

Neural Coding via Transmission Delay Coincidence Detectors: An Embodied Approach

by

Francis Jeanson

A thesis submitted to the Faculty of Graduate
and Postdoctoral Affairs in partial fulfilment of
the requirements for the degree of
Doctor of Philosophy

supervised by

Dr. Anthony White

Institute of Cognitive Science
Carleton University

© Copyright
Francis Jeanson, 2013

Abstract

This thesis aims to contribute to the field of cognition via a careful investigation at the mesoscopic level of neural organization and activity via spatiotemporal dynamics to subserve reactive and adaptive behaviour. In particular, coding via coincidence detection with propagation delays is adopted as the primary mechanism to investigate complex neural dynamics. Initially hypothesized by Moshe Abeles in the early 1980's, coincidence detection enables neurones to respond by emitting an action potential (a spike) only when other input neurones spike in a precise order. In contrast to the traditional interpretation of neural behaviour whereby cells integrate inputs over long periods of time, coincidence detectors are sensitive to changes at millisecond or sub-millisecond time scales. However, networks of coincidence detectors that incorporate propagation delays remain poorly understood with respect to dynamics and functional application for cognitive tasks. After introducing the functional and biological evidence that motivated this research, we explore existing work related to spatiotemporal coding and the neurone and network parameters that influence their behaviour. We then present a discrete neural network model that is used to expose the relationship between structural parameters and network dynamics. After applying these to a reactive light-seeking robot task, we find network parameters that enable dynamic memory storage of input spike patterns. Making use of these dynamics, we then introduce a method for decoding these memories in a modular network architecture and test this approach for robot control in memory maze tasks. Furthermore, the limits of the potentially high memory capacity of these networks is then tested by empirically evaluating both the noise tolerance and the memory capacity of these networks. To complement this empirical work, we develop a set of formal expressions which attempt to approximate analytically the amount of activity of these networks and the probability for any given spike pattern to be expressed by them.

Acknowledgements

This work was made possible by the relentless and continued support of many people and organizations. First and foremost, I would like to thank my outstanding and tireless thesis advisor Dr. Anthony White. Without his curiosity, interest, expertise, and willingness to review countless documents this thesis would not be before you today. Tony has provided significant practical, insightful, and scientific contributions to this work for which I will always be grateful. Tony has also donned the role of mentor and instilled energy into my work when I needed it the most. The itinerary of this doctoral degree has been a unique and immensely rewarding experience thanks to him.

Supportive and motivational even before I entertained the thought of pursuing doctoral studies is my partner Ginnie Chen. Despite many periods of long absence over the years, Ginnie has stood solidly by my side with tireless support. Not only did she motivate me to go further but also took interest in my work which stimulated my inquiry. Without her support this thesis would not have been possible and I owe much of the energy and dedication that was required for this work to her. I will always be grateful for the sacrifices she had to make and hurdles she had to overcome to support me in this endeavour.

Fundamentally, my mother and father, Michèle and Patrick, are at the root of the curiosity for nature, complexity, and intelligence that has bourgeoned within me since childhood. By fostering my endless inquiries, passionate discussions, and stimulating voyages, they have opened my perspectives on life and science in ways that nurtured reason, creativity, and compassion. Furthermore, they have provided substantial moral, financial, and intellectual support. This thesis is, in large part, the product of the countless years of their backing and engagement.

My sisters Béatrice, Isabelle, and Nathalie have also played a key role in my passion for science and dedication to work. Their love and support has provided a solid ground upon which I built my academic interest. Each have been unique role models, teaching me how to care, how to be strong, and how to become what our

heart guides us to become. Our unique solidarity is something that I will always cherish. Their children and life partners are also amazing joys that have enriched my life and I thank them all for their presence in moments of gathering.

My extended family has also played a key role in the undertaking of this research and thesis work. Critical thinking and intellectual inquiry have often stemmed from the times we spend together. In particular, my uncles Jean-Pierre and Jean-Luc Pilaprat as well as my cousins Jean-Christophe and Michael who have been amazing influences with respect to my curiosity and direction of research.

Many friends have been amazingly supportive over the years. In particular I would like to thank Hany Fahim and Kris Kostecky and their company VM Farms for providing valuable computational resources at times when neural network simulation and analysis became intensive. Their virtualized cloud based service is outstanding and their technical support unmatched.

I would also like to acknowledge Dr. Jun Luo and Dr. Luc Berthouze who have guided me intellectually during my B.Sc. and M.Sc. research years. The common passion and interests that we shared along with their outstanding academic guidance have stimulated many ideas and research directions within this work.

Finally, I would like to thank Carleton University and, in particular, the Institute for Cognitive Science and its faculty and staff members who have continuously supported and guided me through the process of completing this doctoral degree and research. The intellectual and human environment of the institute is outstanding, with a program that genuinely compels students to become unique researchers in the disciplines integrated within cognitive science.

Contents

1	Introduction	1
1.1	Neural Coding for Cognition	2
1.2	Neural Coding via Firing Rate	4
1.3	Spatiotemporal Coding	7
1.4	Biological Basis of Coincidence Detection	11
1.5	Goal	14
1.6	Objectives	14
1.7	Scope	15
1.8	Problem Statement	16
1.9	Contributions	17
1.10	Organization of the Document	17
2	Related Work	19
2.1	Coding via Precise Spike Timing	20
2.1.1	Synfire Chains and Braids	20
2.1.2	Discrepancies in the Rate Coding Paradigm	23
2.1.3	Dynamical Cell Assemblies Hypothesis	24
2.1.4	Practical Aspects of Dynamical Cell Assemblies	26
2.1.5	Alternative Codes	28
2.1.6	Time Coding and Polychronization	31
2.2	Neural and Network Parameters of Influence	34
2.2.1	The Membrane Time Constant	35
2.2.2	The Firing Threshold	37
2.2.3	Connectivity and Synaptic Strengths	38
2.2.4	Refractory Period	41
2.2.5	Input, Feedback and Noise	42
2.2.6	Transmission Delays	44

2.2.7	Informational Capacity	47
2.3	Summary	49
3	Basic Properties of Delay Coincidence Detection Neural Networks	51
3.1	Distinguishing Behavioural Complexities	52
3.1.1	From Reaction to Adaptation	52
3.1.2	Towards a Formal Characterization	54
3.1.3	Final Considerations	56
3.2	Computational Simulation of Reactive Behaviour	56
3.2.1	Dynamical Analysis of Delays in Simple Networks	57
3.2.2	Delay Coincidence Detection Network Model	58
3.2.3	Preliminary Test Model Parameters	59
3.2.4	Results	60
3.2.5	Evolving Delays for a Light Seeking Task	66
3.3	Early Theoretical Considerations	71
3.3.1	Sustained and Stable Firing	74
3.4	Summary	76
4	Dynamic Memory: Beyond Sensorimotor Coordination	78
4.1	Why Dynamic Memory?	78
4.2	Delay Mapping	80
4.2.1	A Procedure for Delay Mapping	80
4.2.2	Collision Removal	83
4.3	T-maze	84
4.3.1	Sensorimotor Memory Task	84
4.3.2	Sensory-Motor Network	86
4.3.3	Memory and Inhibitory Networks	87
4.3.4	Results	88
4.4	Ambiguity	91
4.4.1	Assigning Decoding Thresholds	91
4.4.2	Defining Ambiguities	92
4.4.3	Finding Ambiguities	93
4.5	F-Maze	94
4.5.1	Procedure	94
4.5.2	Results	95
4.6	Summary	98

5	Dynamic Memory: Noise Tolerance and Memory Capacity	100
5.1	Noise Tolerance	100
5.1.1	Degrees of Freedom	101
5.1.2	Noise Tolerance Estimation	103
5.1.2.1	Noise with Ambiguity	104
5.1.3	Noise in Correlated Spiking Activity	107
5.1.3.1	Noise via Prolonged Spike-Removal	107
5.1.3.2	Noise via Temporary Spike-Removal	115
5.1.3.3	Noise via Prolonged Misfiring	121
5.1.3.4	Discussion	132
5.2	Effective Memory of Delay Coincidence Detection Networks	133
5.2.1	Input Correlation	134
5.2.1.1	Preliminary Investigation	134
5.2.1.2	Input Ambiguity	135
5.2.2	Effective Memory	137
5.2.2.1	Spatiotemporal States	137
5.2.2.2	Empirical Evaluation of Effective Memory	138
5.2.2.3	Probability of Failure	143
5.2.2.4	Weight of Degree of Ambiguity	145
5.2.3	Effective Memory and Noise Interaction	148
5.2.3.1	Estimation of Effective Memory with Noise	148
5.2.3.2	The Decoding Threshold Tradeoff	150
5.3	Summary	151
6	Formal Investigations	154
6.1	Estimating Network Activity	155
6.2	Probability of Spiking and the Number of Firing-chains	159
6.2.1	Effective Connectivity	160
6.2.2	Temporally Effective Afferent Spikes	162
6.2.2.1	Temporal Breadth	162
6.2.2.2	Missing Afferent Spikes	164
6.2.2.3	Spatial Breadth	165
6.2.2.4	Application and Limitations	167
6.2.3	Spiking Probability with Refractory Periods	169
6.2.4	Probability of Firing States	171
6.3	Summary	174

7	Conclusion and Future Work	176
7.1	Conclusion	176
7.2	Future Work	180
7.2.1	Conductance Based Modelling of Pyramidal Coincidence De- tection	180
7.2.2	Hebbian Plasticity to Regulate Capacity vs. Noise	180
7.2.3	Stimulus Convolution	182
7.2.4	Noise and Efficient Coding	183
	References	186

List of Figures

1.1	Simple two cell assembly of coincidence detectors.	10
2.1	Synfire chain: Synchronized firing in a population of cells projects to a second population in which cells fire synchronously which in turn trigger synchronous firing in a third population.	21
2.2	Synfire braid: Transmission times between cells can vary such that the conduction times for one path between any two cells is equal to all other paths between those two cells. Here path $d_1 + d_2 = d_3$	22
2.3	Approximate reproduction of the spatial pattern A, B, C, and D as well as temporal patterns Program 1 and Program 2 used by Watanabe et al., 1998.	29
2.4	Left: A CD neurone with threshold $T = 2$ will still fire if jitter occurs on redundant afferents (here the 3^{rd} afferent). Right: A CD neurone with threshold $T = 2$ can fire despite loss of an input (2^{nd} afferent) if noisy input (3^{rd} afferent) can fill-in for the missing spike.	31
2.5	Phase-locking: Early arbitrary spikes become entrained by the background oscillatory activity.	32
2.6	Illustration of two different neural activation responses. Rapid rise and fall occurs for cells with short membrane time constants. A longer response time takes place for longer membrane time constants. In reality this difference can vary 10 folds or more.	37
3.1	Ultrastability: an agent's reactive system R can react to environmental changes autonomously but only within strict bounds. If vital internal variables (represented by the gauge) are sufficiently threatened then a secondary system S must actively regulate the reactive response system R to maintain these variables within viable bounds.	55

3.2	Narrow delay range activity. From top to bottom delay ranges 0ts, 3ts and 6ts. Left plots correspond to spike rasters and right plots correspond to spike count histograms.	61
3.3	Broad delay range activity. From top to bottom delay ranges 9ts, 12ts and 15ts. Left plots correspond to spike rasters, center plots correspond to spike count histograms, right plots correspond to the population power spectrum with the horizontal line representing 3 standard deviations of the total power.	63
3.4	Average last time of spiking amongst various delay range conditions. .	64
3.5	Spiking activity divergence between two similar input conditions on the same network.	65
3.6	Illustration of the simulated robot with two light sensors and two wheels. In total, 25 neurones controlled the robot with 4 input neurones per light sensor and 4 neurones per motor effector.	67
3.7	Illustrated genotype for a single agent controller. Genes with value 0 defined the non-existence of a connection between the corresponding two cells. A non-zero value was transcribed as an existing connection with a delay scaled to the delay range of interest during the experiment and an added minimum delay.	68
3.8	Average population fitness for the three tested delay conditions. Lower values represent greater fitness since fitness was directly measured as the average distance from the light during each individual's life-time.	68
3.9	Example light-seeking behaviour for an agent with narrow delays (0 to 3 time steps).	69
3.10	Example light-seeking behaviour for an agent with broad delays (15 to 24 time steps).	70
3.11	Example light-seeking behaviour for an agent with narrow delays but high delay offset (60 to 63 time steps).	70
3.12	Memory: a sub-set of three cells (in purple) fire within a specific temporal structure leading to a firing-chain that triggers at a later point in time the identical initial three cell spike pattern. This results in a self-sustaining spatiotemporal spike pattern in the form of a limit-cycle.	73
3.13	Association: Two stimulus driven firing-chains A and B trigger a third self-sustaining firing-chain C.	74

3.14	Left: Broad delay spike activity with low connectivity of 60% leads to unpredictable spiking. Right: Broad delay spike activity with high connectivity of 100% leads to stable limit-cycle response patterns. . .	75
3.15	Left: Spiking stability s for each cell in a network response with 60% connectivity and broad delays. Right: Spiking stability s for a network with 100% connectivity and broad delays.	76
4.1	Memory network and decoding network configuration. A DCD memory network responding via broad delay limit cycles to some inputs (A ... Z) can be decoded into spatial patterns via delay mapping to a target network of coincidence detectors.	80
4.2	Collision removal and delay mapping. Collision removal is a rudimentary approach to reducing ambiguous decoding because only connections from cells that spike at the same time (here marked by an x) have efferent connections removed to prevent ambiguous spikes from reaching the decoding network. Despite this, many ambiguities can remain using this approach and it also leads to pattern weakening requiring selective ‘filling-in’.	81
4.3	T-maze working memory paradigm.	85
4.4	Illustrated network model with Memory, Inhibitory, and Sensorimotor networks.	87
4.5	Left: spike raster plot of activity in the INH network during presentation of context stimulus A and cue stimulus X. Right: robot path in the AX stimulus condition. Lines originating from the robot mark the 120° field of view.	88
4.6	Left: spike raster plot of activity in the INH network during presentation of context stimulus B and cue stimulus Y. Right: robot path in the BY stimulus condition. Lines originating from the robot mark the 120° field of view.	89
4.7	Two spike patterns with a single ambiguity between both patterns. The circled spike pairs lead to two inner-deltas that are equal for both patterns, hence the ambiguity.	91

4.8	Second method for finding the maximum ambiguity between two patterns. Between pattern spike-time differences are calculated for each of the N neurones. Here, the largest set has two differences (cells 4 and 3) of time step +1, followed by a set with one difference (cell 2) of -1 and a set with one difference (cell 1) of +3. Hence a maximum degree of ambiguity of two corresponding to the size of the largest set.	93
4.9	F-maze.	95
4.10	Left: Robot path on a trial in the F-maze with stimuli AX-CV. Right: Inhibitory network activity.	96
4.11	Left: Robot path on a trial in the F-maze with stimuli AX-DW. Right: Inhibitory network activity.	97
5.1	Maximum (optimistic) and minimum (pessimistic) bounds for the degrees-of-freedom of a DCD network of 25 neurones with a single spike per cell for period λ . As the number of ambiguities of degree 2 increases the degree-of-freedom (df) decreases in a log like fashion.	102
5.2	Optimistic and pessimistic number of available sub-patterns (S) with at least two spike times in a network of 25 cells as ambiguities increase. 105	
5.3	Optimistic and pessimistic noise tolerance in a network of 25 cells with fixed decoding threshold $T = 2$ as ambiguities increase.	106
5.4	Sample Spike removal raster plots. No-noise raster (blue) is superimposed by noisy raster (red) when spike-times coincide. a. Neurone 0 is prevented from spiking – this led to 7 total misfiring neurones (including cell 0). b. Neurone 0 and 1 are prevented – led to no additional neurone misfiring. c. Neurones 0 to 2 are prevented – led to 8 neurones misfiring. d. Neurones 0 to 3 are prevented – led to 9 neurones misfiring. e. Neurones 0 to 4 are prevented – led to 23 neurones misfiring.	108
5.5	Spiking stability in prolonged spike-removal via <i>averaged instantaneous spike time differences</i> for a single network with identical input between noise scenarios. a.: no noise removal. b.: neurones 0–3 are inhibited from time step 800 onwards. c.: neurones 0–7 are inhibited. d.: neurones 0–11 are inhibited. e.: neurones 0–15 are inhibited.	110
5.6	Degree of ambiguity (y-axis) as the number of neurones prevented from spiking (spike removal) increases (x-axis) amongst the 30 different networks. Critical reductions in ambiguity above 8 are circled.	113

5.7	Average maximum degree of ambiguities (AMA) and normalized AMA as the number of spiking neurones prevented from spiking is increased in the prolonged condition. Higher ambiguity denotes greater noise tolerance. Error bars display standard error across 30 trials.	114
5.8	Normalized average ambiguity in noise via spike removal for a prolonged period. A comparison is made between DCD networks of 25 neurones, 100 neurones, and 1000 neurones. Error bars show standard error across the average 30 networks.	116
5.9	Spiking stability in temporary spike-removal via averaged instantaneous spike time differences for a single network. a. : no noise removal. b. : neurones 0–7 are inhibited only once in the time step interval 800 and 811. c. : neurones 0–15 are inhibited. d. : neurones 0–24 are inhibited.	117
5.10	Degree of ambiguity (y-axis) as the number of neurones temporarily prevented from spiking increases (x-axis) amongst the 30 different networks. Critical reductions in ambiguity above 8 are circled.	118
5.11	Average and normalized maximum degree of ambiguities as the number of spiking neurones prevented from spiking is increased in the temporary condition. Higher ambiguity denotes greater noise tolerance. Error bars display standard error across 30 trials.	119
5.12	Normalized average ambiguity in noise via single spike removal (temporary). A comparison is made between DCD networks of 25 neurones, 100 neurones, and 1000 neurones. Error bars show standard error across the average 30 networks.	120
5.13	Spiking stability during prolonged misfiring via averaged instantaneous spike time differences for a single network. a. : no misfiring. b. : neurones 0–7 are set to misfire from time step 800 onwards. c. : neurones 0–15 set to misfire. d. : neurones 0–24 are set to misfire. . .	123
5.14	Degree of ambiguity (y-axis) as the number of neurones misfiring continuously increases (x-axis) amongst the 30 different networks. Critical reductions in ambiguity above 8 are circled.	124
5.15	Average and normalized maximum degree of ambiguities as the number of spiking neurones set to misfire is increased in the prolonged condition. Higher ambiguity denotes greater noise tolerance. Error bars display standard error across 30 trials.	125

5.16	Normalized average ambiguity in noise via prolonged spike misfiring. A comparison is made between DCD networks of 25 neurones, 100 neurones, and 1000 neurones. Error bars show standard error across the average 30 networks.	126
5.17	Spiking stability in temporary misfiring via averaged instantaneous spike time differences for a single network. a. : no misfiring. b. : neurones 0–7 are set to misfire within time step 800 to 811 with probability $1/\lambda$. c. : neurones 0–15 set to misfire. d. : neurones 0–24 are set to misfire.	128
5.18	Degree of ambiguity (y-axis) as the number of neurones misfiring only once increases (x-axis) amongst the 30 different networks. Critical reductions in ambiguity above 8 are circled.	129
5.19	Average and normalized maximum degree of ambiguities as the number of spiking neurones set to misfire is increased in the temporary condition. Higher ambiguity denotes greater noise tolerance. Error bars display standard error across 30 trials.	130
5.20	Normalized average ambiguity in noise via single spike misfiring. A comparison is made between DCD networks of 25 neurones, 100 neurones, and 1000 neurones. Error bars show standard error across the average 30 networks.	131
5.21	Sample raster showing divergence of network response despite similarity of the input. Both inputs are run on the same DCD network. Input to six neurones only differs by one cell as seen at time step 0.	135
5.22	Response correlation with respect to input correlation averaged across 10 random DCD networks. This reveals the degree of ambiguity between any two response patterns for inputs that have no identical input cells up to 5 identical input cells. Bars indicate standard deviations across the 10 trials. The number of comparisons with 5 identical cells is much less likely than with fewer common input cells.	137
5.23	Average maximum degree of ambiguity over 30 random networks of 25 neurones as the number of patterns retrieved increases up to a total of 88,550 patterns. Vertical error bars are standard deviations of the ambiguities. The solid line corresponds to the natural logarithmic regression with coefficients a, b and c. Other lines correspond to the regression of the ambiguities plus or minus multiples of the standard deviations.	141

5.24	Average maximum degree of ambiguity over 30 random networks of 50 neurones as the number of patterns retrieved increases up to a total of 88,550 patterns. Vertical bars show standard deviations. The solid line corresponds to the natural logarithmic regression with coefficients a, b and c. Other lines correspond to the regression of the ambiguities plus or minus multiples of the standard deviations.	143
5.25	Failure probability (0 to 1) averaged over the 30 networks. This measure shows an increase in the probability that two separate input patterns produce fully ambiguous results, i.e., $\Lambda = 25$. Vertical bars indicate standard deviations.	145
5.26	Averaged total proportional weights of degree of ambiguity amongst three different sizes of response subsets: 10, 100, and 1000 patterns. .	147
5.27	Relationship between noise and number of identifiable patterns in DCD networks of 25 neurones based on independent analysis of noise and effective memory.	150
6.1	Illustrated delay of incidence for 1, 2, and 3 time steps. Illustration a. shows that a_1 (Eq. 6.3) can be computed when the proportional activity a_0 is known at time step 0 and taking into account delays of one time step I_{D1} . Illustration b. shows how the proportional activity a_2 (Eq. 6.4) is dependent on delays of incidence I_{D1} and I_{D2} . Illustration c. shows how activity from incident connections delays with 1, 2, and 3 time steps away are used to compute a_3 (Eq. 6.5). Note: the number of incident connections is arbitrarily illustrated here as we assume low variability between neurones in the present activity estimation. . . .	156
6.2	Analytical estimate of network activity following Equation 6.8 in the narrow delay condition (plot a.) and the broad delay condition (plot b.).	159
6.3	The sum of combinations allows to conceptually reduce the number of afferent combinations to a single afferent.	161
6.4	Four afferents to a target. With a delay range $D_{range} = 3$ and assuming uniform distribution of delays, then each projection are one delay time step apart.	163
6.5	Four temporally distinct afferents and one spatially distinct afferent to a target with a delay range $D_{range} = 3$. Afferents a,b,c,d have distinct delays but afferent e has identical delay to afferent cell b.	166

6.6 Partial set of permutations of the temporal order of the four cells a, b, c, d, and the spatial cell e. Effective pairs are circled. Pairs of cells firing one step apart and in order with a single propagation delay time step difference are circled because they will lead to a coincidence of 2. Also spatial cells that have equal propagation delays and fire synchronously are circled (here only b and e since both have propagation delays of 39).167

6.7 Spiking probability distribution z as a function of the overall percentage of activity in the network a for the different effective temporal spiking probability conditions: p_{orders} , $p_{missing_tail}$, $p_{missing_middle}$, and $p_{spatial}$, for parameters $I = 24$, $T = 2$, $R = 10$, and $D = 30$. By taking into account the refractory scaling factor f , probability of spiking is reduced despite increasing activity. 171

7.1 A simple correlation table relating three problems (p1, p2, p3) to three symptoms (s1, s2, s3) and the translated network representation. . . . 184

Chapter 1

Introduction

Intelligent agency will ultimately have to be understood in its whole such that adequate low level and high level functions are integrated into a coherent framework to satisfy our inquiry into the mind, develop new technologies, and discover new therapies. A scientific approach to understanding intelligent agency aims to rationally and empirically decompose the nature of these levels until all have been revealed. The principal content of this thesis focusses on a neural mechanism which, we believe, exhibits promising dynamic, informational and ultimately functional characteristics for reactive and adaptive cognitive function in biological nervous systems. In particular, this mechanism presents features which, through careful arrangement could account for a multitude of cognitive phenomena; thus satisfying bottom-up criteria of parsimony, limited assumption, and compositionality. As such, this mechanism may one day constitute a fundamental principle of possible neural function. We believe, however, that important lessons beyond neural application can be learned from the neural mechanism advanced herein. Before delving into the biological and theoretical neural foundations that have spurred the focus of this thesis, a greater appreciation for the overarching approach taken through this thesis should be motivated.

