight, yet limited enough to be manageable.

The ability to separately consider individual questions within the HD system without needing to tackle everything at once is of great benefit to the computational approach. In this way, we can expand the scope of simulation work, either on the basis of what we already know or have learnt from the previous model iteration, from a position of clearer understanding. It is this iterative process of model refinement, based on empirical evidence, which helps to progress our theoretical understanding of the brain. In this sense, all models are wrong to some degree. In the process of understanding exactly how and why a model differs from reality, we gain knowledge about the brain itself.

Much like hypothesis-building, computational modelling is best suited to explaining one question at a time, on the basis that it is better to understand a smaller aspect of the brain thoroughly than to understand bigger questions incompletely. This hypothesis-refinement approach is central to this thesis, and the work contained herein. The models in this thesis will be strictly focused, and will be accompanied by a discussion of how we might take insight gleaned from them and guide future work, both experimental and computational.

2.1 Continuous Attractor Neural Networks (CANNs)

From close to the beginning of HD cell research, the dominant theory has been the attractor hypothesis (Skaggs et al., 1995), which states that the HD cell system can be modelled using a Continuous Attractor Neural Network (CANN). A CANN is pictured in Figure 2.2. It consists of a ring of HD cells, arranged by similarity of preferred firing direction (i.e., adjacent cells have the most similar preferred firing direction). Cells are reciprocally connected to one another by recurrent collateral connectivity, with connection strengths also based on similarity of preferred direction. That is, the more similar

the preferred firing directions of the pre and postsynaptic cells, the stronger the connection. Formally, this can be expressed as the weights between cells being a Gaussian function of the distance between their preferred firing directions

$$w_{ij} = e^{-(s_{ij})^2/2(\sigma)^2} \tag{2.1}$$

Where w_{ij} represents the synaptic weight from presynaptic head direction cell j onto postsynaptic head direction cell i and σ is the standard deviation of the Gaussian profile. s_{ij} is the difference between the preferred head directions of the post- and pre-synaptic head direction cells i and j respectively.

It should be noted that the topographical ring arrangement is for the sake of simplicity of visualisation: in reality there appears to be no such physical arrangement of HD cells in the brain. The crucial aspect of CANNs is the connectivity, and there is no constraint on the physical arrangement of cells within a CANN, as long as connectivity is based on preferred firing directions. One other caveat of CANN models is that they must incorporate some aspect of inhibition, in order to balance the effect of excitatory recurrent collateral connectivity.

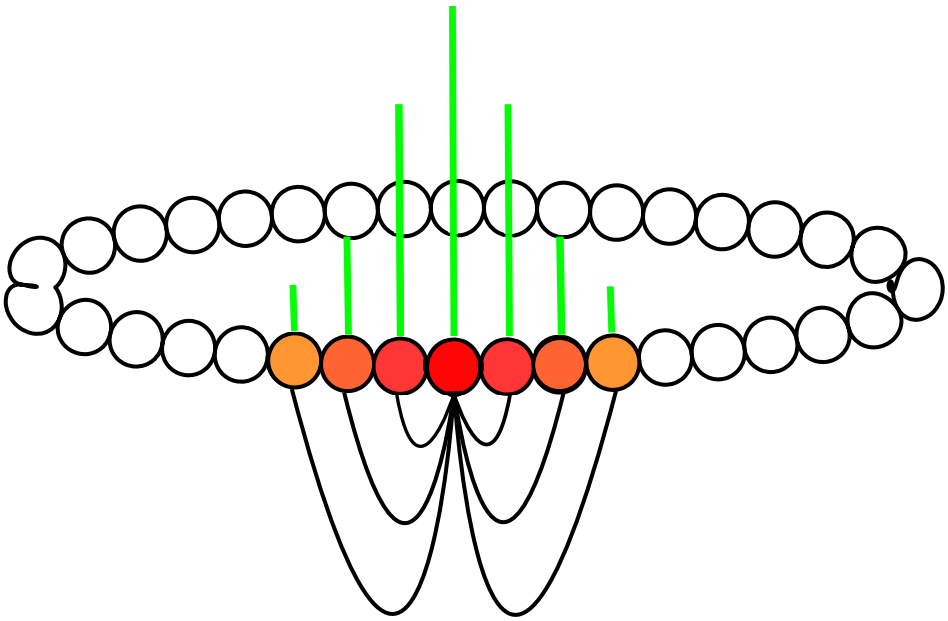


Figure 2.2: Stylised CANN architecture. Weight of the connection from the central cell (red) onto other cells ranges from high (red colours) to low (orange colours). For clarity, only partial connectivity is shown. In reality, every cell is connected to each other cell. If this central cell is stimulated, it will drive all other cells in the layer, with projection strength, and therefore the strength of this input, based on proximity. Given that every cell excites all others, some inhibition must be included to avoid all cells becoming active. With the inclusion of this inhibition, the stimulation to cells at a certain distance to the central cell will not be driven enough to overcome this inhibition and a gaussian packet of firing (green lines) will be formed.

CANNs have a number of properties that make them suitable as a model of HD cell systems. Firstly, they are capable of indefinitely and stably sustaining a Gaussian packet of activity at any point around the ring of cells, by virtue of their specific connectivity. Given cells which have overlapping Gaussian tuning curves and which have preferred directions broadly distributed in the range $0 - 360°$, as is the case with HD cells, this activity packet can be situated anywhere in the HD ring to represent any true head direction.

The response of CANNs to external inputs is of critical importance. If a packet of activity receives any asymmetrical external input, it will shift within the ring. This latter point allows for theoretical discussion of path integration to be focused on the concept of some idiothetic input, possibly of vestibular

origin, acting to update the current HD representation. For example, an activity packet would be stable at a location representing current allothetic head direction, so long as there is no movement of the head. When the head does turn, idiothetic signals, most likely the output of AHV cells, will provide asymmetrical input to the activity packet. This will act to shift the activity packet through the ring, thereby tracking true HD.

It seems, then, that CANNs are able to account for key aspects of the HD cell system: continuous representation of head direction with a discrete set of cells, ability to sustain a stable activity packet in the absence of external input, and influence of idiothetic input in path integration. This has led to a contemporary view of CANNs as a viable HD system model (Touretzky, 2005; McNaughton et al., 2006; Taube, 2007; Clark et al., 2012).

2.2 Competitive Neural Networks

Some models in this thesis (see Chapter 6) make use of an additional neural network architecture, a competitive neural network. A competitive network can be thought of as a self-organising system which learns to categorise inputs. Initially, connectivity between input and output neurons is randomly pre-specified. An input is fed in and generates, by virtue of the connectivity, a random pattern of output firing. However, output is modulated by some competition between neurons, usually an approximation of lateral inhibition between output neurons mediated by inhibitory connections. This inhibition results in a situation in which those neurons firing most strongly will inhibit other neurons with weaker firing. This typically means that a subset of neurons in the output layer will fire in response to an input pattern. Associative learning then takes place between the input pattern and the

competitively-enhanced output firing pattern.

Over multiple presentations of different inputs to the network, the synaptic weight structure will act to categorise inputs. Input patterns that are similar, i.e., that activate the same or close to the same set of input neurons, will result in the same output pattern firing. Inputs that are dissimilar, activating different sets of input neurons, will result in different patterns of output firing. It should be noted that the nature of this categorisation is not specified by the experimenter. Instead, the system learns for itself to categorise inputs. In Chapter 6, cells within a competitive layer learn to respond to particular combinations of the HD and AHV inputs.

2.3 Requirements of HD cell models

In addition to core HD cell functions, CANN models should strive to adequately explain several other key aspects of the HD cell system.

1. **HD system structure:** The effects of lesion studies, and the hierarchical dependence of various HD regions. This also extends to the effects of vestibular lesions, and the dependence of the HD cell system on vestibular input to remain stable and functioning in the dark. In particular, models must account of the bursty firing of LMN/DTN HD cells following loss of vestibular input (Clark et al., 2012).

2. **Cell responses:** The variety of cell responses within the component HD system areas should be accounted for. In particular, the responses of AHV cells ought to be addressed. What is the functional significance of symmetric, asymmetric, fast, and slow AHV cells? Any model should also account for the existence of cells responding to combinations of AHV and HD. In

addition, the responses of simulated neurons should accurately reflect those of cells found in vivo.

3. **Multimodality:** The relationship of allothetic and idiothetic information sources, specifically visual and vestibular information. In particular, the process of visual landmarks gaining control over HD cell preferred firing directions and the dominance of visual input over idiothetic information once this control has been established. Any explanation of this process must include the idea of distance-modulated resolution of conflicts between the two information sources (Knierim et al., 1998; Knight et al., 2014). In addition, it must be explained how HD cell preferred firing directions are able to maintain the angular separation of their preferred firing directions (isomapping) following updates of the HD system.

4. **Path integration:** Not only should a CANN explanation of path integration address how time-accurate path integration might be achieved, but also how errors could accumulate in this process.

5. **Self-organisation:** The development of HD cell responses. In particular, how the HD system self-organises to produce the firing properties observed, and how these firing properties relate to connectivity. Such an explanation should account for how a CANN might be set up within the brain, and account for the development of cell preferred firing directions relative to one another.

6. **Recurrent collateral connectivity:** Although CANNs show reciprocal connections within a layer, there is a lack of evidence for such connectivity within HD cell regions in the brain. Thus, any CANN models of the HD system must make a comment on this situation, either through spreading reciprocal connectivity across multiple HD cell layers or by avoiding recip-

rocal connectivity of any kind.

7. **Effects of vestibular lesions:** As mentioned above (Section 2.1), CANN architectures have the ability to sustain activity in the absence of external input. However, the effects of vestibular lesions show that vestibular input is vital for HD signal generation. Any theoretical account of the HD system should incorporate this fact. Why might HD cells become bursty cells following vestibular lesions?

   Clarity: Given the above-stated purpose of modelling, a key requirement of any useful HD system model is clarity of explanation. It is vital to clearly outline both the hypothesised system proposed and any predictions which might be tested experimentally.

2.4 HD cell models

The following section provides a summary of the major HD system models from the last 25 years. In many cases, it might seem that descriptions of previous simulation work seem somewhat critical. This should not be taken as an attempt to undermine previous computational models: it is only by being frank about the limitations and assumptions of simulations that we make further our understanding of the brain. Indeed, limiting the scope of a model is often advisable, as it allows fuller explanation of specific aspects of HD function without muddying the waters by addressing several issues at once. In laying out what issues previous models have and have not addressed, it is hopefully made clear what the work in this thesis contributes to the field. In the interest of fairness, the Discussion (Chapter 7) gives the same treatment to the simulation work carried out in this thesis. The majority of these models omit at least one of the above criteria. Many of these models are concerned principally with the behaviour of a mature system, and thus do not address questions of developmental trajectory

such as self-organisation. The majority of these models make use of direct within-layer HD to HD connectivity, which is of debatable biological plausibility (Chapter 5 demonstrates the computational advantage of avoiding this type of connectivity).

2.4.1 McNaughton, Chen and Markus (1991)

The first model of the HD cell system, from McNaughton, Chen, and Markus (McNaughton et al., 1991), was not based on CANN architecture. The authors proposed a path integration system whereby combinations of current HD and AHV are associatively mapped onto a future HD after a given time interval. This concept of heading updated via AHV signals has been hugely influential, and has been incorporated into a number of later path integration models. The authors go on to specify that this linear mapping would be best achieved by utilising cells which respond to both AHV and HD, a prediction that has been borne out in the experimental literature (Taube, 1995; Stackman and Taube, 1998; Blair et al., 1998; Bassett and Taube, 2001; Sharp et al., 2001). By utilising these combination cells, the need for mathematical integration within the brain is replaced with associative learning. Also presented was the idea of visual landmarks as a mechanism for correcting accumulated errors in path integration. This was theorised to be achieved using local view cells, which respond to particular views of objects from a specific location and heading. Such local views could over time be associated with HD, and their influence could be overriding.

Although these basic ideas have proven in time to be hugely influential, the specifics of the implementation require further thought. The model requires a rather specific weight profile, yet does not address how such a profile of connectivity would arise through learning. A second issue is how exactly a model such as this would cope with multiple speeds of path integration. This is also related to the

authors lack of consideration of cell response properties: HD, AHV and AHV by HD cells in the brain have tuning curves, rather than one specific response value as in this paper. Indeed, the model does not account for observed cell tuning curves, nor does it account for the various AHV cell subtypes. The idea of visual landmarks being learnt over time is presented but there is no nuanced discussion of how visual-vestibular discrepancies might be resolved. Additionally, there is no relationship between this model and the hierarchical structure of HD areas; and no attempt is made to explain the functional significance of HD cell presence in multiple brain areas. Interestingly, this paper implies that it is vestibular information which allows for sustained HD cell firing in the dark. Whilst this would account for the effects of vestibular lesions abolishing the HD signal, it would not explain the appearance of bursty HD cell firing following such lesions.

2.4.2 Skaggs et al. (1995)

The first usage of CANN models was in the attractor hypothesis of Skaggs and colleagues (Skaggs et al., 1995); a refinement of McNaughton et als (1991) initial model. In this newer model, HD cells exist as a CANN, with ordered reciprocal connectivity. This allows the model to replicate those aspects of HD cells reflected by a CANN network. Included in the model are two layers of rotation cells, which fire to combinations of HD and input from AHV cells. HD cells project to rotation cells, which project back to the HD layer with an offset. During head turns, rotation cells become active via combined HD and AHV input and stimulate the edge of the HD activity packet via their offset connections. This causes the HD activity bump to shift around the ring, achieving path integration. Rotation cells are divided into clockwise and counter-clockwise populations, each being active for one direction of rotation.

Whilst this model has proven theoretically influential, especially as most following models contain some of its essence, the authors did not provide simulation results and the model was not formally described with appropriate equations. This means that it is hard to assess, although it seems this models chief weaknesses lie more in omission than incorrect inclusion of details. The majority of issues present in McNaughton et als model persist in this update from Skaggs and colleagues. However, with the addition of within-layer HD to HD connectivity, the system is now capable of sustaining itself in the absence of vestibular input. This is in disagreement with the experimental literature. Additionally, such within-layer recurrent collateral connectivity is likely to interfere with packet motion during path integration. No mention is given still of how such a system may self-organise, nor how such a system may cope with multiple speeds. The full complexity of AHV cell responses is not considered. Although there is a theorised role of visual input, this is not adequately explained: proximal-distal distinctions are made of the basis of edge detection, rather than distance as is observed experimentally.

2.4.3 Zhang (1996)

The first mathematical formulation of a HD cell model (Zhang, 1996) yielded simulation results detailing how a CANN model can both maintain a stable activity packet, and shift this activity packet without deformation. Whilst a rigorous mathematical formulation is of value for understanding the HD system, this appears to be at the cost of biological realism. The author suggests for example, that the activity packet could be shifted by an instantaneous change in the recurrent weight profile; without making clear how this might be achieved in a biologically plausible manner. However, it is mentioned that an effective asymmetry in the weight profile might be introduced from an external source such as AHV and rotation cells as described by Skaggs et al., (1995), and this would sidestep the need for instantaneous weight changes. Zhang suggests an error-based learning rule, and demonstrates how

the HD system might converge on the correct weight profile using such a rule. However, it is not clear how this error might be computed in vivo. Whilst information is provided on the conditions under which packet motion might perfectly match or even anticipate true HD changes, this not related back to the biology of the HD system. This lack of biological realism also extends to a lack of explanation for why multiple HD areas exist and how these areas interact. In general then, this model valuably clarifies the core mathematical principles of CANN-based HD cell simulation in an idealised theoretical system, which is likely to need to be expanded to fit the known biology of the HD system.

2.4.4 Sharp et al. (1996)

Sharp and colleagues (Sharp, 1996; Blair and Sharp, 1996) attempted to capture the anticipatory nature of some HD cell regions. Their model consists of a population of PoS HD cells, a population of ADN HD cells forming a CANN, a population of AHV cells, and a population of AHV and HD combination cells. ADN HD cells excite their neighbours directly, as well as exciting AHV by HD cells. The latter project offset inhibitory connections back to the ADN HD cell layer. Thus ADN HD cells indirectly inhibit their neighbours via the AHV by HD cell layer.

When the head is not rotating, the ADN HD cell activity packet will remain stable due the balance of excitation and inhibition within the ADN layer. When the head rotates in a given direction, AHV cells for that direction become active, in turn driving up AHV by HD cells for that direction. This will shift the balance of inhibition and excitation in the ADN HD cell layer, causing the activity packet to move. Also included is an AHV cell that fires in inverse proportion to AHV regardless of direction. This cell inhibits the ADN layer, with the result that propagation speed in the ADN layer will be influenced by absolute speed of head rotation, as faster rotations will reduce this inhibition (and slower speeds will

increase inhibition). ADN HD cells project to PoS HD cells, driving them up. However, this drive is delayed by a combination of postsynaptic summation time and conduction delay.

This model makes inroads towards explaining path integration in a manner more specifically related to known HD system characteristics. However, it does not incorporate HD cell tuning curves, thus resulting in a square-wave activity packet rather than the expected Gaussian shape. Although it addresses the idea of anticipation of head direction in terms of between-areas anticipation, it does not elaborate on why ADN HD cells might anticipate true HD. Path integration is achieved, but only over a narrow range of velocities. Although it is not an explicit aim of the model, no aspect of visual/allothetic input is incorporated. Additionally, AHV cells are not modelled in their entirety, with the AHV cell representation missing in particular symmetric cells. Whilst these are more failures of omission, this model also makes the crucial error of depending on within-layer reciprocal connectivity between ADN HD cells, which at present is thought unlikely to exist in the brain. As with previous models, a specific configuration of weights is required for optimal model functioning; yet no attempt is made by the authors to explain how such connectivity might arise or be maintained.

2.4.5 Redish et al. (1996)

Redish (Redish et al., 1996) asserted that previous models, excluding the work of Zhang (1996), did not utilise cells with realistic HD tuning curves and thus functioned in a discrete manner rather than representing continuous HD space. This, he proposes, limits such models to a small range of velocities, as slower rotations require more cells to represent space on the finer scale required. This also limits processing, as more intermediate cells need to become active within the same time interval for slower rotations.

His solution utilises cells with biologically-accurate tuning curves. The model consists of two CANNs, one in PoS and one in ADN. PoS projects to ADN with symmetrical, right-offset, and left-offset connections. When the head is not rotating, symmetrical connections sustain an activity packet in the same position. When the head is rotating, offset connections of the correct direction are strengthened. This provides asymmetric input to the ADN activity packet, which shifts. This will in turn shift the PoS activity packet, with ADN leading PoS. As in Sharp et al., (1996), input from mammillary bodies provides inhibitory modulation to the ADN based on absolute speed of head rotation.

This model is able to accurately match real HD cell data during path integration, incorporating realistic HD cell tuning curves and addressing the multi-region nature of HD cells. However, it has a number of weaknesses. Both the theorised PoS and ADN networks require within-layer recurrent collateral activity, for which there is currently no evidence. The authors acknowledge this, but do not propose how attractor-like behaviour might be generated in the absence of such connectivity. Although the modulation of offset connectivity on the basis of AHV allows for accurate path integration, the authors do not propose how such online changes in weight, both at the start and end of head rotation, might take place in the brain. Indeed, no mention is made of how the configuration of weights between areas might arise through learning. Furthermore, this model does not adequately address why PoS HD cells are demonstrated by lesion studies to depend on intact ADN. This is seemingly at odds with the assumption of this model that PoS forms its own HD cell CANN layer, and does not therefore adequately address the hierarchical dependence of HD cell regions.

2.4.6 Goodridge and Touretzky (2000)

Goodridge and Touretzky (Goodridge and Touretzky 2000) created a model which consisted of PoS and ADN HD cells, as well as LMN cells responding to combinations of AHV and HD. PoS acts as a CANN, sustaining a packet of activity representing current HD. AHV by HD cells in the LMN receive HD information from PoS, as well as AHV input of unclear origin, and project with offsets to ADN cells. ADN HD cells are not modelled as a CANN. Path integration is achieved by offset LMN to ADN shifting the ADN activity packet, which in turn shifts the PoS activity packet.

This model achieves realistic cell responses and connectivity, as well as good path integration tracking ability. Interestingly, the lack of ADN attractor properties allows for the simulation of experimentally-observed modulation and distortion of the ADN activity packet during path integration. Also replicated is the anticipatory nature of ADN relative to PoS. However, many of the criticisms applied to previous models are still applicable to this updated model. PoS still contains within-layer recurrent collaterals, and the specific inter and intra region connectivity in this model is not considered from the perspective of self-organisation. As with many HD cell models, AHV cells are not accurately replicated in terms of the variety of AHV cell responses. Whilst ADN and PoS have the correct temporal relationship, LMN, which is observed in vivo to anticipate HD firing in ADN, does not show anticipatory firing.

Zeidman and Bullinaria (2008) present an extension based on this work. They extend the model in two ways. Firstly, they add optic flow information as an input to the system. Secondly, they allow for object-recognition input to PoS, allowing for visual landmark control of PoS HD cell firing. These

additions allow the system to accurately track HD. Interestingly, this work addressed the idea of multimodal integration in the HD system, with the authors suggesting that optic flow information is of less importance when in conflict with vestibular information. However, this model retains almost all of the previous issues from Goodridge and Touretzky (2000). Additionally, the consideration of optic flow was somewhat limited, being transmitted directly to ADN. In reality, this optic flow information appears to be of less importance to the HD cell system, and is likely integrated into AHV cell firing in NPH. Motor efference copy, another HD cell input of similar importance to the system as optic flow, is omitted from the model altogether.

The idea of conflict, too, is limited in that it appears to be based on parameter tuning rather than some biologically-relevant mechanism: conflict behaviour is simply based on the relative weightings of the two inputs. The focus on robot implementations of HD models draws the discussion away from biology and more towards simple parameterised reliability weighting. In reality, cue conflict is more nuanced (see Knight et al., 2014 for an experimental investigation of visual-vestibular cue conflict in the rat HD system). Similarly, conflict between landmark and vestibular information is not addressed, despite both inputs being present in the model. The authors simply note that visual information should be trusted less, but it is actually more likely the case that this trusting process is dynamic and based on the stability of landmarks relative to idiothetic information. That is, landmarks can be learned as stable, and then later override conflicting vestibular inputs. A discussion of cue integration necessitates some explanation of how cues gain their relative weightings, and how these weightings might change dynamically.

2.4.7 Rubin et al. (2001)

In the first model to branch away from excitatory recurrent collateral connectivity, Rubin et al. (2001) presented a model of thalamic HD cells which is capable of sustaining an activity packet without the need for excitatory recurrent collaterals. Their model consists of a network of thalamocortical relay (TC) and thalamic reticular (RE) cells. RE cells fire due to excitation from TC cells, which are themselves inhibited by RE cell firing. TC cells then fire through postinhibitory rebound (an intrinsic property whereby neurons depolarise once a hyperpolarising input has been removed), with inhibition being thus converted into cell firing.

This model is therefore extremely valuable as it demonstrates how the attractor-like property of stable and self-sustaining neuronal firing can be attained without the need for excitatory recurrent collaterals. However, realistic HD cell tuning curves are not reproduced, with the activity profile representing a square wave rather than a smooth Guassian. No attempt is made to explain how this novel system might shift the activity packet as required for path integration, although the model does demonstrate that external input can act to shift an activity packet in this network architecture. The authors acknowledge that the model does not address in detail the effects of vestibular and visual cues, or the interaction of the two, nor does it address how ADN might interact with other HD areas. Indeed, some input is required to initialise simulation, but no input source is suggested. By extension of these facts, the model is unable to address the effects of vestibular lesions on downstream HD areas. Although the post-inhibitory rebound mechanism is novel and interesting, it is of uncertain status in vivo (Song and Wang, 2005). Regardless, this paper demonstrates how attractorlike behaviour could exist in the absence of true CANN architecture.

2.4.8 Arleo and Gerstner (2000; 2001)

Arleo and Gerstner (Arleo and Gerstner 2000, 2001) developed a multi-layer model of the HD cell system, which was then implemented on a robotic platform. The model is based on populations of ADN, LMN and PoS HD cells. Both allothetic and idiothetic information is used to update the HD representation, but the role of idiothetic input is taken to be primary. A population of AHV cells contains three cells, one firing proportionally to clockwise AHV, one firing proportionally to anticlockwise AHV, and one firing only when the head is not rotating. LMN HD cells respond to conjunctions of HD and AHV. ADN cells respond to HD with an anticipatory effect during rotations. PoS HD cells are driven by ADN cell output, but also receives allothetic visual information. The model is capable of accurately tracking HD during movements of the robotic agent. Interestingly, the authors discuss the idea of calibration and how visual information must be learned with respect to idiothetic signals.

In the model, idiothetic information is projected directly to ADN and LMN. In reality, this information is likely to enter the HD cell system via DTN, a region which is not included in this model. Due to the omission of DTN, the nature of the LMN-DTN complex and its criticality to HD signal generation is not accounted for. Almost all of this models connectivity and cell firing properties are pre-specified. All HD cells have intrinsic preferred directions, and complex inter-connectivity. Thus, the model lacks significant explanatory power as to how such connectivity might develop, and how it gives rise to the observed HD network firing properties. Whilst much of the models operation depends on ADN anticipation of the HD signal, this is an inherently pre-specified property of simulated ADN cells. Although LMN in this model does not anticipate true HD, this is known to be case in vivo, with LMN even anticipating ADN firing. Another major shortcoming of this model is the assumption

that there is reciprocal interdependence of the place and HD cell systems. In reality, the relationship between place and HD cells is complex, but there is likely a degree of parallelism inherent within these systems.

2.4.9 Xie et al (2002)

Xie and colleagues (Xie et al., 2002) proposed that previous models of the HD system (e.g., Goodridge and Touretzky, 2000) had used unrealistically fast time constant values in order to achieve good path integration accuracy. Xie et als model consists of two rings of cells with a specific connectivity: each cell receives offset connections within the same ring, and connections with an opposite offset from the other ring. When the head is not rotating, these connections balance one another. During head rotation, the two rings receive different AHV inputs, unbalancing the pattern of firing across the two rings and causing a shift in activity packets.

This model is an interesting approach to simulating HD cell path integration. It achieves path integration using biologically realistic time constants, and generates anticipatory responses as seen in ADN. However, the requirements for both sustaining and shifting an activity packet are very strict in terms of both parameters and connectivity. A stationary activity packet when the head is not rotating requires perfectly balanced connectivity between and within the two rings. This is troubling given that no explanation is given for how such connectivity might develop through learning. Another limitation of this model is that visual input is not considered, although the authors note that this is a potential avenue of further research. There is, as is common to many models, a lack of accurate AHV input. The model requires within-layer synapses, although connectivity is offset, in the ADN. This is in contradiction with a lack of experimental evidence for such connectivity.

2.4.10 Stringer et al. (2002)

The first model to address the issue of self-organisation of the HD cell system (Stringer et al., 2002) demonstrated how CANN architecture could be generated through associative learning. The model consists of a population of HD cells, receiving input from rotation cells signalling clockwise and counter clockwise head rotations. A trace learning rule was employed, whereby learning takes place between memory traces of both pre and postsynaptic firing. It was found that appropriate connectivity could self-organise based on visual input in order for a stable HD cell activity packet to be sustained, and for input from rotation cells to shift this packet when the head is turning, in the absence of visual input. This is an important finding, given that previous HD cell models have not addressed how the necessary synaptic structures for full HD system functionality could arise.

Despite these important results, the model has a number of issues. Although path integration happens, it does not happen at the correct speed. Synaptic learning is detailed as higher order, with associations being made between multiple presynaptic inputs and a postsynaptic firing. Such Sigma-Pi synapses are of uncertain biological plausibility, as it is not clear how such multiplicative synapses might operate in the brain. Recurrent collateral connectivity within a layer does indeed self-organise. It is possible, however, to say that the model is not truly self-organising. HD cells start with a preferred firing direction relative to the visual world. Despite this, no attempt is made to explain how preferred firing directions are generated in the first place. There is no inclusion of how visual landmarks may gain control of the HD system, and thus no discussion of the relationship between allothetic and idiothetic information sources. As with previous models, this work assumes within-layer recurrent collaterals in HD regions, whereas in reality this connectivity is unlikely. Although rotation cells sig-

nal AHV in a given direction, they are described simply as monotonically increasing with respect to head velocity in a given direction. Thus, the subtleties of the AHV cell representation in vivo are not replicated with regards to the variety of AHV cell responses (e.g., symmetrical vs asymmetrical AHV cells).

An update of this model is presented by Kyriacou (2011), who implemented it on a LEGO robotic platform and extending it to incorporate vestibular and kinesthetic inputs. This provides validation that the model of Stringer et al. (2002) can self-organise from real data, and replicates experimental observations on the contribution of multiple information sources to HD cell stability. However, this updated model still employs Sigma-Pi neurons, and therefore is also not biologically plausible. Additionally, the system can be said not to be fully self-organising as HD cells are still pre-specified with preferred firing directions. A brief discussion is made of conflict between information sources but, much as with the model of Zeidman and Bullinaria (2008) (Section 2.4.6), the treatment of cue conflict is possibly overly simplistic with regards to dynamic changes in salience and remapping of visual information (e.g., Knight et al., 2014). Whilst less biologically plausible, it demonstrates that the general style of HD cell models reviewed in this thesis can indeed guide navigation in a real-world agent.

2.4.11 Song and Wang (2003; 2005)

Song and Wang (2003, 2005), with later analysis by Boucheny et al. (2005), provide an alternative to Rubin and colleagues (2001) method of avoiding excitatory recurrent collaterals. They focus on the reciprocal connectivity between LMN and DTN as a generative HD cell circuit, in strong agreement with experimental evidence outlined in Chapter 1 of this thesis. The model consists of a layer of LMN

HD cells, receiving offset input from DTN AHV by HD cells. There is also consistent external input to LMN cells, to ensure firing.

The model is capable of generating a Guassian activity packet, which can be sustained in any position, and replicates HD cell tuning curves. It is also capable of performing path integration. Thus, attractor and integrator behaviour is achieved without full CANN architecture of excitatory recurrent collaterals within a layer. However, the system is not truly self-sustaining, instead using inhibition to shape a tonic input source. The identity of this input is not adequately elaborated. If vestibular, then this model would not account for bursty firing following vestibular lesions, and would predict a complete loss of HD cell firing. In addition, the system does not self-organise and thus does not clarify how the specific connectivity required for this models operation might arise through learning.

2.4.12 Hahnloser (2003)

Another self-organising model was proposed by Hahnloser (2003). Similar to other self-organising models (Stringer et al. 2002; Stringer and Rolls 2006), the assumption is made that visual information acts as a calibrating input for the HD system. In this model, two HD cell rings each receive visual input and vestibular input. Each HD ring receives input from AHV cells signalling head rotation in a single direction only. The two rings are interconnected with offset connectivity. Once trained, the model is capable of both sustaining a stable activity packet when the head is stationary, and shifting this activity packet during head rotation.

A principal issue with this model, as with other self-organising HD system models, is the assumption that visual information is used to calibrate the HD system. However, for this to work HD cells

must already have preferred directions, and thus the HD relationship with the visual world must already be formed. However, observations of HD cells in rat pups before eye opening (Bjerknes et al., 2015; Tan et al., 2015) suggest that, whilst visual information likely calibrates the system, an at least partially formed HD cell system must exist without visual input. Another major issue is the use of an error-correcting learning rule. In this case, learning is driven by the difference between actual neuronal firing and desired neuronal firing. It is not clear how error might be calculated at, or conveyed to, each synapse in a biologically plausible manner. Vestibular information is not self-organising, as a specific prior relationship between HD cells in the two rings and vestibular input exists. This can be considered analogous to the issues with visual input: how can inputs have an ordered relationship to a system before that system is itself ordered? In addition, vestibular input is not as nuanced as seen in the brain, lacking the full range of AHV cell responses.

2.4.13 Degris et al. (2004)

Degris et al. (2004a) provide a model focusing on the speed of HD system reorientation to visual landmark cues. Using a spiking network, they are capable of reproducing very short (80ms) latencies observed in ADN. Crucially, this work is based on a CANN architecture formed by reciprocal LMN-DTN connectivity. It is likely that any HD cell attractor network would exist between layers, and thus it is of note that these low latencies are achieved with a CANN architecture spread across two layers in this manner. This model makes use of a biologically-realistic architecture, with visual information entering the HD system via PoS, which transmits to ADN, itself in turn projecting to and receiving from a reciprocally-connected LMN-DTN complex. This same model is extended to address idiothetic update (Degris et al., 2004b), and again shows short update latencies (approx.80ms).

This model contributes an understanding of how intrinsic cellular properties (i.e., spiking dynamics) can contribute to short response latencies to both visual and vestibular information, and maintains a good level of biological accuracy in terms of regional connectivity. However, the issue of self-organisation is not addressed, and it is therefore not clear how the required connectivity might arise for successful path integration and HD memory states. This model could well be extended by addressing not just visual and vestibular information separately, but how they might interact. In particular, the learning of visual landmarks, and the behaviour of the HD system during cue conflict ought to be considered as a next step. Due to the dependence of the attractor network on AHV cells which fire for zero head rotation, the model is dependent on vestibular input for HD signal generation. However, it is not clear why HD cells might become bursty cells without vestibular input. It is possible that the answer to this lies in how allothetic cues are learnt in relationship to idiothetic information, and how this learning might break down in the absence of such information. Additionally, path integration should be achieved with a model making full use of the AHV cell subtypes, and with realistic response properties.

2.4.14 Stringer and Rolls (2006)

Stringer and Rolls (2006) present another self-organising HD cell model, in an update to the work of Stringer and colleagues (2002). In this two layer model, a CANN network of HD cells receives inputs from a layer of cells which respond to a combination of HD and AHV. This additional layer of combination AHV by HD cells avoids the need for biologically implausible multiplicative Sigma-Pi synapses. These combination cells learn to respond to specific combinations of HD and AHV through competitive learning. The model employs a variable firing threshold based on current cell firing. This enhances the firing of currently-active cells and thus acts to stabilise an activity packet in the HD cell

layer. One possible mechanism for this is persistent firing based on intrinsic Ca2+ currents Yoshida and Hasselmo (2009), although this warrants further investigation.

Symmetrical connectivity is learnt between HD cells, setting up a CANN architecture. Offset connectivity between combination cells is learnt due to employment of a temporal trace rule: postsynaptic HD cell firing is associated with past presynaptic combination cell firing. This offset connectivity associates a combination of current HD and AHV with HD some time interval in the future, and provides the necessary drive to shift an activity packet around the HD layer.

Although this model is an update of Stringer et al (2002), it still does not achieve accurate path integration. Whilst it does enhance biological plausibility by avoiding multiplicative synapses, it retains several key omissions: no consideration of HD cell preferred firing direction development, and therefore no consideration of learned visual influences, as well as an assumption of within-layer excitatory connectivity. In addition, AHV cells are still modelled in an oversimplified manner as simply asymmetric.

2.4.15 van der Meer et al. (2007)

van der Meer et al. (2007) proposed a model to explain anticipatory firing in the HD cell system. Specifically, they sought to explain the variations in anticipatory firing observed both within and between areas; an aspect of anticipatory firing not considered by previous HD cell models. The model consists of a CANN ring of HD cells, receiving offset connectivity from AHV cells. They show how biophysical mechanisms, such as spike rate adaptation and postinhibitory rebound, in a spiking neural network leads to the HD system receiving a high-pass filtered version of AHV cell firing.

The model provides an interesting spiking perspective on the causes of anticipatory time intervals, accurately replicating both within-session and between-session variability in anticipation. However, the authors admit that some variability in real HD cell data remains unexplained. Regardless, this model demonstrates how anticipation may arise from system dynamics rather than circuitry. However, it is limited in that it is not self-organising and does not model fully visual input. Nor does it accurately model idiothetic inputs, given that AHV cells are taken to be asymmetric only.

2.4.16 Stratton et al. (2010, 2011)

Stratton et al. (2010, 2011) have elaborated on how self-organisation might occur without a specific relationship between HD cells and vision, and thus specific preferred firing directions, being pre-established. Their proposed mechanism makes use of symmetric AHV cells, an aspect of the HD system that is frequently overlooked in the modelling literature. In their model, HD cells receive offset connectivity from two populations of asymmetric HD by AHV cells, one population for each direction. HD cells also receive input from a single symmetric AHV cell projecting equally to all cells. The HD layer recurrent weight profile is roughly initialised to be approximately correct, but containing random error and systematic bias.

It was found that symmetric AHV cell input to the HD layer could be used to update recurrent HD connectivity. As symmetric AHV cells effectively signal speed of rotation regardless of direction, this could be related to HD layer packet motion. Following training, HD layer connectivity allowed for accurate path integration tracking during head rotation, and a stable packet with minimal drift when the head is stationary. However, tracking speed during path integration was not accurate. In their

second study (Stratton et al., 2011), this same model utilises landmark information (possibly tactile in origin) to calibrate tracking speed to real world head rotation speed. This calibration can be achieved through usage of any single landmark, at any bearing from the agent, and is shown to be successful.

Whilst this work provides valuable addition to theoretical understanding by suggesting a role for symmetrical AHV cells in HD layer calibration, it is limited. Most obviously, connectivity between HD cells and AHV cells of all types are pre-specified. It is unclear how these connections might self-organise. Given that this paper is focused on self-organisation, this seems somewhat of an omission. Further, the calibration of HD layer recurrent weights depends on a learning rule that involves symmetric AHV cell firing affecting a synapse between two HD cells. It is not clear how this might occur, although the authors suggest that synaptic proximity might have an effect. In addition, the authors assert that symmetric AHV cells must therefore be different to asymmetric AHV cells in terms of connectivity, and morphology, as well as theorised use within the HD cell system. It is not yet clear whether this is the case.

The idea of a rough initial set of connections, which are then calibrated within the HD system is an interesting one. This certainly seems to be possible given the experimental observations of HD cell firing in rat pups prior to eye opening, and improvements to cell firing stability after eye opening (Bjerknes et al., 2015; Tan et al., 2015). The question remains, however, how such initial connectivity might be set up. The authors suggest a possible developmental mechanism based on concentration gradients of diffusing chemicals, which contains stochastic variations. However, this is not the only plausible scenario, and the question of HD system development in relationship to self-organisation

will be unpacked in more detail in the discussion of this thesis.

2.4.17 Massoud and Horiuchi (2011)

Massoud and Horiuchi (2011) present a VLSI (very large scale integration) model of the HD system, which is constructed of physical chips simulating neurons. They use a CANN architecture as the basis for this model, with each neuron containing an intrinsic current biasing them towards spiking. Rotational signals act to shift a packet activity as the head moves, with the shift mechanism being based on disinhibition. This model is capable of both sustaining and shifting an HD layer activity packet.

This model seems somewhat abstracted from the real brain. Indeed, the authors do not relate their model back to the known physiology of HD system areas. Since the rotational input, and indeed visual input, is somewhat abstract in manner, it is hard to comment on the biological plausibility of such a system. Given that the focus of the authors is not towards making predictions about, or giving detailed explanations of, known properties of the HD system in vivo, this paper will not be considered further.

2.4.18 Fortenberry et al. (2012)

Fortenberry et al. (2012) detail a model they called the HeadMoVVes (Head direction from Motor, Visual and Vestibular signals) model. They explicitly state their objectives as explaining how cues from various sources generate a HD signal which correctly tracks true HD in both the light and the dark. They also seek to explain how motor and vestibular inputs might generated anticipatory time intervals, and explain how specifically distal visual landmarks can be learned to reset HD cell preferred directions. In doing so, they explain specific roles for each HD cell subregion.

In this model, the LMN-DTN generates the HD signal from a mixture of delayed AHV input and anticipatory motor efference copy. These two sources are averaged in ADN, and this averaging is taken to be underlying mechanism of HD cell anticipatory firing. The HD signal is then projected to PoS, where it is integrated with visual information conveyed via RSP. RSP processing allows for the learning of visual landmarks based on their statistical relationship with HD. This allows, for example, a separation of distal and proximal visual landmarks. This model is able to replicate many aspects of the HD system, and several important predictions are generated for future experimental testing.

However, this model has several limitations. It is not fully self-organising and whilst it has very specific multi-layer connectivity, there is no nuanced consideration of how such detailed connectivity might arise. In particular, visual landmarks are learned in reference to HD cells which already have preferred firing directions. Within layer excitatory connectivity exists but, as discussed previously, there is no evidence for this type of connectivity in HD cell areas. Although AHV and motor input are considered in detail, no role is proposed for symmetrical AHV cells.

2.4.19 Walters and Stringer (2010)/Walters, Stringer and Rolls (2013)

Walters and colleagues (Walters and Stringer 2010; Walters et al. 2013 provided a self-organising path integration model based on the work of Stringer et al (2006), aimed at improving path integration speed. In this updated model, a CANN layer of HD cells receives input from cells which respond to a combination of AHV and HD; termed COMB cells. HD cells also receive visual inputs, with each HD cell being pre-specified with a preferred firing direction. Crucially, connections between the HD and COMB cell layers contain a delay in both directions. This delay can be attributed to a

number of causes, such as axonal conductance, or synaptic mechanisms. During training in the light, an activity packet is propagated through the HD cell layer at a set constant rotation speed. Hebbian learning coupled with delay between layers means that a given HD cell projects, via the COMB cell layer, to a predicted future HD cell after time interval 2t, where t is the delay between layers. This allows the system to predict a future HD after rotation at a constant speed over a known time interval. COMB to HD connectivity thus provides the necessary asymmetric input for path integration. Simple Hebbian learning without delay allows for the formation of recurrent connectivity based on similarity of preferred directions in the HD layer, and Hebbian learning enhanced by competition allows COMB cells to fire only to a specific HD and AHV.

Once trained, this model produces approximately accurate path integration. However, this is not achieved at 100% accuracy, reaching a maximum of approximately 80%. Several criticisms are retained from the work of Stringer et al (2006), most notably pre-specification of the relationship between HD cells and allothetic input. This is an outstanding question, as outlined in the discussion of Stratton et al (2010; 2011). Thus, a possible extension would be to approach the development of a relationship between HD cell firing and allothetic spatial cues, and the relationship of this process to idiothetic information. The model assumes within-layer HD to HD connectivity, but this has been strongly suggested not to exist by the experimental literature. Whilst this model serves as an approximation of biology, AHV and COMB cell responses are overly simplified. AHV cells respond to a single AHV in a binary manner, and are uniformly asymmetrical. As a result, COMB cells respond to a single HD and a single AHV. In reality, the firing of COMB cells is much more complex and likely represents a population encoding. It would be interesting to see whether a richer COMB layer

representation might allow the model to cope with multiple path integration speeds, as at present the model cannot generalise to new speeds following training.

2.5 General notes on the models in this thesis

When constructing a computational model, decisions must be made about the level of biological detail to simulate. Too abstract a model can provide little insight, whereas too much detail can obscure the forest for the trees. With regards to this thesis, this decision has been made on the basis of not just whether we can model in detail, but whether we should. That is, does an extra level of detail in the neuronal and synaptic model used generate a corresponding extra insight? A decision has been made not to make use of fully spiking neural networks. That is, individual action potential spikes are not simulated, nor are biophysical processes such as membrane currents. Instead, neuronal activity is taken as a time-averaged instantaneous firing rate.

In part, this decision was based on computational expense. The number of equations required to simulate an individual neuron in a firing rate model is much smaller than in a spiking neural network. For example, action potential generation, as described in a Hodkin-Huxley style neuronal model, is not modelled in detail, as this requires equations describing the effect of ion channels. Reduction of model complexity in terms of the number of equations required to simulate individual neurons greatly reduces the time taken for simulations to run. This was important given the computational resources available, and the requirement that simulations be run with many neurons and with dynamically changing weights between these neurons. Shorter simulation runtimes allow for fuller exploration of model parameters and deeper investigation of model operation principles.

This decision was also based on considering what extra information is conveyed in the relative timing of spikes in the HD system as opposed to rate of spiking. Given that most aspects of HD firing are adequately captured in a firing rate model, this was deemed an appropriate level of biological detail. In modelling the HD system, the timing of neuronal firing is important. For example, path integration must happen at the correct speed. The models employed in this thesis are therefore differential, operating in continuous time.

This thesis makes use of self-organising connectivity. In many of the simulations presented here, self-organisation allows for model functions to be adaptive. That is, a model may alter connectivity in order to respond to changes in the external environment. An example of this is given in Chapters 3 and 4, in which visual information is dynamically remapped onto the HD system. Whilst investigating the behaviour of a fully mature system with ideal pre-specified connectivity is informative and an important proof of concept for any hypothesis, it is also vital to see whether a system can arise naturally. An example of this is given in Chapter 5, in which a differential effect of neuronal rise time is observed during model training and model testing. More generally, the developmental trajectory of the HD cell system is not completely clear. It is however unlikely that the precise architecture, needed for any functioning HD model to stably maintain or accurately shift an activity packet, is pre-specified in the genetic code. Whilst a rough hard-wired architecture could perhaps exist initially (e.g., see the discussion of this thesis), to achieve full system functionality connectivity would require fine tuning through learning. Therefore, a mature system must arise in some form through dynamic changes in synaptic connectivity over a period of time.

Chapter 3

Cue conflict resolution

A number of studies have explored the intricacies of cue conflicts in the head direction cell system, with mixed results as to which types of cue dominate preferred directions. One possible explanation for this is the idea of distance-dependent cue integration, with visual landmarks having a dominant influence for small conflicts, and with a switch to idiothetic dominance for larger conflicts (Knight et al., 2014). This chapter presents details of the architecture, governing equations, implementation and simulation results for a continuous attractor neural network that is capable of resolving conflicts between visual and path integration cues. This is performed in a distance-dependent manner which closely replicates the previous experimental work cited above, and thus provides a working theory of how cue conflict resolution is performed in head direction cell systems. To date, no other existing model has directly addressed cue integration in this manner.

3.1 Model Description

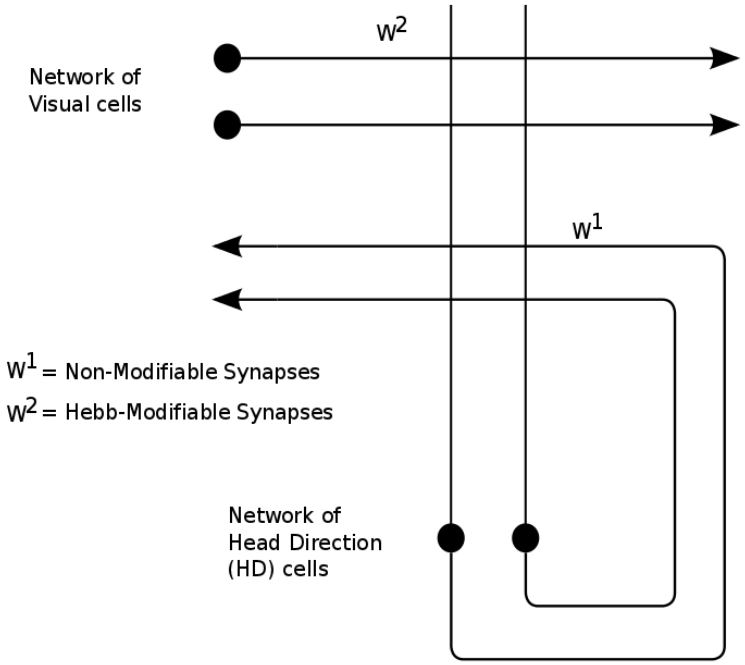


Figure 3.1: Simplified network architecture, demonstrating connectivity both from the visual layer into the head direction cell layer w^2, and within the head direction cell layer itself w^1

A model schematic and network architecture are displayed in Figures 3.1 and 3.2. The network consists of a ring of head direction cells, and a ring of visual input cells. The ring of visual cells encode the egocentric bearing from the agent to a visual cue in the environment. Each visual cell responds maximally for a specific bearing to the cue within the interval $0 - 360\,^{\circ}$, with a Gaussian decline in firing rate as the visual cue departs from the preferred bearing for that cell. The visual cells are arranged topologically with an even distribution around the ring. The egocentric bearing to the visual cue, which is encoded by the location of the activity packet within the ring of visual cells, is determined by the position of the cue in the environment and the head direction of the agent.

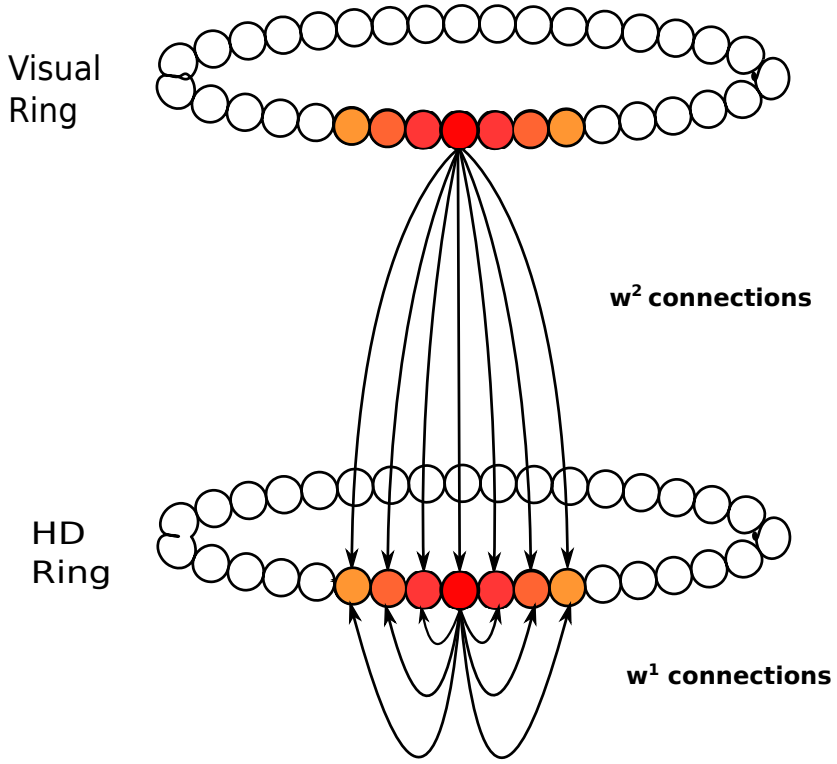


Figure 3.2: Stylised network schematic diagram, showing location correspondence between visual and head direction cell rings.

Although the rat has a more limited 270 ° field of view (Rolls, 1999) this model contains a full 360 ° ring of visual cells. Natural environments usually contain multiple visual cues distributed along different bearings from the animal. Hence, every head orientation would result in a unique visual representation corresponding to a particular subset of cues and their bearings from the rat. However, rather than implementing a biologically detailed vision network that can process input from natural scenes containing multiple cues (Wallis and Rolls, 1997; Rolls et al., 1998; Rolls and Milward, 2000) the visual representation has been simplified by incorporating a single ring of vision cells representing the bearing to a single cue. Nevertheless, the localized activity packet within the vision ring could be thought of as a combined representation of the different bearings to a number of cues. Thus, it is expected that the mechanisms discussed in this paper would still operate effectively within a more

biologically realistic network which can process visual input from natural scenes.

Visual cells send synaptic connections to the ring of head direction cells. The head direction cells are also arranged topologically with an even spread around the ring. The response preferences of the head direction cells are initially established by defining a topological connection relationship from the ring of visual cells to the ring of head direction cells. Each head direction cell receives inputs from all of the visual cells via the w^2 synaptic connections. However, before learning, the w^2 weights are initialized so that each head direction cell receives maximal w^2 weight from the topologically corresponding visual cell, with a Gaussian drop off in the w^2 weight profile across the remaining visual cells. Thus, each head direction cell responds maximally to the head direction that the corresponding visual cell responds to, given the current location of the visual cue within the environment. Consequently, the preferred head directions of the head direction cells are distributed evenly around the ring from $0 - 360\,^{\circ}$. The head direction of the agent is encoded by the location of the activity packet within the ring of head direction cells. During subsequent cue conflict situations, the w^2 connections may be modified by learning. As demonstrated below, this may lead to a remapping between the visual cells and head direction cells.

Individual head direction cells are simulated with a leaky-integrator firing rate model, described in more detail in Section 3.2.1. Neuronal firing is represented as an instantaneous average firing rate for each cell, which does not model individual action potentials. Cells within the head direction ring influence one another via non-modifiable w^1 excitatory recurrent collateral connections, whose profile is based on similarity of cell preferred head directions. That is, connections are strongest between cells

with similar preferred directions and weakest between cells with dissimilar preferred directions. This connectivity allows the head direction ring to behave as a continuous attractor, sustaining a Gaussian packet of activity whose centre represents the current head direction of the agent. Visual cells are not modelled as leaky integrators. Instead, Gaussian activity packets are explicitly specified within the visual ring as part of the experimental protocol. Visual activity packet locations are set depending on where the landmark stimulus is to be located with respect to the agent.

3.2 Model Operation

This section details the central equations of the model, their numerical solutions, and the simulation protocols used.

3.2.1 Model Equations

The activation level h_i^{HD} of head direction cell i is given by

$$\begin{aligned}\tau^{HD}\frac{dh_i^{HD}(t)}{dt} = &- h_i^{HD}(t) + e_i(t) \\ &- \frac{1}{N^{HD}}\sum_j \tilde{w}^{HD} r_j^{HD}(t) \\ &+ \frac{\phi_1}{C^{HD\to HD}}\sum_j w_{ij}^1(t) r_j^{HD}(t) \\ &+ \frac{\phi_2}{C^{VIS\to HD}}\sum_j w_{ij}^2(t) r_j^{VIS}(t)\end{aligned} \tag{3.1}$$

where $h_i^{HD}(t)$ is the activation level of head direction cell *i* at time *t*. The term $-h_i^{HD}(t)$ represents a decay term, whereby in the absence of presynaptic input, the activation of *i* will fall to zero as determined by the time constant τ^{HD}. The term $e_i(t)$ represents external activation of head direction cell i at time t, which may be due to a variety of sensory signals such as tactile, olfactory or visual. This input term is required in order to initialise an activity packet in the correct location within the head direction cell layer before being removed and activation being sustained through recurrent collateral activity alone. It is determined by a Gaussian profile described in Equations (3.12) and (3.13). The term $\frac{1}{N^{HD}} \sum_j \tilde{w}^{HD} r_j^{HD}(t)$ represents inhibitory feedback within the head direction cell layer, summed over all presynaptic head direction cells *j*. $\tilde{w}^{HD}$ is a global constant describing the effect of inhibitory interneurons, and N^{HD} is the total number of head direction cells in the layer. The term $\frac{\phi_1}{C^{HD \to HD}} \sum_j w_{ij}^1(t) r_j^{HD}(t)$ represents excitatory input from the head direction cell layer back onto itself via recurrent collateral synapses, summed across all presynaptic head direction cells connected to *i*. It consists of $r_j^{HD}(t)$, the firing of presynaptic cell *j* at time *t*, and $w_{ij}^1(t)$, the synaptic weight from presynaptic head direction cell *j* to postsynaptic head direction cell *i* at time *t*. It is important to note that this weight is not updated during simulation, and thus remains at the same initial value throughout. The entire term is scaled by the factor $\frac{\phi_1}{C^{HD \to HD}}$, which controls the overall strength of recurrent synapses with the head direction cell layer where ϕ_1 is a constant and $C^{HD \to HD}$ is the number of synapses each postsynaptic head direction cell receives from presynaptic head direction cells.

The term $\frac{\phi_2}{C^{VIS \to HD}} \sum_j w_{ij}^2(t) r_j^{VIS}(t)$ represents input from the visual layer into the head direction cell layer summed over all presynaptic visual cells, where $r_j^{VIS}(t)$ is the firing of presynaptic visual cell *j* at time *t* and $w_{ij}^2(t)$ is the connection strength from presynaptic visual cell *j* onto postsynaptic

head direction cell *i*. Again, a scaling term $\frac{\phi_2}{C^{VIS \rightarrow HD}}$, where ϕ_2 is a constant and $C^{VIS \rightarrow HD}$ is the total number of connections each postsynaptic head direction cell receives from presynaptic visual cells, is used to scale the overall strength of the visual inputs to the head direction cell layer.

The firing rate $r_i^{\mathrm{HD}}(t)$ of head direction cell *i* at time *t* is a sigmoid function of the activation level of cell *i*

$$r_i^{HD}(t) = \frac{1}{1 + e^{-2\beta(h_i^{HD}(t) - \alpha)}} \tag{3.2}$$

where α and β are the sigmoid threshold and slope respectively. As this model is firing-rate based, no individual action potentials are simulated, and neuronal activity is represented as a spontaneous average firing rate.

Both recurrent weights within the head direction cell ring and feed-forward weights from the visual ring onto the head direction cell ring are initially pre-wired in a specific configuration. In both cases, this is a Gaussian configuration, for example determining the recurrent weight w_{ij} for synapses onto postsynaptic head direction cell i from presynaptic head direction cell j

$$w_{ij}^1 = e^{-(s_{ij}^{RC})^2/2(\sigma^{RC})^2} \tag{3.3}$$

where σ^{RC} is the standard deviation of the Gaussian profile. s_{ij}^{RC} is the difference between the pre-

ferred head directions s_i and s_j of the post- and pre-synaptic head direction cells i and j respectively. It is given by

$$s_{ij}^{RC} = MIN(|x_i - x_j|, 360 - |x_i - x_j|) \tag{3.4}$$

which creates a wrap-around effect, with weights between head direction cells as a population remaining continuous across the $360/0\,^{\circ}$divide. Feed-forward weights from the visual layer onto the head direction cell layer in an equivalent manner to Equations (3.3) and (3.4), with σ^{RC} substituted for σ^{FF}, and s_{ij}^{RC} for s_{ij}^{FF}.

Recurrent weights within the head direction cell layer are static and thus retain their initial configuration. The feed-forward synaptic weights are updated by the local associative Hebbian learning rule

$$\frac{dw_{ij}^2(t)}{dt} = kr_i^{HD}(t)r_j^{VIS}(t) \tag{3.5}$$

where k is a learning rate constant controlling the speed of weight changes, $r_i^{HD}(t)$ is the firing rate of postsynaptic head direction cell i at time t and $r_j^{VIS}(t)$ is the firing of presynaptic visual cell j at time t. This learning rule will act to strengthen the synaptic weights between head direction and visual cells which are co-active.

In some simulations, where explicitly stated, homosynaptic LTD is employed. In this case, the weight update equation takes the form

$$\frac{dw_{ij}^2(t)}{dt} = k(r_i^{HD}(t) - \langle r_i^{HD} \rangle) r_j^{VIS}(t) \tag{3.6}$$

where $\langle r_i^{HD} \rangle$ is the average firing of postysnaptic head direction cell i. This has the effect of weakening synapses between active presynaptic visual cells j and inactive or weakly active postsynaptic head direction cells i. It is termed homosynaptic LTD, as the synapses which weaken are the same through which the postsynaptic cell is stimulated. This learning rule will still act to strengthen synapses between head direction and visual cells which are co-active.

In order to ensure that synaptic weights remain bounded, weight normalisation is employed. Feed-forward weights $w_{ij}^2(t)$ are rescaled every timestep, so that for every postsynaptic head direction cell *i*

$$\sqrt{\sum_j (w_{ij}^2(t))^2} = 1 \tag{3.7}$$

where this summation is performed over all presynaptic visual cells j. This ensures that, during learning, the feed-forward weights are convergent, in that they eventually settle to steady values. This process could well be achieved in the brain by synaptic weight decay (Oja, 1982; Rolls et al., 1998).

3.2.2 Numerical Methods

The differential equations given above cannot be solved analytically. Instead, they are implemented in the computer model by making discrete approximations of their solutions. A Forward Euler finite difference scheme is used to approximate all differential equations during simulation. This can be illustrated by taking a simplified form of the differential activation Equation (3.1)

$$\tau^{HD} \frac{dh_i^{HD}(t)}{dt} = -h_i^{HD}(t) + f_i(t) \tag{3.8}$$

where $f_i(t)$ represents the sum of all presynaptic inputs to postsynaptic head direction cell i at time t. The Forward Euler approximation of this simplified equation is given as

$$h_i(t + \delta t) = h_i(t) - \frac{\delta t}{\tau^{HD}} h_i(t) + \frac{\delta t}{\tau^{HD}} f_i(t) \tag{3.9}$$

where $h_i(t + \delta t)$ is the value of h_i to be approximated at time $t + \delta t$, and δt is the timestep of the Forward Euler method. Thus the differential equation is replaced by a Forward Euler method which approximates a solution after the timestep δt based on current values of $h_i(t)$ and $f_i(t)$.

A Forward Euler approximation can also be derived for the differential Equation (3.5) updating synaptic weight $w_{ij}^2(t)$ from presynaptic visual cell j to postsynaptic head direction cell i. The original equation

$$\frac{dw_{ij}^2(t)}{dt} = kr_i^{HD}(t)r_j^{VIS}(t) \tag{3.10}$$

is converted to a Forward Euler approximation

$$w_{ij}^2(t+\delta t) = w_{ij}^2(t) + \delta t[kr_i^{HD}(t)r_j^{VIS}(t)] \tag{3.11}$$

where $w_{ij}^2(t+\delta t)$ is the synaptic weight approximated at time $t+\delta t$ based on current values of $w_{ij}^2(t)$, $r_i^{HD}(t)$ and $r_j^{VIS}(t)$.

In this manner, it is a simple case of expanding the simplified equations given above to the application of the more complex differential equations governing cell activation and feed-forward weight update given previously.

3.2.3 Experimental Protocol

Experiments were conducted in blocks of 18 simulations. At the start of each simulation firing rates for every head direction cell r_i and every visual cell r_j are at zero. A packet is then stimulated in the head direction ring. This is equivalent to a visual landmark cue, such as a light, being shown in a location and then being removed. This is represented in Equation (3.1) as $e_i(t)$, which for each head direction cell i is calculated according to a Gaussian profile

$$e_i = \lambda^{HD} e^{-(s_i^{HD})^2/2(\sigma^{HD})^2} \tag{3.12}$$

where λ^{HD} is a scaling factor determining the strength of this input to head direction cell i and σ^{HD} is the standard deviation of the Gaussian profile. s_i^{HD} is the difference between the true head direction x and the preferred direction x_i of postsynaptic head direction cell i. It is given by

$$s_i^{HD} = MIN(|x_i - x|, 360 - |x_i - x|) \tag{3.13}$$

which creates a wrap-around effect, with response profiles of head direction cells as a population remaining continuous across the $360/0\,^{\circ}$divide. This input is removed after 40ms, and the packet is subsequently sustained through recurrent collateral activity alone.

The packet remains stationary, equivalent to the simplifying assumption that the agent is not free to explore its hypothetical environment. Shortly afterwards, a packet is generated in the visual ring, analogously to Equations (3.12) and (3.13) substituting λ^{HD} for λ^{VIS}, σ^{HD} for σ^{VIS} and s_i^{HD} for s_i^{VIS}. This packet is centered on a location that does not correspond to the current head direction cell activity. This simulates a conflict between the memory location in the HD ring and visual landmark information. The locations of the conflicting visual landmark, relative to the initial head direction representation, are systematically varied to investigate the effect of distance on cue conflict resolution in this model system. In all the simulations presented below, head direction packets were always gen-

erated at $180\,^{\circ}$, whilst the landmark cue was simulated at $10\,^{\circ}$ from this location in the first simulation of the block, and at distances increasing by $10\,^{\circ}$ each time for further simulations. The distance was increased up to a maximum conflict between head direction representation and visual landmark cues of $180\,^{\circ}$.

Once a conflict is introduced, the activation level $h_i^{HD}(t)$ and firing rate $r_i^{HD}(t)$ of each head direction cell i are updated every timestep t as according to numerical approximations of Equations (3.1) and (3.2). Feed-forward weights $w_{ij}^2(t)$ from the visual ring onto the head direction ring are updated every timestep according to Equation (3.5). Simulations end after 2 seconds, as the activity packet within the head direction cell ring is observed as stable by this point.

3.3 Resolving cue conflict: A hypothesis

Several core features of the data presented by Knight et. al Knight et al. (2014) are approached. Firstly, their data shows an initial dominance of visual landmarks upon head direction cell preferred directions, albeit with a lack of complete dominance. Secondly, this dominance decreases with increasing conflict size, with a shift towards internal sources of information. Finally, their data show that visual landmark cue dominance is dependent on experience, with total dominance possible even for large conflict sizes; provided a visual cue has not previously been experienced as unstable relative to ideothetic information.

Initially, the feed-forward weight structure between the visual and head direction rings is consistent with cell preferred directions. That is, each visual cell projects most strongly to the head direction

cell representing the head direction at which that visual cell responds maximally in the light. During a conflict between the bearing of the visual cue and the location of the packet of head direction activity, a packet generated in the visual ring does not exactly correspond to that present in the head direction ring. The edge of this visual activity packet will directly provide asymmetric stimulation to the existing head direction cell packet. In the absence of learning, this would result in the initial head direction packet moving to centre itself over the maximum of this visual input, which is the location of the visual landmark cue as specified on the visual ring (Amari, 1977; Ermentrout and Cowan, 1979; Stringer et al., 2002; Touretzky, 2005). However, experimental studies have shown that this does not happen. Instead, the memory packet within the HD ring appears to stabilise short of the location in the HD ring associated with the new location of the visual landmark. We have been able to explain these experimental observations by introducing learning into the connections from the visual cells to the HD ring.

The crucial aspect of this model is the plastic feed-forward connectivity between the visual and head direction rings, which allows for variation in the final packet location within the head direction cell ring. As the initial head direction activity packet shifts to centre itself over the source of asymmetrical visual input, Hebbian learning is taking place; remapping the visual input onto the head direction cell ring. As the activity packet moves through the head direction cell ring, a strengthening occurs of the connections between the region of the visual ring that is active and the current location of the head direction activity packet. This has the effect of changing where in the head direction ring the current visual activity maximally excites. The net effect of this remapping is to cause the packet to undershoot the location it would have reached were the original weight structure preserved, and to stabilise in a new position. This is a direct result of remapping, with the currently active visual location now pro-

jecting maximally to the location of the final, stable, head direction activity packet. This hypothesised mechanism is represented graphically in Figure 3.3.

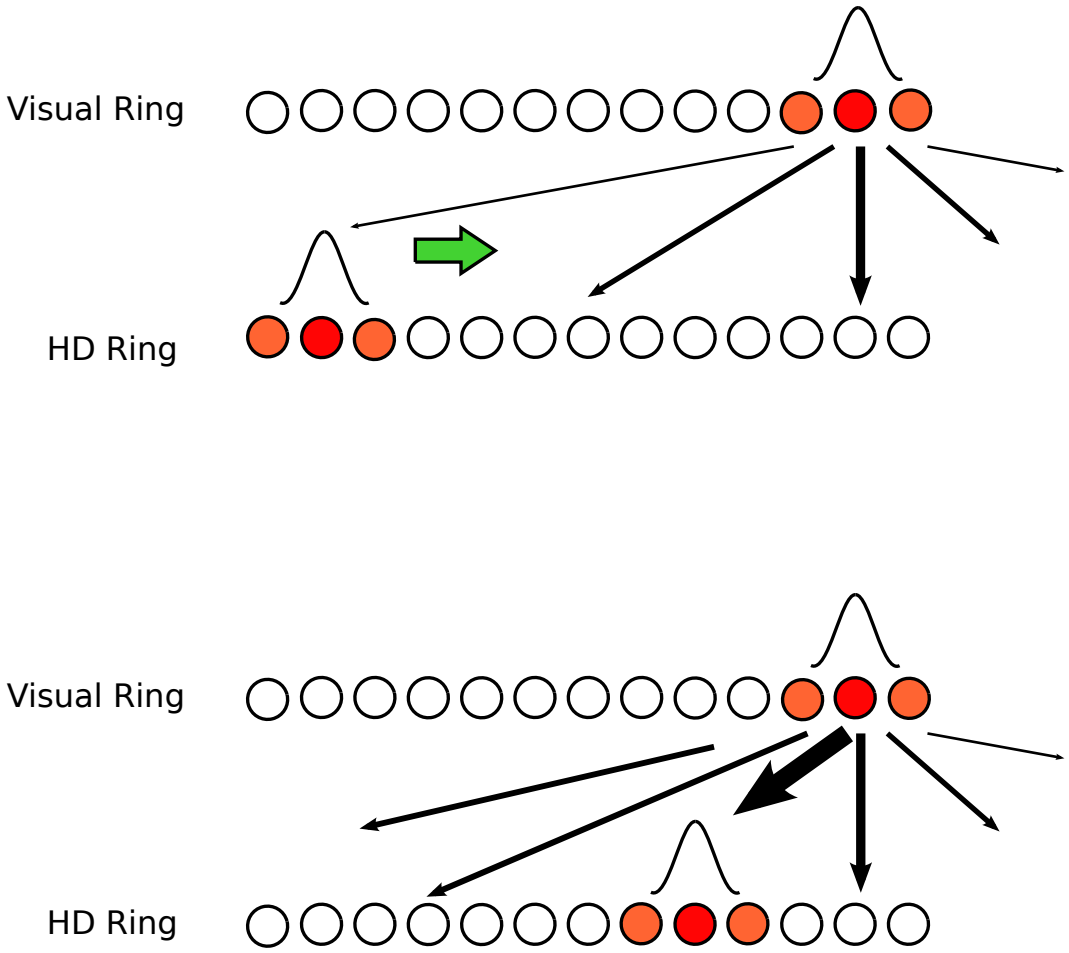


Figure 3.3: Diagram demonstrating the theory behind model operation. Visual landmark cues provide asymmetric input to the head direction cell packet, drawing it through the network. Feed-forward w^2 weight strengths are represented by line thickness. Before learning (top), the strongest weights project to a corresponding location on the HD ring. After learning (bottom), the weights are strengthened between the visual cue and every intermediate HD ring location, with the net result being that the strongest feed-forward connections project to a new location on the HD ring.

3.4 Simulation Results

This section presents simulation results using the model and protocol described above. As well as a demonstration of the main effect, an exploration of the parameter space is presented. For all simulations given, parameters are as shown in Table 3.1 unless explicitly stated to be otherwise. In order to generate summary graphs detailing final packet locations, it is necessary to calculate a population vector statistic to determine the centre of the head direction cell layer activity packet at its final position. This was calculated using an established population vector scheme (Georgopoulos et al., 1986; Song and Wang, 2005) and is as follows

$$\theta_{pop}(t) = \arctan\left(\frac{\sum_i r_i \sin(\theta_i)}{\sum_i r_i \cos(\theta_i)}\right) \tag{3.14}$$

However, because the ring effectively wraps around, it can provide problems for taking the mean of a set of locations spanning the 360/0 mark. For example, the mean of three cells firing maximally at 310, 330 and $350\,^\circ$ would be correctly calculated by the above formula as $330\,^\circ$. However, the mean of three cells firing maximally at 350, 10 and $30\,^\circ$ would be incorrectly calculated as $130\,^\circ$, rather than the correct value of $10\,^\circ$. To account for this, the following corrected formula (Batschelet, 1981;

Fisher, 1993) is used instead

$$\theta_{pop}(t) = \begin{cases} \arctan\left(\dfrac{\sum_i r_i \sin(\theta_i)}{\sum_i r_i \cos(\theta_i)}\right) & \text{if } \sum_i r_i \sin(\theta_i) > 0, \sum_i r_i \cos(\theta_i) > 0; \\ \arctan\left(\dfrac{\sum_i r_i \sin(\theta_i)}{\sum_i r_i \cos(\theta_i)}\right) + 180 & \text{if } \sum_i r_i \cos(\theta_i) < 0; \\ \arctan\left(\dfrac{\sum_i r_i \sin(\theta_i)}{\sum_i r_i \cos(\theta_i)}\right) + 360 & \text{if } \sum_i r_i \sin(\theta_i) < 0, \sum_i r_i \cos(\theta_i) > 0; \end{cases} \quad (3.15)$$

Table 3.1: Network parameter values for simulations reported in this chapter. All simulations, except where explicitly stated, are run using these parameters.

Network Parameters	
Head Direction Cells	500
Visual Cells	500
$C^{HD \to HD}$	500
$C^{VIS \to HD}$	500
ϕ_1	200.0
ϕ_2	1.0
σ^{FF}	65.0
σ^{RC}	35.0
λ^{HD}	1.0
λ^{VIS}	2.0
σ^{HD}	20.0
σ^{VIS}	50.0
τ^{HD}	0.001
$\tilde{w}^{HD}$	0.4
α	0.0
β	0.3
I^{EXTERN}	0.0
δt	0.00001s
k	0.05

3.4.1 Main Result

Figure 3.4 presents a summary of a block of 18 simulations run using the model and experimental protocol described in this chapter with parameters given in Table 3.1. As can be seen, this provides a close match with previous experimentally-derived results picture in Figure 1.6 (Knight et al., 2014). For conflicts below $120\,^{\circ}$, the final head direction packet location is determined principally by the visual landmark cues. However, Hebbian learning within the feed-forward weights partially remaps the visual cue onto each intermediate location through which the packet moves. Due to the effect of this remapping, the packet motion is slowed and the true location of the visual cue is consistently undershot. This undershoot increases with larger conflict sizes, as the packet is in motion over a larger distance, and thus there is a longer time interval over which more extensive remapping may occur. After a breakpoint at $120\,^{\circ}$, the conflict is large enough such that the influence of the visual landmark is weaker than the effect of the memory packet. This causes the final packet location to be closer to the original head direction cell layer representation than the new location indicated by the visual cue. This is shown as a change in structure from the original perfectly Gaussian weight profile to one that reflects the remapping that has occurred during simulation. This occurs in weights that connect the visual input to head direction cell locations between the initial and final packet locations, and thus is manifested as an asymmetrical distortion of the weight profile. Figure 3.5 shows this change in the weight structure during a typical simulation at $260\,^{\circ}$ of conflict. At $180\,^{\circ}$ separation, the visual cue is providing symmetrical input to the initial head direction cell layer packet, which therefore does not move. Remapping does not have an effect at this stage, because Hebbian learning simply reinforces a relationship between the visual cue and the current packet location, and thus no asymmetrical input is presented to the packet (Figure 3.6).

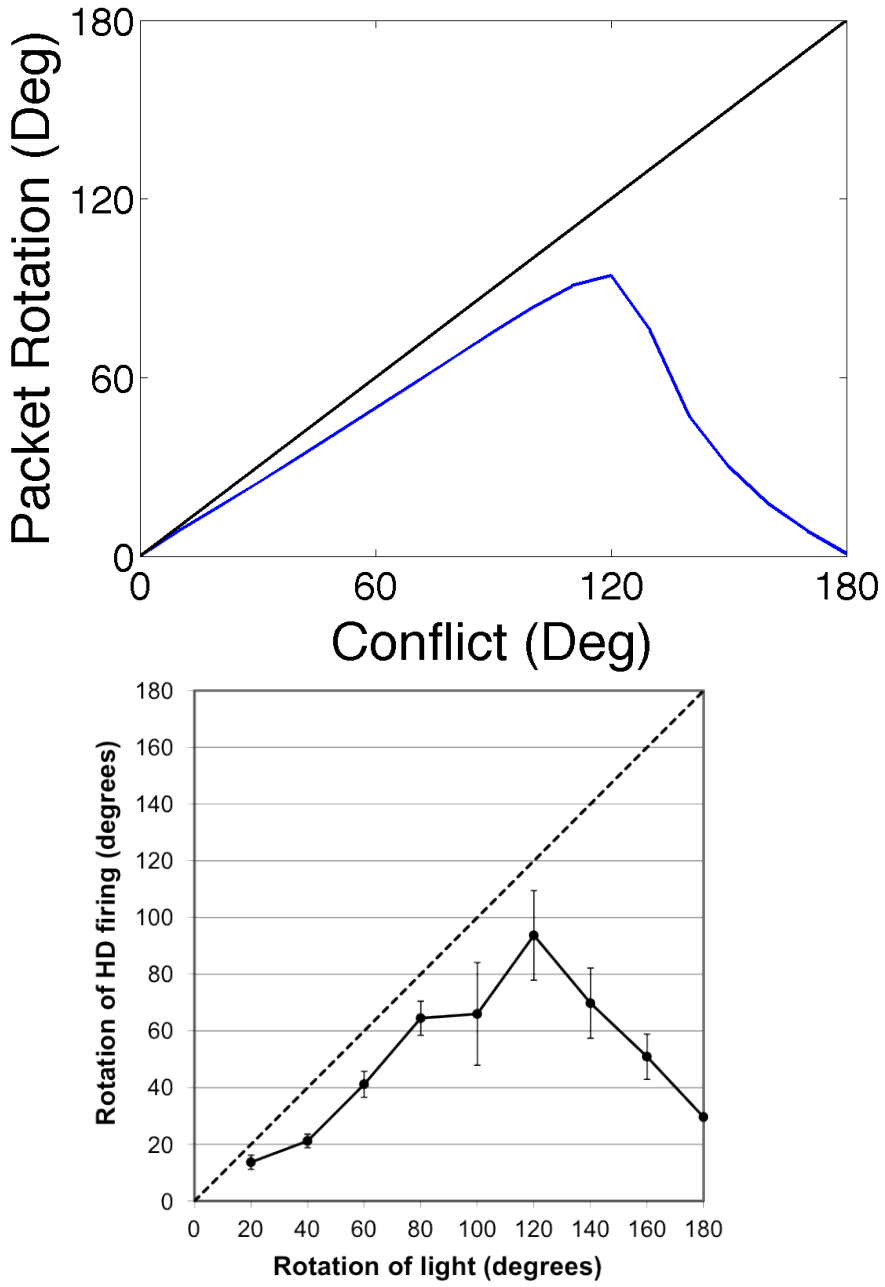


Figure 3.4: Summary graph (top) showing final head direction cell packet locations (blue line) for increasing conflict sizes, in contrast with expected packet location given total visual dominance (black line). Initially packet location is determined mostly by the visual cue, with an undershoot which increases up until 120 ° conflict, after which packet location is determined principally by the memory packet. For ease of comparison, the results of Knight et al. (2014) are included (bottom, with permission).