Fundamental principles of intelligence have been advanced to date which aim to illustrate the structures, rules, or processes that may be used to construct intelligent systems. For example, notions such as learning, decision making, thinking, reasoning, sense making, planning, goal directed behaviour, problem solving, adaptation, and many more have been proposed as fundamental notions depicting intelligence. While many investigators in the field of cognition and intelligent agency may hold these with distinctive regard, we adopt the view that many, if not all, of these notions are, with varying degree, reconcilable. As is the objective of a cognitive science approach, relationships between neuroscience, adaptation, computer science, dynamical systems

theory, information theory and robotics are partially established in order to fulfill the goal of understanding how a biologically plausible neural mechanism could, in fact, play a substantial role in cognition. The achievement of this work was thus made possible by adopting a conceptual framework founded on many principles taken from these disciplines. Perhaps most fundamentally, we believe intelligent agency can be reasonably understood under the guise of adaptation. Without, failing to acknowledge the often criticized vacuous application of the term adaptation in the context of explaining the cause for particular structural or functional traits – such as spandrels (Gould & Lewontin, 1979), we adopt the notion of adaptation instead as the overarching *raison d'être* of intelligent agency. Among numerous constitutive mechanisms tied to the adaptation principle, including prediction and control, we pay particular attention to the simpler but more fundamental requirements for embodiment, reaction, and memory. Indeed, embodiment requires that all systems, living or not, coordinate their physical structure to initially compensate and later control the physical environment for survival. Indeed, by adopting an embodied approach to cognitive science, we not only achieve greater validation of a model's robustness, but also demonstrate genuine purpose for the intended function.

Here, we adopt an embodied approach to neural modelling to demonstrate the capacity of the neural mechanism under focus in simulated conditions eliciting not only specific dynamic characteristics but also functional properties for neural coordination and behavioural control. While, the embodied experiments presented herein are not meant to reflect the actual biological applications of the neural processing offered here, they help construct a more coherent and validated view of the functional properties of this neural mechanism.

1.1 Neural Coding for Cognition

The large majority of research in the fields of both cognitive science and neuroscience have approached neural coding via the classical paradigm known as *rate coding*. In this view, neurones fire a spike or a sequence of spikes based on the frequency and amplitude of action potentials it receives over a period of time. While different neurones may be more or less sensitive to particular input frequencies, the information of a network is contained in the frequency of the signals emitted by cells and is controlled via the connection strengths between neurones. Alternatively, neural coding can take place in radically different forms generally regrouped under the name *spatiotemporal coding*. Here, rather than locating information in the frequency of a neural

population's signal, information is located in distributed patterns of spikes across a population of neurones but also over a specific period of time. While certain classes of spatiotemporal codes are increasingly better understood, such as oscillations and synchrony, other forms, such as temporal coding via coincidence detection, remain elusive with respect to their dynamical, functional and informational characteristics. Despite some progress in the past 20 years, which have demonstrated promising applications of this principle, no work has validated their use in the context of embodied and situated tasks which is essential for an adequate account of cognition from a bottom-up approach.

Our understanding of a neural code is generally related to the notion of transforming a state of activity into an alternate state of activity within a set neural population. There are numerous dimensions of complexity that biological neural networks bear which have functional significance in this transformational process. These can be stratified as: structural characteristics of different neural tissue¹, neurone types², neurone dynamics³, and overall network dynamics⁴. Often, many of these characteristics have unintuitive consequences. For instance, developmental stages of the brain during childhood and adolescence will lead to synaptic pruning and neural death which, in fact, favours the formation of specialized function (M. C. Stevens, 2009). Despite significant loss of neuronal mass in the cortex during childhood, Luciana and Nelson, for instance, observed a significant improvement of prefrontal working memory in children between the ages of 5 and 7 years (Luciana & Nelson, 1998). At a more fundamental level, however, change in network activity can arise from either the *independent change* of individual cell activity or *correlated change* of activity in sets of cells. Interestingly, from the perspective of a single cell, incoming firing rates from different neurones may be completely independent, while for spatiotemporal cells, incoming spikes across different neurones but also across time must uphold minimum correlation.

In the following we decompose in greater detail the characteristics of both coding paradigms and motivate, with particular emphasis on coincidence detection, the importance of acquiring a better scientific understanding of its role in cognition.

¹This may include the connection topology, the propagation delays, vascularization, etc.

²Hundreds of different types of neurones are known based on differences in physiology, receptors and other characteristics.

³Such as refractory periods, threshold sensitivity, spike response curves, etc.

⁴For instance, sub-group activation, synchrony and other attractor states.

1.2 Neural Coding via Firing Rate

Since early recordings of neurone activity (Adrian, 1928) the dominant view of information coding in the brain sees rates of neural firing as the almost exclusive neural mechanism underlying perception, thought and action. Typically, experimental recordings focussed on the activity of cells over several hundred millisecond time scales. This made possible the correlating of statistical firing rates with psychological performance. This view was strongly supported by work on motor neurone activity where firing rates were found to have a clear correspondence with muscle flexion (Robinson, 1975; Shidara, Kawano, Gomi, & Kawato, 1993). Shidara et al., for instance, showed that high frequency firing Purkinje cells of the cerebellum behave as integrators of spike numbers while firing rate modulation controls eye position and acceleration (Shidara et al., 1993). Two distinct perspectives take rate coding as the central means for information coding in the nervous system: *place coding* and *distributed coding*.

The place coding view sees information as being encoded by a single cell's mean spike rate over a particular psychological time during an experiment (Barlow et al., 1972). This notion is characterized by the interpretation of cell function based on both their objective meaning: the spatial location of the firing cell, and subjective meaning: the rate of cell firing (Krüger & Becker, 1991). Here firing rates are averaged over the duration of psychological experiments which often last several hundred milliseconds. Furthermore, experiments are adjusted to elicit maximum firing rates of the recorded cells. From this view the functional role of a cell is inferred from an optimum stimulus: a stimulus that elicits a maximum average firing rate. Place coding thus motivates the hypothesis that subordinate cells transmit their signals to grandmother cells. These grandmother cells are thus understood to be feature specific and thought to play the principle role in pattern detection or action initiation. Face recognition cells, for instance, have been hypothesized to reside in the inferotemporal (IT) lobe (Rolls & Treves, 1990). Also in IT an icon alphabet is claimed to be implemented where single cells respond selectively to perceived characters (Tanaka, 1992). Although experimental recordings show reliable firing rates given specific stimuli, place coding remains restrictive regarding its view of the mode of operation of brain function. As noted by Fujii et al., averaging spike rates over entire experiments can make it impossible to detect intricate firing patterns (Fujii, Ito, Aihara, Ichinose, & Tsukada, 1996). For instance, temporal fluctuations in firing rates are commonly observed during psychological time scales, yet, statistical variations disappear when

averages are computed over the time of an experiment. Furthermore, such studies typically average over many experiments to obtain significance and therefore can not explain how the brain solves problems from a single presentation of the stimulus. Krüger and Becker have provided strong evidence against the notion of the functional significance of mean firing rates over psychological times (Krüger & Becker, 1991). While replicating the experimental findings of Hubel and Wiesel, which originally provided support to the place coding view with orientation-selective cells in the primary visual cortex (V1) (Hubel & Wiesel, 1968), Krüger and Becker found that spike averages over a 300ms time slot of a psychological experiment were insufficient to predict the activation of the recorded cells given a single stimulus. However, smaller time slots in the range of 20ms to 80ms could accurately do so with a performance of up to 80%, thus reinforcing the view that firing rate averages over long periods and many experiments does not capture the informational content and/or functional principle of cognition. More problematic to the place coding view is the assumption that grandmother cells exist for each representation of every object and each of those object's features. Such a hierarchical view thus implies that dedicated cells should be formed or tuned for each novel stimulus. According to von der Malsburg, such a mechanism would lead to a combinatorial explosion where cardinal cells would have to exist or be created and respond preferentially to a specific stimulus by integrating information coming from all modalities (von der Malsburg, 1994). The implausibility of such a mechanism and lack of evidence that such spatial specificity in the brain exists is unsupportive of the place coding view.

The second alternative to rate coding takes into account the combined effect of a large number of cells. The original concepts surrounding the distributed coding view emerged from Donald Hebb's cell assembly hypothesis (Hebb, 1949). According to this hypothesis the fundamental brain encoding is based on the group activity of neural cells. Furthermore, the assembly serves to produce a mean firing rate and is identified by it. Finally, assemblies are formed by the Hebbian learning rule: the simultaneous increase in firing rate of a set of cells (von der Malsburg, 1994). This view sees the group activity of an ensemble as corresponding to a "symbol of the mind" and has served as the foundational principle for the connectionist approach to artificial intelligence (Rumelhart & McClelland, 1986). In contrast to the place coding view, distributed coding does not appeal to the single cell equal single function theory. Here a cell may play more than one role by participating in the coding of multiple patterns. This suggests that a wide variety of inputs can be encoded in various subsets of a single assembly thereby escaping the issue of combinatorial explosion. However, cell

assemblies are still understood to encode stimuli using first order firing rate statistics. As a result of this, the ability for such networks to detect variation between two inputs depends essentially on the change of firing rate of a cell assembly. This coding mechanism can thus lead to a “superposition catastrophe” and fails to account for the problem of binding stimuli (von der Malsburg, 1994). Superposition for this view is problematic because it is impossible to distinguish two assemblies activated in the same area by two simultaneously presented stimuli if only the change in average firing rate is detectable from an assembly’s activity. In other words, it is impossible to determine which cell belongs to which assembly if individual cells fire during the same period of time at arbitrary rates. Although grandmother cell hierarchies have been proposed to resolve this controversy, they do so at the cost of making stark claims regarding the architectural nature of the brain. Indeed, selective response to assembly activity by superordinate cells would require a hierarchical organization of knowledge within cell assemblies which goes against the Hebbian cell assembly hypothesis and brings us back to a place coding view (Fujii et al., 1996). Related to the superposition catastrophe, the binding problem also introduces difficulties to distributed processing views (von der Malsburg, 1994). The binding problem can be briefly introduced by the following question: How is it that the shape and colour of objects are bound in a reliable way and not intermixed when more than one object is perceived? The visual receptive field can selectively respond in a distributed manner to input stimuli such that the colour green is detected and the colour red is detected simultaneously from parallel processes. It can also be the case that an apple is detected by a ‘shape-detecting’ assembly and a pear in another. But distributed coding cannot account for a mechanism that can integrate from the simultaneous presentation to the receptive field the bond between red and apple and between green and pear based on average firing rates alone. Because both red and green are active in the colour area and pickle and apple in the object area super-imposing cell assemblies are co-active and indistinguishable given that firing rates increase simultaneously. Finding how the brain solves this problem is important however because it is generally accepted that the binding of parallel informational fragments must take place to form an integrated whole for higher cognitive processing.

Underlying the problem of co-active populations lies the anatomical fact that neural assemblies are formed within static connections between neurones. However, selective binding of active patterns cannot be achieved by fixed connections because changes in stimuli elicit variation in the neural activity (Gerstner, Kreiter, Markram, & Herz, 1997). Instead, Gerstner et al. suggest that the physically connected pop-

ulation should be seen as the super-set of possible connections that may potentially need to interact. Thus according to this view some dynamical mechanisms of neural coordination must exist which not only allow the formation of co-active populations within the subset of connected parts, but also elicit selective responses of either a target population or change within the active population based on the current firing pattern. Amongst the first to take note of such a requirement within his Correlation Hypothesis framework was von der Malsburg who introduced the idea that subsets of neural populations within a connected network may contribute functionally distinct roles from other subsets through the correlated firing of neurones (von der Malsburg, 1981). More generally, the notion that a second order statistics in neural population codes may reveal alternate mechanisms of signal integration and cell recruitment has been gaining ground. The earliest most influential ideas stem perhaps from von der Malsburg as well as Bienenstock and the theoretical exploration in the capacity for this correlation theory to harness the complexity of mental representation for cognition (von der Malsburg & Bienenstock, 1986). These efforts constitute the early predecessors of the particular mode of coding which will be advocated in this thesis known as *spatiotemporal coding*. In particular, the focus of this thesis is directly motivated by the dynamical, informational, and functional features that spatiotemporal codes may represent for neural based cognition.

1.3 Spatiotemporal Coding

The most established and traditional form of spatiotemporal coding is based on Hans Berger's initial discovery that oscillatory firing is pervasive across different areas of the brain (Berger, 1929). Beyond broad area recordings still used today in EEG studies, precisely localized oscillations have also been recorded using implanted micro-electrodes in many areas of animal and human brains. Functionally, these oscillations of synchronously firing cells have been proposed to solve the superposition problem by offering a mechanism to distinguish groups of active cells within a population (Singer, 1994). If cells in assembly M fire synchronously as do cells in assembly N but with different rates from one another, then the identification of the corresponding stimuli within the same receptive field is possible. Beyond the classical problem of binding, synchronous activity within numerous areas of the brain have been identified experimentally and suggested to underly visual feature integration as well as sensorimotor coordination (C. Gray & Singer, 1987; Eckhorn et al., 1988; C. M. Gray, König, Engel, & Singer, 1989; Engel, König, & Singer, 1991; Murthy & Fetz, 1992). Furthermore,

debates regarding the origin of these oscillations were in part reconciled by realistic neural simulations. Pauluis et al., for instance, demonstrated numerically that synchrony could emerge from networks of cortical cells alone (Pauluis, Baker, & Olivier, 1999). At a similar time, Whittington et al. found that membrane decay constants on inhibitory cells in populations of pyramidal cells determined the frequency of emerging gamma oscillations (Whittington et al., 2000). Similarly, synchronous group activity is claimed to resolve the binding problem (Krüger & Becker, 1991; Gochin, Colombo, Dorfman, Gerstein, & Gross, 1994). Here, colours and shapes are thought to be bound by inter-areal synchrony of the assemblies in visual areas via their connected pathways. Despite the continued debate as to whether or not oscillations are simply epiphenomena, many contemporary researchers in neuroscience are discovering new properties of coherent oscillations and are suggesting new ways in which they could play an important functional role. Based on experimental evidence of synchrony, for instance, Burgess et al. and Hasselmo propose that hippocampal grid cells fire in a phase locked manner while sub-sets shift frequencies transiently based on head direction in rodents (Burgess, Barry, & O’Keefe, 2007; Hasselmo, 2008). High frequency oscillations such as gamma in the cortex have also been proposed to play a critical role in preparatory attention by increasing input sensitivity and suppressing distractors (Börger, Epstein, & Kopell, 2005).

Despite clear advantages that synchronous dynamics can provide in neural processing, its functional role as a precise coding mechanism remains unclear. More specifically, groups of synchronously firing cells are unlikely contenders for the task of complex transformation given some input and historical activity such that precise spatial (which cells) and temporal (when do they fire) patterns can be decoded, interpreted and responded to so as to accomplish specific cognitive tasks. Coincidence detection, however, presents itself as a potentially significant mechanism for accomplishing this while complementing, and perhaps at times underpinning, coherent oscillations in the nervous system. A neurone behaves as a coincidence detector if it fires only when a threshold number of quasi-simultaneous incoming signals reach its soma (Abeles, 1982; Braitenberg, 1988). Indeed, Fujii et al. maintain that assemblies of coincident detecting cells can form sufficiently complex circuits based on four propositions: (1) the network structure and spatiotemporal nature of the input spikes spontaneously and dynamically organize a cell assembly; (2) dynamic cell assemblies can be identified by their cross-correlation; (3) dynamic cell assemblies can serve as detectors of spatiotemporal patterns; (4) multiple cell assemblies can be instantiated within the same neuronal pool (Fujii et al., 1996). Also interested in solving the clas-

sical problems of superposition and feature binding, they suggest that proposition 3 and 4 actually address these without recall to synchrony. Indeed, proposition 4 states that assemblies within the same receptive field can be instantiated, while proposition 3 suggests that firing patterns from incoming pathways can be detected as preferred patterns through a particular active configuration of an assembly within the target population. This view thus offers an alternate perspective which states that features need not be bound by clearly identifiable first order properties such as firing rates, cell assembly hierarchies, or grandmother cells but instead that binding and identification may be implicit to the fine temporal structures of a dynamical cell assembly.

Fujii et al. argue that when looking for the effect of stimuli on firing rates only, an experimenter may overlook the fine temporal structures that arises from the correlated activity of cells (Fujii et al., 1996). In fact, the potential temporal and spatial patterns of excitation or inhibition of a neurone through synaptic or direct somatic connections is very large. With approximately 10^4 synaptic connections on average per cell, cortical neurones are subjected to a bombardment of signals often originating from more than one pre-synaptic cell. This suggests that the total synaptic input to a cell, even over a short period of time, will most likely have equal, if not greater, effect on the overall response of the post-synaptic cell than does the firing rate of individual inputs. Hence, the cooperative timing and location of these incoming signals have a significant impact on the activation or inhibition of the target cell. There are a number of interesting physiological simplifications that can arise from such a mechanism. For instance, a cell may be recruited (caused to fire) by many weak synaptic connections. Whereas rate coding requires that the incoming signal be transmitted through sufficiently strong synaptic connections, coincidence detectors require no minimal synaptic weights but can instead rely on the large amount of coincident stimuli. Furthermore, the frequency of incoming signals has no direct effect on the target cell's firing pattern. Although increased firing frequency may heighten the chance of coincidence of taking place, finely tuned spike times can arise from sparse activity within appropriate circuitry.

One of the most appealing aspect of coincidence detection comes from the fact that the organization of networks of coincidence detecting cells enable a rich set of dynamics while the physiological structure remains fixed. As Fujii et al. illustrate, a simple network, as shown in Figure 1.1, with two coincidence detection neurones A and B that are both excitable via channels p and q can exhibit varying dynamics depending on the temporal structure of the spike trains that are emitted from p and q (Fujii et al., 1996). For instance, it is possible that the superposition of spike trains

from p and q when arriving at cell A and B are coincident and make both cells fire. However, if the phase of the spike trains p and q that reach A is different from the phase that reaches B it is possible to have A fire but not B or vice versa. Change in phase would occur from varying onset of the incoming spike train from p or q. This, for instance, can be accomplished by a difference in propagation delay between channels and cells. Hence, depending on the pattern of the spike trains p and q it is possible to dynamically reconfigure the state of firing of both cells A and B. This however can be accomplished by virtue of the spike firing patterns alone without the requirement for physiological change within the network. If A and B were firing rate integrating neurones, however, it would be impossible to have A fire but not B without strengthening or weakening connections. In this scenario spike rates from p and q that are below threshold would trigger neither cell yet both would fire past a certain excitatory threshold. Hence, in contrast to firing rate neural networks, coincidence detection networks allow for a *dynamic* recruitment or inhibition of cells.

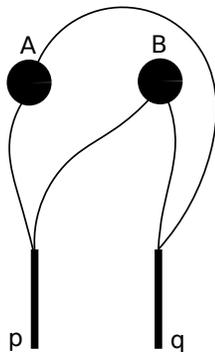


Figure 1.1: Simple two cell assembly of coincidence detectors.

By virtue of this dynamic recruitment it is possible to see how coincidence detection networks may allow neural assemblies to temporarily form by the recruitment of neurones into an active group while other connected cells remain ‘dormant’. Whereas dormant cells would exhibit sub-threshold activity, recruited cells would play a role by emitting spikes within its assembly and/or to other regions. Importantly, the functional specificity of these networks could be transient implying that neurones may not only be recruited by more than one assembly but also be actively engaged in multiple assemblies during the same time period. This suggests that activity emanating from two distinct active assemblies may bind their informational content as long as they belong to a connected superset of cells.

By offering a localized form of spatiotemporal coding capable of performing com-

plex processing on different input patterns, coincidence detection presents itself as a potentially essential form of coding in the central nervous system. Despite these properties, coincidence detection as well as other spatiotemporal coding forms remain poorly understood in terms of dynamics, functional capacity and information bearing properties. In addition, the biological basis for these coding mechanisms remain scarce. As suggested by Fujii et al., this is likely due to the difficulty of analyzing spiking data so as to extract second order statistics such as fine millisecond fluctuations in the spike patterns of cell populations. Despite these limitations, some experimental evidence brings principled support to the ambition of elucidating the functional capacity of coincidence detection for cognition.

1.4 Biological Basis of Coincidence Detection

The earliest most convincing evidence that coincidence detection is present in biological nervous systems comes work on the barn owl's auditory system by Carr and Konishi (Carr & Konishi, 1990). Knowing that barn owls can discern the location of prey from auditory cues alone, Carr and Konishi sought to explore the sensory mechanism responsible for such fine heading. In an elaborate experimental setup, they measured the propagation delay of signals from the cochlear nucleus magnocellularis (NM) from both hemispheres of the bird to the target nucleus laminaris (NL) on both ipsilateral and contralateral sides. While ipsilateral projections reach NL dorsally over similar distances of several hundred microns, contralateral afferents reach NL ventrally a varying distances across the mediolateral plane of NL. They found that while spikes remained phase locked to the stimulus frequency from NM to NL for each individual path, varying propagation times on ventral pathways led to NL cells firing selectively based on the delay between presentations of auditory stimuli to the left and right ear. This led the authors to infer that along with the varying delays on the ventral pathways, neurones in NL must play the role of coincidence detectors so as to enable interaural time differences (ITDs) to be detected by the owl. This, according to Carr and Konishi, validates Jeffress' original sound place theory (Jeffress, 1948). In particular, they found that NL cells derive ITD from interaural signal phase difference such that stimulus time differences as small as $30\mu s$ could be detected by NL coincidence detection cells. Such a small time difference cannot be explained via firing rate integration since neurones in either NM or NL would have to respond selectively between spike trains above 30kHz, yet the membrane time constants of cells in the auditory system of the barn owl are known to be approximately

0.5ms making these cells incapable of detecting any spiking frequency difference above 2kHz. Hence coincidence detection appears to be an essential form of coding in auditory processing. The question remains, however, whether this mechanism is involved in other processes within the central nervous system (CNS).

While the evidence from the barn owl auditory context is compelling, it simply supports the claim that coincidence detection does exist, but perhaps only for a small number of dedicated cells with fixed physiological function tuned for coincidence detection behaviour. Naturally, this may mitigate the perspective that coincidence detection could play a broader role in cognitive function. Interestingly, however, support for precise temporal sensitivity to incoming spikes was shown shortly after via realistic biological simulation of cortical neurones. Bernander et al., indeed showed in the early 1990's that a simulated layer V cortical pyramidal cell under realistic levels of stimulation exhibited a number of dynamic properties. These included variability of the so called membrane time constant from 80ms to 7ms, related variability of the input resistance of pyramidal cells from $110M\Omega$ to $14M\Omega$, and an increase of the electronic distance of synapses to the soma by an average factor of 3 (Bernander, Douglas, Martin, & Koch, 1991). In particular, a decrease in membrane time constant and in input resistance took place when background activity from 4000 excitatory and 1000 inhibitory synaptic inputs was present at frequencies of $1Hz$ and more. This increased background activity also led to the lengthening of the electronic distance for signal propagation within the cell. Their results thus suggest that pyramidal neurones can be dynamically modulated to function as integrators of spikes over time when background activity is low, while cells behave as coincidence detectors when background activity increases. Indeed, they further demonstrated that when background activity was high, the synchronous stimulation of 150 excitatory synapses led to spiking on nearly every occasion. However, by spreading the synaptic stimulation over time (asynchronous stimulation) then only 1 in 8 cycles of stimulation led to a spike. These findings brought considerable support to the view that layer V pyramidal neurones could behave as coincidence detectors under dynamic background conditions. In addition, by demonstrating that this is the case by varying the background activity of cells in associative cortical areas, their results lend great support to the hypothesis that coincidence detection may play a significant role in cortically dependent cognitive phenomena. Furthermore, their computer simulation approach made it possible to question what was difficult to measure in vivo and difficult to do in vitro; i.e., what behaviour do neurones have when contextual activity is taken into account?