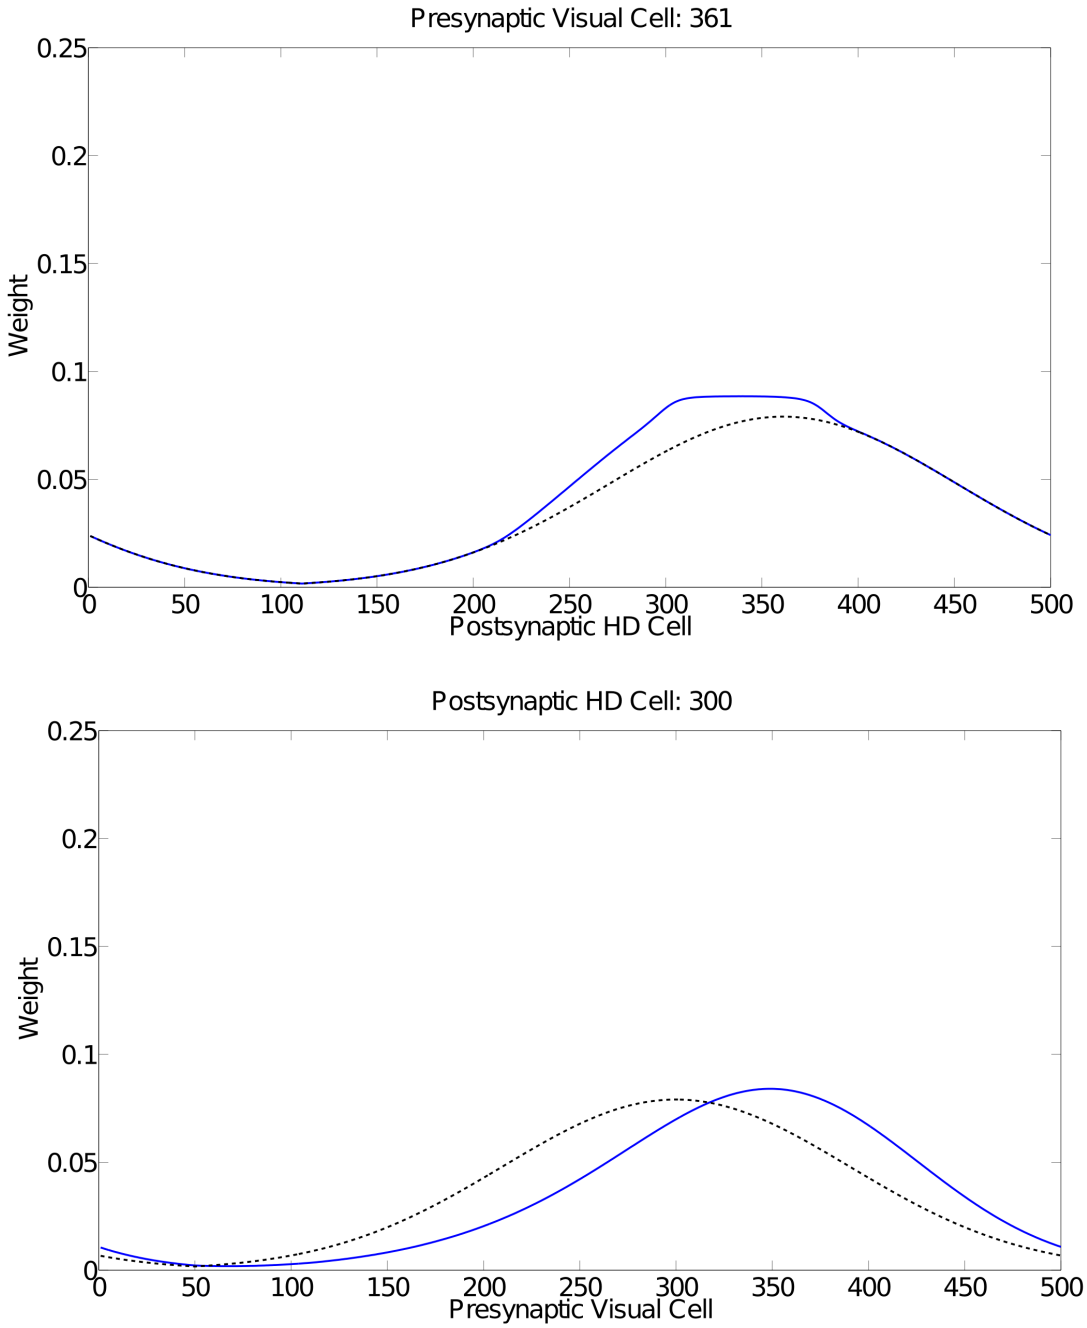


Figure 3.5: Plots showing feed-forward connectivity from the perspective of the projecting visual cell at the center of the conflicting visual landmark (top) and the receiving HD cell at the center of the final HD packet location (bottom), both before (black dashed) and after learning (blue). During simulation, learning causes a strengthening of weights from those visual cells currently active to the current head direction cell location, shown as a distortion of the weight profile. This remapping causes the final HD location to receive maximal input from this new visual location.

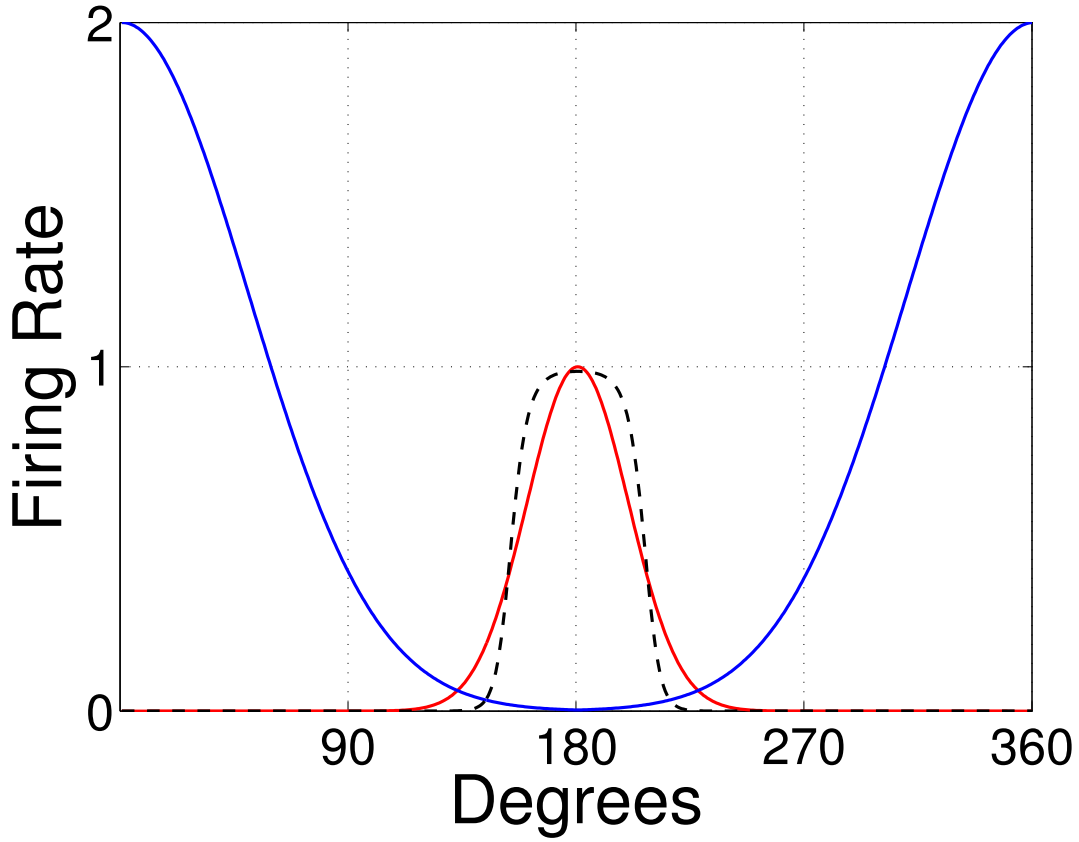
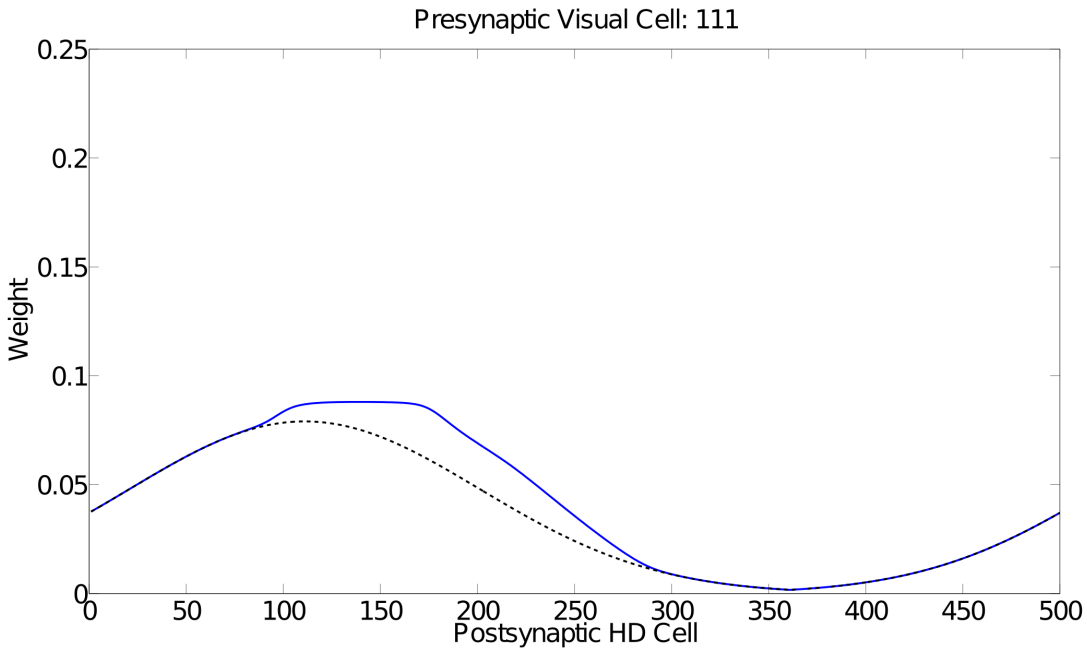


Figure 3.6: Packet locations on the head direction cell ring, showing the initial head direction cell packet (red line), the final head direction cell packet (black dashed line) and an overlay of the equivalent visual packet location (blue line). At this stage, the visual packet is located such that it provides symmetrical input of the same gradient to either edge of the head direction cell packet, which therefore does not move from its initial location.

Changes in the feed-forward weight structure can be more directly linked to final packet location by examining the final weight structure obtained by running a block of simulations with conflicts of the same size as the results in Figure 3.4, but in opposite directions. As can be seen (Figure 3.7), the weight structure is distorted in the locations the packet moves through, in this case the opposite half of the weight structure to that altered when moving the packet in the opposite direction around the ring.



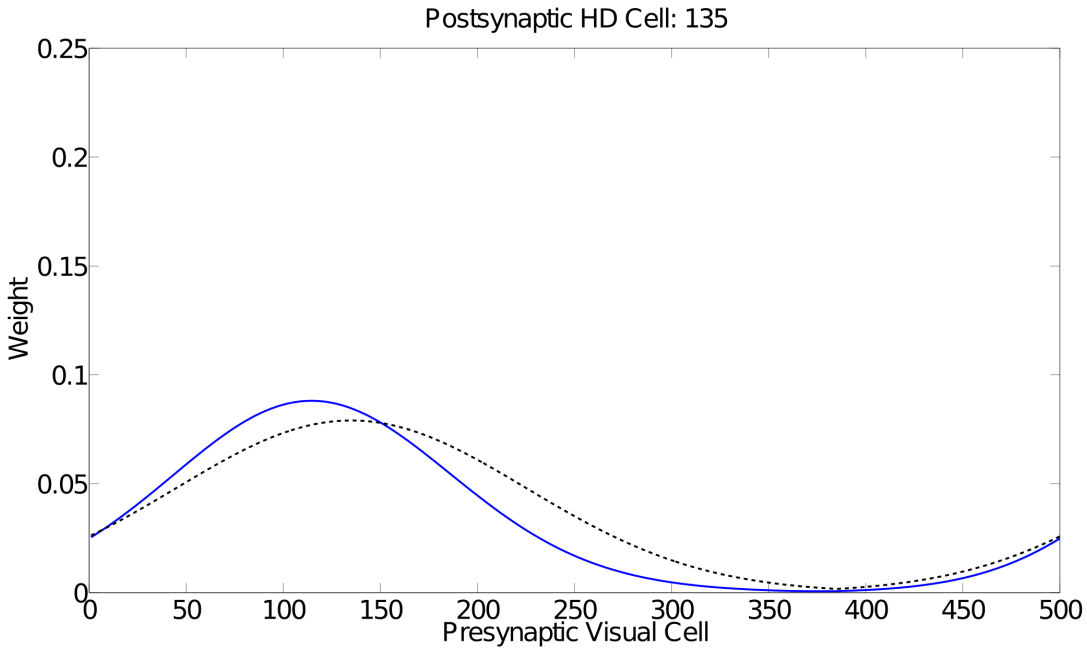


Figure 3.7: Plots showing feed-forward connection weight, analogously to Figure 3.5. As the relative locations of the visual landmark and head direction cell packet are reversed, so too is the distortion of the weight profile relative to that seen in Figure 3.5.

Both of these results can be compared to those obtained from a block of simulations run with static feed-forward weights, which do not update (Figures 3.8 and 3.9). As the visual landmark cue is not remapped onto any of the intermediate head direction cell layer packet locations, the feed-forward weight structure remains unchanged, and, as expected, the packet simply moves to centre itself directly over the input provided by the visual landmark cue. However, at $180\,^{\circ}$ conflict, just as with simulations run with plastic feed-forward weights, the visual cue is located such that it provides symmetrical input to the packet, which therefore does not move.

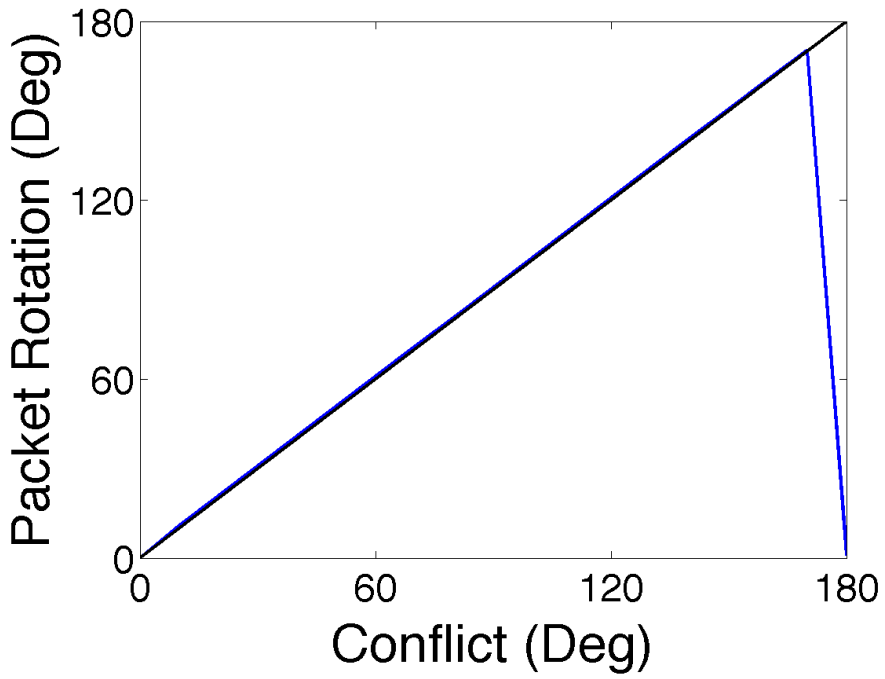


Figure 3.8: Equivalent summary graph to that seen in Figure 3.4, run without learning in w^2 synaptic connections. The abolishment of learning prevents any remapping occurring, with the head direction cell packet always moving to center itself directly over input from the visual cell layer, unless this input is symmetrical as is the case at $180\,^{\circ}$ conflict.

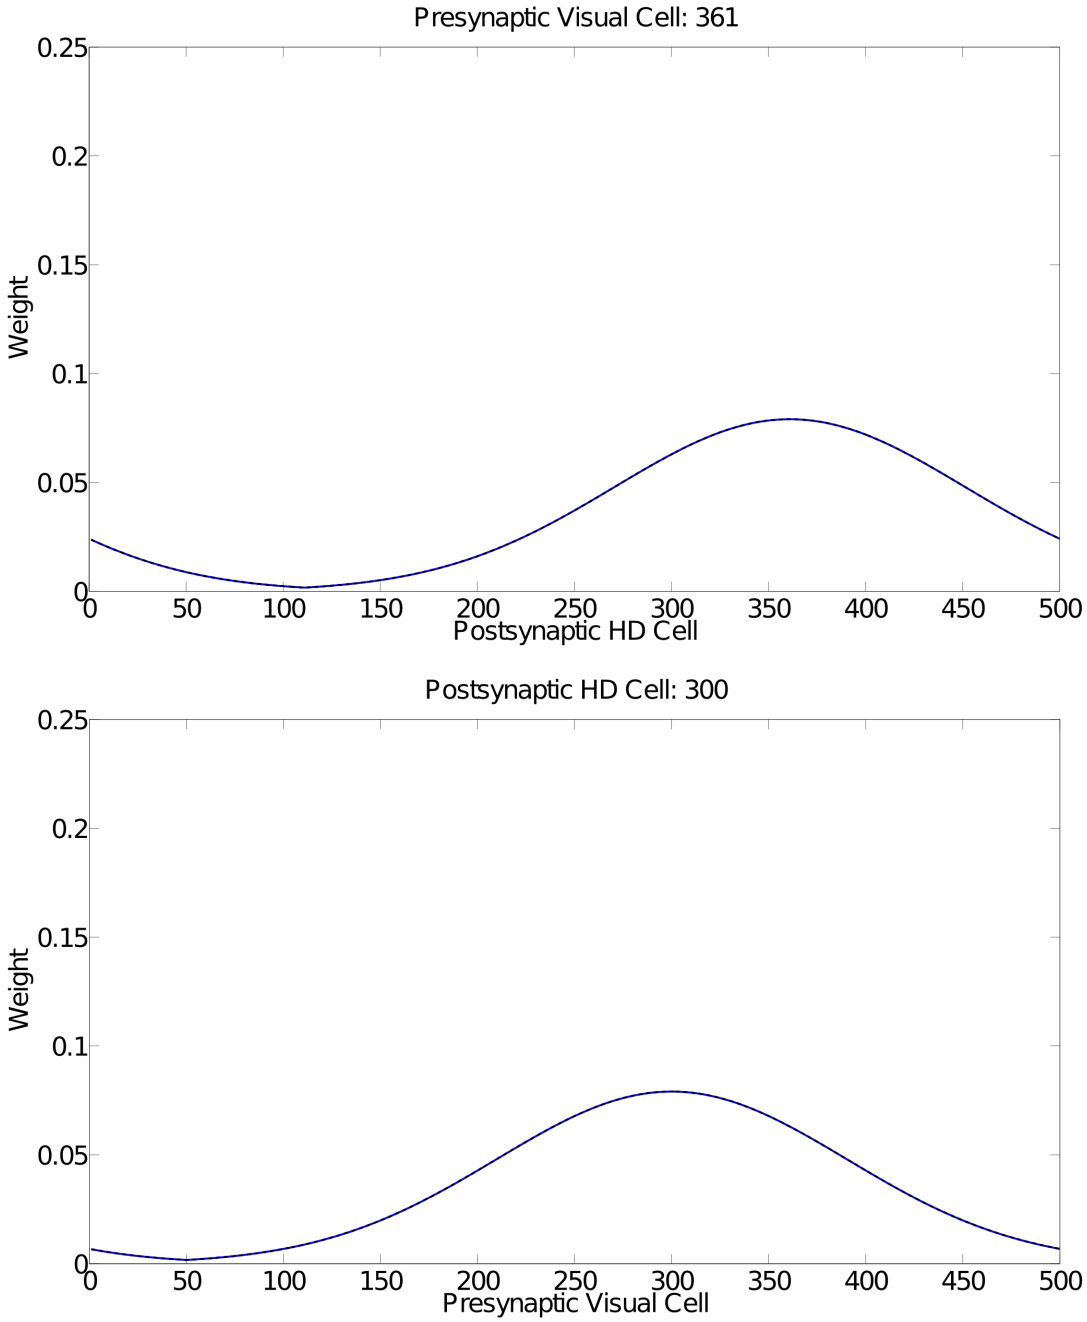


Figure 3.9: Plots showing feed-forward connection weight analogously to Figures 3.5 and 3.7. This confirms that no learning has occurred, and thus weights are the same before and after simulation.

3.5 Saliency of visual information

Figure 3.10 shows the effect of varying ϕ_2, which controls the overall scaling of w^2 feed-forward weights. This is equivalent to a change in salience of visual landmark cues, which more salient cues having a larger value of ϕ_2 and less salient cues a smaller value of ϕ_2. Lowering ϕ_2 causes visual input to have less influence on the head direction memory packet. Thus, the effect of remapping is relatively strong compared to visual input, and final packet location is closer to that of the original memory packet. This weakened visual influence therefore results in a breakpoint at larger conflict sizes and a larger undershoot.

For increasing values of ϕ_2, visual information has a stronger influence on the head direction activity packet. This means the slowing effect of remapping is relatively weakened, resulting in smaller undershoot sizes and a breakpoint at larger conflict sizes. Eventually, for $\phi_2 = 8$ onwards, visual information has such a strong effect that the undershoot and breakpoint are totally eliminated, and final packet location is determined solely by visual cues.

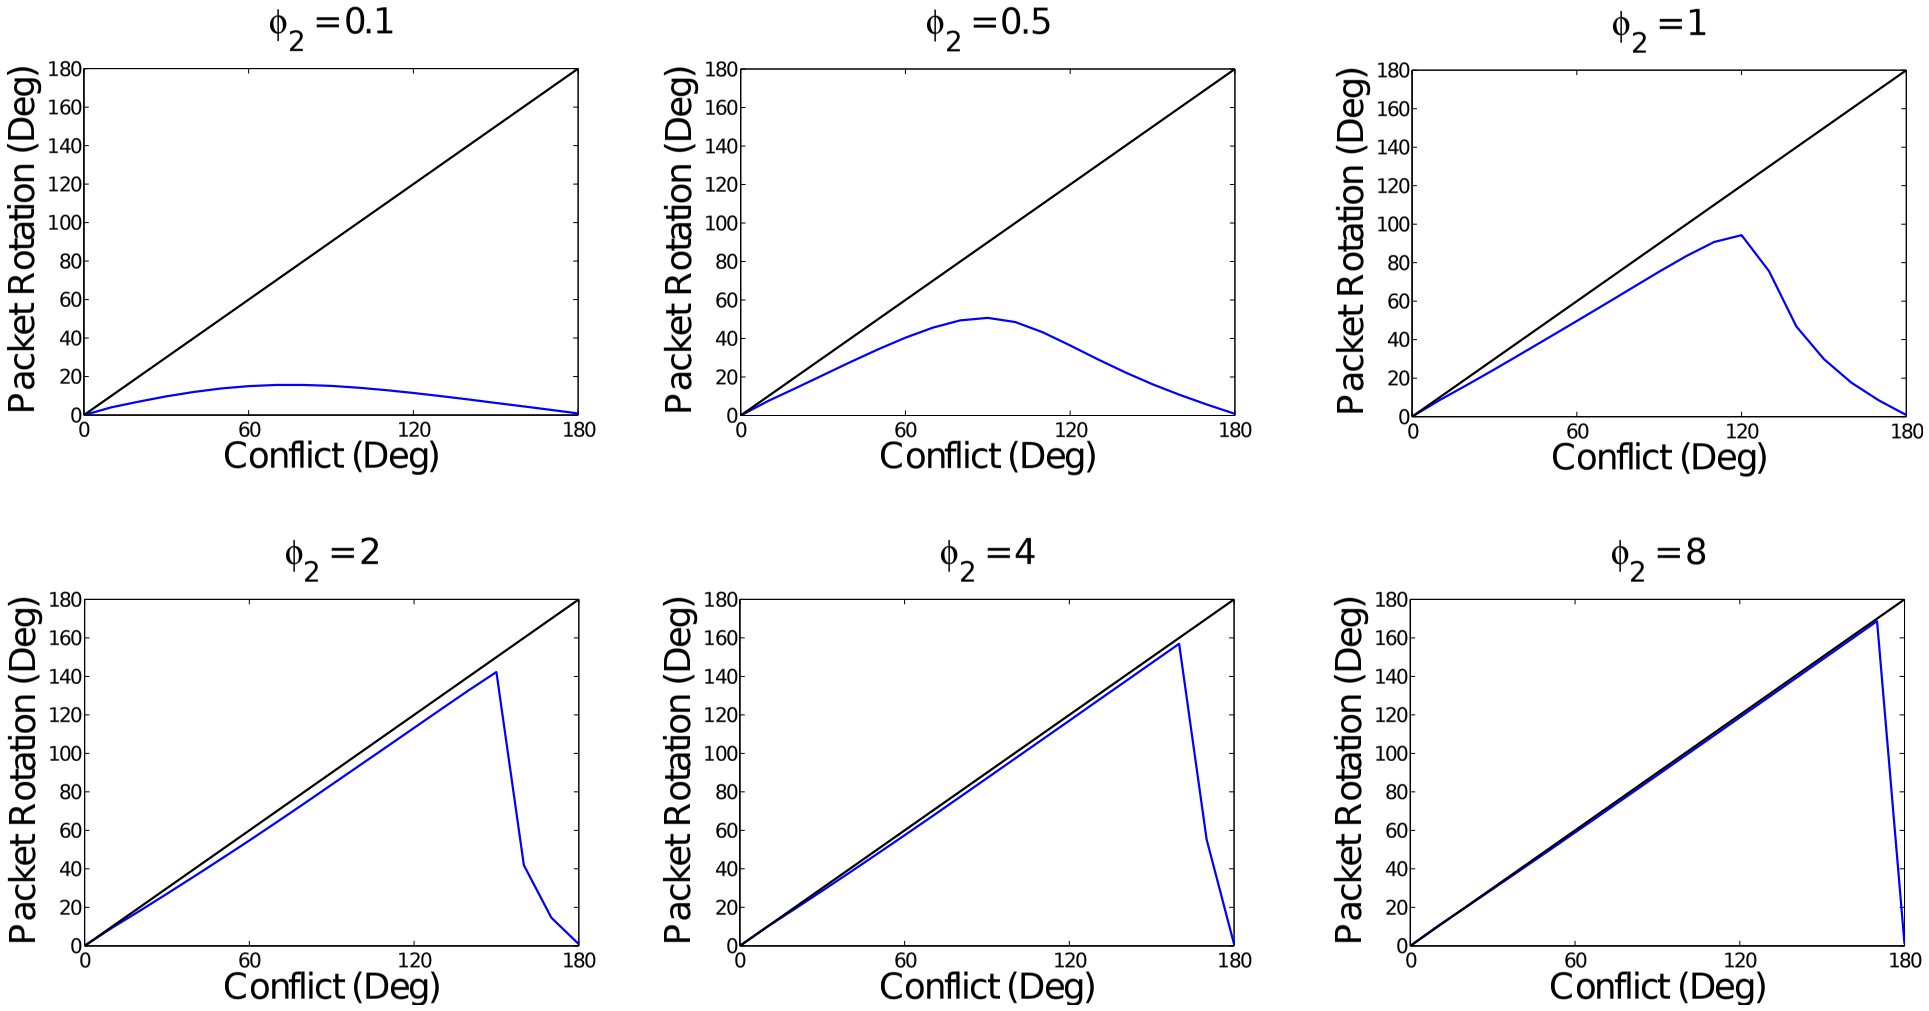


Figure 3.10: Plots showing results for simulations in which the salience of visual input was varied by altering ϕ_2. At lower salience, the system favours ideothetic information, with larger undershoots and earlier breakpoints. As salience increases, visual cues come to dominate cue conflict, with larger breakpoints and smaller undershoots, until the final packet location is determined almost entirely by visual input for $\phi_2 = 8$ and above.

3.5.1 Introduction of homosynaptic LTD

Figure 3.11, shows a summary graph of for a block of simulations run with identical parameters to those previous presented, but updating feed-forward weights via homosynaptic LTD rather than simple Hebbian learning (see Equation (3.6)). This has the effect of not only strengthening connections between co-active head direction and visual cells, but also weakening the connections from active visual cells to inactive or weakly active head direction cells.

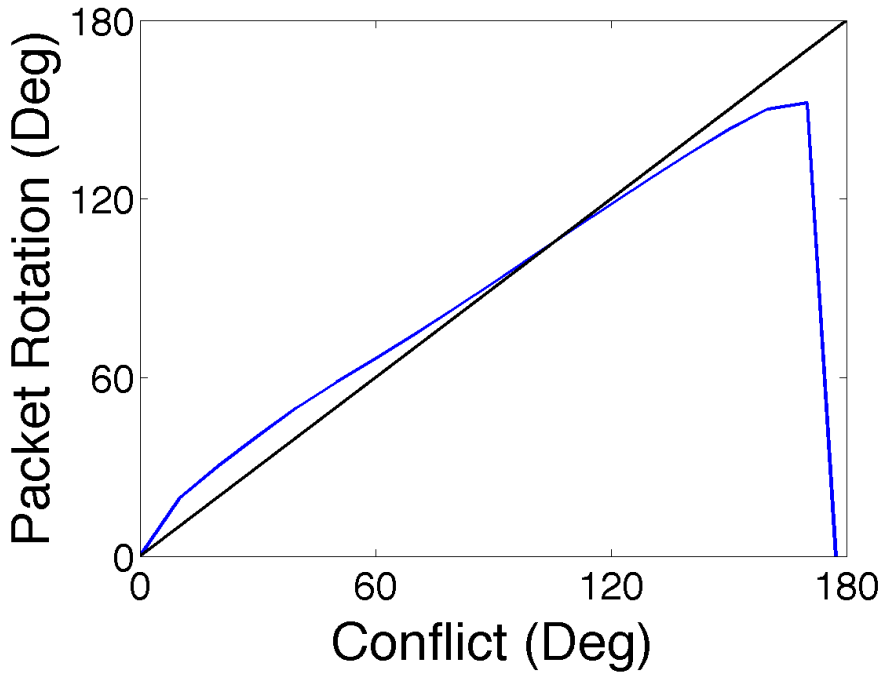


Figure 3.11: Summary graph showing final head direction cell packet locations (blue line) for increasing conflict sizes, in contrast with expected packet location given total visual dominance (black line). Compared to a summary graph for simulations run without LTD (Figure 3.4), we see an initial overshoot of the visual packet location, followed by an undershoot and finally a reversion to reliance on path integration mechanisms.

As seen in the summary graph, the initial trend is for the final packet location to overshoot, rather than undershoot, the location of the visual stimulus. For larger conflicts, this overshoot consistently reduces, before reverting to undershooting behaviour. As the packet moves to centre itself over the vi-

sual input, the connections between currently active visual cells and head direction cells whose firing is below average will be weakened (see Equation (3.6)). This will effectively weaken the connections between the currently active visual cells and the head direction cells at the edges of the activity packet, whose firing is weaker than those cells towards the centre of the packet. Thus, the weights from the currently active visual cells to the leading and trailing edges of the head direction cell packet will be weakened. However, weights will be strengthened between currently active visual cells and head direction cells towards the centre of the packet, whose firing is strongest.

The net effect of this change in weights ought to be to stabilise the packet in its current position. However, at smaller conflict sizes, the influence of the visual landmark will provide strong asymmetric input to the activity packet, drawing it towards a corresponding location on the head direction cell ring. Thus, as the packet moves towards this location, the weights from the visual landmark location onto each position on the ring it passes through will experience a combination of weakening and strengthening. Initially the weights to this location will be weakened, as they are on the leading edge of the packet where head direction cell firing is below average. They will then be strengthened as they are in the centre of the packet, where cell firing is above average, and finally strengthened when they are on the trailing edge of the packet where cell firing drops below average. The net effect is to weaken the connections from the current visual landmark location to each location through which the head direction cell packet moves. This weight profile distortion can clearly be seen in Figure 3.12. As the packet reaches the location within the head direction cell ring which corresponds to the visual landmark location, the impetus to move through the ring will be decreased and overpowered by the effects of weight update encouraging the packet to remain stationary. This will have an effect of halt-

ing packet motion. However, there is still an impetus for the packet to continue moving through the ring, provided by the combined weakening of weights onto all locations the packet has been through. Thus, the slowing effect of strengthening weights onto the packets centre, and weakening those onto the packets edges is ameliorated somewhat, and the packet therefore overshoots the expected location before it eventually comes to a halt.

As the size of conflict increases, the influence of the visual cue is reduced, and thus the relative strength of the influence of weight updates increases, and there is a greater impetus for the packet to remain stationary. Thus the distance the packet travels through the ring is reduced and, as conflict sizes increase, the overshoot decreases before eventually becoming an undershoot. At 180 degrees of conflict, the visual cue provides symmetric input to the head direction cell activity packet. This means that there is nothing causing the packet to move, and weight updates causing it to remain stable. Thus, the packet does not move from its original location.

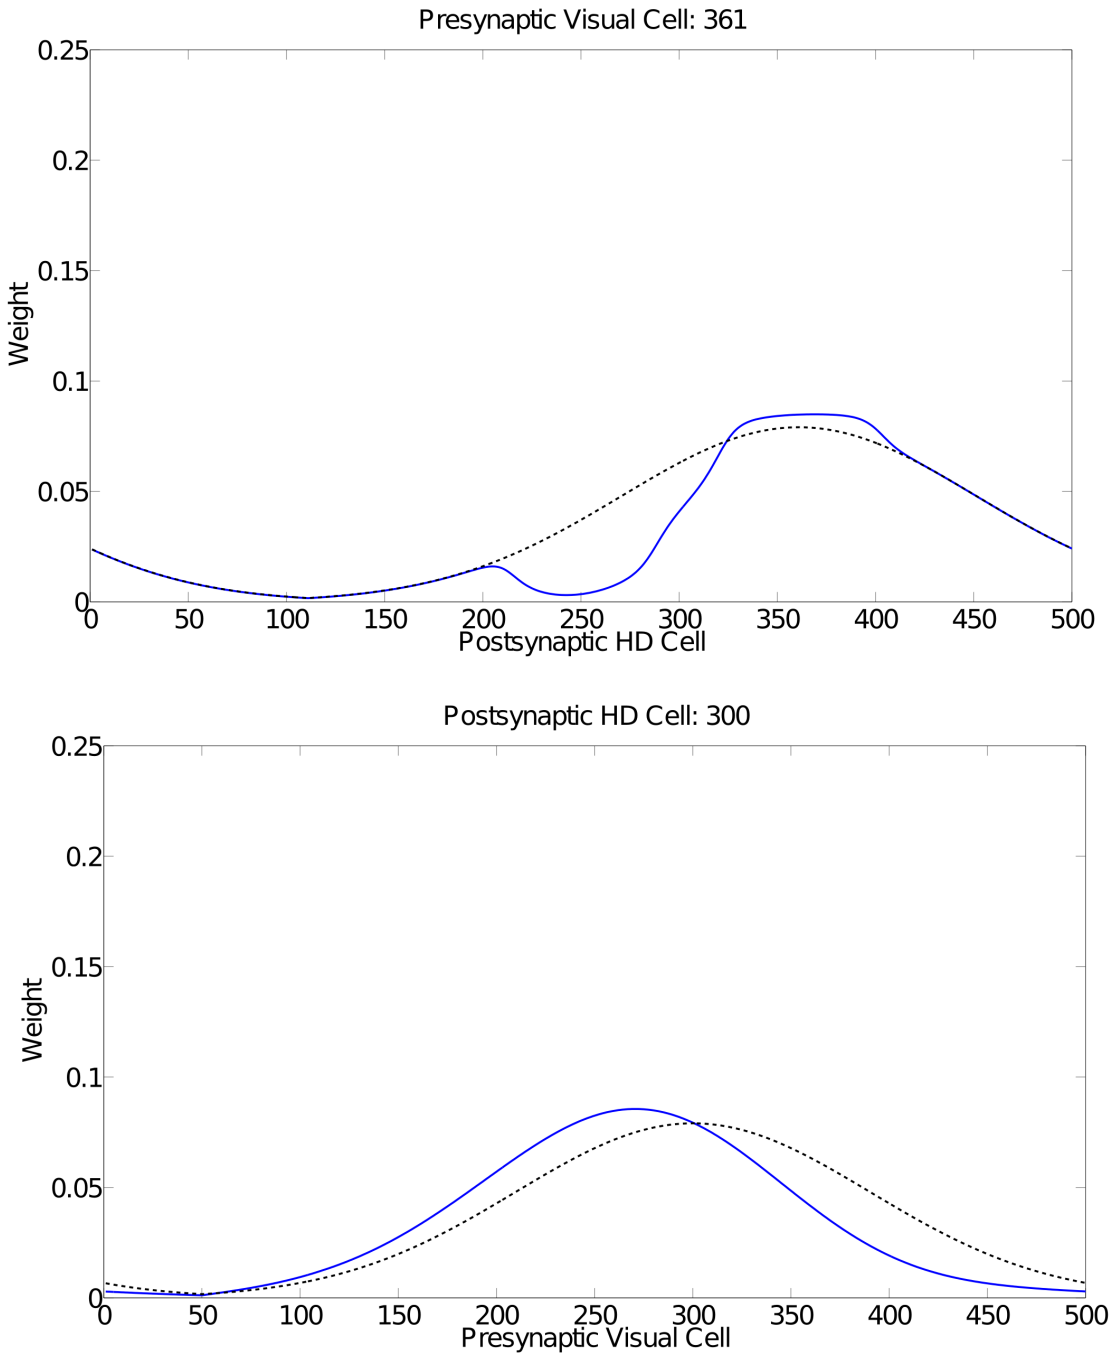


Figure 3.12: Plots showing feed-forward connection weight, analogously to Figures 3.5, 3.7 and 3.9, for simulations with homosynaptic LTD summarised in Figure 3.11. Of particular note is the atrophy of the weight profile from the visual landmark location onto the HD layer, seen on the left hand side of the top profile. Connectivity from this cue to intermediate HD locations has been weakened, a net effect of initial weakening, followed by strengthening, and finally by another weakening. This causes the HD packet to overshoot the visual landmark location.

3.5.2 Effect of learning rate manipulations

Given that this models performance appears to be critically dependent on changes in weight structure, a full investigation into the effects of the learning rate parameter k in the range [0.01, 0.1] was carried out. This investigation was performed with a Hebbian learning rule, according to Equation (3.5) with weight normalisation as in Equation (3.7). This reveals that model behaviour is fairly robust to changes in learning rate, in that it does not break down entirely. However, changes in learning rate have a significant effect on core model operation, and thus it seems that the learning rate parameter needs to be carefully tuned in order to replicate experimental data.

Figure 3.13 shows summaries for blocks of simulations run with learning rates of, from left to right 0.01, 0.025, 0.05, 0.075 and 0.1 respectively. For lower learning rates, the extent of remapping that occurs during the time that packet is in motion is relatively less. This is reflected in a smaller undershoot for small conflicts, which increases with increasing learning rate. Another important effect of increased remapping extent via increased learning rate is that the location of the breakpoint, beyond which remapping is stronger than the pull of the visual cue, changes. Specifically, increasing learning rate increases the effect of remapping thus causing the breakpoint to occur earlier. For example, the breakpoint occurs at a conflict of 160° for a low value of LR = 0.01, at 120° as already seen with an intermediate value of LR = 0.05 and at conflicts as small as 90° for a high learning rate of LR = 0.1. As learning rate decreases and the breakpoint gets closer to the maximum possible value of 170° conflict, the average gradient of the data beyond the breakpoint increases. These two effects combine to give the trend shown in Figure 3.13, where increasing learning rate increases the undershoot size, whilst also shifting the breakpoint to smaller conflict sizes and lessening the post-breakpoint gradient.

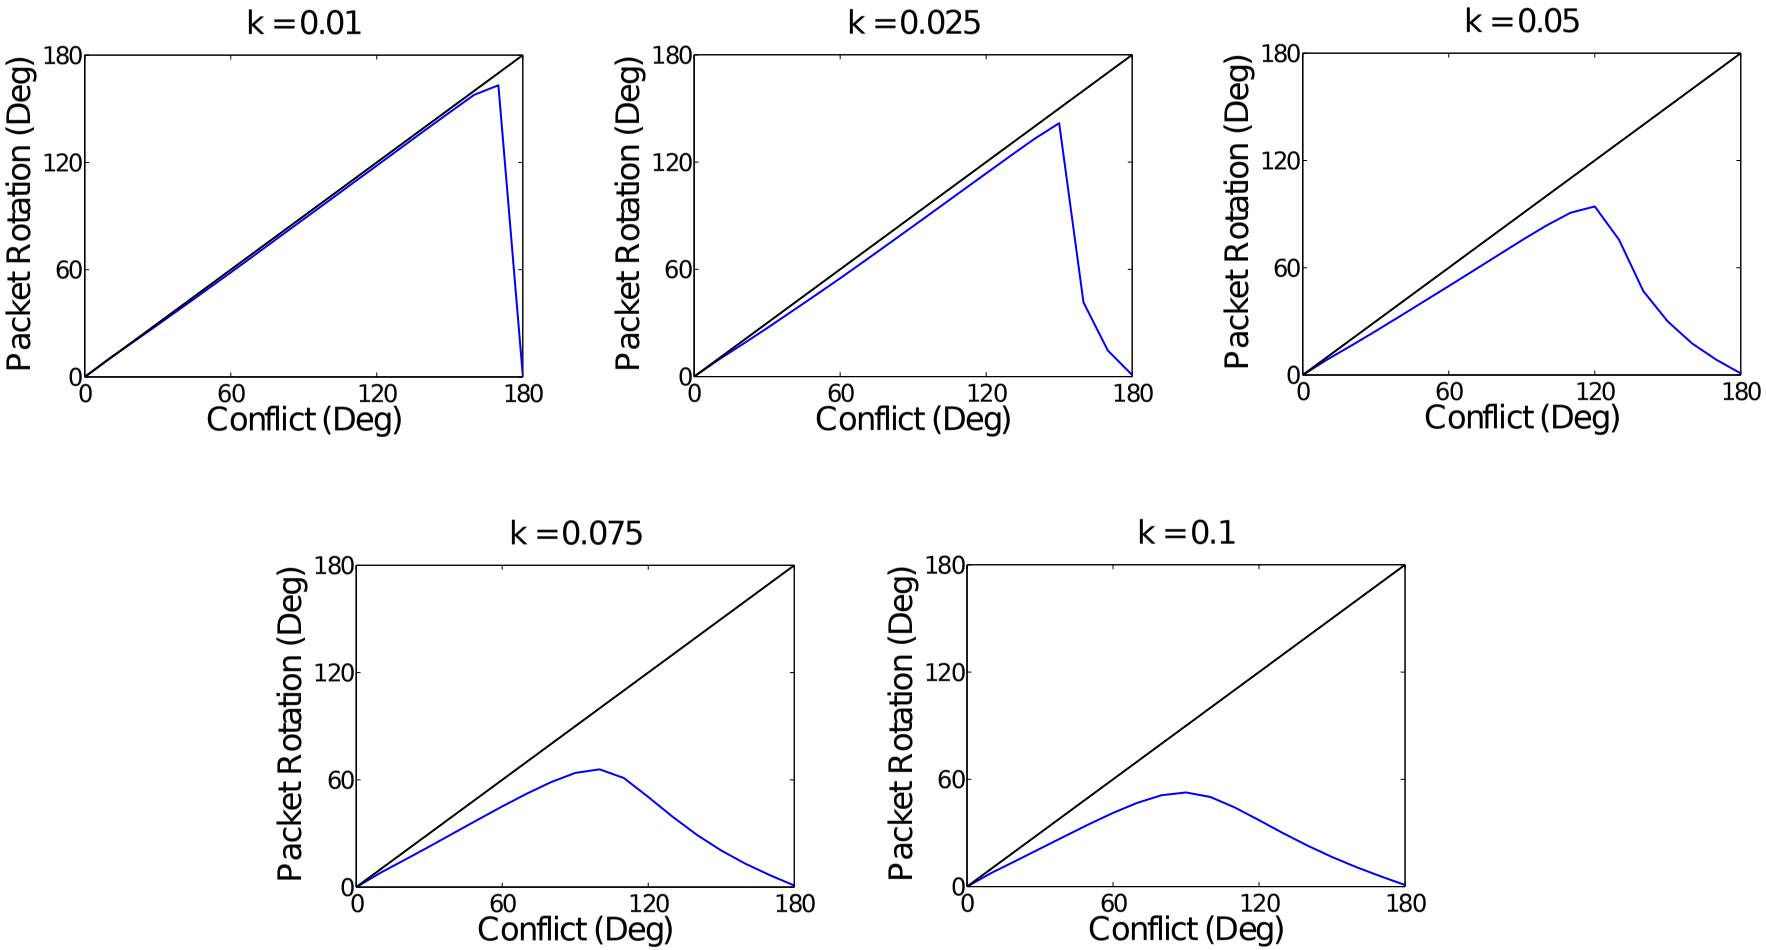


Figure 3.13: Summary plots for blocks of simulations run with learning rates of (from left to right) 0.01, 0.025, 0.05, 0.075 and 0.1. The summary for LR = 0.05 is a replication of Figure 3.4, included for ease of comparison. Increases in learning rate results in a strengthening of the remapping effect, manifesting as larger undershoot sizes and breakpoints at smaller conflict sizes.

3.5.3 Varying spread of feed-forward and recurrent weights

Summary results obtained by holding the width of the recurrent weights constant at $\sigma^{RC} = 35$ and increasing the width of feed-forward weights σ^{FF} are shown in Figure 3.14. Increasing σ^{FF} beyond the initial value of $\sigma^{FF} = 65$, we see an increase in the size of undershoot for smaller conflict sizes, and a small shift in the breakpoint towards smaller conflicts. This could initially be seen as contradictory: increasing the spread of feed-forward weights ought to increase the influence of the visual landmark cue the head direction cell activity packet, thereby reducing undershoot size and shifting the breakpoint to larger conflict sizes. However, broadening the feed-forward weights also widens the range over which plasticity can occur, increasing the effect of remapping. This stronger remapping effect outweighs the increased visual influence, producing the observed effect. It should be noted that, at $\sigma^{FF} = 100$, the width of visual influence creates a final head direction cell packet that is too wide to accurately represent head direction cell firing profiles as seen *in vivo*. This result is presented solely in order to be able to provide enough data to extrapolate a trend for the networks performance beyond $\sigma^{FF} = 65$.

Decreasing σ^{FF} below 65, as seen in Figure 3.15, we see a switch to complete visual dominance, beyond which further reductions in σ^{FF} have no effect. Once feed-forward weights are this narrow, they no longer provide any input to the current packet location within the head direction cell ring. Instead, a second packet is driven up in a new location within the head direction cell ring, which perfectly corresponds to the centre of the visual packet, and inhibits the previous head direction cell packet. With the parameters used previously, this new packet is not fully driven up in the timecourse of simulation and a weak diffuse packet is seen. However, increasing the strength of the visual input

drives this packet up fully (Figure 3.16). This is classic continuous attractor behaviour (Zhang, 1996; Redish and Touretzky, 1997; Samsonovich and McNaughton, 1997; Redish, 1999; Touretzky, 2005).

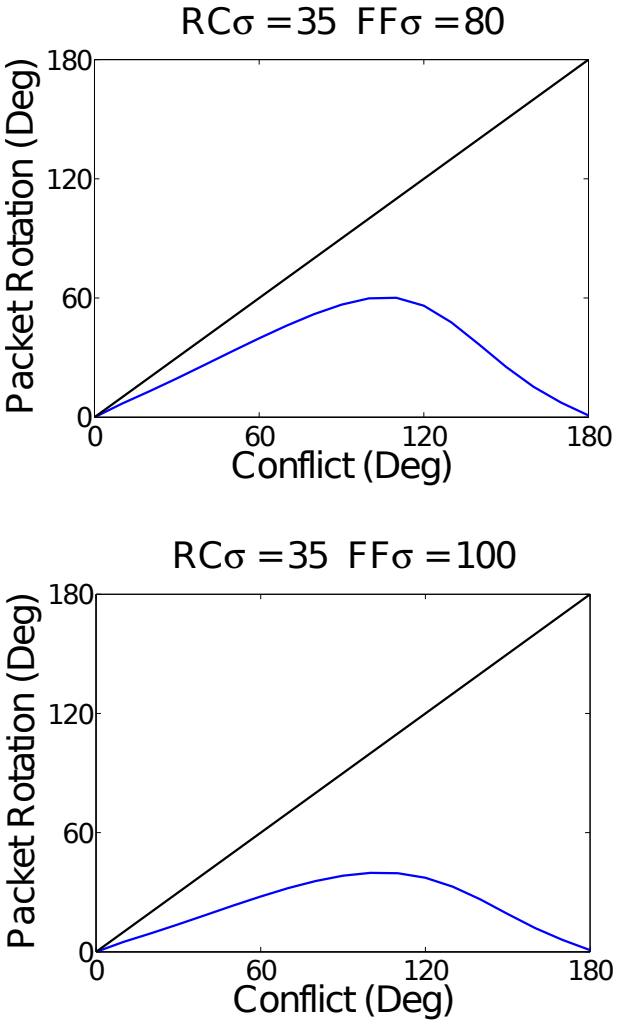


Figure 3.14: Summary plots for blocks of simulations run with increasing values of σ^{FF}. Remapping width outweighs the stronger visual input, resulting in an increased undershoot and breakpoints at earlier values, relative to original results in Figure 3.4.

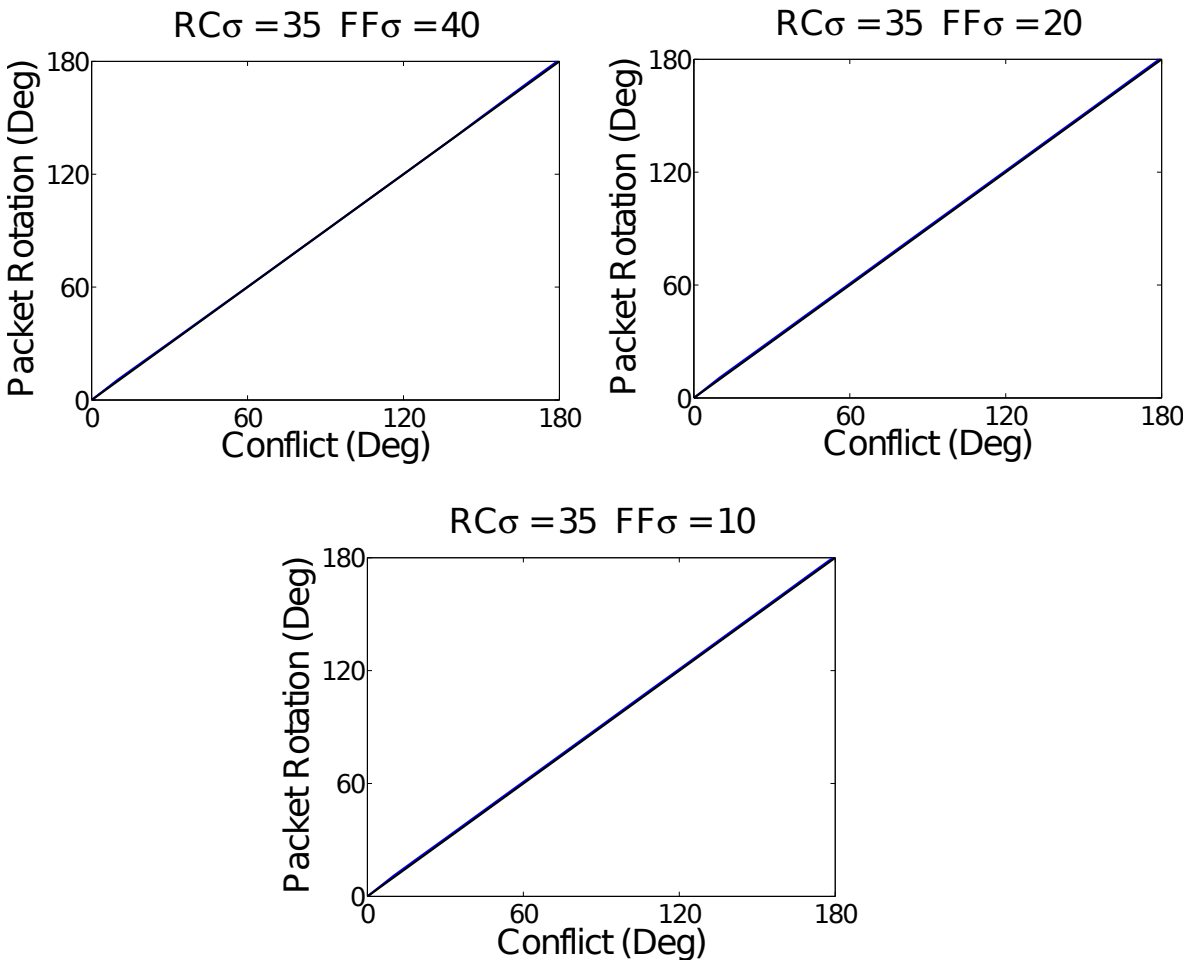


Figure 3.15: Summary plots for blocks of simulations run with decreasing values of σ^{FF}. As feed-forward weights become too narrow to provide direct stimulation to the head direction cell activity packet. Thus a new packet is driven up within the head direction cell ring in a location corresponding exactly to the location of the visual input, extinguishing the old packet.

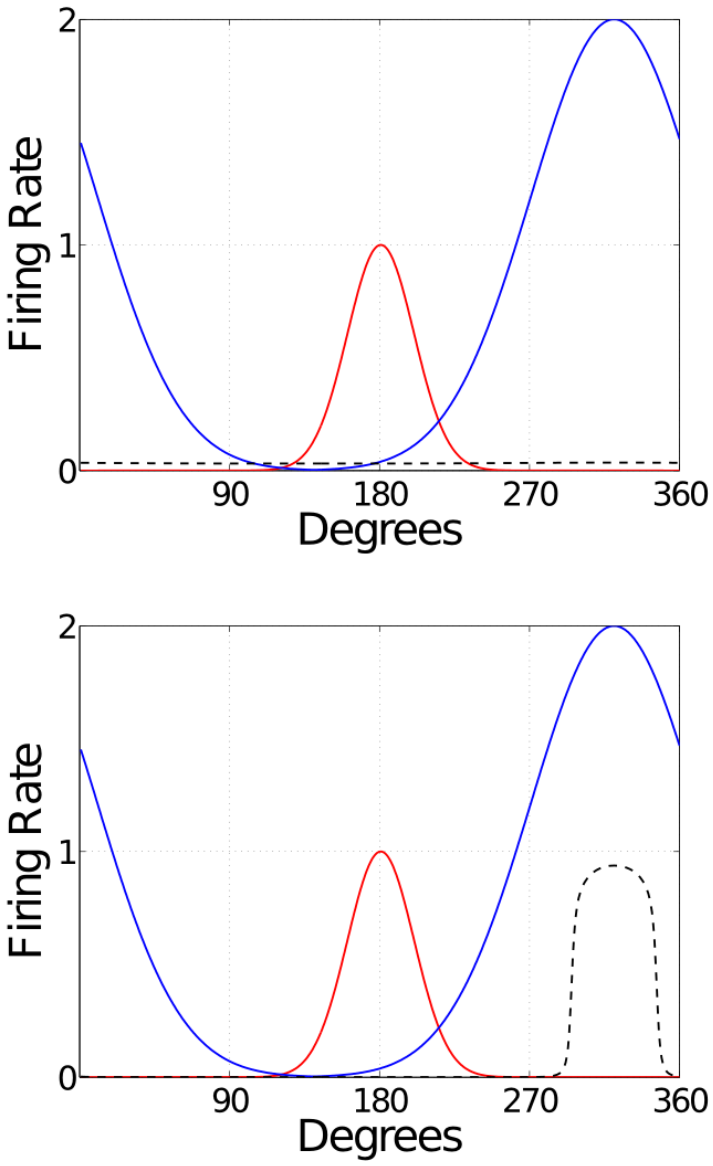


Figure 3.16: Final packet locations for simulations run with weaker (top) and stronger (bottom) visual inputs, showing initial head direction cell activity packet (red line), the final head direction cell packet (black dashed line) and an overlay of the equivalent visual packet location (blue line). A very weak and diffuse final head direction cell packet can just be seen with weaker visual input, whilst strong visual input allows a packet in the new location to be driven up fully.

The trend for low values of σ^{FF} is similar, but not perfectly the same for those with no learning: the final head direction cell packet location always follows the location of the visual landmark cue. This is because the packet no longer moves smoothly due to visual input to the original packet, and thus no

remapping will occur. Instead, a packet is driven up in the new location, simultaneously suppressing the previous head direction cell packet. This behaviour is independent of distance between the original and new locations, and thus the head direction cell packet jumps to the new location.

Summary results are also shown, in Figure 3.17, for simulations with the feed-forward weight parameter held constant at $\sigma^{FF} = 65$ and varying the width of the recurrent weights, σ^{RC}. Smaller values of σ^{RC} result in a larger undershoot and a breakpoint at smaller conflict sizes, whilst increasing sigma reduces undershoot size and shifts the breakpoint up to larger conflict sizes. This is because the extended range of recurrent collateral connections resulting from larger σ^{RC} values facilitates the ease with which a head direction cell packet may move through the ring. Therefore, a larger value of σ^{RC} accentuates the effect of the visual landmark cue on packet motion and increasing it increases the visual influence, reducing undershoots and shifting breakpoints back towards larger conflict sizes. Once weights are so broad that they span the entire ring ($\sigma^{RC} = 80$ and beyond), the final head direction cell activity packet will be so broad that any stimulation has the effect not of shifting the packet, but distorting the currently active packet to have a center that corresponds to the location of the visual input. Thus no shifting or accompanying remapping occurs, and the packet location will always appear to be centered over the visual input.

However, this effect has an upper limit, in that increasing σ^{RC} will also cause the current activity packet to be broadened and flattened. Beyond $\sigma^{RC} = 60$ this results in a packet that is sufficiently diffuse and weak to no longer be acceptable as biologically plausible. This flattening of the head direction cell activity packet in response to increasing values of σ^{RC} can be seen in Figure 3.18.

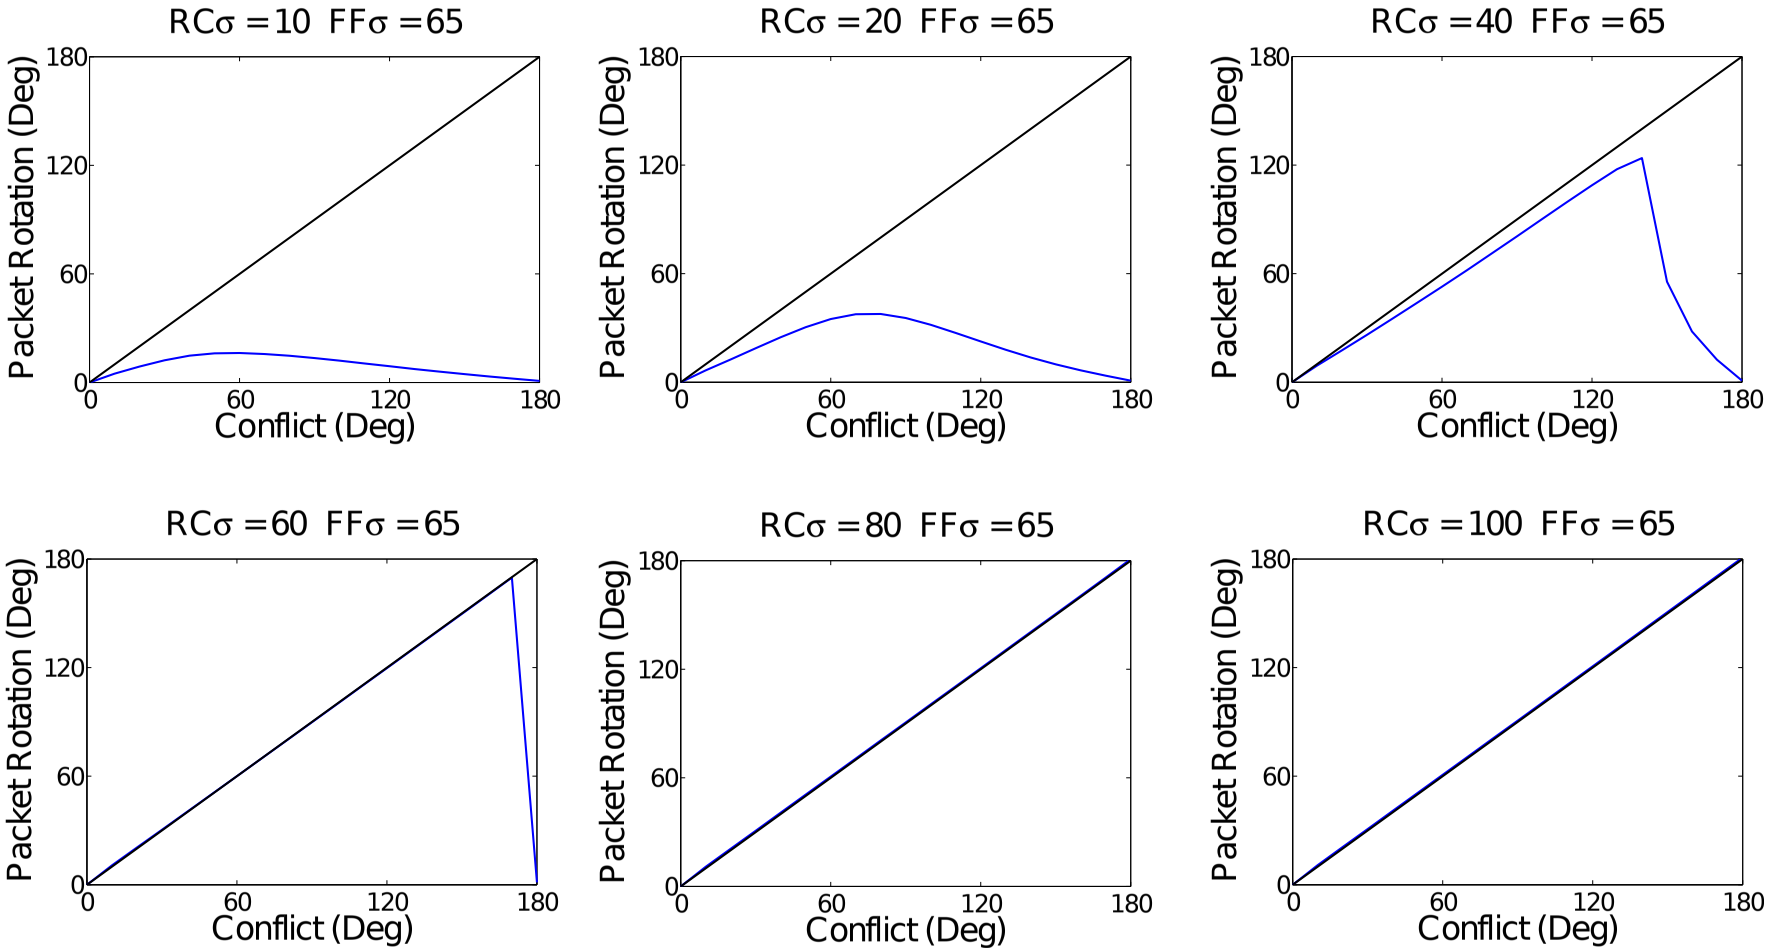


Figure 3.17: Summary plots for blocks of simulations run with a range of values for σ^{RC}. Increasing the width of recurrent collateral connectivity allows for easier head direction cell packet motion, causing a reduction in the undershoot and breakpoints at larger conflict sizes. This trend continues until $\sigma^{RC} = 80$, beyond which the packet is no longer moving through the ring, but is distorting to shift it's center over the location of the visual input.

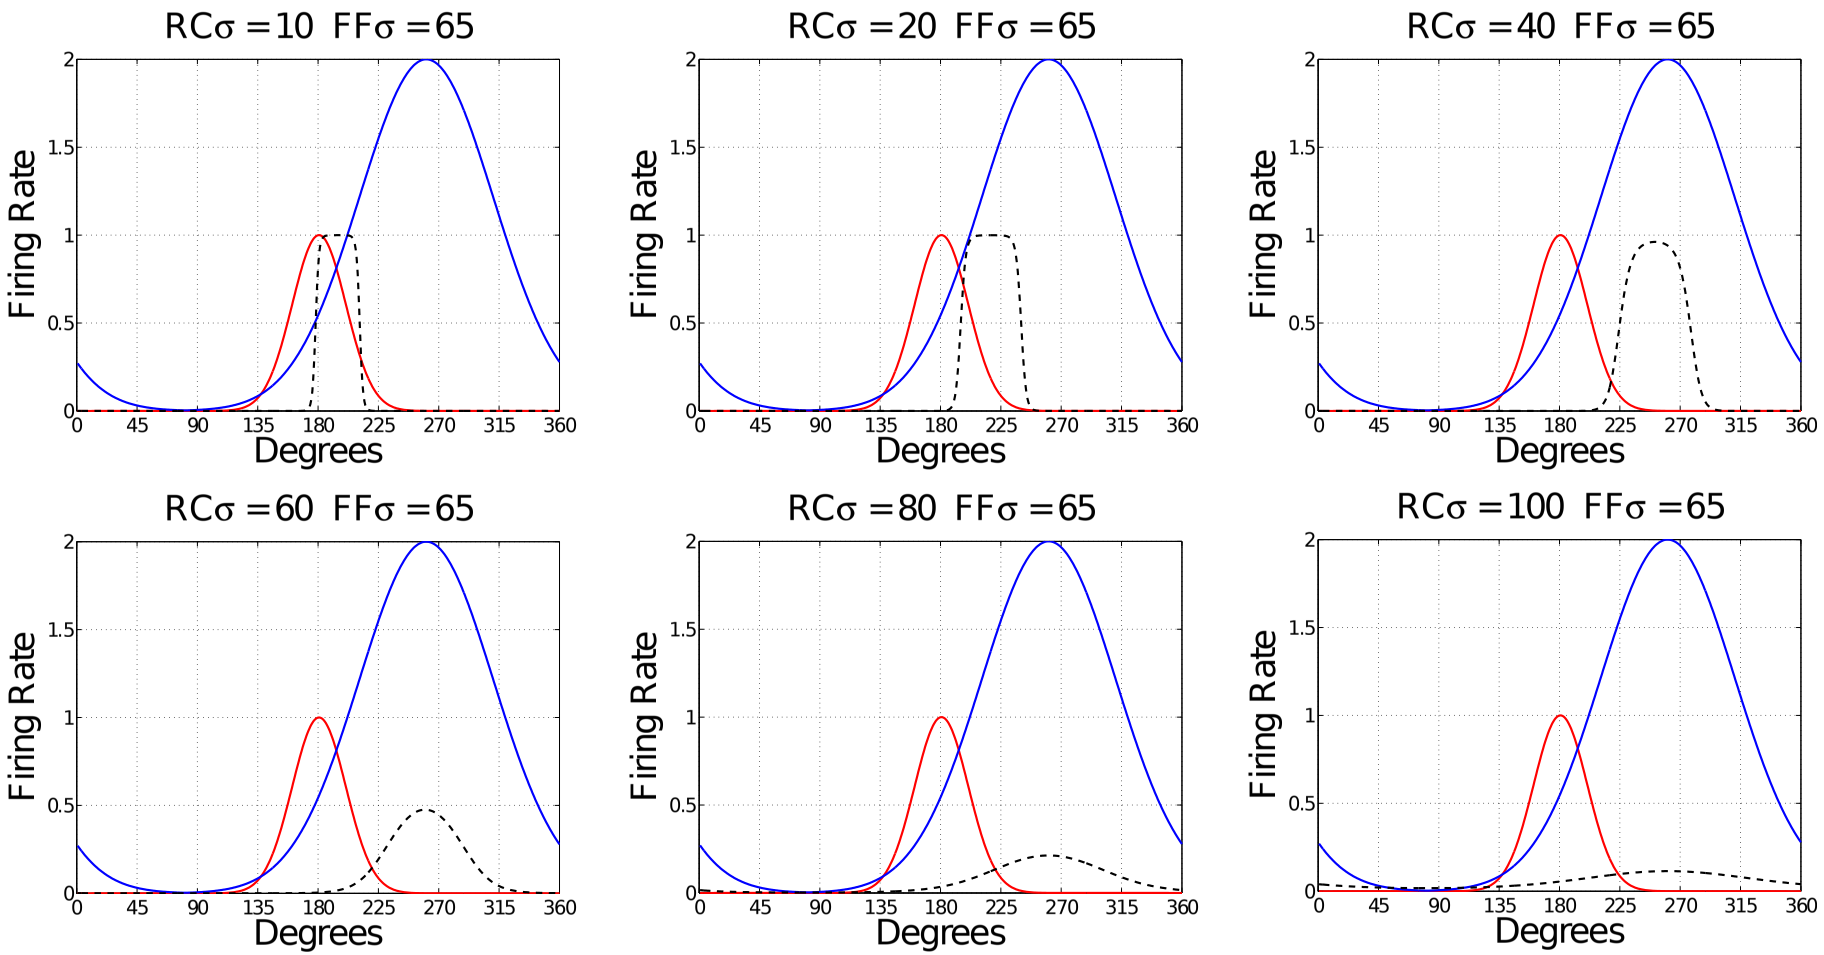


Figure 3.18: Final packet locations for increasing values of σ^{RC}, showing initial head direction cell activity packet (red line), the final head direction cell packet (black dashed line) and an overlay of the equivalent visual packet location (blue line). Final head direction cell packet can be seen becoming more diffuse with increased width of recurrent collateral connectivity.

3.5.4 Varying packet sizes

Figure 3.19 shows summary results for simulations with σ^{HD} held constant, whilst increasing σ^{VIS}, which directly controls the width of the visual landmark cue packet. For smaller packets, there is less influence on the head direction cell packet, and thus remapping is relatively stronger. This is revealed through a larger undershoot, and breakpoints at smaller conflict sizes. As the width of the visual packet increases, so too does its influence on the head direction cell packet, and thus we see a reduced undershoot with breakpoints at greater conflict sizes. However, broadening the visual packet also broadens the range over which remapping takes place. Eventually, the effect of a broader remapping outweighs the effect of an increased visual influence. This is manifested when $\sigma^{VIS} > 60$, in a reversal of the general trend of results, with an increasing undershoot and breakpoints occurring at smaller conflict sizes than were seen at $\sigma^{VIS} = 60$.

It is noteworthy that this trend is similar to that seen when increasing σ^{FF}. However, the trend in results for σ^{VIS} is relatively weaker. This is because although the packet is broadening, which increases the strength and range of visual influence, it is the feed-forward weights that more directly determine the extent of visual landmark cue influence, and the width of the remapping effect. Figure 3.20 shows summary results for simulations with σ^{VIS} held constant, whilst σ^{HD} is varied. Relatively fewer results exist, because although visual cues can theoretically vary in size, the packet in the attractor ring must be simulated at a size that is within the range of packet sizes seen *in vivo*.

In contrast to the clear effects of varying σ^{VIS}, σ^{FF} and σ^{RC}, it seems that varying σ^{HD} has little to no obvious effect. This demonstrates that the crucial factors in determining head direction cell

packet location are the strength and range of visual landmark cue projection into the head direction cell ring. This is in particularly stark contrast to the effect of σ^{RC}, which also acts on the head direction cell packet. However, this can be explained in terms of the recurrent weight structure, which is unaffected by σ^{HD}, and thus no effect of encouraging packet motion is seen when varying the latter parameter.

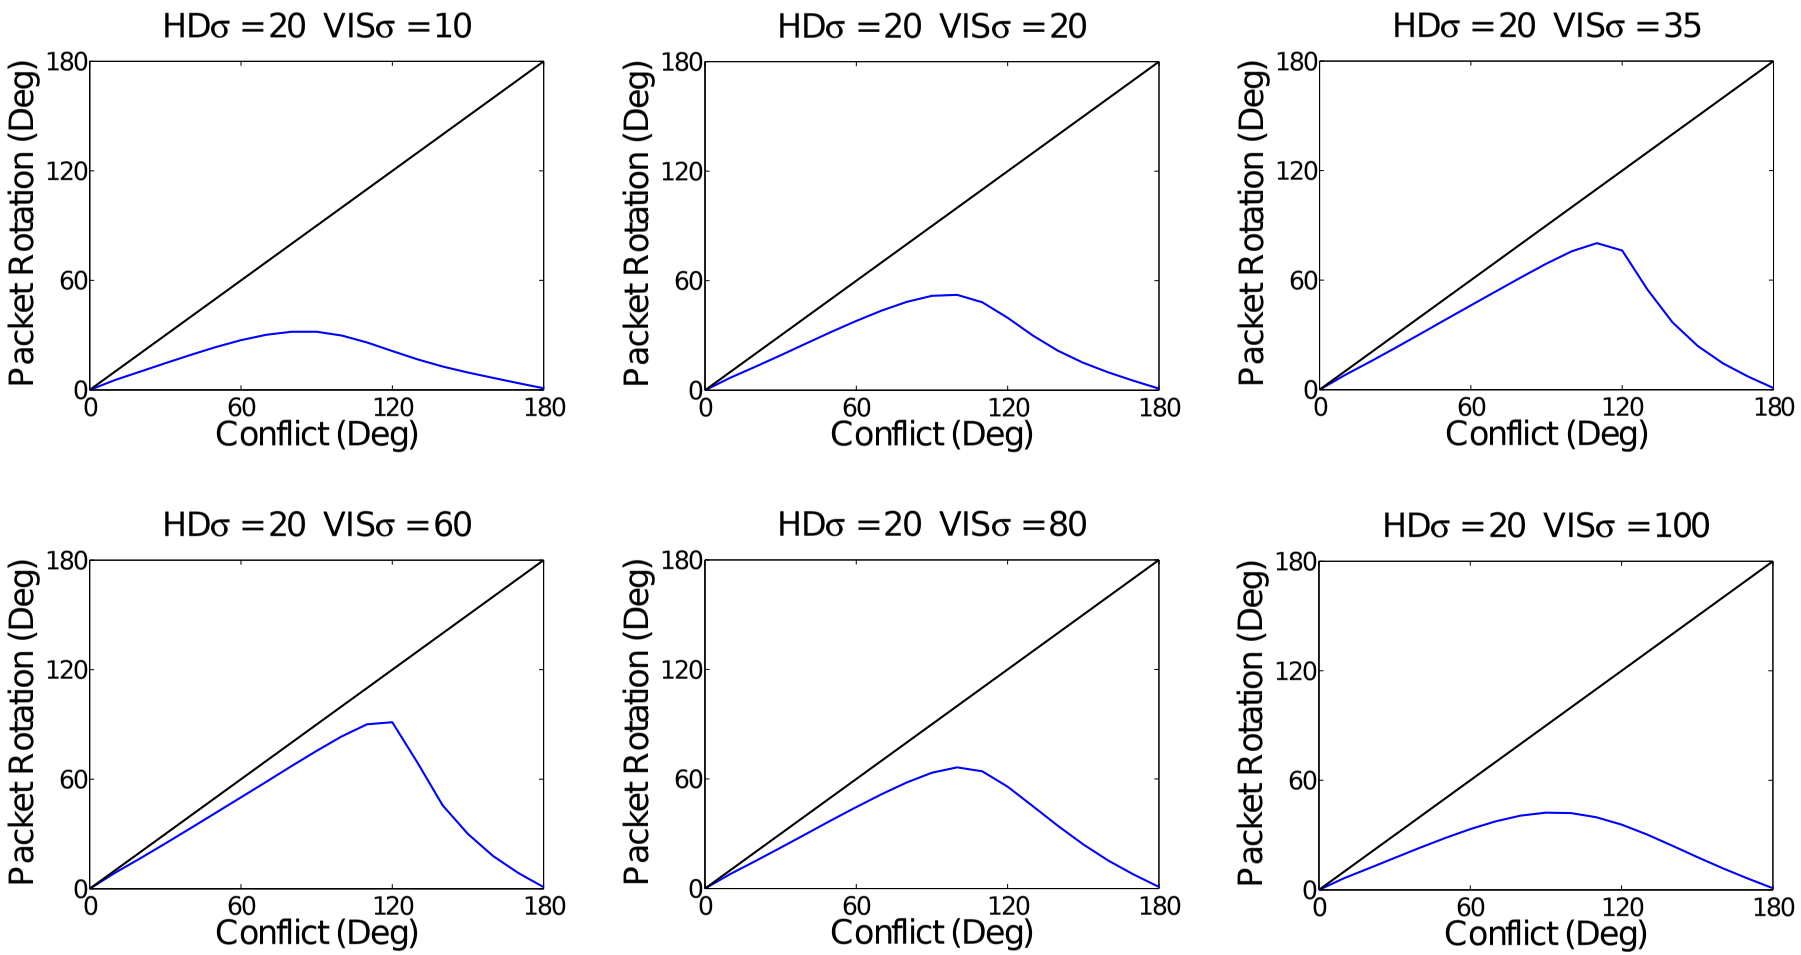


Figure 3.19: Summary plots for blocks of simulations run with a range of values for σ^{VIS}. Initially, increasing the width of the visual activity packet causes smaller undershoots and later breakpoints. However, beyond $\sigma^{VIS} = 80$ the effect of increased visual breadth is outweighed by the increased distance of remapping. Thus undershoots begin to increase and breakpoints begin to appear earlier.

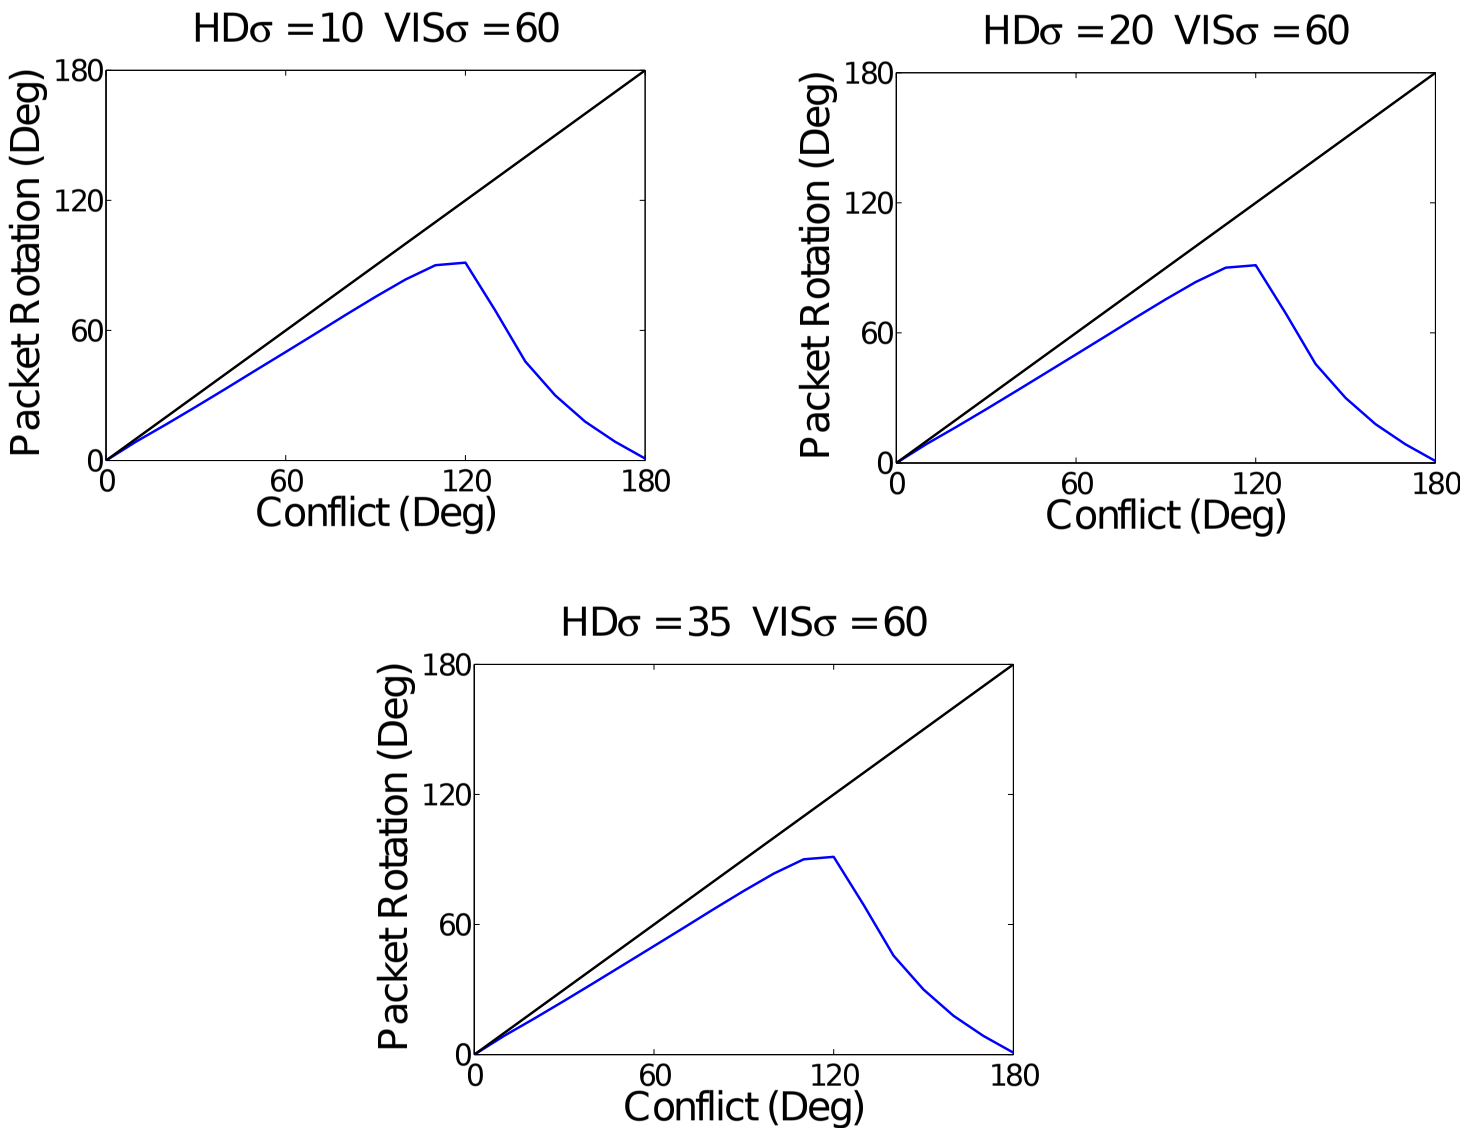


Figure 3.20: Summary plots for blocks of simulations run with a range of values for σ^{HD}. This appears to have little effect on overall model functionality.

3.6 Discussion

This chapter provides details of and simulation results from a model which is able to resolve conflicts between ideothetic and visual landmarks cues in an ecologically valid and biologically plausible manner. Experimental evidence demonstrates that visual landmarks and ideothetic information are integrated into a single reliable head direction signal. The exact nature of this integration has been explored via situations in which these cues are in conflict. A strong case is made for visual landmarks exerting a dominant influence over head direction cell preferred directions for small conflict sizes, with a reversion to reliance on ideothetic information for larger conflict sizes.

A model is presented consisting of a ring of visual cells coupled to a head direction cell continuous attractor layer by plastic feed-forward connections. The model is tested at various sizes of conflict between the activity packet in the head direction cell ring and that in the visual ring. Previous work within continuous attractors has shown that, during conflict, the head direction cell activity packet will centre itself over the source of any visual input during simulation. However, the addition of plastic remapping of visual input allows for the head direction activity packet to reach a stable state that does not correspond to the centre of visual input.

Simulation results show that the introduction of plasticity allows for close replication of previous experimental work on cue conflict situations. Specifically, a visual dominance with an undershoot is seen for small conflicts, but beyond conflicts of 120° the influence of the network's memory state increases steadily up to a complete dominance at the maximum conflict size of 180°. This result is shown to be critically dependent on feed-forward plasticity between the visual and head direction cell

rings, and the subsequent remapping of visual information onto the head direction cell ring. When no learning takes place, model behaviour reverts to standard attractor dynamics, whereby the final packet location centres itself over input received from the visual layer without undershooting. This occurs for all conflict sizes except $180\,^{\circ}$, at which point the packet receives symmetrical input from the visual layer and therefore does not move.

Remapping is evident within the feed-forward weight structure, seen as an asymmetric distortion of the weights between the visual landmark cue location and each location in the head direction cell ring through which the packet has moved. With conflicts of an identical size in the opposite direction, this asymmetry is reversed. By examining the weight structure produced by introducing homosynaptic LTD, the link between changes in weight structure and resulting model behaviour is further strengthened. Introducing LTD has a net effect of weakening the connections from currently active visual cells to each intermediate head direction cell location through which the activity packet has moved. This creates an impetus for the packet to continue moving in the same direction through the ring which, when combined with the drawing influence of the visual landmark cue, results in an overshoot of the expected final location. For larger conflict sizes, the strength of the visual landmark is reduced and thus the combined influence of remapping and the cue no longer produces an overshoot. These results demonstrate very clearly that the final packet location is a result of plastic feed-forward remapping, as a change in the nature of this remapping changes the nature of the models performance.

Manipulations of learning rate reveals that although model behaviour is sensitive to learning rate changes, losing the ability to directly and specifically replicate experimental findings, it is still ca-

pable of functioning within the range of learning rate values tested. Lowering learning rate reduces the size of the undershoot seen with smaller conflicts and increases the range of conflict sizes over which visual landmark cues exert dominance, the latter resulting in a steeper gradient as final packet locations revert back to the original position. As learning rate increases, visual landmarks lose their dominance over successively smaller conflict sizes and final packet locations undershoot that of the visual landmark by successively larger amounts.

Further evidence that feed-forward weights, and accompanying plasticity, are central to model behaviour comes from manipulating their spread. Narrowing the width of the feed-forward weight profile initially reduces visual influence on the head direction cell packet, increasing the undershoot and shifting the breakpoint to smaller conflict sizes. However, once the weight profile is sufficiently narrow that the visual packet does not directly stimulate the head direction cell packet, the model switches to a new quality of behaviour. Provided visual input is strong enough, a second packet is driven up in the head direction cell layer at a location corresponding to the visual input. As it is driven up, this second packet inhibits the original head direction cell activity packet, which gives the appearance of an abrupt jump in the location signalled by head direction cell activity. This behaviour is independent of conflict size, and thus the final packet location is always directly over the location signalled by the visual landmark cue. Broadening the width of feed-forward weights initially has the effect of increasing visual influence, reducing the undershoot and pushing the breakpoint up to larger conflict sizes. However, this effect is limited, as broadening feed-forward weights also allows for remapping over larger distances. Beyond a certain upper level, broadening weights therefore begins to increase the undershoot and bring the breakpoint back to smaller conflict sizes.

In contrast, the effects of manipulating the spread of recurrent weights within the head direction cell layer are much less profound. It seems the width of the recurrent weight profile controls the readiness of the head direction cell activity packet to move. Thus broadening recurrent weights emphasises the effect of visual input on the activity packet, reducing the undershoot and causing the breakpoint to occur at larger conflict sizes. It is noted, however, that this effect is limited given that broadening recurrent weights too far necessarily flattens out the head direction cell activity packets to proportions that are not biologically realistic.

An exploration was also made of parameters controlling the width of packets initially generated in the head direction and visual rings. It was found that varying the width of the visual packet had a similar effect to varying the spread of feed-forward weights. However, the effect of varying the former was much weaker than seen when varying the latter. Varying the width of the head direction cell packet had no effect on final packet locations. Together, these results highlight yet again the central role that feed-forward and, to a lesser extent, recurrent weights play in this models operation.

Taken as a whole, these results clearly show that a continuous attractor-based network architecture is capable of closely replicating experimental results on cue conflicts, and that this is crucially dependent on plastic remapping of visual input. As such, this model provides a biologically plausible working theory of how cue conflicts may be resolved in vivo. In doing so, it supports the notion that cue conflicts are resolved *in vivo* in a distance-dependent manner. It also goes on to present the testable hypothesis that this is achieved through a plastic remapping of visual information. A related

prediction of this model is of state transition dynamics; namely that activity in the head direction cell system must move smoothly between locations, rather than jumping abruptly between states. This is demonstrated by results in Figures 3.14 and 3.15, where the feed-forward weight profile is sufficiently narrow to no longer provide direct input to the HD memory packet. In this situation, the conflicting visual stimulus instead stimulates a new HD packet at a directly corresponding location. Due to inhibition with the HD layer, this new packet eliminates the memory packet, such that activity in the HD layer jumps abruptly between locations. As the HD packet does not move smoothly between locations, no remapping can take place, and thus the model is no longer able to replicate experimental results. Although the question of state transition dynamics is still debated (Taube, 1998), Zugaro et al. (2003) suggest that smooth shifts occur for small conflict sizes, whilst large conflict sizes encourage abrupt jumps. However, this is based primarily on the speed of transition within the HD system, which is observed on the timescale of hundreds of milliseconds. This model contradicts this hypothesis, as smooth transitions must occur regardless of conflict size. However, short transition latency does not necessarily imply abrupt state transitions: the model as it stands can transition states smoothly by approximately 1000-1500ms, and this time could be reduced significantly by the usage of a spiking model, in which state transitions can occur over a timecourse of 100ms (Brunel and Wang, 2001).

A secondary aspect of these simulation results is the effect of cue salience. It appears that more salient visual cues exhibit stronger cue control, with cue conflict therefore being resolved more in favour of visual landmarks and head direction cells subsequently shifting their preferred directions further in favour of the visual cue. This is demonstrated in data from Knight et al (2014), who demonstrate that visual cues which have a consistent relationship with internal sources of information have a dominant

influence on cell preferred directions. The idea that only visual cues with a stable relationship to ideothetic information can form landmarks is evidenced in the literature to date. Some researchers have demonstrated that visual cues become landmarks only after having been experienced by an animal over a period of exploration in a given environment, where visual cues retain a stable relationship with ideothetic sources of information (McNaughton et al., 1995; Goodridge et al., 1998). For example, Knierim and colleagues (Knierim et al., 1995) demonstrated that periods of disorientation stopped the establishment of stimulus control by visual landmarks. Interestingly equivalent findings are available for place cells (Jeffery, 1997). This is entirely consistent with the idea of salience presented in this paper: less salient visual cues are weakened as inputs to the model, resulting in weaker cue control.

Interestingly, this appears to be one mechanism by which the head direction cell system is able to differentiate between distal and proximal visual landmark cues (Cressant et al., 1997; Zugaro et al., 2001; Yoganarasimha et al., 2006). Importantly, it seems that disruption of visual motion information affects the distinction between proximal and distal cues, as distal-proximal distinctions in stroboscopic lighting, which eliminates such information, are no longer present (Zugaro et al., 2004). This distal-proximal distinction has been successfully modelled, and demonstrated to depend on the fact that distal cues maintain a stable relationship with ideothetic information as the agent explores and environment, but proximal cues do not (Walters and Stringer, in preparation). These findings could be explained in our model in terms of a longer-term plasticity down-weighting visual information, which corresponds to a reduction in salience. However, such changes in salience are not limited to landmark stability, and indeed may be modified on behavioural or motivational terms e.g., a visual landmark which has previously been associated with reward may demonstrate cue control more readily than a

landmark of neutral valence. This is an interesting potential future avenue of research, which to the author's current knowledge has not been explored.

In order to achieve a level of explanatory clarity and present a clear hypothesis, this model is a simplification in that it provides a bare-minimum approach to solving the problem at hand. In this manner, part of this chapters aim is to provide a firm and well-understood theoretical foundation upon which further work can be based. With this in mind it is prudent to point out what this model does not address, on the basis that this is not what the model is oriented towards explaining. Although these simulation results pertain to data obtained from specific brain regions, the solution presented is a generally one. Therefore, this model does not propose to model any specific brain region, particularly with respect to the 'visual cells' employed during simulation, which remain a somewhat abstract entity. The model does not self-organise, as each head direction cell has a preferred direction and both weights structures are initialised to a specific configuration. It is of interest to note that the remapping seen in this model is not universal, in that only the weights from the visual ring onto the regions through which the head direction cell activity packet shifts are altered. Although this allows for an undershoot, it also shifts the firing properties of only a subset of head direction cells. This would be in contradiction to the established finding that head direction cells retain the angular difference between their preferred firing directions, a property known as isomapping (Yoganarasimha et al., 2006; Goodridge and Taube, 1995; Taube et al., 1996). Isomapping represents a further challenge to the model presented in this chapter. Further work on isomapping presenting a method, by which all head direction cells within a layer can inherit the shift in preferred firing direction of an initial subset, is detailed in the following chapter.

This model is limited in that it is a pre-wired network, which does not self-organise. However, previous work has shown that continuous attractors of this sort can self-organise (Stringer et al., 2002; Hahnloser, 2003; Stringer et al., 2004; Stringer and Rolls, 2006; Walters and Stringer, 2010) their recurrent connectivity. However, these papers have been criticised as not truly self-organising, due to their inclusion of head direction cells which already have a clear directional preference. That is, specific connections exist between head direction cells and areas of a theorised retina. Thus a chicken and egg situation arises, whereby head direction cells require directional specificity to correctly set up their recurrent connectivity, which they in turn need to correctly establish directional specificity. Stratton and colleagues (Stratton et al., 2010, 2011) propose a solution, whereby an initial approximately correct rough weight matrix, possibly set up during development via chemical gradients, is calibrated through input from symmetric angular head velocity cells. These symmetric angular head velocity signals correctly reflect speed of head rotation, and thus act to stabilise an activity packet to move correctly within the head direction cell ring. Their results show that such a network can self-organise to demonstrate continuous attractor behaviour.

Finally, although the model tolerates changes to a wide range of several parameters, showing nuanced behaviour, some parameters affect the size the head direction cell packet. However, this would imply that head direction cell tuning curves are variable, when the reality is that they are roughly constant under a variety of experimental conditions (Taube et al., 1990b). It is possible that explicit modelling of inhibition may allow for shaping of the activity packet, giving a consistent size (Jeffery, private communication).

To summarise, this chapter presents a model capable of replicating experimental results on distance-dependent cue integration in typical head direction cell system. Supported by a full exploration of parameters, it is hypothesised that the basis of this cue integration is a plastic remapping of visual input onto head direction cell systems. This provides a novel addition to experimental data, with the prediction that plastic remapping of visual information on the head direction cell system is responsible for complex cue conflict resolution behaviour seen in vivo, and a more tentative prediction that state transition dynamics within the head direction cell system are continuous.

Chapter 4

Isomapping of Head Direction Cell preferred firing directions following environmental manipulations

In Chapter 3, it was demonstrated how, following cue conflict a small portion of head direction cells may shift their preferred firing directions in response to a conflict between visual landmarks and ideothetic information. Although the model presented generated simulation results that had an excellent concordance with experimental data, it was noted that not all head direction cells within the layer shifted their preferred directions. Due to the nature of Hebbian learning, only those head direction cells which have been active during simulation will have a different profile of presynaptic weights after simulation. Therefore, only those head direction cells through which the activity packet has moved during simulation will have a different weight profile, and thus have changed their preferred firing directions. This is pictured in Figure 4.1, which shows simulation protocol used in Chapter 3 with a caricatured network below with feed-forward weight changes detailed for cells inside and outside the region of the head direction cell layer active at any point during simulation.

This chapter presents further simulation work, using an identical model to that of Chapter 3 in an extended simulation paradigm. This allows for the use of path integration mechanisms to allow all head direction cells within a layer to inherit the shift in preferred directions of the initial subset that were remapped during initial cue conflict. Thus, a mechanism is presented which allows for isomapping of head direction cell preferred directions following cue conflicts. This represents a novel theoretical perspective on isomapping, which to date has been little explored. A full exploration is made of various model parameters, including agent rotation speed, accuracy of path integration, saliency of visual information and learning rate, in order to fully explain the operational principles of the model. A discussion is made of possible future directions based on these initial results.

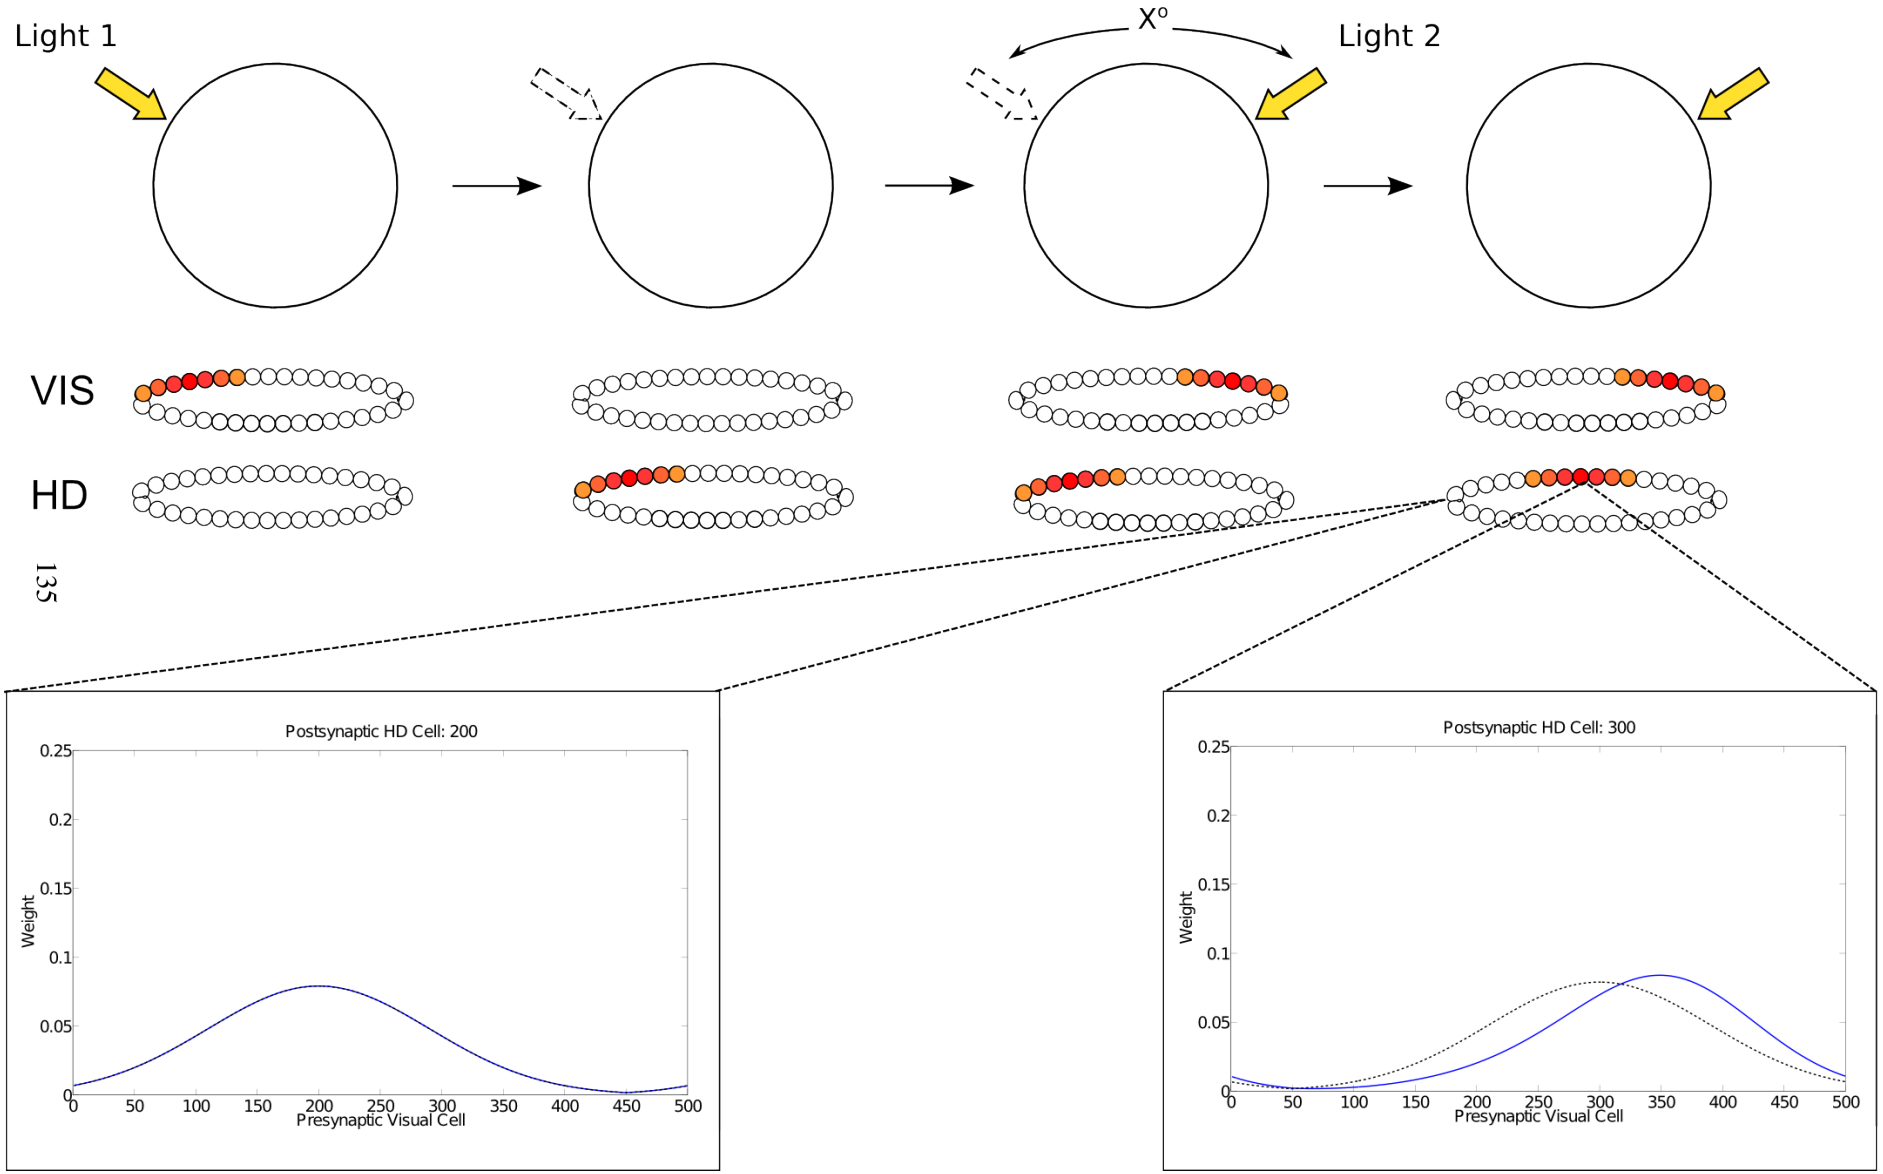


Figure 4.1: Demonstration of partial HD layer remapping. Simulation protocol from Chapter 3 is presented (top), with a caricatured network diagram showing relative positions of VIS and HD activity packets during simulation. Feed-forward weights before (black dotted) and after (solid blue) simulation are given for an HD cell outside the region of the HD layer active during cue conflict (cell 200, bottom left), and for an HD cell that has been active during cue conflict (cell 300, bottom right). Only cell 300 shows a remapped profile of receiving weights from the VIS layer, showing that only cells active during cue conflict remap.

4.1 Model description

The model used in this chapter is identical to that described in Chapter 3, Sections 3.1 and 3.2 and thus equations are not replicated again here.

4.1.1 Experimental protocol

The first stage of these simulations employed the experimental protocol for cue conflict detailed in Chapter 3, Section 3.2.3. After this cue conflict period there followed a period of rotation and path integration, during which the simulated agent rotates on the spot clockwise at angular velocity *V*. The duration of this rotation period was determined by the feed-forward weights: once the w^2 weight profile had reached a stable asymptotic state, the simulation was ended. Activity is moved through the visual ring at angular velocity *V*, with the location of the activity packet within the visual ring being updated each timestep according to *V*.

During this path integration session, the term e_i is updated to represent vestibular input to the head direction cell layer. This is represented in Equation (4.1) as $e_i(t)$, which for each head direction cell i is calculated according to a Gaussian profile

$$e_i(t) = \lambda e^{-(s_i^{HD}(t))^2/2(\sigma^{HD})^2} \tag{4.1}$$

where λ is a scaling factor determining the strength of this input to head direction cell i and σ^{HD} is the standard deviation of the Gaussian profile. s_i^{HD} is the difference between the estimate of current head

direction $x(t)$ computed by path integration mechanisms and the preferred direction x_i of postsynaptic head direction cell i. It is given by

$$s_i^{HD}(t) = MIN(|x_i - x(t)|, 360 - |x_i - x(t)|) \tag{4.2}$$

which creates a wrap-around effect, with response profiles of head direction cells as a population remaining continuous across the $360/0^\circ$ divide.

Activations and firing rates of each head direction cell are updated as according to Equations (3.1) and (3.2) in Chapter 3, and feed-forward weights are updated as according to Equation (3.5) in Chapter 3.

After this period of rotation, plasticity in the feed-forward layer is stopped, and there is no further updating of synaptic weights. There follows a period of 1 second of simulated time, during which all external inputs to the head direction cell layer are removed, and activity is sustained through recurrent collateral connections alone. Finally, there is a two second period during which activity is systematically moved through the visual ring, and activity of the head direction cell layer is observed. This provides a test as to whether the head direction cell layer has retained a stable relationship with the visual ring, which will be the case if each head direction cell has remapped by a similar amount. For clarity, the entire simulation protocol is represented graphically in Figure 4.2.

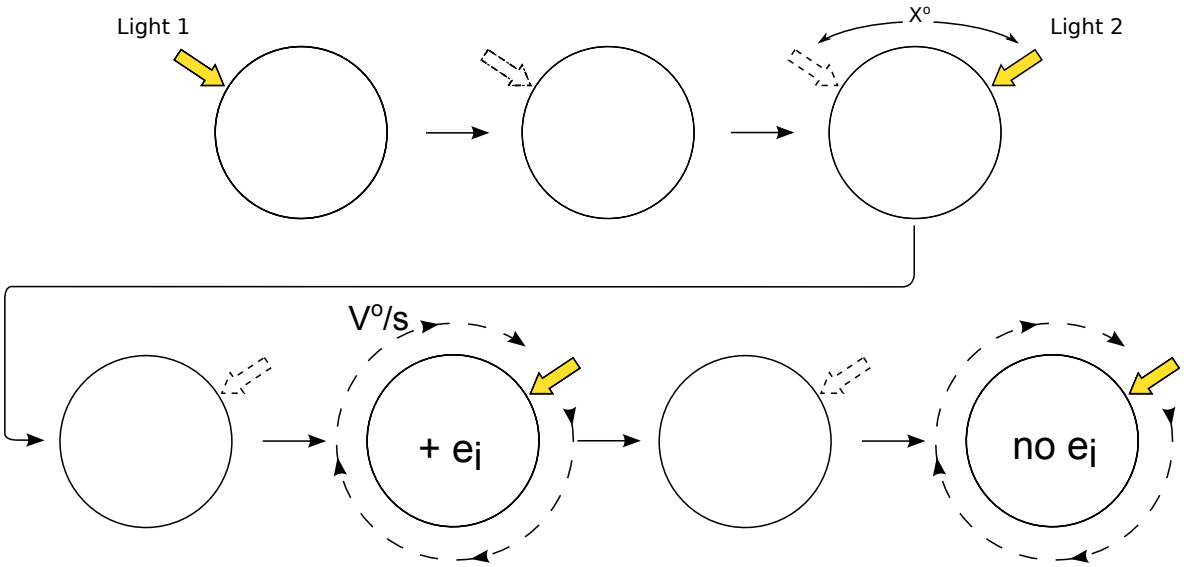


Figure 4.2: Graphical representation of simulation protocol used. Top row shows protocol used in Chapter 3. Following this cue conflict, there is a brief pause (bottom left), followed by a re-presentation of the light stimulus moving through the visual layer at a constant velocity of $V\,^{\circ}/s$, with path integration mechanisms, e_i, also active (bottom left middle). There follows another brief pause (bottom right middle), before a final stage in which an activity packet is systematically shifted through the visual layer, without any path integration (bottom right). This final stage acts to test the coherence of remapping from the VIS layer onto the HD layer.

4.2 Isomapping: A hypothesis

After the initial cue conflict, simulated in Chapter 3, weights have been updated from cells representing the location of the visual landmark to cells in the subregion of the head direction layer through which activity has propagated. This results in these head direction cells responding maximally to different egocentric bearings to the visual cue, and resulting in a shift in preferred direction. However, these changes have only occurred for a subset of head direction cells within the ring, and most cells retain their original configuration with visual space. The problem, then, is how other head direction cells also shift their preferred directions to reflect this change.

Consider just two head direction cells, $H1$ and $H2$. $H1$ is the center of the final head direction packet location after cue conflict, whilst $H2$ is an arbitrary cell that has not been active at any point

during cue conflict. Before simulation, $H1$ and $H2$ receive maximally from visual locations $Vis1$ and $Vis2$ respectively (Figure 4.3A). After cue conflict, feed-forward weights have been updated such that $H1$ now points to location $Vis1 + \gamma$, where γ reflects the disparity between the head direction packet location and the visual packet location (i.e., the amount of undershoot). Thus, $H1$ has shifted preferred direction by γ. Conversely, $H2$ has never been active and thus no update has been made to the weights projecting to this cell (Figure 4.3B).

If the agent then rotates on the spot, the activity will continue to propagate through the visual ring. Assuming then, that path integration system input to the head direction cells is capable of driving activity through the head direction cell layer at the same speed as the agent's rotation, the head direction activity packet will move at the same speed as the visual packet (Figure 4.3C). This will preserve the disparity of γ between head direction and visual locations. When cell $H2$ becomes active, it will be coactive with visual location $Vis2+\gamma$, and thus Hebbian learning will occur to strengthen projections from $Vis2 + \gamma$ to $H2$. This means that $H2$ will have shifted preferred direction by γ, identically to $H1$. As the agent rotates on the spot, continued Hebbian weight updates performed on w^2 connections will result in head direction cells across the entire layer learning to receive maximal input from a visual location displaced by γ from their original preferred direction, exactly as for the initial subset of head direction cells whose preferred directions shifted after initial cue conflict. This results in all head direction cells inheriting a change in preferred directions, and thus leads to isomapping (Figure 4.3D).

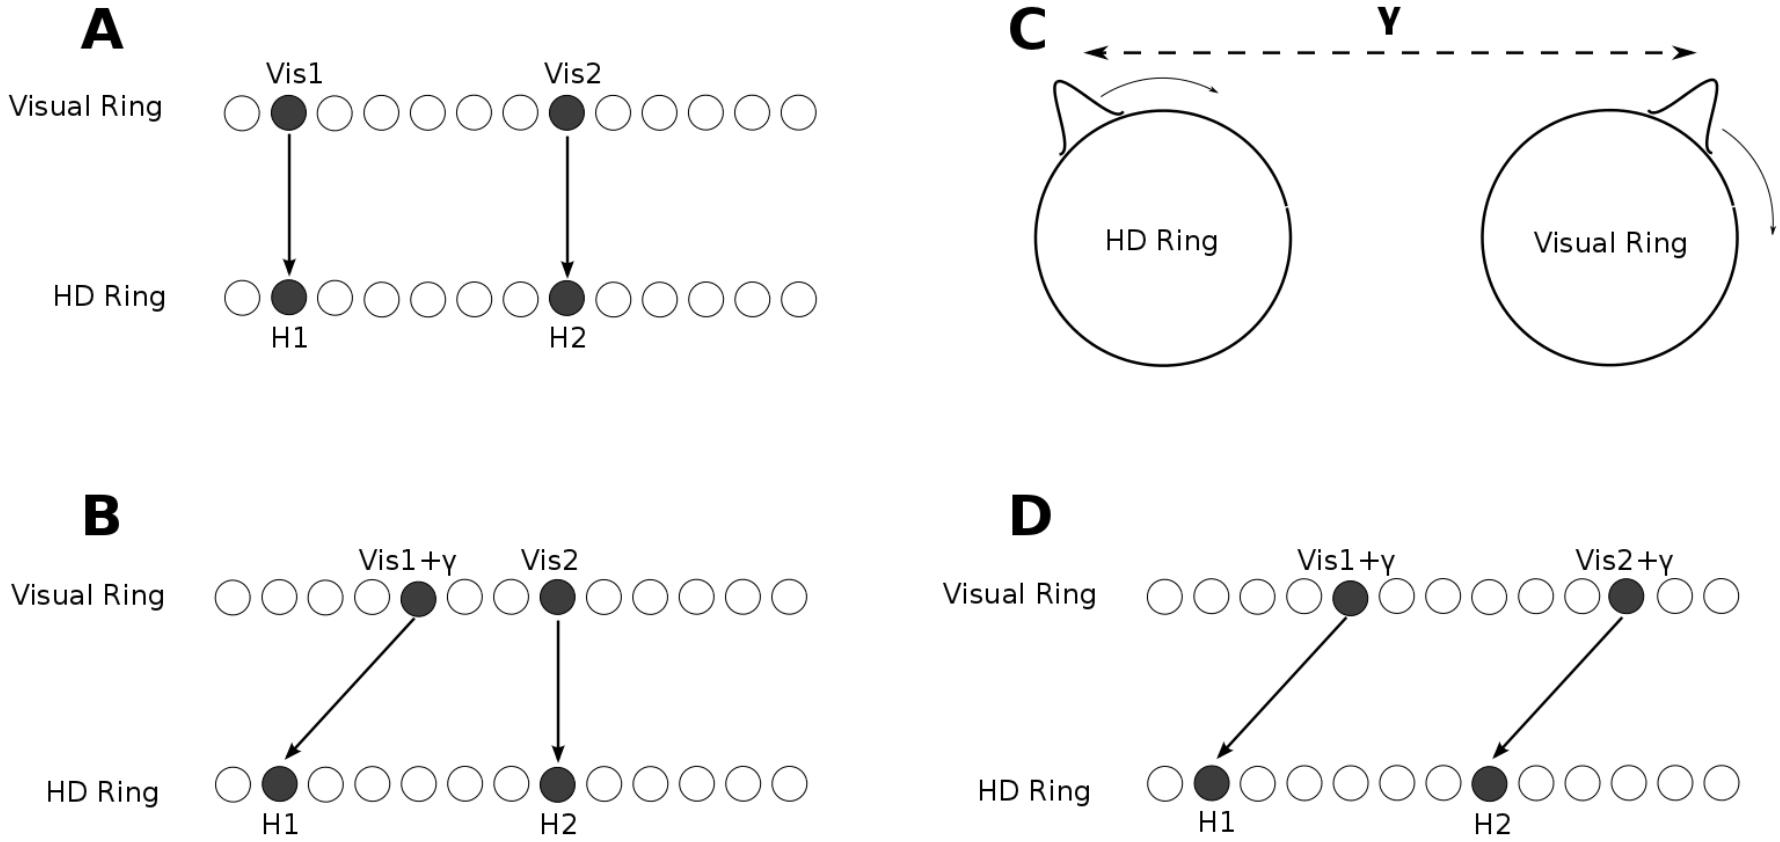


Figure 4.3: Illustration of isomapping theory. A: Both head direction cells received from corresponding locations in the visual ring. B: After initial cue conflict, cells within the cue conflict region (H1) have changed preferred location in the visual ring, but other cells (H2) have not. C: During subsequent rotation of the agent and update of HD cell firing by path integration inputs, the new relationship between visual and head direction packets is learned around the whole HD ring as activity packets in the visual and HD rings shift in lockstep. D: After the agent has rotated several times, all cells across the HD layer have shifted their preferred directions by similar amounts.

4.2.1 The relationship between the change in feed-forward weight structure and head direction cell preferred firing direction

The results presented in this chapter require a brief theoretical explanation in terms of their applicability to experimental data. Changes in feed-forward weight structure do not directly correspond to a change in the preferred compass direction for a particular HD cell. Rather, changes in preferred compass direction of a particular HD cell are linked to changes in feed-forward weight structure in the following manner

$$\text{preferred compass direction change} = \text{cue conflict size} - \text{mean weight vector shift} \tag{4.3}$$

where the change in preferred compass direction is also equal to the amount by which the head direction cell activity packet has shifted following initial cue conflict, and mean weight vector shift is, by elimination, equal to the amount by which the final packet location has undershot visual landmark cue location. Thus, there is a fixed systematic relationship between the shift in mean feed-forward weight and the change in preferred compass direction of a particular HD cell.

In the model detailed in this chapter, visual cells are encoded in a different frame of reference, in that they are rat-centered rather than room-centered (as visual landmarks were conceived in the Knight et al (2014) paper). Thus, the model's visual cells do not represent absolute location of visual landmarks within the environment. Rather, they represent the egocentric bearing from the agent to a particular visual landmark cue.

Consider the case in which the visual activity packet is switched on in a new location that is 80 degrees clockwise. Regardless of the rats current allocentric representation of head direction, which is correct from a room-centered point of view, visual information now tells the rat that it has rotated 80 degrees further counterclockwise than it actually has. In this example, each head direction cell represents 1 degree of allocentric space, with 360 cells in the system and visual information totally dominates head direction cell responses. Say that head direction cell 1 is active in the rat, which is correct from a room-centered frame. Visual information is now telling the rat that the head direction cell representation is incorrect by 80 degrees clockwise. Thus, the head direction cell system updates based on visual information to have head direction cell 81 become active. Whilst absolute, room-centered, head direction has not altered; the location of peak head direction cell activity has changed by 80 degrees. Thus, the preferred firing direction of the head direction cell system has changed by 80 degrees. In this manner, a change in preferred firing direction could be thought of, in this model, as analogous to the initial movement of an activity packet during cue conflict simulation. From this logic, we can extract the systematic relationship expressed in Equation (4.3).

4.3 Simulation results

This section gives simulation results using the model and simulation protocol described above. Similarly to Chapter 3, this chapter makes use of a population vector statistic, which is described in more detail in Chapter 3, Section 3.4. For the purposes of this analysis, an analogous statistic, the weight vector, is computed. The weight vector tells us the average direction of receiving synaptic weights for a given postsynaptic cell and, for each postynaptic head direction cell, i, is calculated as

$$\theta_{wi}(t) = \arctan\left(\frac{\sum_j w_{ij}\sin(\theta_j)}{\sum_j w_{ij}\cos(\theta_j)}\right) \tag{4.4}$$

As with population vector calculations, the following corrected formula (Batschelet, 1981; Fisher, 1993) was used to account for the wrap-around

$$\theta_{wi}(t) = \begin{cases} \arctan\left(\frac{\sum_j w_{ij}\sin(\theta_j)}{\sum_j w_{ij}\cos(\theta_j)}\right) & \text{if } \sum_j w_{ij}\sin(\theta_j) > 0, \sum_j w_{ij}\cos(\theta_j) > 0; \\ \arctan\left(\frac{\sum_j w_{ij}\sin(\theta_j)}{\sum_j w_{ij}\cos(\theta_j)}\right) + 180 & \text{if } \sum_j w_{ij}\cos(\theta_j) < 0; \\ \arctan\left(\frac{\sum_j w_{ij}\sin(\theta_j)}{\sum_j w_{ij}\cos(\theta_j)}\right) + 360 & \text{if } \sum_j w_{ij}\sin(\theta_j) < 0, \sum_j w_{ij}\cos(\theta_j) > 0; \end{cases} \tag{4.5}$$

where w_{ij} is the strength of the synapse from presynaptic visual cell j onto postynaptic head direction cell j and each visual cell has the preferred firing direction θ_j.

A weight shift was calculated for each cell within a layer, by comparing the weight vector before and after simulation. To calculate a mean weight shift, averaged across all cells within a layer, weight shifts for each cell were converted to unit vectors in a two-dimensional plane

$$m_i = \begin{pmatrix} cos(\theta_{wi}) \\ \\ sin(\theta_{wi}) \end{pmatrix} \tag{4.6}$$

These vectors m_i are then averaged

$$\bar{m} = \frac{1}{N} \sum_i m_i \tag{4.7}$$

where N is the total number of postsynaptic head direction cells i. The mean weight shift μ is then given by transforming m_i using the four quadrant inverse arctangent function

$$\mu = \arctan 2\, \bar{m} \tag{4.8}$$

In addition to this mean weight shift, several other statistics were taken to measure the spread of this circular data. The resultant vector length M is measure of circular spread, with values of M closer to 1 reflecting data which is concentrated around the mean direction. It is calculated as

$$M = \|\bar{m_i}\| \tag{4.9}$$

In addition to resultant vector length, variance is calculated as

$$\text{Var} = 1 - M \tag{4.10}$$

This variance is not analogous to variance on a linear scale, and is instead bound in the interval $[0, 1]$. Given its relationship with M, will tend towards zero when all data samples have similar directions. Standard deviation, however, is analogous to standard deviation on a linear scale, and is calculated as

$$SD = \sqrt{2(1 - M)} \tag{4.11}$$

which is bounded in the interval $[0, \sqrt{2}]$. These circular statistics are based on the operation of the MATLAB CircStat tool box (Berens, 2009; Berens and Velasco, 2009). Further details on the relationship between these statistics and similar measures of spread can be found in Mahan (1991); Zar (1999).

All simulations are, unless explicitly stated, run using the parameters given in Table 4.1.

4.3.1 Main result

Figure 4.4 is a summary of results obtained using the model, simulation protocol and parameters given above and varying the size of conflict between initial head direction activity packet and visual packet locations. It graphs, in blue, head direction cell packet undershoot relative to the location of the visual cue after the initial cue conflict. This is overlaid in dashed red with remap in feed-forward receiving weights (calculated as the difference between the initial and final receiving weight vectors

Table 4.1: Network parameter values for simulations reported in this chapter. All simulations, except where explicitly stated, are run using these parameters.

Network Parameters	
No. Head Direction Cells	500
No. Visual Cells	500
Runtime	25.0s
Cue Conflict Time	2.0s
Input Time	0.1s
Path Integration Time	20.0
$C^{HD \to HD}$	500
$C^{VIS \to HD}$	500
ϕ_1	200.0
ϕ_2	6.5
σ^{FF}	65.0
σ^{RC}	35.0
λ^{HD}	1.0
λ^{VIS}	2.0
σ^{HD}	20.0
σ^{VIS}	50.0
τ^{HD}	0.001
$\tilde{w}^{HD}$	0.4
α	0.0
β	0.3
λ	10.0
δt	0.0001s
k	0.05
V	180

for a given cell) averaged across all head direction cells in a layer. This demonstrates an excellent correspondence between undershoot following cue conflict and change in preferred direction of head direction cells. This demonstrates the fact that remapping in the head direction cell layer following path integration is, as hypothesised, dependent on the head direction cell activity packet shift after initial cue conflict. A very strong correlation $r = 0.999$ between initial undershoot and weight remap suggests that remapping following path integration is dependent on undershoot following cue conflict.

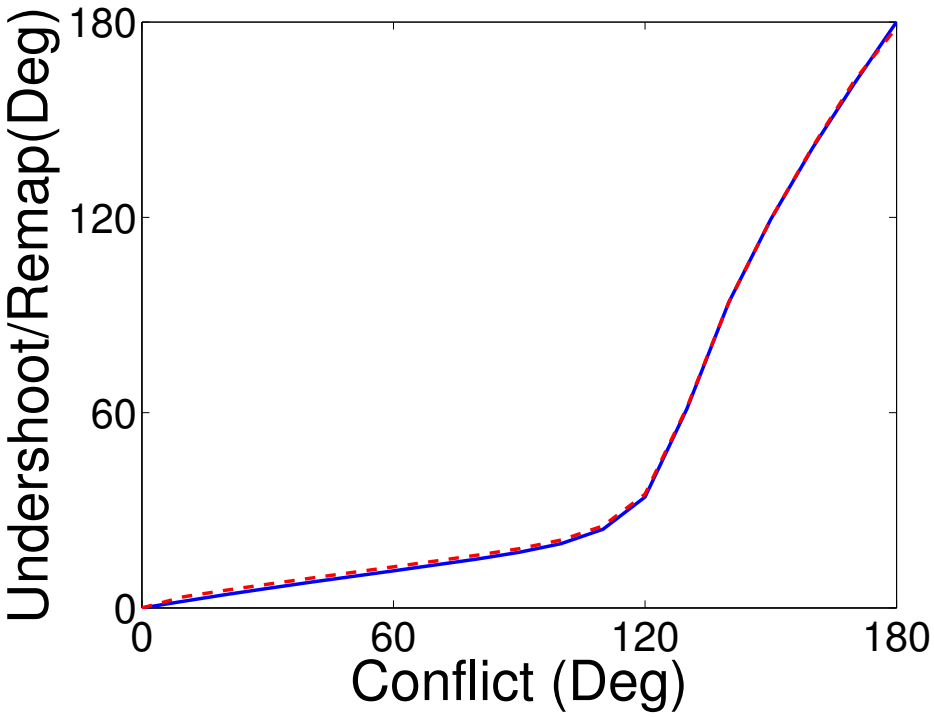


Figure 4.4: Head direction layer packet undershooting of true visual landmark location during initial cue conflict (solid blue), overlaid with mean remap in receiving feed-forward weights following path integration, averaged across all head direction cells within a layer (red dashed). Mean remap after path integration shows excellent concordance with initial packet undershoot during cue conflict, showing a change in head direction cell preferred directions based on cue conflict resolution.

These results can be unpacked further by examining a simulation of a single cue conflict size in detail. Figure 4.5 presents summary graphs of such a simulation, from the block of results summarised in Figure 4.4, run for a conflict size of $80\,^\circ$ between the visual cue and initial head direction packet location. In Figure 4.5A, we see the firing rates of the entire head direction cell population over time. Figure 4.5B shows the locations of the head direction (blue) and visual (red) packets over time. Together, these show that simulations run as hypothesised, with the average remap across all HD cells in close agreement with packet undershoot as seen in Figure 4.4.

During the first two seconds of cue conflict, we see cue conflict resolved in a similar manner to the previous chapter: a final head direction cell packet location which undershoots that of the visual packet by a similar amount ($15.02\,^\circ$) to the main result given in the Chapter 3 ($14.05\,^\circ$). This is reflected in Figure 4.5B, with the head direction packet location (blue) pulling towards, yet stopping short of, the visual packet location (red). From 2 seconds to 22 seconds, we see the packets move through the visual and head direction rings at angular velocity V. During this period, we see in Figure 4.5B the fact that both the visual and head direction packets appear to be moving, and that the disparity between them introduced by cue conflict is preserved through all locations. Following this rotation period, all inputs are removed and the head direction cell activity packet is maintained in position through recurrent collateral synapses. From seconds 23 to 25, during a period where no further learning takes place and head direction activity is influenced solely by visual input, we see that the head direction cell activity accurately tracks visual input, with a consistently-sized separation of head direction and visual packets in Figure 4.5B. This indicates that the remapping of feed-forward weights has been coherent.

Although Figure 4.4 shows that the feed-forward weight remapping after path integration matches the original packet undershoot after cue conflict when averaged across all head direction cells within a layer, it is possible that individual head direction cells show variable degrees of remap centered on the mean values observed. This would be in contradiction to the head direction cell population isomapping observed in the brain. However, in the model simulations reported here, the head direction cells remap by similar amounts. A coherent head direction tracking of a moving light, seen in seconds 23-25 of Figure 4.5, confirms that remapping in feed-forward weights has been coherent across the entire head direction cell layer.

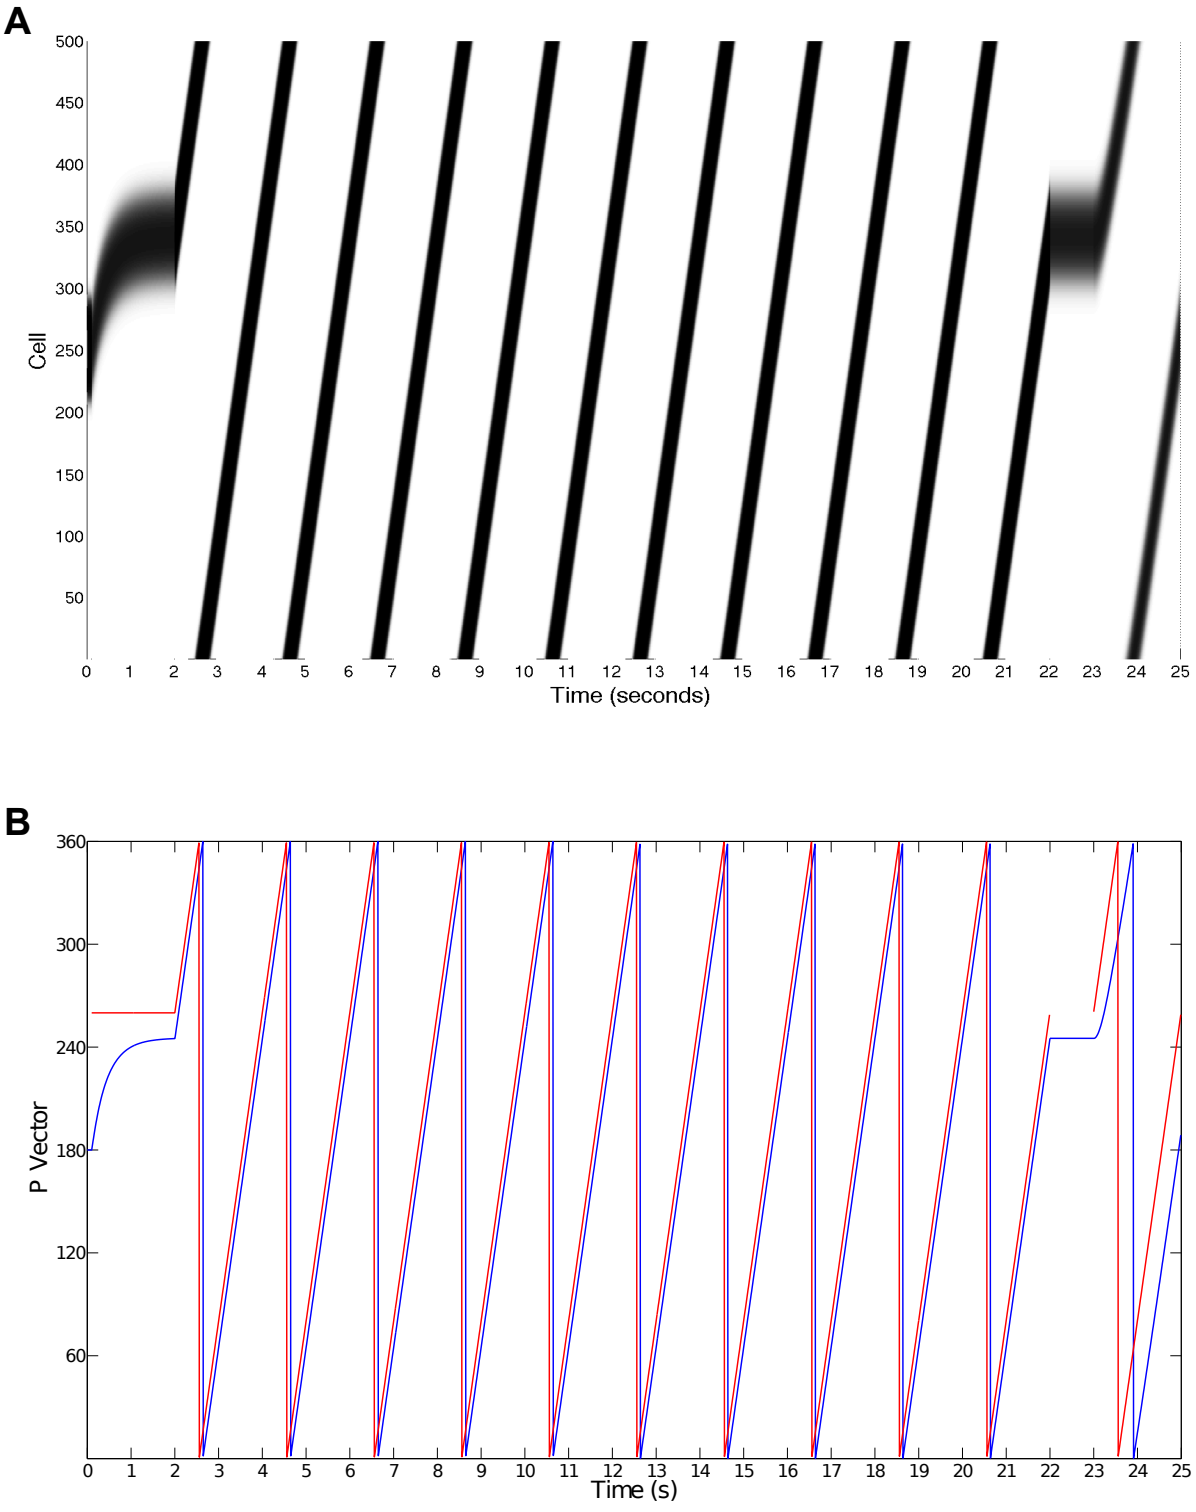


Figure 4.5: Simulation of 80° cue conflict followed by HD cell remapping during path integration. A: HD layer firing rates (black indicates high firing) B: respective HD (blue) and visual packet (red) locations through the course of simulation. Seconds 0-2 show cue conflict resolution, 2-22 path integration and 23-25 test phase.

Figure 4.6 confirms the coherency of remapping within the head direction cell layer. It shows a rose plot of change in weight vector for all head direction cells before and after the simulation depicted in Figure 4.5 for a single conflict size of $80\,^{\circ}$. The distribution, across all head direction cells, of changes in weight vector, reflecting remapping in receiving feed-forward weights, is tightly clustered around a mean value of $16.19\,^{\circ}$, with a resultant vector length of $M = 0.999$, yielding a variance of $\mathrm{Var} = 0.001$ and standard deviation of $\mathrm{SD} = 0.77\,^{\circ}$. As expected given the results in Figure 4.4, this mean value closely matches undershooting of final HD packet location during cue conflict, $15.02\,^{\circ}$, for this conflict size. A resultant vector length close to 1, and correspondingly small variance and standard deviation together show that this population is tightly clustered around this mean value, with head direction cells therefore remapping by very similar amounts as *in vivo*. Results for other conflict sizes (not shown) demonstrate similarly small S.Ds, with the pooled standard deviation across all conflict sizes in this block of simulations being $\mathrm{SD} = 0.734\,^{\circ}$. Thus, results show that path integration causes all cells in the head direction cell layer to remap by a similar amount in response to cue conflict.

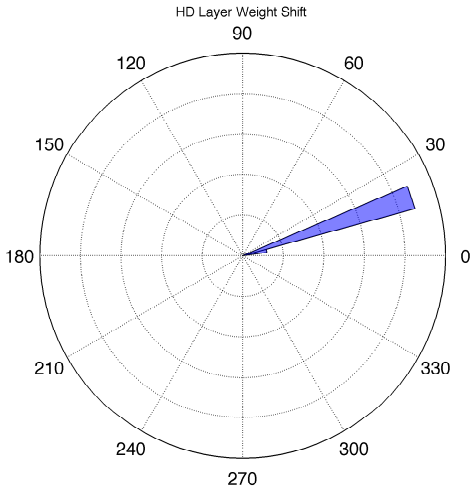


Figure 4.6: Rose Plot with bin size of $7.2\,^{\circ}$, showing change in average weight vector across the entire HD layer after simulation with a conflict size of $80\,^{\circ}$ followed by a driving of the HD packet through the layer by agent rotation and path integration. The distribution of changes in weight vector is tightly clustered around the mean value of $16.19\,^{\circ}$, demonstrating a coherent isomapping of HD cell receiving weight profiles.

Together, these results tell us that, following cue conflict, a shift in the preferred directions of a subset of head direction cells can be inherited by the entire head direction cell layer through path integration. This can be supported by examining the changes in feed-forward weight profiles for a head direction cell involved in the initial cue conflict, and a head direction cell not involved in this initial cue conflict. Figure 4.7 shows the receiving feed-forward weight profiles at an initial stage before simulation (black dotted), an intermediate stage after cue conflict (red solid) and a final stage after path integration (blue solid). These are shown for both a cell within the cue conflict region (top, cell 350) and outside the cue conflict region (bottom, cell 100). Whilst both cells show a remapping over the entire course of simulation, reflected as a difference between initial and final weights, only the cell within the cue conflict region shows an intermediate remapping after initial cue conflict. This is reflected by a clear difference between initial and intermediate weight profiles for cell 350, whilst cell 100 shows no change in weight profile over this time period.

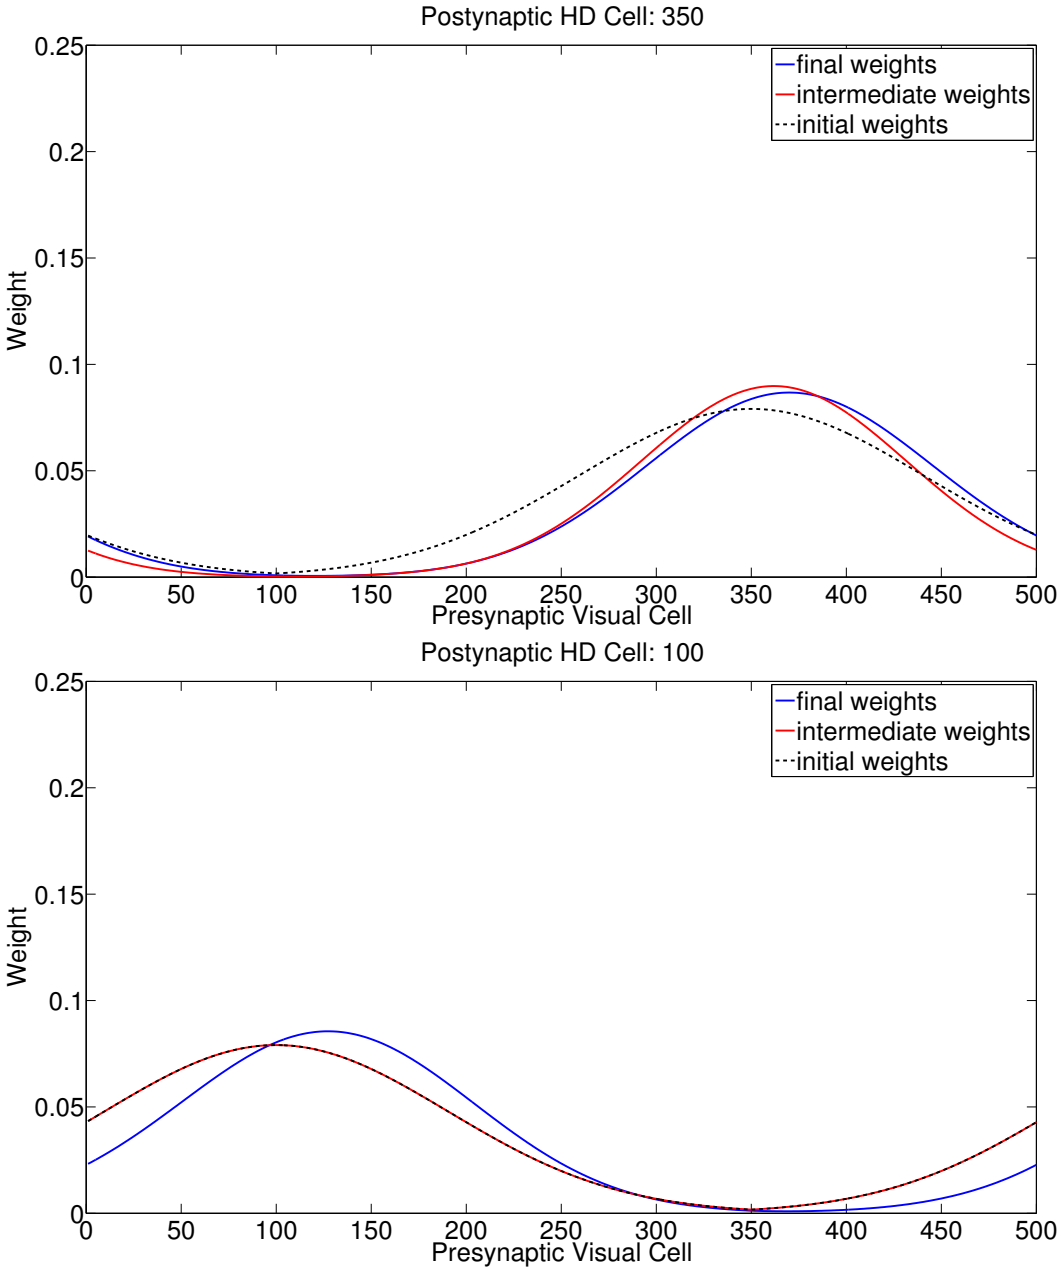


Figure 4.7: Effect of learning during isomapping on the w^2 synaptic weights onto an HD cell within the cue conflict region (top) and outside the conflict region (bottom). Weights are shown before learning (black dashed), after cue conflict (red) and after isomapping (blue). This demonstrates, for HD cells involved in cue conflict only, an intermediate remapping during cue conflict. Remapping is consolidated during path integration and inherited by all HD cells with the layer; including those not involved in the initial cue conflict.

4.3.2 Effect of cue conflict direction

Figure 4.8 shows results obtained for identical conflict sizes as above, but in the opposite direction. Besides this difference in conflict direction, all other simulation protocols and parameters remained the same. A very strong correlation of $r = 0.999$ shows that, as with previous results, weight remap following path integration is closely linked to packet shift following cue conflict. Weight remap appears to be coherent across all HD layer cells.

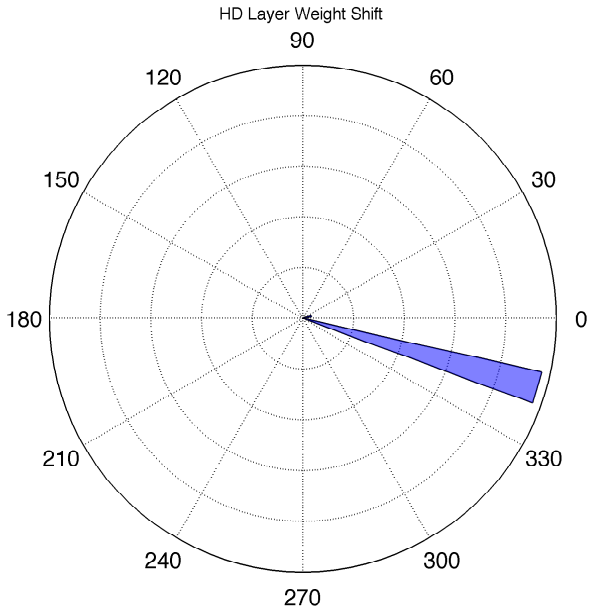


Figure 4.8: Rose Plot with bin size of $7.2\,^\circ$, showing change in average weight vector across the entire HD layer after simulation with a conflict size of $80\,^\circ$ in the opposite direction to Figure 4.6. As in Figure 4.6, isomapping has occured, reflected by a tight distribution around the mean value of $-13.36\,^\circ$, which is plotted on a circular scale as $346.64\,^\circ$.

Analogously to Figure 4.6, Figure 4.8 shows a rose plot of HD layer weight remaps for a conflict size of $80\,^\circ$. Again, the distribution of changes in weight vector is tightly clustered around a mean value of $-13.36\,^\circ$ (value in negative to denote a remap in the opposite direction), with mean resultant vector length of $M = 0.999$, yielding a variance of $\mathrm{Var} = 0.001$ and standard deviation of $\mathrm{SD} = 0.917\,^\circ$ (N.B. values of M and Var appear identical to the statistics accompanying Figure 4.6. However, their difference is beyond the level of 3 significant figures as presented here, which yields a different value for SD). This mean value for weight remap closely matches the HD activity packet undershoot of $-15.01\,^\circ$. Results for other conflict sizes in this direction (not shown) were similarly tightly distributed around their mean values, with a pooled standard deviation of $\mathrm{SD} = 1.002\,^\circ$.

Figure 4.9, similarly to Figure 4.7, shows the receiving weight profiles for a head direction cell inside and outside the cue conflict region before simulation, after cue conflict, and after path integration. Again, there is a qualitative difference in how cells remap: inside the conflict region, there is an intermediate remapping, whilst outside the conflict region there is not. Of equal importance is the observation that, relative to Figure 4.7, the remap seen in Figure 4.9 is in the opposite direction. This demonstrates a very clear effect of conflict direction, as well as magnitude, on subsequent changes in HD cell preferred firing directions.

Given these results, it is concluded that the link between weight remapping and activity packet undershooting during cue conflict is not affected by conflict direction, and that this remapping always reflects a coherent isomapping across all HD cells.

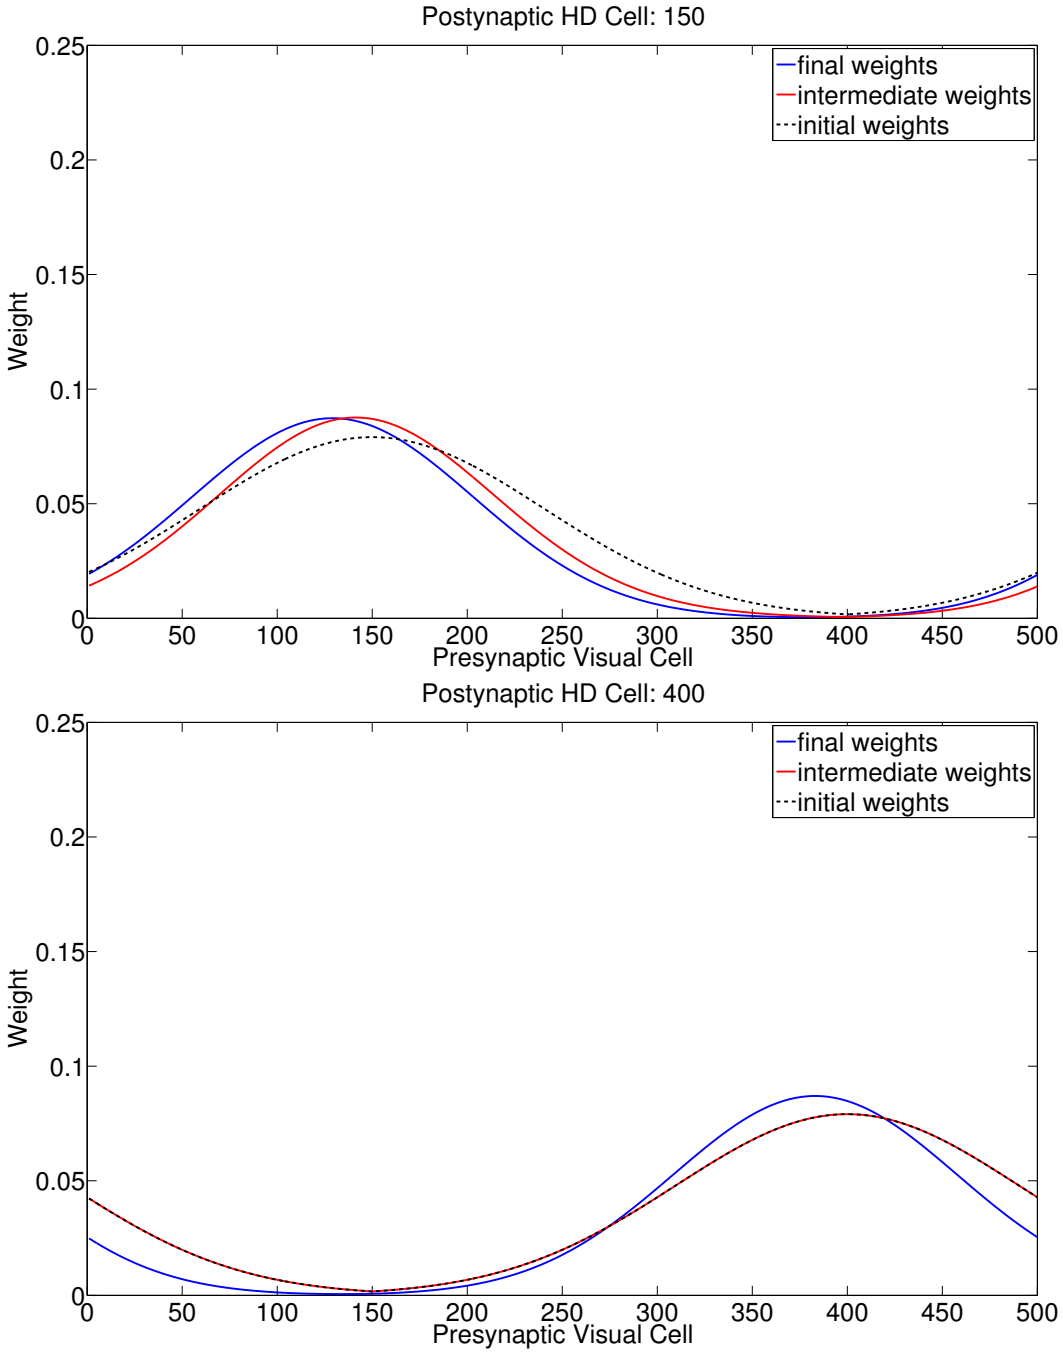


Figure 4.9: Effect of learning during isomapping on the w^2 synaptic weights onto an HD cell within the cue conflict region (top) and outside the conflict region (bottom). Conflict in opposite direction, conventions as in Figure 4.7. This demonstrates that isomapping occurs in the same two-stage process regardless of conflict direction. Of note is the fact that isomapping has occurred in the opposite direction to that seen in simulations yielding Figure 4.7, further demonstrating the link between packet undershoot during cue conflict and weight remap during agent rotation.

4.3.3 Accuracy of path integration

Section 4.2 of this chapter outlines the hypothesied isomapping mechanism, and makes clear that a critical aspect of this mechanism is the accuracy of path integration. Specifically, the hypothesis states that in order to preserve the visual-HD disparity generated by initial cue conflict across all cells, the two activity packets in the visual and head direction layers must move through their respective rings at the same speed during path integration. Without this preserved disparity, there is not a consistent relationship between activity in the visual and HD rings, and therefore stable isomapping cannot occur.

In order to investigate the effects of inaccurate path integration, the model was repeatedly simulated for a single conflict size of $80\,^{\circ}/s$ with the visual packet moving at a constant velocity of $180\,^{\circ}/s$ and the HD packet moving in the same direction at a range of different velocities in order to simulate inaccurate path integration.

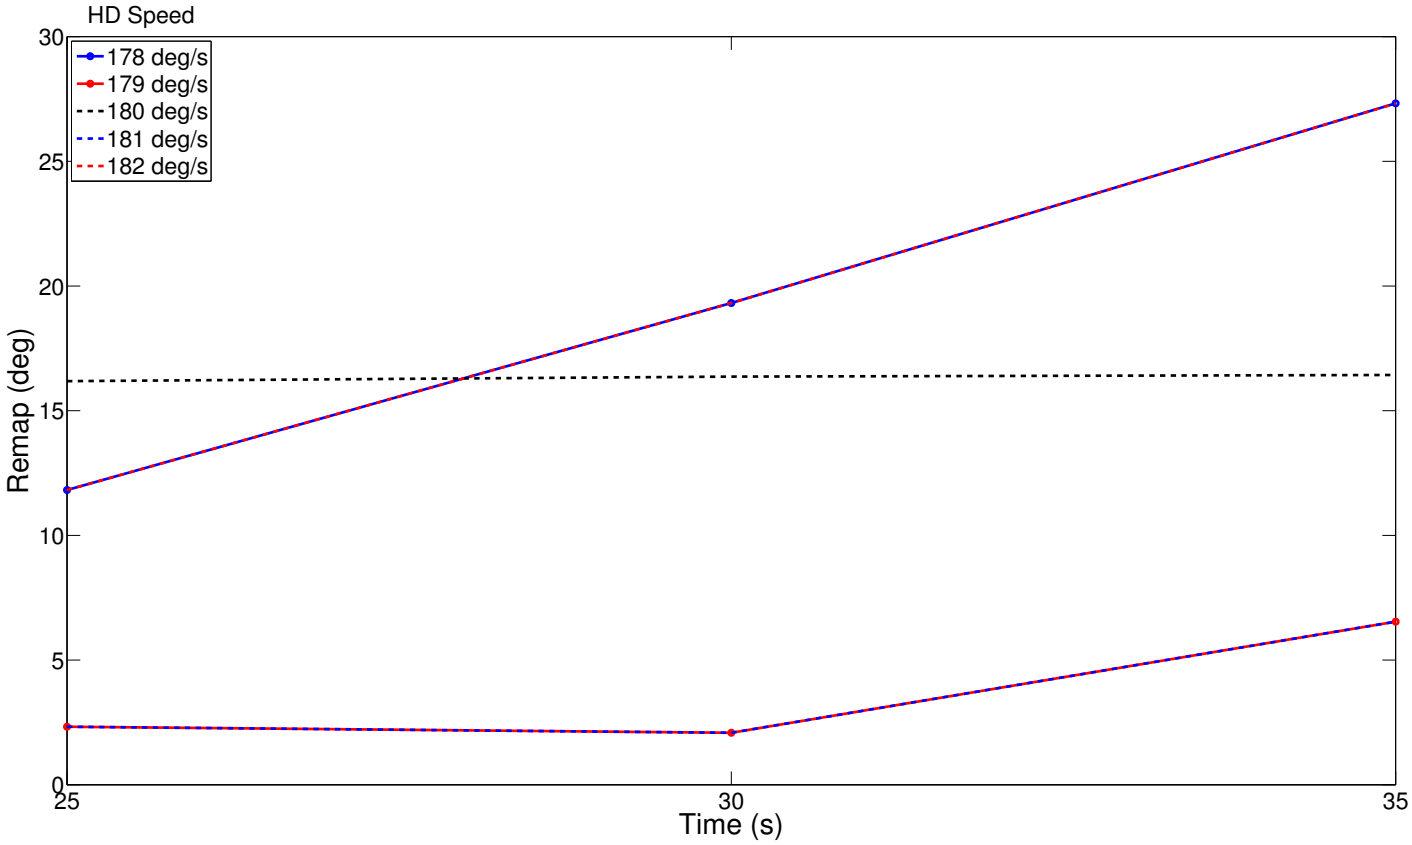


Figure 4.10: Summary graph showing mean remap in HD layer after path integration. Results are shown for a range of HD packet velocities in the interval [178,182] $^\circ/s$. Whilst a HD packet shifting at the same speed as the visual packet causes a remap which is stable over time ($180\,^\circ/s$, represented by the black dotted line), a mismatching of speeds causes unstable remapping (all other plotted lines).

Figure 4.10 presents a summary of these data, showing mean remap in the HD layer after path integration at a range of HD packet speeds. This reveals several interesting facts about the model's operation. Firstly, it tells us that when path integration is accurate at $180\,^\circ/s$, shown by the black dashed line, changes in feed-forward weight profile are stable over time. This is important, as it confirms that weights have reached a stable convergence. Secondly, observing the other plots tells us that when path integration is inaccurate, the weights do not reach a stable convergent point. This is because the relationship between the HD and visual activity packets is in a constantly-changing state, which the weights then continually remap to reflect. Thirdly, the simulations show that it is the absolute discrepancy between HD packet and visual packet speeds which is of consequence rather than whether the

HD packet is moving faster or slower than the visual packet. This is made clear by the overlap of lines for $178\,^{\circ}/s$ (solid blue with circular markers) and $182\,^{\circ}/s$ (dashed red), as well as for $179\,^{\circ}/s$ (solid red with circular markers) and $181\,^{\circ}/s$ (dashed blue).

As well as having an effect on mean HD layer cell shift, inaccurate path integration also seems to affect isomapping, reflected as an increase in the standard deviation of feed-forward weight shifts across the HD layer (Figure 4.11).

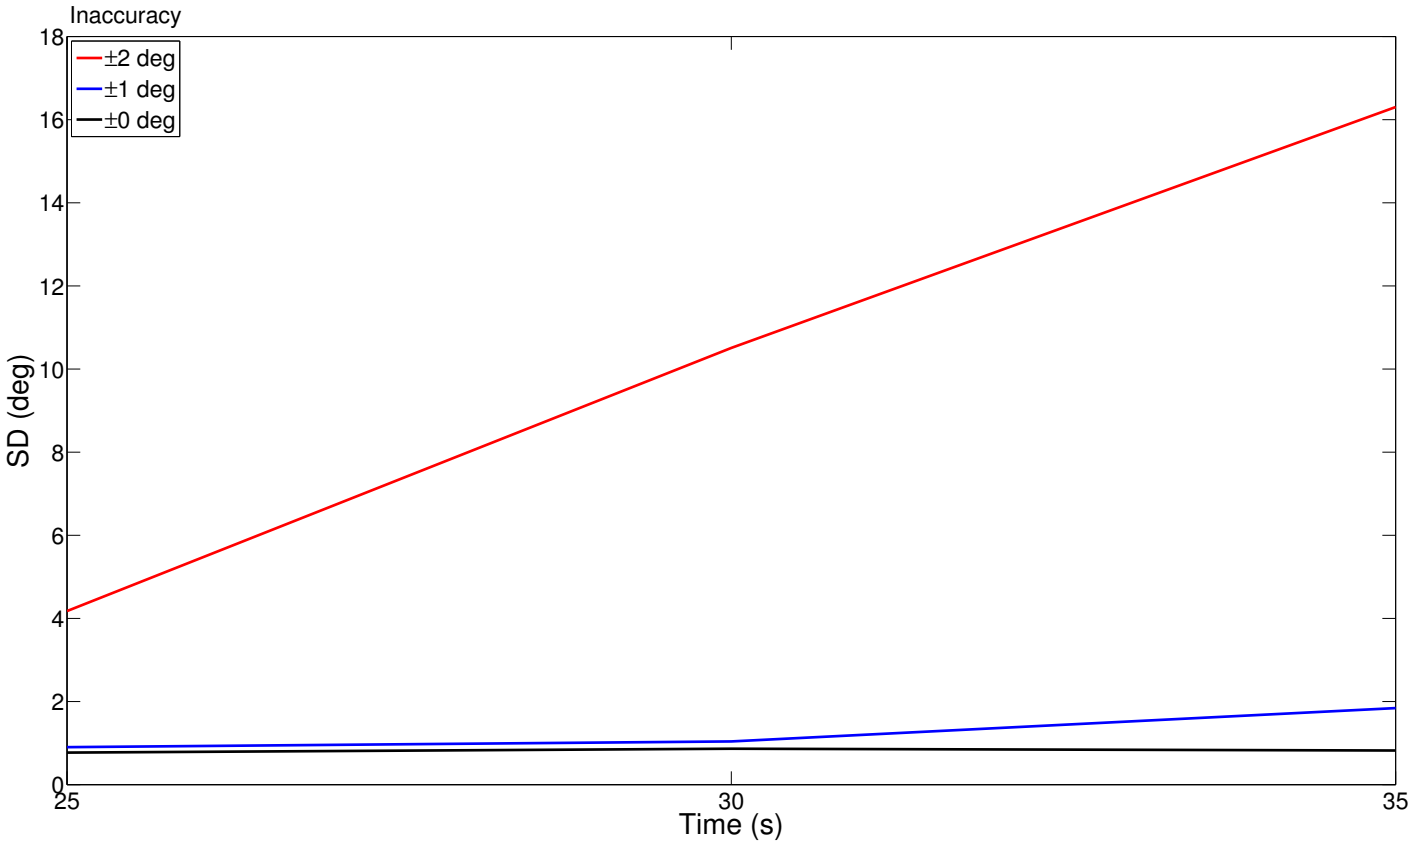


Figure 4.11: Summary graph showing standard deviation of shifts in feed-forward weights across the entire HD layer over time, for absolute path integration inaccuracies between HD and visual activity packet speeds of 2 degrees (red), 1 degree (blue) and 0 degrees (black). Standard deviation increases with path integration inaccuracy, and this effect is enhanced over time.

As the discrepancy increases between HD and visual activity packet speeds during path integration, the amount of remap in the w^2 feed-forward weights becomes less uniform across the HD layer (red vs blue vs black lines). This effect is exacerbated over time. Due to the inconsistent relationship between HD and visual activity packets, feed-forward weights are constantly in a state of flux, and not converging on a stable profile. This shifting weight profile means that the weight shifts cannot converge on a similar value, resulting in a greater spread of weight shifts for cells in the HD layer. Thus, inaccurate path integration serves not only to cause mean weight shift for the entire HD layer to be incorrect, but also causes weight shifts for individual HD cells to be more widely distributed about this mean. Overall, this yields a layer of HD cells which have changed their preferred direction by inaccurate and inconsistent amounts. By comparison, totally-accurate path integration with zero discrepancy between HD packet and visual packet speeds results in relatively uniform HD layer remapping that has shifted in accordance with initial packet undershooting.

4.3.4 Salience of visual information

Similarly to Chapter 3, an exploration was made of the ϕ_2 parameter, which controls the overall scaling of w^2 feed-forward weights and can thus be thought of as a change in the salience of visual information. Whilst a full reproduction of the original parameter exploration is not necessary, Figure 4.12 demonstrates full isomapping following path integration with a maximally-salient visual cue ($\phi_2 = 52.0$). Similarly to Figure 4.4 of this chapter, it shows both head direction activity packet undershooting of visual cue location following cue conflict (solid blue) and mean feed-forward weight remap following agent rotation and path integration (red dashed).