Similarly, based on biological knowledge of cerebellum neurones and their anatom-

ical connections, Kistler and De Zeeuw demonstrated from accurate numerical simulation that granule cells of the cerebellar cortex can be regulated by inhibitory Golgi cells (also in the cerebellar cortex) only if the inhibitory inputs arrive synchronously on the granule targets (Kistler & De Zeeuw, 2003). More specifically, each granule cell is innervated by approximately four Golgi cells that fire regularly with inter-spike intervals of 20ms. Granule cells also receive excitatory input from mossy fibres originating from the deep cerebellar nuclei. They found however, that if inhibitory activation from the afferent Golgi cells reached the granule cells asynchronously, the hyperpolarization was spread too widely preventing any excitatory input from triggering granule cells to spike. However, if inhibition reached the granule targets synchronously, there was a short time window where the mossy fibre excitatory action potentials could trigger the granule target just after the hyperpolarization had passed. This suggests that precisely timed action potentials on behalf of the Golgi cells are necessary to regulate granule cell activity. Furthermore, while granule cells do not implicitly behave as coincidence detectors, the narrowing of the time window during which they can be stimulated indeed results in the dynamic tuning on behalf of Golgi cells of the granule cells into coincidence detectors. In addition, these cerebellar cells form an intricate network along with the inferior-olive (IO) which, with the synchronous firing of Golgi cells, will result in a 100ms reverberating loop due to afferent connections back from IO to the cerebellar cortex. Such loops, present in many areas of the brain, including thalamo-cortical networks, have been proposed by the authors and others as potential bedrocks for spatiotemporal pattern storage and processing (Gerstner & Kistler, 2002).

More recently, compelling biological evidence that coincidence detection is a significant mechanism of the central nervous system (CNS) function was put forth by in-vivo recordings by Prescott and colleagues (Prescott, Ratté, De Koninck, & Sejnowski, 2006). In particular, they showed from rat CA1 hippocampal pyramidal neurones that had an induced increase in membrane conductance (shunting) led to an average drop in input resistance from 158.5 to 69.7 $M\Omega$ accompanied by a 56% reduction in membrane time constant leading to a diminishing in cell excitability. They showed how this occurs due to the interaction between the shunting input and adaptation leading to an increase in the voltage threshold. The increased membrane conductance from shunting allowed adaptation to prohibit repetitive spiking which would normally occur under constant input. Hence, under this condition, brief stimuli will produce precisely timed spikes; but without repetition in virtue of adaptation. Hence, not unlike Bernander et al's. results, these findings show that contextual

change will alter pyramidal neurone behaviour from integration to coincidence detection. This provides further biological support for the claim that, under dynamical conditions, pyramidal neurones can behave as coincidence detectors such that they will not only be sensitive to precisely timed input spikes but will also produce precisely timed action potentials.

Overall, biological networks are likely to employ various neural codes simultaneously and perhaps even in a multiplexed fashion (Panzeri, Brunel, Logothetis, & Kayser, 2010). While biological observations as well as biologically accurate simulations can provide important insights as to the actual operation of particular neural tissue, we believe that a more profound theoretical appreciation with respect to coding via coincidence detection is needed. Indeed, the complexity of biological networks may be functionally decomposed more elegantly if we gain an understanding of the functional potential of fundamental, yet biologically realistic, coding mechanisms. In particular, the present thesis aims to demonstrate not only the dynamic properties of networks of interacting coincidence detectors but also that these networks can offer appealing coding and decoding capabilities for fundamental cognitive mechanisms.

1.5 Goal

The aim of this thesis is to motivate research in the area of theoretical neuroscience with emphasis on the dynamical and functional properties of propagation delay coincidence detection neural networks (DCD networks). These networks constitute a critical neural paradigm for complex brain function with potentially significant implications for rapid sensory processing and other critical cognitive processes.

1.6 Objectives

The primary objectives of this thesis are to demonstrate the viability of neural coding through biologically plausible delay coincidence detection neural networks for embodied reactive and basic adaptive cognition. More specifically, we aim to demonstrate from a bottom-up approach that sensorimotor coordination and memory can be achieved by such networks which we validate in part by adopting a neuro-robotics approach. Furthermore, we aspire to investigate and validate empirically the memory capacity and noise tolerance features of these networks via computer simulations. Finally, we aim to make formal contributions regarding the informational qualities

that delay coincidence detection networks possess with respect to spiking activity and spiking states from a spatiotemporal perspective. Overall, we believe that the successful undertaking of this research will provide important insights regarding aspects of brain function and establish precedence in the area of biological modelling of spiking neural networks for situated cognition via spatiotemporal coding.

1.7 Scope

The research undertaken for the completion of this thesis includes:

- Computer simulations of spiking neural networks employing simple binary threshold neurones with discrete transmission delays.
- Computer simulations of environments and robots based on the Khepera and ePuck robot simulation platforms/libraries (Michel, 1995; Magnenat, Weibel, & Bayeler, 2007), as well as the implementation of evolutionary algorithms aimed at searching for effective network structures, delays and other properties which will lead to the completion of given behavioural tasks.
- Behavioural task design based on existing embodied cognition problems such as the T-maze (Blodgett & McCutchan, 1947) and their computational implementation for agent simulation.
- Development of analytical tools and procedures aimed at identifying the functional characteristics of the implemented neural networks.
- Empirical investigations of spiking characteristics in delay coincidence detection neural networks with particular focus on the information capacity of these networks and their tolerance to noise.
- Formal investigations based on combinatorics and probability theory that approximate analytically the amount of activity in networks over time, spike state probabilities, and early directions for an estimation of the number of firing-chains available.
- Proposals for directions of potential future work.

This thesis will not include:

- Comparative studies with alternative coding paradigms in either biological or artificial neural networks or any other form of coding/control theory.
- Comparative analysis or implementation of more than a single discrete spiking neurone model.
- An exhaustive dynamical characterization of the different firing patterns possible within propagation delay neural networks under various parametric conditions.

- An exploration of the functional characteristics of different network topologies beyond what is required to accomplish the behavioural tasks undertaken.
- An exhaustive search for the various forms of reactive and memory mechanism which may or may not be possible in propagation delay neural networks including learning mechanisms relying on synaptic plasticity (e.g., Hebbian learning, STDP, etc.).
- The development of mathematical formalisms that require more detailed accounts of the biological properties of neural tissue or formalisms with exact numerical and parametric depictions, or a precise estimate of the number of firing-chains available in delay coincidence detection networks.

While some broad analogies to specific brain regions and neuronal properties and functions may be made, no explicit region of the mammalian brain will be explicitly modelled. In addition, while we explore the ability for propagation delay coincidence detection neural networks for enabling simple robots to solve reactive and memory tasks these are not claimed to be directly representative of any real biological agent sensory, neural, and motor coupling. Our investigations are provided as a demonstration of the feasibility of the use of a delay-based neural network in a robot controller. In general, the work proposed here will constitute a theoretical endeavour with respect to what biological neural networks *may* be implementing as solutions to the proposed cognitive tasks rather than a confirmation that these processes actually do take place. This work will thus inform of the findings as proofs of concept and not as proofs of existence.

1.8 Problem Statement

Information and function is traditionally thought to be implemented in the brain via neural communication by spike firing rates. From sensory stimulus integration, to reasoning, to language, and motor control, the individual tuning of neurones to emit action potential at particular frequencies is regarded as the most obvious form of signal discrimination and transformation in the brain. While there is a large body of evidence supporting this form of coding with respect to specific mechanisms such as motor control, there is growing experimental evidence and informed theoretical opinions (see Chapter 2) supporting the claim that spatiotemporal coding is a critical form of information processing in the brain. In particular, coincidence detection could occupy a central role in this coding alternative via its ability to implement different forms of neural dynamics and high information capacity. Our problem statement is: To determine the characteristics and the ability of coincidence detection networks with

propagation delays to solve elementary cognitive tasks from an embodied standpoint and to evaluate the capacity, noise tolerance features, and spiking characteristics that these networks possess.

1.9 Contributions

This thesis makes the following contributions:

- Formalization of binary DCD networks (Section 3.2.2).
- Identification of classes of dynamical behaviour for DCD networks characterized by delay distribution (Section 3.2.4 and Section 6.1) and mapping onto reactive and memory-based cognitive functions (Section 3.3.1).
- An architecture and method for the creation of DCD network controllers for reactive tasks (Section 3.2.5).
- An architecture and method for the creation of DCD network controllers for memory-based tasks (Section 4.3).
- A method to decode spatiotemporal spike patterns using coincidence detection neurones (Section 4.2).
- A method to identify the amount of ambiguity between spatiotemporal patterns and the minimum coincidence threshold required for pattern decoding cells (Section 4.4).
- A characterization of the tolerance to noise of DCD networks (Section 5.1) and their memory capacity (Section 5.2).
- Development of mathematical approximations for the spiking states of DCD networks (Section 6.2).

1.10 Organization of the Document

Chapter 2 presents related work in the area of spatiotemporal coding via coincidence detection and highlight important theoretical considerations based on experimental observations as well as neuronal parameters of influence for exploring the dynamics and function of delay coincidence detection neural networks. The essential neural and network properties of spiking neurones are explored individually in order to form a more comprehensive view of the physical properties that will likely impact neural coding function. Chapter 3 begins with theoretical considerations to approach the decomposition of fundamental cognition into the areas of the reactive and the adaptive. We then present empirical work which explores the dynamical and reactive control

features of delay coincidence neural networks in a light-seeking task. This is accomplished using evolutionary computation so as to impose a minimal set of assumptions and better explore the space of effective connectivity and delay structure. After introducing early theoretical insights which guided the itinerary of this work, we show connectivity conditions that can satisfy these insights. Chapter 4 is then aimed at developing a better understanding of the ways in which delay coincidence detection networks can serve as controllers in agent tasks. In this chapter we make use of the complex firing properties of these networks for robot dynamic memory tasks by developing a multi-network controller including a memory network, sensorimotor network, and a coordinating network which allows memories to guide behaviour. By scaling this task we show that memory networks can dynamically store a large number of input signals for sensorimotor control. Subsequently, we explore via empirical computer simulations the noise tolerance and memory capacity of these networks in Chapter 5. Finally, Chapter 6 makes use of combinatorial and probabilistic methods to identify analytically the activity and spiking properties of these networks. In particular, after introducing an approximation method for the level of activity across time based on structural parameters and past activity, we develop a set of probabilistic methods aimed at capturing the number of states in which a network can be in over time, as well as preliminary efforts aimed at identifying steps towards an analytical and structurally based estimate of the number of firing-chains available in delay coincidence neural networks. This thesis ends with a conclusion and direction for future work in Chapter 7.

Chapter 2

Related Work

Coincidence detection (CD) in spiking neural networks is the theoretical principle by which neuronal cells fire when sufficient temporally synchronized input reaches their dendritic and/or somatic compartments. This notion was originally proposed by Abeles in the early 1980's (Abeles, 1982). Supported by biological evidence, CD is today attracting increased attention as a possibly significant form of neural coding in the nervous system. Traditionally contrasted with the well established rate-coding paradigm, CD is being proposed as an efficient neural mechanism in sensory integration, memory function, but is also hypothesized as playing an important role in higher cognitive processes such as learning, thinking and reasoning. Notably, it was shown that the synchronous input of a few afferent cells will match the depolarization induced by a large number of asynchronous inputs (Bernander et al., 1991; Gerstner & Kistler, 2002). More generally, CD fits within the spatiotemporal coding paradigm of neural coding processes because it relies specifically on the precise timing of neuronal events across a spatially distributed population of neurones. Furthermore, CD has also been proposed as a contributing factor to macro-scale spatiotemporal dynamics such as oscillations and synchrony in neural populations. While CD, as we will show in the following, offers a promising area for scientific inquiry on the neuronal basis of cognition it remains important, however, that we not attribute a-priori relevance to the notion of CD without careful consideration for the aspects that make it conceptually and functionally distinct from the competing rate-coding paradigm. Notably, with support from existing work exposed in Chapter 2, the present thesis will introduce in Chapter 3 a set of experiments which should bring embodied cognition support to the claim that CD is indeed a crucial neural coding contender that is not only biologically realistic but also functionally significant.

In the following, we begin by describing existing theories based on spatiotemporal

coding via coincidence detection such as synfire chains, dynamical cell assemblies, rate-order coding, and polychronization, and the significance of their results from a neural coding perspective via precise spike timing. We then examine existing work showing the influence of various biological characteristics of neurones and their networks on the dynamics and coding efficiency of temporal coding with particular emphasis on coincidence detection. These include: the membrane time constant, the firing threshold, the refractory period, connection topologies, transmission delays as well as external factors such as input, noise and feedback. We subsequently introduce existing work emphasizing the importance of transmission delays on spatiotemporal coding and show that an important theoretical and experimental gap remains with respect to the role that precisely tuned axonal and synaptic delays may play on neural information processing.

2.1 Coding via Precise Spike Timing

Theoretical and biological investigations since the early 80's have suggested that the precise timing of incident spikes onto target neurones may play a significant role in neural coding in the brain. While this view still remains controversial to date, a number of theoretical models have explored the possible population level mechanisms that could result in effective information representation, transmission, and transformation with successful results. As an added impetus, these theoretical models have often found partial support from in-vivo experimental evidence. In the following we present these alternate model-based theories in a loose chronological order.

2.1.1 Synfire Chains and Braids

Using a biologically realistic model of cortical neurones, Abeles hypothesized in the early 80's that groups of neurones can fire synchronously from one pool of cells to another in order to transmit sensory information in a reliable fashion (Abeles, 1982, 1991). These so called synfire chains suggest that sensory stimuli trigger synchronous firing in a subset of cells within a pool of receptive neurones. This synchronized population could then transmit stimulus related activity to a subsequent pool which would in turn fire synchronously and so on as depicted in Figure 2.1. Despite temporal jitter due to noise or fluctuation of post synaptic membranes potentiation, Abeles showed how synfire chains remain remarkably stable. Overall, this approach depicts a network of feed-forward activity, whereby cells in one pool transmit excitatory information to

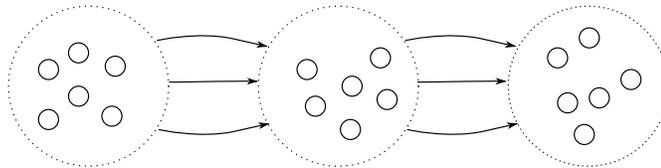


Figure 2.1: Synfire chain: Synchronized firing in a population of cells projects to a second population in which cells fire synchronously which in turn trigger synchronous firing in a third population.

cells in another pool. Abeles, however, further suggested that pools need not be necessarily physically distinct, but that cells which participate in the firing of one pool could also fire in another pool, so long as synchrony amongst pools is maintained. Importantly, cells in a receiving pool initiate synchronous firing at a subsequent time step due to the convergent connections from many cells from the afferent pool and their transmission of coinciding spikes – hence the principle of coincidence detection. This work prevails today as possibly one of the most important insights in coding theory in the neural sciences as it set the stage for the radically alternate view that the activity of neurones may be much more precisely tuned spatially and temporally than had been previously thought. Furthermore, corroborating experiments at both theoretical (Nelken, 1988), computational (Diesmann, Gewaltig, & Aertsen, 1999; Trengrove, 2005), and biological levels (Douglas, Martin, & Whitteridge, 1991; Prut et al., 1998; Ikegaya et al., 2004) have supported Abeles’ synfire chain theory, demonstrating that neural coding can occur as an alternate method to the well established rate coding and statistical coherence theories such that the precise timing of spikes will deterministically impact the activity of post synaptic cells.

Synfire chain theory, however, makes a number of assumptions which may not always hold in cortical networks or networks in other regions of the brain. In particular, information transmission is believed to be present only in the case of synchronous firing and transmission delays are simply assumed to be of equal duration (typically a single delay step). Indeed, in a system of interconnected elements, the prevalence of equal transmission delays has a limiting effect on the dynamical diversity that such networks can exhibit which we expose in Chapter 3. Furthermore, the assumption of equal transmission delays is generally unlikely within a biological context of heterogeneous transmission delays and dendritic/somatic integration times. Indeed, local and especially long distance cortico-cortico connections will inevitably result in variations in action potential transmission times. While Abeles initially suggested that synfire

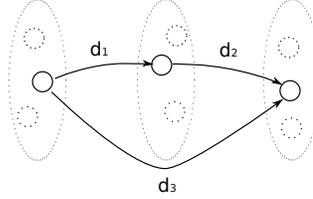


Figure 2.2: Synfire braid: Transmission times between cells can vary such that the conduction times for one path between any two cells is equal to all other paths between those two cells. Here path $d_1 + d_2 = d_3$.

chains would apply more specifically to small localized circuits, there is neither an a priori reason to dismiss the idea of precise spike timing and coincidence detection between long distance areas, nor are there any a priori reasons to believe that fine tuned cell properties in certain regions of the brain can not exploit millisecond scale delay differences in small localized networks. Despite the important coding characteristics that synfire chains offer, Bienenstock modified synfire chains to incorporate variable conduction times, in what he termed ‘synfire braids’ (Bienenstock & Doursat, 1990). As a result, a subset of spikes propagating from one pool effectively ‘skips’ the immediate following pool and project instead to pools of synchronized firing cells that are τ number of steps away (Bienenstock & Doursat, 1990; Bienenstock, 1995). According to this alternative theory, a population only behaves as a synfire braid if the sum of conduction times between any two neurones for one path is equal to another path between these two neurones – what he defined as the consistency rule. This is illustrated in Figure 2.2. Hence synfire braids can incorporate long distance connections wherein conductance heterogeneities are present so as to make synfire chain like dynamics possible. Insightfully, Bienenstock proposed that hebbian learning could be used as a means to reinforce those connections which promote the synchronized firing within synfire braids (Bienenstock, 1995). Furthermore, binding mechanisms have been proposed which have been argued to enable the compositionality of spiking structures via synfire activity into dynamic representations to subserve potential higher order cognition including concept formation and language (Bienenstock, 1996; Bienenstock, Geman, & Potter, 1997; Abeles, Hayon, & Lehmann, 2004).

While synfire braids offer a biologically appealing solution to the lack of temporal heterogeneity in the original synfire chain models, their principle mode of computation remains based on the notion of synchronous firing within neuronal pools. Their functional purpose thus remains focussed, as do synfire chains, to that of stable signal transmission, dynamical entrainment, and binding. While synfire theories are

in substantial consonance with the spatiotemporal activity under investigation in this thesis, we aim to direct our focus on the time locked spiking properties of delay coincidence detection as a general form of coding to subserve reactive and adaptive cognition.

2.1.2 Discrepancies in the Rate Coding Paradigm

The notion Abeles introduced, that the coordination of precise spike timing should play a fundamental role in neuronal activity, has led to considerable reconsiderations by a number of researchers regarding the effectiveness of the rate coding view as a model of brain function. In addition to these new time coding ideas, discrepancies with respect to the accuracy that rate code analysis offered surfaced by the early 90's. In particular, Softky & Koch noted that well established biophysical model parameters of single cells (still widely used today) assume low variability in the inter-spike-intervals (ISI) of model neurones, while physiological data show the contrary in cortical V1 and MT cells (Softky & Koch, 1992). By implementing their own leaky-integrate-and-fire (LIF) model, they demonstrated that high ISI variability would only take place if either the membrane time constant τ_m of a cell was much lower than the mean ISI of their simulated neurone model, or if a high degree of inhibition was applied to the cell in conjunction with stimulation. As they argued that such high degrees of inhibition have not been found in cortex¹, it must be the case that realistic membrane time constants should typically be lower than the assumed values (usually $\tau_m \approx 10ms$). In fact, the high ISI-to- τ_m ratio is considered to be the defining criterion for coincidence detection to occur (König, Engel, & Singer, 1996). By further validating their model using a precise Hodgkin-Huxley compartmental model with random Poisson excitatory input, they suggest that cortical cells will more realistically exhibit very short membrane time constants. Consequently, fast membrane time constants cannot support slow integration times needed for rate coding to occur, but instead precise millisecond scale temporal coding via coincidence detection must take place. A few proposals have been made to account for this discrepancy between the typically longer membrane time constant derived from in-vivo cell somatic capacitance (C) and resistance (R) properties (where $\tau_m = RC$) and the short time constant inferred by these coincidence detection supporting experiments. While it is possible that intrinsic cell properties will be sufficient for coincidence detection to occur, it has also been

¹Another possible contributor to high variability is the presence of balanced inhibition and excitation which make a cell firing dependent on membrane potential fluctuations (Shadlen & Newsome, 1998; Gerstner & Kistler, 2002).

suggested that certain network dynamics will also facilitate short integration times. For instance, Softky demonstrated via the simulation of a precise biophysical neuron compartment model that balanced inhibitory conductance could dynamically reduce the membrane time constant of individual cells resulting in a shorter effective time constant τ_{eff} (0.67ms in his model) than passive time constant t_m (20ms in his model) (Softky, 1995). Bernander et al. similarly showed that synaptic background activity could dynamically reduce the effective membrane time constant and other cellular features resulting in a much stronger firing responses from synchronous synaptic activation as opposed to asynchronous input when taking contextual activity into account (Bernander et al., 1991). A more careful investigation of the various cellular properties and network parameters favouring coincidence detection is thus warranted.

2.1.3 Dynamical Cell Assemblies Hypothesis

Inspired by Abeles' notion of coincidence detection and the subsequent characterization of the neuronal and network properties leading to the sensitization of neurones to temporally narrow events, an important set of theoretical insights ensued. Based on the early work on pulse propagation neural networks by Judd & Aihara who demonstrated the informational and dynamic complexity of a single dimensional model with direct sensitivity to the time of last spike, the *dynamical cell assembly* hypothesis was proposed (Judd & Aihara, 1993; Fujii et al., 1996). This hypothesis is based on the notion that coincidence detection is the central organizing principle of dynamics in the cortex and that it enables the formation of functionally cooperating cell assemblies on a level above physical connectivity. Marking a departure from the synfire chain paradigm, the dynamical cell assembly hypothesis does not make the restrictive claim that neurones must fire synchronously within an assembly but instead exhibit correlated firing at a fine (millisecond) temporal scale. Part of the justification for this move by Fujii et al. is that early experimental data on neural coding can be fitted by various traditional coding theories because the analytical methods used by theorists under-constrain their interpretations. In particular, Vaadia et al. noted that existing theories lacked the capacity to adequately address noticeable variations in spiking regimes over short periods of time despite stable input and fixed synaptic connectivity (Vaadia et al., 1995). Furthermore, near random firing observed from ISI histograms (leading to the assumption of a poisson process) has led most to believe that no information is carried in individual spike trains. However, Fujii et al. argued that this sort of analysis can only reflect first order statistics of a spike train which cannot

capture the temporal accuracy and correlation between spikes. Instead, they propose that the sensitivity of cells to temporally coinciding input spikes leads to *functional connectivity* whereby cells that fire at the right time become transient members of an active group of cells (an assembly). Because spikes that do not coincide with others loose causal relevance only timely spiking cells belong to the formation of a dynamic assembly. This arises for a given spatiotemporal input spike pattern, the current spiking state of the network, and a given set of network properties such as connectivity, temporal and spatial characteristics. Interestingly, individual neurones participating in one cell assembly during a certain period of time can also participate in another cell assembly during the same period of time. To detect the fine temporally correlated activity between cells, Fujii et al. propose both cross-correlograms as well as joint peri stimulus time histograms (J-PSTH) methods (Aertsen, Gerstein, Habib, & Palm, 1989). While cross-correlograms provide a quantitative mean to detect the correlated activity between two spike trains, its significance is dependent on the choice of an adequate time window and assumes that the measured spike process is overall stationary. Instead, J-PSTH provides a means to observe variation in correlated spiking between any two neurones over time. Noise tolerance in these networks is also considered by Fujii et al. to be an important feature since relatively high degrees of noise will not interfere with an assembly's activity due to a lack of adequate timing. Instead, noise may actually help dynamic assemblies correct erroneous 'misfirings' which may occur to interfering assemblies. With respect to cognition, this cell assembly hypothesis is thought to have many functional repercussions leading to the capacity for both *representation* and *processing*. These include the binding of inputs into new dynamic assemblies, relaxing the superposition bottleneck, and satisfying important criteria for the realization of hierarchical knowledge representation (Fujii et al., 1996).

The dynamical cell assembly hypothesis pushes the coding principle via coincidence detection a step further by proposing that global assembly synchrony, as opposed to synfire chains, is not necessary for spatiotemporal coding to take place. In addition, dynamic cell assemblies are demonstrated, as a proof of concept, to possess the necessary features for handling both representation and processing via spiking activity. Fujii et al., however, make a limited case for the applicability of such networks in realistic behavioural/cognitive task applications and provide limited experimental support for their proposal. While it may be the case that significant spike variation takes place in experimental data, it remains difficult to assert that these variations are due to functional reorganization of dynamic cell assemblies without proper experimental recordings or realistic simulation support. Furthermore, no adequate formalism

with respect to the informational content of these dynamical cell assemblies is offered. However, the notion of functional connectivity which exists above structural connectivity and beyond synchronous firing makes the dynamical cell assembly hypothesis a significant framework upon which we aim to build.

2.1.4 Practical Aspects of Dynamical Cell Assemblies

To support the idea that neurones form dynamic cell assemblies based on coincidence detection and network properties such as propagation delays, a number of simulation based experiments were undertaken by collaborators. Watanabe & Aihara first began by exploring the chaotic dynamics of small networks of simple deterministic neurones with non-linear spike rise time (Watanabe & Aihara, 1997). More specifically, neurones add up input pulses leading to activation $a(t)$ which decays exponentially with time constant τ . When this activation becomes larger than the firing threshold T plus negative feedback $r(t)$ then the neurone fires (sends an output spike) with time delay $g(h(t))$ where $h(t)$ is the amount of supra threshold activity defined by:

$$h(t) = a(t) - (T + r(t)) \quad (2.1)$$

While $g(x)$ is a first order inverse function such that a small supra threshold activation leads to a longer delay before pulse propagation. A small network of 64 neurones fully connected with equal weights and random delays with normal distribution (M:10, σ : 0.2) was then implemented. Watanabe & Aihara showed that by stimulating the network with a small regular input (simultaneous stimulation of 3 neurones with period 20 time steps) two distinct regimes emerged in the network depending on the firing threshold. With higher threshold ($T = 1.4$) a subset of neurones in the network spiked reaching a limit cycle with a period of 120 time steps. With lower threshold ($T = 1.2$), however, no limit cycle was reached but a chaotic state of activation emerged instead. Hence the length of so called "chain firing" shifted from a finite length to infinite duration. To characterize this firing-chain length formally, they defined the recursive notion of *pulse generation*: for a given set of effective incident pulses to a neurone pick the pulse with the largest pulse generation and add 1. Suppose a neurone i fires when it receives its N^{th} pulse P_N at time t_i^N . The effective pulses to a neurone are pulses P_{N-k} to P_N for the smallest k which satisfy the following inequality:

$$\sum_{j=0}^k \exp\left(-\frac{t_i^N - t_i^{N-j}}{\tau}\right) \geq T + r(t_i^N) \quad (2.2)$$

Intuitively, the effective pulses to a neurone i are those which arrive within a sufficiently small time-window $t_i^N - t_i^{N-k}$ to trigger a spike on this neurone. Hence when the sum of the exponential expression approaches zero, the neurone behaves more as a coincidence detector, while the further away from zero (usually bounded at 1) the neurone behaves as an integrator. In addition, Watanabe & Aihara investigated the sensitivity of these networks to perturbation. They showed that small perturbation during normal network activity will dissipate quickly when the firing threshold is higher since all signs of perturbation are lost by the end of the pulse generation (after 120 time steps). However, for a low threshold the infinite pulse generation allowed a small perturbation to propagate through the network throughout the duration of the simulation with amplifying consequences. This confirmed that small changes in the state of a network with a chaotic attractor leads to important divergence in activity state. Despite this sensitivity the authors suggest that chaotic dynamics may be an important feature for cognitive processes such as dynamic associative memory, novelty filtering, and task switching. In particular, the ability for small local changes to propagate to a network’s global behaviour via chaotic attractors could work interchangeably with a state of phase locked oscillations to suppress any ‘distracting’ perturbations within dynamical cell assemblies.

In later work, Watanabe et al. demonstrated that similar pulse propagation networks can lead to either ‘conventional synaptic connectivity’ for longer membrane decay time constants τ while ‘functional connectivity’ is possible with shorter membrane decay time constants (Watanabe et al., 1998). Similar to the previous experiments, they first demonstrated that a cyclical input with fixed frequency on a network of cells with large $\tau = 0.3$ led to strong firing and highly synchronous spike trains with mean frequency matching the mean synaptic delay. However, with small $\tau = 0.05$ synchronous firing was replaced by sparse irregular spiking. Under their interpretation, this later condition is the result of coincidence detection by neurones and should lead to functional connectivity: the process by which variation in the spatiotemporal input will modify the effective connectivity between cells leading to multifunctional processing. As illustrated in Figure 2.3, they designed 4 *spatial* input patterns A, B, C, and D where A & B share common cells, A & C share different cells, B & D share

a separate set of cells, and C & D share another set of cells. They then presented each individual spatial patterns to a network of coincidence detectors ($\tau = 0.05$) according to two different *temporal* input patterns, which they call programs. The first program presented a spatial pattern at regular time intervals of 2.0 simulated seconds, while the second program presented spatial patterns in 1.0 second separated doublets presented every 3.0 simulated seconds as shown in Figure 2.3. Results showed that when presented under the conditions of Program 1, the network activity responded with equal repeating spatiotemporal patterns for spatial inputs A and B, and equal spatiotemporal patterns for inputs C and D. When presented under temporal conditions of Program 2, however, the network activity displayed equal repeating spatiotemporal patterns for inputs A and C, and equal patterns for inputs B and D. Hence the temporal consequence of presenting separate patterns sharing common spatial substructures with different temporal signatures enabled the network to categorize each input pattern. This captures the essence of functional connectivity in that no physical change had to be applied to the network connections or to individual neurones so as to achieve categorization. Interestingly, this experiment may have demonstrated a working principle by which background oscillatory phases could interplay with coincidence detecting networks, since background oscillations could entrain spiking at a particular temporal scale without affecting the spatial characteristics of the activity thus playing the role of ‘program’. Overall, the capacity for coincidence detection to perform non-trivial computations under appropriate network conditions without the need for structural alteration makes it a potentially critical form of coding in the central nervous system. Given these appealing coding properties of coincidence detection within simulated networks incorporating propagation delays, a new perspective on neural information coding which distances itself from the traditional rate coding paradigm becomes apparent.

2.1.5 Alternative Codes

Before exploring more recent developments on functional connectivity and dynamic spatiotemporal coding it is important to note that alternative neural coding theories based on the precision of spike times have been proposed and tested numerically. In particular Thorpe and colleagues proposed an elegant argument claiming that neurones participating in visual perception cannot be (solely) integrating multiple spikes over time but more likely compute visual scenes with a single spike per neurone (S. J. Thorpe & Imbert, 1989; S. Thorpe, Delorme, & Van Rullen, 2001). Their argu-

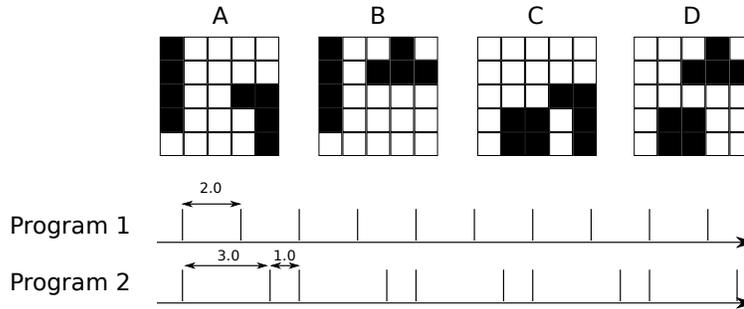


Figure 2.3: Approximate reproduction of the spatial pattern A, B, C, and D as well as temporal patterns Program 1 and Program 2 used by Watanabe et al., 1998.

ment is based on experimental data which shows that humans perform natural scene categorization within approximately 150ms (Antal, Keri, Kovacs, Janka, & Benedek, 2000). Furthermore, biophysical data regarding visual pathways from the retina to the cortical areas responsible for scene categorization established that approximately 10 layers of neurones are traversed during visual scene spike propagation, yet average neuronal firing rates are estimated to occur at about 10Hz. This implies that within 150ms neurones in each layer can spike on average only once during the task. Overall, this provides a strong argument for sparse activity coding via a particular spatiotemporal code in visual processing over a firing rate dependent mechanism. Coincidence detection, however, is not the favoured coding mechanism of this group. Instead, they propose that neurones are likely to encode retinal information via *rank-order coding*, that is, that the temporal order with which spikes arrive on a target is the carrier of information. The first component of their argument is based on the idea that throughout the visual system, cells fire with differing latencies based on various factors such as local luminescence and contrast for ganglion cells, orientation for cortical V1 cells, etc. (Gauthrais & Thorpe, 1998). They explain this latency by the dynamic properties of leaky integrate-and-fire neurones which show latent responses depending on strength of activation: stronger input activation leads to shorter spike response times. Hence rather than depending on synchronous inputs, which they claim other spatiotemporal codes depend on, rank-order coding is merely sensitive to the latency order at which spikes are propagated to a target population. The second aspect of their argument has to do with the informational properties of this coding method in terms of the number of firing states it can produce. First, they show that most of the information can be stored in the first few spikes from the afferent population therefore providing even greater coding efficiency with as little as 1% to 2% of cells

coding a visual scene (Van Rullen & Thorpe, 2001). This method, however, depends on a stronger weighting of the connections carrying those initial spikes. Second, rank order coding carries more information in a population of cells than does simple spike counting (firing rate) or parallel coding. A firing rate code, implies that no distinction is made between originating afferents; hence in a population of N cells, $N + 1$ states of firing can be detected by the target population (with the added non-firing state). Parallel coding, however, implies that all combinations of afferent cell states can be decoded leading to 2^N possible input states. Rank order coding however, is sensitive to all arrangements of the input spikes leading to $N!$ possible input states. Finally, because rank order coding is only sensitive to the order at which spikes arrive and not the strict value of this latency, this temporal coding method is robust to spike time jitter. That is, variation in spike times due to noise or interference will not affect the decoding within a range of temporal tolerance equal to the time difference between afferent spikes.

Taken together, rank coding offers a powerful set of features which Thorpe and colleagues were able to demonstrate in a number of visual processing tasks. In addition, they support rank order coding as being a more likely candidate for spatial temporal coding than coincidence detection because the later suffers in two different ways. First, they argue that coincidence detection used as a time code is too sensitive to the precise timing of spikes making it highly sensitive to noise and spike jitter. Second, the decoding mechanism that must be implemented by neurones may be “prohibitively complicated” in view of the fact that coincidence detection requires a mechanism to detect the precise latency of each spike. Despite this, they recognized that coincidence detection can potentially lead to much greater information carrying capacity depending on the temporal resolution at which coinciding spikes arrive. Their criticisms, however, may not be vindicated in light of more careful theoretical considerations as well as mounting experimental evidence from numerical simulation. A refutation of coincidence detection (CD) based on sensitivity to noise is only valid if one adopts an idealistic view of the input. Indeed, if each spike within a population must be accounted for then a slight loss in synchrony due to some perturbation would render CD coding impotent. However, it is not unreasonable to expect redundancy of spikes over many neurones enabling CD to preserve its coding efficiency over time. For instance, if the number of afferent inputs from a given layer reach a CD neurone in greater number than what is necessary to elicit a spike, then this cell can afford a degree of jitter proportional to this difference between the minimum number of coinciding inputs to that received (see Figure 2.4). Hence, a CD network with a greater

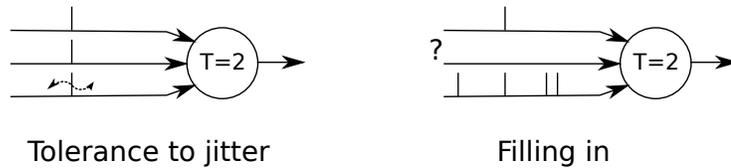


Figure 2.4: Left: A CD neurone with threshold $T = 2$ will still fire if jitter occurs on redundant afferents (here the 3^{rd} afferent). Right: A CD neurone with threshold $T = 2$ can fire despite loss of an input (2^{nd} afferent) if noisy input (3^{rd} afferent) can fill-in for the missing spike.

degree of afferents per cell or a lower coincidence detection threshold should lead to sufficient jitter and noise tolerance. Furthermore, as mentioned above, Fujii et al. suggested that the presence of background noise may actually help correct misfiring within CD networks by ‘filling-in’ for absent spikes as shown in Figure 2.4 (Fujii et al., 1996). These properties are revisited in Section 5.1. The second critique makes the claim that neural coding via CD would suffer from an overly complex spatiotemporal pattern making it too difficult to decode for some arbitrary system. This view, however, fails to acknowledge the exploitative nature of complex adaptive systems such as the brain. If the production of precise spike time activity can be instantiated in some neural tissue given natural biological properties, then it is likely that networks possessing equal complexity should be capable of decoding such patterns in order to translate them into exploitable forms producing an adaptive response. It is true, however, that coincidence detection remains a complex code for which theoretical and experimental development remains scarce. We later address this decoding problem from both an evolutionary approach in Section 3.2.5 and algorithmically in Section 4.2. While, coincidence detection is increasingly recognized as a basic principle of neural processing at the level of individual cells and within particular areas of the the CNS, the coding potential at a network level elicits continued debate. We now discuss more recent developments with respect to coincidence detection and its relationship to other fundamental neural factors.

2.1.6 Time Coding and Polychronization

While we typically consider the input to a neural network to be exclusive and transient, there are reasons to believe that far reaching and more persistent background activity could also modulate precise spike times. This phenomenon could take place in the form of phase-locking if the background activity is oscillatory in nature as

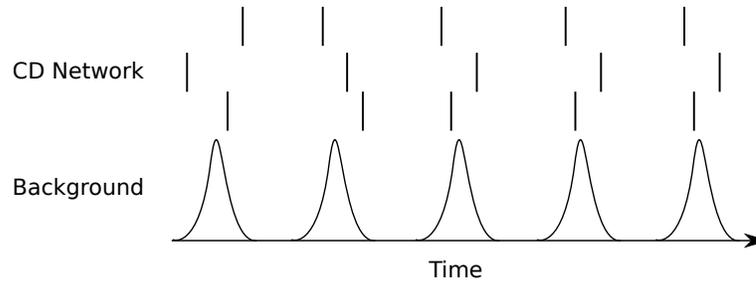


Figure 2.5: Phase-locking: Early arbitrary spikes become entrained by the background oscillatory activity.

shown in Figure 2.5. The olivo-cerebellar system, for instance, is a known circuit linking the inferior olive in the brain stem to the cerebellar cortex in a bi-directional fashion. This connectivity leads to a reverberating loop that induces inferior olive cells, which tend to fire irregularly, to synchronize due to a 100ms feedback from their projection to the cerebellar cortex (Sotelo, Llinas, & Baker, 1974; Gerstner & Kistler, 2002). Interestingly, while the population oscillation was measured at about 10Hz, individual spike firing was sparse and irregular. To account for this, Gerstner & Kistler proposed that if individual cells fire with an overall lower frequency than the population but that their intermittent spikes fell in phase with other cell spikes, then their cumulative activity over time would match the population frequency. For instance, in a population of 4 cells, if cell 1 fired with frequency 5Hz, cell 2 at 2.5Hz, cell 3 at 1.75Hz and cell 4 at 5Hz, and these phases are arbitrarily shifted by π then the overall population frequency would result in 10Hz activity. Crucially, this phase locked activity is driven over a period of activity due to the sub-threshold input reaching cells from the reverberatory feedback. If sufficient synchronized input reaches a cell then it will shift from an out-of-phase behaviour to an in-phase behaviour.

In a different context, Izhikevich proposed that phase coding could take place in a network of cortical neurones with varying delays as long as these delays are selected for by some mechanism (E. M. Izhikevich, 2006b). He demonstrated numerically that balanced inhibition and excitation in a network with broad propagation delays (between 1ms and 20ms) could lead to the emergence of ‘polychronous groups’; that is, groups of cells which fire deterministically in a time-locked fashion in response to a particular spatiotemporal input. These cells have a tendency to form a larger number of groups if axonal delays are selected for via spike time dependent plasticity (STDP). Interestingly, by adopting biologically plausible cortical connectivity with a mixture of both inhibitory and excitatory neurones, this process led to the appearance

of transient synchronous states such as delta (2Hz-4Hz) and gamma (30Hz-100Hz) frequencies which are pervasive in the brain (E. M. Izhikevich, 2006b). Indeed, certain network properties such as interneurone membrane time constants were found to be responsible for the emergence of gamma oscillations in pyramidal populations of the cortex – the so called PING effect (Whittington et al., 2000). These network effects induce STDP to select connection delays between cells that satisfy these regimes, leading to polychronous groups with connection delays matching network periods, i.e., 10ms to 20ms delays for gamma. After reaching a phase-locked state, polychronous groups are hypothesized to account for the lack of correlation between neurones despite the presence of background gamma. Similarly to Gerstner & Kistler’s account, neurones are not likely firing at the same rate but only at multiples of the underlying frequency, in this case recordings from two neurones which do not belong to the same group will likely exhibit little correlation. Furthermore, smeared stimulus onset-to-response time histograms, when averaging over a number of trials, can be explained by phase based coding. As Izhikevich points out, this should not be interpreted as evidence against spike-time coding but can be explained by a mechanism where the presentation of a stimulus onset reaches a target cell population at different phases of the background oscillations between each trial. Izhikevich’s work along with other colleagues contributes significant weight to the notion of spike-time coding in the brain. Perhaps, one of the most attractive features of polychronization is the notion that the number of neuronal groups is much larger than the total number of neurones or edges in the network (up to 8000 groups were found in a network of 1000 cells with only 10% connectivity), thereby serving as a possible platform for active memories, stimulus binding, and more. This more recent work, however, makes little to no reference of the impact of various parameter settings, such as propagation delay distributions, the effect of varying decay constants, refractory period effects, etc. Furthermore, while Izhikevich recognizes the high similitude of polychronization within synfire theories by Abeles, Bienenstock and colleagues, he makes little allusion to extensive work by the proponents of the dynamic cell assembly hypothesis. Indeed, the formation of polychronous groups, other than being facilitated by STDP, resembles closely the notion of functional connectivity forming dynamical cell assemblies (Fujii et al., 1996). While Izhikevich adopts a biologically more realistic modelling paradigm, a limiting factor in his approach is that it is committed to role played by STDP as a driving force behind the selection of neuronal groups (functional cell assemblies). However, the selection of delays based on STDP when constrained by background oscillations, will lead to a narrow range of possible delays within the network to accomplish some

functional task. Indeed, when entrained via phase-locking, spikes will more likely occur at the same period of the phase rather than across the entire available time period; this consequently reduces the encoding/decoding potential of the network. Never the less, STDP and other forms of plasticity are likely to play a critical role in achieving complex spatiotemporal dynamics in diluted network structures so that reliable pathways can be established. Self-organizing mechanisms from ontogenesis to synaptic strengthening have also been reflected in synfire neural theories for the purposes of signal to noise enhancement and effective dynamics (Hertz & Prügel-Bennett, 1996; Doursat & Bienenstock, 2006; Doursat, 2013).

Despite these limitations, the polychronous neural group theory embodies one of the most recent approaches which demonstrate interesting features and some of the functional implications of precise spike time coding in biologically plausible neural networks. In the following we present existing work which explore the influence of various neurone features and network properties on neural coding via precise spike timing. In particular, emphasis will be given to transmission delays which, we argue, constitutes a key factor in neural organization and function.

2.2 Neural and Network Parameters of Influence

The possibility that spatiotemporal coding via precise spike timing is possible remains an open question to this day because the spatial and temporal resolution of our experimental (in vivo and in vitro) methods remain limited. Indeed, it is difficult to identify the characteristics of a spatial pattern over time taking place amongst 10^{10} or more neurones, while not knowing with high precision the individual temporal properties of each cell, their synaptic connectivity, and strengths. However, by knowing the critical features of individuals cells that exist in a particular tissue we can explore, with computational modelling, the space of parameters which are most likely to elicit interesting coding properties in the biological brain. To serve as reference for the present document we introduce the following descriptive neurone model activation equation:

$$\begin{aligned}
 a_i(t_n) &= \sum_{j=1}^{I_i} w_{ij} a_j(t_n - d_{ij}) + f(a_i(t_n - 1)) \quad \text{if } t_n - (R + t_{s_i}) > 0 \\
 a_i(t_n) &= f(a_i(t_n - 1)) \quad \text{otherwise}
 \end{aligned}
 \tag{2.3}$$

Where $a_i(t_n)$ is neurone i 's activation level in millivolts at time t_n , I_i the number of incident connections to neurone i , w_{ij} the connection weight from neurone j to i , d_{ij} the transmission delay in milliseconds from neurone j to i , and $f(x)$ a leak activation function with the approximate form $e^{kx} + l$. In this simplistic conditional model, the cell activation receives input from incident cells until it spikes at time t_{s_i} after which no input occurs until the refractory period R is passed. During this period the cell activation decays. In the case of coincidence detection, cells should only fire when a minimum number of incident connections T_i reach the neurone i ; hence we generally assume equal weights between all incident connections and set $w_{ij} = \max(a_i)/T_i$. While this model remains highly simplistic and does not capture the continuous dynamics of simple response models or the more complex conductance based models, it captures most of the essential components of neural response that we will focus on. In the following we describe relevant work on the fundamental role of neurones as well as network properties leading to coincidence detection behaviour and functional connectivity; these include: the membrane time constant, the firing threshold, the refractory period, connectivity, noise, and propagation delays.