Even at a very large conflict size of $170\,^\circ$, visual information dominated cue conflict, with the head direction activity packet undershooting by as little as $2.33\,^\circ$. Subsequent agent rotation and path in-

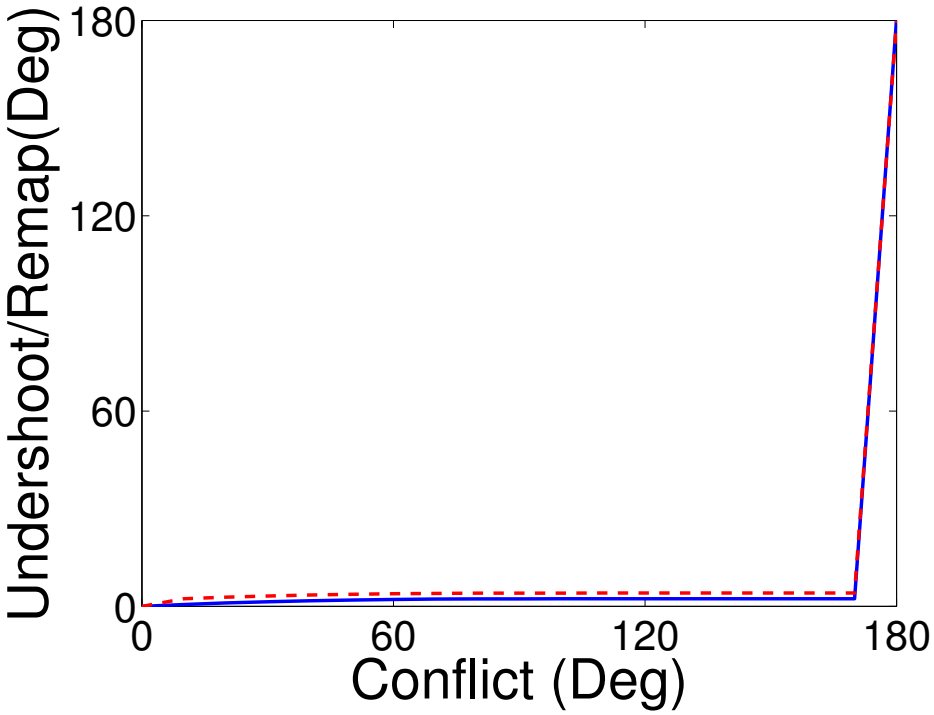


Figure 4.12: Head direction layer packet undershooting of visual landmark location during initial cue conflict(solid blue) overlaid with mean remap in feed-forward weights following path integration, averaged across all HD cells (red dashed). Similarly to Figure 4.4 of this chapter, there is an excellent concordance between HD packet undershoot during cue conflict and weight remap after path integration. With very salient visual information, as in this case, visual landmarks dominate cue conflict even for large conflicts, and this total capture is reflected accordingly by the feed-forward weight profile changes.

tegration caused weights to shift concordantly, with a mean weight shift value of $4.10\,^{\circ}$, giving a change in HD cell preferred firing directions of $165.9\,^{\circ}$. This shift in weights represents isomapping, and across all HD cells demonstrated a mean resultant vector length of $M = 0.999$, yielding a variance of $\mathrm{Var} = 0.001$ and standard deviation of $\mathrm{SD} = 0.969\,^{\circ}$.

An interesting result is generated at $180\,^{\circ}$ conflict. As the visual activity packet is equally distant from the HD activity packet in either direction, it provides equal stimulation to either side of the HD activity packet, which therefore does not move (for more information on this, see Chapter 3, Figure 3.6). This would result in a change in preferred direction of $0\,^{\circ}$, which does not demonstrate complete

cue capture at 180° conflict. However, theories of attractor dynamics show that when the HD layer is strongly stimulated outside the activity packet, a new packet can be driven up over the center of this stimulation and the old activity packet can be simultaneously extinguished. This model is no exception, and this effect can be seen in Chapter 3 Figure 3.16 of this thesis. In theory then if, at 180 degrees separation, the visual activity packet was very strong and did not provide any direct input to the HD activity packet, a new HD activity packet would simply be driven up in a location directly corresponding to that of the visual packet. In this case, the HD activity packet will have changed location by the full conflict size and thus cells in the HD layer will, after agent rotation and accompanying path integration, have changed their preferred firing directions by the same amount.

4.3.5 Effect of learning rate

Figure 4.13 shows summary results of simulations in which the learning rate parameter k was varied in the intveral [0.001, 0.01]. Figure 4.13A shows the correlation, across cue conflict sizes, between initial packet undershoot during cue conflict and weight remap after rotation and path integration. Figure 4.13B gives the standard deviation, pooled across conflict sizes, of the shift in weight vector of each cell in the HD cell layer. Figure 4.13C gives the mean absolute discrepancy between packet undershoot and mean weight shift, averaged across all conflict sizes. At low values of learning rate, we see a deterioration in all measures of remapping. With small values of k, weights have not trained until convergence and thus there is a large mean discrepancy between mean weight remap after path integration and initial packet undershoot (13C), with an accompanying weakening of the correlation between the initial packet undershoot and mean weight shift (13A). A low learning rate also means that isomapping breaks down, as the HD cells' weight profiles have not updated enough to be brought into line with initial changes following cue conflict, and this is reflected in larger pooled SD (13B).

Interestingly, this negative effect on pooled SD, mean deviation and correlation is not equally distributed across all conflict sizes. As would make logical sense, the average weight remap after path integration is in fact very close to undershoot size for the majority of conflict sizes (in the region of $1\,^\circ$ discrepancy) and has a tight distribution (SD in the region of $0.75\,^\circ$). At $180\,^\circ$ conflict, the average weight remap is much more discrepant, with a difference between the initial undershoot and mean weight remap of $156.69\,^\circ$, and weight remaps are more widely disributed with $SD = 48.62\,^\circ$. This is simply an effect of the learning rate being too low for the weight profile to have fully reached a stable remapping when this remapping is particularly distant from the original weight profile. As proof of this concept, the simulation for $180\,^\circ$ conflict was re-run over 50 seconds rather than the usual 25. A pronounced effect of training time was seen, with discrepancy dropping to just $0.66\,^\circ$ and SD to $0.64\,^\circ$. Thus, it can be concluded firstly that the effects of a low learning rate are, as is logical, limited to situations in which a large weight profile remapping is required and secondly that this effect can be ameliorated with increased training.

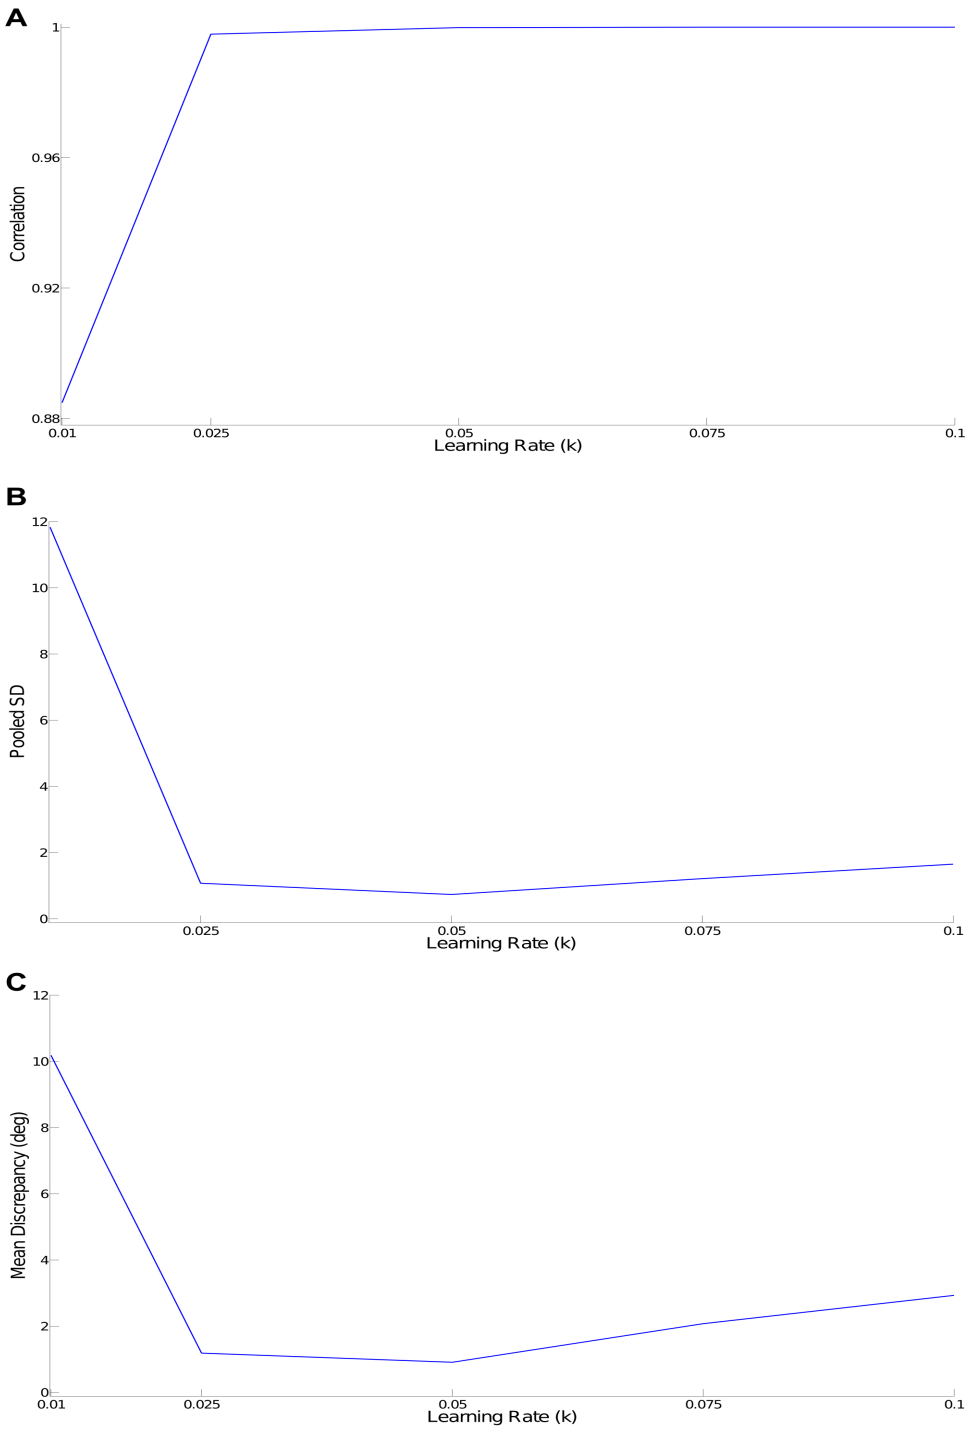


Figure 4.13: Effect of varying learning rate k, averaged across all conflict sizes. A: Correlation between initial packet undershoot following cue conflict and mean weight shift across all HD cells following agent rotation and path integration. B: Pooled SD of weight shift across all cells in the HD layer. C: Mean absolute discrepancy between initial packet undershoot following cue conflict and mean weight shift across all HD cells following agent rotation and path integration.

As learning rate increases, there are improvements in all measures of remapping accuracy and uniformity. However, beyond $k > 0.05$, we see a mean discrepancy and SD begin to worsen. Unlike the effects of low values of k, this detriment does not seem to be specific to any particular conflict size (results not given). Figure 4.13 presents, for $80\,^\circ$ conflict and three different values of k, mean weight shifts for the receiving profile to each cell in the HD layer. Looking at the right hand plot for $k = 0.1$ and comparing to the other two plots for $k = 0.01$ and $k = 0.05$, it is clear that not only is there a larger general discrepancy between weight shift and packet undershoot, but that there is also a larger internal discrepancy in the remap within the HD cell layer. At high values of k, the weight profile is exceptionally labile. This means that, as the agent rotates and activity packets are propagated through their respective rings, rather than a small change in the feed-forward weight structure that gradually results in isomapping, the weight profile learned after cue conflict is simply relearned in its entirety at a slightly transposed location every timestep.

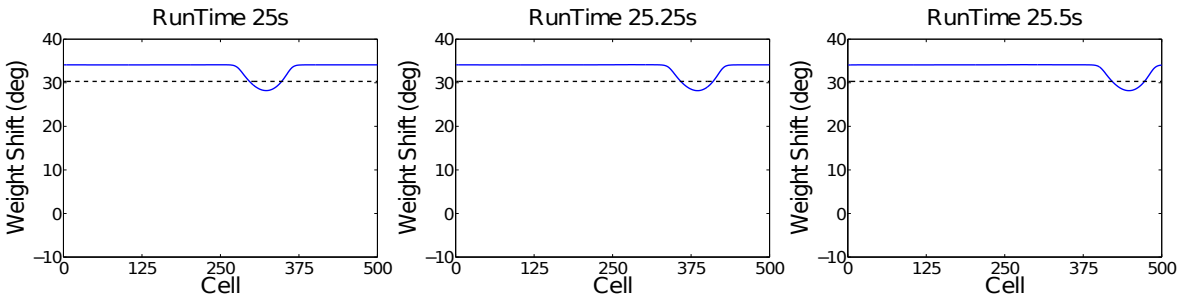


Figure 4.14: Mean shifts in the feed-forward weight profile for all cells in HD layer following simulation with learning rate of $k = 0.1$ after varying amounts of simulated time. A high learning rate has resulted in transposition of the initial weight profile. This is in contrast to gradual overwriting, leading to isomapping, at low values of k.

This effect is illustrated by Figure 4.14, which shows feed-forward mean weight shifts before and after simulation for every cell in the HD layer at different points during simulation with a learning rate value of $k = 0.1$ and conflict size of $80\,^{\circ}$. At each time, the post-conflict weight profile has simply been learned transposed into a new location. This is a direct result of a high value of k causing 'flipping' of the weight profile. Retention of this initial weight profile thus results in an increased SD and mean discrepancy.

4.3.6 Effect of path integration strength

A full exploration of the parameter λ, controlling the strength of path integration input during agent rotation, was made. Similar to Figure 4.13, Figure 4.15 presents summary results from simulations run with varying values of λ in the range [0.1,40]. Below $\lambda = 0.3$, there is a pronounced deterioration in pooled SD, mean discrepancy and correlation. When path integration strength is weak, the HD activity packet location is no longer principally determined by path integration input, but is overridden by the now relatively strong visual input. However, following cue conflict the feed-forward weight profile has been partially remapped and the relationship between the visual and HD rings is now inconsistent. This results in a HD activity packet motion that does not accurately track visual activity packet motion. As feed-forward weights continue to update during this attempted path integration, the overall weight profile becomes increasingly erratic with continued simulation.

Figure 4.16 provides data evidencing this process for a conflict size of $170\,^{\circ}$ and path integration input strength of $\lambda = 0.1$. On the top, we see HD layer activity over the course of simulation, showing a highly erratic HD activity packet motion during agent rotation. Beneath this, we see how much the feed-forward w^2 weights have shifted from their initial configuration at various points during simula-

tion (N.B. this shift is always measured relative to the initial weight configuration). After cue conflict, the initial shift in weight is exactly as would be expected, localised to the HD layer subregion involved in the initial cue conflict. Indeed, this initial post-cue-conflict remapping is identical to that seen in simulations reported in other sections. However, in this simulation weak path integration input is not enough to force the HD packet to accurately track the visual ring activity packet despite this initial distortion of the weight profile, as the agent rotates on the spot. As can be seen from the motion of the HD layer activity packet, this results in a breakdown of path integration and an erratic HD activity packet motion. After the full session of agent rotation and attempted path integration, the final weight profile (bottom right) is significantly distorted, resulting in the observed breakdown in measures of remapping accuracy and coherence.

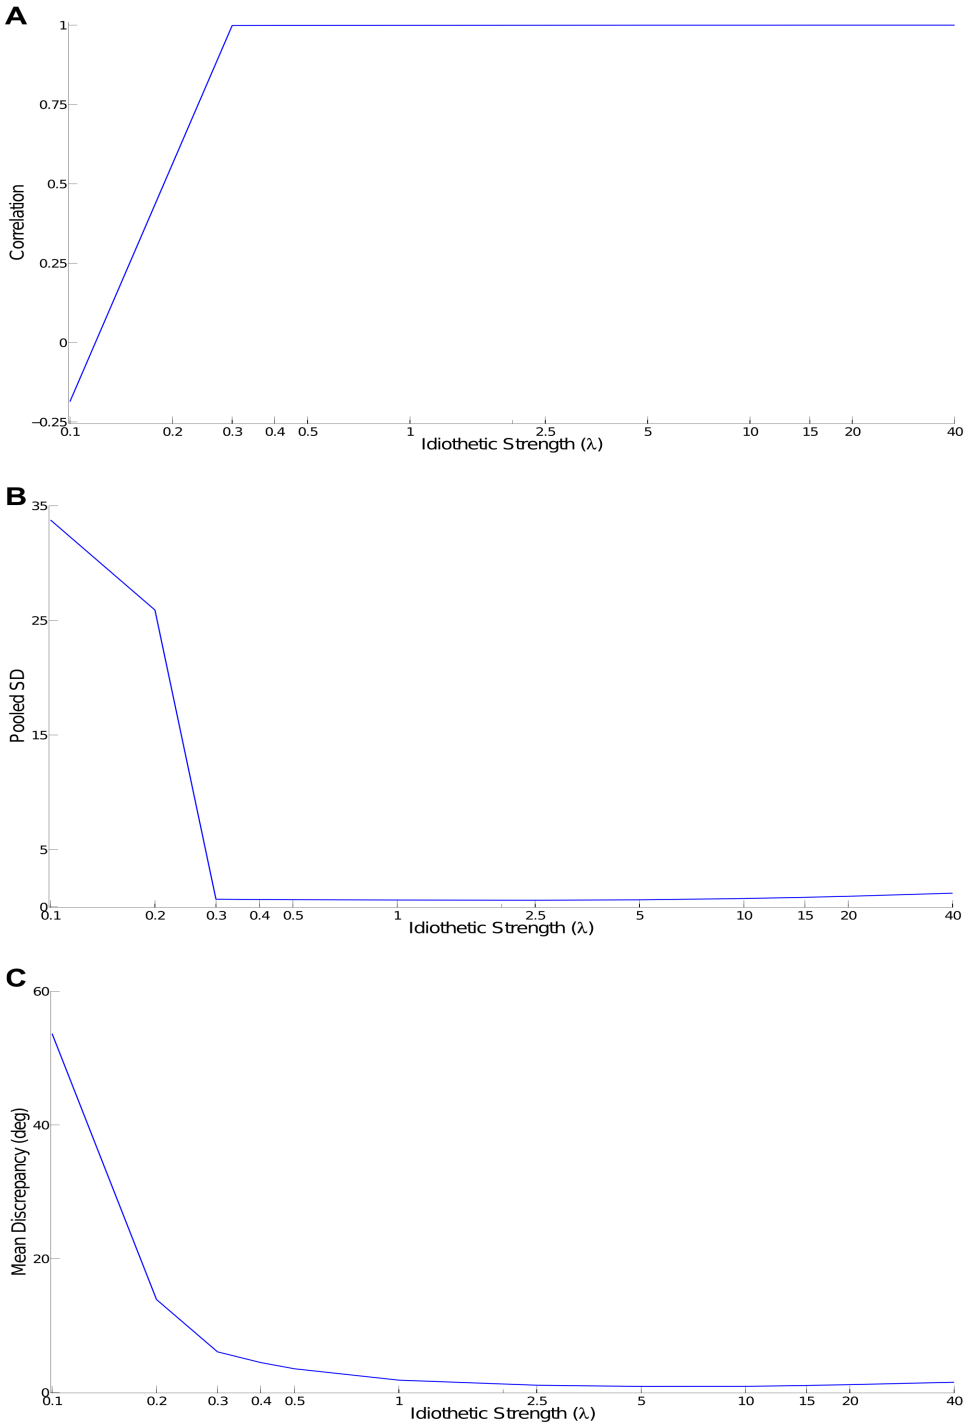


Figure 4.15: Effect of varying strength of path integration input λ, averaged across all conflict sizes. Conventions as in Figure 4.13. Note: as the data covers a large range of λ values, the x-axis is on a base 10 logarithmic scale.

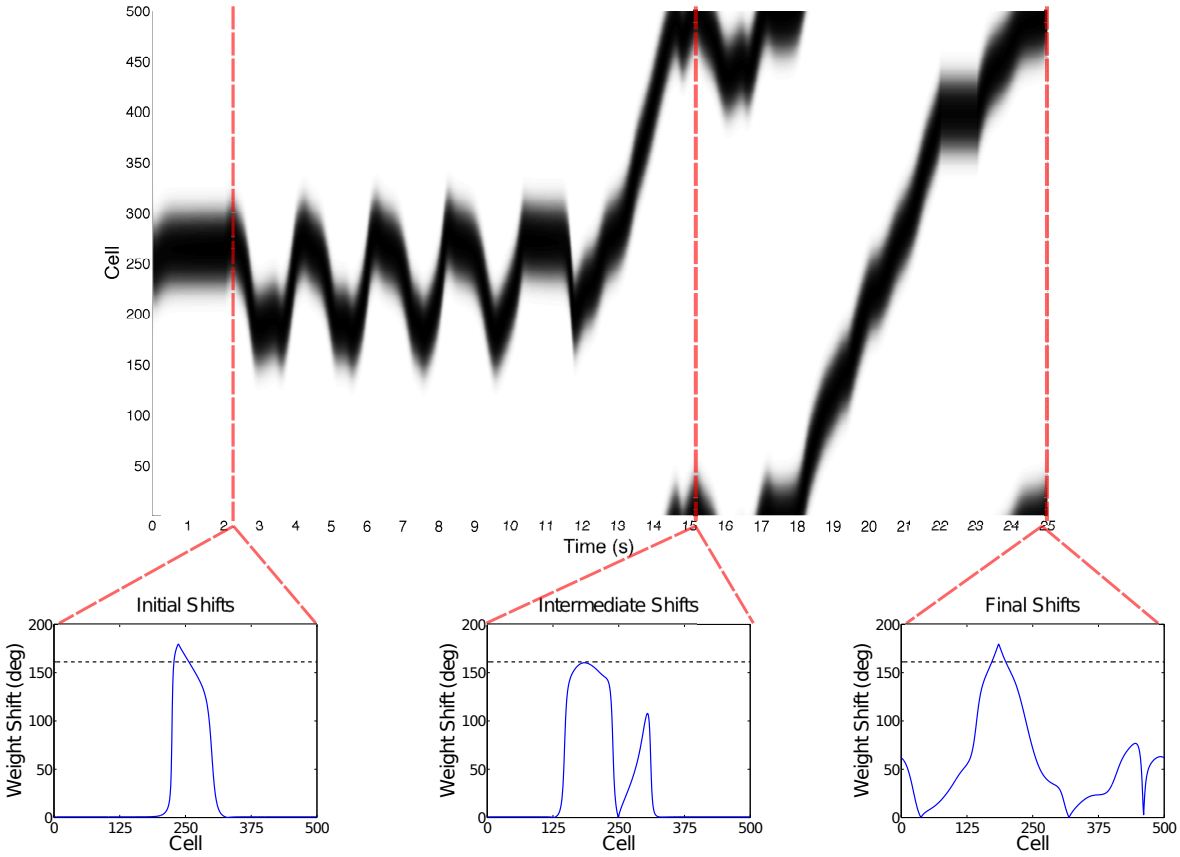


Figure 4.16: Demonstration of the effect of low values of λ on both HD layer activity packet motion (top), and mean shift in the feed-forward weight profile for all HD cellsat three different points in simulation (bottom). After cue conflict (bottom left), shift in weights is as expected. However, as the simulation progresses into the rotation phase, path integration input is too weak to override visual input. Therefore, HD activity packet location is determined solely by visual input, which arrives via the distorted weight profile. As a result, the HD packet fails to track the visual activity packet. Hebbian learning associates visual activity with the incorrect HD ring locations, causing erratic, incorrect, weight changes (bottom middle). This exacerbates the loss of HD packet tracking, which in turn magnifies erratic weight changes. The final result is a disordered mapping of feed-forward weights onto the HD layer (bottom right) and a total breakdown of HD packet tracking (top).

Interestingly, this disruption of the feed-forward weight profile appears to have most of an effect on the largest and smallest conflict sizes: for $\lambda = 0.1$, conflicts in the ranges 10° to 60° and 130° to 180° show the largest SD values (not shown). This is most likely due to the breadth of weight profile distortion combined with the magnitude of this distortion. For small and large conflict sizes, the absolute packet motion during cue conflict is less than for intermediate conflict sizes (this fact is evident from a combination of Equation (4.3) and Figure (4.4) of this chapter). This means that any initial

changes in feed-forward weight structure during cue conflict take place over a smaller area. As a result of weight normalisation , weight changes over a larger area (i.e., weight changes following cue conflict for intermediate conflict sizes) are weaker in magnitude, whilst those over a smaller area are stronger. Therefore, small and large cue conflict sizes have a concentrated region of weight changes, whilst intemediate conflict sizes, with a similar magnitude of undershoot (and therefore remapping) to small conflicts, have a larger spread of weights and thus a reduction in feed-forward weight strength.

Since the amount of feed-forward weight remapping is directly influenced by packet undershooting during cue conflict, this means that weight changes for large undershoots are also of a larger magnitude. For conflict sizes beyond $130\,^{\circ}$, packet undershooting following cue conflict begins to increase significantly. These conflict sizes have a concentrated region of weight changes of a large magnitude due to their large undershoot. Since the initial feed-forward weight profile disruption is the cause of path integration disruption, large and small conflict sizes with their relatively stronger weight changes are more able to overpower the relatively weak path integration inputs. This disruption in the path integration phase causes an erratic final weight profile which is significantly less homogenous than for intermediate conflict sizes, reflected by a breakdown of model performance on all measures of remapping homogeneity and accuracy seen in Figure 4.15.

Eliminating the path integration input completely ($\lambda = 0$), as expected, causes a total deterioration of isomapping across the board. With no path integration input, any initial weight changes will have a detrimental effect on packet motion during the rotation phase of simulation. Learning during this erratic HD packet motion will further affect the weight profile, resulting in a final feed-forward weight

profile which has not isomapped or remapped by the correct amount. Thus, weights do not shift by the same amount, and as a whole do not shift by the correct amount. This is reflected by a high pooled SD across all conflict sizes of $45.29\,^\circ$ for weight remapping (with all conflict sizes equallty affected, giving SDs in the range $[44.5\,^\circ, 51.5\,^\circ]$), and a mean discrepancy between undershoot and mean cell shift of $55.21\,^\circ$. The lack of relationship between initial packet undershooting and mean cell shift can be best shown in Figure 4.17, which is analagous to Figure 4.4.

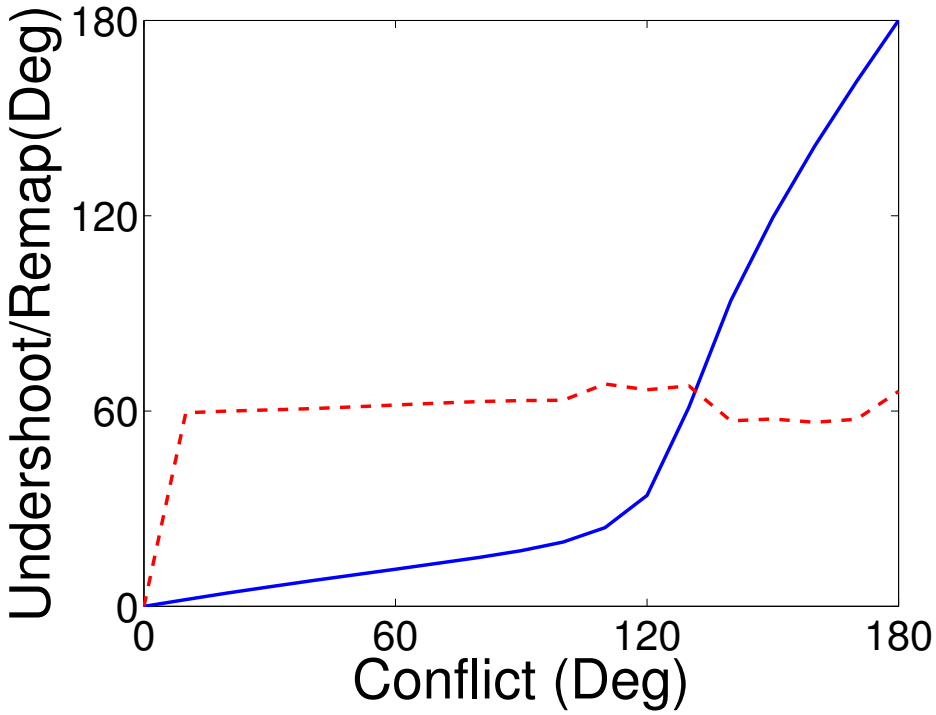


Figure 4.17: Head direction layer packet undershooting of true visual landmark location during initial cue conflict (solid blue), overlaid with mean remap in feed-forward w^2 weights following path integration, averaged across all head direction cells within a layer (red dashed), for simulations without any path integration input ($\lambda = 0$). Mean remap shows little to no relationship to packet undershooting, demonstrating a breakdown of model performance without path integration.

4.3.7 Effect of rotation speed

Figure 4.18 gives a summary, similar to Figure 4.13, of results from simulations in which the rotation speed V is varied in the range [18,360]°. This range was not exceeded, as it was considered implausible for a biological agent (i.e., a rat) to rotate consistently at a velocity in excess of $V = 360\,^{\circ}/s$. As seen in Figure 4.18A, the correlation between activity packet undershoot after cue conflict and mean weight shift after agent rotation is always very high, remaining around 0.999 for all rotation speeds. However, although the relatonship between packet undershoot and mean weight shift is very strong, it may well be that individual cells in the HD layer show considerable variation and therefore do not isomap.

Looking at Figure 4.18B shows us that the standard deviation of weight shifts, pooled across all conflict sizes, tends to be much larger for slower rotation speeds. This shows that isomapping breaks down for slower rotation velocities, and is improved with faster rotation speeds.

Figure 4.18C demonstrates that the correspondence in terms of absolute discrepancy between packet undershoot after cue conflict and weight shift after path integration is affected by rotation velocity. Overly-slow rotation speeds increase the discrepancy between weight shift and packet undershoot, which can be explained as a break down in isomapping, thereby pushing the mean weight shift away from the value determined by packet undershooting.

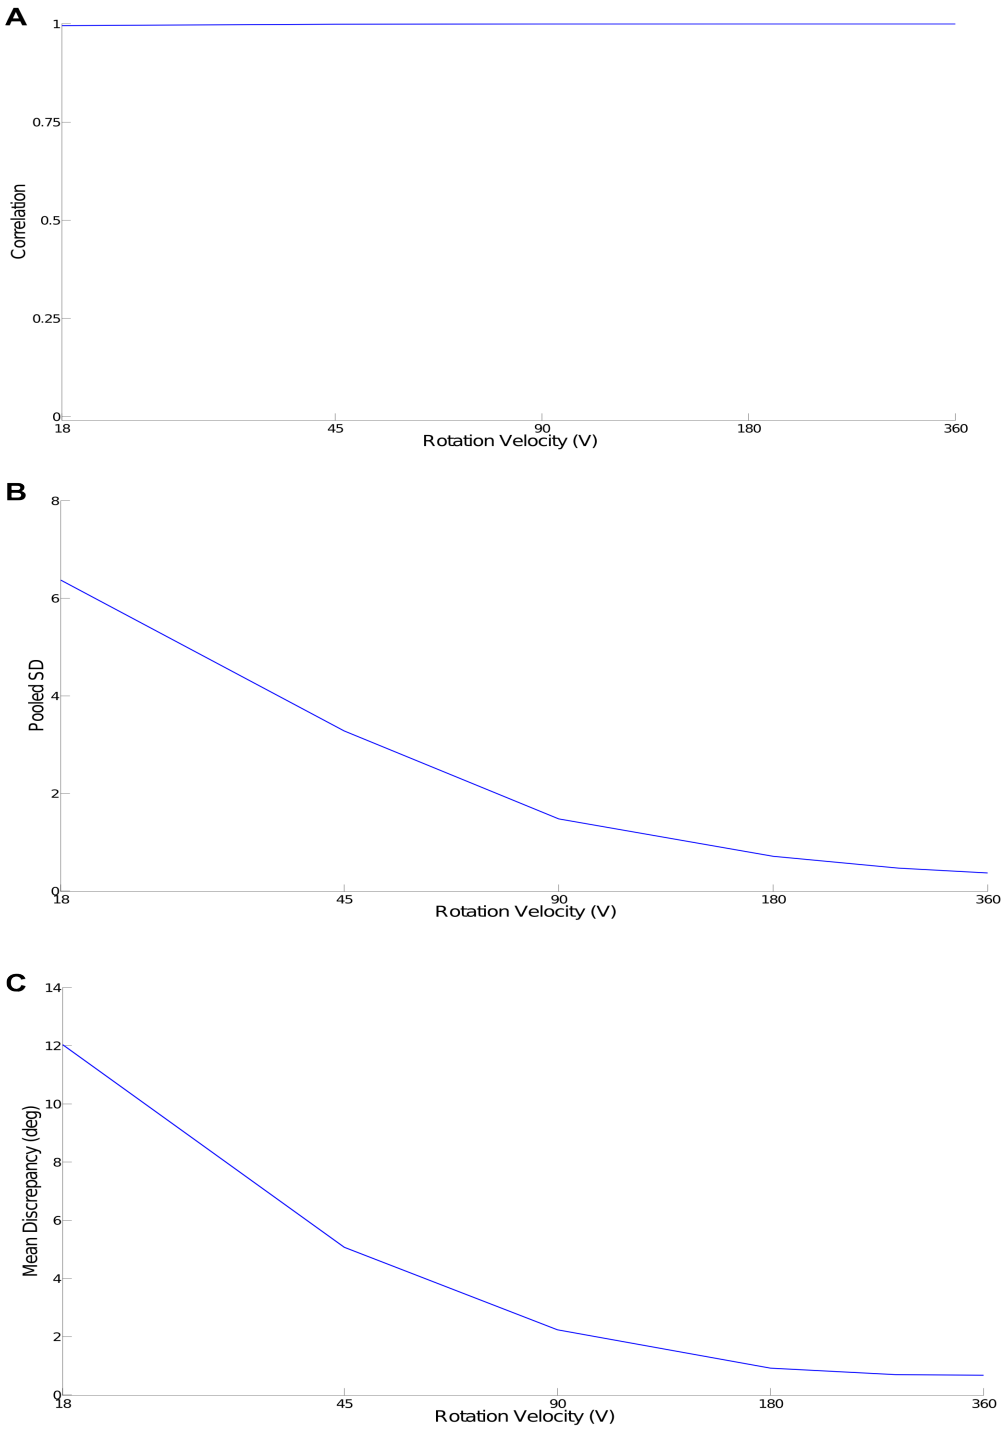


Figure 4.18: Effect of varying agent rotation speed V, in degrees/second, averaged across all conflict sizes. Conventions as in Figures 4.13 and 4.15. As in Figure 4.15, the x-axis is on a base 10 logarithmic scale.

Figure 4.19 shows the amount of remapping demonstrated by individual HD cells after simulation at $80\,^\circ$ conflict, with rotation at, from left to right, $18\,^\circ/s$ and $180\,^\circ/s$. This is compared to the 'target' remapping value representing how much the HD layer should have remapped by, based on initial packet undershoot following cue conflict. This figure shows that simulation at $18\,^\circ/s$ has failed to show the isomapping seen at $180\,^\circ/s$, demonstrating a lack of uniform weight shifts across the HD layer for slower rotation speeds during path integration.

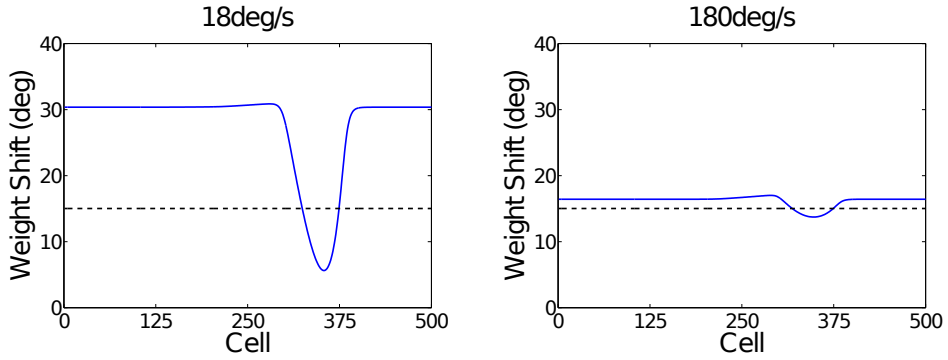


Figure 4.19: Cell shifts after simulation with rotation speeds of $18\,^\circ/s$ (left) and $180\,^\circ/s$ (right) are shown by unbroken blue line. The initial HD packet undershoots after cue conflict are shown by black dashed line. Path integration during agent rotation at $18\,^\circ/s$ has failed to cause isomapping, while rotation at $180\,^\circ$ demonstrates successful isomapping.

It is possible that this degradation in isomapping is due to fewer rotations during training at slower rotation speeds. For example, rotation at $18\,^\circ/s$ over $20s$ of path integration time would have the agent rotating just $360\,^\circ$, giving just 1 complete revolution. In comparison, rotation at $180\,^\circ/s$ over the same time period would have the agent rotating through $3600\,^\circ$, or 10 complete revolutions. At $18\,^\circ/s$ the majority of cells have only had one opportunity to update their weights, which would cause isomapping to break down. Running at faster speeds results in cells having more opportunities to adjust their weights, which allows cells to remap by a similar amount, thereby reducing the SD. Thus, pooled SD reduces with increasing values of V.

In order to test this hypothesis, simulations with a rotation speed of $18\,^{\circ}/s$ were run with double-length path integration sessions. However, no discernible change was seen in terms of the quality of isomapping, as measured by SD or mean discrepancy. However, this may well be due to an interaction effect with the learning rate parameter k. As shown in Section 4.3.5, the k parameter must be balanced so that learning allows cells to shift their weights by similar amounts gradually over successive learning opportunities. Excessively high values of k resulted in overlearning, and the post-cue conflict weight profile, with it's unequal distribution of weight changes, was simply relearnt in it's entirety at each location around the ring (Figure 4.14, this chapter). Although the current value of $k = 0.05$, did not show overlearning at a rotation speed of $180\,^{\circ}/s$, it is possible that a lower rotation speed of $18\,^{\circ}/s$ results in an HD activity packet that spends longer at each location in the HD layer, causing overlearning at $k = 0.05$.

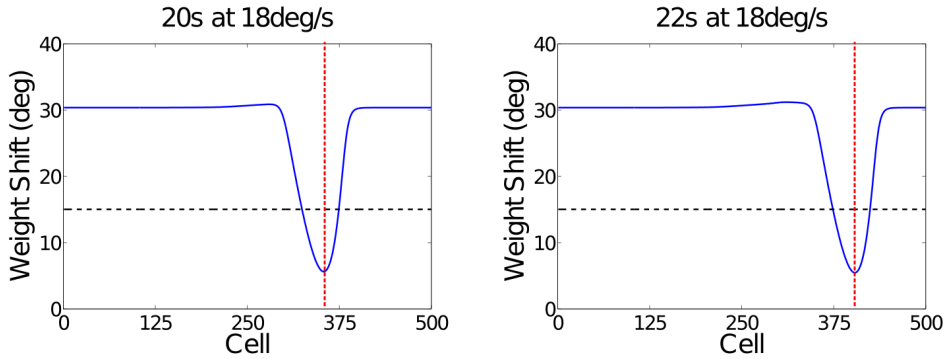


Figure 4.20: HD cell shifts after simulation with rotation speeds of $18\,^{\circ}/s$ for $20s$ of path integration (left) and $22s$ of path integration (right). Weight profile appears to be transposed into each new location, an effect of overlearning. Red dotted lines show locations of visual landmarks in weight profile, for ease of comparison. Conventions as is Figure 4.19.

This theory of overlearning becoming more apparent is supported by examining changes in the feed-forward weight structure at different points during simulation. Figure 4.20 shows HD layer cell shifts for simulations at a rotation speed of $18\,^{\circ}/s$ with $20s$ of path integration (left) and $22s$ of path integration (right). These plots clearly show that at this speed, a transposition of the weight profile occurs,

and thus no change in SD or mean discrepancy will happen over time. This is in contrast to the sort of gradual learning required to generate isomapping.

Simulations with a rotation speed of $18\,^{\circ}/s$ and conflict size of $80\,^{\circ}$ were re-run with a lower value of $k = 0.01$. The strength of visual input was also changed to $\phi_2 = 2.5$, in order to resolve cue conflict in the same manner as simulations with higher values of k. All other parameters remained the same. Figure 4.21 shows cell weight vector shifts after $20s$ of path integration (left) and $60s$ of path integration (right) using these new parameter values. At this lower value of k we no longer see interaction effects of slow rotation and learning, and thus the feed-forward weight profile evolves normally rather than transposing into each new location as seen in Figure 4.20. After $60s$ of simulation, the standard deviation for this conflict was reduced to just $1.43\,^{\circ}$. The mean remap across the HD layer was $16.05\,^{\circ}$, providing a good match with the observed cue conflict undershoot of $15.13\,^{\circ}$. Thus, this result demonstrates that interaction effects of low rotation speeds and high learning rates can be ameliorated by lowering the learning rate. This allows the model to cope with rotation speeds as low as $18\,^{\circ}/s$.

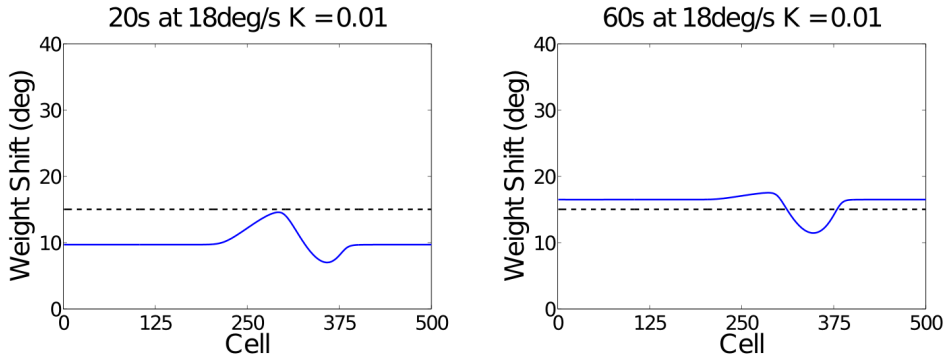


Figure 4.21: HD cell shifts after simulation with rotation speeds of $18\,^{\circ}/s$ with a lower learning rate parameter value of $k = 0.01$ for $20s$ of path integration (left) and $60s$ of path integration (right). Isomapping is beginning to take place, with the weight profile become more uniform, and the amount of remap per cell converging on the original undershoot value of $15.13\,^{\circ}$. Conventions as in Figure 4.19.

4.4 Discussion

This chapter provides operational details and simulation results of a neural network model which is not only able to resolve conflicts between ideothetic (self-motion) and allothetic (visual) information streams, but also allows the entirety of a simulated HD layer to coherently remap (isomap) in response to a cue conflict between these information streams. Using this model, it is suggested that cue conflict is resolved in a two-step process, with an initial subset of HD cells remapping during cue conflict, and the rest of the HD layer remapping in line with this subset over the course of a short period of rotation path integration. This represents a biologically plausible mechanism by which HD cells may change their preferred directions following cue conflict.

The model in this chapter is very similar to that presented in Chapter 3. It consists of a ring of visual cells coupled to a head direction cell continuous attractor layer via plastic feed-forward connections. The principal difference between the two chapters is the simulation paradigm used. The paradigm used is initially identical to that used in Chapter 3. However, following this cue conflict resolution, the agent is simulated as rotating on the spot at a constant speed. During this period, both the visual and HD activity packets move in lock-step through their respective layers at the same constant speed. Thus, the remapped relationship between the head direction cell packet and visual landmarks is preserved across all locations in the head direction cell layer, and Hebbian updates to the feed-forward weights allow other cells within the layer to shift their preferred directions in a topologically consistent manner. This eventually results in a coherent change in preferred direction across the entire head direction layer.

This two stage process represents a testable hypothesis. Due to isomapping, shifts in the firing properties of the entire head direction cell layer are typically inferred experimentally from the shifts of an individual cell or small group of cells (Goodridge et al., 1998; Taube et al., 1996). In future, recordings of ensembles of head direction cells may well prove informative. It is noted that there has been little study of head direction cells in terms of entire populations (Stackman, 2011; Johnson et al., 2005), and certainly not in the context of the fine-scaled dynamics of remapping. Thus, it remains to be seen experimentally whether this hypothesis truly holds, although for now it reflects accurately recorded data on cue conflict in vivo.

An extensive parameter exploration helps to clarify the model's operational principles. The main result is shown to be independent of cue conflict size or direction. Concordantly with the hypothesised mechanism elaborated in this chapter, accurate path integration is shown to be crucial for allowing HD layer weights to converge on the correct value, reflected by a breakdown in isomapping coherence and accuracy during inaccurate path integration. Interestingly, it seems to be only the magnitude and not the direction, of path integration accuracy which is important. It was also shown that weakening or eliminating this path integration input caused a breakdown in model performance.

Salience of visual information is shown to be an important factor in cue conflict resolution, with maximally-salient visual landmarks having complete control over HD packet locations, reflected by changes in HD cell preferred firing directions. Learning rate is also shown to be important, with too-high values resulting in a failure of isomapping. The model is more tolerant of low learning rates, simply requiring extra training for weights to converge on the correct profile. The model is shown

to be able to cope with a range of agent rotation speeds between $18\,^{\circ}/s$ and $360\,^{\circ}/s$. Although there was an interaction effect of slow rotation speeds causing overlearning, it was shown that this can be completely elminated through the use of a slightly lower learning rate.

The theory of frames of reference, elaborated in Section 4.2.1, underpins observed model behavior. Simulation results show that a conflict of 80 degrees causes the head direction cell activity packet to move 64.98 degrees. Given our simplifying assumption that actual head direction remains fixed, a shift in head direction cell activity of 64.98 degrees means that cell preferred compass directions have change by an amount equal to this shift. This matches the preferred direction change seen experimentally (Knight et al., 2014).

In the model, this initial activity packet shift is accompanied by an approximately corresponding shift in mean weight vector for an individual HD cell of 16.43 degrees. This shift in mean weight vector, corresponding to a change in relationship of head direction cells to rat-centered representations of visual space, thus causes an accompanying change in head direction cell responses to absolute direction in a room-centered, i.e., allocentric, frame of reference. Indeed, this change in weight vector acts to preserve the altered relationship between head direction cells and rat-centered visual representations, such that the initial change in preferred direction is inherited by all cells within the head direction layer.

The issue of different frames of reference can be further clarified by considering experimental results with cue cards. In the original rotation study (Taube et al., 1990b), a rat was placed in a paradigmatic

cylindrical enclosure, with a single white cue card on the wall. After experiencing the cue card as a visual landmark during navigation, the rat is removed from the arena and systematically disorientated. It is then reintroduced to the arena with the cue card rotated. It was observed that the shift in head direction cell preferred directions was approximately equal to the shift in the visual landmark cue.

In terms of the model, Taube and colleagues' results can be best understood when thinking of the shift in preferred direction in rat-centric, rather than room-centric, terms. When the rat is reintroduced to the arena, it is to the rat as if the visual landmark has not changed: there are no other orienting features in the environment. This means that the same cells respond maximally to the egocentric direction in which the landmark cue lies. For example, if cell 1 responded maximally to the direction of the light cue before disorientation and cue card shift, it will also do so afterwards. There has been no change in the relationship between head direction cells and the egocentric bearing to which they respond to the visual. However, in terms of the room, the visual cue has moved from, for example, North to East. Thus, if HD cell A responds when the rat is facing the landmark cue both before and after cue rotation, it has changed which room-centered direction it responds to from North to East.

The idea of the head direction system reorienting over a short period of time in response to changes in environmental manipulations is consistent with experimental evdience, showing that a visual landmark cue, once learned, can influence head direction cell activity in as little as 80ms (Zugaro et al., 2003). Although the results presented in this chapter represent conflict resolution within a few seconds, it is more than plausible that such a mechanism could work over a reduced time course in a spiking, rather than rate coded, model (Degris et al., 2004a). It should be noted, for clarity, that this

addresses specifically the speed at which HD cell activity can be influenced by an established visual landmark, rather than time taken to establish a stable visual landmark as a reliable cue. The latter, in terms of this model, would represent learning of the w^2 connectivity, which is instead explicitly initialised in our simulations.

Considering the structure of feed-forward w^2 connectivity, it is important to note that the model discussed in this chapter has explicitly initialised this weight matrix. Thus, this model does not comment specifically on the learning of a visual cue as a stable landmark in relation to idiothetic information. However, changes in feed-forward weight structure during a relearning of the relationship between vision and idiothetic information are addressed by this model, and could be thought of as deriving from the same mechanism involved in the initial learning of this relationship. Experimental evidence regarding the time scale over which a prominent visual landmark can gain stimulus control over the head direction cell system following environmental manipulations is informative in this regard. Several studies suggest that visual landmark cues could be learned, and subsequently used to reset head direction cell activity upon reintroduction, after a period of time in the region of 8 minutes (Taube, 1995; Dudchenko et al., 1997). A study by Goodridge and colleagues (Goodridge et al., 1998) demonstrated that, stimulus control was achieved after an 8 minute exposure to a prominent landmark, but that this control was greatly weakened if exposure was reduced to 1 minute.

The results presented in this chapter represent changes in feed-forward weight structure within a minute, which is somewhat in opposition to the experimental literature. There are a multitude of reasons why this might be so. For example, a short simulation time and higher learning rate were actively

used to reduce runtime to manageable proportions. It is more than plausible that such a mechanism could work over a slightly extended time course in vivo, particularly with a lower learning rate. Indeed, Section 4.3.7 demonstrates the desirability of a lowered learning rate, allowing the model to cope with slower agent rotation speeds without overlearning. Additionally, a caveat is raised that initial learning of a visual landmark cue and then modifications of that cue's relationship to idiothetic information need not be over the same timescales even if deriving from the same mechanisms. As will be discussed below, it may be that shorter-term changes in weight profile are responsible for isomapping, whilst longer-term changes in weight profile are involved with modulating salience. Thus the conclusions relating to speed of w^2 learning discussed in this paragraph and that preceeding ought to be taken as conjecture.

Another important result from this chapter, as in the previous chapter of this thesis, is the effect of visual cue salience. It appears that more salient visual cues exhibit stronger cue control, with cue conflict therefore being resolved more in favour of visual landmarks and head direction cells subsequently shifting their preferred directions further in favour of the visual cue, as demonstrated in data from Knight and colleagues (2014). The idea that only visual cues with a stable relationship to ideothetic information can form landmarks is well-established in the literature, and has already been detailed in the discussion of the previous chapter (McNaughton et al., 1995; Goodridge et al., 1998; Knierim et al., 1995; Jeffery, 1997).

Results of this and the preceeding chapter concur with these experimental observations: less salient visual cues are weakened as model inputs, resulting in weaker cue control. Salience affects how well

a visual cue is learned as a stable landmark. In terms of the model, this would affect the strength and topopgraphical organisation of the w^2 weight profile. Whilst we simulate the effect of salience via a uniform scaling of prespecified w^2 connectivity, visual cue stability likely has an effect on how the initial w^2 connectivity is formed. These changes in salience represent w^2 plasticity over a longer timescale. In contrast, visual to HD remapping and isomapping is relatively more labile, occuring within a trial rather than across trials. This shorter-term plasticity would happen on top of an initialised w^2 profile. A future modeling direction would be to explicitly simulate the development of the w^2 weight profile in a novel environment. This would in essence encapsulate the concept of salience without, as in the current model, the need for an explicit parameter: only stable visual cues would form ordered w^2 connectivity and thus act as landmarks.

The model presented in this chapter is a bare-bones substantiation of a hypothesis, and thus contains several simplifications that bear mentioning. Firstly, it makes use of an abstracted term, $e_i(t)$, which represents idiothetic input to head direction cells, and is used to shift activity through the head direction layer at the correct speed during simulated path integration. This provided a simple and uncluttered explanation of the hypothesised remapping mechanism in action during cue conflict. Future simulations, however, could well incorporate detailed neuronal structures for performing path integration. Walters and Stringer (2010); Walters et al. (2013) have modelled path integration in a similar architecture, with the path integration inputs derived from a layer of cells responding to combinations of angular head velocity and head direction. These proposed combination cells are plausible given head direction cells in the lateral mammillary nucleus which have been shown to modulate their firing on the basis of angular head velocity (Stackman and Taube, 1998) and populations of angular head

velocity cells in the dorsal tegmental nucleus whose firing is modulated by head direction (Bassett and Taube, 2001; Sharp et al., 2001). Such cells have been shown to be able to shift head direction cell activity packets at approximately the same speed as the animals head rotation during path integration (Walters and Stringer, 2010).

A related issue is the implementation of recurrent collateral connections within the head direction cell layer, used to sustain activity in the absence of any external input. However, there seems to be little evidence that head direction cell regions display such connectivity (Jeffery, private communication). Provisional simulation results demonstrate that these recurrent connections could instead exist across multiple layers. In the brain, reciprocal connectivity between LMN and DTN (Allen and Hopkins, 1988, 1989) could support this architecture. DTN is known to contain cells which fire for angular head rotations, and which have some firing even for no head rotation (Bassett and Taube, 2001). Thus, a combination cell layer could develop with cells that fire for a combination of head direction and no rotation, and connections from this layer back onto head direction cell layers could plausibly provide attractor-like behaviour whilst avoiding reliance on recurrent collateral connections within a layer.

The simulation paradigm used involves only rotation in a single direction at constant speed, and distinctly separated cue conflict and path integration phases, by virtue of the simplifying assumption we make that the agent remains stationary during cue conflict. Of course, this paradigm is hardly ecologically valid. However, separating cue conflict and path integration when simulating changes in cell preferred direction is necessary to elucidate the mechanism by which cell preferred directions are updated. *In vivo*, these two stages would overlap. One can see how the mechanism described above

could be adapted to more ecologically-valid simulation paradigms, with remapping due to cue conflict being inherited across an entire head direction cell layer even as this cue conflict is being resolved, and with the final shift in preferred direction remaining a compromise between visual and idiothetic information sources.

To summarise, this chapter presents a biologically plausible neural network architecture capable of allowing HD cells to isomap by the correct amount following resolution of a conflict between idiothetic and allothetic information. An extensive parameter exploration clarifies key aspects of this hypothesised mechanism, and demonstrates its robustness. This acts as an extension to the hypothesis outlined in Chapter 3, and together these chapters form a novel theoretical contribution, underpinning experimental data. During a discussion of the results of this chapter's simulations, several interesting future avenues of research, both theoretical and empirical, are outlined.

Chapter 5

Factors affecting path integration speed

Results from the previous chapter of this thesis demonstrate that isomapping following cue conflict is only accurate when path integration input is fully accurate. Indeed, this inaccuracy, with a target packet speed of $180\,^{\circ}/s$, need only be $1\,^{\circ}$ to result in unstable isomapping. Although this does seem a very harsh constraint, it should be noted that whilst path integration inaccuracies of this magnitude cause negative effects in isomapping, it may be that these effects are less pronounced *in vivo*. For example, the simulations in the previous chapter were over a shorter timescale than seen experimentally, principally for practical reasons with respect to simulation runtimes and result file sizes. This resulted in the use of a learning rate parameter value of $k = 0.05$. However, it may be that learning occurs more slowly in the rat brain, and thus the effects of speed mismatches on isomapping are less exaggerated. Regardless, it is prudent to perform a full investigation of the accuracy contraints of path integration in a rate-coded neural network.

The previous model utilised an abstracted input term, e_i in order to shift the activity packet at precisely the correct speed. Although there are some models employing recurrent weight changes to shift the

HD cell activity packet (e.g., Zhang 1996), most models use some variety of external input to shift activity (for a summary, see Touretzky 2005). A major recent theory of how time-accurate path integration is achieved (Walters and Stringer, 2010; Walters, 2011)) is based around the idea of forming associations over delays naturally found in the brain. This would allow an association to be made between a starting head direction and a predicted head direction after head rotation at a known speed over this delay. A diagram explaining the general mechanism of Walters and Stringer (2010); Walters et al. (2013) is presented below in Figure 5.1:

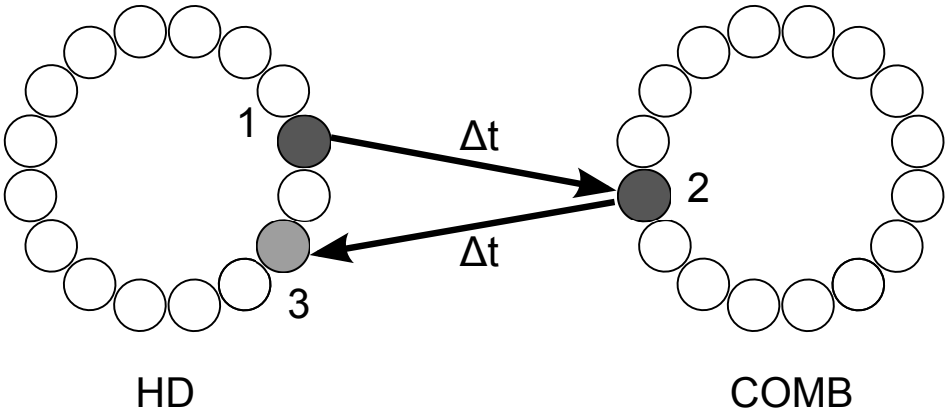


Figure 5.1: Mechanism of time-accurate path integration from Walters et al. (2010,2013). During training in the light at a constant velocity, activity packets will move through HD and COMB rings at the same speed. However, due to delays in the firing of neurons (Δt) postsynaptic cells will be associated with presynaptic activity Δt in the past. As a result, COMB cell 2 will receive maximally from HD cell 1, and HD cell 3 will receive maximally from COMB cell 2. During path integration in the dark, rotation of the head from starting location 1 will thus drive up COMB cell 2, and in turn HD cell 3.

This model has a layer of HD cells, connected to a layer of COMB cells, which respond to a combination of head direction and angular head rotation. It is first trained in the light with explicit visual and head rotation information. Crucially, there is a delay in the transmission of activity between neurons - this is either attributed to axonal conductance or neuronal time constant. Effectively, this causes postsynaptic activity to be associated with presynaptic activity some delay in the past. This means that whilst activity packets move at the same speed through the HD and COMB layers, their connectivity will be offset. Each HD cell projects maximally , via the COMB layer, onto a postsynaptic HD cell

with an offset based on the delay and the target velocity. This can be thought of as HD cells predicting the future location, after a particular delay, of an activity packet moving at a particular velocity. This is a similar idea to the theories of McNaughton et al. (1991). During path integration, this offset connectivity will act to shift an activity packet through the HD layer at the correct speed.

In these studies however, path integration speed was found again to be inaccurate, with packet speed during testing reaching a maximum of 81% of the trained speed. A logical next step to extending the previous chapter's model would be to incorporate such a biologically-accurate path integration system. Therefore, it is important to understand how such a sytem can operate with improved accuracy to avoid errors in isomapping. Two potential sources of path integration error have been identified: neuronal rise time and symmetric recurrent connections within layers.

Path integration error due to neuronal rise time

One potential source of error is the time between when a neuron receives input and when it starts to fire, known as rise time. As described above, a network trained with an activity packet moving through a ring of HD cells at a constant speed will generate an offset profile of synaptic weights which can later act to shift HD layer activity at the training speed. During the self-organisation of this weight profile, rise time will not have an effect. This is because both pre- and post-synaptic HD cells will be driven up, dynamically, by the same visual input, generating the same rise time. Thus, neurons will always be associated with synaptic connectivity that contains an offset equal the amount by which the visual activity packet has moved during the delay.

Once an offset activity profile of connectivity is learnt, this will produce an activity packet which moves around the HD ring, replaying approximately the activity packet speed during training. However, if we consider a given learned profile of connectivity with an offset, rise time will cause an undershoot in packet speed. In effect, the delay is now too long by the amount rise time adds. Therefore, by the time an activity packet is driven up, it already ought to be in a slightly different location (indeed, the difference in this location being the amount the activity packet would have moved, travelling at the target velocity, over the course of the rise time). This has the effect of causing packet speed to be slower than it ought to be for that particular learnt weight offset.

This chapter demonstrates the slowing effect of neuronal rise time on packet speed. Even with a correct offset, whether learnt or pre-specified, we see that the network cannot reproduce the target packet speed. Simulation results suggests that rise time in relation to conduction delay is a crucial factor in determining replay inaccuracy. This represents a major contribution to inaccuracy in path integration models to date, and an important consideration by which to improve future models.

Path integration error due to symmetric recurrent connections

Another potential source of inaccuracy may come from symmterical recurrent collateral connectivity, which is required for continuous attractor neural networks (CANNs). Let us consider a network of HD cells with non-offset recurrent input. That is, each postsynaptic cell receives a Gaussian profile of recurrent collateral weights centred on itself. In addition to this, cells also have offset recurrent input, whereby each cell receives a second Gaussian profile of recurrent collateral weights which are centred on a location of the HD ring offset by some amount from its own location. For a given postsynaptic

cell, both receiving weight profiles can be thought of as Gaussian distributions in the form

$$ae^{-\frac{(x-b)^2}{2c^2}}$$

where e is Euler's number, a is the curve height, b is the central location of this peak, and c is the standard deviation, or width, of the curve. Figure 5.2 shows, for postsynaptic HD cell 250, the net effect (black dashed line) of receiving both non-offset Gaussian weights (blue line) and Gaussian weights with a variety of offset sizes. (red line). From left to right, this latter profile is offset by $10\,^\circ$, $60\,^\circ$ and $120\,^\circ$ respectively.

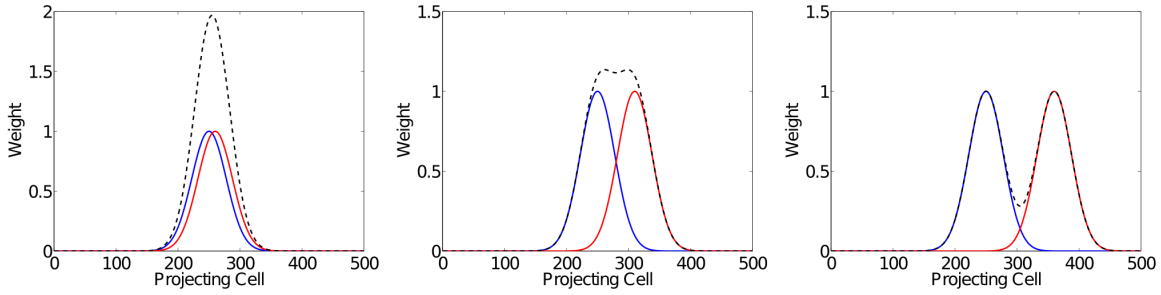


Figure 5.2: The effect of adding Gaussian profiles. A non-offset component (blue) is added to an offset component (red) of varying offset sizes: $10\,^\circ$ (left), $60\,^\circ$ (middle) and $120\,^\circ$ (right). This produces a resultant profile(black dashed). Whilst the exact effect of the addition varies with offset size, it will always act to disrupt the desired offset.

Adding a symmetrical non-offset profile to an offset weight profile has a different effect depending on their separation (i.e., on how large the offset is). For smaller offsets, the resultant profile is still approximately Gaussian with single peak, but with a reduced offset size (Figure 5.2, left plot). As offset size increases (Figure 5.2, middle and right plots), the resultant profile is no longer Gaussian, and eventually has two peaks (Figure 5.2, right plot). In all of these cases, the addition of a symmetrical non-offset component disrupts the desired offset. It is clear that any situation involving the addition of a symmetrical non-offset weight profile due to symmetrical recurrent collateral connections is po-

tentially troublesome for CANN models of head direction cells.

The neural network models in this chapter utilise a single layer of neurons, with recurrent connections that can contain both an offset component, representing path integration input, and non-offset recurrent collaterals, as have been included in the majority of HD models thus far to allow cells to maintain a representation of head direction in the absence of visual input. For simplicity and clarity, these connections have been reduced to a single recurrent layer. The simulations in this chapter present evidence that the inclusion of non-offset recurrrent collateral connectivity has a slowing effect on packet motion, preventing a network from accurately reproducing a trained target velocity. This theoretical contribution suggests that current continuous attractor-type path integration architectures containing non-offset recurrent collateral connectivity need to be re-evaluated as a biophysically and neurophysiologically viable model of time-accurate path integration.

5.1 Experiment 1: Pre-wired network

In order to investigate the factors governing packet motion, a bare-bones model is required in order to make clear-cut conclusions. Thus, initially a pre-wired neural network is used, consisting of only a single layer of cells connected with itself. These connections can contain, as described in the introduction offset components and non-offset components. This is used, in an abstract manner, to represet the driving force of path integration input in the HD system (offset input) and the effect of symmetric recurrent collateral weights within the HD system (non-offset input). This gives an ideal ”toy” system in which to simply and clearly investigate the relative effects of these two sources of input to the HD system, as well as providing a platform from which to study the parameters affecting packet

speed. In particular, the relationship of the neuronal time constant in relation to conduction delay is hypothesised to represent a crucial governing factor.

5.1.1 Model description

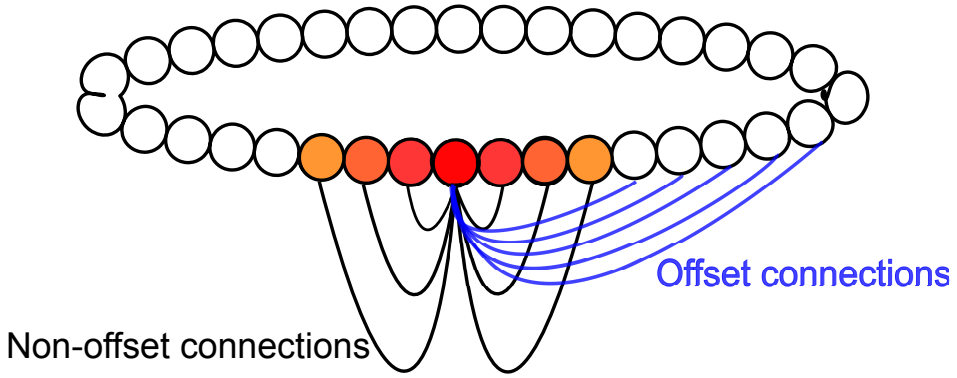


Figure 5.3: Stylised model schematic diagram, showing offset and non-offset connectivity profiles.

The model uses an architecure based on a continuous attractor neural network (CANN), and is represented as a schematic in Figure 5.3. It consists of a single ring of head direction cells. Each cell is prespecified with a preferred firing direction, with these preferred firing directions being distributed evenly around the ring from $0 - 360\,^{\circ}$. Head direction cells are arranged topographically: cells with similar preferred firing directions are located next to one another. Thus, a Gaussian packet of activity in the head direction cell layer represents the current head direction of the agent. Individual head direction cells are modelled using a leaky-integrator firing rate model, described in more detail below. Individual spikes or spike trains are not modelled, and instead neuronal firing is represented as an instantaneous average firing rate for each cell.

All cells influence one another via excitatory recurrent collateral connections. Recurrent weights are pre-wired with Guassian profiles to have either an offset profile or a combination of an offset profile

with an added non-offset component of varying strength. These are constructed in a Gaussian configuration: offset weights have an included term, such that each cell projects maximally not to itself but to an offset cell in the same layer. This offset acts to drive activity through the network at a fixed speed, intended to match a target speed, V°/s. In some simulations, this offset profile will have added to it a non-offset Gaussian weight profile. This additive profile is initialised in the same way, but is centered so that each cell will project maximally back onto itself. It should be noted then, that the combined profile from adding an offset and non-offset weight profile will result in a profile where individual cells project maximally neither onto themselves or onto the expected cell given the offset weight profile alone, but rather onto an intermediate cell between these two locations. The exact location of this intermediate cell is determined by the relative strength of the offset and non-offset weight components.

Recurrent collateral connections contain a delay Δt. This represents a delay inherent in axonal conduction, the time it takes for action potentials to be conveyed along the length of an axon from near the soma to axon terminals. Incorporating an axonal conduction delay into the model means that postsynaptic cells are activated by presynaptic firing only after a time delay of Δt. This delay is crucial for the model to shift a packet of activity around the HD ring.

5.1.2 Model operation

5.1.2.1 Model equations

The activation level h_i^{HD} of head direction cell *i* is given by

$$\begin{aligned}\tau^{HD}\frac{dh_i^{HD}(t)}{dt} =& -\ h_i^{HD}(t) + \ e_i(t) \\ & -\ \frac{1}{N^{HD}}\sum_j \tilde{w}^{HD} r_j^{HD}(t) \\ & +\ \frac{\phi_1}{C^{HD\to HD}}\sum_j w_{ij}^1(t) r_j^{HD}(t-\Delta t)\end{aligned} \tag{5.1}$$

where $h_i^{HD}(t)$ is the activation level of head direction cell *i* at time *t*. The term $-\,h_i^{HD}(t)$ represents a decay term, whereby in the absence of presynaptic input, the activation of head direction cell *i* will fall to zero as determined by the time constant τ^{HD}. The term $e_i(t)$ represents external visual input to head direction cell i at time t. This input term is used to generate an initial packet within the head direction cell layer. It is determined by a Gaussian profile described in the simulation protocol below. The term $\frac{1}{N^{HD}}\sum_j \tilde{w}^{HD} r_j^{HD}(t)$ represents inhibitory feedback within the head direction cell layer, summed over all presynaptic head direction cells *j*. $\tilde{w}^{HD}$ is a global constant describing the effect of inhibitory interneurons, and N^{HD} is the total number of head direction cells in the layer. The term $\frac{\phi_1}{C^{HD\to HD}}\sum_j w_{ij}^1(t) r_j^{HD}(t-\Delta t)$ represents excitatory input from the head direction cell layer back onto itself via recurrent collateral synapses with an axonal conduction delay of Δt, summed across all presynaptic head direction cells connected to cell *i*. It consists of $r_j^{HD}(t)$, the firing of presynaptic cell *j* at time *t*, and $w_{ij}^1(t)$, the synaptic weight from presynaptic head direction cell *j* to postsynaptic head direction cell *i* at time *t*. It is important to note that this weight is not updated during simulation, and thus remains at the same initial value throughout. The entire term is scaled by the factor $\frac{\phi_1}{C^{HD\to HD}}$, which controls the overall strength of recurrent synapses with the head direction cell layer where ϕ_1

is a constant and $C^{HD \rightarrow HD}$ is the number of synapses each postsynaptic head direction cell receives from presynaptic head direction cells.

The firing rate $r_i^{\text{HD}}(t)$ of head direction cell *i* at time *t* is calculated as the hyberbolic tangent function of the activation level of cell *i* at time *t*. This function is naturally bounded in the interval [-1, 1]. However, real neurons cannot have a firing rate of < 0. Thus, firing rate is bounded at zero, giving a firing rate equation of

$$r_i^{HD}(t) = \begin{cases} tanh(h_i^{HD}(t)) & \text{if } tanh(h_i^{HD}(t)) \geq 0 \\ 0 & \text{otherwise} \end{cases} \tag{5.2}$$

Recurrent collateral weights from head direction cells back onto the head direction cell layer are pre-wired with slight variations depending on whether they are to be purely offset or a mixture of offset and non-offset components. Offset weights are initialised using the Gaussian function

$$w_{ij}^1 = e^{-(s_{ij}^{RC})^2/2(\sigma^{RC})^2} \tag{5.3}$$

where w_{ij}^1 represents the synaptic weight from presynaptic head direction cell j onto postsynaptic head direction cell i and σ^{RC} is the standard deviation of the Gaussian profile. s_{ij}^{RC} is the difference between the preferred head directions x_i and x_j of the post- and pre-synaptic head direction cells i and j respectively. It is given by

$$s_{ij}^{RC} = MIN(|x_i - (x_j + O)|, 360 - |x_i - (x_j + O)|) \tag{5.4}$$

which creates a wrap-around effect, with weights between head direction cells as a population remaining continuous across the $360/0\,^{\circ}$divide. Note that the presynaptic preferred directions are incrememted by a fixed offset, O. This offset O is calcuated based on the target packet speed $V\,^{\circ}/s$ and the axonal conduction delay Δt as

$$O = V\Delta t \tag{5.5}$$

which effectively pre-wires connectivity with an offset that matches the amount by which a packet ought to have moved over the course of the conduction delay, given a particular target packet speed.

In some simulations, the weight profile is initialised with a combination of offset and non-offset components. In this case, the resulting weight profile is a simple addition of both offset weights, as calculated above, and non-offset weights. These non-offset recurrent collateral weights are pre-wired using the same method, with the change that presynaptic preferred firing directions are no longer incremented by O, such that the previous equation is now calculated as

$$s_{ij}^{RC} = MIN(|x_i - x_j|, 360 - |x_i - x_j|) \tag{5.6}$$

The relative strengths of the non-offset and non-offset weight components are modified by the parameter λ^{NO}, which acts as a multiplier on the non-offset weight component. Thus, the λ^{NO} parameter can be thought of as determining the degree of asymmetry via its effect on the non-offset weight component: as λ^{NO} decreases, the overall asymmetry of the weight profile tends towards O. In both situations (offset weights and mixed weights), the weight profile is normalised after initialisation. This is achieved by ensuring that the square root of the sum of the squares of afferent weights, for a

particular postsynaptic cell *i*, is limited to 1. This is computed as

$$\sqrt{\sum_{j}(w_{ij}^{1}(t))^{2}} = 1 \tag{5.7}$$

The differential equations given for this model cannot be solved analytically. As in the previous chapters, they are solved numerically by a Forward Euler finite difference scheme. The interested reader should refer to Chapters 3 and 4 for full details.

5.1.2.2 Simulation protocol

At the beginning of each simulation, firing rates r_i of all head direction cells are set to zero. Recurrent collateral connections are pre-wired as offset or a mixture of offset and non-offset, as detailed in Equations (5.3), (5.4) and(5.6). Synaptic weights are then normalised as in Equation (5.7). An external input, e_i in Equation (5.1), is applied for $200ms$, and both the activation $h_i^{HD}(t)$ and firing rate $r_i^{HD}(t)$ of all head direction cells are updated according to Equations (5.1) and (5.2). e_i is calculated for each head direction cell i by a Gaussian profile

$$e_i = \lambda^{HD} e^{-(s_i^{HD})^2/2(\sigma^{HD})^2} \tag{5.8}$$

where λ^{HD} is a scaling factor determining the strength of this input to head direction cell i and σ^{HD} is the standard deviation of the Gaussian profile. s_i^{HD} is the difference between the true head direction x and the preferred direction x_i of postsynaptic head direction cell i. It is given by

$$s_i^{HD} = MIN(|x_i - x|, 360 - |x_i - x|) \tag{5.9}$$

which creates a wrap-around effect, with response profiles of head direction cells as a population remaining continuous across the $360/0\,^{\circ}$divide.

This input acts solely to initialise an activity packet within the head direction cell layer. Due to the nature of these equations, the initial location of the head direction cell activity packet is determined by x. After $200ms$, external input is removed, and the network is simulated for a further $2s$. During this latter phase, head direction cell activations and firing rates continue to be updated as in Equations (5.1) and (5.2). It is important to note that, during this second phase of simulation, head direction cell activity is sustained solely through recurrent collateral activity. The speed at which the head direction cell activity packet moves in this latter phase is the critical behaviour to be observed, and to be compared to the target speed.

5.1.3 Results

In situations where offset and non-offset weight profiles have been combined, it is necessary to check the new value of the offset O_{new} for each presynaptic cell j. This is calculated in stages as follows.

First, the center of mass of the final weight vector is calculated for each head direction cell within the layer by a similar method as seen in Chapter 4 of this thesis. However, unlike the previous weight vector measure, this is done from the perspective of a presynaptic cell. In previous chapters, head direction cells were postysnaptic to visual input and therefore it was necessary to calculate a weight vector from that perspective. In this chapter, head direction cells are both pre- and post-synaptic, and thus it is possible to take the more intuitive measure of a presynaptic weight vector. Thus, the weight vector in this chapter tells us the overall direction to which the projecting weights point for a given

presynaptic cell. The center of mass of the weight vector x_j^w is calculated for each presynaptic cell j as

$$x_j^w(t) = \arctan\left(\frac{\sum_i w_{ij}\sin(x_i)}{\sum_i w_{ij}\cos(x_i)}\right) \tag{5.10}$$

where the sum is over presynaptic head direction cells *i*.

As with previous chapters, a corrected forumula (Batschelet, 1981; Fisher, 1993) was used to account for data circularity

$$x_j^w(t) = \begin{cases} \arctan\left(\frac{\sum_i w_{ij}\sin(x_i)}{\sum_i w_{ij}\cos(x_i)}\right) & \text{if } \sum_i w_{ij}\sin(x_i) > 0, \sum_i w_{ij}\cos(x_i) > 0; \\ \arctan\left(\frac{\sum_i w_{ij}\sin(x_i)}{\sum_i w_{ij}\cos(x_i)}\right) + 180 & \text{if } \sum_i w_{ij}\cos(x_i) < 0; \\ \arctan\left(\frac{\sum_i w_{ij}\sin(x_i)}{\sum_i w_{ij}\cos(x_i)}\right) + 360 & \text{if } \sum_i w_{ij}\sin(x_i) < 0, \sum_i w_{ij}\cos(x_i) > 0; \end{cases} \tag{5.11}$$

where w_{ij} is the strength of the synapse from presynaptic head direction cell j onto postynaptic head direction cell i and each postsynaptic cell has the preferred firing direction x_i.

The new offset value is then calculated using the final projecting weight vector x_j^w and the preferred

firing direction x_j for each presynaptic cell j

$$O_{new} = MIN(|x_j^w - x_j|, 360 - |x_j^w - x_j|) \tag{5.12}$$

which, similarly to initial weight profile and input activity packet calculations, accounts for the circular nature of the data.

Table 5.1: Network parameter values for simulations reported in Experiment 1. All simulations, except where explicitly stated, are run using these parameters. All measures of time are given in seconds.

Network Parameters	
Head Direction Cells	500
$C^{HD \to HD}$	500
ϕ_1	200.0
σ^{RC}	10.0°
τ^{HD}	0.001s
$\tilde{w}^{HD}$	0.005
δt	0.0001s
λ^{HD}	10.0
σ^{HD}	20.0°
V	$180.0^\circ/s$
Δt	0.01s

5.1.3.1 Weight profile effects

Figure 5.4 shows summary results for simulations run with parameters in Table 5.1. For these simulations, the model is wired with an offset profile with offset O, to which is added a non-offset component. The value of the non-offset component's relative strength λ^{NO} varies in the range [0,1], giving resultant weight profiles with varying degrees of offset.

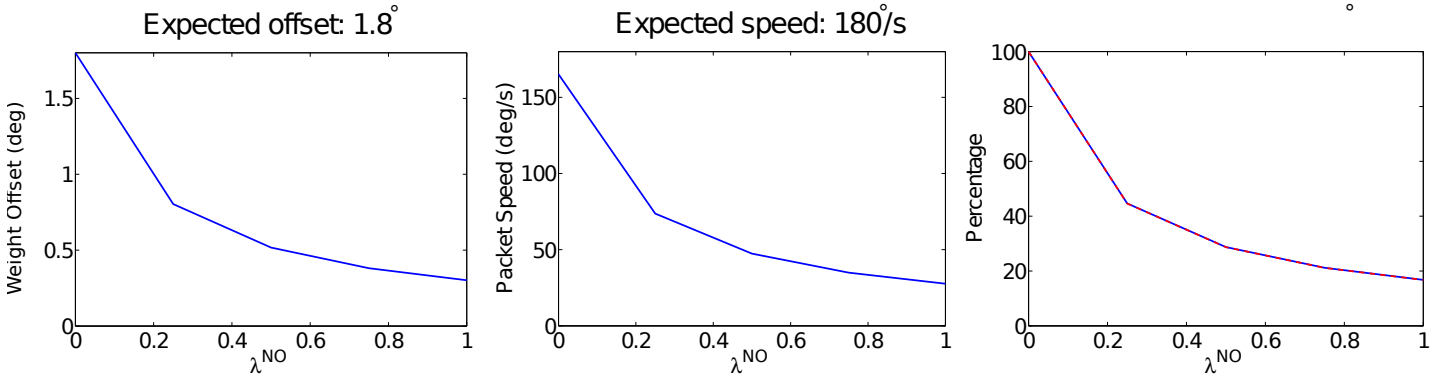


Figure 5.4: The effects of altering λ^{NO} on weight offset (left) and test phase packet speed (center). With just the offset weights, the expected weight offset is $1.8\,^{\circ}$ and the expected packet speed is $180\,^{\circ}/s$. To the right, we see both packet speed (solid blue) and weight offset (dashed red) as a percentage of their respective values at $\lambda^{NO} = 0$. These two lines directly correspond, showing that a reduction offset always causes a proportional reduction in speed.

To the left of Figure 5.4, we see the resulting offset with a variety of values of λ^{NO}. As expected, adding non-offset components of increasing relative strength to an initial offset has the effect of reducing the overall offset present in the weight profile. The middel plot shows the effect that this increase in λ^{NO} has on packet speed. As non-offset strength increases, there is a concurrent reduction in packet speed. Interestingly, this reduction in speed appears to be directly linked to the reduction in offset. This can be seen in the right hand plot of Figure 5.4, which shows the weight offset as a percentage of the expected value of $1.8\,^{\circ}$, overlaid with the packet speed as a percentage of the speed achieved at the correct offset value. These two plots directly correspond, highlighting the fact that speed is always reduced in proportion to the non-offset weight component.