2.2.1 The Membrane Time Constant

Based on the molecular physiology of their membrane, neural cells will have specific electric properties for the amount of charge they can store (their capacitance C) and the resistance to electrical flow on the surface of the membrane (the resistance R). Together these impact the time it takes for a cell to fluctuate electrically given a rate of charge transfer from inward and outward ionic flow. Traditionally, this dampening factor is estimated to be a constant known as the membrane time constant τ_m , where $\tau_m = RC$. This time constant plays a critical part in defining the role of neurones as either being sensitive to a number of time elapsed input pulses (integrator) or to a set of near-synchronous input spikes (coincidence detector). Indeed, this affects the rise-time and decay-time of the cell's activation which we capture in Equation 2.3 by function $f(x)$. Figure 2.6 illustrates the response time of two neurones at both sub-threshold and super-threshold activation levels, one with short membrane time constant and the other with a longer membrane time constant. Typical neurones require more than a single input spike to reach the threshold to initiate their own action potential. If the membrane time constant is slow (e.g. 20ms) then a single incoming spike will induce sub-threshold activity that will sustain a longer transient (longer time to recover) providing a sufficient time window for additional incoming

spikes to be added to the existing sub-threshold activity. If enough spikes reach the target then a post-synaptic spike will likely be produced. This slow membrane time constant, however, does not prevent near-synchronous incoming spikes from triggering a post-synaptic spike. The mere detection of coinciding spikes can thus take place in *integrator neurones*. What they lack, however, is the ability to eliminate spurious asynchronous spikes. Instead, a short time constant (e.g., 1ms) will make a cell's sub-threshold fluctuate much faster thereby losing any sub-threshold activity more quickly making it impossible for a sequence of incoming spikes to trigger it. Only when enough near-synchronous incoming spikes reach the soma will sufficient depolarization take place leading to a post-synaptic spike. This 'filtering' process is the defining factor of *coincidence detection* cells, making them more selective with respect to the source of their input. This is because only afferent cells firing with the adequate temporal offset to each other with respect to their conduction delays will reach a common target synchronously. Typically, this offset window can be defined as the ratio between the average inter-spike interval (ISI) and membrane time constant: $\frac{ISI}{\tau_m}$. Given this ratio it was initially proposed that we define a cell as a coincidence detector $\frac{ISI}{\tau_m} \gg 1$ (König et al., 1996). Since the recovery of cells is faster than the rate of spike input, this input train will not accumulate post-synaptically. A biologically realistic membrane time constant of 20ms suggests that average input firing rates of 50Hz and lower would not be sufficient to trigger a spike if the cell behaved as an integrator. However, the target will fire if a sufficient number of coincident spikes reach the target in what would otherwise be measured as a slow ISI. Interestingly, a large number of important processes happen at low frequencies, such as slow theta waves (2Hz-8Hz) in hippocampal function during exploration tasks (Colgin et al., 2009), alpha oscillations (9Hz-15Hz) to suppress distractions (Jensen, Gelfand, Kounios, & Lisman, 2002), beta oscillations (15Hz-30Hz) in motor control (Baker, 2007) and post visual stimulus processing (Vázquez Marrufo, Vaquero, Cardoso, & Gómez, 2001), and gamma (30Hz-60Hz) in perception, attention and planning tasks (C. M. Gray et al., 1989; Murthy & Fetz, 1992; Börgers et al., 2005) to name a few. Higher frequency firing is also known to take place however. While their presence is typically interpreted as meaning that encoding by first order statistics is taking place (integration), it is possible that high rate spiking regimes can, in certain circumstances, promote the chances of coincidence detection to occur. As we introduced in Section 1.4, it is actually likely that the continued presence of background activity results in greater sensitivity on behalf of cells to coinciding inputs. Recall how Bernander et al., for instance, showed that background activity, even of relatively low frequency could

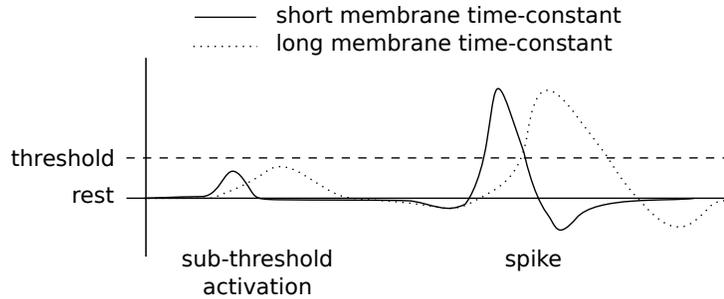


Figure 2.6: Illustration of two different neural activation responses. Rapid rise and fall occurs for cells with short membrane time constants. A longer response time takes place for longer membrane time constants. In reality this difference can vary 10 folds or more.

reduce the membrane time constant by one order of magnitude from 80ms down to 7ms (Bernander et al., 1991). Hence, these experimental results by Bernander and colleagues as well as Softky, suggest that the membrane time constant is not, in fact, a constant at all. The dynamic reduction of the membrane response time via background activity and/or inhibitory bombardment could lead to a faster activation response making the cell sensitive to spikes coinciding within a short time window only. This is often called the effective time constant of the membrane (τ_{eff}). Such mechanisms would enable coincidence detection at rates much higher than 50Hz. For instance, a 1ms window of sensitivity to incoming spikes would enable coincidence detection with input up to 1000Hz. Of course, if this sensitivity window is too small, then the probability that sufficiently near spikes will take place becomes decreasingly likely. Consequently, coincidence detection will more probably occur if other neural, network, and input conditions are adequately balanced.

2.2.2 The Firing Threshold

The threshold of depolarization at which a neurone produces a spike will significantly constrain its propensity to fire as a coincidence detector. The value of this threshold θ is only meaningful with respect to the synaptic strengths w_{ij} of individual inputs. If the threshold θ is very low then a single afferent input could trigger a post synaptic spike. This condition voids the relevance of CD as a coding mechanism in the network and the cell plays the role of a simple spike relay. On the other hand, a very large threshold could render a neurone extremely selective of its input and would require a high degree of synaptic connectivity from its neighbourhood. Fundamentally, coincidence detection becomes a relevant coding mechanism when the firing threshold (and

assuming a sufficiently short membrane response time) requires *at least* two coinciding input pulses. Biologically, we know that cortical neurones receive on average 10^4 synaptic connections and that most cells need between 20 to 100 PSPs to fire a post synaptic spike (Abeles, 1991; Gerstner & Kistler, 2002). If the integration period is broad, then a small number of cells could fire at relatively high frequency to induce a post synaptic spike. In the case of short membrane response times, however, this will require that 20 to 100 PSPs reach a target within a very short period of time to be efficacious.

Similar to the membrane response time, a neurone's firing threshold θ is also susceptible to change with respect to contextual electronic fluctuations. Within a similar set of experiments previously discussed, Bernander et al. showed that input resistance lowered 10 fold from $110M\Omega$ to $14M\Omega$ in the presence of background activity in their biophysical model of a pyramidal cell (Bernander et al., 1991). This leads to an effective lowering of the membrane threshold suggesting fewer spikes are required to trigger a post synaptic action potential. Such dynamic properties of biological neurones with respect to their environment can have potentially profound effects on information coding via coincidence detection.

The assumption that synaptic connectivity strengths are constant and equal between afferents to a cell's dendritic tree and soma can be an oversimplification if the objective is to address the parameters of influence on coincidence detection and functional connectivity. We present recent related work on the influence of connection topologies and synaptic weights on network dynamics.

2.2.3 Connectivity and Synaptic Strengths

Traditionally, neural network models were implemented within two general classes of networks: feed-forward models (Rumelhart & McClelland, 1986) and symmetrical connectivity models (Hopfield, 1982). Despite their functional value, their biological realism remains questionable. Indeed, projections from a large number of origins including feedback connections are prominent in the CNS (Abeles, 1982; Amit, 1989). Formally, biologically realistic networks will fall under the general category of attractor neural networks (ANNs). These are networks with asymmetric diluted connectivity, temporal delays, and disposing of mechanisms for learning static as well as sequences of patterns. Nützel et al. were amongst the first to study the dynamics of these networks using random connectivity (Nützel, Kien, Bauer, Altman, & Krey, 1994). In particular, they suggest that the degree of connectivity (dilution) for a cell

should not be determined by ‘proportional parametrization’ but should instead be set by a ‘power-law parametrization’. Hence, rather than attributing a number of incident connections I from N cells via some linear factor γ such that $I = \gamma N$, one should use a power-law scale such that $I = N^\gamma$ (where $0 \leq \gamma \leq 1$). The selection of this method stems from Derrida et al.’s careful consideration of the impact over time of preceding spikes within the population of neurones (Derrida, Gardner, & Zippelius, 1987). Indeed the state of a neurone at time t depends on the state of I neurones at time $t-1$, which all depend on I^2 neurones at time $t-2$, etc. Hence a neurone’s current state depends on the previous state of I^t cells. However, if $I^t \ll \sqrt{N}$ then no two neurones will have common ancestors (Nützel et al., 1994). This condition can be met if I is chosen such that $I < \log N$. The advantage of uncorrelated history between cells enables gradual degradation of stored patterns versus sudden degradation in highly correlated cases (large I) (Amit, Gutfreund, & Sompolinsky, 1985). In support of the power-law connectivity, Nützel et al. show that if linear scaling was used proportionally to the known number of neurones (about 10^{10}) and synapses per neurones (about 10^4) (Abeles, 1982), then we would have $\gamma = 10^{-6}$ making it an ineffective factor to use in small scale simulations. For instance, a simulated network of 10,000 cells would suggest that each cell receives at most 0.01 incident connections (or only 1 incident in every 100 cells). However, by knowing I and N from some system, one can estimate a power-law scale where $\gamma = 1 + \frac{\ln(I/N)}{\ln N}$. Using $I = 10^4$ and $N = 10^{10}$ we thus obtain a constant $\gamma = 0.4$. Hence, in a reduced network of 10,000 cells, we can estimate $10,000^{0.4}$ incident connections, i.e., $I = 39.8$. Importantly, Nützel et al. demonstrate that networks of various sizes preserve the same dynamics using this power-law scaling for the number of average incident connections (Nützel et al., 1994). Of course, the estimation of this constant will vary based on the neural tissue being modelled. This method, however, provides a sound basis for the a realistic setting of the degree of connectivity between cells when the biological proportions are known. Hence, while interesting dynamics may appear at a particular connectivity ratio, these dynamics may be unrealistic if the used γ is too removed from the tissue’s realistic γ .

Most modelling approaches assume uniform random connectivity between neurones in a network. While neural growth factors in the CNS during development likely promote a uniform diffusive arboration of axonal and dendritic projections, it is possible that developmental factors such as critical periods and plasticity lead to non-uniform connectivity in the early years of the brain. In fact, recent proposals suggests that certain networks in the brain possess small-world like connection

properties such as in layers 2 and 3 of the cortex (Hellwig, 2000), the brainstem reticular formation (Humphries, Gurner, & Prescott, 2006), and in small cortical networks (Song, Sjöström, Reigl, Nelson, & Chklovskii, 2005). Small-world networks like random networks, possess a short average path length between any two nodes; in addition, they also display the scale-free characteristic of high clustering (Watts & Strogatz, 1998). Simulations with these kinds of topologies have revealed a number of interesting properties. Using polychronous networks (c.f. Section 2.1.6), Vertes and Duke demonstrated that small-world topology improved recall of stored memories, in the form of polychronous neural groups, despite highly degraded input cues in comparison to randomly connected networks and scale-free networks (Vertes & Duke, 2010). In addition, they showed that memory capacity was much larger in small-world topologies versus random networks since an overlap of 70% of the input neurones in the small-world case retrieved memories just as well as completely separate input in the random network case. In a random network, however, degradation was significant as soon as the inputs overlapped. While scale-free networks showed some improvement in recall, they remained much more susceptible to degraded input than did small-world networks. These results suggest that small-world connectivity, if biologically realistic in neural tissues, could improve memory capacity in networks employing coincidence detection. Their effectiveness with respect to other cognitive functions such as categorization, sequential learning, response inhibition, etc. remains unknown, however.

In 2002, Barahona and Pecora showed that small-world networks of oscillators facilitate synchronization in comparison to random topologies (Barahona & Pecora, 2002). Small-world connectivity was also shown to promote persistent activity without continued input as well as display various attractor states based on the number of random long-distance connections p (Roxin, Riecke, & Solla, 2004). In particular, they showed that activity with a low p led to a bistable state of quiescence and sustained period activity. At intermediate p the system lost periodicity and was more prone to dissipation (failure to sustain). However, at high levels of p the system displayed again long periods of firing (quasi-sustained) but in a disordered (non-periodic) fashion. Furthermore, they showed that an increase in the propagation delay between cells will reduce the probability of failure to sustain for a broader range of p with a tradeoff of reduced periodicity.

Unfortunately, little is known with respect to the encoding relationship between an increased degree of incidence I on cells and the number of possible firing-chains within coincidence detection networks. While Nützel et al.'s work on the importance of scale-

free parametrization of the degree of incidence should be carefully considered when working on reduced models of biological spiking cells, their work originally applied to firing rate models where cells could continuously integrate over broad periods of time. Coincidence detection networks, however, may require a sufficient degree of synchronous incident spikes to maintain a signal. Because small-world topologies can enhance this it may be important to consider appropriate values for the scaling parameter γ in coincidence detection networks with both random and small-world connectivity.

2.2.4 Refractory Period

Coincidence detection at millisecond scales suggests that many neural factors may significantly influence coding at this time scale. Beyond the membrane time constant, we also know that biological neurones exhibit a substantial refractory period after emitting a spike during which no other action potential can be produced. This refraction is due to the inactivation of ionic channels after depolarization of a post-synaptic spike and is manifest in two forms: the absolute refractory period R_a which completely prevents a cell from firing no matter what input it receives, and a relative refractory period R_r which make it difficult for a cell to fire unless greater input is applied. The duration of the refractory period may vary depending on the type of neurone. Typically, cortical pyramidal cells will have R_a of a few milliseconds and R_r of a few centiseconds (Pauluis et al., 1999; Hasenstaub et al., 2005), while fast spiking neurones will likely have up to three folds shorter refractory periods than pyramidal cells. The absolute refraction thus imposes a hard limit on the frequency with which cells can fire while the relative refraction reduces the probability that firing takes place given synaptic input. Often, refraction in simplified integrate and fire models are ignored because the inter-spike interval is usually assumed to be greater than the time it takes for cells to return to their resting potential. While this may be a reasonable assumption in rate coding models, coincidence detection assumes that individual input channels to a cell can take place at any given moment. This suggests that refractory periods should not only be carefully considered in neurone models but also taken into account when investigating the dynamical and informational characteristics of coincident detection networks.

Functionally, refraction can also play an important role. Gerstner and Kistler showed that noise and broad postsynaptic potentials tend to disturb well defined spike packets (Gerstner & Kistler, 2002). If, in addition, the synaptic coupling is

strong enough then cells will start firing during the rising phase. However, a strong refractoriness will cease firing before the PSPs reach their maximum and restores sharp spiking. Hence, refractoriness was found to counteract the effects of noise and synaptic transmission and help maintain precise timing. This finding reinforces the the importance of refraction in coding mechanisms based on spike timing such as coincidence detection.

2.2.5 Input, Feedback and Noise

The purpose of neural coding is essentially to perform an operation based on the input provided to the network and its history of activation. These operations may include storage for memory purposes, retrieval of previously stored patterns based on a cuing input, transmission of possibly amplified spiking activity to other neuronal populations, and more generally transformations of the input data to achieve some form of representation to be used by other neural groups or to trigger some physiological response(s). While sparse or phasic² input can lead to sustained activity for dynamic storage under particular network conditions, increased input via tonic³ activity could be essential for continued processing tasks. Typically, a neural network will possess properties of interest when they operate in an intermediate regime of activity. While quiescent to low activity may suggest little to no information processing, exaggerated activity is likely to indicate saturation of firing with low information content.

Achieving balanced activity, however, has been explored thoroughly. For instance, early spiking models such as Aihara and colleague's Pulse Propagation Networks made use of a global feedback signal to limit run-away excitation in networks of purely excitable cells. For instance, Watanabe and Aihara used a global negative feedback computed from the network output at each time step which dynamically increased the membrane threshold of cells when global activity became too large (Watanabe & Aihara, 1997). This enabled a form of self-regulation of the network given some decay constant for the negative feedback preventing uncontrolled saturation. Importantly, this was an essential component in their simulation of functional connectivity in coincidence detection networks. Unfortunately, this method assumes that individual cells have access to some global activity value which is unrealistic from a biological point of view. At about the same time, Shadlen and Newsome demonstrated how the effect of inhibitory cells can give rise to 'balanced-inhibiton' leading to the balanced

²Activity that is temporally isolated.

³Activity that is regular.

yet irregular spiking often observed in biological recordings (Shadlen & Newsome, 1994, 1998). This suggests that models which include realistic amounts of inhibitory interneurons will result in regulated activity without the need for a global signal of spiking activity. Indeed, Wang and Zochowski recently re-emphasized the importance of balancing the more prominent excitatory cells (typically 80% of cells in the cortex) with inhibitory activity (Wang & Zochowski, 2012). They note that inhibition can have a regulatory role at both a local level such as in dendritic branches of hippocampal neurons (Liu, 2004) but also in global cortical dynamics (Haider & McCormick, 2009). Inhibition can also be effective for threshold control and activity stabilization (Hasselmo, Schnell, & Barkai, 1995). Furthermore, the balance of excitatory and inhibitory currents were also shown to play a role in signal gating (Vogels & Abbott, 2009). Wang & Zochowski thus conclude that inhibitory feedback can increase the dynamic range of a network's ability to respond to an input and improve its informational capacity.

Noise is another important factor to take into account when considering the dynamics and information capacity of spiking neural networks. Rate coding schemes are considered to be resilient to noise since misfiring will not significantly change the overall average firing frequency. In contrast, noise is thought to be highly problematic with respect to precise spike time codes since every spike can have a significant effect on the firing pattern over time. Under appropriate conditions, however, this may not necessarily be the case. Three types of noise can be considered in spike time coding models: additive spiking interference, subtractive spiking interference, and misfiring spiking interference. If excitatory, additive spiking can be handled quite well by CD networks. Because existing spikes will be sufficient to cause a target to fire, additional coincident spikes will have no effect on spiking. Additional spikes will only be problematic if a sufficient number arrive synchronously to induce an unwarranted post-synaptic spike. In fact, as suggested by Fujii et al., additive spiking, as seen in Figure 2.4, could help complete a volley of spikes by 'filling-in' when pre-synaptic cells misfire (Fujii et al., 1996). If inhibitory, then additive spikes may only become problematic if the number of excitatory synchronous cells reaching the target is barely over the necessary activation threshold and the additional inhibition prevents the target to depolarize sufficiently. Misfiring will lead to spikes that arrive asynchronously with respect to other volleys. If a sufficient number reach the a target cell, however, additional asynchronous spikes will not affect target cell firing since the coincidence criterion will not be met. CD networks thus 'filter-out' untimely spikes. Furthermore, the continued presence of sub-threshold noisy activation can lead to the effect

of stochastic resonance. Stochastic resonance takes place when input to a neurone is too low to transmit temporal information but the presence of noise enables spiking which matches the temporal structure of the input (Gerstner & Kistler, 2002). As demonstrated by Gerstner & Kistler, however, too much noise can lead to the undesirable state of constant rate spiking. Hence an intermediate noise level is likely optimal for transmission. Interestingly, it was also shown that by regulating the amount of additive gaussian noise, networks could be driven between states of temporal coding and rate coding thus leading to a dual coding paradigm (Masuda & Aihara, 2002; van Rossum, Turrigiano, & Nelson, 2002). More problematically, over-activation of inhibitory cells will subsequently reduce the possibility for post-synaptic targets to fire despite normal synchronous activation. The prevention of spikes is likely to induce a severe under-activation of cells leading to the destruction of potentially functional neural assemblies that form useful spatiotemporal patterns. However, subtractive spiking of this sort could be remedied via a dynamical decrease of the firing threshold or increase of the resting potential of the target cell. Interestingly, the latter was one of the effects found by Bernander et al. when context activity was applied to their pyramidal neurone model (Bernander et al., 1991). To gain a better understanding of the impact of noise on DCD networks, we investigate the effects of misfiring and subtractive spiking in Section 5.1.

Input, feedback and noise are critical to the functional investigation of spiking neural networks that make use of coincidence detection for information storage, transmission and processing. The balanced interplay between these factors is likely to differ with respect to spatial location in the CNS but also over time within the same areas of the brain. It is thus critical that a careful exploration of these factors be undertaken when developing more accurate models of spatiotemporal coding.

2.2.6 Transmission Delays

The central nervous system is comprised of billions of cells with high degrees of variation, intricate structure, and specific connectivity across many different tissues. These occupy precise localities with respect to each other where transmission can take place within small localized networks but also between regions separated by long distances. This intricate spatial structure is likely a critical factor underlying the complexity of cognition. However, while this structure bears significance with respect to the degree of functional specification of various areas of the brain, this spatial organization introduces an additional effect typically regarded as spurious or inconvenient,

that is, temporal delays. Indeed, the three dimensional nature of the brain imposes transmission delays from cell to cell that, we argue, are non-negligible. Such axonal or synaptic delays have been found in cortical layers to reach speeds of approximately 0.18 m/s for excitatory projections and 0.06 m/s for inhibitory connections (Murakoshi, Guo, & Ichinose, 1993; Salin & Prince, 1996). This suggests that within small local cortical networks where single projections in layer II/III pyramidal neurons reaching neighbouring cells as close as $300\mu\text{m}$ (Abeles, 1991) could take up to 1.7 ms to reach a target and about 5ms for an inhibitory signal to cover the same distance. Inter region connections are likely to take even longer, such as the 50ms delay measured between the inferior olive of the brain stem and the cerebellum (Gerstner & Kistler, 2002). Beyond distance, other biophysical factors such as myelination, electronic distances and other resistive properties of neural processes will also highly influence these transmission speeds.

While investigations of the role of temporal delays have made important contributions regarding their relationship to oscillatory and synchronized behaviours (Crook, Ermentrout, Vanier, & Bower, 1997; Pauluis et al., 1999; Jirsa, 2008; Campbell & Kobaevskiy, 2012) it remains unclear how physiological characteristics of neural networks relate to their coding and function for cognition. As presented above, the dynamic cell assembly and polychronous neural group views support the notion that transmission delays can lead to functional connectivity (or polychronous group activation) by virtue of introducing a means to coordinate in a time locked fashion the spiking patterns of interconnected cells within and across tissues. This property enables an exponential increase in the number of spike patterns with each additional neuron and added connection (Fujii et al., 1996; E. M. Izhikevich, 2006b). In addition, these networks become sensitive not only to the spatial organization of a signal (as do rate coding networks) but also to the temporal structure of this signal, therefore adding the capacity for time based signal discrimination despite possible similarity in its spatial origin. Furthermore, individual cells may participate in many independent spatiotemporal spiking patterns because firing is not simply dependent on physical connectivity and degree of activation but also sensitive to timely coincident action potentials. Similar to synaptic weights, each connection will possess an individual transmission delay. Differences in delays will be negligible if, as mentioned in the previous section, the inter-spike-interval between two separate channels are smaller than the integration time of the target cell, however, for very short integration time constants, small delay differences are likely to be detected.

With respect to tractability, spiking neural networks (SNNs) with transmission de-

lays were proven more powerful than traditional neural networks. Maass and Schmitt showed that *a single* binary spiking neurone with delays could compute any binary disjunctive normal form (DNF) function in virtue of coincidence detection, while binary threshold neuronal networks (generation 1, e.g., Hopfield networks, boltzman machines, perceptrons, etc.) require at least $\frac{n}{\log(n+1)}$ neurones (Maass & Schmitt, 1999; Paugam-Moisy & Bohte, 2009). Furthermore, while binary SNNs cannot outperform continuous threshold neural networks (generation 2, e.g., multi layered perceptrons, sigmoid activation NNs, etc.), they can do so by employing continuous activation functions such that the spike time can vary continuously depending on the activation level. Together with the significant spatiotemporal pattern sensitivity that they offer, this suggests that delay based coding via precise spike timing in continuously activated cells could be a very efficient form of coding exploited by the brain under particular conditions and within precisely tuned circuits. Indeed, the work previously discussed by Aihara and colleagues, Gerstner and Kistler, as well as Izhikevich and related groups have begun to substantiate the functional characteristics of these networks in memory, associative processes, and categorization tasks. This work, however, typically considers transmission delays between neurones or groups of neurones by simply averaging them to a single value or randomly distributing them with some degree of variation. Despite some relative successes, no traditional learning method such as backpropagation have been successfully applied to the learning of functions via the selection of synaptic weights and/or delays in SNNs (Paugam-Moisy & Bohte, 2009). Unfortunately, these methods typically make a number of assumptions such as simplified or regular networks structures and connectivity schemes making them less efficient with respect to biological networks. Furthermore, the attempt to solve formal problems is not a sufficient exploration of the dynamical nature of the joint effect of transmission delays and precise spike timing detection for cognition. Interestingly, more biologically plausible delay selection methods have been proposed such as ontogenetic development (Gerstner, 1996) and spike time dependent plasticity (STDP) (Gerstner & Kistler, 2002; E. M. Izhikevich, Gally, & Edelman, 2004). Gerstner et al. suggest that a form of hebbian learning during development could reinforce connections that match specific delay characteristics for functionally effective spatiotemporal patterns (Gerstner, 1996). Delayed connections, however, could also be continuously selected for via STDP even after development since closely coinciding spikes would be commonly reinforced when reaching a target. This is the primary means of neural group selection proposed by Izhikevich and other proponents of the polychronous approach. While both of these learning mechanisms are likely to play a

role in the CNS, to our knowledge, the phylogenetic approach of selecting delays via evolutionary principles remains unexplored. This approach suggests that genetically and epigenetically determined brain organization at both micro and macro scales may in fact be selected, in-part, for the delay characteristics that they give rise to. Later selection processes such as hebbian learning and STDP could then help refine the temporal scales at which delays operate as well as favour dynamic restructuring during a given agent’s lifetime. In the following section we will further motivate this research direction.