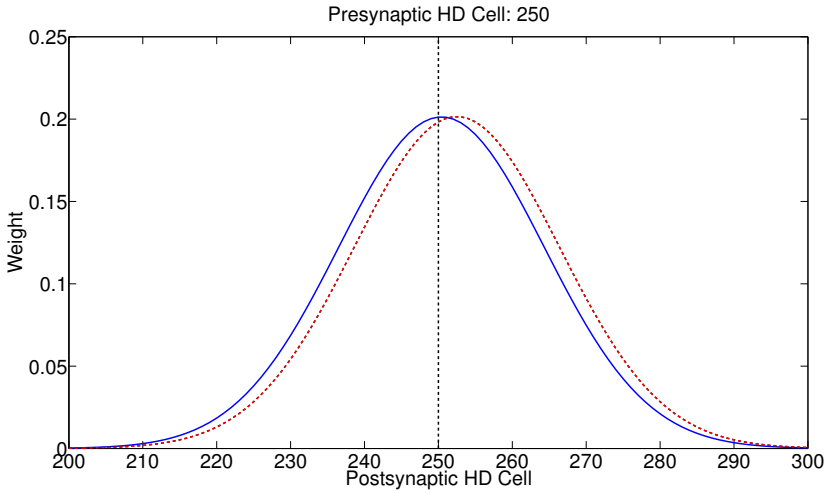


Figure 5.5: Weight profile for presynaptic HD cell 250 onto all postsynaptic cells before overlaying with non-offset component (red dashed) and after overlaying with non offset component with strength $\lambda^{NO} = 1$ (blue). Addition of non-offset component visibly reduces offset relative to the centerline of a perfectly symmetrical weight profile (black dashed line). For clarity, only part of the weight profile (postsynaptic cells 200-300) is plotted.

A reduction in offset can be seen in Figure 5.5, which shows for $\lambda^{NO} = 1$ both the desired offset profile (red dashed line) and the resultant additive profile (blue solid line) for a single presynaptic HD cell. This shows graphically the effect of adding a non-offset weight component which, in turn, results in a slower packet speed. These results demonstrate the negative effect non-offset weights can have for packet speed in a continuous attractor neural network. Such non-offset weights are found in almost all extant neural network models of the HD cell system, in the form of recurrent collateral connectivity acting to sustain HD cell layer activity. Therefore, these results suggest that alternative models not using such symmetric recurrent connectivity are crucial for achieving time-accurate path integration.

One more, very important, point arising from these data is that even when the offset is exactly at the correct value, we do not see a packet motion of exactly the correct speed. Indeed, top packet speed

in these simulations was $165.14\,^{\circ}/s$, which is only 91.8% of $180\,^{\circ}/s$. Certainly, it seems that there are factors other than weight profile shape affecting packet speed in these simulations. It is possible that this is, as hypothesised, a rise time effect which cannot be completly avoided in a rate-coded network.

5.1.3.2 Evidence of neuronal rise time effects

As mentioned above, one possible candidate for a source of path integration error, other than weight profile shape, is the rise time of neurons. One way to manipulate rise time, in order to determine whether it has an effect on packet speed, is via the τ^{HD} parameter. If rise time does affect packet speeds, we should see that increasing τ^{HD}, which will cause a concurrent increase in rise time, will yield lower packet speeds.

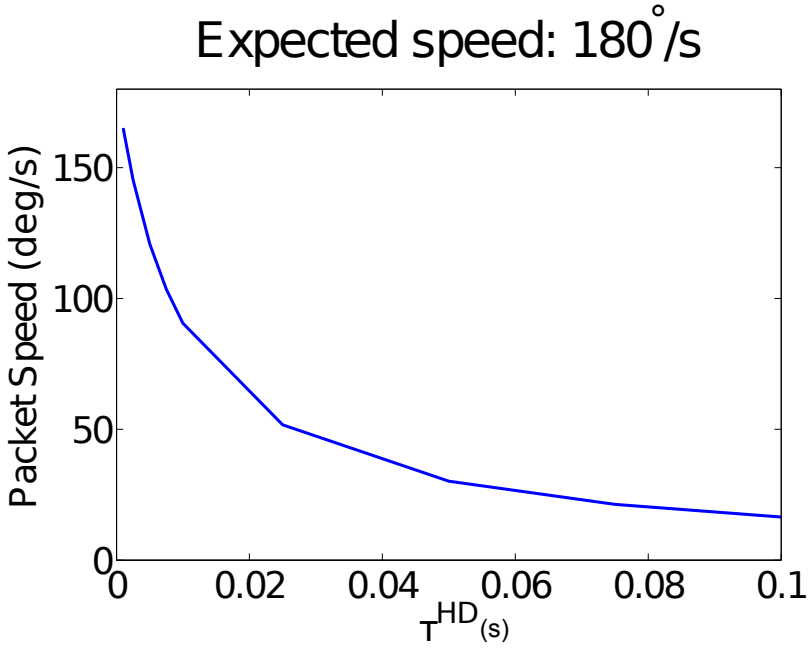


Figure 5.6: The effect of increasing values of τ^{HD} on test phase packet speed. As τ^{HD} increases, and with it neuronal rise time, there is a reduction in packet speed. As τ^{HD} tends to zero, the packet speed tends to the expected speed of $180\,^{\circ}/s$.

Figure 5.6 shows the packet speed for simulations with purely asymmetric weights and a range of values of τ^{HD}. All other parameter values are as listed in Table 5.1. Clearly, increasing rise time by

manipulating τ^{HD} has a negative effect on packet speed. This is because the activity packet is slowed by the increased time it takes for a neuron to start firing above threshold: by the time a neuron has become active, the activity packet will already be required to have moved to a new location. Thus the packet lags behind the expected speed, and this effect is increased with longer rise times.

However, the effect of rise times should be understood in the context of axonal conduction delays which are present in the brain. In particular, the error in packet speed depends on the relative magnitudes of the rise time and axonal conduction delay. The offset weights are set up to reflect learned associations between successive head directions over the axonal delay period Δt. However, during testing, postsynaptic neurons respond to inputs from presynaptic neurons only after the axonal delay period plus the additional rise time, t_R. In this case, the accuracy of packet speed is proportional to the ratio of the conduction delay to the sum of the conduction delay plus the rise time. That is,

$$\text{Actual packet speed} = \frac{\Delta t}{\Delta t + t_R} \times \text{Expected packet speed.}$$

The coefficient $\Delta t/(\Delta t + t_R)$ is maximised for either large Δt or small t_R. In particular, the actual packet speed will converge to the expected packet speed either as Δt increases or as t_R tends to zero.

So, a key factor in this process is how large the rise time t_R is relative to the conduction delay that already exists. For example, a constant inaccuracy due to rise time of 1ms (note: this figure is quoted purely as an example) is less severe in the context of a conduction delay of 100ms than it is with a conduction delay of 10ms.

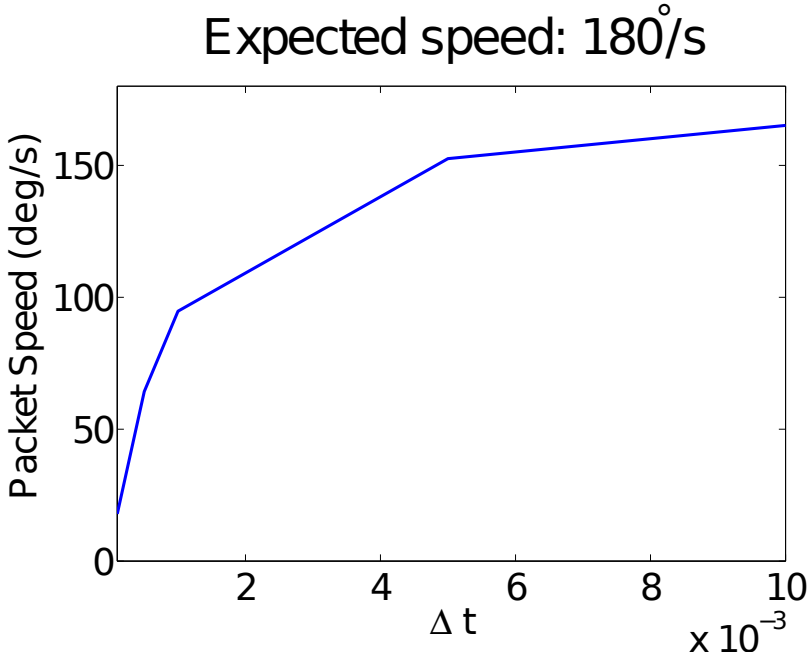


Figure 5.7: The effect of reducing values of the axonal conduction delay Δt on test phase packet speed with a constant value of $\tau^{HD} = 0.001$. As Δt gets smaller relative to rise time, packet speed is reduced. As Δt increases, the packet speed tends to the expected speed of $180\,^{\circ}/s$.

Figure 5.7 shows final packet speed for simulations with purely offset weights, run with a variety of axonal conduction delays Δt, in the range [0.0001,0.01]. This verifies experimentally the intuition that longer conduction delays will be associated with less error, whilst short conduction delays will give more error. This is reflected in much slower packet speeds with shorter conduction delays than with longer ones.

5.1.3.3 Varying target speeds: Clarifying rise time effects

As already seen, changing conduction delay size affects the top packet speed the network is able to produce. However, in the above simulations the conduction delay and weight offset were varied together to achieve a fixed expected speed of $180\,^{\circ}/s$. Therefore, it was not clear whether the observed changes in packet speed were a direct effect of the delay itself, or the weight offset size as determined by conduction delay. Since the weight offset is calculated as based on both the conduction delay and

target speed, this question can be answered by holding delay constant and manipulating target speed.

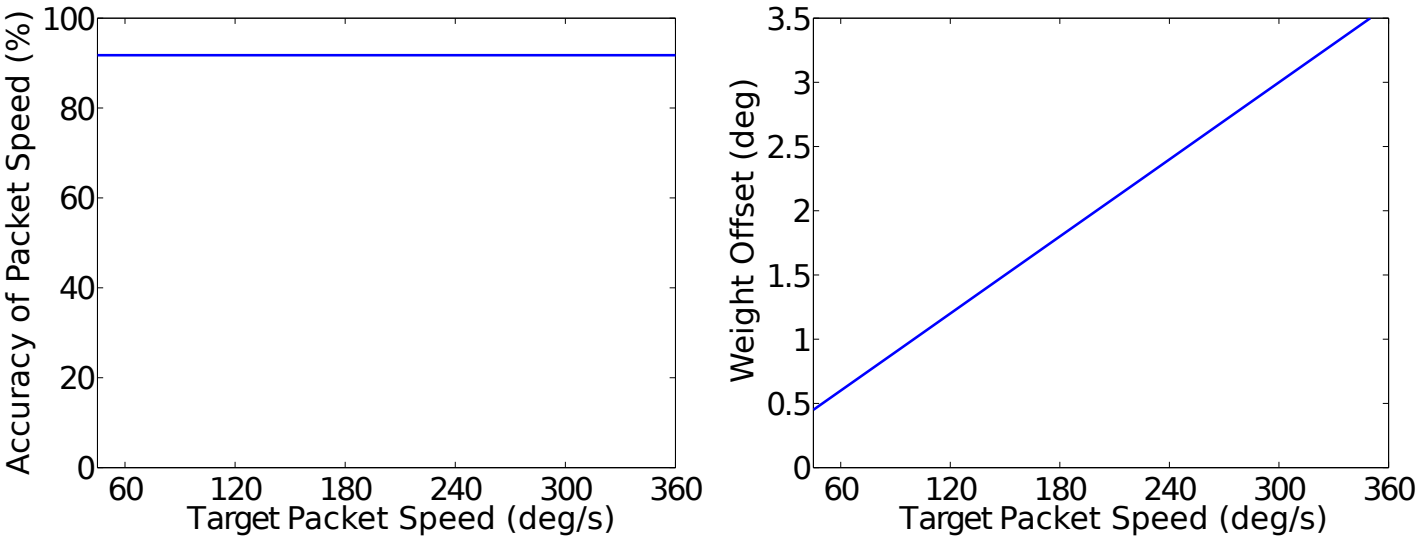


Figure 5.8: Simulation results in which the conduction delays are kept fixed at $0.01s$ while the weight offset is varied to give different target packet speeds. The accuracy of the actual packet speed remains constant at about 91.8%, regardless of the target packet speed and hence weight offset (left). Also plotted is how the weight offset variues linearly with the target packet speed (right). These results demonstrate that, with a fixed conduction delay, the accuracy of the actual packet speed does not depend on the weight offset itself.

Figure 5.8 shows summary graphs for simulations run with purely offset weights, for a variety of target velocities in the range $[45, 360\,^\circ/s]$. On the left, we see the actual packet velocity, as a percentage of the target velocity, and on the right we see the accompanying weight offset. Curiously, we see on the left that, regardless of target velocity, the final packet speed is always a certain percentage of that target (specifically, 91.8%). This is despite a weight offset which is of varying sizes, as can be seen on the right. Together with Figure 5.7, these results demonstrate that the relationship of the rise time with conduction delay size, rather than offset size, is critical.

5.1.4 Estimating rise time

Given the hypothesised effects described in the introduction and supported by the data thus far, it should be possible to estimate the effective rise time for a given set of parameters. In the following

simulations the weight offset is fixed at a constant value corresponding to a particular combination of conduction delay Δt and target velocity. If the rise time is effectively, for a given weight offset, the amount by which signal transmission is too long, it is possible to keep the weight offset constant and reduce Δt until there is no undershooting of target speed. At this point, the effective rise time will be equal to the difference between the Δt value used and the original true value of Δt.

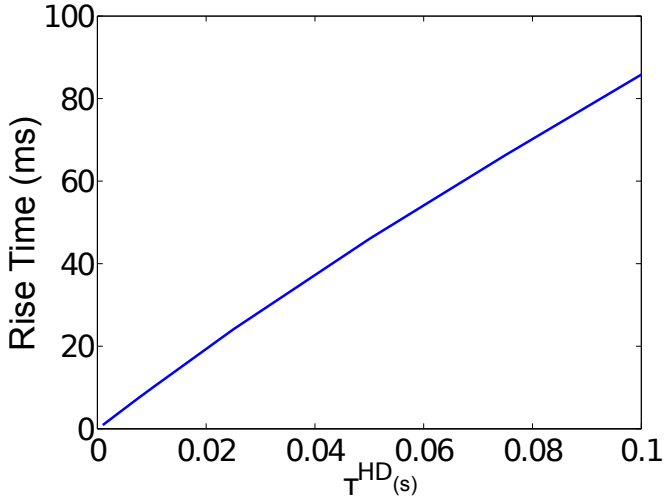


Figure 5.9: Estimates of rise time for a range of τ^{HD} values. This demonstrates that rise time is an observable effect, varying with τ^{HD} and affecting network performance via an interaction with Δt.

Figure 5.9 shows, for a range of values of τ^{HD}, the estimated rise times obtained by using this inferential method. Parameters are as given in Table 5.1. It was found that the effective rise time was approximately linearly related to τ^{HD}. If we shorten Δt by this amount, we see that it is possible to perfectly replicate the target velocity, demonstrating that previous results in which target velocity has been undershot were due to this rise time effect.

5.2 Experiment 2: Self-organising network

Thus far, we have seen that the accuracy of path integration speed is affected by two major factors: the presence of a symmetric non-offet component to the recurrent weight layer, and neuronal rise time (specifically, the relationship of rise time to axonal conduction delay). It is of value then, to continue to investigate these factors in a self-organising network. This is vital in order to see the effects of these factors when weights are self-organised, and to determine whether there are any other important factors affecting packet speed in a self-organising network as opposed to a pre-wired network.

5.2.1 Model description

The model used for Experiment 2 is very similar to that used for Experiment 1. The sole difference is that recurrent collateral weights are not pre-wired to a specific profile, but instead self organise during a training regime using a local Hebbian learning rule. Recurrent collateral weights are initialised during training to a uniform, low, value. The model is then trained with external visual input $e_i(t)$ at a particular rotation speed, until weights have converged on a final profile. Then model is then tested in the absence of the visual input $e_i(t)$as if it were a pre-wired network as in Experiment 1. As in the previous model, each recurrent collateral connection contains a conduction delay Δt. During training, the effect of this delay is to associate postynaptic HD cells with presynaptic HD cells active Δt in the past, thus generating an offset weight profile.

5.2.2 Model operation

This section provides details of the model governing equations and simulation protocols. As this model is very similar to that used for Experiment 1, only those aspects which differ are discussed.

5.2.2.1 Model equations

In order to maintain a similar level of network activation during training and testing phases, feed-forward inhibition is used. This takes the form of a constant I^{FF}, which is subtracted from HD cell activation during the training phase only, such that the new activation equation is given as

$$\begin{aligned}\tau^{HD}\frac{dh_i^{HD}(t)}{dt} = &- h_i^{HD}(t) + e_i(t) \\ &- \frac{1}{N^{HD}}\sum_j \tilde{w}^{HD} r_j^{HD}(t) \\ &+ \frac{\phi_1}{C^{HD\to HD}}\sum_j w_{ij}^1(t) r_j^{HD}(t-\Delta t) - I^{FF}\end{aligned} \tag{5.13}$$

This both prevents a strong visual training input $e_i(t)$ flooding the network with activation, and allows the network to sustain activation through a recurrent weight profile alone in the absence of strong feed-forward inhibition. Otherwise, all governing equations are as detailed for Experiment 1, with the addition of weight updates every timestep. Synaptic weights are updated every timestep according to a local associative Hebbian learning rule

$$\frac{dw_{ij}^1(t)}{dt} = k r_i^{HD}(t) r_j^{HD}(t-\Delta t) \tag{5.14}$$

where w_{ij}^1 is the synaptic weight from presynaptic HD cell j with firing rate r_j^{HD} onto postysnaptic

HD cell i with firing rate r_i^{HD}, and k is the learning rate constant, which determines speed of weight change. Unlike Experiment 1, weight update changes are not followed by normalisation.

5.2.2.2 Simulation protocol

The model is initialised with all cell firing rates and activations at zero. Recurrent collateral weights are initialsed to a uniform negligibly small value of $w_{ij} = 0.0001$. Simulation is split into two phases: training and testing. During the training phase, a Gaussian packet of external activation is applied to the head direction cell layer through the input $e_i(t)$. This packet of activation moves through the ring clockwise at a constant velocity of V°/s. During this phase, cell activations and firing rates are updated, and recurrent collateral weights are updated and normalised as according to the relevant equations. After this training phase, the external input $e_i(t)$ is removed. The test phase then begins, with the network continuing to update in the absence of any external input. The recurrent collateral weight profile is fixed during the test phase, and thus no further changes in synaptic strengths occur. In some simulations, a symmetric non-offset weight profile is overlaid onto the existing learnt weight profile at the beginning of the test phase. This is done identically to the overlaying of non-offset and offset weight profiles when initialising synaptic weights in Experiment 1.

Table 5.2: Network parameter values for simulations reported in Experiment 2. All simulations, except where explicitly stated, are run using these parameters. All measures of time are given in seconds.

Network Parameters	
Head Direction Cells	500
Training Time	$298.5s$
Testing Time	$2.0s$
$C^{HD \to HD}$	500
ϕ_1	60.0
τ^{HD}	$0.001s$
$\tilde{w}^{HD}$	0.01
I^{FF}	50.0
δt	$0.0001s$
λ^{HD}	70.0
σ^{HD}	$30.0°/s$
V	$180.0°/s$
Δt	$0.01s$
k	0.01

5.2.3 Results

This section details simulation results from running with the new self organising model and simulation protocol described above. Methods of analysis used are the same as detailed in the results section for Experiment 1. All parameters are as in Table 5.2, unless stated otherwise.

5.2.3.1 Weight profile effects

Figure 5.10 shows, similarly to Figure 5.4, results from simulations run with an overlaid symmetric non-offset component at a variety of values of λ^{NO} in the range [0,1]. To the left, we see the resulting weight offset with various values of λ^{NO}. As in the previous simulation, we see that adding

non-offset components acts to reduce the weight offset size present, with this reduction being dependent on λ^{NO}. In the center plot, we see the resulting speed given these varying values of λ^{NO} (and concurrently varying sizes of weight offset). Again, a result from the pre-wired network is replicated in that the stronger the non-offset component, the smaller the weight offset and therefore the slower a packet is able to move through the network. In the right hand graph of Figure 5.10, we see plotted both packet speed as a percentage of the packet speed at $\lambda^{NO} = 0$ (blue solid line) and weight offset as a percentage of the weight offset at $\lambda^{NO} = 0$ (red dashed line). As in Experiment 1, these two proportions have an extremely close correspondence across a range of values of λ^{NO}, indicating that reductions in packet speed are caused by weight offset reductions. These weight offset reductions are themselves directly caused by changes in the strength of the non-offset recurrent weights via the λ^{NO} parameter.

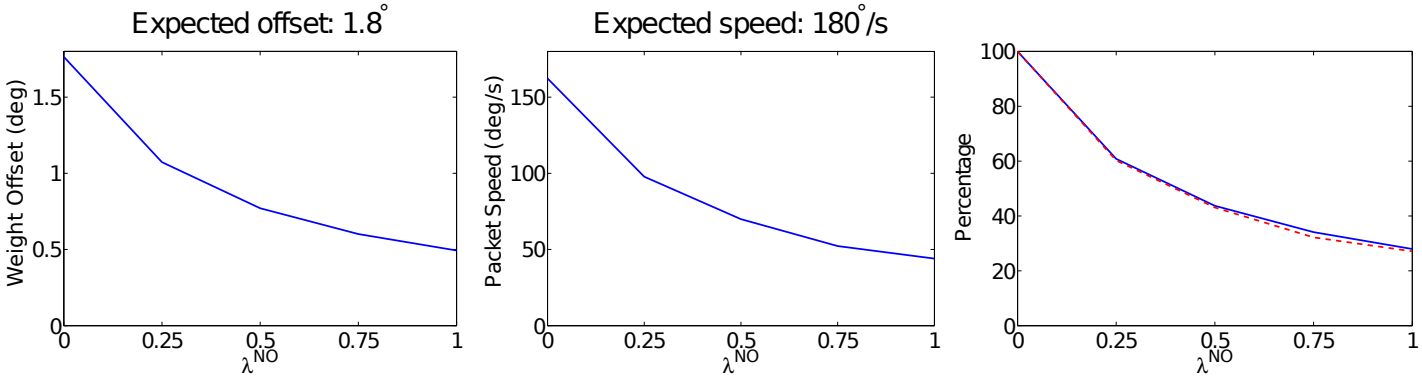


Figure 5.10: Effect of varying λ^{NO} on weight offset (left) and test phase packet speed (center). As λ^{NO} is decreaed to zero, the speed at testing apprproaches the target speed of $180°/s$ used during training. Increasing the λ^{NO} both reduces weight offset and packet speed. Similarly to Figure 5.4's results from the pre-wired network, these two reductions are shown to correspond directly when represented as a percentage of their value at $\lambda^{NO} = 0$ (right, conventions as in Figure 5.4).

However, there are some minor differences, not elucidated in the figures above, between the results given here and those seen in the pre-wired network. Even with $\lambda^{NO} = 0$, we see that the top speed for

this self-organising network ($162.26°/s$) is slightly less than that achieved in the pre-wired counterpart ($165.14°/s$). This difference is relatively small: as a proportion of the target speed of $180°/s$, these two speeds represent 90.1% and 91.7% respectively. It is likely attributable to the fact that the self-organised weight profile is not as coherent as the pre-wired one. The weight offset value given here is not a perfect $1.8°$, but rather $1.76°$. This statistic is calculated as an average of the weight offsets for all receiving HD cells in the layer. Their distribution has a standard deviation of $0.017°$, showing that whilst it is fairly uniform, it is not perfectly homogenous as in the pre-wired case. Naturally, this produces a packet which moves during the testing phase at a speed with some minor variations, and the speed statistic here is calcuated as an average across this stage. This represents another possible explanation for this difference. Thus, minor differences between self-organised and pre-wired results can be expected.

In the interests of full disclosure, it should be noted that the simulation at $\lambda^{NO} = 1$ was run with a reduced ϕ_1 parameter during the testing phase ($\phi_1 = 30$). This is because the addition of a strong non-offset component to the learned offset weight profile resulted in an overly strong recurrent weight profile and uncontrolled spread of activation through the HD layer. This adjustment was considered acceptable, considering exploratory results (not pictured) demonstrating that the ϕ_1 parameter does not have a significant effect on packet speed.

5.2.3.2 Evidence of rise time effects

As with the pre-wired network, the effect of τ^{HD} was investigated. Figure 5.11 shows summary results from simulations run with a variety of values for τ^{HD} in the range [0.001, 0.1]. In these simulations, λ^{NO} was set to zero. To the left we see the resulting packet speed during testing at each

of these τ^{HD} values, and to the right we see the accompanying self-organised weight offset. First, the rise time does not have any noticeable effect on the self-organisation of the offset weight profile. This is reflected by the constant weight offset for all values of τ^{HD} shown on the right. Secondly, the graph to the left, when compared to Figure 5.6, shows that the effect of varying the neuronal time constant on packet speed is the same in both the pre-wired and self-organising networks. In particular, as τ^{HD} is reduced to give smaller rise times, the speed of testing approaches the target speed of $180°/s$ used during training.

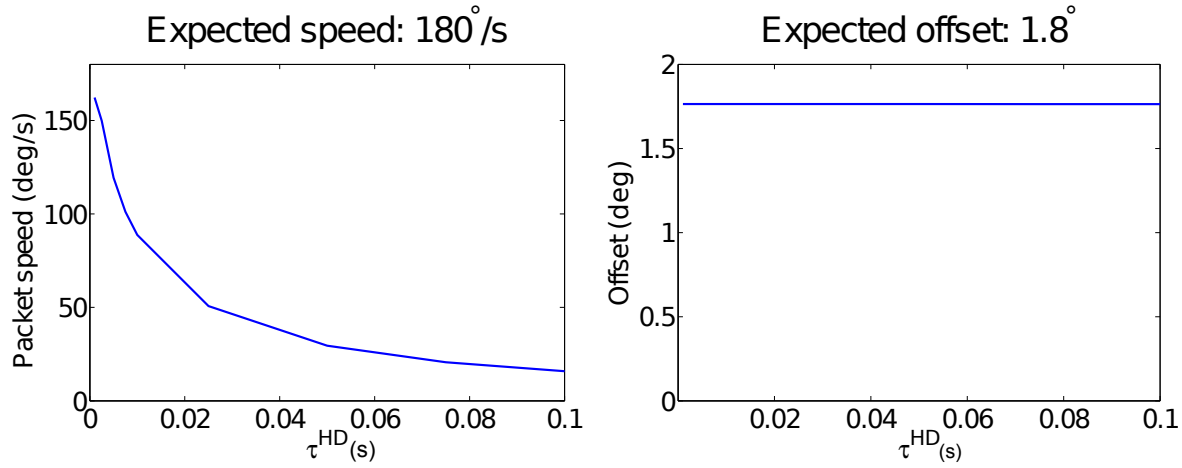


Figure 5.11: The effect of increasing values of τ^{HD} on test phase packet speed (left) and self-organised offset (right). Identically to the pre-wired network, we see that increasing τ^{HD} results in slower packet speed. However, offset does not seem to be affected by τ^{HD}, indicating that rise time does not have an effect during learning.

Interestingly, we see that the weight offset for all simulations is not the desired value of $1.8°$, but rather is a systematic undershooting in this self-organising model ($1.76°$). The weight profile for presynaptic HD cell 250 after simulation with $\tau^{HD} = 0.001$ was found to have skewness of 1.88 ($SE = 0.11$), showing that that the weight profile was positively skewed (towards the trailing edge). This skewness figure was similar for all HD cells examined (results not given). These self-organising results thus give a naturally skewed weight profiles, which are not guaranteed to produce exactly the correct speed

during testing.

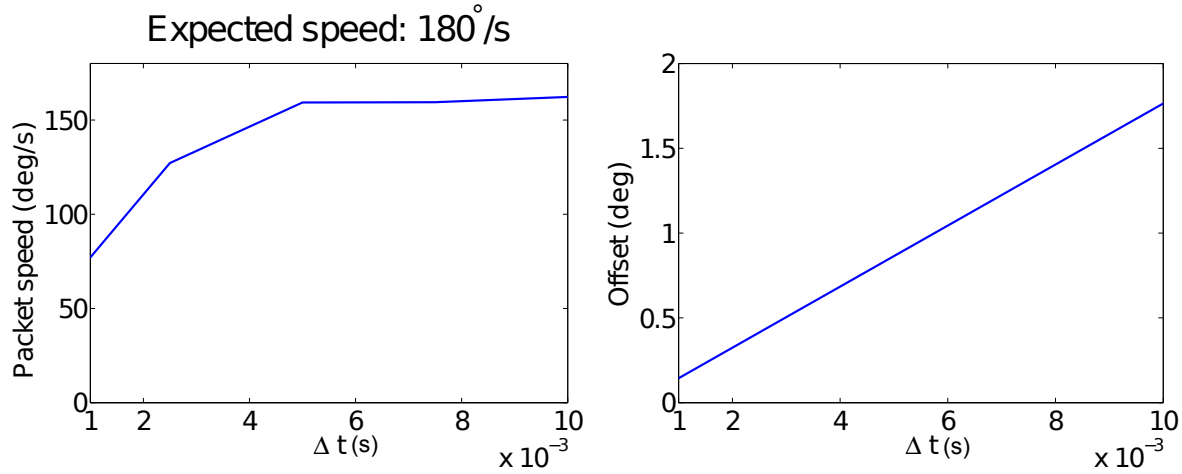


Figure 5.12: The effect of reducing the axonal conduction delay Δt on test phase packet speed (left) and weight offset (right) with a constant value of $\tau^{HD} = 0.001$. Similarly to Figure 5.7 of Experiment 1, smaller Δt relative to rise time gives a reduced packet speed. This seems to be a testing effect rather than training, as offset is still being learnt correctly based on values of Δt (right)

Similarly to Experiment 1, more evidence of the negative effect of long rise times comes from manipulating the axonal conduction delay. Summary results from simulations run with different values of Δt in the range [0.001,0.01s] are presented in Figure 5.12. To the left we see packet speed during testing, whilst to the right we see the weight offset generated by various values of Δt during training. Identically to results from Experiment 1, we see that reducing Δt, for a constant value of $\tau^{HD} = 0.001$, causes a drop in packet speed. This is because the constant inaccuracy resulting from rise time is more severe in the context of the shorter conduction delays at smaller values of Δt. This reduction in speed does not seem to be due to any adverse effects of the Δt parameter on training: to the right we see that the weight offset reduces linearly with reducing values of Δt, exactly as expected.

5.2.3.3 Further evidence of rise time effects: Varying target speed

As with Experiment 1, we have seen that packet speed is negatively influenced by smaller axonal conduction delays. However, we must once more ascertain whether this is the relationship between

τ^{HD} and Δt or simply τ^{HD} and absolute weight offset size. If, as with the pre-wired network, it is the relationship between τ^{HD} and Δt that is critical, then for a given combination of τ^{HD} and Δt the test phase packet speed should always be the same constant proportion of a variety of target velocities on which the model is trained.

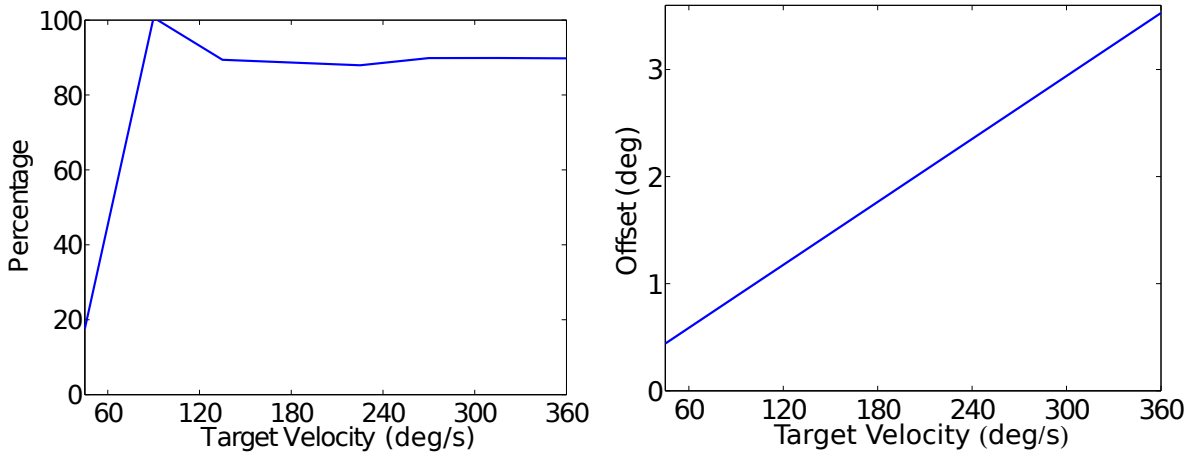


Figure 5.13: The effect of varying target velocity on test phase speed as a percentage of target (left) and on weight offset (right). Packet speed as a percentage of the target is not much influenced by target velocity for $V > 135°/s$. However, some irregularities exist, which is likely to be due to a secondary effect of target velocity during training.

Figure 5.13 shows summary results from simulations run with varying target speeds, V during training, $\tau^{HD} = 0.001s$ and $\Delta t = 0.01s$. To the right, we see that the weight offset is, as expected, being correctly learnt at different sizes for the different target velocities. To the left we see the test phase packet speed as a percentage of target velocity. For higher target velocities, we see that our hypothesis is confirmed: there is little effect of target velocity, thus of absolute offset size, on proportional test phase packet speed.

However, below $V = 135°/s$ we see that this relationship breaks down. It is possible that, at slower training speeds, the weight profile has not yet fully self-organised. An immature weight profile might well have a correct mean offset value (Figure 5.13, right) but the distribution of individual weight

offsets for all HD cells in the layer could show variation. This would be due to the fact that slower packets have completed fewer complete rotations through the HD layer. This is supported by Figure 5.14, which shows weight offsets after training for the entire HD layer, at three different speeds. Clearly we see a relationship between rotation speed and the uniformity of HD layer offsets, with slower rotation speeds resulting in visible distortions of the distribution. Further confirming training effects at slower target rotations, Figure 5.15 shows the effects of doubling the amount of training to $597s$ for a target rotation velocity of $90°/s$. Doubling training time has reduced the deformity in the HD layer weight offset distribution.

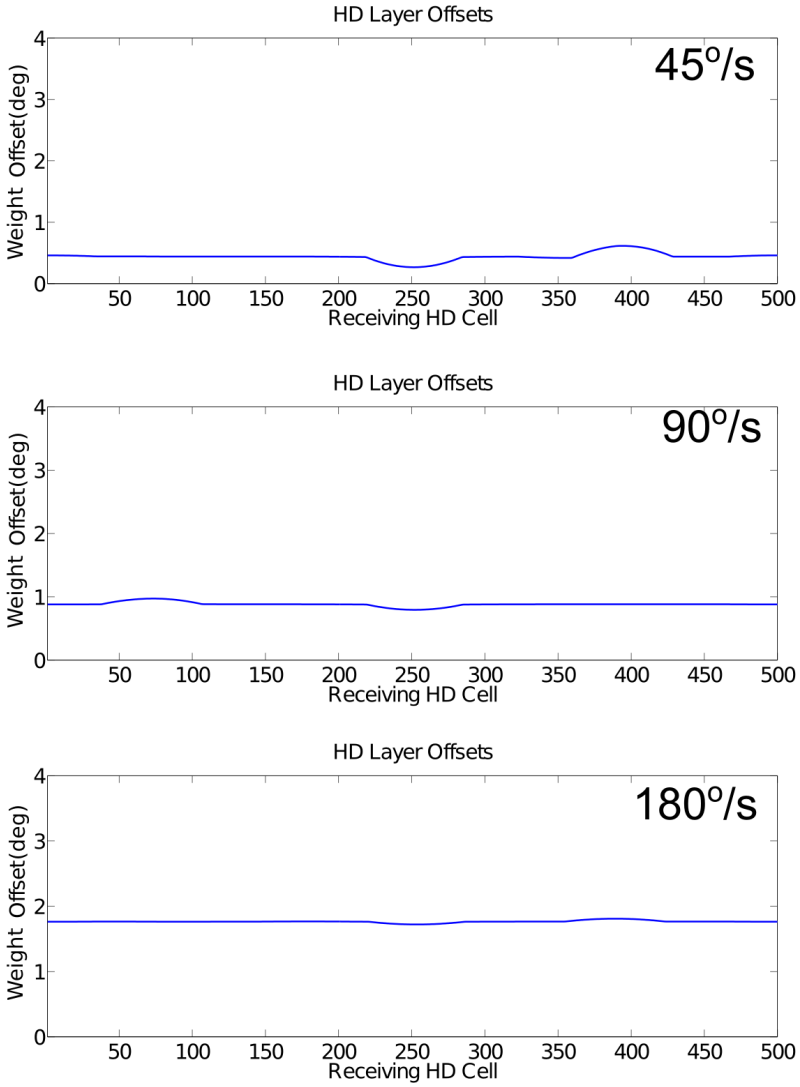


Figure 5.14: Offset in receiving weights for all HD layer cells after training at target velocities of $45°/s$ (top), $90°/s$ (middle) and $180°/s$ (bottom). The distribution of weight offsets becomes more uniform as target velocity during training increases. At lower target velocities, the activity packet has moved around the HD layer fewer times during training, resulting in a weight profile which has not trained out all inconsistencies. In contrast, higher rotation velocities result in more full rotations around the HD layer during training and thus a more trained, more uniform distribution of offsets.

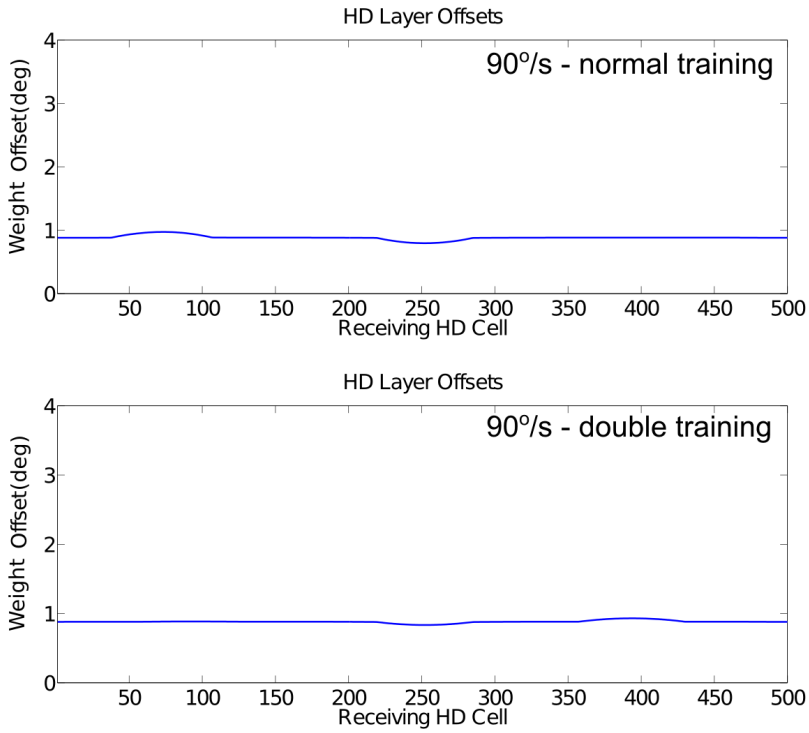


Figure 5.15: HD layer receiving offsets showing the effect of training at $90°/s$ for a normal training period $298.5s$ (top) and for double length training $597s$ (bottom). Longer training has ameliorated the effect of training at a lower target velocity, with the distribution of HD layer weight offsets becoming much more uniform.

5.3 Discussion

This chapter presents details of the architecture and operation of a simple one-layer neural network model, designed to investigate potential computational factors contributing to inaccuracy in path integration mechanisms. Two key hypothesised sources of error are identified through simulation. Firstly, the antagonism between symmetrical (non-offset) recurrent collateral weights and asymmetrical (offset) weights. Secondly, neuronal rise time. Both of these factors are investigated in two variants of the model: one with pre-wired synaptic weights and one which undergoes training and accompanying self-organisation of the weight profile. This represents a simple, yet powerful, approach to uncovering key factors affecting accuracy of path integration speed.

The models utilised in this chapter consist of a ring of HD cells, connected back onto itself with recurrent collateral synapses which contain a conduction delay Δt. The profile of recurrent collateral weights can be offset, or a combination of offset and non-offset compnents. The relative strengths of these components is varied by altering the strength of the non-offset component. These two weight components represent the two types of weights that are typically required for path integration in the vast majority of extant HD models. The non-offset component represents symmetrical recurrent collateral weights, as seen in continuous attractor models of the HD cell system, which act to sustain HD layer firing in the absence of any external input. The offset component represents external input to the HD cell layer, which acts to update the HD cell representation in response to changes in head direction. This external input to the HD cell layer is typically the output of some path integration system which shifts the HD cell activity packet to a new location within an attractor network. The crucial aspect of path integration addressed, then, in this chapter is what are the factors that prevent it

from occurring at the correct speed.

In the first experiment, the model's connectivity is pre-wired, but in the second experiment, offset connectivity is learnt through training with an activity packet moving at a constant target rotation velocity V°/s. An offset is learnt due to the influence of Δt, which causes a postysnaptic activity packet to be associated with a presynaptic activity packet Δt in the past. This past location will, given constant packet motion, be an offset of a constant amount determined by Δt and V. Once this offset weight profile has correctly self-organised, it can be overlaid with a symmetric non-offset component identically to the pre-wired model.

Parameter exploration helps to explain the key factors governing packet speed during testing in both varieties of model. In both the pre-wired and self-organising varities of the model, we see that packet speed is reduced in direct proportion to the strength of the non-offset weight component. This non-offset component represents symmetrical recurrent collateral connectivity, which is a key aspect of most continuous attractor neural networks (CANNs). Following the attractor hypothesis of Skaggs and colleagues (Skaggs et al., 1995), the majority of extant head direction cell system models are based on CANN architectures and therefore employ such symmetrical connectivity.

Head direction cells have previously been considered an ideal system to model using CANN approaches (Touretzky, 2005). The ability of head direction cells to sustain firing in the absence of visual input (Taube et al., 1990a) and, provided the rat continues to face in the cell's preferred direction, to fire persistently without adaptation (Taube, 1998) is reminiscent of attractor dynamics as

applied to short term memory functions in various parts of the brain e.g., dorsolateral prefrontal cortex (Goldman-Rakic et al., 1996), parietal cortex (Duhamel et al., 1992), hippocampal CA1 (Wills et al., 2005) and lateral interparietal cortex (Albantakis and Deco, 2011). In addition, CANN-based models have been used to explain a large range of HD cell phenomena beyond those core properties listed above (see Knierim and Zhang 2012).

The fact that symmetrical connections have been shown in this model system to impede path integration speed suggests a major shift for HD modelling approaches away from the current form of CANN architectures. Certainly, there has been a conspicuous lack of any hard evidence beyond the inferential observations given above for recurrent collateral connectivity in HD cell regions such as lateral mammillary nucleus (Allen and Hopkins, 1989) and anterodorsal thalamic nuclei (Rubin et al., 2001).

If the simulation results in this chapter suggest moving away from CANN models what, then, are the alternatives? There have already been forays into HD cell model architectures which do not rely on symmetric recurrent collateral connectivity within a layer but can both sustain an activity packet and perform path integration (Boucheny et al., 2005; Rubin et al., 2001; Song and Wang, 2005). However, such models are based on inhibition shaping some kind of tonic input whose source remains unidentified. A potential candidate for this tonic input source lies with the discovery of intrinsic cellular mechanisms allowing graded persistent firing. Using slice preparations of rat spinal cord, Morisset and Nagy (2000) demonstrated that afferent fibre stimulation caused dorsal horn neurons to fire persistently at a tonic rate, which was increased by further stimulation and decreased by inhibition.

Further investigations suggested that an intrinsic mechanism, involving Ca^{2+} dependent non-specific cationic currents, is responsible for this behaviour (Egorov et al., 2002, 2006; Fransen et al., 2006; Winograd et al., 2008). Crucially, this behaviour has been observed in head direction cell areas, such as the postsubiculum (Yoshida and Hasselmo, 2009) and the anterodorsal thalamic nucleus (Kulkarni et al., 2011), and is thought to be dependent on similar intrinsic cellular mechanisms. Additonally, this persistent firing has been seen in other closely related regions involved in spatial cognition, such as entorhinal cortex (Yoshida et al., 2008; Tahvildari et al., 2007). Although it is therefore conceivable that sustained activity could be driven in head direction cell areas in this manner, persistent firing of this type is too low in frequency (5-10Hz as opposed to 100Hz) to be solely responsible. Despite this, persistent firing remains an interesting prospect, and may act to stabilise activity within head direction cell networks.

Another potentially more plausible approach lies with the use of angular head velocity cells to provide an input to the HD cell system. A possible architecture could be based on neurophysiological evidence demonstrating reciprocal connectivity between excitatory lateral mamilliary nucleus (LMN) cells and inhibitory dorsal tegmental nucleus (DTN) cells (Allen and Hopkins, 1988, 1989); with this LMN-DTN complex being a likely candidate substrate for the generation of the head direction signal (Song and Wang, 2005; Taube and Bassett, 2003; Taube, 2007).

The proposed architecture would have a layer of LMN HD cells which are reciprocally connected to a DTN layer of combination (COMB) cells, in the manner of Walters and Stringer (2010) and Walters et al. (2013), which fire to a combination of head direction and rotation. This COMB layer

would receive input signalling head rotation from angular head velocity (AHV) cells: it is known that the DTN receives input conveying an AHV signal from the nucleus prepositus hypoglossi (Khalsa et al., 1987), which itself receives projections from vestibular nuclei (R.A. and R., 1985). Single-cell recordings confirm that DTN contains both AHV cells and cells responding to combinations of head direction and angular head velocity (Bassett and Taube, 2001; Sharp et al., 2001). However, it should be noted that the source of this AHV signal may not be tied specifically to the DTN; AHV cells have also been found in LMN (Stackman and Taube, 1998).

A crucial property of AHV cells is that they have some firing even when the head is not rotating ((Bassett and Taube, 2001). This would allow the some cells in the COMB layer to associate a head direction with no rotation. Such COMB cells would represent a particular head direction during no rotation. These COMB cells could then form reciprocal connections with HD cells representing the same head direction. These linked cells could then stably support themselves as an attractor state during periods of no rotation. During no rotation, head direction cells would be driven by corresponding COMB layer cells which respond to a lack of rotation, allowing these cells to sustain firing without symmetric recurrent connectivity within the HD layer or intrinsic cellular mechanisms. During path integration these no-rotation COMB cells would become quiescent, whilst COMB cells responding the head direction and rotation would begin firing and act to shift the HD cell activity packet. This architecture would allow sustained head direction cell firing without hindering the motion of an activity packet during path integration. The idea of reliance on ascending vestibular information to generate a sustained HD representation is potentially a more plausible scenario than either recurrent connectivity within the HD layer or intrinsic persistent firing, given that damage to the vestibular system eliminates

the head direction cell signal (Stackman and Taube, 1997; Stackman and Herbert, 2002; Bassett and Taube, 2005; Valerio and Taube, 2012).

It is therefore hypothesised that individual layers of the head direction system (e.g., in rats) do not contain recurrent connections to help stabilise activity packets representing current head direction, because these connections would slow down path integration as soon as the animal begins to rotate. The fundamental problem is that the recurrent connections within a layer could only learn one specific rotation speed, whether this is either no rotation or some fixed speed of rotation. If the animal rotates at any other speed, these recurrent connections will introduce error into the path integration. However, consider the two-layer architecture of Walters et al. (2013) with plastic bi-directional connections between the HD and COMB layers, as shown in Figure 5.1. Within the COMB layer, neurons learn to respond to specific combinations of head direction and rotation speed. At any moment, only the correct subset of COMB cells are activated corresponding to the current rotation speed. This means that the bi-directional connections between the two layers learn to encode specific rotation speeds, with different bi-directional connections encoding different speeds. Consequently, the bi-directional connections between the two layers can implement accurate path integration across all of the trained speeds, including stabilising the HD activity packet during no rotation.

Another key finding of this chapter is that, for a given weight offset, long neuronal rise times have a detrimental effect on activity packet speed during path integration. This is true in both pre-wired and self-organising networks. Rise time appears to reduce packet speed in relation to the axonal conduction delay used: the shorter the conduction delay for a given rise time, the more severe the slowing

effect. This is in contrast to the lack of functional interaction of rise time and absolute weight offset size, demonstrated by the fact that the test phase packet speed undershot by a constant proportion a range of target velocities. Another important related finding is that rise time does not seem to have an effect on the self-organisation of offset connectivity, as demonstrated in Experiment 2. These findings support the intuition elaborated in this chapter's introduction that rise time causes signal transmission between HD cells to be too long during testing, without playing a significant role during learning.

Whilst rise time is a particular issue for a rate-coded neural network such as the one utilised in this chapter, it may be ameliorated somewhat by the greater degree of biophysical and temporal accuracy of a spiking neural network. Spike timing is well established as a means of neuronal information encoding (Maas and Bishop, 2001; Izhikevich, 2006) and has been suggested to convey spatial information in the hippocampus (O'Keefe and Recce, 1993), being important for the development of spatial maps (Blum and Abbott, 1996; Tsodkys et al., 1996). Incorporating neuronal spiking with spike-timing dependent plasticity (Bi and Poo, 1998) may well help with model accuracy at the level of individual neurons.

Spiking networks are known to be able to update their representations very rapidly (Brunel and Wang, 2001), and there has been in the last few decades a desire to overcome the speed constraints of rate-coded systems in order to match speeds of sensory processing in the brain (Thorpe et al., 1996). This has primarily been due to observations of very rapid visual processing which suggests the existence of a feed-forward mechanisms which operates faster than is accounted for in rate-coded models (Antal et al., 2000). Similar evidence of fast processing is available in the auditory modality (Jen et al.,

1989). There are a number of strategies for rapid processing within spiking networks, such as Rank Order Coding (Thorpe et al., 2001).

Another solution to be applied to this particular problem may be that neuronal state-switching can occur rapidly thanks to background levels of neural noise. Essentially, neurons may have a background level of activation that is very close to firing threshold (Fourcaud-Trocme et al., 2003), thus reducing the time taken for input to drive them up to firing. This has been demonstrated in neural network models (Tsodkys and Sejnowski, 1995; Gerstner, 2000; Brunel and Wang, 2001). It could thus be the case that a spiking network will allow rise time to be minimised to an extent that no longer interferes with path integration.

The models used in this chapter are simplifications designed to illustrate some general, yet vital, principles of speed limitations in path integration. It is therefore important to be explicit about the simplifications employed. Firstly, HD cells in these networks are explicitly pre-designated with a preferred direction. In reality, the preferred direction of individual HD cells must be calibrated in some way to establish the specific connectivity, both between HD cells, and from visual areas onto HD areas. This chapter however, has not been focused on the development of HD cell responses, and is not focused on answering how HD cells accquire a preferred direction.

On the topic of connectivity, these models employ a single layer of recurrent connections with varying degrees of asymmetry. The degree of asymmetry is controlled by manipulating the relative strength of two components, one offset and one non-offset component. The final weight profile represents the ad-

ditive combination of these two components. This represents a simplification of two sources of input: non-offset synapses from the HD layer back onto itself, and offset output of some path integration system back onto itself. However, reducing these two sources to one layer is considered a necessary simplification to demonstrate clearly the interactions of two competing weight profiles. Without recurrent connections of some kind, there would be no sustained activity packet to shift through the HD ring, and this model therefore represents the simplest possible architecture to observe path integration speed in relation to current models. It is believed that the case in which HD cells receive offset input from a different source, such as a layer of COMB cells described by Walters and Stringer, would not be qualitatively different.

Interestingly, the results in this chapter suggest that the slowing effects on path integration speed of non-offset recurrent collateral connections can be ameliorated by having offset input strong enough that the relative slowing effect is negligibly small. However, this kind of overwhelmingly strong path integration input would need to always overwhelm to the point where recurrent activation is negligibly small during path integration. As it stands, there is no evidence that path integration system input to the HD system is overwhelmingly strong relative to any as yet unverified recurrent input. Certainly, any scenario still involving recurrent collateral connectivity would struggle to explain the dependence of the HD signal on ascending vestibular input, as described above. Thus, this alternative seems relatively unlikely.

Another modelling simplification employed in this chapter is the form of the inhibition model used. A global inhibition constant $\tilde{w}^{HD}$ is used to describe the feedback effect of inhibitory interneurons

within the HD layer. This is present in both pre-wired and self-organising models, and acts to stabilise the network, with the firing of each cell increasing the feedback inhibition back onto that cell. However, in the self-organising model, the global feed-forward inhibition term I^{FF} is also included during training in the light. This constant is substracted from the firing of each HD cell in the layer when training input is applied. Although this does not explicity model interneurons, it is designed to capture the core functionality of feed-forward inhibition: training input will excite interneurons, which will in turn suppress the firing of cells within the HD layer. This is qualitatively different to feedback inhibition, in which the firing of HD layer neurons themselves excite interneurons. The effect of feed-forward inhibition is to prevent strong visual training input from flooding the network with activation without relying on stronger feedback inhibition, which would prevent network activity being sustained by recurrent collateral connectivity alone during testing in the absence of visual input.

This mixture of feed-forward and feedback inhibition was therefore necessary for the self-organising model to function correctly. The modelling simplification described above was used as only the general properties of interneuron action were required, and thus it was considered unnecessarily complex to model individual interneurons. These distinct kinds of inhibition are well-established phenomena in the hippocampal formation (Buzsáki, 1984). While specific physiological evidence has been found for feed-forward inhibition in the brain (Pouille and Scanziani, 2001; Gabernet et al., 2005), the existence of inhibitory interneurons in HD regions is by no means certain. However, it is possible that similar inhibitory effects can be achieved across several layers (Taube and Bassett, 2003; Xie et al., 2002; Rubin et al., 2001). For example, inhibitory GABAergic connectivity from DTN to LMN (Wirtshafter and Stratford, 1993), believed to be the origin of the HD signal.

The training paradigm used in these simulations was a highly idealised one, with the training activity packet moving through the HD layer at a constant velocity in a single direction. This paradigm was designed for clarity, and in order to clearly demonstrate the operational principles in action, and how they govern path integration speed. There is no reason why these models could not function in a different direction. However, a constant velocity is required in order to test how well the network reproduces this velocity during the test phase. Models of path integration based on an offset input can suffer from a lack of generalisation to different speeds (Walters et al., 2013; Walters and Stringer, 2010). This model demonstrates concepts of path integration accuracy relative to a single target speed, rather that how the HD cell system can cope with a variety of different speeds and directions, which is an ongoing research question. Future explorations should expand these principle to more realistic training paradigms in terms of varied speed and direction of head rotation.

To summarise, this chapter has investigated the issue of path integration speed in continuous attractor models of the head direction cell system. Two major factors were discovered to affect the replay speed of path integration systems: the presence of symmetric non-offset recurrent collateral connectivity within the layer, and neuronal rise time. In the case of rise time, it appears that this factor does not adversely affect the self-organisation of these path integration systems. Further simulations highlight the fact that it is possible to have perfect replay of target packet velocity if the rise time is negligibly small and recurrent connectivity within the layer does not contain a non-offset component. These findings represent a major contribution to theoretical understanding of the factors governing path integration speed. Several future approaches to coping with these speed-limiting factors are suggested.

Chapter 6

Path integration in a two-layer system: Solving the speed problem

In the previous chapter of this thesis, the issue of path integration speed was investigated. Two major factors governing path integration speed were discovered. Firstly, the time it takes for a neuron to begin firing once it has received input, known as neuronal rise time, is shown to slow path integration speed proportionally. HD cells in the previous Chapter's model project back onto themselves with a delayed offset weight profile. This offset is determined by the axonal conduction delay present in the weight profile and the target path integration speed, such that HD cells project maximally to the offset HD cell which ought to be active after the conduction delay at the current target speed. However, if the postsynaptic HD cells take some time to start firing, then by the time it is active, the activity packet ought to be elsewhere. Thus, neuronal rise time slows path integration speed.

Secondly, the presence of within-layer non-offset connectivity is also shown to slow path integration speed. This sort of connectivity is used to allow persistent HD firing without adaptation, even in the dark and without head rotation, as is seen *in vivo*. However, during path integration, this non-offset

connectivity interferes with the offset connectivity required for time-accurate path integration. The issue is that a single layer of recurrent collateral connections can learn only one specific rotation speed, whether this is no rotation or some fixed speed of rotation. At any other speed, recurrent connectivity will introduce error into path integration. Whilst the previous chapter used a single layer of recurrent HD connectivity, this problem will still exist with offset connectivity arriving into the HD layer from a different population of cells, provided that HD cells receive input from offset and non-offset sources at the same time. Thus, recurrent connectivity within the HD layer will always cause errors in path integration.

This chapter presents a solution to the problem caused by the second of these factors, which can be succinctly stated as how to allow the HD layer to both sustain activity at a given location and shift activity between locations at the correct speed without these two functions interfering with one another. In order to do this, a path integration model without recurrent HD layer connectivity, referred to henceforth as the NORC ("no recurrent collaterals"), is presented. First a pre-wired version of the NORC model is demonstrated to be able to both sustain a stationary HD layer activity packet and shift this activity packet during path integration. The model is demonstrated to have excellent ($> 99\%$) accuracy across a range of target velocities. Interrogation of the model reveals that it operates as theorised. Rise time is found to negatively affect packet speed during path integration and this is, as with the previous chapter, found to be in relation to conduction delay. The NORC model is also shown to operate when conduction delays are assigned randomly rather than taking some uniform value. A self-organising variant of the model is also demonstrated to function well, achieving 91% accurate path integration. The reasons for the differences in self-organised and pre-wired NORC model per-

formance are discussed. Together, this work represents an important proof of concept that an updated version of traditional CANN models of the HD cell system is not only capable of performing path integration, but does so with a high level of accuracy.

6.1 Experiment 1: Pre-wired network

6.1.1 Model description

The model in this chapter draws from the existing literature of HD cell modelling, and can be considered an extension of the standard CANN approach. Its architecture is pictured in Figure 6.1. It consists of three layers, one of HD cells, one of COMB cells and one of AHV cells. The HD cell layer is very similar to other models in this thesis so far: it consists of a ring of HD cells, with each cell *i* having a preferred firing direction x_i in the range $0 - 360\,^{\circ}$. These directions are distributed evenly around the HD ring. Cells are arranged topographically such that adjacent HD cells have adjacent preferred firing directions. This arrangement is for ease of visualisation and analysis, and does not reflect HD cell organisation in the brain; where no such topography has been observed. The current model would function just as well without HD cell organised by preferred firing direction. In this manner, a contiguous Gaussian packet of activity in the HD layer represents current head direction. A crucial aspect of this model is that HD cells do not have direct excitatory influence on one another. Instead, HD cells only project to (w^1) and receive from (w^2) the COMB cell layer. This HD to COMB and COMB to HD connectivity contains an axonal conduction delay Δt in both directions.

Cells in the COMB layer receive connections from the HD and AHV layers, and project back onto the HD layer. The AHV layer consists of two kinds of cell: a ROT cell, which is active during head

rotation, and a NOROT cell which is active when the head is not rotating. Accordingly, COMB cells are conceptually split into two subpopulations: those that receive from the ROT cells (ROT-COMB cells) via w^3 connectivity and those that receive from the NOROT cells (NOROT-COMB cells) via w^4 connectivity. Connectivity from the AHV to COMB layer does not contain any delays. Both NOROT and ROT variants of COMB cells have preferred directions similarly to HD cells, by virtue of the connectivity they receive from the HD cell layer, as described below.. These two COMB cell variants have separate preferred direction distributions, each evenly distributed from $0 - 360\,^\circ$.

Whilst ROT and NOROT cells are programmed as distinct in this model, it should be kept in mind that this is a simplification of the known physiology: in reality they originate from the same layer of cells, and do not exist as distinct populations. Indeed, the binary nature of the ROT and NOROT cells simply signals whether the head is rotating at the target velocity. This too is a simplification of angular head velocity (AHV) cells in the brain, which have complex firing responses to angular head velocity (see Chapter 1 of this thesis for a fuller description of AHV cell firing properties). The aim with these ROT and NOROT cells is simply to capture the fact that the AHV cell population as a whole will provide information about when the head is rotating at any given AHV, including 0, to the COMB cell layer.

Connectivity between the HD and COMB, as well as that between COMB and HD layers, is based on COMB subpopulations. NOROT-COMB cells receive and project from the same HD cell. These NOROT-COMB cells thus respond to a combination of head direction and a lack of head rotation. ROT-COMB cells receive maximally from an offset region of the HD layer, and project back onto

a further offset region of the HD layer. These offsets are based on the conduction delay Δt in each direction and the speed of head rotation signalled by the corresponding ROT cells. The operating principles of this connectivity will be fully detailed below..

Individual HD and COMB cells are modelled, as in previous chapters, using a leaky-integrator firing rate model. Consequently, individual spikes are not modelled, and instead cell firing is measured as an instantaneous average firing rate. ROT and NOROT cells are modelled as binary units, being either on (firing rate of 1) or off (firing rate of 0).

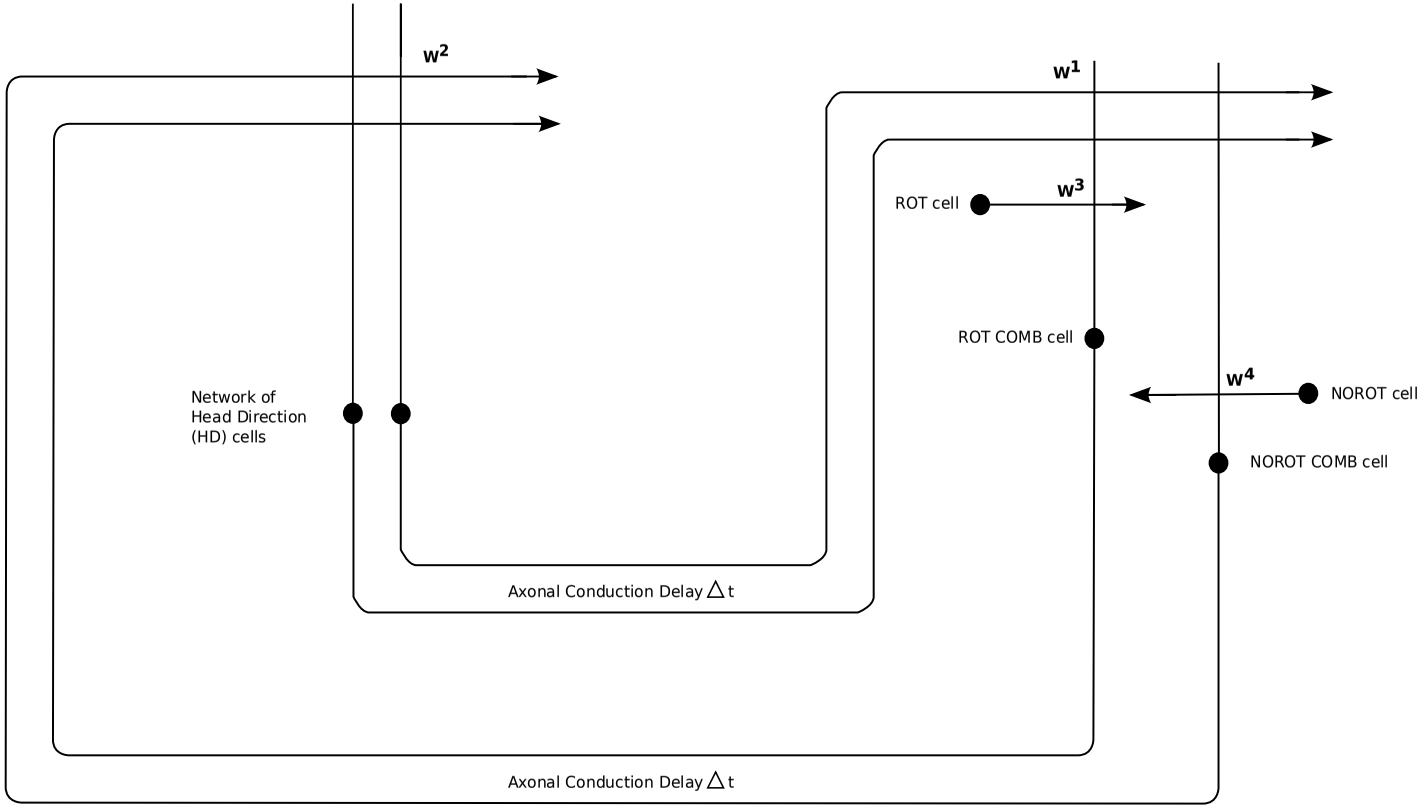


Figure 6.1: Network architecture, demonstrating w^1 connections from HD to COMB cells, w^2 connections from COMB to HD cells, w^3 connections from ROT cells to ROT-COMB cells, and w^4 connections from NOROT cells to NOROT-COMB cells.

6.1.1.1 Model equations

The activation level h_i^{HD} of head direction cell *i* is given by

$$\begin{aligned}\tau^{HD}\frac{dh_i^{HD}(t)}{dt} = &- h_i^{HD}(t) + e_i(t) \\ &- \frac{1}{N^{HD}}\sum_j \tilde{w}^{HD} r_j^{HD}(t) \\ &+ \frac{\phi_2}{C^{COMB\rightarrow HD}}\sum_j w_{ij}^2 r_j^{COMB}(t-\Delta t)\end{aligned} \tag{6.1}$$

where $h_i^{HD}(t)$ is the activation level of head direction cell *i* at time *t*. The term $-h_i^{HD}(t)$ represents a decay term, whereby in the absence of presynaptic input, the activation of head direction cell *i* will fall to zero as determined by the time constant τ^{HD}. The term $e_i(t)$ represents external visual input to head direction cell i at time t. This input term is used to generate an initial packet within the head direction cell layer. It is determined by a Gaussian profile described in Section 6. 1.1.2 below. The term $\frac{1}{N^{HD}}\sum_j \tilde{w}^{HD} r_j^{HD}(t)$ represents inhibitory feedback within the head direction cell layer, summed over all presynaptic head direction cells *j*. $\tilde{w}^{HD}$ is a global constant describing the effect of inhibitory interneurons, and N^{HD} is the total number of head direction cells in the layer. The term $\frac{\phi_2}{C^{COMB\rightarrow HD}}\sum_j w_{ij}^2 r_j^{COMB}(t-\Delta t)$ represents excitatory input from the COMB cell layer back onto the HD layer, summed across all presynaptic COMB cells connected to cell *i*. It consists of $r_j^{COMB}(t-\Delta t)$, the firing of presynaptic COMB cell *j* at time $t-\Delta t$, where Δt represents axonal conduction delay, and w_{ij}^2, the synaptic weight from presynaptic COMB cell *j* to postsynaptic HD cell *i*. It is important to note that this weight, nor any nother weight, is not updated during simulation, and

thus remains at the same initial value throughout. The entire term is scaled by the factor $\frac{\phi_2}{C^{COMB \to HD}}$, which controls the overall strength of COMB synapses onto the head direction cell layer where ϕ_2 is a constant and $C^{COMB \to HD}$ is the number of synapses each postsynaptic head direction cell receives from presynaptic COMB cells.

The firing rate $r_i^{\text{HD}}(t)$ of head direction cell i at time t is a sigmoid function of the activation level of cell i

$$r_i^{HD}(t) = \frac{1}{1 + e^{-2\beta^{HD}(h_i^{HD}(t) - \alpha^{HD})}} \tag{6.2}$$

where α^{HD} and β^{HD} are the sigmoid threshold and slope respectively. As this model is firing-rate based, no individual action potentials are simulated, and neuronal activity is represented as a spontaneous average firing rate.

The activation level h_i^{COMB} of ROT-COMB cells and NOROT-COMB cells indexed by i differ slightly. The ROT-COMB cell activation is given by

$$\begin{aligned}\tau^{COMB}\frac{dh_i^{COMB}(t)}{dt} = &- h_i^{COMB}(t) \\ &- \frac{1}{N^{COMB}}\sum_j \tilde{w}^{COMB} r_j^{COMB}(t) \\ &+ \frac{\phi_1}{C^{HD\rightarrow COMB}}\sum_j w_{ij}^1 r_j^{HD}(t-\Delta t) \\ &+ \frac{\phi_3}{C^{ROT\rightarrow COMB}}\sum_j r_j^{ROT}(t)\end{aligned} \tag{6.3}$$

whilst NOROT-COMB activation is given by

$$\begin{aligned}\tau^{COMB}\frac{dh_i^{COMB}(t)}{dt} = &- h_i^{COMB}(t) \\ &- \frac{1}{N^{COMB}}\sum_j \tilde{w}^{COMB} r_j^{COMB}(t) \\ &+ \frac{\phi_1}{C^{HD\rightarrow COMB}}\sum_j w_{ij}^1 r_j^{HD}(t-\Delta t) \\ &+ \frac{\phi_4}{C^{NOROT\rightarrow COMB}}\sum_j r_j^{NOROT}(t)\end{aligned} \tag{6.4}$$

where $h_i^{COMB}(t)$ is the activation level of COMB cell i at time t. The term $-h_i^{COMB}(t)$ represents decay, where in the absence of any input the activation of COMB cell i will fall to zero as determined by the time constant τ^{COMB}. The term $\frac{1}{N^{COMB}}\sum_j \tilde{w}^{COMB} r_j^{COMB}(t)$ represents inhibitory feedback within the COMB cell layer, summed over all presynaptic COMB cells *j*. $\tilde{w}^{COMB}$ is a global

constant describing the effect of inhibitory interneurons, and N^{COMB} is the total number of COMB cells in the layer. The term $\frac{\phi_1}{C^{HD\rightarrow COMB}} \sum_j w_{ij}^1 r_j^{HD}(t-\Delta t)$ represents excitatory input from the HD cell layer back onto the COMB layer, summed across all presynaptic HD cells connected to cell *i*. It consists of $r_j^{HD}(t-\Delta t)$, the firing of presynaptic HD cell *j* at time $t-\Delta t$, where Δt represents an axonal conduction delay, and w_{ij}^1, the synaptic weight from presynaptic HD cell *j* to postsynaptic COMB cell *i*. It is scaled by the factor $\frac{\phi_1}{C^{HD\rightarrow COMB}}$, which controls the overall strength of HD synapses onto the COMB cell layer where ϕ_1 is a constant and $C^{HD\rightarrow COMB}$ is the number of synapses each postsynaptic COMB cell receives from presynaptic HD cells.

The final term in the activation equation differs for NOROT-COMB and ROT-COMB cells. For ROT-COMB cells it is $\frac{\phi_3}{C^{ROT\rightarrow COMB}} \sum_j r_j^{ROT}(t)$, where $\sum_j r_j^{ROT}(t)$ is the summed firing of presynaptic ROT cells indexed by j at time t. This is scaled by the factor $\frac{\phi 3}{C^{ROT\rightarrow COMB}}$ where ϕ_3 is a constant, and $C^{ROT\rightarrow COMB}$ is the number of synapses each postsynaptic COMB cell receives from presynaptic ROT cells. These terms are similar for NOROT-COMB cells, with the substitution of NOROT for ROT and ϕ_4 for ϕ_3. In these simulations, $N^{ROT} = 1$ and $N^{NOROT} = 1$, and as such the terms $\frac{\phi_3}{C^{ROT\rightarrow COMB}} \sum_j r_j^{ROT}(t)$ and $\frac{\phi_4}{C^{NOROT\rightarrow COMB}} \sum_j r_j^{NOROT}(t)$ can be simplified to $\phi_3 r^{ROT}(t)$ and $\phi_4 r^{NOROT}(t)$ respectively.

Note, then, that the activations of NOROT-COMB and ROT-COMB cells differ only by whether they receive input from ROT or NOROT cells. In all other cases, i.e., where inhibition is concerned, COMB cells function as a homogenous population.

Similar to HD cell firing rates, the firing rate $r_i^{\text{COMB}}(t)$ of COMB cell i at time t is a sigmoid function of the activation level of cell i

$$r_i^{COMB}(t) = \frac{1}{1 + e^{-2\beta^{COMB}(h_i^{COMB}(t) - \alpha^{COMB})}} \tag{6.5}$$

Note that both the sigmoid thresholds α^{COMB} and α^{HD}, as well as slopes β^{COMB} and β^{HD} may differ in value between COMB and HD cell layers.

In this pre-wired network, all connectivity is set before, and does not change over the course of simulation. w^1 connectivity from HD to COMB cells is set up in two distinct stages. Firstly, non-offset connectivity is set up from HD to NOROT COMB cells. This is calculated as a Gaussian configuration, with the weight w_{ij} for synapses onto postsynaptic NOROT-COMB cell i from presynaptic head direction cell j being

$$w_{ij}^1 = e^{-(s_{ij}^{HDCOMB})^2/2(\sigma^{HDCOMB})^2} \tag{6.6}$$

where σ^{HDCOMB} is the standard deviation of the Gaussian profile. s_{ij}^{HDCOMB} is the difference between the preferred head directions s_i and s_j of the post- and pre-synaptic cells i and j respectively. It is given by

$$s_{ij}^{HDCOMB} = MIN(|x_i - x_j|, 360 - |x_i - x_j|) \tag{6.7}$$

where x_i and x_j are the preferred head directions of cells i and j respectively. This creates a wrap-around effect, with weights between head direction cells as a population remaining continuous across the $360/0\,^{\circ}$divide.

Offset connectivity from HD to ROT-COMB cells is calculated similarly, with the weight w_{ij} for synapses onto postsynaptic ROT-COMB cell i from presynaptic head direction cell j being calculated as in Equation 6. However, the difference between preferred head directions of post- and pre-synaptic cells i and j now includes an additional offset O, and is calculated as

$$s_{ij}^{HDCOMB} = MIN(|x_i - (x_j + O)|, 360 - |x_i - (x_j + O)|) \tag{6.8}$$

where the offset O is calculated, based on the target speed $V\,^{\circ}/s$ and the axonal conduction delay Δt, as

$$O = V\Delta t \tag{6.9}$$

This effectively causes presynaptic HD cells to be connected to an offset region of the ROT-COMB cell layer.

w^2 connectivity from COMB to HD cells is pre-wired in an identical fashion to w^1 connectivity.

The sole difference is that the perspective is now switched, with the HD layer becoming post-synaptic cells indexed by i and COMB cells are presynaptic cells j. The resulting combination of w^1 and w^2 connectivity means that individual HD cells both project symmetrically back onto themselves via NOROT-COMB cells, as well as asymmetrically back onto a region of the HD layer with an offset of $2O$ via ROT-COMB cells.

As this model makes use of a single ROT and a single NOROT cell, w^3 connectivity from the ROT cell to ROT-COMB cells and w^4 connectivity from the NOROT cell to NOROT-COMB cells is not explicitly modeled. Instead, both COMB cell subtypes receive input from their resective AHV cell weighted directly by the parameters ϕ_3 for the ROT cell onto ROT-COMB cells and ϕ_4 for the NOROT cell onto NOROT-COMB cells.

The differential equations given for this model cannot be solved analytically. As in the previous chapters, they are solved numerically by a Forward Euler finite difference scheme with a timstep size of δt. The interested reader should refer to Chapters 3 and 4 for full details.

6.1.1.2 Simulation protocol

At the beginning of simulation, firing rates of all cells are set to zero. Connectivity from HD to COMB, from COMB to HD and from AHV cells to COMB is set. An external input, represented as the term e_i is applied to each cell in the HD layer for $100ms$ to initialise network activity. e_i is calculated for a given postsynaptic HD cell i according to the Gaussian function

$$e_i = \lambda^{HD} e^{-(s_i^{HD})^2/2(\sigma^{HD})^2} \tag{6.10}$$

where λ^{HD} is a scaling factor determining the strength of this input to head direction cell i and σ^{HD} is the standard deviation of the Gaussian profile. s_i^{HD} is the difference between the true head direction x and the preferred direction x_i of postsynaptic head direction cell i. It is given by

$$s_i^{HD} = MIN(|x_i - x|, 360 - |x_i - x|) \tag{6.11}$$

which creates a wrap-around effect, with response profiles of head direction cells as a population remaining continuous across the $360/0\,^{\circ}$divide. e_i thus generates in the HD layer a Gaussian activity packet centered on location x. During this period, HD activations h_i^{HD} and firing rates r_i^{HD} are updated as in Equations 1 and 2. COMB activations h_i^{COMB} and firing rates r_i^{COMB} are also updated as in Equations 3, 4, and 5, with $r^{NOROT} = 1$ and $r^{ROT} = 0$.

After this initialisation period, e_i is removed. There then follows a $2s$ period of zero head rotation. All activations and firing rates continue to update, and AHV cells remain set with $r^{NOROT} = 1$ and $r^{ROT} = 0$.

Next there is a $2s$ period of head rotation. AHV cell firings are changed to reflect the fact that head rotation is occurring, with $r^{NOROT} = 0$ and $r^{ROT} = 1$. HD and COMB cell activations and firing rates continue to update.

Finally, there is another $1s$ period of zero head rotation, in which HD and COMB cell activations and firings continue to update and AHV cell firings are set to reflect zero head rotation with $r^{NOROT} = 1$ and $r^{ROT} = 0$.

6.1.1.3 Operating principles

A diagram of the model's operation is given in 6.2. When the agent's head is not rotating, NOROT-COMB cells, responding to a combination of zero angular head velocity and the current head direction, will become active. These cells project directly back to the HD cell from which they receive. This non-offset reciprocal connectivity across two layers will allow the HD activity packet to be sustained in its current location. Crucially, this connectivity will not remain active during head rotation, as the angular head velocity cells will no longer signal a lack of head rotation, and thus NOROT-COMB cells will become quiescent.

When the agent rotates its head at arbitrary velocity V°/s, ROT-COMB cells will become active. These cells respond to a combination of current head direction and angular head velocity V°/s. They project back onto the HD layer with an offset based on the target velocity and the conduction delay Δt. This offset connectivity means that a given HD cell will project, via the ROT-COMB layer, most strongly to the HD cell which ought to be active given the total conduction delay and the target velocity. It should be noted that the effective total conduction delay in this case will be $2\Delta t$, as the HD layer projects back onto itself via two sets of connections (HD to COMB and COMB to HD). This will act to shift the activity through the HD layer to the correct location, allowing the system to accurately track true head direction given current angular head velocity. Effectively, this HD to ROT-COMB to HD connectivity allows the system to know where a final HD activity packet ought to be, given a particular starting location, following head rotation at a specific angular head velocity over a known time interval.