2.2.7 Informational Capacity

To date, the informational capacity of transmission delay spiking networks that are sensitive to coincident input is not well understood. Recently, Cessac et al. derived an upper bound estimate of the bitwise content present in a network of coincidence detection neurones with refractory periods for a given time window of activity (Cessac, Paugam-Moisy, & Viéville, 2010). By taking into account the refractory time R of cells after a spike and the narrow time window δt during which any two incident spikes are not distinguishable (typically $\delta t \approx 0.1ms$), they define the total information content in bits in a population of N spiking cells for a time period P to be:

$$N \frac{P}{R} \log_2 \left(\frac{P}{\delta t} \right); r > \delta t \tag{2.4}$$

Since there are $P/\delta t$ choices for the first spike, there are then $P/\delta t - 1$ choices for the next spike, etc. However, there are P/R total possible spikes during the time period P . Hence, Cessac et al. propose the upper bound $(P/\delta t)^{P/R}$ possible spike patterns during P , leading to the information content measure of equation (2.4). However, this measure assumes that $P/\delta t$ possible spikes can take place for each cell, yet the refractory period is precisely a limiting factor for this. Indeed, it is not the case that $P/\delta t$ states exist for a single combination of spikes during period P . As we revisit in Section 6.2.3, a cell can in fact achieve $R \times 2^{P/R}$ states of firing. Furthermore, while Cessac et al.’s measure attempts to take into account the activity dependent effect of refractory periods in a closed form, it ignores the combinatorial influence of multiple incident projections and the minimum required number of coincident inputs to reach the spiking threshold.

Related to the problem of information bearing, determining the probability with

which a spike will occur would provide a good indicator of the degree of activity to be expected in a network of coincidence detection neurones. Early estimates by Abeles led to the well established poisson spike probability distribution (Abeles, 1991). More recently, Gerstner and Kistler took into account connectivity λ (the ratio of incoming to outgoing connections to a cell), prior network activity (a_0), and coincidence detection threshold T to establish an accurate estimate, based on the cumulative distribution, of the probability that any given neurone will be active in a subsequent step (Gerstner & Kistler, 2002). This probability is expressed by the following:

$$a_1 = \sum_{k=T}^N \binom{N}{k} (a_0 \lambda N^{-1})^k (1 - a_0 \lambda N^{-1})^{N-k} \quad (2.5)$$

One sees how the minimum number of afferents to trigger a post-synaptic spike can range between T inputs up to N inputs, hence the sum of all possible afferent input to the target cell. As Gerstner and Kistler make note of, however, the ability to estimate this activity over time becomes increasingly limiting since this expression does not account for potential accumulation of correlations in the firing pattern. As discussed in Section 2.2.3 on connectivity, low degrees of incidence can remove correlation at the cost of limiting our understanding of a more complex class of networks. What this spiking probability estimate lacks, however, is a consideration for the delay and refractory features of networks and neurones. Indeed, spiking probabilities should vary broadly based on the ratio of the transmission delays to the refractory period of cells and the degree of variation in delays themselves.

Interestingly, related work on asymmetric attractor neural networks with delays can provide a number of insights which may be translatable to networks that use coincidence detection rather than synaptic weighting for signal propagation. In particular, this work identified two types of attractors for delayed neural networks: point attractors where cells, given a particular input, reached a terminal on or off state; and limit cycle attractors where cells continuously changed states over long periods of time (Amit et al., 1985; Derrida et al., 1987; Nützel et al., 1994). Interestingly, the length of these limit cycles was shown to increase with a larger number of cells but also with a broad distribution of delays. Furthermore, using a hebbian learning rule affecting connection weights, they showed that these networks could store input patterns as well as pattern sequences under appropriate connectivity conditions. Using a so called ‘loading-parameter’ $\alpha = p/I$, where p is the number of stored patterns and I

the number of input connections per cell, they argued that high input correlation due to a large I led to sudden memory failure at $\alpha \leq 0.14$ while a large number of patterns led to memory deterioration for $\alpha \geq 0.64$. A smaller degree of input to cells I with respect to the number of neurones ($\alpha \rightarrow 0$) resulted in the case where exactly one attractor per input was found in numerical simulations. Nützel et al. further showed via simulation that sequential versus static input did not vary network behaviours (Nützel et al., 1994). Interestingly, while spiking activity revolved around particular attractors corresponding to a learned input pattern, this activity may move away to another learned pattern spontaneously when the network is sufficiently diluted (i.e., for a smaller I). Generally, it may be the case that some dynamical outcomes found in these experiments will also be found in coincidence detection networks with transmission delays. However, we should also expect to find other dynamical features due to the much stricter temporal sensitivity of neurones. Counting the number of attractors and their type for a given set of inputs may be a good empirical strategy in determining the informational capacity of networks. Similarly, this approach can be applied to networks of coincidence detectors. Empirical evaluation of information storage is limiting however, since different networks will possess different informational properties new evaluations will have to be undertaken for every new network. Instead, it would be beneficial to have at our disposal a formula or set of formulae that capture this information content a priori given network and neurone parameters.

2.3 Summary

We have seen that existing work can be analyzed to form a cohesive outlook with respect to neural information processing via spatiotemporal codes. The original postulation by Abeles regarding the unique coding characteristics of coincidence detectors introduces a paradigm shift with respect to how we should interpret the functioning of many neurones in the CNS. We have also seen how other original contributions emphasizing the importance of transmission delays in biological networks brings this paradigm within a more complete physiological context with potentially significant functional outcome for cognition.

Crucially, important mechanisms in cognition such as sensorimotor coordination, stimulus binding, memory, learning, prediction, problem solving, and action selection may prove to be implementable and/or controllable via networks of transmission delay coincidence detectors. The current body of work, however, fails to incorporate important elements to provide such a proof. First, there lacks a cohesive perspective

as to the individual roles of the relevant neural and network properties on spiking dynamics. In particular, a cohesive assessment with respect to the potential of these networks must incorporate neurone properties including the sensitivity to coinciding inputs which is determined by the membrane integration time, the firing threshold, the refractory period as well as network properties which include the degree of interconnection between cells, the topology of this connectivity, and spatial/delay characteristics. Hence an adequate synthesis as to the role and relationships between neural and network features needs to be developed. Second, the existing body of work focusses exclusively on the local properties of networks from either informational, logical or local dynamics perspectives. Instead, it has become clear from theoretical literature in cognitive science that an adequate functional assessment of any cognitive paradigm should be validated via embodied and embedded tasks. Hence, while these traditional forms of assessment should help direct experimental trials, neural networks of delay coincidence detectors must be examined within the context of behavioural tasks which put to test target cognitive capacities. Finally, current theoretical work on these networks remains largely empirically driven. This limits predictability given a set of results to networks bearing very similar conditions. Instead, it would be advantageous to develop a set of equations which accurately capture the informational qualities of these networks. This kind of development could then shed light as to other functional attributes, yet unknown, that these networks may provide for cognition.

In the following chapter, we begin our investigation into the properties of networks of coincidence detection neurones connected with propagation delays (DCD networks). We first establish a conceptual framework which can serve to validate the functional application of DCD networks in embodied cognitive tasks. In line with this framework, we then present preliminary dynamical characteristics of these networks with respect to structural parameter values and explore the functional application for some neural mechanisms in a simple reactive embodied agent task. Finally, this leads to a distinction of the dynamical conditions which can be applied to higher functional tasks.

Chapter 3

Basic Properties of Delay Coincidence Detection Neural Networks

The goal of this chapter is to investigate, from an interdisciplinary perspective, how fundamental dynamic properties and basic coding potential of functional connectivity in delay coincidence detection neural networks (DCD networks) can apply to a simple cognitive task. We approach this from the perspective of embodied cognition, which we see as a means to substantiate the dynamical and informational qualities of neural processes into realistic functional tasks. In particular, we aim to identify the key dynamic regimes in DCD networks that are most amenable to fundamental reactive behaviour and identify those dynamics that could serve as scaffolds for more incrementally complex cognitive function. Before exposing details of the DCD neuronal model and the behavioural task, we begin by motivating the overall embodied approach to cognition and the manner in which neural coding via DCD networks fits within this scientific agenda.

Despite the valuable contributions by existing work on the spiking properties of coincidence detection networks, no known efforts have attempted to reconcile the informational and dynamical properties of this spatiotemporal paradigm as a means to enable agents to form an intelligent response to particular stimuli given restrictive bodily and environmental conditions. Two general behavioural categories are recognized for the classification of cognitive function. The first, known as *reactive behaviour* (or reactive system), designates behavioural function whereby a stimulus or set of stimuli elicit a direct response without the requirement for indirect complex processing – which we characterize in what follows. Reactive behaviours may be a posteriori adaptive in that they have been selected for over time as suitable solu-

tions for solving a particular set of problems (e.g., via evolution) and may also, at times, be pre-adaptive in that they are able to solve problems for which they had not originally been selected for. However, they lack any self-contained capability to self-regulate or modify initial response trajectories. In other words, for a particular input condition, the same output condition will always be generated. The second behavioural category, which we denote as *adaptive behaviour* (or adaptive system), are those sets of behaviours driven by systems which possess the self-contained means to produce a selective response to external stimuli. Importantly, adaptive behaviours will be generated by systems which exploit some form of memory¹. However, while memory may be necessary for adaptive behaviour, it is not sufficient. Indeed, other mechanisms must be in place such as selection or coordination of sensorimotor mechanisms. Distinguishing between reactive and adaptive behaviours can thus serve as an overarching method to classify and identify the types of structures and functions constitutive of agency, the characteristics of the environment affecting this agency, but also the nature of the task which the agent accomplishes as a function of that behaviour.

In the following, we begin by exposing different accounts which distinguish reactive from adaptive behaviours and argue that functional connectivity in DCD networks should in principle be sufficient to not only serve in reactive behavioural control but in adaptive behaviour as well. We then present recent work that we have conducted which demonstrates the ability of DCD networks at performing a reactive task when embedded as a robot controller. We then propose early theoretical considerations regarding in-principle functional applications for DCD networks including dynamic memory, dynamic association making and stability analysis. This provides for preliminary insights regarding DCD networks and sets the stage for the investigations we have conducted on more advanced adaptive behaviour in Chapter 4.

3.1 Distinguishing Behavioural Complexities

3.1.1 From Reaction to Adaptation

One of the most intuitive distinctions between reactive behaviour and adaptive behaviours is found by considering memory and learning mechanisms as clearly indicative of adaptive cognition. Indeed, the capacity to store useful information and

¹While memory could be self-contained, such systems could also exploit external forms of storage such as via stigmergy.

relationships to be used again during the agent's lifetime contributes significantly to an agent's capacity to cope with environmental change by facilitating the process of generating an internal model of the external world. By doing so, an agent then gains access to a cognitive foundation upon which prediction, problem solving, and planning can occur. Storage and learning, however, suggests that a form of dynamic internal organization takes place that is necessary for intelligent behaviour (Nolfi & Floreano, 2000). By building associations between stimuli, for instance, an agent can provide a delayed response to a stimulus, it could associate proprioceptive with external stimuli to form landmarks for spatial navigation, it could make use of stored stimuli for conditional reasoning when presented with novel problems, etc. On the one hand, the capacity for memory alone is not enough to provide a sufficient account of complex adaptive responses. While complex internal organization must take place for the storage of information, the capacity for memories of stimuli does not imply that knowledge about *how* to use information is acquired, i.e., how the selection of outcomes is processed. On the other hand, learning should provide for the relationships necessary to acquire this knowledge so long as a memory system is also available. In this context, learning can be understood as both the construction of sufficiently accurate internal models of sensorimotor outcomes and the improvement of response selection processes which satisfy some constraints based on sensorimotor and internal model states. Importantly, the construction of internal models should lead to a gain in capacity for information interpretation, while the mechanisms for selection of adequate responses processes addresses the problem of information use.

Learning should thus naturally appeal as a defining factor for adaptive behaviour. Interestingly, learning mechanisms can be thought of as existing at various timescales including phylogenic inheritance, ontogenetic development, and learning during an agent's lifetime. Indeed, each lead to the adaptation of the agent (or its characteristics) to its evolving environment. However, evolution and development lead to (transient) fixed structures. While inseparable from the function these structures give rise to, lifetime based learning requires dynamic organization of electrochemical activity as an ongoing process. For our purposes, we will focus on this latter form of learning. Learning in this sense depends on a multitude of constitutive functions such as memory, sensory processing, etc. Because of this, it may not be limited to self-contained physiological processes alone, but could encompass the extended environment of the agent. From stigmergy to note taking, suggestions of such an 'extended mind' have been made philosophically (A. Clark & Chalmers, 1998). This view, however, removes the location of the agent from the organizationally closed system under observation

to a system that includes environment and time (Adams & Aizawa, 2001; Somers & Jeanson, 2011). Instead, we focus here on systems which, conservatively, possess all self-contained features required for learning leading to either reactive or adaptive behaviours.

3.1.2 Towards a Formal Characterization

Ross Ashby was perhaps one of the first to offer a formal distinction between both reactive and adaptive behaviours (Ashby, 1956). His interpretation related more specifically to the inner function of an agent which should provide for a set of reactive functions (system R), meant to compensate for environmental changes to a degree proportional to these changes (albeit with limitations), and a set of second order functions (system S) which respond proportionally to the degree of change of vital internal variables and regulate these reactive functions. A reactive response is accomplished when reactive functions are sufficient to maintain equilibrium of the vital internal variables, while an adaptive response is realized when system S regulates the reactive system R to maintain internal variables within viable bounds. Ashby named this mechanism ‘ultrastability’, see Figure 3.1. Thus, reactive mechanisms part of R receive input only from the environment while secondary mechanisms part of S receive input from the environment and from internal indicators. Without having to cite explicit behaviours, Ashby’s account offers a strict distinction via implicit processes to the system. Interestingly, most problems in cognitive science are likely concerned with system S and the indirect internal regulation of behaviour. This view, however, suffers from two difficulties. First, it is not clear how to dissociate reactive structures from indirectly regulating ones since, arguably, from a deterministic viewpoint, internal vital states driving system S which in turn drive system R form together nothing but a complex reactive system. This criticism, however, may be rebutted by an argument of ‘degree of complexity’. That is, system S operates with a degree of complexity that remains qualitatively distinct from what any reactive system may look like. Second, this Ashbian dichotomy could impose too great a criterion for the attribute of adaptive behaviour since it would necessarily require that vital functions be threatened for an agent to perform intelligently. Indeed, it seems possible that advanced cognition can take place under non-threatening circumstances but should instead require that we identify the fundamental processes involved in behavioural complexity. Furthermore, while Ashby’s approach advances important questions with respect to cognition and functional hierarchies, his view makes no commitment with

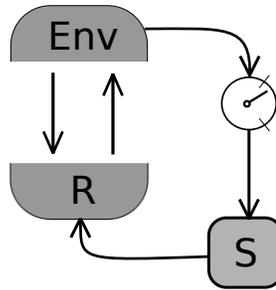


Figure 3.1: Ultrastability: an agent’s reactive system R can react to environmental changes autonomously but only within strict bounds. If vital internal variables (represented by the gauge) are sufficiently threatened then a secondary system S must actively regulate the reactive response system R to maintain these variables within viable bounds.

respect to the qualitative principles of the internal dynamics that accomplish these functions. This is, however, what we aim to apprehend.

Beyond control hierarchies, cognitive and/or behavioural distinctions can be in part defined by certain properties of a task to be accomplished. By the late 90’s, Clark and Thornton made an important contribution in this respect (a. Clark & Thornton, 1997). In their view, agents are confronted with complex tasks when the sensory system must find marginal or hidden regularities in an otherwise noisy sensory space. They denoted these as type-2 problems, while type-1 problems refer to simple sensory spaces for which a direct mapping of regularities are readily available from perception to action. Importantly, it was shown that certain control architectures and/or sensorimotor couplings could enable agents to transform type-2 problems into type-1 problems (Nolfi & Floreano, 2000). While Nolfi and Floreano do not recognize this ability as distinguishing between reactive from adaptive behaviour, there may be grounds to do so. In particular, solutions to the problem of transforming type-2 to type-1 problems will typically require either non-trivial forms of sensorimotor coordination and/or internal/external reconfiguration. For instance, Elman showed that complex tasks could not be solved by training feedforward neural networks with back-propagation (Elman, 1993). Instead, these networks first had to be trained on a more simple subtask so as to enhance regularities of the sensorimotor patterns. After doing so networks could be trained on complex tasks since non-trivial regularities could now be transformed from type-2 to type-1 in the sensory space. Hence, external reconfiguration of the problem served as a scaffolding method for advanced behaviour. In a different setting, autonomous agents were also shown to be capable

of transforming type-2 problems into type-1 in a cylinder classification task (Scheier, Pfeifer, & Kuniyoshi, 1998). When a simple Khepera robot with two wheels and 8 proximity sensors was placed in an arena that contained small and large cylinders, the sensorimotor controlling neural network of the agent could not learn to distinguish between these objects when placed statically at various locations in the arena. However, when the agent was allowed to control its own motion around the arena, it could learn behavioural patterns that would enable it to classify both object types. In contrast to Elman’s method, this demonstrated the effectiveness of autonomous restructuring of the sensorimotor space via *active perception*.

3.1.3 Final Considerations

Overall, approaches that distinguish reactive from adaptive behaviours should play a valuable role in assessing the functional features of neural networks. In particular, we will show how reactive behaviours can be best achieved under simple dynamic regimes in DCD networks, while Chapter 4 will reveal how more complex dynamics can serve sophisticated memory functions. In consonance with Ashby’s insights, we show that increased network dynamic complexity enables increased complex task solving. These results will show that functional connectivity in DCD networks provides a unique capacity for internal dynamic restructuring, possibly with some analogy to the external form of body-environment restructuring found in active perception.

While frequency coding is a firmly established form of neural coding it remains unclear what cognitive advantages spatiotemporal coding via DCD networks can achieve from an embodied perspective. Crucially, embodiment can provide a valuable form of assessment for this coding paradigm by validating neural dynamics for behavioural based functions that require the coordination of sensory input with motor output to solve both reactive and adaptive cognitive tasks. By testing these via computer simulation we can verify current hypotheses regarding their potential for neural coding as well as provide insights into their potential for reactive behaviour (in what follows) and adaptive behaviour (in Chapter 4).

3.2 Computational Simulation of Reactive Behaviour

We begin our investigation of neural activity in DCD networks by adopting an empirical approach which evaluates computationally the dynamic features of a phenomenological neural network model of spiking neurones that are uniformly connected with

propagation delays, coincidence detection thresholds, and refractory periods. We then conduct preliminary experiments to evaluate their spiking behaviours under various parametric and input conditions. These experiments are then applied in a robot simulation environment to test their capacity at solving reactive sensorimotor tasks. The goals of this work are to establish characterizations of DCD networks in which classes of behaviour are observed. Furthermore, these classes of behaviour are intended to support the requirements of reactive and adaptive (memory-based) controllers designed in this Chapter and Chapter 4.

3.2.1 Dynamical Analysis of Delays in Simple Networks

As originally suggested by McCulloch and Pitts, neural spiking can be essentially captured by binary threshold units where cells are either in a spiking state or not (McCulloch & Pitts, 1943). More specifically, when a cell receives enough input at time t and reaches a threshold value, this cell changes from a quiescent state to a spike state, leading to the initiation of a spike propagation along this cell's axon. At time $t + 1$ this cell then returns to a quiescent state and remains in this state for a fixed time period R which simulates the absolute refractory period of neural cells. If the input is not sufficient for a cell to spike then it is ignored, thus fulfilling the approximation of a short membrane time constant necessary for coincidence detection to occur. Neurone thresholds should be set to the minimal condition where a spike occurs only if two or more incident spikes arrive at the same time on a target. This fulfills the basic requirement for coincidence detection to take place. In addition, propagation delays are assigned to each connection to simulate the biologically realistic characteristic of spatial distances between cells and action potential transmission along axons, across synapses, and along dendritic trees. As an alternative to the binary threshold neurone, we could adopt more realistic conductance based models such as Hodgkin-Huxley models or continuous time integrated models such as a FitzHugh-Nagumo model. Indeed, in addition to improved biological fidelity these may improve coding capacity via continuous activation of post-synaptic potentials as shown by Maass and Schmitt and discussed in Section 2.2.6. If such models are adopted, it is essential that coincidence detection is preserved by ensuring short membrane time constants. Although a number of interesting dynamical properties arise from differentially integrated spiking neurone models, we believe simple discrete state neurones will allow a clear isolation of the essential coding characteristics that DCD networks possess. To exhibit the essential spiking characteristics of DCD networks

binary threshold neurones are sufficient in that they incorporate the essential characteristics of neuronal signalling where delays, refractory periods, coincidence threshold and connectivity are concerned. We describe this model in the following sections.

3.2.2 Delay Coincidence Detection Network Model

Although a number of interesting dynamical properties arise from differentially integrated spiking neurone models, we believe simple discrete state neurones will allow a clear isolation of the essential coding characteristics that propagation delay networks possess. For this reason, cells in our DCD network model have integer activation values with a negative or positive value. Axons have discretized lengths and either carry a value of 0 or 1 representing inactivity or a spike respectively. A single cell in the network behaves as an integer-valued coincidence detection and fire unit. When the activation of the node reaches its threshold value, a spike of value 1 is added to the beginning of the axonal paths leading to target nodes. If a node does not reach its threshold at given time step t then a 0 signal is set on the axonal spike train. Because nodes follow a basic coincidence detection scheme, we allow incoming signals to be either excitatory or inhibitory. Hence, an incoming spike can either cause the coincident activation to go up by 1 if the connection is excitatory or down by 1 if the connection is inhibitory. Excitation and inhibition are associated with connections rather than cells; hence any given cell can have both excitatory or inhibitory effects on their targets depending on the connection type². Importantly, a distinction is made between ‘strict coincidence detection’ where a target cell fires at time t only if the right number (the threshold) of incident spikes reach this cell at time t , and ‘loose coincidence detection’ where sub-threshold activation of the target cell may also contribute to total activation at time t . In the following description we restrict our model to the strict coincidence detection scenario only. If a node’s activation reaches a threshold T and spikes then the node’s activation is reset to 0 for the duration of a refractory period R during which no change of activation can take place. Activation of cells thus obey a deterministic threshold model which we more precisely describe here as:

²While this approach violates Dale’s principle that a cell’s target connections are either all excitatory or all inhibitory, it is consistent with recently observed biological findings in cortical pyramidal cells (Ren, Yoshimura, Takada, Horibe, & Komatsu, 2007) and has been found to significantly attenuate any global residual synchrony otherwise found in Dale-like networks (Brunel, 2000; Vogels, Rajan, & Abbott, 2005; Kriener, Tetzlaff, Aertsen, Diesmann, & Rotter, 2008).

$$\begin{aligned}
a_i(t_n) &= \Theta \left[\sum_{j=1}^{I_i} w_{ij} a_j(t_n - d_{ij}) \right] && \text{if } t_n - (R + t_{s_i}) > 0 \\
&\text{but } a_i(t_n) = a_i(t_{n-1}) && \text{if } f(a_i(t_n - 1)) > 0
\end{aligned} \tag{3.1}$$

Where $a_i(t_n)$ is the activation level of neurone i at time t_n such that $a_i \in \{-1, 0, 1\}$, I_i the number of incident connections to neurone i , w_{ij} the connection weight from neurone j to i which we set to 1 for all connections, and d_{ij} the discrete transmission delay in time steps from neurone j to i . Also, t_{s_i} is the time of the last spike for cell i , hence the cell activation receives input from incident cells until it spikes at time t_{s_i} after which no input occurs until the refractory period R is passed. Furthermore, $f(x)$ is the leak activation state of the cell which maintains neurones in a spike state for a period proportional to the simulation time step resolution. In our simulations this period was only a single time step. To detect thresholds we make use of Θ which is the Heavyside step function such that $\Theta(x) = 1$ if $x \geq T$ and $\Theta = 0$ if $x < T$. Finally, when a source cell is inhibitory its value is set to -1 instead of 1.

In the following, we first introduce preliminary simulation results of our DCD network activity given a simple input, which we expect to elicit non-trivial sustained firing with the ability to support temporal coding. We then adapt this network model to work within the simulated context of a Khepera robot performing a light seeking task.