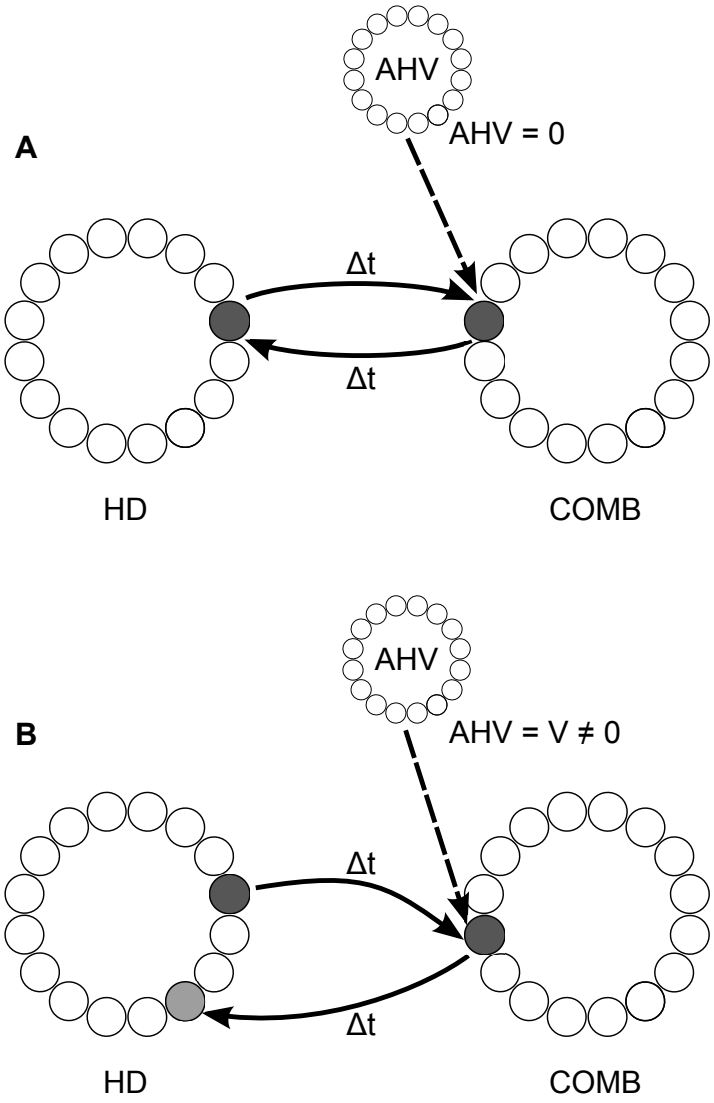


Figure 6.2: Operation of the NORC model, with HD activity both sustained and shifted by reciprocal connectivity with a layer of combination (COMB) cells. This connectivity contains a conduction delay Δt. COMB cells respond to a particular combination of HD and angular head velocity (AHV). During no head rotation (A), HD activity is stabilised through reciprocal connectivity with COMB cells representing the combination of current head direction and no angular head velocity (AHV = 0). When the head is rotating (B), activity is shifted through the HD layer via connectivity with different COMB cells representing the conjunction of current head direction with angular head velocity V (AHV = V, where V is some non-zero value). Connectivity from these COMB cells back onto the HD layer projects with an offset from the original HD activity of $2\Delta tV$; causing HD cell activity to accurately track true head direction given current angular head velocity V.

This mechanism will be easily observable, as there is a very clear pattern of firing across the two layers. Consider the model's operation during a head turn. At time $t = 0$, the agent begins rotating its head. At this stage only an initial HD cell and a ROT cell will be active. At time $t = \Delta t$, the ROT-COMB cell which receives maximally from the initial HD cell will become active, as firing from the initial HD cell will have arrived after a conduction delay. At time $t = 2\Delta t$, the HD cell which receives maximally from the ROT-COMB cell which was active at time $t = \Delta t$ will become active, again because of the conduction delay between layers. Thus, any changes in state in the HD or COMB layer will have a specific temporal ordering: HD layer changes in state will be separated by $2\Delta t$ from one another, as will COMB layer changes in state. However, HD layer changes in state will be separated from COMB layer changes in state by Δt.

6.1.2 Results

This chapter makes use of a population vector statistic, which is described in more detail in Chapter 3, Section 3.4. The population vector (Pvector) gives the centre of mass (i.e., location) of an activity packet within that particular layer. For all HD cells, this statistic is calculated exactly as detailed in Chapter 3. COMB cells are, however, split into ROT-COMB and NOROT COMB populations, with each population each having it's own distribution of cell preferred firing directions in the interval $[0-360\,^{\circ}]$. This is to simplify plotting of Pvector: any plots of COMB layer Pvectors are collated and thus will show NOROT-COMB Pvectors during stationary phases, and ROT-COMB Pvectors during rotation. Pvectors are also used to calculate path integration speed, which was calculated by noting the total distance travelled during the rotation phase of simulation and diving by rotation time.

This chapter also analyses the precise timings of packet movement during simulation. This was done by taking the gradient of the Pvector values using inbuilt *MATLAB* software gradient function, and recording the time of changes in this gradient. HD-HD, HD-COMB and COMB-COMB shift intervals were then calculated using these shift times, and the first 10 readings were averaged to give the final shift interval statistics.

Table 6.1: Network parameter values for simulations reported in Experiment 1. All simulations, except where explicitly stated, are run using these parameters. All measures of time are given in seconds.

Network Parameters	
Head Direction Cells	500
ROT-COMB Cells	500
NOROT-COMB cells	500
Simulation Time	$4.1s$
$C^{NOROT-COMB \rightarrow HD}$	500
$C^{ROT-COMB \rightarrow HD}$	500
ϕ_1	700.0
ϕ_2	4500.0
σ^{HDCOMB}	20.0°
σ^{COMBHD}	20.0°
ROT Cells	1
NOROT Cells	1
$C^{ROT \rightarrow COMB}$	1
$C^{NOROT \rightarrow COMB}$	1
ϕ_3	80.0
ϕ_4	80.0
α^{HD}	0.0
β^{HD}	0.2
α^{COMB}	16.0
β^{COMB}	0.3
τ^{HD}	$0.0001s$
τ^{COMB}	$0.0001s$
$\tilde{w}^{HD}$	0.2
$\tilde{w}^{COMB}$	0.35
δt	$0.0001s$
λ^{HD}	2.0
σ^{HD}	20.0°

6.1.2.1 Main Result

Figure 6.3 shows the model operation, with parameter values given in Table 6.1, with a variety of target velocities in the range $[0 - 360\,^{\circ}/s]$. We see clearly that the path integration velocity attainable by the model is extremely close to the target velocity. Indeed, this value is in the region of 99.5% for all target velocities in the range tested.

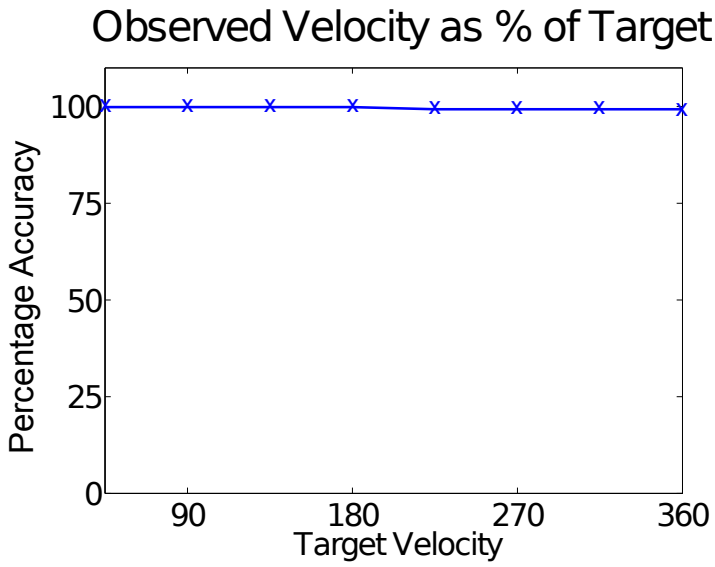


Figure 6.3: Simulation results of pre-wired network model. Path integration speed as a percentage of target for a range of target velocities, with parameters as given in Table 6.1. Model performance is excellent (observed speed $> 99\%$ of target) at all target velocities tested.

Examining the motion of HD and COMB activity packets during path integration yields some interesting results. Figure 6.4 shows HD layer packet firing over the entire course of simulation with $V = 180\,^{\circ}/s$, whilst Figure 6.5 shows the position (Pvector) of HD (blue) and COMB (red) activity packets during seconds 1 to 1.4 of simulation with $V = 180\,^{\circ}/s$. Rather than moving smoothly through their respective layers, we instead see both HD and COMB activity packets shifting in a periodic manner. This is particularly visible in Figure 6.5. This is likely an effect of having two linked layers separated by a single value of conduction delay: an HD activity packet will shift itself, via the

COMB layer after $2\Delta t$. Indeed, this periodic behaviour ought to match the operating principles given above: HD shifts will be separated from one another by $2\Delta t$, which is also the separation between COMB shifts. HD shifts will be separated from COMB shifts by Δt. This matches what was observed in these simulations, in which $\Delta t = 0.01s$: recording the shift times of the two activity packets and averaging the HD to HD, COMB to COMB and HD to COMB shift intervals reveals that HD to HD and COMB to COMB shift intervals were both $0.02s$ $(2\Delta t)$, whilst the HD to COMB shift interval was $0.01s$ (Δt).

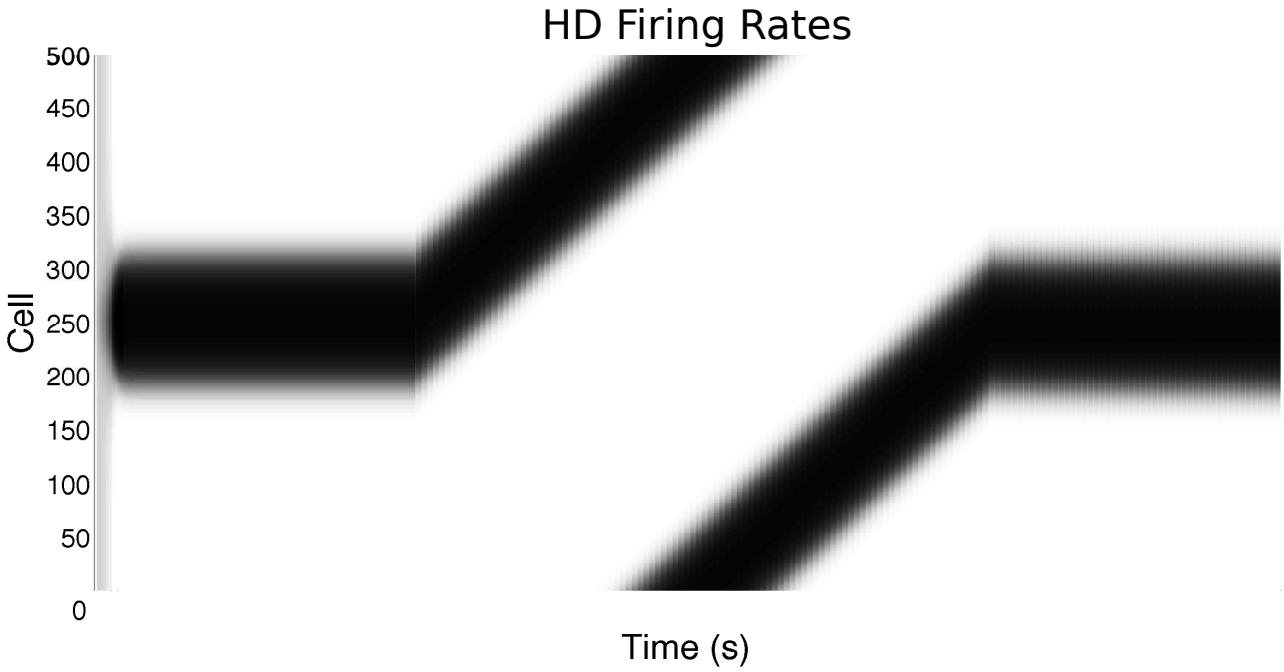


Figure 6.4: HD firing rates during simulation with parameters given in Table 6.1 and target velocity $V = 180\,^{\circ}/s$. A consistent HD layer activity packet is achieved. This packet is sustained in the absence of external input, both stationary and during path integration. Overall packet speed during path integration $(179.73^{\circ}/s)$ closely matches target velocity. However, path integration tracking is not smooth, but is instead comprised of a multitude very quick position jumps followed by stationary points.

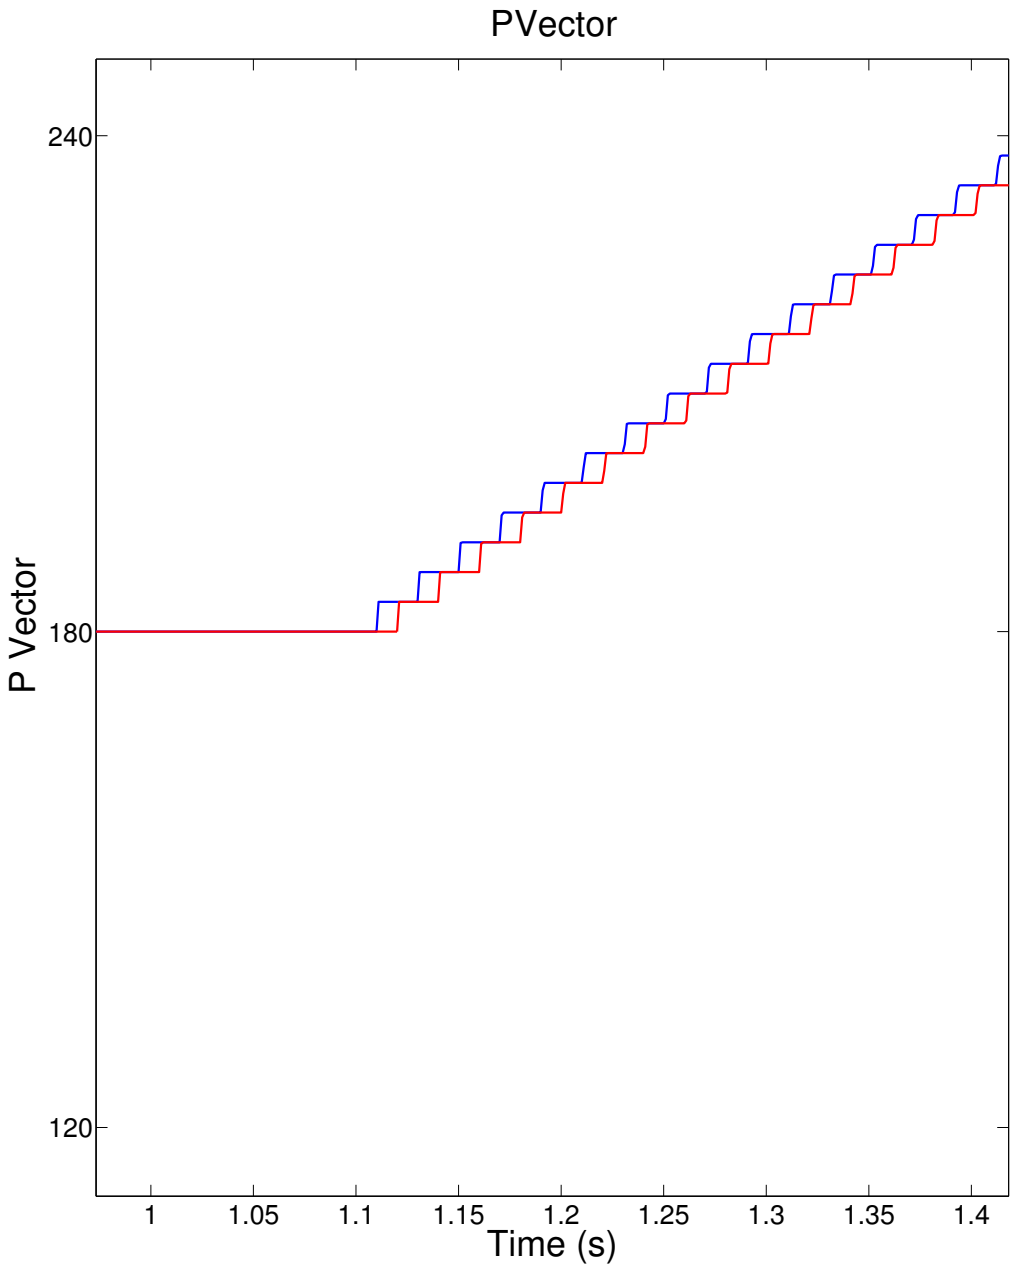


Figure 6.5: Position (PVector) of both HD (blue) and COMB (red) activity packets from seconds 1 to 1.4 of simulation. Whilst both packets are stationary up to $1.1s$, beyond this the two activity packets shift in a manner consistent with the model operating principles. The HD packet shifts first, followed by the COMB packet. The separation between HD and COMB shifts is $0.01s$, equal to the conduction delay $\Delta t = 0.01s$. The separation between HD shifts is $0.02s$, as is the separation between COMB shifts, with both equal to twice the conduction delay.

6.1.2.2 Effect of neuronal time constant

A key factor in the path integration speed, as discussed in the previous chapter, is the neuronal time constant τ^{HD} (and in this model τ^{COMB}). For this reason, an investigation was made into the effects of time constant across values of τ^{HD} and τ^{COMB} in the range [$0.0001s$,$0.01s$]. The results of this investigation are given in Figure 6.6. In all cases, target velocity was $V = 180\,^{\circ}/s$, conduction delay was fixed at $\Delta t = 0.01$, and $\tau^{HD} = \tau^{COMB}$. In line with the results of the previous chapter, we see that increasing rise time via increasing the value of both τ^{HD} and τ^{COMB} reduces packet speed. A top speed of $179.43\,^{\circ}/s$ was achieved in these simulations with $\tau = 0.0001s$. In other regards, this pattern is very similar to that shown in Figure 6.6 of the previous chapter. This suggests that, whilst this model is architecturally distinct from that of the previous chapter, it still retains core path integration functionality.

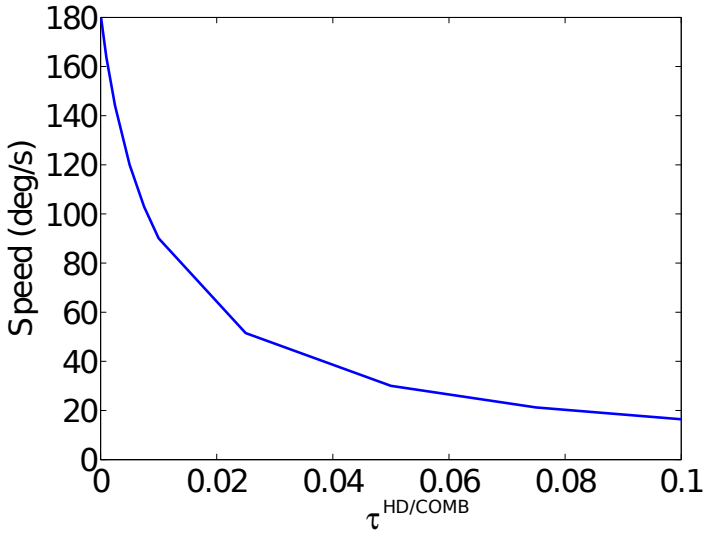


Figure 6.6: Effect of neuronal time constant on path integration accuracy. Top speeds are shown for a range of values of $\tau^{HD} = \tau^{COMB}$ with a target velocity $V = 180\,^{\circ}/s$. As τ^{HD} and τ^{COMB} increase, so too does the effect of rise time, and thus path integration speed is reduced.

The effect of rise time on packet speed can be examined further by considering the periodic behaviour observed in Figures 6.4 and 6.5. According to the hypothesis, these changes in position should be very closely linked to the conduction delay Δt between layers, as described in Section 6.1.1.3. If rise time causes neurons to fire after they are supposed to, then we should see this link break down with increasing values of $\tau = \tau^{HD} = \tau^{COMB}$. Table 6.2 shows summary results detailing the average time interval between successive HD activity packet shifts, COMB activity shifts, and between HD and COMB activity shifts. All simulations are run with the same range of values of τ [$0.0001s$,$0.01s$], whilst conduction delay is held constant at $\Delta t = 0.05$. Although the expected HD-HD, COMB-COMB and HD-COMB intervals for $\Delta t = 0.05$ of $0.1s$, $0.1s$ and $0.05s$ respectively are observed with smaller values of τ, we see that increasing τ causes this relationship to break down. As τ increases we see that neurons are beginning to fire after an increasing delay, disrupting the temporal relationship of packet motion in both HD and COMB layer, thus causing inaccurate path integration.

Table 6.2: Summary of results for successive HD-HD, COMB-COMB, HD-COMB average packet shift intervals across a range of values of τ with conduction delay $\Delta t = 0.05$. All time values are given in seconds.

	Shift Intervals		
τ	HD-HD	COMB-COMB	HD-COMB
0.0001	0.1	0.1	0.05
0.001	0.0999	0.0999	0.0500
0.0025	0.0999	0.0999	0.0503
0.005	0.1047	0.1047	0.0505
0.0075	0.1088	0.1091	0.0530
0.01	0.1136	0.1146	0.0549

6.1.2.3 Effect of conduction delay

Another key factor in path integration speed, as discovered in the previous chapter, is the conduction delay Δt. For this reason, a variety of different conduction delays were used in this model. Figure 6.7 shows summary results from simulations run with the parameters given in Table 6.1 and Δt varying in the range [$0.001s$,$0.05s$]. In all simulations $\tau^{HD} = \tau^{COMB} = 0.001s$ and $V = 180\,^{\circ}/s$. As expected given the results in the previous chapter, we see that shorter conduction delays are associated with reduced packet speed, whilst longer conduction delays show more accurate path integration. This is because, as elaborated in the previous chapter, a constant error introduced by neuronal rise time will be proportionally more severe in the context of a smaller conduction delay rather than a larger one.

As with the effect of τ, the effect of varying Δt on model operation can be shown in detail by considering its effect on periodic HD and COMB layer packet shifts. Table 6.3 shows a summary of the average time interval between successive HD packet shifts, COMB activity shifts, and between HD and COMB activity shifts. All simulations are run with varying Δt values in the range [$0.001s$,$0.05s$], whilst τ has been reduced to a negligibly small value $\tau^{HD} = tau^{COMB} = 0.0001s$ in order to see more clearly the effect of changing conduction delay without interference from neuronal rise time. As would be expected given model operating principles, the inter-shift interval is always $2\Delta t$ for HD-HD and COMB-COMB intervals, and is always Δt for HD-COMB intervals.

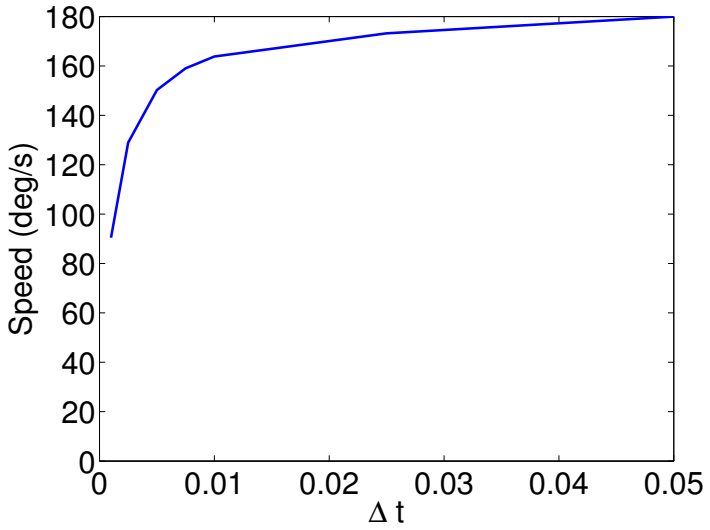


Figure 6.7: Effect of conduction delay on path integration accuracy. Top speeds are shown for a range of values of conduction delay Δt with a target velocity of $V = 180\,^{\circ}/s$ and a fixed neuronal time constant of $\tau^{HD} = \tau^{COMB} = 0.001s$. As seen in the previous chapter, a constant error introduced by rise time will be proportionally more severe in the context of shorter conduction delays. Thus, the slowing effect on path integration of rise time is increased for small Δt and decreased for large Δt.

Table 6.3: Summary of results for successive HD-HD, COMB-COMB, HD-COMB average packet shift intervals across a range of values of Δt with $\tau = 0.0001$. This value of τ has been reduced to a negligibly small value to more accurately observe the effects of Δt on network dynamics). All time values are given in seconds.

Shift Intervals			
Δt	HD-HD	COMB-COMB	HD-COMB
0.0025	0.005	0.005	0.0025
0.005	0.01	0.01	0.005
0.0075	0.015	0.015	0.0075
0.01	0.02	0.02	0.01
0.025	0.05	0.05	0.025
0.05	0.1	0.1	0.05

6.1.2.4 A distribution of conduction delays

The results so far demonstrate model functionality when all HD to COMB w^1 and COMB to HD w^2 connections involve the same axonal conduction delay Δt. However, in the brain it is much more likely that such a system would operate with varying values of Δt across different synapses. Therefore, simulations were carried out in which all w^1 and w^2 synapses have their own value of Δt, drawn randomly from a uniform distribution in the range [$0.001s$,$0.1s$].

Figure 6.8 shows summary results giving the speed from simulations run with different values of τ^{HD}, which was in all cases equal to τ^{COMB}. Target velocity was held constant at $V = 180\,^\circ/s$. All w^1 and w^2 weights had conduction delays randomly drawn from a uniform distribution as described above. However, these values remained fixed across all simulations with different values of τ, in order to accurately compare the effects of τ on simulations with a distribution of Δt values. Although affected by rise time in a similar manner to models with a single value of Δt, the model operates well across a wide range of τ values. Indeed, a distribution of conduction delays does not prevent the model from achieving very accurate path integration: a top speed of $179.24\,^\circ/s$ was recorded with $\tau = 0.001s$. In fact, the results shown in Figure 6.8 with distributions of axonal delays are significantly more accurate, across a range of values of τ, than the results shown in Figure 6.6 with a single value of conduction delay.

Interestingly, using a distribution of conduction delays appears to eliminate the periodic path integration behaviour seen in previous simulations using a single value of Δt. Figure 6.9 shows the HD layer firing rates during simulation with a distribution of Δt values in the range $0.001s$,$0.1s$] and a

neuronal time constant of $\tau = 0.001$. Figure 6.10, comparably to Figure 6.5, shows HD and COMB activity packet locations (P vectors) for simulation with the same parameters as Figure 6.9, focusing on seconds 1 to 1.4. Unlike simulations with a single Δt value, packet motion during path integration is smooth, and no periodic behaviour is observed. However, packet speed is not perfectly equal to target velocity at all times, particularly when transitioning from a stationary activity packet to a moving one. This is due to the distribution of delays causing the HD layer to receive conflicting COMB-layer inputs: for a short time the HD layer will receive input from NOROT-COMB cells with longer conduction delays and ROT-COMB cells with shorter delays.

Another feature of these simulations is that when stationary, the HD activity packet shows a very slight drift. This was found across multiple simulations (not pictured). Although the amount and direction of the drift varied, it was typically in the region of $1 - 4\,^{\circ}/s$ (2% error). This drift is likely a quirk of the random distribution of axonal conduction delays. If a situation arises in which the current HD activity location receives more stimulation on one side or the other, the activity packet will drift. At any given timestep, this situation might occur if the distribution of delays is uneven. Although conduction delays are drawn randomly from a uniform distribution, the sample taken will only approach uniformity for larger sample sizes (i.e., much larger numbers of cells than 500 of each variety). This is demonstrated in Figure 6.11, which shows a histogram of delays drawn randomly from a uniform distribution in the range [$0.001s$,$0.1s$] with varying sample size N. For small N, the sample does not accurately reflect a uniform distribution. As N increases however, the sample begins to look more and more uniform. Although we could thus infer that a larger network might well function with less error, full simulations with larger number of cells have proven to be computationally too demanding.

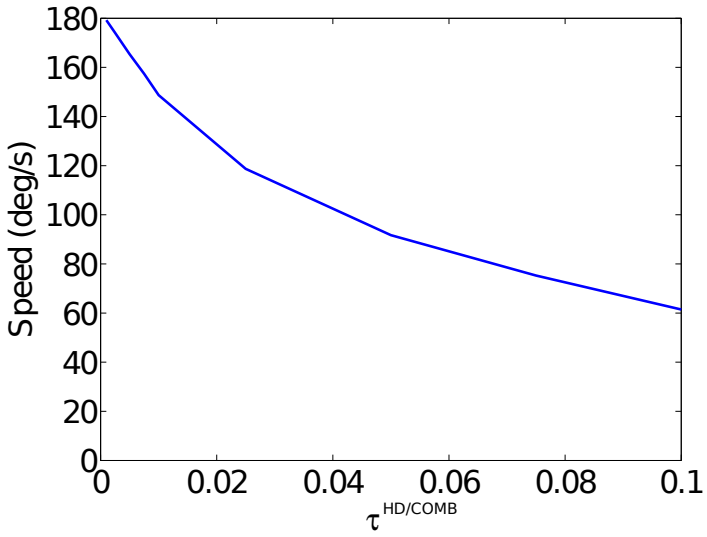


Figure 6.8: Effect of neuronal time constant on a model with a distribution of conduction delay values in the range $[0.0001s, 0.1s]$. Importantly, this model maintains significantly better path integration accuracy across the range of τ values when compared to a model with a single value of Δt. Identically to Figure 6.6 of this chapter, we see that increased rise time via increasing $\tau^{HD} = \tau^{COMB}$ reduces path integration accuracy.

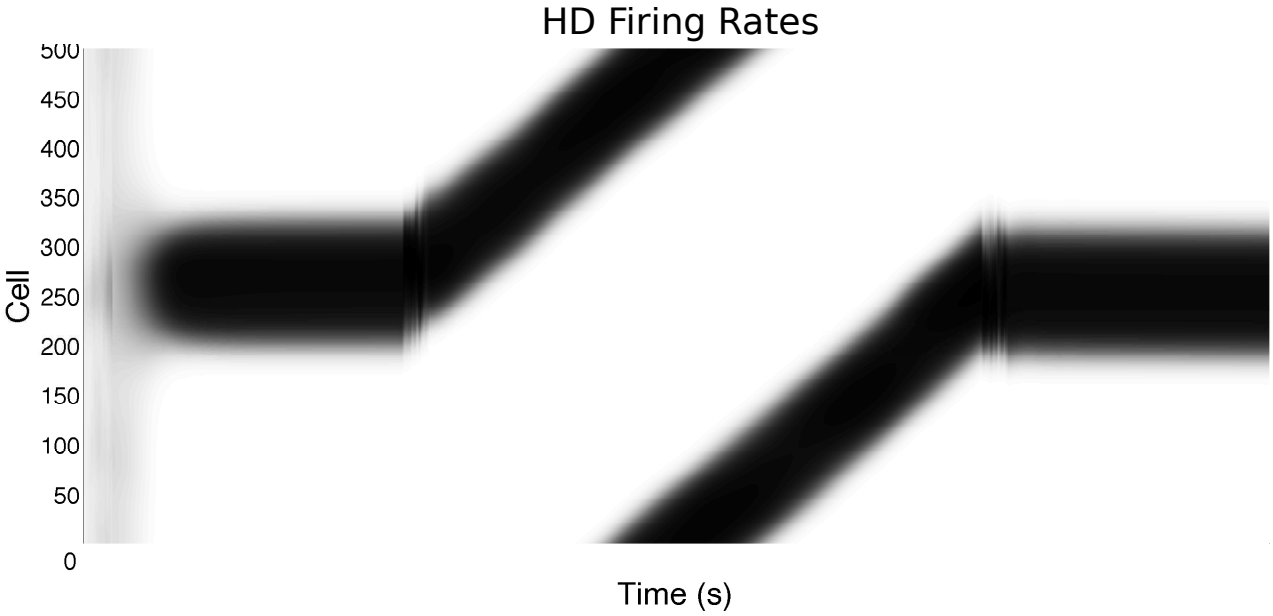


Figure 6.9: HD firing rates during simulation with parameters given in Table 6.1, and a distribution of conduction delays. A target velocity of $V = 180\,^{\circ}/s$ is used. Similarly to Figure 6.4, a consistent HD packet is seen both when the agent is stationary and during path integration. Thus it seems that a distribution of Δt values is not detrimental to model performance. However, we do see a loss of periodic path integration behaviour as seen in simulations with a single Δt value.

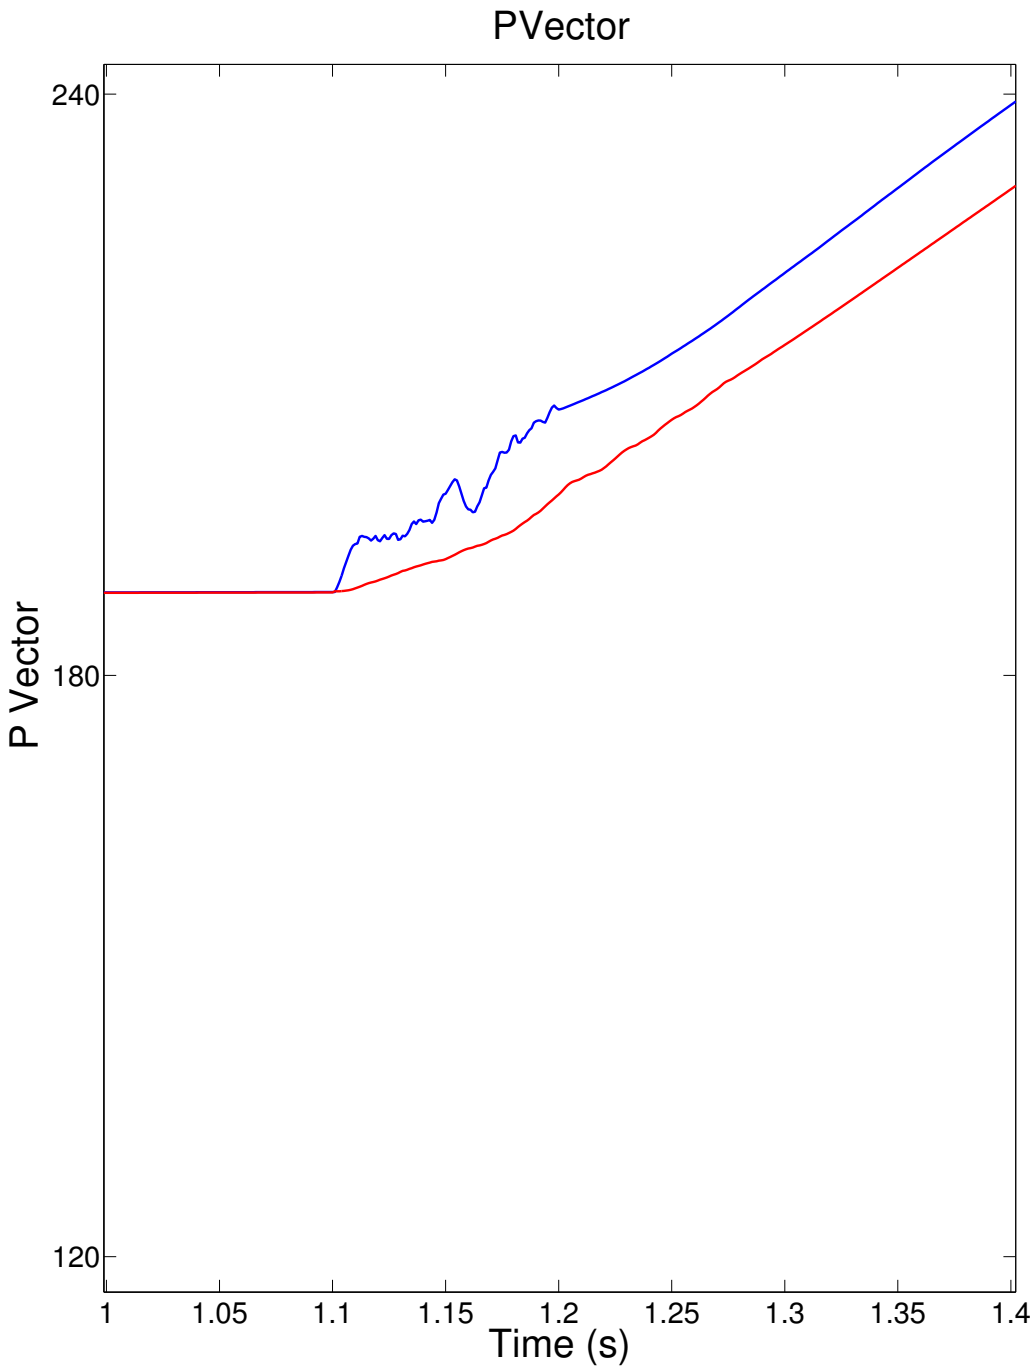


Figure 6.10: Position (PVector) of both HD (blue) and COMB (red) activity packets from seconds 1 to 1.4 of simulation with a distribution of conduction delays. Compared to Figure 6.5, we see clearly the loss of periodic path integration behaviour in the absence of a single Δt value.

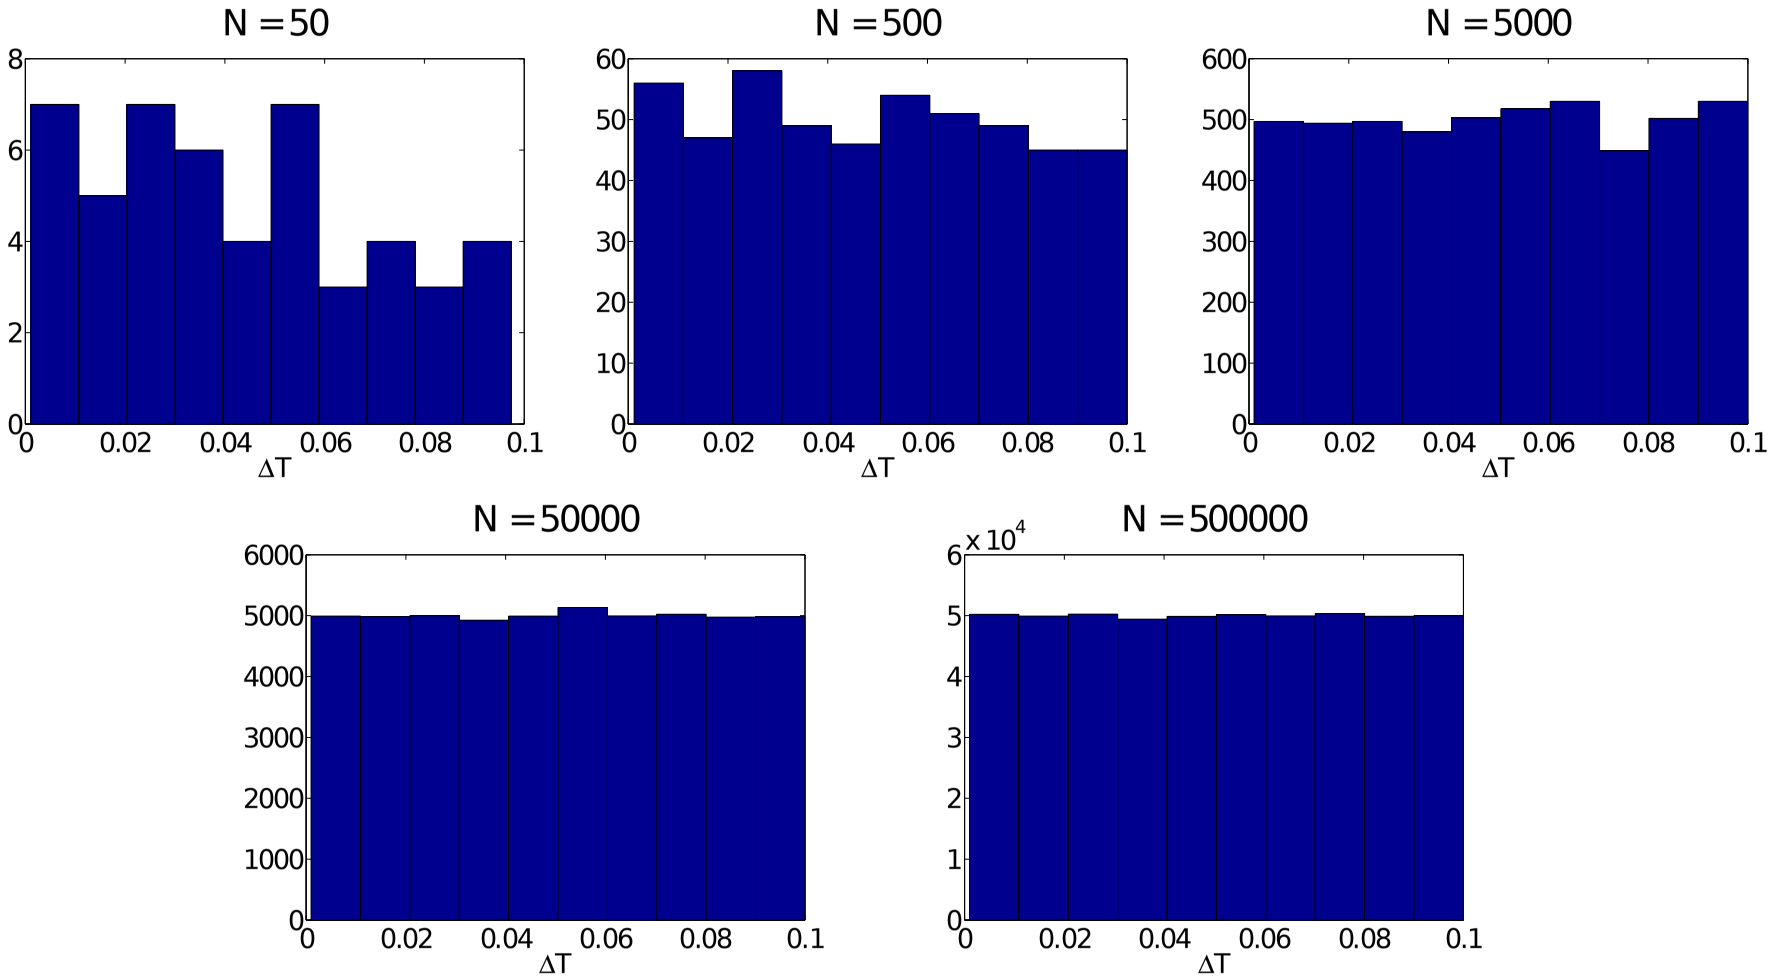


Figure 6.11: Histograms of Δt values in the range [$0.001s$,$0.1s$] drawn randomly from a uniform distribution with varying sample sizes (N). We see clearly that smaller sample sizes (smaller N) are associate with less sample uniformity. This is a possible explanation for drift in simulations with a distribution of Δt values.

6.2 Experiment 2: Self-organising network

Given that pre-wired simulations in Experiment 1 demonstrate the core functionality of the hypothesised system with idealised connectivity, a next step is to check whether such a system can self-organise. Can the self-organised model, in which the synaptic connections have been set up by learning, perform accurate path integration? In particular, can ROT-COMB and NOROT-COMB cells be formed without this distinction being imposed on the model by artificial pre-wiring of the synaptic connections? This section details the architecture, governing equations and results from a self-organising variant of the model demonstrated in in Experiment 1.

6.2.1 Model description

The model used for this experiment is almost identical to that used in Experiment 1. The principle difference is that connectivity is not pre-wired. All connections (HD-COMB, COMB-HD, ROT-COMB) are initialised with random values. During simulation, all weight profiles self-organise via Hebbian learning rules over the course of model training in the light, detailed below, with the final weight profile then being tested in the manner as Experiment 1.

Crucially, HD to COMB connectivity is diluted. That is, each COMB cell receives projections from a limited number of HD cells. In the simulations detailed below, each of the 1000 COMB cells receives from 25 of the 500 HD cells, giving 5% connectivity. This is in order to avoid the effects of continuous transformation (CT) learning (Stringer et al., 2006). In brief, the CT effect is a result of the fact that the activity packet moves with spatio-temporal continuity during training. Say that an arbitrary pattern of COMB cell firing is driven up by an HD activity packet at a given position. Connections

will be strengthened between all active HD cells and all active COMB cells. If, at the next timestep, the next HD activity packet has enough overlap with that of the previous timestep, the same or very similar pattern of COMB cells will be driven up by virtue of similar HD cells being involved, and the strength of their projections to the COMB layer. However, some additional HD cells will become active, and their connections onto active COMB cells will be strengthened. Over multiple iterations of this process, it is possible for connections to be strengthened between individual COMB cells and all HD cells, thus eliminating COMB specificity to HD. By reducing the number of HD cells from which a COMB cell receives connections, the CT effect is avoided.

Also present in the model of Experiment 2 is feed-forward HD layer inhibition during training. This is to ensure that the effect of excitatory visual input during training does not flood the network with activation when combined with HD-COMB reciprocal connectivity.

Unlike the previous experiment, no *a priori* distinction is made between ROT-COMB and NOROT-COMB cells. Instead, the COMB layer receives connections from a layer of 1000 ROT cells, half of which are active during head rotation, and half of which are active when the head is stationary.

6.2.1.1 Model equations

All model equations are as in Experiment 1, with the exception of the HD layer activation equation, which now incorporates feed-forward inhibition. This maintains a similar level of network activation across both training and testing, and is represented by the constant I^{FF}. The activation level h_i^{HD} of head direction cell i is therefore now given by

$$\tau^{HD}\frac{dh_i^{HD}(t)}{dt} = -\ h_i^{HD}(t) + e_i(t) - \frac{1}{N^{HD}}\sum_j \tilde{w}^{HD} r_j^{HD}(t) + \frac{\phi_2}{C^{COMB\to HD}}\sum_j w_{ij}^2(t) r_j^{COMB}(t-\Delta t) - I^{FF} \tag{6.12}$$

where $h_i^{HD}(t)$ is the activation level of head direction cell *i* at time *t*. The term $-\ h_i^{HD}(t)$ represents a decay term, whereby in the absence of presynaptic input, the activation of head direction cell *i* will fall to zero as determined by the time constant τ^{HD}. The term $e_i(t)$ represents external visual input to head direction cell i at time t. As in Experiment 1, this input term is used to generate an initial packet within the head direction cell layer. It is determined by a Gaussian profile described for Experiment 1 The term $\frac{1}{N^{HD}}\sum_j \tilde{w}^{HD} r_j^{HD}(t)$ represents inhibitory feedback within the head direction cell layer, summed over all presynaptic head direction cells *j*. $\tilde{w}^{HD}$ is a global constant describing the effect of inhibitory interneurons, and N^{HD} is the total number of head direction cells in the layer. The term $\frac{\phi_2}{C^{COMB\to HD}}\sum_j w_{ij}^2(t) r_j^{COMB}(t-\Delta t)$ represents excitatory input from the COMB cell layer back onto the HD layer, summed across all presynaptic COMB cells connected to cell *i*. It consists of $r_j^{COMB}(t-\Delta t)$, the firing of presynaptic COMB cell *j* at time $t-\Delta t$, where Δt represents axonal conduction delay, and $w_{ij}^2(t)$, the synaptic weight from presynaptic COMB cell *j* to postsynaptic HD cell *i* at time *t*. It is important to note that unlike Experiment 1, this weight is updated during simulation, and does not remain at the same initial. The entire term is scaled by the factor $\frac{\phi_2}{C^{COMB\to HD}}$, which controls the overall strength of COMB synapses onto the head direction cell layer where ϕ_2 is

a constant and $C^{COMB \rightarrow HD}$ is the number of synapses each postsynaptic head direction cell receives from presynaptic COMB cells.

The firing rate $r_i^{\text{HD}}(t)$ of head direction cell i at time t is a sigmoid function of the activation level of cell i

$$r_i^{HD}(t) = \frac{1}{1 + e^{-2\beta^{HD}(h_i^{HD}(t) - \alpha^{HD})}} \tag{6.13}$$

where α^{HD} and β^{HD} are the sigmoid threshold and slope respectively. As this model is firing-rate based, no individual action potentials are simulated, and neuronal activity is represented as a spontaneous average firing rate.

All COMB cells receive from all ROT cells, via w^3 weights. The activation level h_i^{COMB} of COMB cell i is given by

$$\begin{aligned}\tau^{COMB}\frac{dh_i^{COMB}(t)}{dt} = &- h_i^{COMB}(t) \\ &- \frac{1}{N^{COMB}}\sum_j \tilde{w}^{COMB} r_j^{COMB}(t) \\ &+ \frac{\phi_1}{C^{HD\to COMB}}\sum_j w_{ij}^1(t) r_j^{HD}(t-\Delta t) \\ &+ \frac{\phi_3}{C^{ROT\to COMB}}\sum_j w_{ij}^3 r_j^{ROT}(t)\end{aligned} \tag{6.14}$$

where $h_i^{COMB}(t)$ is the activation level of COMB cell i at time t. The term $-h_i^{COMB}(t)$ represents decay, where in the absence of any input the activation of COMB cell i will fall to zero as determined by the time constant τ^{COMB}. The term $\frac{1}{N^{COMB}}\sum_j \tilde{w}^{COMB} r_j^{COMB}(t)$ represents inhibitory feedback within the COMB cell layer, summed over all presynaptic COMB cells *j*. $\tilde{w}^{COMB}$ is a global constant describing the effect of inhibitory interneurons, and N^{COMB} is the total number of COMB cells in the layer. The term $\frac{\phi_1}{C^{HD\to COMB}}\sum_j w_{ij}^1(t) r_j^{HD}(t-\Delta t)$ represents excitatory input from the HD cell layer back onto the COMB layer, summed across all presynaptic HD cells connected to cell *i*. It consists of $r_j^{HD}(t-\Delta t)$, the firing of presynaptic HD cell *j* at time $t-\Delta t$, where Δt represents an axonal conduction delay, and $w_{ij}^1(t)$, the synaptic weight from presynaptic HD cell *j* to postsynaptic COMB cell *i* at time *t*. It is scaled by the factor $\frac{\phi_1}{C^{HD\to COMB}}$, which controls the overall strength of HD synapses onto the COMB cell layer where ϕ_1 is a constant and $C^{HD\to COMB}$ is the number of synapses each postsynaptic COMB cell receives from presynaptic HD cells. The final term is $\frac{\phi_3}{C^{ROT\to COMB}}\sum_j w_{ij}^3(t) r_j^{ROT}(t)$, which represents ROT cell influence on the COMB cell layer, where

$\sum_j w_{ij}^3(t) r_j^{ROT}(t)$ is the summed firing of presynaptic ROT cells indexed by j at time t multiplied by the weight $w_{ij}^3(t)$ from presynaptic ROT cell j to postsynaptic COMB cell i at time *t*. This is scaled by the factor $\frac{\phi 3}{C^{ROT \rightarrow COMB}}$ where ϕ_3 is a constant, and $C^{ROT \rightarrow COMB}$ is the number of synapses each postsynaptic COMB cell receives from presynaptic ROT cells.

SImilarly to HD cells, the firing rate $r_i^{\text{COMB}}(t)$ of COMB cell *i* at time *t* is a sigmoid function of the activation level of cell *i*

$$r_i^{COMB}(t) = \frac{1}{1 + e^{-2\beta^{COMB}(h_i^{COMB}(t) - \alpha^{COMB})}} \tag{6.15}$$

Unlike Experiment 1, the synaptic weights are updated during training based on Hebbian learning rules. These rules take the same form regardless of whether connections are HD-COMB, COMB-HD, or ROT-COMB. The weight between presynaptic HD cell j and postsynaptic COMB cell i is updated according to a local associative Hebbian learning rule

$$\frac{dw_{ij}^1(t)}{dt} = k r_i^{COMB}(t) r_j^{HD}(t - \Delta t) \tag{6.16}$$

where w_{ij}^1 is the synaptic weight from presynaptic HD cell j with firing rate r_j^{HD} onto postysnaptic COMB cell i with firing rate r_i^{COMB}, and k is the learning rate constant, which determines speed of weight change.

Similarly, the weight between presynaptic COMB cell j and postsynaptic HD cell i is updated by the equation

$$\frac{dw_{ij}^{2}(t)}{dt} = kr_{i}^{HD}(t)r_{j}^{COMB}(t - \Delta t) \tag{6.17}$$

Finally, the weight between presynaptic ROT cell j and postsynaptic COMB cell i is updated by the equation

$$\frac{dw_{ij}^{3}(t)}{dt} = kr_{i}^{COMB}(t)r_{j}^{ROT}(t - \Delta t) \tag{6.18}$$

At each timestep all weights are also normalised after updating. This is achieved by ensuring that the square root of the sum of the squares of afferent weights, for a particular postsynaptic cell *i*, is limited to 1. The weight normalisation equation for HD to COMB weights is given as

$$\sqrt{\sum_{j}(w_{ij}^{1}(t))^{2}} = 1 \tag{6.19}$$

and for COMB to HD weights

$$\sqrt{\sum_{j}(w_{ij}^{2}(t))^{2}} = 1 \tag{6.20}$$

and for ROT to COMB weights

$$\sqrt{\sum_{j}(w_{ij}^{3}(t))^{2}} = 1 \tag{6.21}$$

6.2.1.2 Simulation protocol

At the beginning of simulation, the firing rates of all cells are set to zero. Connectivity from HD to COMB, from COMB to HD and from ROT to COMB is set. In these self-organising simulations, all weights are initialised with random values. This is achieved using the GNU C Library's *srand48()* and *drand48()*. Together, these functions generate non-negative random numbers drawn from a uniform distribution in the interval $[0.0, 1.0]$. Following random initialisation of weight values, all weight profiles are normalised. In the case of diluted HD to COMB connectivity, 25 connections for each COMB cell are randomly selected from the 500 possible presynaptic HD cells using the *gsl_ran_choose()* function.

Following this initialisation, training begins. A single training epoch is comprised of rotating and stationary periods. During the rotation period, the network undergoes a period of an HD activity being imposed by external visual input (as described for Experiment 1), and shifted through the HD ring at the target velocity. The firing rates of half of the ROT cells are set to 1, whilst the rates of the other half of the ROT cells have firing rates of 0. Following the rotation period, ROT cells switch states and there follows a stationary period. During this latter period, the activity packet moves to be centered on each HD cell in turn for a fixed period of time. The amount of time spent during the rotation period and the stationary period is kept equal, such that the amount of time spent rotating is 500 (the number of discrete HD cell locations) times the amount of time spent stationary at each location. This is to ensure equal training with the agent rotating and stationary. This process constitutes a single epoch of

training, and can be repeated as much as is required.

After training, the four second testing period begins. External visual input and feed-forward inhibition are both removed. In the first second of testing, the HD packet remains where it ended up after training, and those ROT cells active during stationary training are active. This holds the HD activity packet at the current location. During the next two seconds, ROT cells switch states (i.e., those signalling rotation become active, whilst stationary ROT cells are quiescent), and the HD activity packet is observed. In the final second of testing, ROT cells switch states once more, and the HD activity packet is held at its final location.

6.2.2 Results

This section demonstrates the results of training and testing the self-organising model as described above. Simulation parameters are as given in Table 6.4. Figure 6.12 shows the HD layer firing rates and position during testing. As we can see, the self-organising network has made a good job of replicating the results of a pre-wired network, with approximately two stationary phases and a rotation phase, with the switch between the two phases governed by the activity of ROT cells. This demonstrates clearly that model functionality can be self-organised, with a single network being able to smoothly transition between holding and shifting an activity packet. However, it should be noted that these results are not perfect. During the first second of testing, there is some activity packet drift. Additionally, the top speed achieved during head rotation is $164.4\,^{\circ}/s$, which represents 91% of the target velocity. Whilst this is an improvement on previous path integration models relying on within-layer recurrent collateral connections, it falls short of the ideal result achieved in the pre-wired model of Experiment 1.

Table 6.4: Network parameter values for simulations reported in Experiment 2. All simulations, except where explicitly stated, are run using these parameters. All measures of time are given in seconds.

Network Parameters	
Head Direction Cells	500
COMB Cells	1000
ROT cells	500
Rotation Time (per Epoch)	$12.0s$
Stationary Time (per HD cell)	$0.024s$
Epochs	50
$C^{COMB \to HD}$	1000
$C^{HD \to COMB}$	25
ϕ_1	2400.0
ϕ_2	10000.0
ROT Cells	1000
$C^{ROT \to COMB}$	1000
ϕ_3	5000.0
α^{HD}	0.0
β^{HD}	0.2
α^{COMB}	113.0
β^{COMB}	0.3
τ^{HD}	$0.0001s$
τ^{COMB}	$0.0001s$
$\tilde{w}^{HD}$	1.0
$\tilde{w}^{COMB}$	1000.0
δt	$0.0001s$
λ^{HD}	450.0
σ^{HD}	20.0°
I^{FF}	420.0
Target Velocity	$180\,^\circ/s$
k	0.1
Δt	$0.01s$

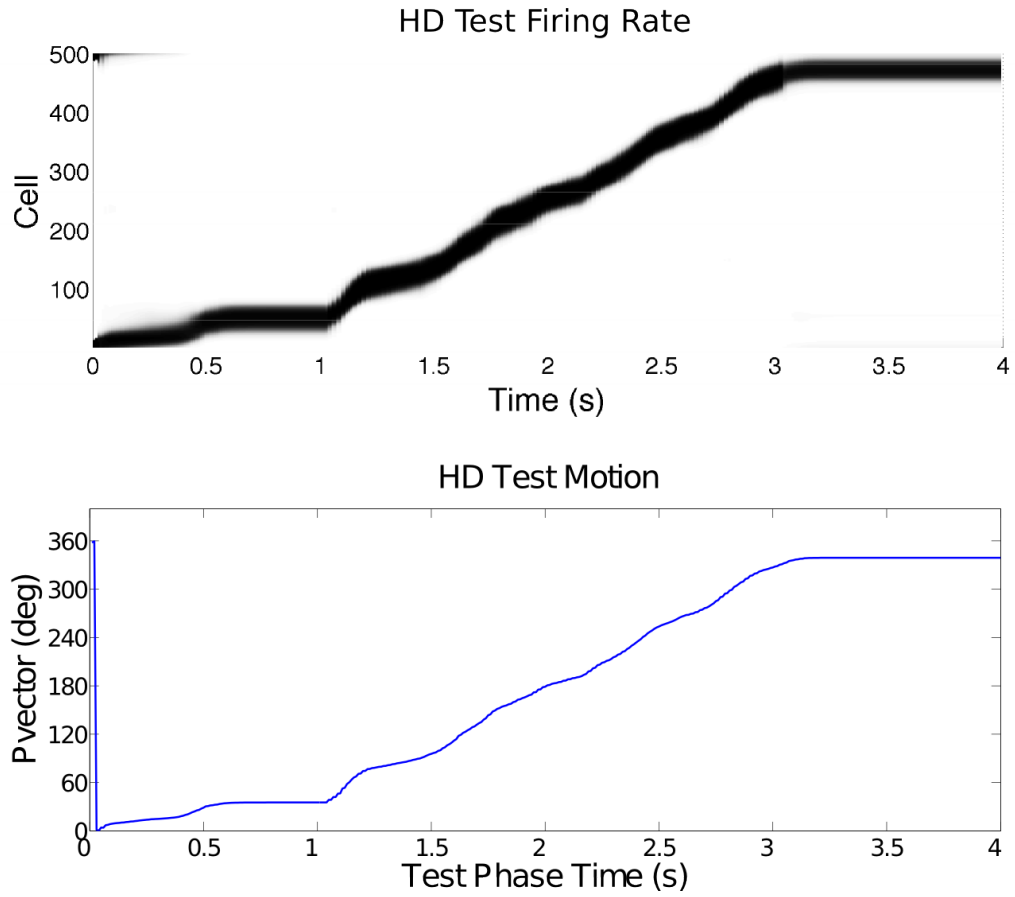


Figure 6.12: Simulation results from the self-organising model. HD layer firing (top) and packet location (PVector, bottom) during testing phase of simulation. Model performance is approximately correct, with a 2s period of packet motion during path integration, bookended by two 1s periods of no packet motion whilst HD does not change.

Model function will be sub-optimal if either or both of two facts becomes true: If COMB cells are active when they are not supposed to be and/or if COMB cells do not stimulate the correct region of HD space. However, it should be remembered that this latter fact is related to the former. If a COMB cell is active for multiple regions of HD space, then a single COMB projection will be correct for one space, and not for the other. As the number of COMB cells for which either fact is true increases, model performance is expected to degrade.

Observing the learned weights can help us to understand why the model does not achieve perfect

functionality, as it tells us which regions of HD space and which types of head motion drive COMB cells. Figure 6.13 shows weight profiles for COMB cell 125, with final learned weight profile (blue) overlaid on initial pre-learning random weights (red). The top row shows the weights received from the HD layer, the middle row shows the weights projecting back to the HD layer, and the bottom row shows the weights received from the ROT cell layer. These profiles are always from the perspective of a given COMB cell, since it is the COMB representation that is critical to model functioning.

Here we can see that, after learning (blue weights), this COMB cell receives projections from a specific region of the HD layer. This COMB cell also projects back to a specific region of the HD layer with a similar Gaussian profile. We can therefore say that this COMB cell has good HD specificity. In the middle plot showing COMB to HD weights, we see a vertical black line, which represents the overall direction of the w^1 weights this COMB cell receives from the HD layer. We can see therefore that this COMB cell projects back to the HD layer with an offset compared to the region of the HD layer from which it receives. Despite promising HD cell specificity for this cell in line with the hypothesised weight profiles required for model operation, ROT specificity is poor. We can see in the bottom plot, showing receiving weights for this COMB cell from the ROT layer, that COMB cell 125 is not receiving exclusively from ROT cells either responding to head rotation or lack of rotation. Instead, this distribution of receiving weights, with the COMB cell receiving approximately equal projections from ROT and NOROT cells, suggests that this particular COMB cell is almost equally driven up by either sort of head motion. Since we see from the weights to and from the HD system that this COMB cell projects back with an offset, if this cell were active when the head was stationary it would provide asymmetric drive to the HD layer and cause an activity packet to drift. Cells such

as this could be the reason for the HD activity packet drift during the first, stationary, second of testing.

Similarly to Figure 6.13, Figure 6.14 shows weights for COMB cell 475. This COMB cell does not have good HD layer specificity, both receiving from and projecting to two distinct Gaussian packets. We do, however, see that this cell has good ROT specificity, responding only to those cells active for head rotation (the last 250 cells). Cells such as this could, during head rotation, project conflicting information back onto the HD layer, possibly affecting packet speed during path integration. It should be noted, however, this is not definitely the case. Although all weight profiles are not pictured for the sake of brevity, the majority of COMB cells appear to respond specifically to ROT cell input in the manner of cell 475, responding to either head rotation or no head rotation, but project and receive from multiple less specific regions of the HD cell layer. An example of this is shown in Figure 6.15 which, similarly to Figures 6.13 and 6.14, shows weights for a specific COMB cell, 925. This cell responds to the first 250 ROT cells, specific therefore to no head rotation, but with noisy weight profiles to and from the HD layer.

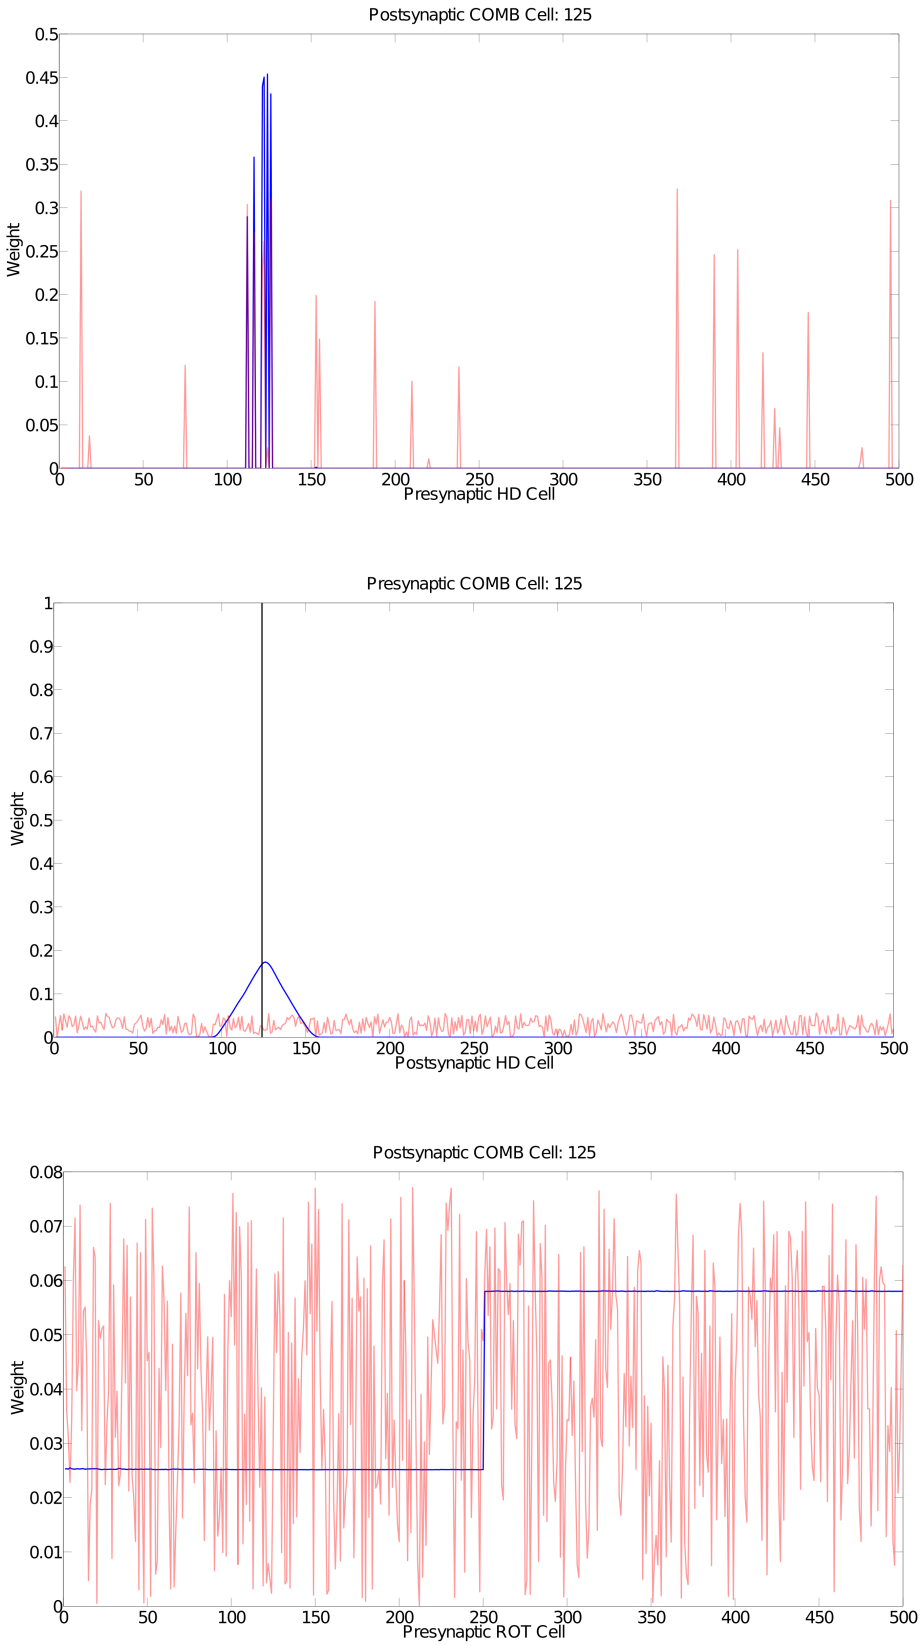


Figure 6.13: Learned synaptic weights in the self-organising model. HD to COMB (top), COMB to HD (middle), and ROT to COMB (bottom) weights for COMB cell 125. Final weight profiles in blue, initial (pre-training) weights in red.

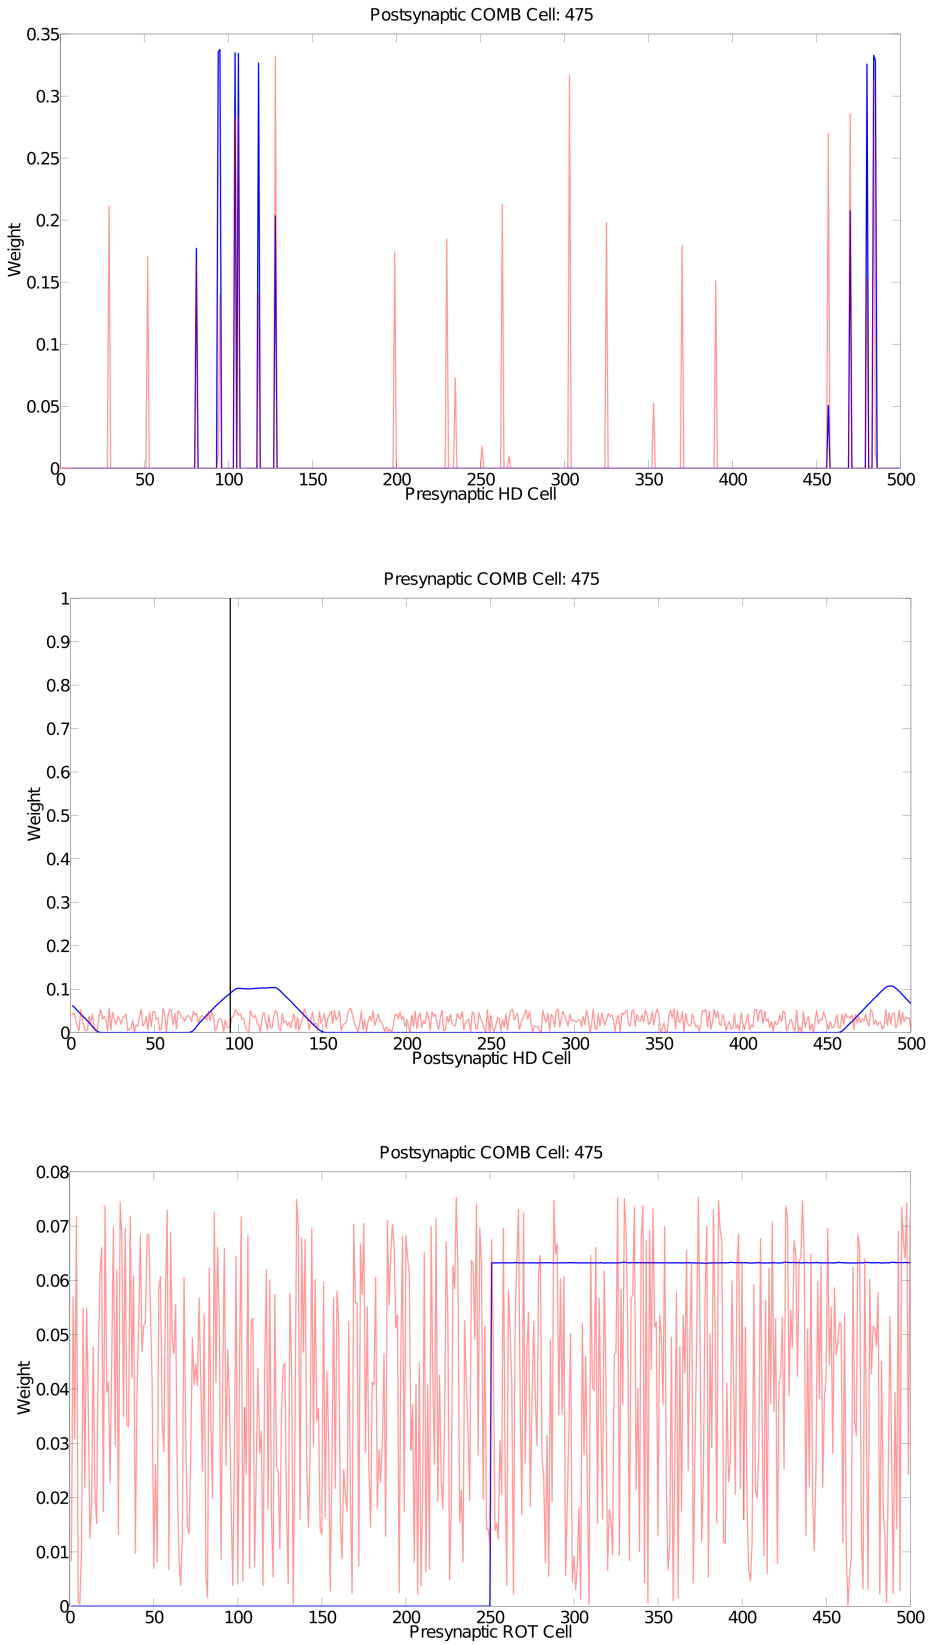


Figure 6.14: Learned synaptic weights for cell 475 in the self-organising model, conventions as in figure 6.13.

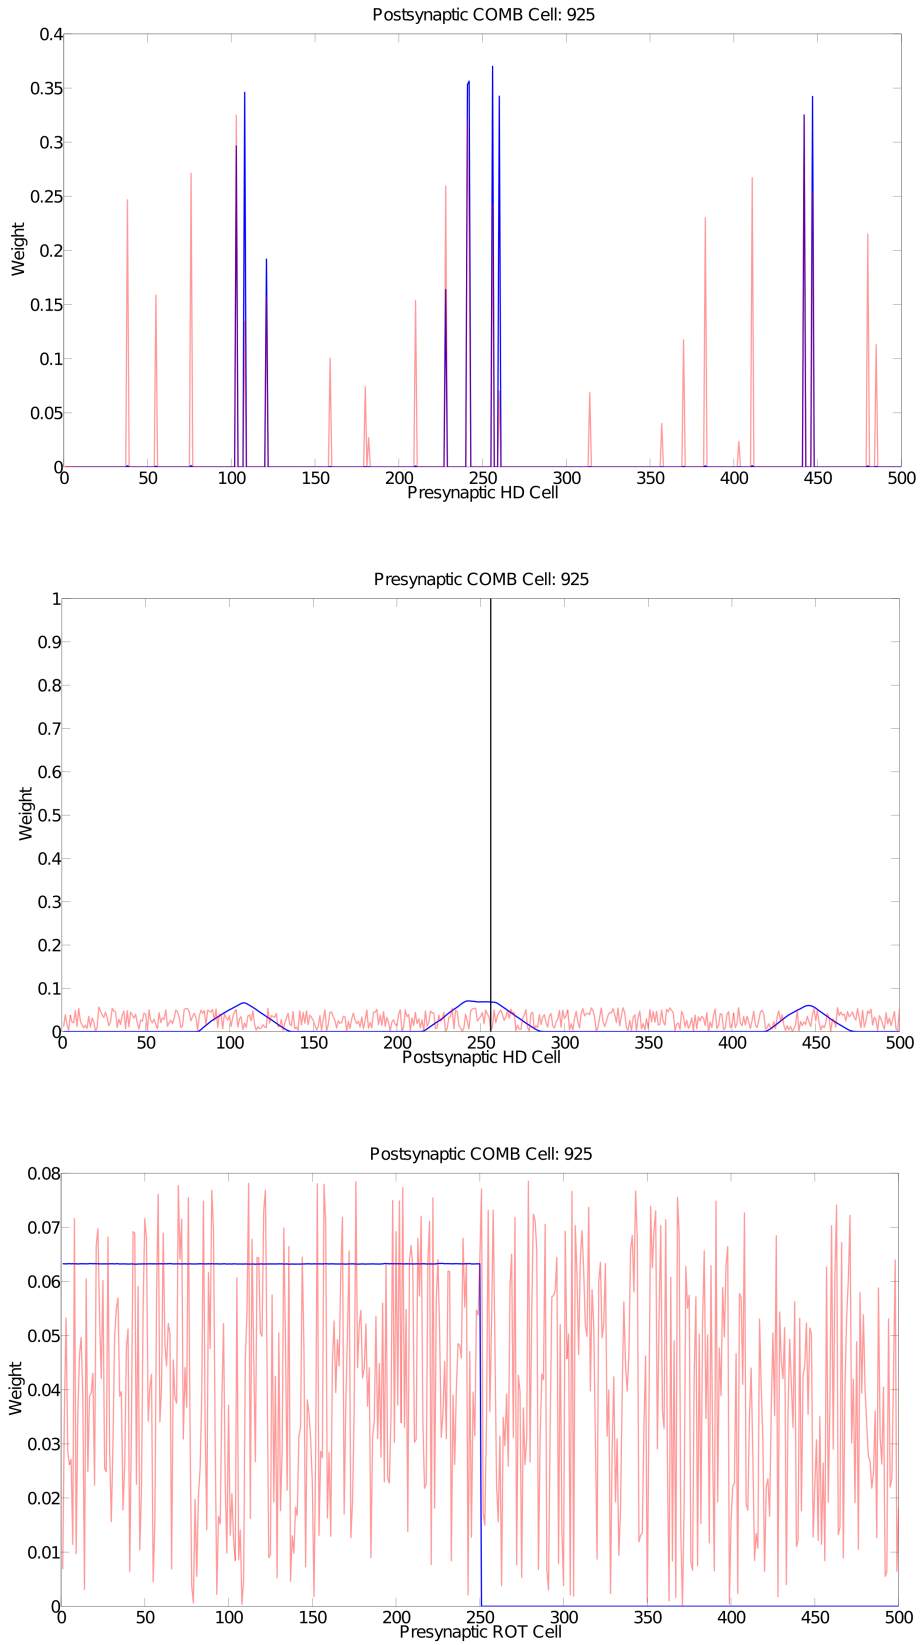


Figure 6.15: Learned synaptic weights for cell 925 in the self-organising model, conventions as in figures 6.13 and 6.14.

We should consider the activity of COMB cells from three different perspectives. Firstly, is a given COMB cell active for a specific region of HD space? Secondly, is its activity specific to a particular head rotation speed? Finally, does it stimulate a specific region of HD space?

To help unpack the first of these questions, Figure 6.16 plots HD activity packet position (Pvector, blue) overlaid with COMB cell firing rate (red) for COMB cells 125, 475, 925, and 850. Clearly we can see that whilst cells 125, 475, and 925 have good HD specificity, COMB cell 850 does not. In order to assess the specificity of the COMB cell layer as a whole to HD, Figure 6.17 is a histogram showing the number of COMB cell firing peaks (i.e., contiguous periods of elevated firing) during the testing phase. Given that the HD activity packet does not make a full movement through $360\,^{\circ}$, COMB cells would only be expected to have one firing peak if they were specific to HD. Here instead, we see that COMB cells have any number of peaks between 0 and 4. Thus, we must assume that a good number of COMB cells are firing when they should not be (i.e., to more than one HD).

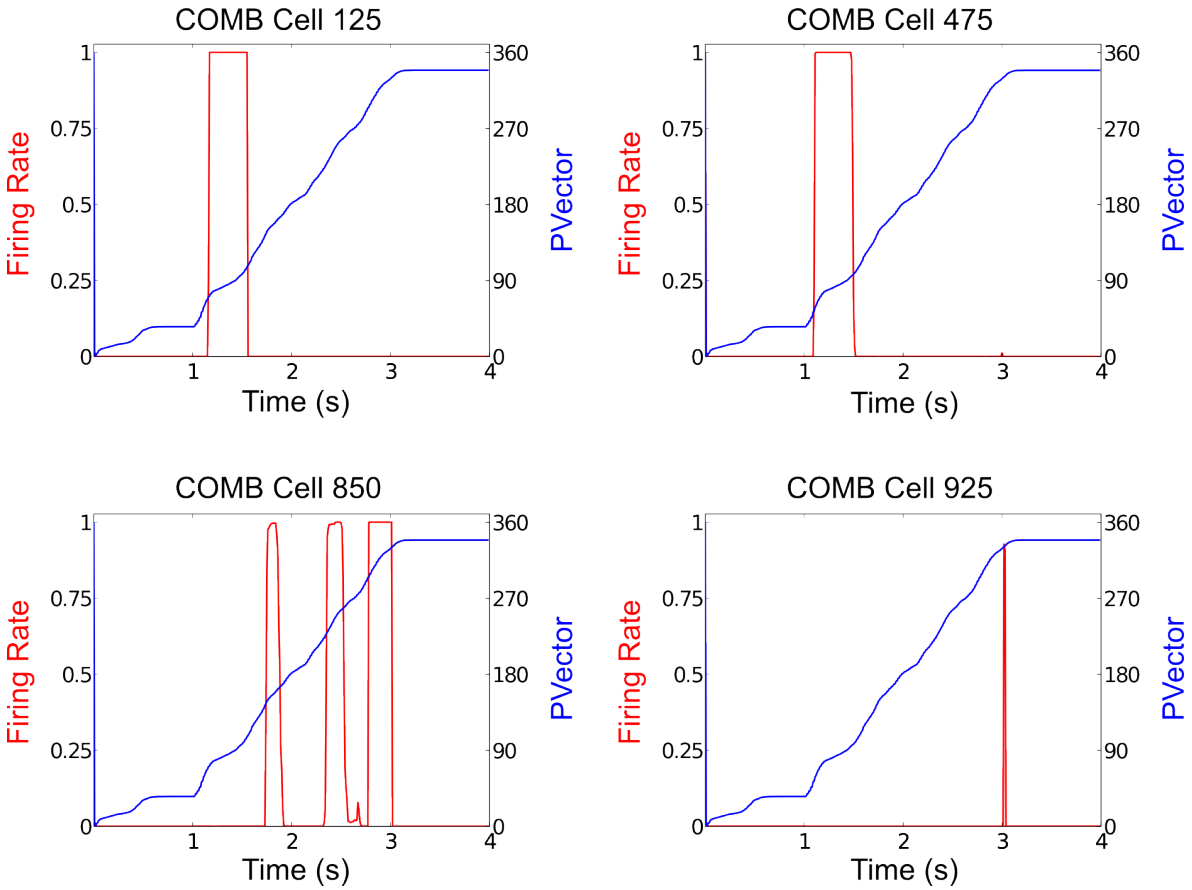


Figure 6.16: Firing rate (red) for COMB cells 125 (top left) 475 (top right) 850 (bottom left) and 925 (bottom right). HD layer Pvector is overlaid (blue) in order to determine to which regions of HD space each COMB cell responds.

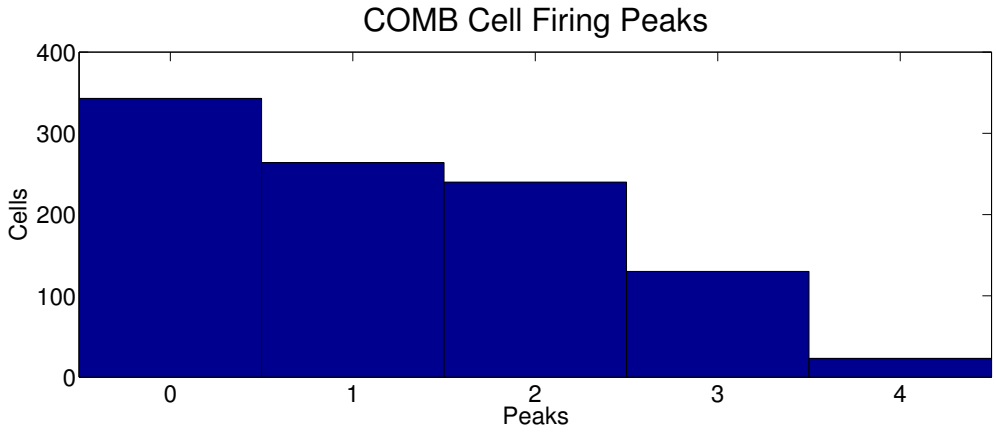


Figure 6.17: Histogram detailing the number of firing rate peaks (contiguous periods of elevated firing, such as those seen in figure 6.16) for each cell in the COMB layer during testing.

The second of these criteria, whether COMB cells are specific to head rotation or no head rotation, cannot be assessed from COMB activity during testing. Whilst the test phase of simulation does involve a rotation through just less than $360\,^{\circ}$, the HD packet is only held stationary at two distinct locations. It is therefore difficult to assesss COMB specificity to ROT signals by observing COMB cell activity during test phase alone. In order to do so, a ROT index statistic was constructed, whereby for each COMB cell, the proportion of the total ROT to COMB weights for that cell was calculated for the first 250 ROT cells (i.e., those ROT cells firing when the head is not rotating) and the second 250 ROT cells (i.e., those ROT cells firing when the head is rotating). The ROT index is then the proportion for head rotation minus the proportion for no head rotation. Thus, each COMB cell has a ROT index ranging from $+1$ (all weights from ROT cells firing to head rotations) to -1 (all weights from ROT cells firing when the head is stationary). A ROT index of 0 would indicate no specificity to head rotation.

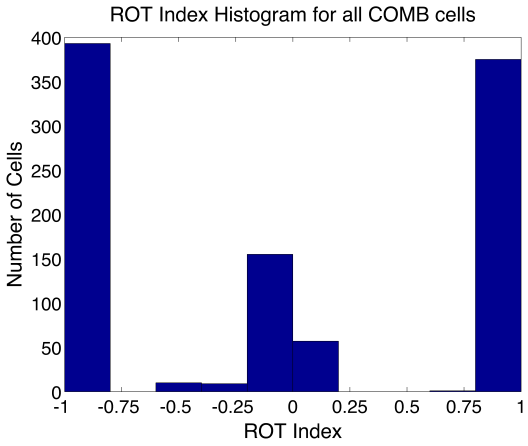


Figure 6.18: Histogram detailing which part of the ROT cell layer COMB cells receive projections. A ROT index of $+1$ indicates the majority of weights for that cell comes from ROT cells signalling head rotation, whilst a ROT index of -1 indicates the majority of weights for that cell comes from ROT cells signalling a stationary head. ROT index values in between this range indicate mixed ROT cell specificty for a given COMB cell. This figure demonstrates that, whilst the majority of COMB cells are specific to head rotation, there is a noticeable proportion which are not.

Figure 6.18 shows a histogram of ROT index values for the trained COMB cell layer. 768 of the 1000 COMB cells (76.8%) had an absolute ROT index greater than 0.8. Of these cells, approximately equal proportions were specific to head rotation (37.5%) or no head rotation (37.5%). However, the majority of the remaining 23.2% of COMB cells had a ROT index in the range -0.6 to 0.2, indicating that COMB layer specificity to head rotatiation, whilst quite good, was not perfect.

The third of these criteria, whether COMB cells stimulate a specific region of HD space, is more difficult to assess. Figures 6.13, 6.14, and 6.15 do tell us about COMB to HD weights (cell 125 has good specificity in this regard, whilst the other cells do not). However, weight profiles are of limited value in demonstrating how specifically COMB cells drive the HD layer. This is because the effect of COMB input on HD packet motion depends on exactly where a given COMB cell projects to with respect to the current HD packet location. For example, if a COMB cell projects two (or more) distinct localised synaptic weight packets onto the HD layer, only the localised weight components which interact directly with the current HD packet location will affect the motion of the HD activity packet. The other projections will have an effect on the HD layer in a distance-dependent manner: the closer the center of mass of these projections to the HD packet location, the greater the effect on HD packet motion. The complexity of this response is further increased by the fact that a given COMB cell is quite capable of firing for multiple regions of HD space, and therefore the suitability of the COMB cell response will vary depending on how it is driven up (i.e., at which HD does a given COMB cell become active?). Yet further increases in complexity come when it is realised that the influence of a given COMB cell on the HD layer does not exist alone, but rather in the context of all the other COMB cells influencing the HD layer at that time. To investigate this, we can look at COMB to HD weights

across the population of COMB cells. Figure 6.19 shows a histogram of the number of peaks in the w^2 synaptic weight profile projected onto the HD layer by a given COMB cell, recorded over the entire population of COMB cells. This clearly shows that for many COMB cells, the weight profile has multiple peaks, suggesting that many of the COMB cells will have an influence on the HD cell layer that is not specific to a single region of HD space.

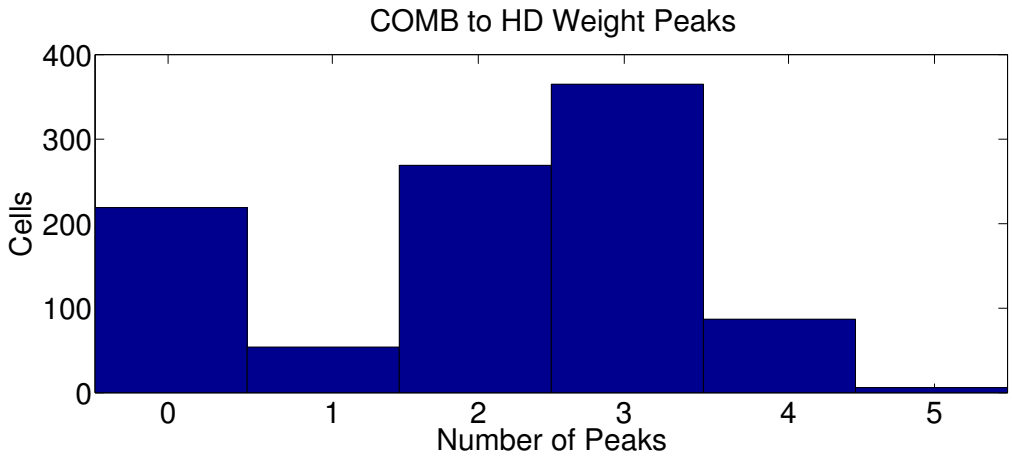


Figure 6.19: Histogram showing the number of COMB to HD weight profile peaks, across the COMB cell layer. It is observed that a given COMB cell may project a weight profile containing anywhere from 0 to 5 peaks onto the HD cell layer.

6.3 Discussion

This chapter presents operational details and simulation results from a head direction cell path integration model which does not contain within-layer symmetric recurrent collateral weights. The model in this chapter consists of a ring of HD cells, reciprocally connected to a ring of COMB cells, with the latter also receiving input from ROT cells. All HD to COMB and COMB to HD connections contain a delay Δt. COMB cells which receive input from ROT cells signalling no head rotation project back to the HD cell from which they maximally receive. These cells are responsible for sustaining a stationary HD layer activity packet when the head is not moving. COMB cells which receive input from ROT cells signalling head rotation at the target velocity project back onto the HD layer with an offset determined by the target velocity and the conduction delay. In this manner, direct HD to HD connections are no longer present, and instead all HD cells influence one another via the COMB cell layer regardless of head motion.

In Experiment 1, it is found that a network pre-wired in the manner described above is capable of performing extremely accurate ($> 99\%$ of target velocity) path integration at a range of target velocities in the interval [0-360 $^\circ/s$], as well as sustaining a stationary activity packet when the head is not rotating. This finding extends to the more biologically-plausible situation in which conduction delays in COMB to HD and HD to COMB connections are randomly varied within a set range. A curious oscillatory effect is discovered across layers, but is only present in the model in which a single value of conduction delay is used, and is thus not thought to be a major failing of the system. In the model using a distribution of conduction delays, there is a very slight drift of HD activity packet location when the head is stationary. It is thought that this is due to the effect of non-uniform sampling of

conduction delay values.

This finding is an important one, both in terms of the work in previous chapters of this thesis and in terms of HD cell models in general. In Chapter 4, an isomapping mechanism was shown to depend on time-accurate path integration, and in Chapter 5 path integration speed was shown to be negatively affected by the presence of within-layer HD to HD connectivity. This chapter completes the picture, demonstrating how an architecture which no longer incorporates such connectivity is capable of extremely accurate path integration. The implications of this finding are discussed in greater detail in the discussion of this thesis.

In a replication of the results of the previous chapter, rise time is found to have a negative effect on path integration speed, and this is again seen to be in the context of a particular conduction delay. One other notable observation from varying conduction delay is that the system is tolerant of a broad range of values of Δt, including very short conduction delays. This is in contrast to results in models which make use of within-layer HD to HD connections (Walters and Stringer, 2010; Walters et al.,2013). This is likely because when shorter conduction delays are employed, a given HD cell will project back to the HD layer with a smaller offset. Depending on the spread of recurrent collateral weights, this could mean that this input is not offset enough to overcome the slowing effect of recurrent collaterals. In this newer model, within-layer recurrent collaterals are no longer present and it is therefore no longer possible that recurrent weights interfere with path integration is this offset-dependent manner.

In Experiment 2, the same model is now initialised with random weights, with the full weight profile being self-organised via Hebbian learning during training. This model performs well, and is able to both sustain an activity packet at a single location when the head is not rotating and shift this activity packet when the head is rotating. However, performance is not at as high a level as the pre-wired network: the top speed achieved is $164.4\,^{\circ}/s$ (91% of the target velocity $180\,^{\circ}/s$). Additionally, there is some minor HD activity packet drift when the head is stationary. Regrettably, time constraints due to long simulation runtimes have limited the exploration of the parameter space of this self-organising model.

What might be the reason for sub-optimal performance in the self-organising model? We can think of performance in terms of COMB specificity to HD (both in receiving and projecting weights), as well as in terms of COMB specificity to different speeds of head rotation. Examination of the weight profiles for individual COMB cells demonstrates that poor performance is in large part due to a loss of COMB specificity with respect to HD. Most COMB cells receive from either the first 250 ROT cells (signalling head rotation) or the last 250 ROT cells (signalling a stationary head), and in this sense COMB cells have correctly learnt to be specific to one type of head motion. However, most COMB cells do not respond to a specific HD, nor do they project back to a specific HD. Indeed, COMB cells can fire for multiple regions of HD space. Depending on the projections from a given COMB cell to the HD layer, this can mean that individual COMB cells influence the HD layer in a manner not consistent with desired path integration speed. Whether the head is rotating or not, this can act to change activity packet speed in an undesired manner, either making the HD layer activity packet move too fast or too slow relative to the target velocity, or moving the packet when it should be stationary. By

examining COMB to HD weights, we see that there are a number of COMB cells projecting back onto the HD layer with weight profiles containing multiple peaks, thus stimulating multiple HD regions. We can see therefore that we have a twofold problem: COMB cells become active when they should not be active, and COMB cells drive up regions of the HD layer which should not be active.

Given the lack of specificity of COMB cell projections to or from the HD layer, it is interesting that the model performs as well as it does. This implies that accurate packet motion can be achieved by a COMB cell layer without perfect HD specificity. The system might function based on the average direction of the weight profiles, with the concordance between these directions determining HD layer offset. However, inspection of the vectors of receiving and projecting weights for COMB cells (not pictured) did not reveal any specific organisation on these terms. More likely is the possibility of a complex interaction of the specific structure of receiving and projecting weight profiles with current HD layer packet location. Put simply, if a COMB cell stimulates an area of the HD layer which is not close to the current HD cell activity packet, this stimulation might be ignored as it it likely not enough to overcome the strength of the inhibitory effect which the currently-active HD cell area exerts on all other HD cells. It is therefore possible that the specific configuration of the weights in combination with HD layer lateral inhibition prevents all but the largest perturbations to normal (i.e., accurate) HD layer packet motion. Conversely, if a COMB cell projects multiple packets of weight onto the HD layer, and these packets are close together, this will negatively effect packet motion.

One question that remains unanswered in the self-organising model is why are so many COMB cells unresponsive? Figure 6.17 shows that a number of COMB cells do not show firing peaks at any stage

during testing. This is likely because these cells have not been driven up during training, and thus have not formed receiving or projecting weight profiles structured by learning. COMB cell firing during training is governed by a complex interaction of a number of parameters, including the strength of ROT cell input, the strength of HD cell input, COMB cell firing threshold and COMB layer lateral inhibition. What is curious however, is that the model still operates with these cells not active during training. Thus it is possible the COMB layer representation is not as simple as that hypothesised and demonstrated in the pre-wired network. It is not possible for so few COMB cells to represent every HD cell location in combination with both zero head rotation and with head rotation at the target velocity. This suggests that it might not be necessary for every individual HD cell to be represented by the COMB layer in this self-organising model. Provided that some COMB layer stimulation is correctly providing asymmetric input to the current HD cell activity packet location, there might well be enough driving input to shift the HD layer activity packet during rotation. When the head is stationary, the packet would, however, need symmetrical stimulation in order to remain stationary. Therefore, irregular input from the COMB layer would result in a drifting activity packet. This effect is perhaps responsible for slight activity packet drift seen in simulation during the first second of test phase. It is also possible that this drift, which is in the same direction as target packet motion, is an artefact of training and testing on one direction of rotation only. It should be noted that test phase does not attempt to hold the activity packet stable in every location around the HD cell layer. For this reason, it is not possible, based on these results, to fully assess the performance of the model when the head is stationary at every possible location, and there is potential for activity packet drift to vary across a layer.

The models in this chapter have several limitations in their design, in order to simplify operation enough to gain insight without being lost in complexity. In particular, the firing of ROT cells is greatly simplified when compared to AHV cells in the brain. AHV cells have a roughly linear firing rate response across a range of head velocities, and vary with their response to head rotation of different directions (for full details of AHV cell responses and their potential importance, see the discussion of this thesis). The ROT cells is this thesis fire in binarised manner, either firing at maximum for the preferred velocity of head rotation, or at zero otherwise. A natural extension of the work in this chapter would therefore be to have ROT cells more accurately reflect AHV cell firing. Given a more biologically accurate ROT cell layer with a less caricatured response to head rotational velocity, it would be interesting to see whether this network could still function. It would also be of interest to see what influence a population encoding in the ROT cell layer would have on the representations formed in the COMB cell layer. It is notable that the COMB cells in this thesis are also simplified relative to biology, in that they respond to a single HD and a single speed of rotation. In the brain, COMB cells show elevated firing across a range of head directions, and this is modulated in a complex manner by head rotation. It would be of interest then, to see if realistic COMB cell representations follow from realistic ROT cell representations. The current model currently deals only with two rotation speeds at a time, with one being zero, and in a single direction. It is possible that a richer ROT/COMB cell representation would allow the system to function across a range of angular head velocities, and it would be interesting to see how the system copes with this increased challenge, and what sort of weight profiles are required.

The simulation protocol used by this network is also a simplification. Naturally, any expansion of

the model to cope with multiple speeds will require training with multiple speeds and in both clockwise and anticlockwise directions. The best way to make a direct comparison with rodent HD cells is to use realistic training data, by replicating the specific manner in which rats explore an environment. Roughly, rats tend to establish a home location, and their exploration of an environment can be split into more circuitous outward paths followed by a more direct return journey. During outward segments, there is an inverse relationship between curvature of path and speed of locomotion, such that the faster the rat travels, the straighter the path they take (Wallace et al., 2006, 2010). This indicates that there is a possible relationship between the specific structure of movements and path integration functioning, and that we should not be ignoring the precise means by which rats explore their environment. The potential significance of this movement structure, and how it breaks down as error in path integration accumulates, is treated in more depth in the discussion of this thesis. Also covered in the overall discussion is the question of how much self-organisation is essential for a network to undergo in order to replicate biology, and this question is therefore not replicated in the current discussion. Whilst questions of training regime are still open, the model should be expanded to function for realistic head movements. At the very least then, the testing protocol should aim to maintain path integration accuracy during complex head movements over a simulated period of realistic locomotion behaviour. This will naturally involve head rotations of multiple speeds in both directions, and periods when the head is stationary whilst facing several different allocentric directions.

The models in this chapter employ a simplified version of inhibition. At each timestep, inhibition within a layer is a uniform constant $\tilde{w}^{HD}$ or $\tilde{w}^{COMB}$, scaled by the firing of all cells in that layer. Thus inhibition is dynamic but not regionally-specific. This mechanism of inhibition is a fair approx-

imation of the global effect of inhibitory interneurons within a layer. Also employed during training only is constant feedforward inhibition I^{FF} acting on the HD layer. This describes the effect of feed-forward inhibition, whereby the presence of visual input increases the amount of HD layer inhibition by stimulating directly HD layer inhibitory interneurons, without the need for a computationally expensive model describing individual interneurons. Although these approximations are considered fair, particularly from the perspective of making simulations computationally viable, it would be interesting to model inhibition in more detail. This is particularly important given that the COMB cell representation which self-organised is not as expected. It might be that a more biologically detailed inhibition model would allow specific COMB cell representations to form more easily. The interested reader should refer to the discussion of the previous chapter, which covers inhibition models in more detail.

Based on the known physiology of the HD cell system (see Chapter 1 for full details), the model in this chapter should be considered in terms of DTN and LMN, and thus we must attempt to match the connectivity of these regions. Roughly speaking, the HD region could be mapped onto the LMN, whilst the COMB cell region could be mapped onto the DTN. Unlike the connectivity in each of this chapter's models, which is uniformly excitatory, DTN to LM connections appear to be inhibitory. In order to have strict biological accuracy, this model should therefore be extended to cover inhibitory connectivity from COMB cells back onto the HD layer. One potential issue with this change, particularly for self-organisation, is the idea of updating weights which are inhibitory. One solution to this issue is the concept of anti-Hebbian learning, whereby weights are increased based on anticorrelation between pre- and post-synaptic firing rather than correlation (Koch et al., 2013; Vogels et al., 2013).

Future work should incororpate inhibitory connectivity and an anti-Hebbian learning rule of this kind.

To summarise, this chapter has presented a two-layer model capable of sustaining a stationary activity packet, and shifting an activity packet using an internal head rotation signal with a high degree of accuracy. Within-layer recurrent collaterals are not included within the HD cell layer, and all HD cells influence each other via the COMB cell layer. This influence can either be symmetrical when the head is stationary, or with an offset when the head rotating. The switch between offset and non-offset connectivity is governed by the firing of ROT cells, with non-offset COMB cells being sensitive to combinations of HD and no rotation, and offset COMB cells being sensitive to combinations of HD and head rotation at the target velocity. This model is shown to be able to accurately shift an activity packet at greater than $> 99\%$ of a given arbitrary target velocity. A self-organising version of the model is also presented, which is capable of less accurate (91% of target velocity) path integration. It should be noted, however, that even this self-organising model is an improvement on previous self-organising path integration models. The reasons for path integration inaccuracy are discussed in relation to how the self-organised model differs from the pre-wired model. Several points are discussed, with the aim of helping to improve future iterations of the model. This work represents a useful contribution to our theoretical understanding of the rat HD cell system, and provides a platform from which to investigate path integration in more detail.

Chapter 7

Discussion

7.1 Summary of findings

This thesis has been focused on computational modelling of Head Direction (HD) cells, which respond to allocentric head direction in the horizontal plane. In particular, self-organising rate coded models have been employed. In most cases, these models were limited in scope, presenting the simplest possible architecture to provide an explanation for a given HD cell phenomenon. In many cases, this has allowed detailed parameter exploration to further understanding of model operation. In brief, this thesis covers two main topics: remapping of visual landmark information onto the HD system following cue conflict, and path integration speed.

HD cells are known to respond to both externally-derived (allothetic) and internally-derived (idiothetic) sources of information. Chapter 3 focuses on understanding how conflict between these two information sources is resolved and explores how visual landmarks are remapped onto the HD cell system following environmental manipulations. Previous work has suggested that conflicts between visual and idiothetic information sources might be always resolved in favour, although incompletely,

of visual landmark information (Goodridge and Taube, 1995). However, later studies have indicated that idiothetic-allothetic distinction s are based on the discrepancy between these two information sources (Knierim et al., 1998). The simulations in this chapter are aimed directly at explaining data from the Jeffery lab (Knight et al., 2014) investigating the effect of discrepancy size on cue conflict resolution. Knight and colleagues recorded HD cells from rats exploring a paradigmatic circular arena containing a single light as a visual landmark cue. Following exposure to this condition, allowing the HD system to anchor itself to available landmark information, the light was extinguished briefly, before another light was turned on at a specific angular separation from the original. The separation between the two light cues directly reflected the conflict being introduced between idiothetic and visual information. It was found that when this discrepancy was small, visual information had incomplete dominance. However, at a critical separation (approximately 120°) the system switched to relying more on idiothetic information, with idiothetic dominance increasing with conflict size. A secondary finding of Knight and colleagues was that visual landmark dominance was dependent on whether the light cue had previously been experienced in conflict with idiothetic information. HD cells recorded from naïve rats were found to rely on visual landmark information even at large conflict sizes. This is taken to mean that, in line with previous HD cell recording studies (Knierim et al., 1998), visual landmarks can be taken as more or less salient based on their relative stability.

Chapter 3 presents a CANN neural network capable of replicating Knight et als (2014) data. The model consists of a central CANN HD cell network, which receives topographic connectivity from a ring of visual cells representing the egocentric bearing from the simulated agent to a visual landmark cue in the environment. These connections are modifiable by Hebbian learning. During simulation,

an activity packet is generated in the visual ring, which in turn generates an activity packet in the HD ring. This represents the influence of the initial landmark cue. A second activity packet is generated in the visual ring, at a specific angular separation from the first. This represents the second landmark cue and the introduction of cue conflict. Input from this second visual activity packet shifts the HD activity packet. However, Hebbian learning maps this second visual packet onto every intermediate HD location, slowing HD activity packet motion and eventually causing it to come to rest in an undershot location relative to the visual packet. The final HD packet location is a function of the strength of two opposing forces: the drawing effect of input from the visual packet, and the slowing effect of dynamic remapping of the visual to HD relationship. Simulating at a variety of conflict sizes showed that this model was capable of replicating experimental findings. At small conflict sizes, the final HD packet location is determined principally by second visual packet location, but with an undershoot due to remapping. As conflict size increases, the effect of remapping begins to outweigh the weaker influence of a visual cue further from the HD packet, from at a critical point onwards the final HD packet locations is determined principally by the initial HD location. The effect of salience was modelled with a uniform weighting parameter on visual to HD weights, providing a uniformly strengthened or weakened influence of visual information on final HD packet location, and replicating findings on the effects of previous cue conflict experience on final HD packet location

The angular separation between preferred firing directions of HD cells is known to remain constant across multiple situations. Chapter 4 extends the work of Chapter 3 to explain how this property of HD cells, known as isomapping, can be sustained after preferred firing direction changes following cue conflict. Following the simulations in Chapter 3, a small proportion of HD cells, those active

during cue conflict resolution, change their preferred firing directions by virtue of their changed relationship to visual information; as reflected by updated visual to HD weights. However, different HD cells in this region will have different changes in preferred firing direction based on the extent of the changes in the receiving visual weight profile. Other HD cells not involved in cue conflict will not change preferred firing directions at all. This represents a break down in isomapping: the angular separation between preferred firing directions of adjacent HD cells is not maintained across cue conflict situations.

The model in Chapter 4 is identical to that used in Chapter 3. However, added to the simulation protocol is a period of the agent navigating in the light. Once cue conflict has been simulated, as described in Chapter 3, the HD and visual activity packets move through their respective rings at the same speed, with Hebbian learning taking place in visual to HD connectivity. This represents a period of simplified navigation, with the agent rotating on the spot. HD activity is assumed to be shifted at the correct speed by the output of path integration mechanisms. This output is not modelled in detail, and it is assumed that speed-accurate path integration is achieved. Given that the simulations in Chapter 3 end with a mismatch in the locations of the HD and visual activity packets, and that these two activity packet shift at the same speed during rotation, the new relationship to vision is learnt by all HD cells in the ring. Thus, isomapping is achieved. Simulations show that this model is capable of resolving cue conflict in the manner observed experimentally, with HD cells all shifting their preferred directions in a manner dependent on conflict size. A full parameter investigation revealed that this isomapping mechanism was largely independent of rotation speed during path integration, and the model was therefore able to isomap fully over a large range of rotation speeds. The effect of saliency

is again investigated as in Chapter 3, and the effect of saliency on final packet location is found to translate into an effect of changes in visual to HD weights. Further investigation of isomapping shows that path integration must be accurate in order to shift the HD activity packet at the same speed as the visual activity packet, maintaining a consistent relationship between HD cells and vision at all locations around the HD ring. By introducing a systematic error into path integration, it was demonstrated that when there is no stable relationship between the visual and HD activity packets, vision to HD weights do not settle into a stable profile and isomapping does not occur.

The isomapping mechanism described in Chapter 4 depends on speed-accurate path integration to preserve the relationship between HD cell and visual cell firing. It is demonstrated that isomapping will occur in such conditions, but does not model the path integration process in full. A logical next step is therefore to investigate the conditions under which speed-accurate path integration can be achieved. Chapter 5 approaches the path integration speed problem. HD cells are typically modelled using CANN architecture, with cells connected to one another via specific topographical recurrent collateral weights. This allows for an activity packet, representing current HD, to be sustained at any arbitrary location in the continuous state space represented by cells; current allocentric HD in this case. Path integration is typically modelled as some asymmetric input to an HD CANN activity packet, causing the location of the activity packet, and thus the current representation of HD, to shift. The major question to answer then is what biologically-plausible mechanism could provide this input, and how such a mechanism could allow the current HD representation to accurately track true HD. Previous work (Walters and Stringer, 2010; Walters et al., 2013) has attempted to model path integration based on forming associations between successive HD cells across naturally-occuring time

delays in the brain. This amounts to a profile of connectivity that predicts, for a given starting HD, where an activity packet ought to be following rotation at a particular velocity over a specific time delay. However, this work has not achieved speed-accurate path integration, achieving a maximum of approximately 80% of the target speed.

The model in Chapter 5 is a minimal one, consisting of a single layer of HD cells. These cells are notionally arranged as a CANN network with specifically wired recurrent connections containing a delay Δt. Connectivity for each cell is pre-wired as a Gaussian profile with an offset, which is based on target speed and Δt. Each HD cell projects a Gaussian recurrent weight profile centred not on its own location, but on a location offset from its own by speed $\times$ Δt. Connectivity can be purely offset in this manner, or it can have added to it a non-offset symmetrical weight profile which is constructed in the manner typical of CANN networks, with each HD cell projecting back onto its own location with a Gaussian weight profile. The final weight profile, therefore, can be anything from completely offset, to completely symmetrical depending on the relative weighting of these two overlaid recurrent weight structures. These two weight profiles can be conceptualised as the symmetrical within-layer recurrents necessary for CANNs, and the offset weights necessary for path integration. This model is therefore designed to provide a testing environment for observing in detail the interactions of these two sources of HD cell input.

Simulation results from this model system indicate two major factors affecting path integration speed: the interaction of symmetrical and offset recurrent collateral connectivity, and the time it takes for neurons to start firing after receiving input (neuronal rise time). It was found that path integration

speed dropped in proportion to the strength of the non-offset weight component. This is interpreted as a clear demonstration that non-offset recurrent weights will, if active at the same time as offset weights, have a slowing effect on the activity packet speed, resulting in inaccurate path integration. Even with fully offset weights, thus negating this weight profile interaction effect, packet speed was not 100% of the target speed. This was explained by the effect of neuronal rise time. Exploration of the neuronal time constant parameter τ was found to affect the discrepancy between target and obtained speeds. Larger values of τ are associated with slower obtained speeds. Given that increasing τ increases rise time, it was shown that the inaccuracy in path integration with fully offset weights is due to the effect of rise time, and this effect is made more severe with increasing τ. This effect was found to be independent of target speed, and therefore absolute offset size, with the ratio between target and observed speeds remaining the same across multiple absolute values. This model was also tested with a simulation paradigm in which an activity packet is forced through the HD layer at the target speed, and recurrent weights are allowed to self-organise based on a Hebbian learning rule containing the delay Δt. It was found that neuronal rise time does not have an effect on offset learning, and is therefore principally a factor in replaying of previously learned speeds.

The results of Chapter 5 suggest that within-layer symmetric recurrent collateral connectivity will hinder path integration. The problem can be elucidated as follows. A given profile of connectivity between HD cells can only accurately shift the activity packet at a particular speed. If two speeds are learned, then the connectivity components associated with each will interfere with one another. This is true of any two speeds, including a speed of zero, as is the case with non-offset recurrent collateral connectivity. Whilst Chapter 5 had two connectivity profiles overlaid within a layer, this finding can

be extended to the case in which offset connectivity and non-offset connectivity originate from different sources. Indeed, it is the co-occurrence of these two inputs which limits path integration speed.

Chapter 6 provides a solution to the issue of recurrent collateral connectivity, and its negative effects on path integration speed. This is achieved by omitting recurrent collateral connectivity, with HD to HD excitatory connections no longer included in the model architecture. Instead, HD cells are reciprocally connected, with delays of Δt in each direction, to a layer of combination (COMB) cells, which respond to combinations of angular head velocity (AHV) and HD. Crucially, some COMB cells respond to combinations of HD and AHV $= 0$ (NOROT-COMB cells), and some cells respond to combinations of HD and AHV $= V$, where V is rotational speed of the head (ROT-COMB cells). When the head is not moving, HD cells have non-offset connectivity with NOROT-COMB cells, and thus a stable activity packet can be sustained at the current head direction. When the head starts to rotate, NOROT-COMB cells are quiescent, and ROT-COMB cells are driven up. These ROT-COMB cells project onto to the HD layer with an offset, such that the HD layer projects back onto itself, via the ROT-COMB layer, with an offset equal to $V \times 2\Delta t$, where V is the target speed, and $2\Delta t$ is the total delay across HD to COMB and COMB to HD connections. Thus, two separate HD to HD connectivity profiles exist: non-offset connectivity via NOROT-COMB cells, which sustains a stationary activity packet, and offset connectivity via ROT-COMB cells, which shifts an activity packet. By spreading connectivity across two layers in this manner, there is no longer a temporal co-occurrence of offset and non-offset input the HD layer, as HD cells cannot receive offset and non-offset input at the same time. Consequently, there is no interference between attractor and integrator modes of model operation.

This new model is tested in Chapter 6. Initially, ideal connectivity as described above is pre-wired into the model. This model is shown to be capable of performing path integration with a high level of accuracy ($> 99\%$ in all cases) across a range of arbitrary target velocities, as well as holding an activity packet stationary when the head is not moving. The model also copes well with a range of conduction delay values, including the biologically realistic situation in which conduction delay values vary randomly for a given connection. As expected from the previous chapter, rise time is still found to be an issue in relation to conduction delay. A self-organising model is also demonstrated to be very adept (91% accuracy) at performing path integration. The reasons for residual error in the self-organising model are discussed.

7.2 Experimental predictions

Computational models provide a hypothetical understanding of certain aspects of the brain. Typically, such models not only collate what we already know, but also make suggestions and predictions to help guide experimental investigations. Indeed, this can be considered the key contribution from computational modelling to our understanding of the brain. The models detailed in this thesis generate some important predictions about the brain. Testing these predictions can guide future research, and for this reason the major predictions, and putative methods of their empirical testing, are listed below. The following section (future directions) presents a discussion of outstanding questions related to the models in this thesis and, where applicable, how these questions might be investigated.

- **Chapter 3:** Cue conflict depends on dynamic remapping of visual information to HD. Transition states occur during cue conflict, with intermediate head directions, between that signalled

by internal HD memory state and that signalled by external allothetic landmark cues, being represented. This prediction could be tested very simply by replicating the Knight et al. (2014) study, whilst preventing synaptic plasticity in feed-forward visual to HD area connectivity. This would involve postsynaptic administration of substances capable of disrupting N-methyl-D-aspartate (NMDA) mediated synaptic plasticity, also known as Long Term Potentiatio (LTP). Such substances include any number of NMDA receptor antagonists (e.g., ketamine). It should be noted that to observe specifically the effect of blocking plasticity on cue conflict resolution, as opposed to learning of landmark salience, blocking of LTP must occur *after* a landmark cue has been learnt as a stable source of directional information. If dynamic reweighting of visual landmarks is responsible for the observed cue conflict resolutions as described in this thesis, blocking LTP *in vivo* should generate results similar to Figure 3.8 of Chapter 3, where we see cue conflict always resolved in favour of the direction signalled by allocentric visual inputs.

- **Chapter 4:** Isomapping is not an instantaneous process, but occurs over a short period of time following initial environmental changes. THis prediction can also be tested empirically. In theory, a a large scale ensemble recording of HD cells could allow fine temporal scale measurements of individual cell's preferred firing directions. This would enable checking of the preferred firing directions of multiple HD cells, to see when exactly individual cells change directions relative to one another. However, HD cell preferred firing directions are calculated over time. For a given cell, this is done by observing which spikes occur at which heading of the animal and dividing the total number of spikes by the time spent at each directional heading in order to extract a relationship between firing rate and HD. Thus, measuring preferred firing directions over a very short period of time, such as that over which isomapping might occur,

does not make sense.

Another approach might be to administer an NMDA receptor antagonist a short period after presentation of a conflicting visual landmark cue. This would be a potential method to pause remapping in relationship to vision such that some cells have altered their relationship to the landmark and others have not. However, it is not clear precisely what time interval is acceptable for some cells to have remapped in relationship. That is, does the time course of model operation correspond to the time course of an equivalent mechanism in vivo? In addition, it should be pointed out that isomapping in this case is only occurring in relationship to a visual landmark cue. This means that even if, theoretically, one could pause the network's remapping state, it might well be that the combined effect of oredered HD to HD connectivity, whether direct or via another layer, coupled with strong path integration input, overrides disorganised visual input making the effect of the partial remapping in response to vision hard to observe in terms of HD cell firing. This is especially likely, given the prediction of Chapter 4 that strong path integration input is a dominant influence on HD layer packet location.

- **Chapters 3 and 4:** Visual landmarks must be learnt as stable with respect to an internal representation of HD. The stability of landmarks can affect how cue conflict is resolved. This prediction might be empirically tested by systematically varying the stability of visual landmarks and observing how cue conflict resolution changes. The work in this thesis would predict that less stable landmarks would be reduced in salience and rotation of such landmarks would therefore cause less rotation of HD cell preferred firing directions. Given that salience is tentatively suggested in this thesis to be reflected in the coherency of visual to HD feed-forward

weights, it would follow that cue control from unstable cues would be degraded.

In theory, it might be possible to systematically vary the salience of two visual landmarks cue in relationship to one another.One could then test how HD cell preferred firing directions would rotate following independent rotation of visual landmark cues of varying previously-experienced stability. Utilising two landmarks could provide an important further test of the effect of landmark stability on visual landmark control, as it would allow discrimination between absolute and relative landmark salience. Would a landmark, previously learnt as unstable and therefore exhibiting weak cue control when experienced alone, exhibit stronger cue control when experienced simultaneously to another even less stable landmark? However, it should be noted that detailed explanation of the influence of multiple visual landmarks is not within the scope of this thesis, and any putative experiments of this type are not direct tests of landmark stability and salience in terms of the cue conflict mechanism outlined.

- **Chapters 5 and 6:** Direct excitatory within-layer recurrent collateral connectivity does not exist in the HD system because it hinders path integration. Instead, HD cells influence one another via another layer, both when exhibiting attractor behaviour of holding a stable pattern of activity and when exhibiting integrator behaviour of shifting a pattern of activity at a specific speed. This assertion can be empirically testing by one or both of two means. The first would be observing (or not) direct within-layer connectivity in those HD regions determined to be involved in the generation of the HD signal (i.e., DTN and LMN). This could be achieved experimentally through the use of synaptic tracers, injecting in a presynaptic HD region and observing where these tracers occur postsynaptically. If within-layer connectivity is present,

HD cells in the same layer should show presence of the synaptic tracer. However, this might be complicated by the fact that it could be difficult to determine whether the presence of tracer in the HD region of interest is via synaptic transmission or initial administration. Even if it were possible to make this distinction, it might be that postsynaptic tracer is present in the HD region of interest via connectivity with another layer of cells. Indeed, simply observing the presence of a given sort of connectivity doesn't tell us much about its role.

For these reasons a second more functional, as opposed to architectural approach, whereby reciprocal connectivity across layers is disrupted would be more informative. This should be done without disrupting possible within-layer connectivity, as disrupting between and possible within layer connectivity at the same time (e.g., with lesions as has been done in the past) does not separate the functional contribution of these sorts of connectivity. Therefore, the best methodology would be an optogenetic inactivation specifically of the reciprocal connectivity between DTN and LMN. If within-layer is all that is needed to generate an HD attractor network, attractor behaviour of either brain area should not be affected by this intervention. Examination of the specific deficit following this manipulation could help to uncover what, if not providing attractor capabilities, reciprocal DTN-LMN connectiivity does.

- **Chapter 6:** A three-layer system, consisting of HD, COMB (AHV by HD) and AHV cells, is capable of performing very accurate path integration in addition to acting as a CANN. This system critically depends on AHV input. This prediction has in fact been tested by disrupting AHV input to the HD cell system. and observing whether the network is still capable of holding a stable activity packet and shifting this activity packet in register with head movements. Whilst

studies outline in Chapter 1 (Section 1.2.4) show unstable HD cell firing, evidence is equivocal as to whether this result represents an attractor network decoupled from path integration inputs, or a degradation of the system as a whole. Further studies could very closely examine HD cell firing after removal of AHV input. Specific attention should be paid to the relationship of individual cells to one another, as this might allow inferences to made about connectivity. Cells with similar preferred firing directions would fire after a monosynaptic delay (2-3ms) in a system with direct within-layer connectivity, or after a longer delay if excitation occurs across layers. However, it should be noted that whilst this thesis refers to AHV input, it is likely the ideothetic input involved in path integration is from several sources (e.g., motor efference copy). A true test of this model's functioning would therefore be to remove all ideothetic inputs and observe HD layer functionality.

Also implicit in the model are certain architectural predictions (although it should be noted that the model is deliberately vague in this regard). Figure 7.1 is a modified version of Figure 1.4 (Chapter 1), showing the rough anatomical regions corresponding to the model layers in Chapter 6. Roughly, visual information areas are RSC and PoS, although it is not clear precisely the individual role of these regions, nor the exact form that visual information takes. It seems likely that the DTN-LMN complex is responsible for the generation of the HD signal, given that ADN HD cell firing is abolished by lesions to either DTN (Bassett et al., 2007) or LMN (Blair et al., 1997; Bassett et al., 2007). However, it is difficult to clarify further how the architecture outlined in Chapter 6 applies to LMN and DTN. Both LMN and DTN contain cells responsive to AHV, (Sharp et al., 2001; Stackman and Taube, 1998; Bassett and Taube, 2001) combinations of AHV and HD (Stackman and Taube, 1998; Bassett and Taube, 2001). These regions also represent the convergence of information seen in the model,

receiving both AHV inputs from NPH and SGN (Winter and Taube, 2014) and visual information via direction PoS to LMN connections (Yoder et al., 2015). Thus, the necessary cellular mechanisms are in place for the model to operate as theorised. However, it seems unlikely that either DTN or LMN represents a region of a single cell type and purpose as seen in the model. More plausible, is that this architecture is spread between DTN and LMN. A method of testing this prediction experimentally might be to trace the postsynaptic targets not of cells within a particular area, but of cells with a particular response type. This would reveal whether connectivity of the sort described in the model exists *in vivo*.

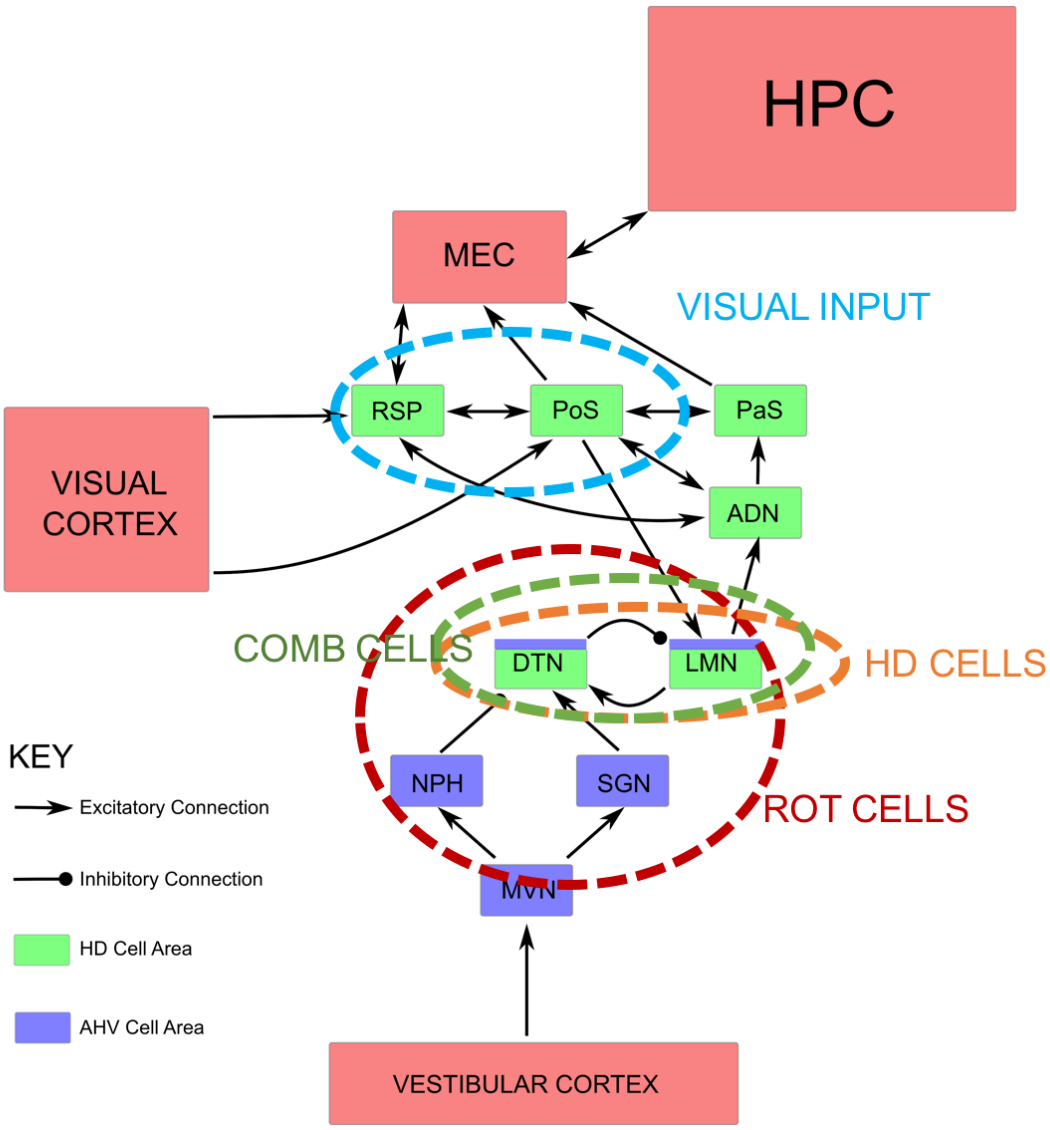


Figure 7.1: Modified version of Figure 1.4 (Chapter 1), showing anatomical regions corresponding to the model layers described in Chapter 6.

7.3 Future directions

Whilst computational modelling informs experimental work, it can also highlight which other issues are yet to be explained with a hypothesis in the form of a revised model. Indeed, the work in this thesis follows this pattern, with each experimental chapter addressing a question raised by the one preceding it. Given this trend it is no surprise, then, that as many questions arise from this thesis as are answered. The following sections discuss a number of outstanding questions regarding the HD cell system, and provide some theoretical context in which to place future work; both experimental and computational.

7.3.1 Origins and nature of visual landmark input

This thesis makes extensive use of the idea of visual information mapping onto the HD system. It has been necessary to simplify the nature of visual input, in order to focus on clarifying the phenomena that make use of such input (i.e., cue conflict). However, richer modelling of visual input is an important next step in understanding the details of how visual information dynamically influences HD cell firing. The model in Chapters 3 and 4 made use of a ring of visual cells, whose firing is taken to represent egocentric bearing of a visual landmark. The simplifying assumptions are made that the agent is stationary, and that true HD does not change. Therefore the visual ring could in fact theoretically encode landmark bearing either independently of, or dependent on, current HD, in either an egocentric or allocentric reference frame. For the cue conflict mechanism in Chapters 3 and 4 to work as outlined, activity in the visual ring must not shift unless the location of a landmark shifts. Therefore, if the representation is to be egocentric, landmarks must be distant enough to the agent such that the egocentric bearing to them varies minimally as the agent explores. It may even be possible that

visual landmark input does correspond exactly to a single reference frame as described in the model and the HD system instead receives a combination of egocentric and allocentric landmark bearing information, or information which does not neatly fall into the egocentric-allocentric categorisation.

What might be the origin of specific landmark bearing information? PoS is a good candidate entry point for visual information to the HD system (Clark and Taube, 2011; Winter and Taube, 2014), given that lesions disrupt landmark control of HD firing in other areas (Yoder and Taube, 2008), and blocking PoS plasticity can affect object-location memory (Bett et al., 2013). It appears specifically that connectivity from PoS to LMN is responsible for conveying landmark information to the HD system (Yoder et al., 2015). RSP is also thought to be involved in landmark processing, and it receives direct connectivity from the visual cortex (Van Groen and Wyss, 2003). It is known to be responsible for processing spatial information during navigation (Knight and Hayman, 2014). Damage to RSP in rats causes impairments in spatial tasks that require allocentric landmark processing (Vann et al., 2009; Hindley et al., 2014; Nelson et al., 2015). However, lesions appear to cause a milder impairment in landmark control of HD firing (Clark et al., 2010). It seems that RSP is ideally placed for transition between different sorts of landmark representations. It is considered an interface between posterior parietal cortex and parahippocampal structures such as PoS (Wilber et al., 2014), converting between allocentric and egocentric representations (Epstein, 2008; Burgess, 2008). Studies using fMRI in humans show RSC responding principally to landmarks that are more permanent in the environment (Auger et al., 2012; Auger and Maguire, 2013; Miller et al., 2014), and this suggests a role in identifying stable landmarks which could provide a navigational fix. It is unclear why, if RSP is an important node for conveying landmark information to the HD system, lesions indicate the most

severe landmark control deficits following PoS damage and relatively milder impairments following RSP damage. Future research should seek to explain more fully the interrelationship of PoS and RSP, and clarify their respective roles in landmark processing.

Beyond these two structures, we are faced with the question of which are the principal inputs conveying landmark information to PoS and RSC. Posterior parietal cortex (PPC) is a possible source of landmark bearing information, and is suggested to process conjunctions of allocentric and egocentric information, as well as having a specific role in representing egocentric coordinates (Spiers and Barry, 2015). It is known to project to RSP (Wilber et al., 2014), and thus could provide input to the HD system via RSP or PoS, with these latter two structures being reciprocally connected (van Groen and Wyss, 1992; Hopkins, 2005). Lesions of PPC impair orienting towards targets (Save and Poucet, 2000), suggesting it is responsible for processing some aspect of orientation relative to a particular location. Recently, cells encoding egocentric landmark bearing have been found in PPC (Wilber et al., 2014). Interestingly, the results of Wilber and colleagues (2014) show that PPC also contains conjunctive cells encoding both egocentric landmark bearing and HD. The authors note that this information could be integrated to produce allocentric representations of landmark bearing, and then combined with distance information to produce landmark vector (LV) cells. LV cells can be characterised as encoding location as a vector from a specific landmark, firing at a set bearing and distance from environmental landmarks (Deshmukh and Knierim, 2013). Despite this promising evidence, PPC lesions have failed to impair landmark control of HD cell firing (Calton et al., 2008).

Another possible source of landmark input to PoS/RSC is hippocampal CA1, a region which has

recently been discovered to contain the theorised LV cells (Deshmukh, 2014), and is known to play a role in encoding goal location during navigation (Dupret et al., 2010). Despite the fact that these cells do not encode bearing alone, instead encoding both bearing and distance, it opens the possibility that hippocampal cells are capable of conveying allocentric landmark bearing information. Although PoS does receive connections from CA1 (Van Groen and Wyss, 1990), it seems that CA1 lesions do not impair landmark control of HD cell firing (Golob and Taube, 1997). This speaks against CA1 input acting as the primary source of landmark information directly to the HD system, and instead might interact with other landmark encoding systems. Another possible locus for an allocentric vector towards a spatial goal is the postrhinal cortex, although lesion studies vary in their assessment of postrhinal cortex importance for landmark navigation (Liu and Bilkey, 2002; Burwell et al., 2004). Postrhinal cortex projects to the entorhinal cortex (Erdem and Hasselmo, 2012), revealing a potential route to PoS/RSC. However, entorhinal cortex lesions have not been shown to affect landmark control in HD cells (Clark and Taube, 2011); although landmark control in place cells is mildly impaired (Van Cauter et al., 2008). To date, evidence of allocentric landmark position cells in PPC specifically has been equivocal (Wilber et al., 2014), and thus far there has been no evidence for allocentric goal direction encoding in the brain (Spiers and Barry, 2015). Despite this, they remain an important theoretical construct. Further recording studies targeting PPC, RSP or even entorhinal cortex will help to clarify the nature of landmark input to the HD system by locating such allocentric landmark representations, and clarifying the importance of these structures for landmark control of HD cell firing.

As mentioned above, the simulation results in Chapters 3 and 4 are generated using a model with the simplifying assumption that the agent is stationary during cue conflict. Whilst this has greatly helped

explanation of cue conflict resolution, this mechanism should be extended to situations in which a simulated agent is freely exploring the environment, as this will more closely match the experimental conditions of Knight and colleagues (2014). As is clear from the above discussion of landmark bearing representations in the brain, it is not yet known exactly what form landmark input to the HD system takes, and what the specific role of different structures are in processing such information. If the model of Chapters 3 and 4 is to exist in the same from, then the visual ring activity packet must not shift as the agent explores the environment. This would suggest an allocentric representation of landmark location, unless landmarks were sufficiently distant for egocentric bearing to be relatively unchanging as the agent moves. Alternatively, it may be that landmark information exists in a more complicated form, and that mixtures of allocentric and egocentric landmark-related information act as inputs to the HD system. In this case, future iterations of the model would look somewhat different but would retain the core hypothesis of the original model presented in this thesis: dynamic plasticity changes in the relationship between visual landmarks and the HD system during cue conflict. Although both CA1 and PPC demonstrate cells firing properties which may be of use in conveying landmark information, lesion studies have demonstrated that neither area is critical for landmark control of HD responses. The role of RSP too is unclear, given that lesions do affect landmark control, but not as severely as lesions of PoS. It is hoped that further experimental work will provide enough information to make an informed computational hypothesis as to the nature of landmark information entry to the HD system.

7.3.2 Proximal-Distal landmark distinctions

On the topic of visual landmark information, there is also the related question of proximal-distal landmark distinctions. It is known that distal cues (i.e., those relatively further from the rat) have a greater

influence over the HD system than proximal cues (i.e., those relatively closer to the rat). This has been both inferred from a variety of behavioural studies indicating a preferential use of distal landmarks for navigation in the Morris water maze (Knierim and Hamilton, 2011) and directly observed with HD cell preferred firing directions shifting with distal, but not proximal, landmarks (Zugaro et al., 2001, 2004; Yoganarasimha et al., 2006). Future work elaborating the nature of visual landmark inputs to the HD system ought to address the topic of proximal-distal distinctions; both in in terms of how the HD system distinguishes between these two types of landmark cue, and why distal cues are preferred.

One approach is based on the idea of visual input statistics. In vision research, input statistics have been shown to affect the representations formed by competitive neural networks. For example, separate representations of multiple objects can be formed when objects are not always seen together (Stringer and Rolls, 2008). Separate representations can even be generated when objects are always seen together but transform in visual space independently (Tromans et al., 2012). This can be explained by virtue of the fact that the features comprising individual objects are more strongly correlated with one another than with the features comprising other objects, and are thus associated with one another to create separate object representations.

This understanding of statistical regularity as a means to separate visual object representations can be applied to proximal-distal distinctions. Let us assume a distal object which is infinitely distant from the agent. As the agent explores the environment, this distal cue will always be viewed in the same manner when the head is facing in a particular direction; regardless of location in the environment. Thus, there is a consistent relationship between head motion and motion of the distal landmark

cue. In contrast, a proximal cue within the environment would not always be at the same egocentric bearing from the agent for a given head direction across multiple locations. Thus no consistent relationship can be learnt between egocentric bearing to a proximal cue and head direction. In short, as the location of the agent varies, the egocentric bearing to distal cues will transform in a manner consistent with head direction transformation, whilst proximal cue bearing will not.

In a sense, this proximal-distal distinction is an arbitrary one imposed by experimenters and theorists. In reality, the only true distinction is between objects which transform in a manner that is consistent with HD and those which do not. In fact, this opens the possibility of a less binarised distinction: could visual objects lie along a continuum of least useful to most useful based on their consistency with head direction, and therefore their usefulness as a landmark cue? This is certainly implicit in the work of Chapters 3 and 4, which presents the idea that visual landmarks can be more or less salient based on their perceived reliability. In both of these chapters, reliability is specifically theorised to be a reflection of the consistency between a visual landmarks egocentric bearing and current head direction. Indeed, real environments do not consist of objects at two set distances from the agent, and objects are likely to vary in position along this proximal-distal continuum.

Discrimination between proximal and distal objects in this manner is a general principle: any statistical difference inherent in the way distal and proximal objects transform can be used to separate these cue types based on distance alone. Provided that this statistical relationship is different for proximal and distal cues, they can be learned separately, and whichever type of cue demonstrates the most concordance with current HD will be of most use to the HD system.

It would be interesting to test the role of input statistics in visual cue discrimination empirically, perhaps with an experiment designed to systematically reduce the salience of distal visual cues relative to proximal ones. Contemporary opinion is shifting towards a view of the hippocampus in general as a structure involved in encoding context, of which spatial information is only one component (e.g., Smith and Bulkin 2014. Is it therefore possible that ambiguous information from proximal cues could be disambiguated using more complete information about the spatial layout of a given apparatus/environment provided from the hippocampus? Top down control is of debatable likelihood given the lack of evidence for CA1 lesions affecting visual control of HD cell firing (Golob and Taube, 1997), although it should be noted that this study did not test landmark control in the specific situation theorised (i.e., in the case of making sense of, or compensating for, unstable landmarks). Even without top-down control, could the arrangement of multiple proximal cues provide enough orienting information? Certainly, the idea of considering individual visual landmarks is somewhat limited, given that the entire visual scenes with large amounts of directional information are available. Or would proximal cues not provide the necessary visual input, causing the HD system to, as in cue conflict simulations of Chapters 3 and 4 with low salience, rely more on an internal sense of direction?

Let us consider the experimental literature in more detail. In the experiments of Zugaro and colleagues (Zugaro et al., 2001), rats explored a cylindrical apparatus containing objects at the edge. With the cylinder in place (the proximal condition), these objects controlled HD preferred firing directions. When the cylinder was removed (the distal condition), these objects no longer controlled HD preferred firing directions. In contrast to the ideas presented in this thesis, Zugaro and colleagues do

not assume that the proximal objects are ignored in the distal condition because they provide ambiguous orienting information, citing the fact that these objects also provided ambiguous information in the proximal condition, where they controlled HD firing. For example, rats were able to view the objects by exploring the narrow space between object and enclosure, and stability of landmarks relative to HD did not change with removal of the cylinder. They state instead that it is distance alone which is important, with the HD system anchoring to the most distal landmark cues available. They support this assertion with a follow up study (Zugaro et al., 2004), demonstrating that strobe lighting disrupts proximal-distal distinctions and suggesting that dynamic visual information is therefore important.

The idea of differentiation based on relative distance, rather than stability, begs the question of how and why distance is important for separating landmark cues. If more distant cues are used as anchors, it is presumably because they are more reliable sources of orienting information. It seems curious then that reliability of orienting information does not seem to affect usage of proximal cues in Zugaro et als proximal condition. More parsimonious is the idea presented here that relatively stability is both the why and how of separating landmark cues in a distance-dependent manner. Thus, the results of Zugaro and colleagues could be explained by the fact that although landmark cues did not stably indicate HD, they were the most stable landmarks available. This could also explain why stroboscopic lighting disrupts this process, as it would disrupt much of the information needed to ascertain cue stability. If stability is the crucial determinant, this raises the further question of how HD firing remained stable in the environment during the proximal condition. This is curious, considering that their variability in relationship to head direction is the same in both conditions. As discussed above, this is potentially related to the overall composition of the environment, rather than the information signalled by a single

proximal cue alone, or potentially the simple fact that the proximal landmarks were relatively stable (i.e., they were the only landmarks available). It is, however, an open question.

In order to address fully visual involvement in the HD system, future work should be focused in two principal directions. Firstly, as described in the above section on the origins and nature of visual input, the precise form of visual input to the HD system should be explored. Future work should seek to clarify not only the origins of visual input to the HD system, but also what form it might take. That is, how is visual information processed before it reaches the HD system? Current theories of proximal-distal landmark separation rely on general principles of input statistics (and indeed this thesis relies on a relatively general mechanism of dynamic reweighting of visual input), but the exact form of this input is unknown. It is likely that the exact manner in which visual information is pre-processed before it reaches the HD system will have an effect on the proximal-distal discrimination process.

Secondly, the rich nature of visual input itself should be considered. The work in this thesis has used idealised visual input, meaning that much of the information available in the transforming visual scene is not considered. The example of Tromans and colleagues (2012) demonstrated how independent rotation could allow for the formation of separate object representations. More generally, independent movement of objects (e.g., movement in multiple planes) through a complex combination of agent movement and rotation in the environment, in addition to other regular differences (e.g., lighting), would provide further means of separating objects by de-correlating their motion relative to the agent. It is therefore important to provide simulations with as rich a visual world as possible.

Realistic visual input would be provided by embedding neural network simulations in a realistic 3D simulated environment. In this way, no information about visual landmark movements is excluded, but full control is retained.

Another solution might be to embed a neural network simulation within a moving robotic agent, in order to sample realistic visual scenes. An important advantage of using realistic visual inputs to the network is to allow more direct replication of experimental paradigms in simulation. This concordance between experimental and simulation results would greatly aid comparison and strengthen hypotheses generated through simulation. One interesting simulation to carry out in this manner would be situations in which multiple objects are available at varying distances, and are systematically rotated independently of one another. Would HD preferred firing directions be anchored to the most distal object? A concrete test of the idea of stability-based landmark distinction would be to systematically reduce the stability of distal landmarks. Could proximal landmarks alone come to dominate the HD cell response when both landmark types are present? Based on the results of Zugaro et al (2001), this is likely. Using a virtual environment setup, real and simulated agents could be presented with the same, systematically varying, arrangements of visual landmarks; and this would greatly aid model concordance with known neural responses.

Future HD cell recording experiments could even be combined with lesions studies to help clarify the neural substrates involved in landmark distinction and general processing. For example, RSP is thought to be involved in placing allocentric information into an egocentric framework (Burgess, 2008), and could therefore have a distinct role in proximal landmark processing. Certainly RSP is

known to respond to permanent landmarks in the environment (Auger et al., 2012), presumably because they are navigationally significant. This would fit well with the conception of landmark stability presented above, as RSP is shown to be engaged in the processing of familiar rather than unfamiliar landmarks, and with increasing involvement as familiarity increases (Epstein et al., 2007). Despite this promising evidence, RSP lesions do not seem to affect landmark control of HD cell firing as severely as PoS lesions (Clark et al., 2010). However, other studies have shown that RSP lesions disrupt the ability of rats to use distal landmarks for navigation (Hindley et al., 2014). Therefore, further investigation of the effect of RSP lesions on the HD system is warranted.

7.3.3 Biological realism of AHV and COMB cells and more general path integration

Similarly to a requirement for richer modelling of visual inputs, another approach to extending the work in this thesis would be to simulate idiothetic input in more detail. Chapter 6 of this thesis makes use of input from AHV cells to generate a layer of combination (COMB) cells, which respond to a unique combination of current HD and AHV. However the form of AHV input, and correspondingly the representations formed in the COMB cell layer, is not strictly biologically accurate. In the simulations detailed in this thesis, AHV input is approximated by rotation (ROT) and no-rotation (NOROT) cells. The simulations in this thesis only require use of one direction of rotation, and thus cells are only responsive to AHV in one direction. ROT cells are binarised, firing at maximum when the agent is rotating at their favoured AHV and being quiescent otherwise. All ROT cells have the same preferred AHV. NOROT cells are also binarised, firing maximally when the agents head is stationary and being quiescent otherwise. Similarly, COMB cells in Chapter 6 are simplified. For every HD cell there are two COMB cells: one responding to a combination of HD + AHV $= 0$, and one responding to HD + AHV $= V$ (where, as before, V is the agent rotation speed used in simulation). This means that each

COMB cell responds to a specific HD, and only during either rotation or when the head is stationary.

Realistic ROT firing might confer additional functionality on the model. ROT cells all have the same preferred AHV in the same direction, thus limiting the model. It should therefore be noted that the model only copes with two speeds of head rotation: AHV $= V$ (where$V$ is some arbitrary favoured AHV of the ROT cell population) and AHV $= 0$. A more biologically realistic AHV representation should allow COMB cells to also develop more biological firing properties.

Whilst it has been of use to simplify matters in order to provide a clear mechanistic explanation for time-accurate path integration, the problem of multiple speeds and directions within a single model system needs now to be addressed. Future simulations ought to encompass a realistic representation of AHV and hopefully generating similarly realistic COMB cell firing. This is possibly with the result of creating a path integration mechanism which is generalizable to realistic agent head movements (i.e., varying speeds and directions).

A first step is to have a clear understanding of how AHV cells behave in vivo. AHV cells are found in several regions, and are generally quite similar. AHV cells in the DTN have been recorded in two main types: symmetric and asymmetric (Bassett and Taube, 2001). Symmetric cells have an approximately linear firing relationship with AHV, regardless of direction. However, this tends to plateau at higher rotation speeds. The slope of this relationship varied between cells. Asymmetric cells have the same approximately linear relationship between firing and AHV, with the same plateauing at higher speeds and the same variations in slope. However, this is only true in one direction. Individual cells

vary in their response in the opposite direction. Some show no correlation at all between firing rate and AHV, whilst some show a negative correlation. Indeed, some asymmetrical AHV cells show a negative correlation in their preferred firing direction and no correlation in the opposite direction. Crucially, all cells show some firing when the head is not rotating (i.e., when AHV $= 0$), and this fact is represented in the models of Chapter 6 by NOROT cells. Diagrams of AHV cell firing can be found in Chapter 1, Figure 1.3.

Other researchers have replicated these findings, recording asymmetric AHV cells with a linear relationship between AHV and firing, and some firing at AHV $= 0$ (Sharp et al., 2001). In their preferred direction of rotation, these cells showed a positive correlation, and showed a negative correlation in their non-preferred direction. Elsewhere in the HD system, the LMN has been found to contain symmetrical AHV cells (Stackman and Taube, 1998). These were further classified as fast cells, which showed increased firing rates for faster head turns in either direction, and slow cells, which showed decreased firing rates for faster head turns in either direction.

Cells responding to combinations of HD and AHV, modelled in this thesis as COMB cells, also exist in the brain, but are again rather more complex in the brain. They do not, as modelled, respond to a specific combination of a single AHV and a single HD. Basset and Taube (2001) report that DTN contains a small proportion of AHV cells which are also modulated by HD. Such cells show an elevated rate of firing across a range of head directions, typically 200o, which is much broader than the tuning of HD cells. One other potentially significant difference is that these cells show a lower peak firing rate than HD cells (7.6 spikes/s compared to 69.5 spikes/s). This observation of AHV cells whose

firing is also modulated by HD has been repeated in DTN (Sharp et al., 2001).

Given this complex representation of AHV, and combinations of AHV and HD, how might future networks function? It is likely that AHV is represented as a population code, such that for a given AHV, there is a unique pattern of firing across a population of AHV cells. In theory a competitive network would be able to separate these patterns, processing different AHVs independently. A proof of concept would be provided by simulating as accurately as possible the firing properties of real AHV cells and observing what type of COMB cell representations are formed. It is possible that realistic AHV input would in turn generate realistic COMB cell profiles. Much as the representation of AHV would be across a population, so too would the representation of AHV and HD in the COMB cell layer. The overall pattern of connectivity from the COMB layer onto the HD layer should allow for correct HD cell to be most strongly stimulated. At a given time, the overall pattern of activation in the COMB cell layer will uniquely specify a single HD and a single AHV. Thus, it is hypothesised that the current path integration mechanism would generalise to a more biologically realistic situation. However, this is with the caveat that often with these kinds of simulations it is very hard to predict exactly what will happen without doing them. That is, after all, part of their purpose.

The idea of COMB cells in particular yields two further, as yet unanswered, questions. Firstly, why are there so few of them? Studies reporting COMB cells make the observation that they are not the dominant cell type in the areas in question. Basset and Taube (2001) report that as few as 11% of recorded cells showed modulation by both AHV and HD. Although Sharp and colleagues (2001) report that 88% of cells recorded were COMB cells, this is by their admission not by a very stringent criterion,

and it is implied (although not stated) that the actual number of COMB cells is much fewer. In the models of Chapter 6, there is a 2:1 ratio of COMB to HD cells, as each COMB cell must represent a single HD plus both ROT and NOROT input. This situation is of course not reflected in the biology. Instead, a smaller proportion of COMB cells, as seen experimentally, might represent combinations of HD and AHV with a complex population encoding. Future simulations with realistic AHV responses should also vary the number of COMB cells available, accurately reproducing the experimentally-derived proportions where possible. It is hypothesised that different COMB cell representations will be formed depending on the number of cells in the layer. Certainly, the self-organising simulations in Chapter 6 hint at the possibility of path integration functionality with a reduced number of COMB cells.

Secondly, why are COMB cells present in both DTN and LMN? One possible explanation is that they are serving different functions. However, another answer might be that they are serving the same function, which is distributed across regions. HD and AHV cells are also seen in multiple areas, and the same principle applies. The concept of different brain regions serving different purposes might well be, to a limited extent, somewhat outdated. Much more interesting is the idea that certain arrangements of cells, and the connectivity between them, serve a particular purpose. Whilst cells that perform similar computational functions are often grouped anatomically, there may be some leeway in this organisational principle.

To ask another outstanding question about idiothetic input to the HD system: is there a specific role for asymmetric vs symmetric AHV cells? They convey somewhat different information, with asymmet-

ric cells signalling speed and direction, whilst symmetric cells signal speed independent of direction. Why would these two functions be separated? As far as path integration in a mature system is concerned, they do not seem to convey much extra information. However, cells signalling the absolute speed of rotation have been theorised to have an important role in calibrating the HD system during development of connectivity capable of performing path integration (Stratton et al., 2010, 2011). The authors speculate that the different roles of asymmetric and symmetric AHV cells require separate morphology and connectivity. It is not yet clear whether this is the case. It is possible, given the variety of cell responses, that asymmetric and symmetric AHV cells are of the same basic type, and simply display stochastic variance in their firing properties. The proposition of Stratton and colleagues (2010, 2011) is that symmetric AHV cells provide information about the speed at which an HD activity packet ought to be moving, which can be used to calibrate direct HD to HD connectivity. Given the direction of this thesis towards an HD system model that relies solely on HD to HD connectivity via a COMB cell layer this proposition would need to be altered somewhat to affect all HD to HD connections, regardless of whether they are used to maintain a stable attractor state or shift an activity packet. Could symmetric AHV input act as a type of error signal, fine tuning system connectivity to ensure path integration is at the desired speed?

7.3.4 HD attractor networks without visual input

Another issue that has not been touched upon in this thesis is the full development of HD area connectivity, and HD cell preferred firing directions. It is highly unlikely that the specific HD area connectivity required for adequate model operation, both in sustaining and in shifting an HD activity packet, is present from birth. Therefore, it is essential to explain how such connectivity can arise. In the models presented in Chapters 5 and 6, HD to HD connectivity, whether direct or via an intermediate cell

layer, self-organises during training. However, this requires each HD cell to already have a preferred firing direction, and thus a specific relationship between HD cells and allothetic landmark information must already exist. As is theorised in Chapters 3 and 4, this relationship is one that develops through experience of an environment, and is likely based on the coherence between visual information and HD cell firing. It is possible to argue, then, that for this relationship to be established, the HD system must already be internally consistent in order to correctly map allothetic information onto HD. We are therefore faced with the interesting developmental question of how HD interconnectivity is formed and related to the external world. In the context of the models presented in this thesis, allothetic input is conceived of purely as visual, and this will be the initial perspective of the below discussion. However, relating the HD system to the allothetic world is a general concept that can be applied to other modalities, as also discussed below.

Experimental results show that the HD system has at least a roughly correct HD to HD profile of connectivity from well before visual information is available. As discussed in the introduction of this thesis, it has been discovered that rat pups show HD cell activity at least 3 days before eye opening (Tan et al., 2015). Prior to eye opening, individual cells showed directional tuning, but preferred directions drifted within trials, and were inconsistent across trials. Importantly, simultaneously recorded cells maintain their relative firing directions (Bjerknes et al., 2015). Following eye opening preferred firing directions quickly become, and remain, stable both within and across trials, implying that minimal additional learning is required to associate vision with HD cell firing. Indeed, visual control of preferred firing directions is established within 24 hours of eye opening. Taken as a whole, these findings indicate that a mostly intact HD system is present independently of visual input.

7.3.4.1 Single-layer models without vision

One possibility for initial rough wiring of HD area connectivity might be based on the intrinsic dynamics of HD cell areas. There are enough questions for another thesis, but it is of interest to present some preliminary thinking on the topic. A good starting point is to consider the case of a single layer of HD cells with recurrent collateral connectivity, receiving AHV input. This is pictured in Figure 7.2. Whilst the validity of direct HD to HD connectivity, as in a standard implementation of a CANN network, is debatable, it is helpful to consider self-organisation in this simplified case before moving on to a multi-layer model without recurrent connectivity within layers, as proposed in Chapters 5 and 6. This will provide enough of a framework to discuss below potential issues in modelling pre-visual HD cell development. A strict CANN hypothesis may also be of some use, given that many researchers still conceive of the HD system as an attractor of some sort, and that standard CANN architectures with recurrent connections within layers might exist elsewhere in the brain.

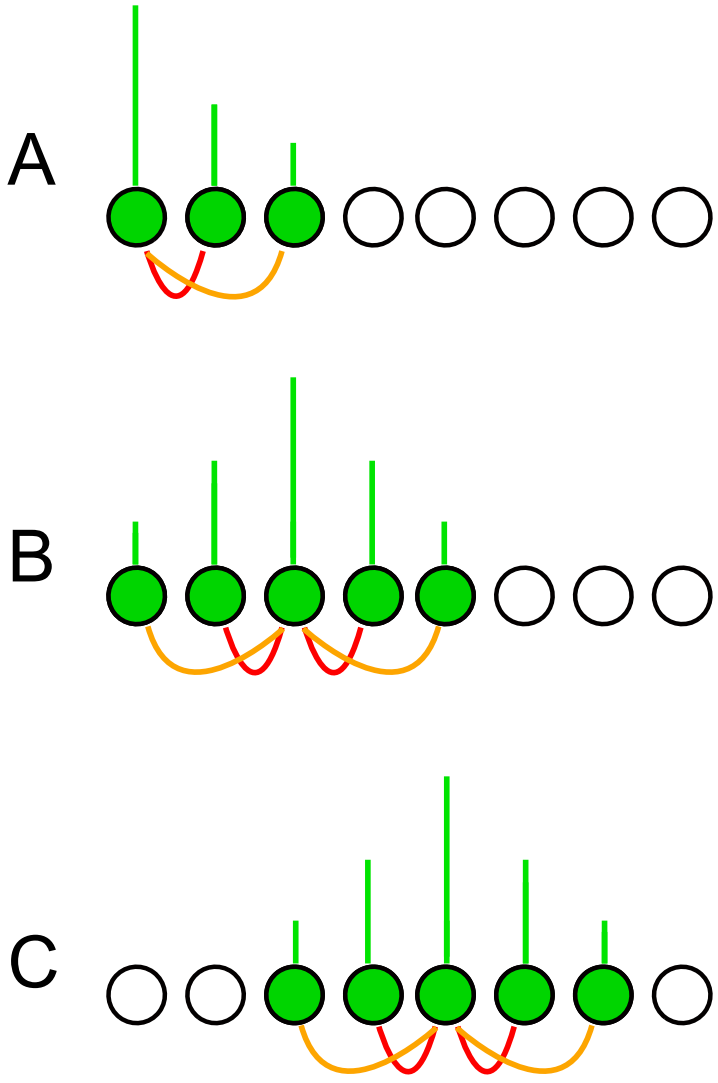


Figure 7.2: Mechanism for 1-layer attractor formation, given constant AHV input. **A:** Initially, a subset of HD cells is driven up by AHV input (green). Connectivity of varying strength (red for stronger, orange for weaker) is formed depending on the firing rate of coactive cells (green lines). For simplicity, only weights projecting from most strongly activated cell are shown. **B:** Some intrinsic mechanism reduces the firing of those cells most strongly activated, changing the pattern of HD layer firing (green cells and lines), allowing different connectivity to be formed (red and orange lines). **C:** This process continues across the HD layer.

For simplicity, we shall first consider the case in which the head is rotating at a fixed velocity in a single direction. This allows us to make the simplifying assumption that AHV input is fixed. Before eye opening, HD cells will be driven by AHV cell input. AHV to HD connectivity and HD to HD connectivity is initially random. AHV input will drive up a subpopulation of HD cells, which will have an inhibitory effect on all other HD cells. This balance of excitation and inhibition will allow a certain proportion of HD cells, those that receive the strongest AHV input, to become active whilst others are quiescent. Connectivity will then be strengthened between the active HD cells.

In this scenario, the current HD activity packet would remain active given the current pattern of AHV firing. If there is some mechanism attenuating the firing rate of currently active HD cells, the balance of excitation and inhibition will change and a different set of HD cells will next become active. This process will repeat multiple times, generating a sequence of firing in the HD layer as the head is turning. Associative Hebbian learning during this process would strengthen connectivity between HD cells that are co-active at any time (i.e., for the same or similar head directions). Depending on the time course of the attenuation, it is possible that the original subset of HD cells will become active again, and Hebbian association, coupled with temporal overlap of different HD firing patterns, would thus allow a ring attractor to be formed (without the original set of HD cells becoming active once more, a line attractor would be formed). Any attractor formed would not be tied to the external world, possibly resulting in a full rotation through all cells representing an amount of real world rotation more or less than 360°, but HD cells would have at least some rough internal coherence.

This would be a relatively simple model to simulate, and would be an important first step in mod-

elling HD system development prior to vision. It is of course a simplified version of reality, and poses several additional questions that ought to be addressed. For example, if following the above process, the same AHV input were to be applied after the above system had already run once, then the same sequence of HD cell firing ought to be generated. However, we have not yet considered whether learning would occur in the AHV to HD connections, and what impact such learning might have on how this simple model works in complex cases. Experimental reports of HD cells prior to eye opening show that cell preferred directions drift within recording trials, and thus it is not clear whether AHV to HD connectivity responsible for path integration has formed. A next experimental step to answer this question might be to compare the magnitude and nature of drift in these pre-vision HD cells with that of adult rat HD cells in the dark. If these drifts are similar, the implication would be that a mature path integration system has formed prior to vision.

A next step would then be to address cases in which AHV input changes between successive sequences. What would be the effect of prior learning on subsequent sequences? Connectivity between HD cells, once established in a single sequence, might influence which HD cells that become active in future. Those cells with reciprocal connectivity will tend to be coactive. This reciprocal effect might be enough to push firing into the same sequence, with varying AHV input patterns simply generating the sequence from a different starting cell. However, it might be that a different pattern of AHV firing stimulates a different initial pattern of cell firing, which is then unrelated to the ring originally formed. In this case, it is not clear how a unified representation of HD would be formed. One possibility is that the time course of firing attenuation is longer that the period of fixed AHV input, and thus a multitude of line attractors are formed. These lines could eventually meet to form a coherent ring.

Natural head motions are not comprised of periods of head rotation at a fixed velocity, and it is also of value for future work to consider more realistic head motions by matching the statistics of experimentally observed rat pup movement. During complex movements, the pattern of AHV firing will be expected to be time varying. This would lead to sequences of HD cell firing being driven up independently without the need of any firing rate attenuation mechanism. However, without the natural time course of firing rate attenuation, there is no guarantee that the originally-stimulate HD cells will become active again, and thus no specific means for a ring attractor to be formed. It is therefore not clear without performing the simulations described whether a regular ring structure could form via the formation of multiple line attractors as described above. Between birth and day 11, experimenters observe developmental milestones in the stereotyped movements that rat pups perform (Golani et al., 1981). The potential functional significance of these movements and how they relate to the connectivity formed should be investigated.

7.3.4.2 Multi-layer models without vision

The mechanism above, as applied to a single layer of HD cells, should also be adapted to explain the self-organisation of a two-layer system, such as the one used in Chapter 6 of this thesis. Given that this thesis emphasises the utility of a multi-layer HD system without recurrent collateral connections within layers, it is important to address how such a system might arise prior to visual input. Consider a multi-layer network, in which HD cells project and receive from a layer of AHV-by-HD cells, and AHV-by-HD cells receive some AHV input in addition to HD input. All initial wiring will be random, and all connectivity will contain some delay Δt, in line with the model of Chapter 6. Recurrent connectivity is no longer present. This schematic is pictured in Figure 7.3.