3.2.3 Preliminary Test Model Parameters

Spiking neural networks with propagation delays, like many other networks, will exhibit a wide number of spiking behaviours based on their parameter settings (E. Izhikevich, 2003). Although discretization of the network helps reduce the model’s state space, a number of parameters will have a significant impact on the network’s behaviour. By basing these settings as closely as possible on experimental knowledge of biological networks, we expected the appearance of sustained firing as well as predictable global spiking powers in the frequency domain. Our initial test DCD network model included the following parameters: number of neurones (N), connection probability (C), activation threshold (T), refractory period (R), minimum delay (D_{min}), and delay range (D_{range}). Connections could also be excitatory or inhibitory which we set to the cortically plausible ratio of 85% excitatory versus 15% inhibitory (Peters & Yilmaz, 1993). Keeping parsimony in mind, we decided to start with a very small

number of $N = 25$ cells in the model. The connection probability was set high given the small number of cells to $C = 0.6$, resulting in an average of 360 total connections in the network³. Thus, each cell received an average of 14 incident connections ($I = 14$). Each connection had a synaptic weight of one, meaning that unlike other neural network models, connection weighting had no variable influence on network behaviour. Activation threshold was in contrast set to the minimally interesting condition where at least two spikes ($T = 2$) had to coincide on a target to trigger a post-synaptic spike; this also promoted a greater chance of observing coincidence detection events given the small number of neurones. Simulation time was discretized where we interpreted 10 time steps to correspond to 1ms; hence temporal resolution was one order of magnitude below the biological time scale of neuronal spiking events⁴. Discretization has the advantage of significantly reducing computation. However, the lack of noise often led cells to spike at identical times. Using this, the period during which cells could not be excited or depressed, including the absolute refractory period, was set to 10 time steps ($R = 10$) which aimed to reflect a biologically realistic duration of 1ms. Connection delays were also scaled in order to reflect realistic delay ranges above 1.5ms ($D_{min} = 15$) but never over 10ms ($D_{max} = 100$). Given that connection delays were of central interest in our investigation, we varied the delay range (D_{range}) that axons were uniformly randomly assigned from 0 to 30 time steps. Thus, a network with $D_{range} = 9$, for example, had a random distribution of connection delays between $D_{min} = 15$ to $D_{min} + 9 = 24$ time steps (i.e., 1.5ms to 2.4ms in our model). Overall connections and their delays were assigned with uniform random probability.

3.2.4 Results

Different ranges of delays within network connections were tested systematically between 15 time steps (15ts) up to 40ts. Three delay range conditions led to the identification of three critical dynamical states within these networks. Networks with narrow delay ranges between 15ts and 21ts reliably led to sustained firing with random inputs to approximately 30% of the population. Furthermore, these networks displayed highly regular synchronized firing. As seen from both the spike raster and spike count histogram of sample runs in Figure 3.2, spike variation slightly increased between each of the three cases: delay ranges 0ts, 3ts and 6ts. The frequency of these firing sequences was primarily driven by the average delay within the network.

³Neurone could not project to themselves and a uniform network was assumed.

⁴No significant differences in dynamics was found in simulations with 100 time steps or more per millisecond.

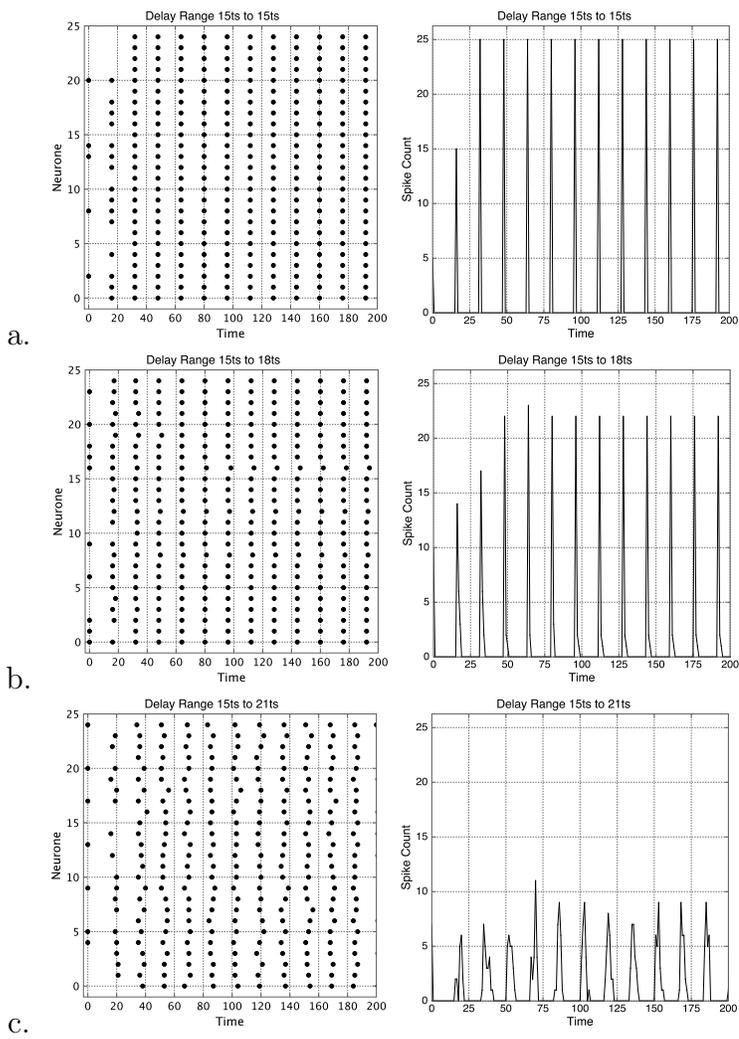


Figure 3.2: Narrow delay range activity. From top to bottom delay ranges 0ts, 3ts and 6ts. Left plots correspond to spike rasters and right plots correspond to spike count histograms.

Interestingly, when delay ranges were slightly increased by another 3ts (delay range 9ts) and more, strong synchronous firing became less apparent from the visual observation of the spike rasters and histograms. Figure 3.3 demonstrates this for delay ranges: 9ts, 12ts and 15ts. Despite the apparent absence of synchrony in these broader delay cases, power spectrum analysis revealed the presence of broad, yet small, frequency components about the mean of the average connection delay for each network category. Power spectra averaged over 30 simulations for each network case are shown in Figure 3.3. The horizontal line indicates 3 times the standard deviation of the signal suggesting subtle, yet significant, synchrony for all three delay ranges. Furthermore, during 1 second simulations (10,000ts) the total number of spikes in the narrow delay range cases reached 15,287 spikes versus 19,769 spikes in the broad delay range cases. This can be explained by a broader window of spiking opportunity for cells in networks with broad delays.

Interestingly, as delay ranges increased the probability of obtaining sustained firing within the network decreased rapidly. Figure 3.4 shows the average last spike time over 80 runs in narrow to very broad delay ranges. This suggests that networks are likely to reach a quiescent state when delays are too broad and the network input too small. This does not mean, however, that no interesting processing can take place in the case of rapid dissipation since continuous input could reveal interesting computational states for this condition. Here, however, we were initially interested in the cases where balanced activity could be reached for simple input conditions.

These results indicate that a critical threshold in the delay variation was reached between delay ranges 6ts and 9ts. In addition, the spiking dynamics are more complex in the broader delay condition. While highly stable oscillatory attractors were present within narrow delay range networks, less stability was present in broad delay range networks. A simple comparison in spiking activity for a network with delay range 9ts revealed that two input conditions that differ only by one neurone led to substantial divergence in activity over a few time steps. Figure 3.5 reveals spiking divergence after simulating the same network on slightly different input conditions. While most spikes coincide except for 4 on the first cell spikes, all but 5 neurones fire at different moments in time by the end of the simulation. This indicates that broad delay networks are more likely to possess chaotic-like characteristics thus supporting the notion that dynamical complexity increases as delay variability increases. Hence, in addition to the fact that a greater number of spikes take place during the same amount of time in broad versus narrow delay conditions, broad delay dynamics make it possible for the network to provide a larger set of firing-chain responses for independent input

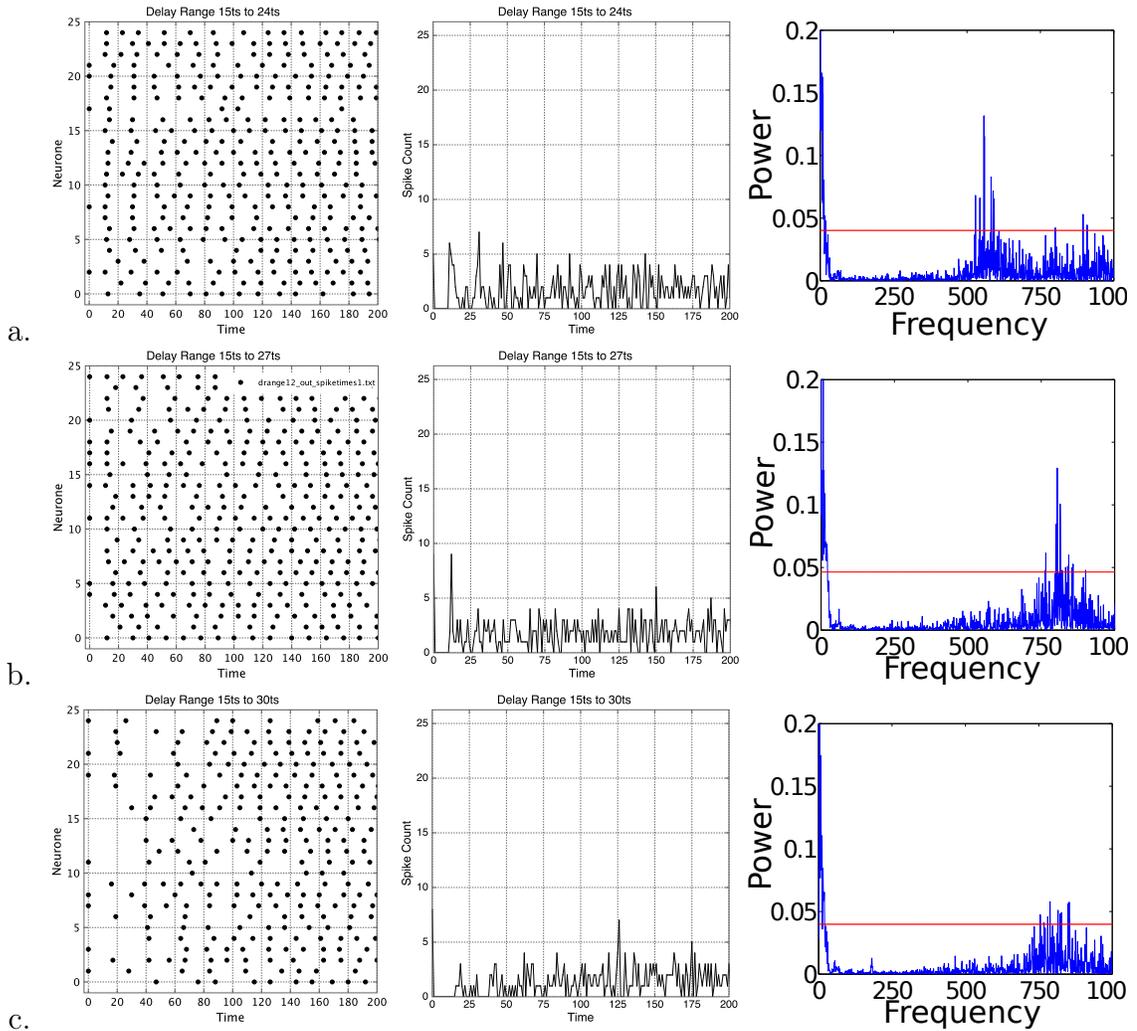


Figure 3.3: Broad delay range activity. From top to bottom delay ranges 9ts, 12ts and 15ts. Left plots correspond to spike rasters, center plots correspond to spike count histograms, right plots correspond to the population power spectrum with the horizontal line representing 3 standard deviations of the total power.

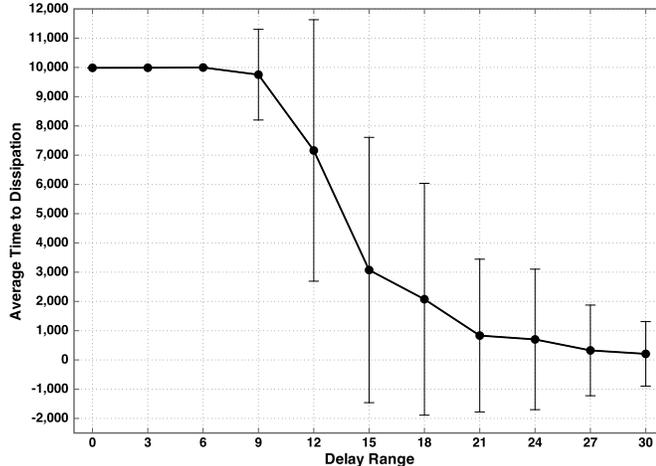


Figure 3.4: Average last time of spiking amongst various delay range conditions.

patterns. Interestingly, this is true for input patterns that are different in either their spatial arrangement, temporal arrangement or both. On the other hand, narrow delay ranges may be well suited for conditions requiring, for example, the recruitment of a large number of cells in comparison to the input, a more direct transfer of signalling between connected cells, the setting of a background frequency within the cell population, and other purposes.

Interestingly, the connection and activation properties that were found to work best in these experiments happen to fit closely, with some modification, to what Nützel et al. established as a power law scaling factor for the degree of incidents to cells (Nützel et al., 1994). According to them, the number of incidents in a firing rate model neurone should be set to $\gamma = N^{0.4}$. With 25 cells this would mean that approximately 3.62 incidents should be used on average ($\gamma = 3.6$). However, because our model makes use of coincidence detection cells, firing cannot be driven by single action potentials. However if we compute the ratio I/T we get close to the power law ratio proposed by Nützel et al. With a connection probability of 0.6 there are approximately 14 incident connections per cell ($I = 14$). For $T = 2$ we obtain $I/T = 7$ which indicate that correlation is relatively high in the network simulations tested here. However, for $T = 4$ we obtain $I/T = 3.5$ which approximates closely the power law ratio of $25^{0.4} = 3.62$. Activity will also be dependent on the amount of input to the network. In our case, 30% input led to sustained activity with adequate delay ranges when the threshold T was set to 2. If on the other hand, a full input to the network was applied, then a higher threshold of 4 may be more suitable. Nützel

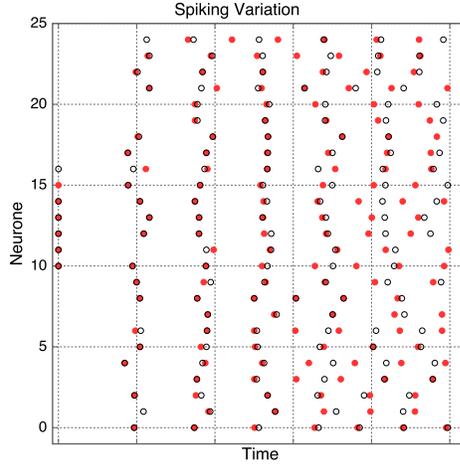


Figure 3.5: Spiking activity divergence between two similar input conditions on the same network.

et al. originally established their power law scale $\gamma = N^{0.4}$ based on general statistics regarding neural connectivity in the brain. For the case of coincidence detection, however, using the ratio $I/T = N^{0.4}$ may support the hypothesis that brain regions with high connectivity are perhaps more likely to function via coincidence detection networks. To compute T knowing N and the connection probability κ we evaluate:

$$T = \frac{(N - 1) * \kappa}{N^{0.4}} \quad (3.2)$$

Where T is the coincidence threshold, N the number of neurones and κ the connection probability between any two cells. N should replace $N-1$ if we admit self-connections in a network. Interestingly, this leads to a sub-linear increase of the coincidence detection threshold T with respect to a linear increase in the number of neurones N in the network which suggests that less variability in neuronal threshold is required despite large variability in network size⁵.

While spiking analysis of this network provided important insights with respect to its dynamics under various parametric conditions, we wished to investigate the potential functional role of this DCD network under various delay conditions for the control of simple embodied agents in a reactive behavioural task.

⁵This may constitute an appealing factor for evolution, development, and membrane dynamics since slight physiological changes due to genetic or developmental variations leading to subtle change in threshold sensitivity will make neurones adaptable to a broad set of network configurations.

3.2.5 Evolving Delays for a Light Seeking Task

To explore the capacity for DCD networks to serve as sensorimotor controllers we implemented this network architecture within the framework of a simulated Khepera robot (Michel, 1995). In particular, we opted to explore the ability for these networks to control a Khepera robot in a light-seeking task (Jeanson & White, 2012). Light-seeking requires that a mobile agent makes use of its sensory perception to respond to the presence of light and activate motor actuators so as to approach the light as closely as possible. While effective solutions, such as Braitenberg vehicles (Braitenberg, 1984), exist for the successful completion of this task, we wished to explore the space of potential connectivity and delays for DCD network based solutions. Indeed, DCD networks are highly sensitive to the precise timing of spikes and we did not wish to assume a specific decomposition from network output to input. Furthermore, it is unclear what set of connection delays would lead to an efficient solution to the task. Hence, to test the capacity for these networks to solve this task we opted for an evolutionary robotics approach where we evolved the connections and delays of the neural controller via a genetic algorithm. By doing so, we aimed to effectively explore the vast space of 3^{180} up to 8^{180} possible network configurations for the narrow and broad delay conditions respectively.

The rich set of dynamics obtained from the simulation of 25 neurones suggested that a similar sized network may be suitable for the light-seeking task. While 25 neurones were used, the number of possible connections was reduced from 60% to 30% to reduce the amount of over excitation in the network from repeated input neurone stimulation. For the same reason, the threshold of cells was increased to $T = 3$. Only excitatory connections were allowed so as to reduce the amount of complexity in the network. Transduction from two light sensors placed at 45° to the left and to the right of the robot was accomplished by stimulating 4 neurones per sensor in an inverse fashion where in the absence of light neurones were stimulated. This promoted exploration of the space when no light could be sensed. Both left and right wheeled motor actuators were activated by 4 neurones each as well. In this case, 3 neurones provided positive input to the wheel while 1 neurone transmitted a negative signal. Hence wheels could vary in speed between -10 and 30 by increments of 10. No minimum delay within the network was set. After running a few preliminary tests we noticed that sensory activation was too sparse to cause the robot to react to the light source. To remedy this we removed the refractory period of sensory cells allowing for much higher firing frequency than the rest of the cells which retained

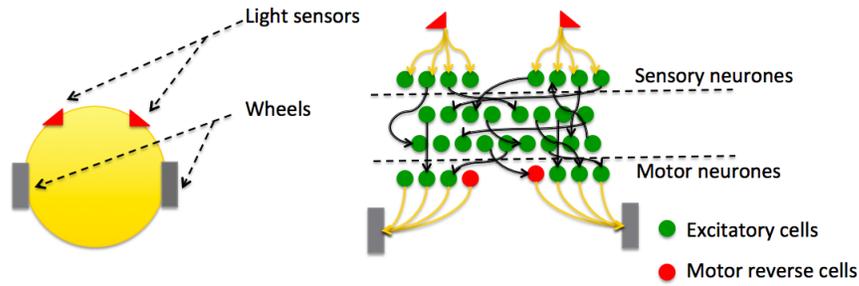


Figure 3.6: Illustration of the simulated robot with two light sensors and two wheels. In total, 25 neurones controlled the robot with 4 input neurones per light sensor and 4 neurones per motor effector.

their 10ts refractoriness. The overall architecture of the robot and its controller is depicted in Figure 3.6. The environment was a square surface of size 900mm * 900mm with a light placed in the centre. No other obstacles were present in the space. The simulated robot was allowed to be instantly repositioned on the opposite side of the arena when it reached the border. This removes the complexity of wall avoidance navigation.

The evolutionary algorithm was designed with a fitness function which computed the average distance of robots from the position of the light over 3 independent trials of the same agent with random repositioning. This prevented behavioural specialization to a single spatial location. The population was composed of 25 agents that were independently tested in the arena. Evolution consisted of 1000 generations where each individual was tested during each generation. The top 20% individuals were selected for their genotype from which two were picked randomly. These two genotypes were crossed-over with 50% probability forming a new genotype which was subsequently mutated with 50% probability using a gaussian distribution with mean 0.3. Genes were encoded in the range 0 to 1. A gene with value 0 represented the special case of ‘no-connection’ while non-zero values were scaled to the delay range plus minimum delay when transcribed to the phenotype of the network. Hence mutation was applied only to non-zero genes so as to restrict the amount of connections within the network. The gaussian mean of 0.3 signified that a delay was shifted by approximately 1ts when mutation took place. An example genotype is illustrated in Figure 3.7.

Results showed that the average population fitness improved over generations and typically reached a plateau by approximately 300 generations. In particular,

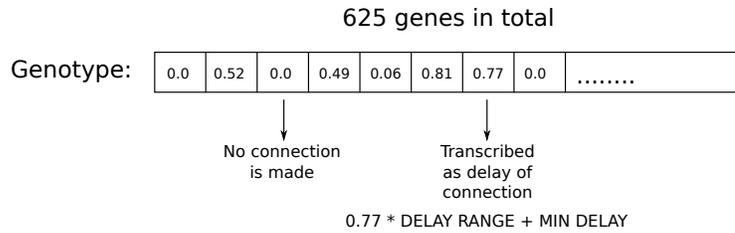


Figure 3.7: Illustrated genotype for a single agent controller. Genes with value 0 defined the non-existence of a connection between the corresponding two cells. A non-zero value was transcribed as an existing connection with a delay scaled to the delay range of interest during the experiment and an added minimum delay.

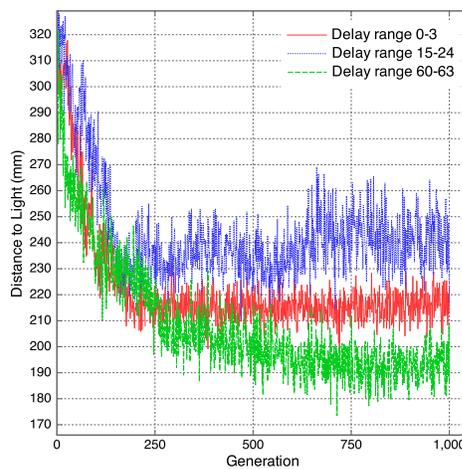


Figure 3.8: Average population fitness for the three tested delay conditions. Lower values represent greater fitness since fitness was directly measured as the average distance from the light during each individual’s life-time.

three delay conditions tested showed a behavioural improvement; these conditions were: narrow delay range (0ts to 3ts), broad delay conditions (15ts to 23ts) and narrow delay range with 60ts offset (60ts to 63ts). We found that in the narrow delay condition (population A) fitness reached an average distance to the light source of approximately 219mm, in the broad delay condition agents did not fair as well with an average fitness (distance from light) of 244mm, while in the narrow distance scenario with a high delay offset (population C) the average fitness was approximately 196mm by the final generation. Figure 3.8 shows the average population fitnesses for each of these evolutionary scenarios. A one-way ANOVA between population fitnesses revealed that the evolved performance between these three types of delay scenarios was significant with $F(2, 2997) = 735.2, p < 0.001$.

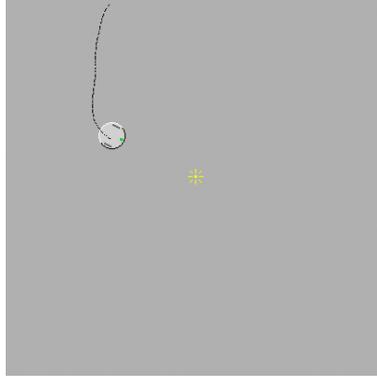


Figure 3.9: Example light-seeking behaviour for an agent with narrow delays (0 to 3 time steps).