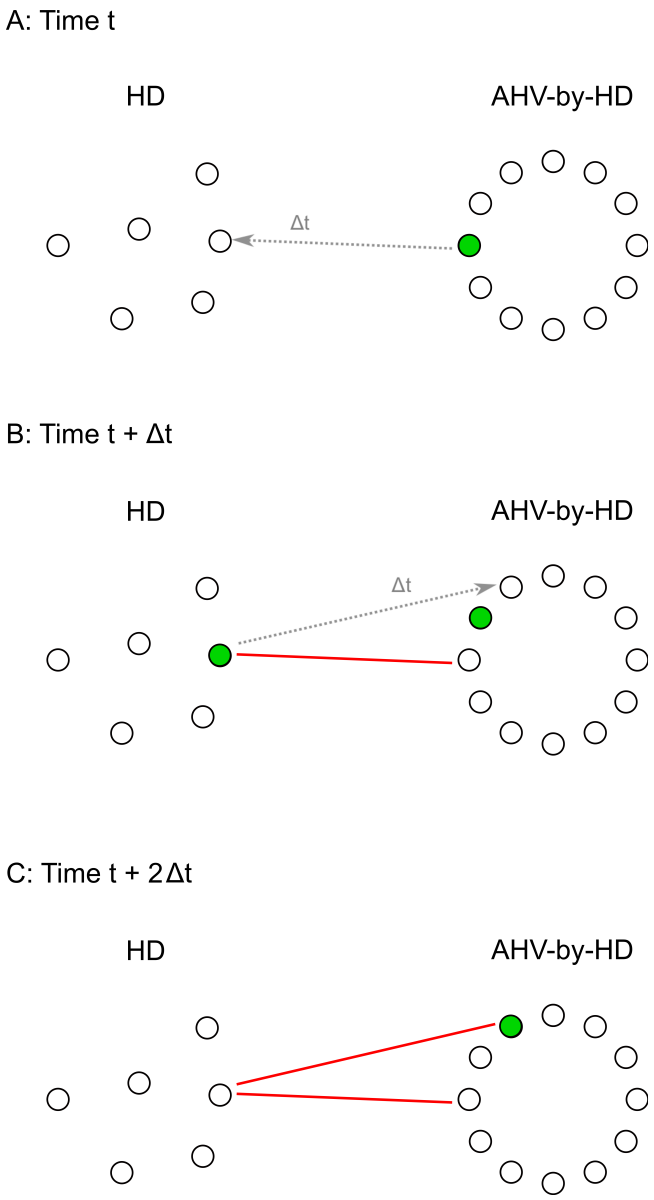


Figure 7.3: Mechanism for 2-layer attractor formation, shown for two cells only (one HD, one AHV-by-HD). Constant AHV input to the AHV-by-HD layer is assumed. For simplicity, AHV-by-HD are arranged in a ring, but it should be noted that no such topographical organisation exists. Only the most strongly firing cells in each layer at any given timestep are shown for clarity. **A:** AHV input drives up a pattern of activity in the AHV-by-HD layer (right). This cell projects to the HD layer with a connection containing a time delay Δt (grey dashed line). **B:** After Δt, a cell in the HD layer (green) is driven up. Due to a delay of Δt in the learning rule, this cell forms a strong connection (red) with the AHV-by-HD cell that was firing Δt in the past. The currently active HD cell projects back to the AHV-by-HD layer with a time delay Δt (grey dashed line). Meanwhile, the dynamics of firing rate inhibition is causing the pattern of activity in the AHV-by-HD layer to change. For simplicity, the maximally firing AHV-by-HD cell shifts one cell to the right every Δt interval. **C:** Due to firing rate inhibition, a new AHV-by-HD cell becomes active. Due to a delay Δt in the learning rule, this cell will form a strong connection with the HD cell firing maximally Δt in the past. In this manner, reciprocal connectivity will form between AHV-by-HD and HD cell layers.

Let us first consider a simplified case of a single fixed AHV signal. Much like the previous mechanism, this multi-layer mechanism will need to be extended to more realistic head motions, but a simple case will be a good start for both theoretical discussion and future simulation work. AHV cell input will drive up the AHV-by-HD layer. As described previously, some firing rate attenuation mechanism will generate a sequence of firing in this AHV-by-HD layer. Each pattern of firing in the AHV-by-HD layer will in turn drive a pattern of HD layer firing after the conduction delay Δt. In this manner, associations will be formed between HD cells and AHV-by-HD cell firing Δt in the past. Considering the connections in the opposite direction, a different pattern of AHV-by-HD cell firing will be associated with HD cell firing Δt in the past. This is because the AHV-by-HD cells will be driven in a sequence determined by firing rate attenuation. By this mechanism, reciprocal connectivity will be generated between the HD layer and the AHV-by-HD layer.

This multi-layer mechanism assumes that HD to AHV-by-HD connectivity does not affect the sequence of AHV-by-HD firing generated by the combination of firing attenuation and AHV input. This would not be detrimental if it did occur, given that the AHV-by-HD layer contains no inherent recurrent connectivity, and thus any HD influence (e.g., any AHV-by-HD cells that are driven up by HD inputs, and which would not be driven up otherwise) would simply alter the sequence slightly. It would require simulation to determine the effect of HD layer input on repeated iterations through this mechanism. At some point the AHV-by-HD layer firing would come to reflect both AHV and HD input, rather than being driven by AHV input alone. From an experimental perspective, it would be very informative to assess whether AHV-by-HD cells are present before eye opening.

By operating with a delay Δt in the learning rule, the connectivity developed from AHV-by-HD to HD, and from HD to AHV-by-HD would be asymmetrical. That is, a given HD cell would receive maximally from an AHV-by-HD cell that was firing Δ t in the past, and would project maximally to an AHV-by-HD cell that will be firing Δt in the future. There is a disconnect with the model of Chapter 6, which also requires additional symmetrical connectivity between these two layers in order to maintain a stationary HD layer activity packet when the head is not rotating. However, the results of Bjerknes and colleagues (2015) as well as Tan and colleagues (2015) indicate that there is a drift of HD cell preferred firing directions, which might well indicate that symmetrical inter-layer connectivity is not self-organised without vision. Indeed, this would then mean that a rough system containing drift could be calibrated by visual input; a hypothesis in line with experimentally observed improvements in preferred firing direction stability following eye opening. Again, light would be shed on this question by comparing preferred firing direction drift in pre-vision HD cells with drift related to path integration inaccuracy in adult HD cells. If drifts are similar in both cases, it would imply that a path integration system matures prior to vision. It would also be of use to characterise the drift in pre-vision HD cells to see if it matches the prediction of this hypothesised mechanism does it occur when the head is stationary?

7.3.4.3 Sensory inputs prior to eye opening

The suggestion of researchers behind experimental observations of pre-vision HD cells (Bjerknes et al., 2015; Tan et al., 2015) is that the HD system is approximately intact prior to vision, and that sensory input other than from the vestibular system is responsible for early HD system self-organisation. Certainly, rat pups have olfactory (Polan and Hofer, 1998) and tactile (Landers and Philip Zeigler, 2006) information available to them. It might well be, then, that the hypothesised mechanisms rely-

ing on intrinsic HD system dynamics, as described above, are not solely responsible for HD system connectivity prior to vision. Instead, it might be that such mechanisms set up an initial rough HD system structure that is then calibrated by other sensory modalities prior to vision. This is in line with the work of Stratton and colleagues (Stratton et al., 2010, 2011), who suggest a computational mechanism by which rough initial connectivity can be calibrated by other sensory modalities prior to vision, relying on stereotypical movements displayed by rat pups (Golani et al., 1981). Whilst they hypothesise pre-calibration connectivity to be formed via concentration gradients of diffusing chemicals during development, it could well be that the hypothesised mechanisms given above are responsible for initial connectivity.

7.3.4.4 Intrinsic firing rate attenuation

These HD network development hypotheses might require some intrinsic cellular mechanism for a reduction in firing rate. One good candidate is short-term depression (STD) (York and Van Rossum, 2009; Fung et al., 2012). This is a pervasive phenomenon, caused by the depletion of presynaptic neurotransmitter (Zucker, 1989; Tsodyks et al., 1998), and is typically on the timescale of 100-1000ms. STD has been shown to produce a variety of interesting behaviours in CANN networks, such as causing an activity packet to shift spontaneously despite constant external input (York and Van Rossum, 2009). This provides a potential mechanism for sequences of HD area activity, with STD presumably affecting both HD to HD and AHV to HD connectivity. However, it may be that other intrinsic properties of HD neurons can lead to spiking variability (Tsanov et al., 2014) and possibly therefore to shifting packets of activity. There is a possibility of this mechanism being coupled with intrinsic cellular mechanisms of AHV input firing, such as spike rate adaptation and postinhibitory rebound (van der Meer et al., 2007). The specifics of attractor self-organisation mechanisms are ideally suited

to being explored through simulation. Future work should focus on providing the most accurate model possible of AHV input to the HD system, as well as the dynamics of intrinsic cellular mechanisms such as STD, in order to provide the most insight into biological HD attractor self-organisation prior to visual input.

7.3.5 Relationship to bursty HD cell firing following vestibular lesions

Vestibular lesions upstream from the DTN-LMN complex have been shown to produce bursty firing of HD cells (Yoder and Taube, 2009; Muir et al., 2009; Clark et al., 2012). Following lesions, HD cell firing was no longer tied to head direction, with individual cells showing intermittent bursts of firing as if the head were passing through that cells preferred direction. Essentially, these bursty cells show the same firing characteristics as HD cells, but their firing is not tied to any specific allocentric head direction. It is generally thought that this represents the activity of an intact HD cell attractor network, which is no longer correctly signalling current head direction. This is supported by several observations. Although HD and bursty cells had similar firing characteristics, they were never observed at the same time. Additionally, bursting activity was related to head rotation, with faster rotations yielding shorter bursts of firing and the temporal order of cell bursting being dependent on direction of rotation.

Firstly, and most significantly, we must ask what is causing bursty firing? Presumably HD cells are driven up by other sources of information than visual input. Visual information is known to be used to reset the HD representation following accumulated error in path integration, and therefore ought to be able to sustain an accurate HD signal without vestibular input. This might be explained by the concept of visual landmark learning. Specifically, it is hypothesised that visual landmarks must be learnt as stable in order to provide directional information to the HD system. By stable, it is meant

that landmarks must be learnt as consistent with vestibular information in order to be used reliably. This is discussed in Chapter 3. If vestibular information were to be removed via lesions, no consistent mapping could be learned between a given visual landmark and current representation of HD as updated by vestibular information. This would mean that HD cells are driven by visual input that is not mapped consistently, causing bursts of firing which are not tied to true head direction.

This concept of a visual-vestibular interrelationship could be investigated experimentally. Does bursty firing occur in environments which have been experienced prior to lesions? This would inform us as to whether prior visual-vestibular mappings (i.e., previously learned stable landmarks) are affected by vestibular inactivation. However, this is with the caveat that the learning of these mappings is dynamic: following loss of vestibular information, even prior mappings would become degraded as the relationship between landmarks and the now-absent vestibular information falls apart. In addition to these experimental directions, it would be very informative to see whether cells responding to combinations of HD and AHV, represented by COMB cells in Chapter 6 of this thesis, also display bursty firing. This would provide valuable insight into the biological accuracy of the model outlined in Chapter 6. If biology is consistent with simulation, it would be expected that cells responding to both HD and AHV would not demonstrate bursty firing, given that they are driven, in this thesis, by both HD and AHV inputs.

Further insight into the causes of bursty firing comes from those studies in which vestibular lesions did not produce bursting cells. For example, an earlier study from Stackman and Taube (1997) did not yield bursting cells. In this study, sodium arsanilate was used to damage hair cells in the labyrinth.

However, it is argued (Muir et al., 2009) that the timing of these lesions relative to recording is an important factor. Stackman and Taube's recordings were performed shortly after lesions over a time period of 96 hours, whilst Muir and colleagues waited several weeks post lesion. This is important because secondary vestibular neurons only return to pre-lesions firing levels after a week (Ris and Godaux, 1998). Yoder and Taube (2009), who also reported bursty firing, experimented using genetically modified mice and thus would have an intact secondary vestibular nerve. Similarly, Stackman and colleagues (2002) reported no bursty firing after transient vestibular inactivation via injections of tetrodotoxin, but recorded in the 48 hours directly post injection. Could it be then that secondary vestibular neuron firing is causing bursting? The pattern of results seems consistent with this hypothesis, but the existing studies removing vestibular input may have other differences that have not yet been considered. Regardless, this train of thought highlights the need for future computational work to consider fully the complex nature of vestibular input to the HD system. The fact that vestibular function needs to be explored in more depth is also revealed by suggestions of differential roles of vestibular organs. Yoder and Taube (2009) observe of a small proportion of normal HD cells in their otoconia deficient mice, whilst studies affecting the semicircular canal do not report any normal HD functioning. Does this mean that the otoliths, whilst clearly contributing to normal HD function, are not essential?

A second major question to ask is could a seemingly intact HD system be formed without vestibular input? It is hypothesised in this thesis that AHV cell firing plays a critical role in self-organisation prior to eye opening. Of particular interest then is the implication, in studies reporting bursty firing, that the HD system is intact. This is hinted at by the fact that bursty firing is not completely random,

with HD cells showing a consistent temporal relationship to one another. The answer to this question is not made clear by those studies removing vestibular input through pharmacological or surgical intervention (Stackman and Taube, 1997; Stackman and Herbert, 2002; Muir et al., 2009) as they do not do so early enough to disrupt development. For this reason, we must assume that results come from an HD system with mature internal connectivity. More informatively, Yoder and Taube (2009) have observed bursty firing in mice congenitally lacking otoconia (a crucial component of the otolith organs in the inner ear). This would presumably result in a lack of vestibular input to the HD system during development, and it is unclear why the HD system would appear to be intact. However, Yoder and Taube also report some normal HD cells. This implies that otoliths are not the only vestibular input to the HD system, and AVH cells could be driven by other vestibular organs during development. Alternatively, a rough attractor might be in place without vestibular input, with a possibility for other idiothetic information sources (e.g., motor efference copy) adaptively compensating during development.

7.3.6 Preferred firing direction shifts and path integration accuracy

This thesis has, in part, focused on achieving accurate path integration. All simulations and subsequent discussion have been aimed towards achieving 100% accurate replay in a model system. Whilst this is a useful stepping-stone towards simulating HD cell path integration, it is worth considering how accurate path integration actually is in vivo. It is known that biological path integration accumulates errors, but quantifying these errors and understanding their nature is not simple. It is of use to first discuss how experimental findings on path integration inaccuracy relate to the models presented in this thesis, as this will demonstrate how difficult it is to tie these experimental works to a theoretical underpinning in order to estimate path integration inaccuracy. This is principally because there is a

disconnect in experimental vs. computational perspective: experimental reports of path integration inaccuracy measure preferred firing direction drift, whilst the models in this thesis consider path integration inaccuracy in terms of HD layer activity packet motion.

Several studies passing comment on path integration accuracy record HD cells in situations in which visual input does not provide clear directional information, as no learned visual landmark is available. Stackman and Herbert (2002) recorded HD cells in rats transitioning from a familiar environment to a novel one. Previous studies have found that preferred firing direction remains relatively stable, presumably reflecting path integration maintenance of the HD representation when faced with a novel environment containing visual landmarks not yet learned as stable. However, an average error of 18° (shift range $\pm 30^\circ$) was introduced (Taube, 1995), possibly representing path integration inaccuracy. Stackman et al (2002) replicated this previous finding, also demonstrating that removing other sources of idiothetic input caused further shifts in preferred firing directions: turning the lights off, thus removing optic flow information, caused a mean shift of 30°, and passive transport, thus removing motor efference copy, a mean shift of 70°. This suggests multiple information sources are important for path integration. A similar approach, removing visual landmark information, was used by Goodridge and Taube (1995). HD cells were recorded as rats explored a cylindrical apparatus containing a single card acting as a visual landmark. The card was removed, and after 8 minutes preferred firing directions were recorded. In PoS, the mean shift was 43.6°, and in ADN 73.2° (but in the latter case, there were only 5 cells sampled). Again, this is taken as indicative of accumulating error in path integration.

These studies imply that absence of learned visual landmark cues is the case, and that other allo-

thetic information is minimised, such that the system relies on path integration to maintain a stable HD representation. The assumption would be that shifts in preferred firing direction are based on inaccuracies in path integration. However, this might not be strictly true, given that some visual information persists. To demonstrate why, let us define what exactly a change in preferred firing direction is. Preferred firing directions are by their very nature allothetic: a given HD cell has a changed preferred firing direction if a new relationship between the internal HD representation and allothetic spatial information has been learnt. However, this relationship could change from either side: could visual information be causing a shift in preferred firing direction?

In the cue conflict models of Chapters 3 and 4, which explains how preferred directions change, it is the relationship between visual input and the internal HD representation updated by path integration which changes the weight of visual information onto the HD system. Thus, a simple explanation for the results given above is that preferred firing directions shift upon re-presentation of the visual cue, due to a mismatch between the idiothetic and allothetic directional information. That is, because path integration has been inaccurate and no longer tracks allothetic head direction 100% correctly. This introduces a conflict between internally-derived and externally-derived sources of directional information. In the cue conflict model, visual input is not salient until learnt, and thus does not affect the HD cell system. However, it might be that other aspects of visual information not captured in the original model are influencing preferred firing directions continuously. For example, room geometry is important to rats when landmarks are absent (Knight et al., 2011), and spatial representations have been thought to generalise across geometrically similar environments (Marchette et al., 2014). It is not clear how visual input to the HD system would change, if not simply removed, in the absence

of visual landmarks. As discussed above, it would be very informative to unravel the form that HD system visual input might take.

Against the idea of the above results being caused by altered visual input is the observation that HD cell preferred firing directions drift over time in the dark. This more directly tests path integration accuracy by removing allothetic visual input completely. Goodridge et al. (1998) recorded HD cells in rats before and after a period of navigation in the dark. HD cells were found to accumulate error during this dark navigation period, with a mean preferred firing direction shift of 23.2°. Given that visual input is absent, this change is likely to be due to inaccurate path integration. This observation of spontaneous drift in the dark during exploration has not been replicated in a later study (Knierim et al., 1998). However, this latter study was performed without attempt to minimise olfactory information, which might well have provided a stabilising allothetic input to HD cells during foraging in the dark.

Path integration error likely accumulates over time, but is it reflected in altered preferred firing directions before a comparison with allothetic input is possible? Within the model system for cue conflict outlined in Chapters 3 and 4, the necessity for conjunctive pre- and post-synaptic activity implies the presence of visual input at the time of change. This would require some visual input, which could be present in studies where that the rat is still navigating in the light. In the dark, however, preferred firing directions are found to shift continuously. Within recording sessions, rat HD cells were found to drift continuously, with preferred firing directions being compared at multiple intervals (Goodridge et al., 1998). In a sense, this could be problematic for the model presented in Chapters 3 and 4, as it is not clear how preferred firing directions shift continuously without conjunctive allothetic and idiothetic

input.

Clearly, future iterations of this model must move towards replicating preferred direction drift in the dark. It is possible that other allothetic information exists in this scenario. For example, olfactory information might persist within the environment, providing an allothetic background against which inaccurate path integration can cause a shift in preferred firing direction. Alternatively a modified learning rule might be used, which allows for changes in synaptic strength in the absence presynaptic firing. Another possibility might be that the visual representation used in the model is a slightly more abstract representation of allothetic space, and therefore persists in the dark. Yet again this highlights the need to fully explore the nature of visual input to the HD system.

Another question raised by a constant shift in preferred firing directions, is how the isomapping mechanism of Chapter 4 might operate continuously. If preferred firing directions are drifting continuously in the dark, then the distinction between the cue conflict process (Chapter 3) and the isomapping process (Chapter 4) will be somewhat blurred. Now that a general algorithm for preferred direction change has been established, it is important to extend it to the case of continuous preferred direction change. The current model would predict that HD cell coherence will break down, with the angular separation between preferred firing directions of arbitrary HD cells changing over time (i.e., a breakdown of isomapping). It is not clear yet whether this is the case, and it would be informative to take some experimental measure of isomapping in the dark.

It is possible that some other mechanism pushes the relationship between HD cells towards a par-

ticular configuration (i.e., maintains isomapping). For example, the model of Chapter 6 employs COMB cells, responding to a combination of a specific HD and angular head velocity. These COMB cells have a specific reciprocal relationship with HD cells. Therefore, it might be that when a given HD cell changes preferred firing direction, the regular relationship between HD and COMB will cause the HD system to shift as a whole, achieving isomapping. Future simulation work should explore this possibility.

In addition to visual input, the theorised isomapping mechanism requires accurate path integration in order to systematically alter the HD relationship with allothetic information. This presents the question of how inaccurate path integration, which is presumably causing preferred firing direction shifts, can allow for isomapping. This is not a trivial question, given that isomapping in the light also depends on very accurate path integration. However, it might be that over a time period short enough for isomapping to occur, path integration is accurate (or negligibly inaccurate) in order to allow effective isomapping. Indeed, drifting preferred firing directions could, at least in part, reflect incomplete isomapping on the basis of accumulated error in path integration

In order to understand how inaccurate path integration is, we must try to relate experimental findings to the theoretical work in this thesis. As is likely clear from the above discussion, this is no easy task. Looking at the results of Goodridge et al (1998), as these are taken from a situation where allothetic cues are minimised, the range of preferred firing direction drift over an 8 minute session is in the range of $0 - 90^{\circ}$, with an average of 23.2°. This means that the drift in preferred firing direction is $0 - 0.19^{\circ}/s$ with an average of $0.05^{\circ}/s$. It is of interest that the isomapping simulations

of Chapter 4 showed an inaccuracy of $1o^{\circ}/s$ in path integration led to a preferred firing direction drift of approximately $0.7^{\circ}/s$. One could be forgiven for thinking that the similarity of this drift speed suggests that path integration error in Goodridge et als experiment is approximately $1^{\circ}/s$. However we should consider this with the caveat that it is very difficult to relate to experimental findings a model which is not yet fully reproducing preferred firing direction change. For example, that finding was not with an ecologically valid simulation protocol in that it was with constant speed of rotation in a single direction and with path integration inaccurate by a constant amount. In reality, rodent navigation involves complex motions of the head, with multiple changes of direction and at multiple speeds. This raises the question of where exactly path integration error arises: is it from idiothetic input that constantly shifts an HD activity packet too slowly during turns (or indeed too quickly), or that does not change the direction of HD activity packet motion quickly enough when changing turn direction, or is it both? Is the error generated in the HD system itself, or is it vestibular in origin? In addition, the simulations in Chapter 4 also reported a breakdown in isomapping, with the SD of HD layer preferred firing directions increasing over time, and preferred firing direction shift thus representing an average of a layer in which isomapping no longer occurs.

In order to have a model which generates predictions about path integration inaccuracy, several improvements must be made. Simulated agents must undergo realistic head movements similar to those seen of rodents in the experiment of Goodridge et al (1998). In addition, the mechanisms by which preferred firing directions change must be clarified. Specifically, a mechanism by which preferred firing directions might change in the dark must be proposed, both in terms of preferred firing direction shift and isomapping. That is, how might visual to HD weights change in the dark? Are we to

assume there is no visual input, or is there some visual short term memory input? There must also be proposed a mechanism for isomapping to occur even in the presence of inaccurate path integration. In expanding the theoretical explanation of preferred firing direction change in these ways, it is hoped that progress will be made towards more clearly relating drifts in preferred firing direction to path integration in accuracy. From an experimental perspective, it would be helpful to provide further insight into how path integration error accumulates in vivo. For example, it would be important to reduce the influence of allothetic information as much as possible, and to observe what behaviours rats are performing that lead to errors. It is also important to provide data, potentially using multi-electrode recordings, on whether isomapping breaks down in the dark.

In the models formulated so far, path integration has been consistently slow. It is more likely that path integration errors are randomly distributed (i.e., both too slow and too fast), given that experimental observations of preferred direction drift in the dark show changes in both directions (Goodridge et al., 1998). This indicates that there is perhaps a noisy element to path integration. Over time, such noise would cause errors to accumulate. Future models should incorporate neuronal action potential generation, in order to see what contribution stochastic neuronal firing has on path integration. Certainly, the rate coded models in this thesis and in previous work (Walters and Stringer, 2010; Walters et al., 2013) should seek to move away from path integration errors involving a systematic undershooting of expected HD packet speed. In this regard, it is perhaps a good start to create a rate coded model system in which fully accurate path integration can occur as a first step towards a noisier spiking network.

It is possible that at least some error in path integration is generated prior to the HD system itself.

It has been shown that rats tend to structure their exploratory movements in a highly specific manner. They typically start at a home base, and perform a number of movements before returning home in a more direct fashion. During their outward journey, they typically segment their movements into paths of low curvature at higher speed, and paths of high curvature at lower speed (i.e., the faster the rat travels, the straighter the path they take). This inverse correlation of speed and curvature likely enhances the gain of angular accelerations, and make them more detectible to the vestibular system (Wallace et al., 2006). For example, slower changes in heading whilst moving across a long curved path might be harder to detect than splitting the same path into a series of more abrupt changes with straighter paths. It has been shown in humans that disruption of this inverse relationship is associated with errors in path integration (Wallace et al., 2006).

It could well be that movement segmentation falls apart due to errors in path integration, rather than causing them. However, it has been shown that when rats are forced to segment their path in a manner that disrupts inverse correlations of speed and path curvature, they are less accurate at returning to home than when they are free to structure their own movements (Wallace et al., 2010). One possibility is that path integration has a baseline level of error due to a noisy biological process. This error can accumulate and cause a breakdown in the movement segmentation described above. This loss of movement segmentation compounds error by reducing the accuracy of vestibular input. This means that there are potentially two interacting sources of error in path integration (Wallace, private communication). In addition to addressing the question of errors inherent in the HD cell system, future work should also address the contribution of vestibular input, and well as the interaction of this input with HD system stochasticity, to path integration error.

7.4 Closing statement

This thesis has investigated firing properties of HD cells via the medium of computational modelling. A novel mechanism explaining weighted integration of allothetic and idiothetic information based on dynamic remapping of visual information is proposed. This process is suggested to be affected by perceived reliability of allothetic cues. A mechanism is proposed to explain how the angular separation in preferred firing directions of different HD cells might be preserved across environmental manipulations. The issue of path integration accuracy is addressed in a simplified model system, and two major factors are proposed to affect path integration speed. Firstly, the time taken for neurons to fire after they have received input, known as rise time. Secondly, the presence of direct HD to HD connectivity within a layer. A new model is proposed to solve the second of these issues, by spreading path integration functionality across multiple layers of cells. In this manner, observed HD system architecture can be directly related to function.

Accompanying this modelling investigation of the HD system has been a discussion of the theoretical implications of the work in this thesis and directions that future work, both computational and experimental, might take. It is hoped that the value of computational modelling is demonstrated, in that it provides a framework in which to make sense of experimental findings, and from which novel hypotheses might be generated.

Bibliography

Aggleton, J. (2005). Cognitive deficits induced by lesions of structures containing head direction cells. In Wiener, S. I. and Taube, J. S., editors, *Head Direction Cells and the Neural Mechanisms of Spatial Orientation*, pages 137–162. MIT Press, Cambridge, MA.

Aggleton, J. and Brown, M. (1999). Episodic memory, amnesia, and the hippocampal–anterior thalamic axis. *Behavioral and Brain Sciences*, 22(03):425–444.

Aggleton, J., Hunt, P., Nagle, S., and Neave, N. (1996). The effects of selective lesions within the anterior thalamic nuclei on spatial memory in the rat. *Behavioural Brain Research*, 81(1):189–198.

Albantakis, L. and Deco, G. (2011). Changes of mind in an attractor network of decision-making. *PLoS Comput. Biol*, 7(6):e1002086.

Allen, G. V. and Hopkins, D. A. (1988). Mammillary body in the rat: A cytoarchitectonic, golgi, and ultrastructural study. *Journal of Comparative Neurology*, 275:39–64.

Allen, G. V. and Hopkins, D. A. (1989). Mammillary body in the rat: Topography and synaptology of projections from the subicular complex, prefrontal cortex, and midbrain tegmentum. *Journal of Comparative Neurology*, 286:311–336.

Amari, S. (1977). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, 27:77–87.

Antal, A., Keri, S., Kovacs, G., Janka, Z., and Benedek, G. (2000). Early and late components of visual categorization: An event-related potential study. *Cognitive Brain Research*, 9(1):117–119.

Arleo, A., Déjean, C., Allegraud, P., Khamassi, M., Zugaro, M., and Wiener, S. (2013). Optic flow stimuli update anterodorsal thalamus head direction neuronal activity in rats. *The Journal of Neuroscience*, 33(42):16790–16795.

Arleo, A. and Gerstner, W. (2000). Modeling rodent head-direction cells and place cells for spatial learning in bio-mimetic robotics. *From Animals to Animats*, 6:236–245.

Arleo, A. and Gerstner, W. (2001). Spatial orientation in navigating agents: Modeling head-direction cells. *Neurocomputing*, 38:1059–1065.

Auger, S. and Maguire, E. (2013). Assessing the mechanism of response in the retrosplenial cortex of good and poor navigators. *Cortex*, 49(10):2904–2913.

Auger, S., Mullally, S., and Maguire, E. (2012). Retrosplenial cortex codes for permanent landmarks. *PLoS One*, 7(8):e43620.

Barry, C., Lever, C., Hayman, R., Hartley, T., Burton, S., O'Keefe, J., Jeffery, K., and Burgess, N. (2006). The boundary vector cell model of place cell firing and spatial memory. *Reviews in the Neurosciences*, 17(1-2):71.

Bassett, J., Tullman, M., and Taube, J. (2007). Lesions of the tegmentomammillary circuit in the head direction system disrupt the head direction signal in the anterior thalamus. *The Journal of neuroscience*, 27(28):7564–7577.

Bassett, J. P. and Taube, J. S. (2001). Neural correlates for angular head velocity in the rat dorsal tegmental nucleus. *Journal of Neuroscience*, 21:5740–5751.

Bassett, J. P. and Taube, J. S. (2005). Head direction signal generation: Ascending and descending information streams. In Wiener, S. I. and Taube, J. S., editors, *Head Direction Cells and the Neural Mechanisms of Spatial Orientation*, pages 83–109. MIT Press, Cambridge, MA.

Batschelet, E. (1981). *Circular Statistics in Biology*. Academic Press, London.

Berens, P. (2009). Circstat: A matlab toolbox for circular statistics. *Journal of Statistical Software*, 31(10):1–21.

Berens, P. and Velasco, M. (2009). The circular statistics toolbox for matlab. *MPI Technical Reports*, 184:1–21.

Berthoz, A. and Viaud-Delmon, I. (1999). Multisensory integration in spatial orientation. *Current Opinion in Neurobiology*, 9(6):708–712.

Bett, D., Stevenson, C., Shires, K., Smith, M., Martin, S., Dudchenko, P., and Wood, E. (2013). The postsubiculum and spatial learning: The role of postsubicular synaptic activity and synaptic plasticity in hippocampal place cell, object, and object-location memory. *The Journal of Neuroscience*, 33(16):6928–6943.

Bett, D., Wood, E., and Dudchenko, P. (2012). The postsubiculum is necessary for spatial alternation but not for homing by path integration. *Behavioral Neuroscience*, 126(2):237.

Bi, G. G. and Poo, M. M. (1998). Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and postsynaptic cell type. *Journal of Neuroscience*, 18:10464–10472.

Biazoli, C., Goto, M., Campos, A., and Canteras, N. (2006). The supragenual nucleus: A putative relay station for ascending vestibular signs to head direction cells. *Brain Research*, 1094:138–148.

Bird, C. and Burgess, N. (2008). The hippocampus and memory: Insights from spatial processing. *Nature Reviews Neuroscience*, 9(3):182–194.

Bjerknes, T., Langston, R., Kruge, I., Moser, E., and Moser, M. (2015). Coherence among head direction cells before eye opening in rat pups. *Current Biology*, 25(1):103–108.

Blair, H., Lipscomb, B., and Sharp, P. (1997). Anticipatory time intervals of head-direction cells in the anterior thalamus of the rat: Implications for path integration in the head-direction circuit. *Journal of Neurophysiology*, 78(1):145–159.

Blair, H. T., Cho, J., and Sharp, P. E. (1998). Role of the lateral mammillary nucleus in the rat head direction circuit: A combined single unit recording and lesion study. *Neuron*, 21:1387–1397.

Blair, H. T. and Sharp, P. E. (1995). Anticipatory head-direction cells in anterior thalamus: Evidence for a thalamocortical circuit that integrates angular head motion to compute head direction. *Journal of Neuroscience*, 15:6260–6270.

Blair, H. T. and Sharp, P. E. (1996). Visual and vestibular influences on head-direction cells in the anterior thalamus of the rat. *Behavioral Neuroscience*, 120:643–660.

Blum, K. and Abbott, L. F. (1996). A model of spatial map formation in the hippocampus of the rat. *Neural Computation*, 8:85–93.

Boccara, C., Sargolini, F., Thoresen, V., Solstad, T., Witter, M., Moser, E., and Moser, M. (2010). Grid cells in pre-and parasubiculum. *Nature neuroscience*, 13(8):987–994.

Boucheny, C., Brunel, N., and Arleo, A. (2005). A continuous attractor network model without recurrent excitation: maintenance and integration in the head direction cell system. *Journal of Computational Neuroscience*, 18:205–227.

Brandon, M., Bogaard, A., Andrews, C., and Hasselmo, M. (2012). Head direction cells in the postsubiculum do not show replay of prior waking sequences during sleep. *Hippocampus*, 22(3):604–618.

Brandon, M., Bogaard, A., Libby, C., Connerney, M., Gupta, K., and Hasselmo, M. (2011). Reduction of theta rhythm dissociates grid cell spatial periodicity from directional tuning. *Science*, 332(6029):595–599.

Brandon, M., Koenig, J., Leutgeb, J., and Leutgeb, S. (2014a). New and distinct hippocampal place codes are generated in a new environment during septal inactivation. *Neuron*, 82(4):789–796.

Brandon, M., Koenig, J., and Leutgeb, S. (2014b). Parallel and convergent processing in grid cell, head-direction cell, boundary cell, and place cell networks. *Wiley Interdisciplinary Reviews: Cognitive Science*, 5(2):207–219.

Brown, J., Card, J., and Yates, B. (2005). Polysynaptic pathways from the vestibular nuclei to the lateral mammillary nucleus of the rat: Substrates for vestibular input to head direction cells. *Experimental Brain Research*, 161(1):47–61.

Brunel, N. and Wang, X.-J. (2001). Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. *Journal of Computational Neuroscience*, 11:63–85.

Burak, Y. and Fiete, I. (2009). Accurate path integration in continuous attractor network models of grid cells. *PLoS Comput Biol*, 5(2):e1000291.

Burgess, N. (2008). Grid cells and theta as oscillatory interference: Theory and predictions. *Hippocampus*, 18(12):1157–1174.

Burgess, N., Barry, C., and O'Keefe, J. (2007). An oscillatory interference model of grid cell firing. *Hippocampus*, 17(9):801–812.

Burwell, R., Bucci, D., Sanborn, M., and Jutras, M. (2004). Perirhinal and postrhinal contributions to remote memory for context. *The Journal of Neuroscience*, 24(49):11023–11028.

Bush, D., Barry, C., and Burgess, N. (2014). What do grid cells contribute to place cell firing? *Trends in Neurosciences*, 37(3):136–145.

Butler, W. and Taube, J. (2015). The nucleus prepositus hypoglossi contributes to head direction cell stability in rats. *The Journal of Neuroscience*, 35(6):2547–2558.

Buzsáki, G. (1984). Feed-forward inhibition in the hippocampal formation. *Progress in Neurobiology*, 22(2):131–153.

Byrne, P., Becker, S., and Burgess, N. (2007). Remembering the past and imagining the future: a neural model of spatial memory and imagery. *Psychological Review*, 114(2):340.

Calton, J., Stackman, R., Goodridge, J., Archey, W., Dudchenko, P., and Taube, J. (2003). Hippocampal place cell instability after lesions of the head direction cell network. *The Journal of Neuroscience*, 23(30):9719–9731.

Calton, J. and Taube, J. (2005). Degradation of head direction cell activity during inverted locomotion. *The Journal of Neuroscience*, 25(9):2420–2428.

Calton, J., Turner, C., Cyrenne, D., Lee, B., and Taube, J. (2008). Landmark control and updating of self-movement cues are largely maintained in head direction cells after lesions of the posterior parietal cortex. *Behavioral Neuroscience*, 122(4):827.

Chen, L. L., Lin, L. H., Green, E. J., Barnes, C. A., and McNaughton, B. L. (1994). Head-direction cells in the rat posterior cortex. I. Anatomical distribution and behavioral modulation. *Experimental Brain Research*, 101:8–23.

Cho, J. and Sharp, P. E. (2001). Head direction, place, and movement correlates for cells in the rat retrosplenial cortex. *Behavioral Neuroscience*, 115:3–25.

Clark, B., Bassett, J., Wang, S., and Taube, J. (2010). Impaired head direction cell representation in the anterodorsal thalamus after lesions of the retrosplenial cortex. *The Journal of Neuroscience*, 30(15):5289–5302.

Clark, B., Brown, J., and Taube, J. S. (2012). Head direction cell activity in the anterodorsal thalamus requires intact supragenual nuclei. *Journal of Neurophysiology*, 108:2676–2784.

Clark, B., Rice, J., Akers, K., Candelaria-Cook, F., Taube, J., and Hamilton, D. (2013). Lesions of the dorsal tegmental nuclei disrupt control of navigation by distal landmarks in cued, directional, and place variants of the morris water task. *Behavioral Neuroscience*, 127(4):566.

Clark, B., Sarma, A., and Taube, J. (2009). Head direction cell instability in the anterior dorsal thalamus after lesions of the interpeduncular nucleus. *The Journal of Neuroscience*, 29(2):493–507.

Clark, B. and Taube, J. (2011). Intact landmark control and angular path integration by head direction cells in the anterodorsal thalamus after lesions of the medial entorhinal cortex. *Hippocampus*, 21(7):767–782.

Clark, B. and Taube, J. (2012). Vestibular and attractor network basis of the head direction cell signal in subcortical circuits. *Frontiers in Neural Circuits 6:7, doi:10.3389/fncir2012.00007*, 6.

Clark, B., Valerio, S., and Taube, J. (2011). Disrupted grid and head direction cell signal in the entorhinal cortex and parasubiculum after lesions of the head direction system. In *Soc Neurosci Abstr*, volume 729.

Cressant, A., Muller, R. U., and Poucet, B. (1997). Failure of centrally placed objects to control the firing fields of hippocampal place cells. *Journal of Neuroscience*, 17:2531–2542.

de Almeida, L., Idiart, M., and Lisman, J. (2009). The input–output transformation of the hippocampal granule cells: from grid cells to place fields. *The Journal of Neuroscience*, 29(23):7504–7512.

Degris, T., Brunel, N., Sigaud, O., and Arleo, A. (2004a). Rapid response of head direction cells to reorienting visual cues: A computational model. *Neurocomputing*, 58–60:675–682.

Degris, T., Lacheze, L., Boucheny, C., and Arleo, A. (2004b). A spiking neuron model of head-direction cells for robot orientation. In Schaal, S., Ijspeert, A., Billard, A., Vijayakumar, S., Hallam, J., and Meyer, J., editors, *Proceedings of the Eighth Int. Conf. on the Simulation of Adaptive Behavior, From Animals to Animats*, pages 255–263. MIT Press, Cambridge, MA.

Deshmukh, S. (2014). Spatial and nonspatial representations in the lateral entorhinal cortex. In *Space, Time and Memory in the Hippocampal Formation*, pages 127–152. Springer.

Deshmukh, S. and Knierim, J. (2013). Influence of local objects on hippocampal representations: Landmark vectors and memory. *Hippocampus*, 23(4):253–267.

Dudchenko, P. and Bruce, C. (2005). Navigation without landmarks: Can rats use a sense of direction to return to a home site? *Connection Science*, 17(1-2):107–125.

Dudchenko, P. A., Goodridge, J. P., and Taube, J. S. (1997). The effects of disorientation on visual landmark control of head direction cell orientation. *Experimental Brain Research*, 115:375–380.

Duhamel, J., Colby, C., and Goldberg, M. (1992). The updating of the representation of visual space in parietal cortex by intended eye movements. *Science*, 255(5040):90.

Dupret, D., O'Neill, J., Pleydell-Bouverie, B., and Csicsvari, J. (2010). The reorganization and reactivation of hippocampal maps predict spatial memory performance. *Nature Neuroscience*, 13(8):995–1002.

Egorov, A., Hamam, B., Fransen, E., Hasselmo, M. E., and Alonso, A. (2002). Graded presistent activity in entorhinal coretx neurons. *Nature*, 420:173–178.

Egorov, A., Unsicker, K., and von Bohlen und Halbach, O. (2006). Muscarinic control of graded persistent activity in lateral amygdala neurons. *European Journal of Neuroscience*, 13:3183–3194.

Epstein, R. (2008). Parahippocampal and retrosplenial contributions to human spatial navigation. *Trends in Cognitive Sciences*, 12(10):388–396.

Epstein, R., Higgins, J., Jablonski, K., and Feiler, A. (2007). Visual scene processing in familiar and unfamiliar environments. *Journal of Neurophysiology*, 97(5):3670–3683.

Erdem, U. and Hasselmo, M. (2012). A goal-directed spatial navigation model using forward trajectory planning based on grid cells. *European Journal of Neuroscience*, 35(6):916–931.

Ermentrout, G. and Cowan, J. (1979). A mathematical theory of visual hallucination patterns. *Biological Cybernetics*, 34:137–150.

Fisher, N. (1993). *Statistical Analysis of Circular Data*. Cambridge University Press.

Fortenberry, B., Gorchetchnikov, A., and Grossberg, S. (2012). Learned integration of visual, vestibular, and motor cues in multiple brain regions computes head direction during visually guided navigation. *Hippocampus*, 22(12):2219–2237.

Fourcaud-Trocme, N., Hansel, D., van Vreeswijk, C., and Brunel, N. (2003). How spike genera- tion mechanisms determine the neuronal response to fluctuating inputs. *The Journal of Neuroscience*, 23(27):11628–11640.

Fransen, E., Tahvildari, B., Egorov, A., Hasselmo, M. E., and Alonso, A. (2006). Mechanisms of graded persistent cellular activity of entorhinal cortex layer v neurons. *Neuron*, 49:735–746.

Frohardt, R., Bassett, J., and Taube, J. (2006). Path integration and lesions within the head direction cell circuit: comparison between the roles of the anterodorsal thalamus and dorsal tegmental nucleus. *Behavioral Neuroscience*, 120(1):135.

Fuhs, M. and Touretzky, D. (2006). A spin glass model of path integration in rat medial entorhinal cortex. *The Journal of Neuroscience*, 26(16):4266–4276.

Fung, C. C. A., Wong, K. Y. M., and Wu, S. (2012). Delay compensation with dynamical synapses. *Advances in Neural Information Processing Systems*, 25:1088–1096.

Fyhn, M., Molden, S., Witter, M., Moser, E., and Moser, M. (2004). Spatial representation in the entorhinal cortex. *Science*, 305(5688):1258–1264.

Gabernet, L., Jadhav, S., Feldman, D., Carandini, M., and Scanziani, M. (2005). Somatosensory integration controlled by dynamic thalamocortical feed-forward inhibition. *Neuron*, 48(2):315–327.

Georgopoulos, A., Schwartz, A., and Kettner, R. (1986). Neuronal population coding of movement direction. *Science*, 233:1416–1419.

Gerstner, W. (2000). Population dynamics of spiking neurons: fast transients, asynchronous states, and locking. *Neural Computation*, 12(1):43–89.

Gibson, B., Butler, W., and Taube, J. (2013). The head-direction signal is critical for navigation requiring a cognitive map but not for learning a spatial habit. *Current Biology*, 23(16):1536–1540.

Golani, I., Bronchti, G., Moualem, D., and Teitelbaum, P. (1981). Warm–up along dimensions of movement in the ontogeny of exploration in rats and other infant mammals. *Proceedings of the National Academy of Sciences of the USA*, 78:7226.

Goldman-Rakic, P., Cools, A., and Srivastava, K. (1996). The prefrontal landscape: Implications of functional architecture for understanding human mentation and the central executive. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 351(1346):1445–1453.

Golob, E. and Taube, J. (1997). Head direction cells and episodic spatial information in rats without a hippocampus. *Proceedings of the National Academy of Sciences*, 94(14):7645–7650.

Goodridge, J. P., Dudchenko, P. A., Worboys, K. A., Golob, E. J., and Taube, J. S. (1998). Cue control and head direction cells. *Behavioral Neuroscience*, 112:749–761.

Goodridge, J. P. and Taube, J. S. (1995). Preferential use of the landmark navigational system by head direction cells in rats. *Behavioral Neuroscience*, 109:49–61.

Goodridge, J. P. and Taube, J. S. (1997). Interaction between the postsubiculum and anterior thalamus in the generation of head direction cell activity. *Journal of Neuroscience*, 17:9315–9330.

Goodridge, J. P. and Touretzky, D. S. (2000). Modeling attractor deformation in the rodent head-direction system. *Journal of Neurophysiology*, 83:3402–3410.

Groenewegen, H., Ahlenius, S., Haber, S., Kowall, N., and Nauta, W. (1986). Cytoarchitecture, fiber connections, and some histochemical aspects of the interpeduncular nucleus in the rat. *Journal of Comparative Neurology*, 249(1):65–102.

Gupta, K., Beer, N., Keller, L., and Hasselmo, M. (2014). Medial entorhinal grid cells and head direction cells rotate with a t-maze more often during less recently experienced rotations. *Cerebral Cortex*, 24(6):1630–1644.

Hafting, T., Fyhn, M., Molden, S., Moser, M. B., and Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436:801–806.

Hahnloser, R. H. R. (2003). Emergence of neural integration in the head-direction system by visual supervision. *Neuroscience*, 120:877–891.

Hargreaves, E., Yoganarasimha, D., and Knierim, J. (2007). Cohesiveness of spatial and directional representations recorded from neural ensembles in the anterior thalamus, parasubiculum, medial entorhinal cortex, and hippocampus. *Hippocampus*, 17(9):826.

Hasselmo, M. (2008a). Grid cell mechanisms and function: Contributions of entorhinal persistent spiking and phase resetting. *Hippocampus*, 18(12):1213–1229.

Hasselmo, M., Giocomo, L., Brandon, M., and Yoshida, M. (2010). Cellular dynamical mechanisms for encoding the time and place of events along spatiotemporal trajectories in episodic memory. *Behavioural Brain Research*, 215(2):261–274.

Hasselmo, M. E. (2008b). Temporally structured replay of neural activity in a model of entorhinal cortex, hippocampus and postsubiculum. *European Journal of Neuroscience*, 28:1301–1315.

Hayakawa, T. and Zyo, K. (1984). Comparative anatomical study of the tegmentomammillary projections in some mammals: A horseradish peroxidase study. *Brain Research*, 300(2):335–349.

Hindley, E., Nelson, A., Aggleton, J., and Vann, S. (2014). The rat retrosplenial cortex is required when visual cues are used flexibly to determine location. *Behavioural Brain Research*, 263:98–107.

Hitier, M., Besnard, S., and Smith, P. (2013). Vestibular pathways involved in cognition. *Frontiers in Integrative Neuroscience*, 8:59–59.

Holdstock, J., Mayes, A., Cezayirli, E., Aggleton, J., and Roberts, N. (1999). A comparison of egocentric and allocentric spatial memory in medial temporal lobe and korsakoff amnesics. *Cortex*, 35(4):479–501.

Hopkins, D. A. (2005). Neuroanatomy of head direction cell circuits. In Wiener, S. I. and Taube, J. S., editors, *Head Direction Cells and the Neural Mechanisms of Spatial Orientation*. MIT Press, Cambridge, MA.

Iaria, G., Bogod, N., Fox, C., and Barton, J. (2009). Developmental topographical disorientation: Case one. *Neuropsychologia*, 47(1):30–40.

Izhikevich, E. (2006). Polychronization: Computation with spikes. *Neural Computation*, 18:245–282.

Jeffery, K. J. (1997). Learning of landmark stability and instability by hippocampal place cells. *Neuropharmacology*, 37:677–687.

Jeffery, K. J. and O'Keefe, J. (1999). Learned interaction of visual and ideothetic cues in the control of place field orientation. *Experimental Brain Research*, 127:151–161.

Jen, P., Sun, X., and Lin, P. (1989). Frequency and space representation in the primary auditory cortex of the frequency modulating bateptesicus fuscus. *Journal of Comparative Physiology A*, 165(1):1–14.

Johnson, A., Seeland, K., and Redish, A. D. (2005). Reconstruction of the postsubiculum head direction signal from neural ensembles. *Hippocampus*, 15:86–96.

Khalsa, S., Tomlinson, R., Schwarz, D., and Landolt, J. (1987). Vestibular nuclear neuron activity during active and passive head movement in the alert rhesus monkey. *Journal of Neurophysiology*, 57:1484–1497.

Knierim, J. and Hamilton, D. (2011). Framing spatial cognition: neural representations of proximal and distal frames of reference and their roles in navigation. *Physiological reviews*, 91(4):1245–1279.

Knierim, J. and Zhang, K. (2012). Attractor dynamics of spatially correlated neural activity in the limbic system. *Annual Review of Neuroscience*, 35:267–285.

Knierim, J. J., Kudrimoti, H. S., and McNaughton, B. L. (1995). Place cells, head direction cells, and the learning of landmark stability. *Journal of Neuroscience*, 15:1648–1659.

Knierim, J. J., Kudrimoti, H. S., and McNaughton, B. L. (1998). Interactions between idiothetic cues and external landmarks in the control of place cells and head direction cells. *Journal of Neurophysiology*, 80:425–446.

Knight, R. and Hayman, R. (2014). Allocentric directional processing in the rodent and human retrosplenial cortex. *Frontiers in Human Neuroscience*, 8:135.

Knight, R., Hayman, R., Ginzberg, L., and Jeffery, K. (2011). Geometric cues influence head direction cells only weakly in nondisoriented rats. *The Journal of Neuroscience*, 31(44):15681–15692.

Knight, R., Piette, C., Page, H., Walters, D., Marozzi, E., Nardini, M., Stringer, S., and Jeffery, K. (2014). Weighted cue integration in the rodent head direction system. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1635):20120512.

Koch, G., Ponzo, V., Di Lorenzo, F., Caltagirone, C., and Veniero, D. (2013). Hebbian and anti-hebbian spike-timing-dependent plasticity of human cortico-cortical connections. *The Journal of Neuroscience*, 33(23):9725–9733.

Koenig, J., Linder, A., Leutgeb, J., and Leutgeb, S. (2011). The spatial periodicity of grid cells is not sustained during reduced theta oscillations. *Science*, 332(6029):592–595.

Kulkarni, M., Zhang, K., and Kirkwood, A. (2011). Single-cell persistent activity in anterodorsal thalamus. *Neuroscience letters*, 498(3):179–184.

Kyriacou, T. (2011). An implementation of a biologically inspired model of head direction cells on a robot. In *Towards Autonomous Robotic Systems*, pages 66–77. Springer.

Landers, M. and Philip Zeigler, H. (2006). Development of rodent whisking: Trigeminal input and central pattern generation. *Somatosensory and Motor Research*, 23(1-2):1–10.

Langston, R. F., Ainge, J. A., Couey, J. J., Canto, C. B., Bjerknes, T. L., Witter, M. P., Moser, E. I., and Moser, M. B. (2010). Development of the spatial representation system in the rat. *Science*, 328:1576–1580.

Lannou, J., Cazin, L., Precht, W., and Le Taillanter, M. (1984). Responses of prepositus hypoglossi neurons to optokinetic and vestibular stimulations in the rat. *Brain Research*, 301(1):39–45.

Lee, I., Yoganarasimha, D., Rao, G., and Knierim, J. (2004). Comparison of population coherence of place cells in hippocampal subfields ca1 and ca3. *Nature*, 430(6998):456–459.

Leigh, R. and Zee, D. (2015). *The neurology of eye movements*. Oxford University Press.

Leutgeb, S., Ragozzino, K. E., and Mizumori, S. J. (2000). Convergence of head direction and place information in the CA1 region of the hippocampus. *Neuroscience*, 100:11–19.

Lever, C., Burton, S., Jeewajee, A., O'Keefe, J., and Burgess, N. (2009). Boundary vector cells in the subiculum of the hippocampal formation. *The Journal of Neuroscience*, 29(31):9771–9777.

Liu, P. and Bilkey, D. (2002). The effects of nmda lesions centered on the postrhinal cortex on spatial memory tasks in the rat. *Behavioral Neuroscience*, 116(5):860.

Liu, R., Chang, L., and Wickern, G. (1984). The dorsal tegmental nucleus: An axoplasmic transport study. *Brain Research*, 310(1):123–132.

Maas, W. and Bishop, C. (2001). *Pulsed Neural Networks*. MIT Press, Cambridge, MA.

Mahan, R. (1991). Circular statistical methods: Applications in spatial and temporal performance analysis. Technical report, DTIC Document.

Mammarella, I., Lucangeli, D., and Cornoldi, C. (2010). Spatial working memory and arithmetic deficits in children with nonverbal learning difficulties. *Journal of Learning Disabilities*, 43(5):455–468.

Marchette, S., Vass, L., Ryan, J., and Epstein, R. (2014). Anchoring the neural compass: Coding of local spatial reference frames in human medial parietal lobe. *Nature Neuroscience*, 17(11):1598–1606.

Massoud, T. and Horiuchi, T. (2011). A neuromorphic vlsi head direction cell system. *Circuits and Systems I: Regular Papers, IEEE Transactions on*, 58(1):150–163.

McNaughton, B., Barnes, C., Gerrard, J., Gothard, K., Jung, M., Knierim, J., Kudrimoti, H., Qin, Y., Skaggs, W., Suster, M., et al. (1996). Deciphering the hippocampal polyglot: The hippocampus as a path integration system. *The Journal of Experimental Biology*, 199(1):173–185.

McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I., and Moser, M. B. (2006). Path integration and the neural basis of the 'cognitive map'. *Nature Reviews: Neurosciences*, 7:663–768.

McNaughton, B. L., Chen, L. L., and Markus, E. J. (1991). Dead reckoning, landmark learning, and the sense of direction: A neurophysiological and computational hypothesis. *Journal of Cognitive Neuroscience*, 3:190–202.

McNaughton, B. L., Knierim, J. J., and Wilson, M. (1995). Vector encoding and the vestibular foundations of spatial cognition: A neurophysiological and computational hypothesis. In *The Cognitive Neurosciences*. MIT Press, Cambridge, MA.

Miller, A., Vedder, L., Law, L., and Smith, D. (2014). Cues, context, and long-term memory: The role of the retrosplenial cortex in spatial cognition. *Frontiers in Human Neuroscience*, 8.

Mittelstaedt, M. L. and Mittelstaedt, H. (1980). Homing by path integration in a mammal. *Naturwissenschaften*, 67:566–567.

Mizumori, S. J., Puryear, C. B., Gill, K. M., and Guazzelli, A. (2005). Head direction cells in hippocampal afferent and efferent systems: What functions do they serve? In Wiener, S. I. and Taube, J. S., editors, *Head Direction Cells and the Neural Mechanisms of Spatial Orientation*. MIT Press, Cambridge, MA.

Mizumori, S. J., Ragozzino, K. E., and Cooper, B. G. (2000). Location and head direction representation in the dorsal striatum of rats. *Psychobiology*, 28:441–462.

Mizumori, S. J. and Williams, J. D. (1993). Directionally selective mnemonic properties of neurons in the lateral dorsal nucleus of the thalamus of rats. *Journal of Neuroscience*, 13:4015–4028.

Monacelli, A., Cushman, L., Kavcic, V., and Duffy, C. (2003). Spatial disorientation in alzheimers disease the remembrance of things passed. *Neurology*, 61(11):1491–1497.

Morisset, V. and Nagy, F. (2000). Plateau potential-dependent windup of the response to primary afferent stimuli in rat dorsal horn neurons. *European Journal of Neuroscience*, 12:3087–3095.

Moser, E. and Moser, M. (2008). A metric for space. *Hippocampus*, 18(12):1142–1156.

Muir, G., Brown, J., Carey, J., Hirvonen, T., Della Santina, C., Minor, L., and Taube, J. (2009). Disruption of the head direction cell signal after occlusion of the semicircular canals in the freely moving chinchilla. *The Journal of Neuroscience*, 29(46):14521–14533.

Muller, R. U. and Kubie, J. L. (1987). The effects of changes in the environment on the spatial firing of hippocampal complex spike cells. *Journal of Neuroscience*, 7:1951–1968.

Muller, R. U., Kubie, J. L., and Ranck, Jr., J. B. (1987). Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. *Journal of Neuroscience*, 7:1935–1950.

Nelson, A., Hindley, E., Pearce, J., Vann, S., and Aggleton, J. (2015). The effect of retrosplenial cortex lesions in rats on incidental and active spatial learning. *Frontiers in Behavioral Neuroscience*, 9(11).

Oja, E. (1982). A simplified neuron model as a principal component analyser. *Journal of Mathematical Biology*, 15:267–273.

O'Keefe, J. (1979). A review of the hippocampal place cells. *Progress in neurobiology*, 13(4):419–439.

O'Keefe, J. and Dostrovsky, J. (1971). The hippocampus as a spatial map: Preliminary evidence from unit activity in the freely moving rats. *Brain Research*, 34:171–175.

O'Keefe, J. and Recce, M. (1993). Phase relationship between hippocampal place units and eeg theta rhythm. *Hippocampus*, 3:317–330.

Owen, A., Iddon, J., Hodges, J., Summers, B., and Robbins, T. (1997). Spatial and non-spatial working memory at different stages of parkinson's disease. *Neuropsychologia*, 35(4):519–532.

Papez, J. (1937). A proposed mechanism of emotion. *Archives of Neurology & Psychiatry*, 38(4):725–743.

Passolunghi, M. and Mammarella, I. (2010). Spatial and visual working memory ability in children with difficulties in arithmetic word problem solving. *European Journal of Cognitive Psychology*, 22(6):944–963.

Peyrache, A., Lacroix, M., Petersen, P., and Buzsáki, G. (2015). Internally organized mechanisms of the head direction sense. *Nature neuroscience*, 18(4):569–575.

Pfeifer, R. and Scheier, C. (2001). *Understanding Intelligence*. MIT press.

Plotnik, M., Marlinski, V., and Goldberg, J. (2002). Reflections of efferent activity in rotational responses of chinchilla vestibular afferents. *Journal of neurophysiology*, 88(3):1234–1244.

Polan, H. J. and Hofer, M. (1998). Olfactory preference for mother over home nest shavings by newborn rats. *Developmental psychobiology*, 33(1):5–20.

Possin, K. (2010). Visual spatial cognition in neurodegenerative disease. *Neurocase*, 16(6):466–487.

Pouille, F. and Scanziani, M. (2001). Enforcement of temporal fidelity in pyramidal cells by somatic feed-forward inhibition. *Science*, 293(5532):1159–1163.

R.A., M. and R., B. (1985). Anatomical connections of the nucleus prepositus of the cat. *Journal of Comparative Neurology*, 237(3):377–407.

Ranck, Jr., J. B. (1985). Head direction cells in the deep cell layer of dorsolateral presubiculum in freely moving rats. In Buzsáki, G. and Vanderwolf, C., editors, *Electrical Activity of the Archicortex*, pages 217–220. Akadémiai Kiadó, Budapest.

Redish, A. and Touretzky, D. (1997). Cognitive maps beyond the hippocampus. *Hippocampus*, 7(1):15–35.

Redish, A. D. (1999). *Beyond the Cognitive Map*. MIT Press, Cambridge, MA.

Redish, A. D., Elga, A. N., and Touretzky, D. S. (1996). A coupled attractor model of the rodent head direction system. *Network: Computation in Neural Systems*, 7:671–685.

Renaudineau, S., Poucet, B., and Save, E. (2007). Flexible use of proximal objects and distal cues by hippocampal place cells. *Hippocampus*, 17(5):381–395.

Ris, L. and Godaux, E. (1998). Neuronal activity in the vestibular nuclei after contralateral or bilateral labyrinthectomy in the alert guinea pig. *Journal of Neurophysiology*, 80(5):2352–2367.

Rolls, E. T. (1999). Spatial view cells and the representation of place in the primate hippocampus. *Hippocampus*, 9:467–480.

Rolls, E. T. and Milward, T. (2000). A model of invariant object recognition in the visual system: learning rules, activation functions, lateral inhibition, and information-based performance measures. *Neural Computation*, 12(11):2547–2572.

Rolls, E. T., Treves, A., Robertson, R. G., Georges-François, P., and Panzeri, S. (1998). Information about spatial view in an ensemble of primate hippocampal cells. *Journal of Neurophysiology*, 79:1797–1813.

Rubin, J., Terman, D., and Chow, C. (2001). Localized bumps of activity sustained by inhibition in a two-layer thalamic network. *Journal of Computational Neuroscience*, 10:313–331.

Samsonovich, A. and McNaughton, B. L. (1997). Path integration and cognitive mapping in a continuous attractor neural network model. *Journal of Neuroscience*, 17:5900–5920.

Sargolini, F., Fyhn, M., Hafting, T., McNaughton, B. L., Witter, M. P., Moser, M. B., and Moser, E. I. (2006). Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science*, 312:758–762.

Save, E., Cressant, A., Thinus-Blanc, C., and Poucet, B. (1998). Spatial firing of hippocampal place cells in blind rats. *Journal of Neuroscience*, 18:1818–1826.

Schultheiss, N. and Redish, A. (2015). The compass within. *Nature neuroscience*, 18(4):482–483.

Sharp, P. (1999). Complimentary roles for hippocampal versus subicular/entorhinal place cells in coding place, context, and events. *Hippocampus*, 9(4):432–443.

Sharp, P. and Koester, K. (2008). Lesions of the mammillary body region severely disrupt the cortical head direction, but not place cell signal. *Hippocampus*, 18(8):766–784.

Sharp, P. E. (1996). Multiple spatial-behavioral correlates for cells in the rat postsubiculum: Multiple regression analysis and comparison to other hippocampal areas. *Cerebral Cortex*, 6:238–259.

Sharp, P. E. (2005). Regional distribution and variation in the firing properties of head direction cells. In Wiener, S. I. and Taube, J. S., editors, *Head Direction Cells and the Neural Mechanisms of Spatial Orientation*. MIT Press, Cambridge, MA.

Sharp, P. E., Tinkelman, A., and Cho, J. (2001). Angular velocity and head direction signals recorded from the dorsal tegmental nucleus of gudden in the rat: Implications for path integration in the head direction cell circuit. *Behavioral Neuroscience*, 115:571–588.

Shibata, H. (1993). Direct projections from the anterior thalamic nuclei to the retrohippocampal region in the rat. *Journal of Comparative Neurology*, 337(3):431–445.

Shibata, H. (1998). Organization of projections of rat retrosplenial cortex to the anterior thalamic nuclei. *European Journal of Neuroscience*, 10(10):3210–3219.

Skaggs, W. E., Knierim, J. J., Kudrimoti, H. S., and McNaughton, B. L. (1995). A model of the neural basis of the rat's sense of direction. In Tesauro, G., Touretzky, D., and Leen, T., editors, *Advances in Neural Information Processing Systems*, volume 7, pages 173–180. MIT Press, Cambridge, MA.

Smith, D. and Bulkin, D. (2014). The form and function of hippocampal context representations. *Neuroscience & Biobehavioral Reviews*, 40:52–61.

Smith, P., Geddes, L., Baek, J., Darlington, C., and Zheng, Y. (2010). Modulation of memory by vestibular lesions and galvanic vestibular stimulation. *Frontiers in Neurology*, 1:1–141.

Song, P. and Wang, X.-J. (2003). A three-population attractor network model of rodent head direction system without recurrent excitation. *Society of Neuroscience Abstracts*, 29:939.3.

Song, P. and Wang, X.-J. (2005). Angular path integration by moving "hill of activity": A spiking neuron model without recurrent excitation of the head-direction system. *Journal of Neuroscience*, 25:1002–1014.

Spiers, H. and Barry, C. (2015). Neural systems supporting navigation. *Current Opinion in Behavioral Sciences*, 1:47–55.

Stackman, R., Golob, E., Bassett, J., and Taube, J. (2003). Passive transport disrupts directional path integration by rat head direction cells. *Journal of Neurophysiology*, 90(5):2862–2874.

Stackman, R. and Herbert, A. (2002). Rats with lesions of the vestibular system require a visual landmark for spatial navigation. *Behavioural Brain Research*, 128(1):27–40.

Stackman, R., Lora, J., and Williams, S. (2012). Directional responding of c57bl/6j mice in the morris water maze is influenced by visual and vestibular cues and is dependent on the anterior thalamic nuclei. *The Journal of Neuroscience*, 32(30):10211–10225.

Stackman, R., Tullman, M., and Taube, J. (2000). Maintenance of rat head direction cell firing during locomotion in the vertical plane. *Journal of Neurophysiology*, 83(1):393–405.

Stackman, R. W. (2011). Behavioural correlates of neuronal activity recorded as single-units: Promises and pitfalls as illustrated by the rodent head direction cell system. In *Electrophysiological Recording Techniques*. Humana Press, New York, NY.

Stackman, R. W. and Taube, J. S. (1997). Firing properties of head direction cells in the rat anterior thalamic nucelus: Dependence on vestibular input. *Journal of Neuroscience*, 17:4349–4358.

Stackman, R. W. and Taube, J. S. (1998). Firing properties of rat lateral mammillary single units: Head direction, head pitch, and angular head velocity. *Journal of Neuroscience*, 18:9020–9037.

Stackman, R. W. and Zugaro, M. B. (2005). Self-motion cues and resolving intermodality conflicts: Head direction cells, place cells, and behavior. In Wiener, S. I. and Taube, J. S., editors, *Head Direction Cells and the Neural Mechanisms of Spatial Orientation*, pages 137–162. MIT Press, Cambridge, MA.

Stratton, P., Milford, M., Wyeth, G., and Wiles, J. (2011). Using strategic movement to calibrate and neural compass: A spiking network for tracking head direction in rats and robots. *PLoS ONE*, 6:1–15.

Stratton, P., Wyeth, G., and Wiles, J. (2010). Calibration of the head direction network: A role for symmetric angular head velocity cells. *Journal of Computational Neuroscience*, 28:527–538.

Stringer, S. M., Perry, G., Rolls, E. T., and Proske, J. H. (2006). Learning invariant object recognition in the visual system with continuous transformations. *Biological Cybernetics*, 94:128–142.

Stringer, S. M. and Rolls, E. T. (2006). Self-organizing path integration using a linked continuous attractor and competitive network: Path integration of head direction. *Network: Computation in Neural Systems*, 17:419–445.

Stringer, S. M., Rolls, E. T., and Trappenberg, T. P. (2004). Self-organising continuous attractor networks with multiple activity packets, and the representation of space. *Neural Networks*, 17:5–27.

Stringer, S. M., Trappenberg, T. P., Rolls, E. T., and De Araujo, I. E. T. (2002). Self-organizing continuous attractor networks and path integration: One dimensional models of head direction cells. *Network: Compuation in Neural Systems*, 13:217–242.

Tahvildari, B., Fransén, E., Alonso, A., and Hasselmo, M. (2007). Switching between on and off states of persistent activity in lateral entorhinal layer iii neurons. *Hippocampus*, 17(4):257–263.

Tan, H., Bassett, J., OKeefe, J., Cacucci, F., and Wills, T. (2015). The development of the head direction system before eye opening in the rat. *Current Biology*, 25(4):479–483.

Taube, J. and Muller, R. (1998). Comparisons of head direction cell activity in the postsubiculum and anterior thalamus of freely moving rats. *Hippocampus*, 8(2):87–108.

Taube, J., Stackman, R., Calton, J., and Oman, C. (2004). Rat head direction cell responses in zero-gravity parabolic flight. *Journal of Neurophysiology*, 92(5):2887–2997.

Taube, J. S. (1995). Head direction cells recorded in the anterior thalamic nuclei of freely moving rats. *Journal of Neuroscience*, 15:70–86.

Taube, J. S. (1998). Head direction cells and the neurophysiological basis for a sense of direction. *Progress in Neurobiology*, 55:225–256.

Taube, J. S. (2007). The head direction signal: Origins and sensory-motor integration. *Annual Review of Neuroscience*, 30:181–207.

Taube, J. S. and Bassett, J. P. (2003). Persistent neural activity in head direction cells. *Cerebral Cortex*, 13(11):1162–1172.

Taube, J. S. and Burton, H. L. (1995). Head direction cell activity monitored in a novel environment and during a cue conflict situation. *Journal of Neurophysiology*, 74:1953–1971.

Taube, J. S., Goodridge, J. P., Golob, E. G., Dudchenko, P. A., and Stackman, R. W. (1996). Processing the head direction signal: A review and commentary. *Brain Research Bulletin*, 40:477–486.

Taube, J. S., Muller, R. U., and Ranck, Jr., J. B. (1990a). Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *Journal of Neuroscience*, 10:420–435.

Taube, J. S., Muller, R. U., and Ranck, Jr., J. B. (1990b). Head-direction cells recorded from the postsubiculum in freely moving rats. II. Effects of environmental manipulations. *Journal of Neuroscience*, 10:436–447.

Thorpe, S., Delorme, A., and Van Rullen, R. (2001). Spike-based strategies for rapid processing. *Neural Networks*, 14(6):715–725.

Thorpe, S., Fize, D., and Marlot, C. (1996). Speed of processing in the human visual system. *Nature*, 381(6582):520–522.

Touretzky, D. and Redish, A. (1996). Theory of rodent navigation based on interacting representations of space. *Hippocampus*, 6(3):247–270.

Touretzky, D. S. (2005). Attractor network models of head direction cells. In Wiener, S. I. and Taube, J. S., editors, *Head Direction Cells and the Neural Mechanisms of Spatial Orientation*, pages 411–432. MIT Press, Cambridge, MA.

Tromans, J., Page, H., and Stringer, S. (2012). Learning separate visual representations of independently rotating objects. *Network: Computation in Neural Systems*, 23(1-2):1–23.

Tsanov, M., Chah, E., Noor, M., Egan, C., Reilly, R., Aggleton, J., Erichsen, J., Vann, S., and O'Mara, S. (2014). The irregular firing properties of thalamic head direction cells mediate turn-specific modulation of the directional tuning curve. *Journal of Neurophysiology*, 112(9):2316–2331.

Tsanov, M., Chah, E., Vann, S., Reilly, R., Erichsen, J., Aggleton, J., and O'Mara, S. (2011). Theta-modulated head direction cells in the rat anterior thalamus. *The Journal of Neuroscience*, 31(26):9489–9502.

Tsodkys, M. and Sejnowski, T. J. (1995). Rapid state switching in balanced cortical network models. *Network: Compuation in Neural Systems*, 6(2):111–124.

Tsodkys, M., Skaggs, W., Sejnowski, T. J., and McNaughton, B. L. (1996). Population dynamics and theta rhythm phase precession of hippocampal place cell firing: A spiking neuron model. *Hippocamps*, 6:271–280.

Tsodyks, M., Pawelzik, K., and Markram, H. (1998). Neural networks with dynamic synapses. *Neural Computation*, 10(4):821–835.

Uchino, Y., Sasaki, M., Sato, H., Bai, R., and Kawamoto, E. (2005). Otolith and canal integration on single vestibular neurons in cats. *Experimental Brain Research*, 164(3):271–285.

Valerio, S., Clark, B., Chan, J., Frost, C., Harris, M., and Taube, J. (2010). Directional learning, but no spatial mapping by rats performing a navigational task in an inverted orientation. *Neurobiology of Learning and Memory*, 93(4):495–505.

Valerio, S. and Taube, J. (2012). Path integration: How the head direction signal maintains and corrects spatial orientation. *Nature neuroscience*, 15(10):1445–1453.

Van Cauter, T., Poucet, B., and Save, E. (2008). Unstable CA1 place cell representation in rats with entorhinal cortex lesions. *European Journal of Neuroscience*, 27(8):1933–1946.

van der Meer, M., Knierim, J., Yoganarasimha, D., Wood, E., and van Rossum, M. (2007). Anticipation in the rodent head direction system can be explained by an interaction of head movements and vestibular firing properties. *Journal of Neurophysiology*, 98(4):1883–1897.

Van Groen, T. and Wyss, J. (1990). Extrinsic projections from area ca1 of the rat hippocampus: olfactory, cortical, subcortical, and bilateral hippocampal formation projections. *Journal of Comparative Neurology*, 302(3):515–528.

van Groen, T. and Wyss, J. (1992). Connections of the retrosplenial dysgranular cortex in the rat. *Journal of Comparative Neurology*, 315(2):200–216.

Van Groen, T. and Wyss, J. (2003). Connections of the retrosplenial granular b cortex in the rat. *Journal of Comparative Neurology*, 463(3):249–263.

Vann, S. and Aggleton, J. (2004). The mammillary bodies: Two memory systems in one? *Nature Reviews Neuroscience*, 5(1):35–44.

Vann, S. and Aggleton, J. (2005). Selective dysgranular retrosplenial cortex lesions in rats disrupt allocentric performance of the radial-arm maze task. *Behavioral Neuroscience*, 119(6):1682.

Vann, S., Aggleton, J., and Maguire, E. (2009). What does the retrosplenial cortex do? *Nature Reviews Neuroscience*, 10(11):792–802.

Vass, L. and Epstein, R. (2013). Abstract representations of location and facing direction in the human brain. *The Journal of Neuroscience*, 33(14):6133–6142.

Vogels, T., Froemke, R., Doyon, N., Gilson, M., Haas, J., Liu, R., Maffei, A., Miller, P., Wierenga, C., and Woodin, M. (2013). Inhibitory synaptic plasticity: spike timing-dependence and putative network function. *Frontiers in neural circuits*, 7.

Wallace, D., Choudhry, S., and Martin, M. (2006). Comparative analysis of movement characteristics during dead-reckoning-based navigation in humans and rats. *Journal of Comparative Psychology*, 120(4):331.

Wallace, D., Köppen, J., Jones, J., Winter, S., and Wagner, S. (2010). Navigating with fingers and feet: Analysis of human (homo sapiens) and rat (rattus norvegicus) movement organization during nonvisual spatial tasks. *Journal of Comparative Psychology*, 124(4):381.

Wallis, G. and Rolls, E. T. (1997). Invariant face and object recognition in the visual system. *Progress in Neurobiology*, 51:167–194.

Walters, D. M. (2011). *The Computational Neuroscience of Head Direction Cells*. PhD thesis, University of Oxford.

Walters, D. M., Stringer, S., and Rolls, E. (2013). Path integration of head direction: Updating a packet of neural activity at the correct speed using axonal conduction delays. *PLoS ONE*, 8(3):e58330.

Walters, D. M. and Stringer, S. M. (2010). Path integration of head direction: Updating a packet of neural activity at the correct speed using neuronal time constants. *Biological Cybernetics*, 103:21–41.

Warburton, E., Morgan, A., Baird, A., Muir, J., and Aggleton, J. (1999). Does pretraining spare the spatial deficit associated with anterior thalamic damage in rats? *Behavioral Neuroscience*, 113(5):956.

Whitlock, J. and Derdikman, D. (2012). Head direction maps remain stable despite grid map fragmentation. *Frontiers in Neural Circuits*, 6:9.

Wiener, S., Berthoz, A., and Zugaro, M. (2002). Multisensory processing in the elaboration of place and head direction responses by limbic system neurons. *Cognitive Brain Research*, 14(1):75–90.

Wiener, S. I. (1993). Spatial and behavioral correlates of striatal neurons in rats performing a self-initiated navigation task. *Journal of Neuroscience*, 13:3802–3817.

Wiener, S. I., Korshunov, V., Garcia, R., and Berthoz, A. (1995). Intertial, substratal and landmark cue control of hippocampal ca1 place cell activity. *European Journal of Neuroscience*, 7:2206–2219.

Wilber, A., Clark, B., Forster, T., Tatsuno, M., and McNaughton, B. (2014). Interaction of egocentric and world-centered reference frames in the rat posterior parietal cortex. *The Journal of Neuroscience*, 34(16):5431–5446.

Wills, T., Lever, C., Cacucci, F., Burgess, N., and O'Keefe, J. (2005). Attractor dynamics in the hippocampal representation of the local environment. *Science*, 308(5723):873–876.

Wills, T. J., Cacucci, F., Burgess, N., and O'Keefe, J. (2010). Development of the hippocampal cognitive map in preweanling rats. *Science*, 328:1573–1576.

Winograd, M., Destexhe, A., and Sanchez-Vives, M. (2008). Hyperpolarization-activated graded persistent activity in the prefrontal cortex. *Proceedings of the National Academy of Sciences of the USA*, 105:7298–7303.

Winter, S., Clark, B., and Taube, J. (2015). Disruption of the head direction cell network impairs the parahippocampal grid cell signal. *Science*, 347(6224):870–874.

Winter, S. and Taube, J. (2014). Head direction cells: from generation to integration. In *Space, Time and Memory in the Hippocampal Formation*, pages 83–106. Springer.

Winter, S., Wagner, S., McMillin, J., and Wallace, D. (2011). Mammillothalamic tract lesions disrupt dead reckoning in the rat. *European Journal of Neuroscience*, 33(2):371–381.

Wirtshafter, D. and Stratford, T. (1993). Evidence for gabaergic projections from the tegmental nuclei of gudden to the mammillary body in the rat. *Brain Research*, 630(1):188–194.

Wyss, J. and Van Groen, T. (1992). Connections between the retrosplenial cortex and the hippocampal formation in the rat: a review. *Hippocampus*, 2(1):1–11.

Xie, X., Hahnloser, R. H. R., and Seung, H. S. (2002). Double-ring network model of the head-direction system. *Physical Review E*, 66:041902.

Yoder, R., Peck, J., and Taube, J. (2015). Visual landmark information gains control of the head direction signal at the lateral mammillary nuclei. *The Journal of Neuroscience*, 35(4):1354–1367.

Yoder, R. and Taube, J. (2008). The postsubiculum provides visual landmark control to the head direction signal at the lateral mammillary nuclei. In *Soc Neurosci Abstr*, volume 90.

Yoder, R. and Taube, J. (2009). Head direction cell activity in mice: robust directional signal depends on intact otolith organs. *The Journal of Neuroscience*, 29(4):1061–1076.

Yoder, R. and Taube, J. (2014). The vestibular contribution to the head direction signal and navigation. *Frontiers in Integrative Neuroscience*, 8:32–32.

Yoganarasimha, D. and Knierim, J. J. (2005). Coupling between place cells and head direction cells during relative translations and rotations of distal landmarks. *Experimental Brain Research*, 160:344–359.

Yoganarasimha, D., Yu, X., and Knierim, J. J. (2006). Head direction cells maintain internal coherence during conflicting proximal and distal cue rotations: Comparison with hippocampal place cells. *Journal of Neuroscience*, 26:622–631.

York, L. and Van Rossum, M. (2009). Recurrent networks with short term synaptic depression. *Journal of Computational Neuroscience*, 27(3):607–620.

Yoshida, M., Fransén, E., and Hasselmo, M. (2008). mglur-dependent persistent firing in entorhinal cortex layer iii neurons. *European Journal of Neuroscience*, 28(6):1116–1126.

Yoshida, M. and Hasselmo, M. (2009). Persistent firing supported by an intrinsic cellular mechanism in a component of the head direction system. *The Journal of neuroscience*, 29(15):4945–4952.

Zar, J. (1999). *Biostatistical Analysis*. Pearson Education India.

Zeidman, P. and Bullinaria, J. (2008). Neural models of head-direction cells. In French, R. M. and Thomas, E., editors, *From Associations to Rules: Connectionist Models of Behavior and Cognition,*, pages 165–177. World Scientific Publishing Company, Singapore.

Zhang, K. (1996). Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: A theory. *Journal of Neuroscience*, 16:2112–2126.

Zhang, S. and Manahan-Vaughan, D. (2015). Spatial olfactory learning contributes to place field formation in the hippocampus. *Cerebral Cortex*, 25(2):423–432.

Zucker, R. (1989). Short-term synaptic plasticity. *Annual Review of Neuroscience*, 12(1):13–31.

Zugaro, M. B., Arleo, A., Berthoz, A., and Wiener, S. I. (2003). Rapid spatial reorientation and head direction cells. *Journal of Neuroscience*, 23:3478–3482.

Zugaro, M. B., Arleo, A., Déjean, C., Burguière, E., Khamassi, M., and Wiener, S. I. (2004). Rat anterodorsal thalamic head direction neurons depend upon dynamic visual signals to select anchoring landmark cues. *European Journal of Neuroscience*, 20:530–536.

Zugaro, M. B., Berthoz, A., and Wiener, S. (2002). Peak firing rates of rat anterodorsal thalamic head direction cells are higher during faster passive rotations. *Hippocampus*, 12(4):481–486.

Zugaro, M. B., Berthoz, A., and Wiener, S. I. (2001). Background, but not foreground, spatial cues are taken as references for head direction responses by rat anterodorsal thalamus neurons. *Journal of Neuroscience*, 21 RC154:1–5.

Zugaro, M. B., Tabuchi, E., and Wiener, S. I. (2000). Influence of conflicting visual, inertial and substratal cues on head direction cell activity. *Experimental Brain Research*, 133:198–208.