In the narrow delay condition, agents displayed moderate behavioural variation during evolution. By the final generation, agents typically moved in a linear fashion with slight jitter leading to an efficient exploration of the space and came to a quick stop as soon as the light source was perceived by either the left or right light sensors. Figure 3.9 shows a typical agent stop near a light and the drawn path of its motion in the space when the controller has a narrow delay distribution. Agents which evolved controllers with broad delays, however, displayed much greater behavioural variation during evolution. By the final generation, most agents displayed a spinning behaviour with either preference for left or right spins with slight drift depending on the evolutionary run. This led to a less time efficient exploration of the space since only a partial area of the arena could be explored by spinning with drift. Figure 3.10 illustrates the typical motion of an agent with broad connection delays in the arena. Finally, in the narrow delay condition but with greater delay offset, agents displayed similar linear motion behaviours as in the first narrow delay condition. However, the presence of longer delays in the network induced a delayed response on behalf of robots when they perceived the light source. This resulted in a delayed stop and allowed agents to approach the light closer than in the simple narrow delay case. Figure 3.11 illustrates this behaviour. This led to an overall greater fitness for agents of this population but did not lead to qualitative behavioural differences beyond a delayed response. Computing the power spectrum for these individual cases did not result in noticeably significant frequency components in any of the delay cases. While we expect heightened synchrony in the narrow delay controlled behaviours, the stimulus induced variation as well as the relatively little amount of spiking during particular actions made it difficult to isolate frequency components.

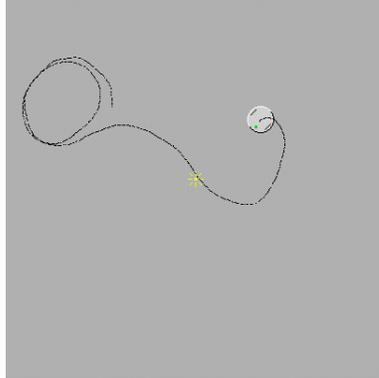


Figure 3.10: Example light-seeking behaviour for an agent with broad delays (15 to 24 time steps).

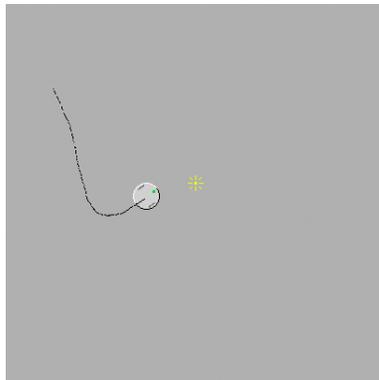


Figure 3.11: Example light-seeking behaviour for an agent with narrow delays but high delay offset (60 to 63 time steps).

Overall, these results suggest that narrow delay conditions led to a more efficient sensorimotor coordination for this simple reactive task. This can be explained by a more effective recruitment of motor neurones when relevant sensory stimuli were produced. By contrast, while the broad delay condition lead to more diverse behaviours, these were not advantageous for the efficient completion of this light-seeking task. From these findings, we expect that narrow delay dynamics may be more suited for the control of simple reactive behaviours as well as facilitating the propagation and communication of structured signals across different areas of the neural network. A high degree of variation in transmission delay coincidence detection networks, however, is more likely to serve the purpose of complex transformation of spike patterns for localized computation.

As a result, these experiments provided a proof of concept for the feasibility of applying DCD networks in robot control. In addition, we isolated a critical condition whereby a slight increase in the variability of delays within these networks led to quantitative and qualitative differences in both spiking dynamics and agent behaviours. What remains to be explored, however, is the degree of cognitive and behavioural control that can be achieved using these networks; particularly in the case of broad delay connectivity. Based on this initial experiment we thus wish to explore the applicability of these networks for functional connectivity in robot tasks so as to provide further leverage to the theoretical claim that delay coincidence detection networks can serve functional roles via the selection of appropriate delays. In the following we present the early theoretical perspective that we initially developed and served to guide the subsequent research.

3.3 Early Theoretical Considerations

Having shown that DCD networks can give rise to interesting dynamics as well as control simple robots in a reactive task motivates further exploration into the capabilities of these networks at solving more complex tasks. Specifically, we would like to explore the ability of DCD networks to provide the cognitive capability of memory. A number of neuronal mechanisms are currently well understood which likely belong to the repertoire of memory mechanisms used by the brain. These include various forms of Hebbian learning (Hebb, 1949) and the more recently established Spike Time Dependent Plasticity (STDP) (Bi & Poo, 1998). Together, these mechanisms are generally understood to primarily affect the connection strength between any two neural cells. Strengthening of the connection typically leads to a local form

of learning while weakening of the connection constitutes loss of learned relationships. While such mechanisms are known to be highly important in a number of brain areas, there remains a degree of uncertainty with respect to mechanisms responsible for short term storage and association making in the associative areas of the brain, such as the parietal, temporal and prefrontal areas of the neocortex. Goldman-Rakic as well as Fuster, for instance, argued that the prefrontal cortex must dispose of active memory mechanisms where neurones fire selectively with respect to stimuli so as to guide behavioural responses via sustained neural activity (Goldman-Rakic, 1987; Fuster, 1989).

As an alternative to the metabolic change of synaptic strengths, ideas on how the dynamic maintenance of stored signals can take place via recurrent connectivity within a population of neurones have been explored since the 1970's (Cowan, 1972). With respect to DCD networks, we have previously argued that synaptic plasticity is likely to be too slow as a mechanism to account for sub-second feature storage and the rapid learning of temporary associations, but also that dynamic memory and dynamic learning is possible in principle in DCD networks (Jeanson, 2011). More specifically, the vast number of possible firing-chains within a DCD network of moderate size suggests that networks of interconnected cells with appropriately tuned delays (via evolution, development, and/or plasticity) may deliver a unique firing response with respect to individual sensory stimuli. If, in addition, these responses have self-sustaining qualities then sensory memory traces become possible within these networks. Two types of dynamic memories are possible in this fashion. In the first case, a more traditional form of a memory can exist as the continued activation of the identical spatiotemporal pattern over a period of time. This can occur if a precise spatiotemporal pattern in the population of neurones at some time t is triggered again by the firing-chain at a future time $t + n$. Figure 3.12 illustrates this principle. Here, a subset of three cells (purple) initially triggered and lead to a unique firing-chain that eventually triggers this same set of initial cells to fire with identical temporal structure. The dynamics of the network would thus reach a limit cycle attractor in the form of a repeating firing-chain. The advantage of this form of dynamic memory is that the identical firing-chain will be triggered repeatedly until sufficient interference or dissipation takes place. This repeated activation could work well in concert with mechanisms of synaptic plasticity so as to reinforce connections within this network, such as proposed by Izhikevich's polychronous neural group theory, or even reinforce projections to third party neural populations. A second memory mechanism, however, may be instantiated without precisely repeating firing-chains, such as a chaotic

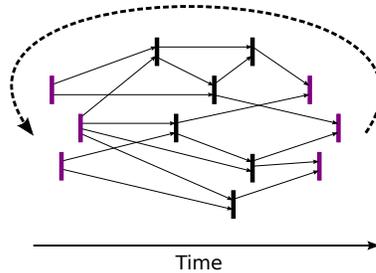


Figure 3.12: Memory: a sub-set of three cells (in purple) fire within a specific temporal structure leading to a firing-chain that triggers at a later point in time the identical initial three cell spike pattern. This results in a self-sustaining spatiotemporal spike pattern in the form of a limit-cycle.

pattern, so long as sustained activity is maintained. Indeed, a population which fires in a sustained fashion but without repeating spiking patterns could remain uniquely identifiable by connected neural tissue over an extended period of time.

Jeanson argued that DCD networks possess, in principle, another valuable characteristic which may be of significant use for adaptive behaviour (Jeanson, 2011). This characteristic is the ability for DCD networks to build dynamical associations between individual stimuli presented during the same time window. Indeed, if two independent firing-chains A and B co-occur in response to two individual stimuli then, by virtue of reaching other target cells coincidentally, they may trigger together a third firing-chain C . If this third firing-chain C is self-sustaining then a stable dynamic association of both stimuli is created. Figure 3.13 illustrates graphically this possible scenario within a population of delay coincidence detection neurones. While this principle could be confirmed via network simulation alone, an embodied task requiring the ability for association would bring stronger validation with respect to its plausibility within the nervous system.

Interestingly, in the preliminary experiments discussed above, sustained activity was achieved in both narrow delay and broad delay conditions. However, narrow delays typically led to global activity of the network with global synchronous firing which suggests a reduced variety of features from the spiking response. This lack of feature variety implies that little differentiation could be made despite differences in the input patterns. Unless firing frequency of the overall population can be controlled, this spiking mode is less suitable as a rich dynamic memory store. Broad delays, however, recruited less cells to fire synchronously and led to highly diverse spiking features suggesting that they could be applied as a rich mechanism for dynamic

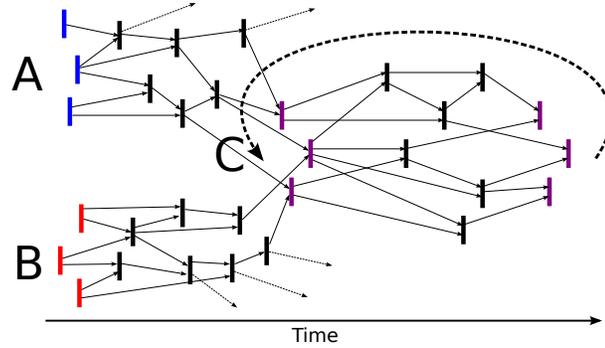


Figure 3.13: Association: Two stimulus driven firing-chains A and B trigger a third self-sustaining firing-chain C.

memory without the need to control firing frequency.

As a natural progression from these early theoretical considerations, the following section and chapter explore the capacity for DCD network to produce self-sustaining memory patterns with chaotic and limit-cycle dynamics (Section 3.3.1), how to decode them (Section 4.2), and associating input stimuli in an embodied dynamic memory task (Section 4.3).

3.3.1 Sustained and Stable Firing

While integrate-and-fire neurones will indistinctly detect structured or unstructured spatiotemporal patterns, delay coincidence detection neurones will be sensitive to the changes of spiking patterns over time. Importantly, the activity which results from DCD networks will also, in turn, lead to spatially and temporally specific outputs which we have referred to as firing-chains. As illustrated in Figure 3.12 these patterns are spatial with respect to the activation across a number of cells in the population and temporal due to the continuous change in activation over time. Furthermore, we, among others, have shown that these deterministic patterns exhibit fixed, cyclical, and chaotic dynamic behaviours (E. M. Izhikevich, 2006a; Jeanson & White, 2012; Suemitsu & Nara, 2003). Despite this, their informational characteristics remain poorly understood.

We have shown in Section 3.2.3 that a simple discrete neurone model implementing a delay coincidence detection neural network can exhibit in-phase synchronous firing with narrow delay ranges while complex firing can take place under broad delay conditions (Jeanson & White, 2012). We now expose an additional dynamic characteristic of the these networks under different connectivity regimes which was published

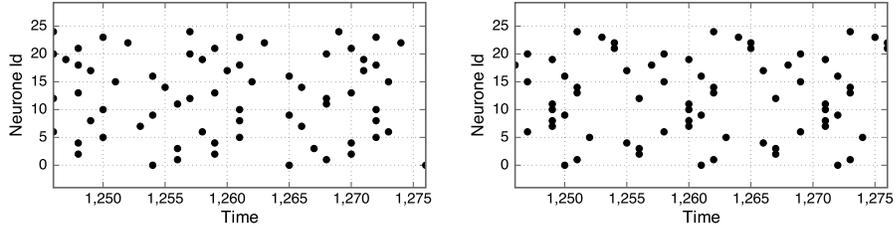


Figure 3.14: Left: Broad delay spike activity with low connectivity of 60% leads to unpredictable spiking. Right: Broad delay spike activity with high connectivity of 100% leads to stable limit-cycle response patterns.

in GECCO 2013 (Jeanson & White, 2013b). Here, we implemented a network of 25 coincidence detection cells ($N = 25$) that fire if two or more spikes reach a cell synchronously ($T = 2$). All connections transmitted excitatory spikes. Furthermore, an absolute refractory period of 10 time steps ($R = 10$) was used along with minimum delays of 20 time steps ($D_{min} = 20$) and maximum delays of 40 time steps ($D_{max} = 40$) randomly distributed; thus resulting in a typical broad delay connectivity structure. Under these network conditions, we found that firing-chains could reach chaotic-like attractors when connectivity remained moderate at 60% ($I = 14$ on average), while highly stable limit-cycle attractors were reached when connectivity was high at 100% ($I = 24$ on average). A sample run is shown in Figure 3.14.

To quantify this stability, we identify the spiking regularity s for a neurone to find how close its period converges to a target period λ over multiple renewal processes. To do so, we sum n consecutive spike time differences, where we express the time difference between two consecutive spike times as $\nu_{t+1} - \nu_t$. This sum is computed for individual neurones across a time interval of interest. We then average this sum over all n measured times the expected limit cycle λ to obtain a measure of divergence from this limit cycle period. This is summarized by the following equation:

$$s = \frac{\sum_{t=0}^n \nu_{t+1} - \nu_t}{n * \lambda} \quad (3.3)$$

Here, spike-time differences were measured from time steps 1000 to 3000 and the target period λ set to 11 time steps (11 ts). This λ corresponds to the refractory period of cells + 1 ts. When s for a particular neurone approaches 0 the spike pattern has period converging to λ . Figure 3.15 shows the spiking stability for each of the 25 neurones in both the low connectivity (leading to unpredictable spiking) versus high connectivity (leading to stable regular spiking).

Importantly, the limit-cycle behaviour of the network under high connectivity con-

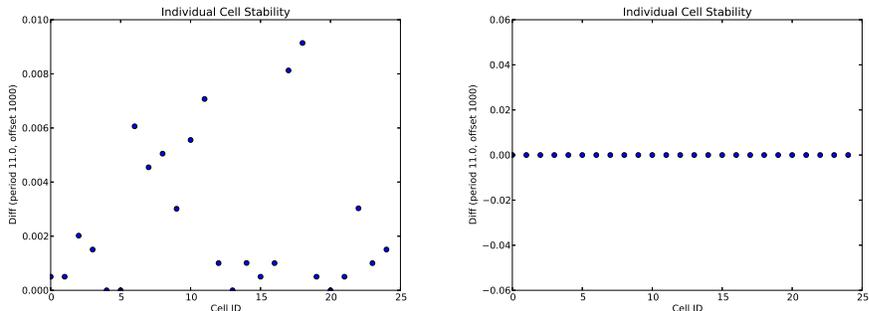


Figure 3.15: Left: Spiking stability s for each cell in a network response with 60% connectivity and broad delays. Right: Spiking stability s for a network with 100% connectivity and broad delays.

ditions suggests that a self-sustaining spatiotemporal spiking pattern with potentially unique response characteristics with respect to independent inputs can be achieved, as initially hypothesized and as seen in Figure 3.12. This mode of spiking thus represents an important special condition of DCD network function that can be exploited for dynamic memory purposes. We explore this potential in the following Chapter 4. Importantly, a network implementing model parameters presented here with $N = 25$, $R = 10$, $I = 24$, $T = 2$, and $D_{range} = 20$ will be referred to, in the remainder of this thesis, as a Standard DCD network.

3.4 Summary

In this chapter we began by outlining the conceptual distinction between reactive and adaptive behaviours, how these relate to distinct cognitive capabilities and system criteria. In particular, we argued that while systems that produce reactive behaviours may implement mechanisms require no internal restructuring, that is, where a direct mapping in sensorimotor coordination occurs. However, systems producing adaptive behaviours must possess the capacity to regulated their internal dynamics such that stored stimuli can be coordinated with novel stimuli to produce an non reactive response. To further complete the notion of adaptive behaviour we then introduced existing theoretical and empirically based approaches which describe adaptability in cognitive agency. We then proceeded to explore via an computer simulation the characteristics for DCD networks at accomplishing reactive behavioural tasks. After describing the basic dynamics of a simple binary threshold DCD network model, it was shown that narrow delay distributions among the network connections led to

strong synchronous responses, while broad delay distributions led to complex firing-chain responses. These distinct modes of connectivity, and consequently firing, were tested in an embodied reactive task requiring a robot to explore the space and move towards a light source. By using an evolutionary robotics method we were able to show that narrow delay distributions were preferred as a result of evolutionary selection, while broad delay connectivity led to complex agent behaviours with less optimal solutions for the reactive task. Finally, we further explored the spiking qualities of DCD networks with broad delay connectivity distributions and showed that low network connectivity led to self-sustained chaotic-like dynamics, while high connectivity led to self-sustained limit-cycle patterns of activation.

While the dynamics found here may represent idealized network responses with no noise or natural dissipation, we believe that they capture the essential characteristics of biological neurones which behave as coincidence detectors and are interconnected with realistic delay propagations.

Chapter 4

Dynamic Memory: Beyond Sensorimotor Coordination

In this chapter we focus on exploring, via additional network and agent simulations, the ability for DCD networks to be applied to storage, retrieval, and coordination of sensory stimuli for robot control. This work thus paves the way for the embodied application of spatiotemporal coding with DCD networks in tasks eliciting behaviours beyond the reactive and towards increasingly adaptive capabilities.

4.1 Why Dynamic Memory?

Learning and memory are primarily attributed to the various functional mechanisms behind connection strengthening and weakening such as Hebbian learning, spike time dependent plasticity, among others. While these are important mechanisms of neural function, alternative mechanisms should not be ignored. Even the quickest synaptic plasticity mechanisms elicit relatively slow metabolic processes that typically take over 500ms (F. Stevens & Wesseling, 1999) when compared to rapid stimulus encoding which can take place in the range of 50ms (Sperling, 1960). The presence of active short-term mechanisms at the millisecond time scale are suggested to account for scene segmentation, rapid categorization, delayed responses, etc. (S. J. Thorpe & Imbert, 1989). In particular, mechanisms based on sustained neural dynamics via re-entrant signalling have been proposed since the 1970's (Cowan, 1972; Zipser, Kehoe, Littlewort, & Fuster, 1993). While a substantial amount of work has focussed on overall spike densities as possible memory signatures, few researchers have investigated their potential in realistic embodied agent tasks. As suggested by Abeles and

others since, spatiotemporal mechanisms are likely to play a significant role for the storage of neural activity for delayed use, transformation, or coordination in order to accomplish cognitive tasks (Abeles, 1991). Here, we aim to investigate the potentially significant role for delay coincidence detection networks (DCDN) to serve as a neural platform for memory function that is not grounded on the modification of synaptic strengths but is instead founded on the precise spike timing amongst a sufficient number of cells to form dynamic spiking patterns that can be combined with other incoming spikes but also decoded for a specific functional purpose. Because stored patterns, in this case, are not encoded by metabolic change but instead via precise spike time sequences that remain active over time, we denote this kind memory mechanism as a *dynamic memory*. As mentioned in previous sections of this thesis, contemporary interest in coincidence detection and transmission delays has led to a number of interesting findings with respect to stability criteria (Arik, 2000), oscillatory dynamics and the emergence of functional groups (i.e., polychronous neural groups) (E. M. Izhikevich, 2006a), as well as neurally based logic circuitry (Fernando, 2011). However, no known effort has been dedicated to apply the distinctive dynamics of DCD networks for the generation and decoding of dynamic memory function to achieve precise control of embodied agents.

In order to classify a system as a memory system we argue that two distinct conditions must be satisfied (Jeanson & White, 2013c). First, a memory system must exhibit a *differential* response with respect to recall cues. That is, responses should be distinguishable from other cued responses. This comes as a requirement since a unique response should be elicited by the memory system so as to distinguish it from the recall of other elicited patterns from the system. Second, a memory system must *reliably* produce a pattern which the decoding mechanism is expecting. That is, while the same cue may not always produce the same response (although for simple memory systems it will), what matters is that the interpretability by some decoding process remains unchanged. In the simple case, memories will be identical in spatial and/or temporal characteristics for the same recall cues repeated in time, but a system could conceivably produce a varying pattern, despite identical input, so long as the decoder varies correspondingly.

Importantly, both of these conditions are obtained in DCD networks without recourse to structural or metabolic change. Interestingly, as seen in Subsection 3.3.1, the presence of distinct spatiotemporal limit cycles in response to varying input patterns makes it possible to use these networks as dynamic memories. Furthermore, we believe this can be accomplished statistically for any spatiotemporal pattern with

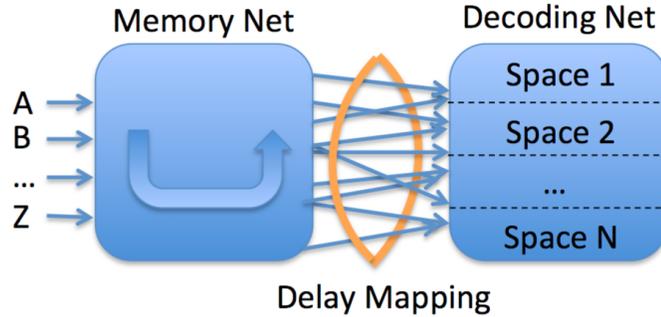


Figure 4.1: Memory network and decoding network configuration. A DCD memory network responding via broad delay limit cycles to some inputs (A ... Z) can be decoded into spatial patterns via delay mapping to a target network of coincidence detectors.

sufficient stability despite the presence of uncorrelated spikes and noise.

4.2 Delay Mapping

To effectively decode spatiotemporal memories into spatial activation patterns in a target network we developed a procedure called ‘delay mapping’ (Jeanson & White, 2013b). After describing this delay mapping method, we present a newly developed procedure which identifies spike timing ambiguities to determine the coincidence detection threshold of decoding cells (i.e., the decoding threshold) as first described in (Jeanson & White, 2013c). This enables a greater decoding accuracy than what was previously obtained from ‘collision removal’ (Jeanson & White, 2013b).

4.2.1 A Procedure for Delay Mapping

In order to decode spatiotemporal patterns into spatial patterns, delays from the memory network to a decoding network must be adjusted. The overall network layout is illustrated in Figure 4.1. To do so, we developed a method to generate specific delays for projections from the memory to the decoding network. Assuming we wish to decode 4 patterns in memory, we first generate a single random memory network to which four different input patterns are applied, namely A, B, C, and D. After determining the limit cycle period λ of the corresponding response patterns, we choose a sufficiently late offset spike-time st_o beyond the settling point of the attractor to mark the start of the region of analysis (in our case $st_o = 1400ts$). For each of the four distinct spike renewal processes we analyzed the temporal region from st_o to $st_o + \lambda$.

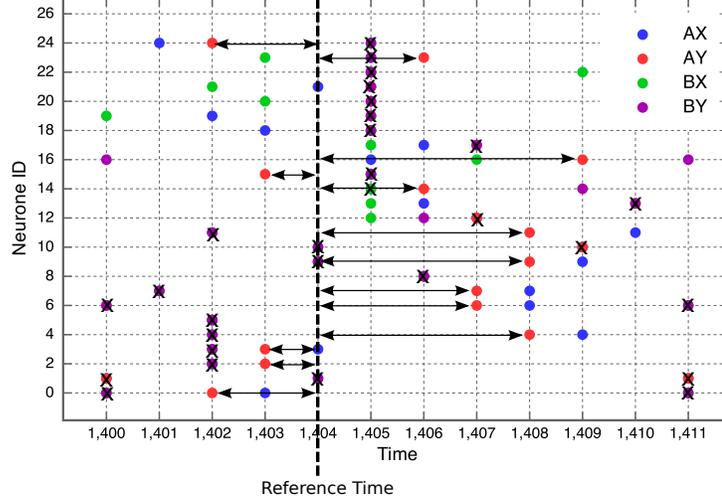


Figure 4.2: Collision removal and delay mapping. Collision removal is a rudimentary approach to reducing ambiguous decoding because only connections from cells that spike at the same time (here marked by an x) have efferent connections removed to prevent ambiguous spikes from reaching the decoding network. Despite this, many ambiguities can remain using this approach and it also leads to pattern weakening requiring selective ‘filling-in’.

We then selected an arbitrary reference time step ts_r within the period interval. Using ts_r we derived the set of spike-time ‘deltas’ which are the time step differences of the spike-time of a given cell from the reference ts_r . We denote this set of deltas for all cells of pattern ρ as Δ^ρ , and the specific delta of a given cell i in this set is referred to as Δ_i^ρ . Figure 4.2 illustrates the reference time step from which spike-time deltas (black arrows) were calculated for pattern A. Because (at most) a single spike-time st_i^ρ exists for a cell i for pattern ρ during period λ in our DCD memory network we can simply compute deltas as:

$$\Delta_i^\rho = ts_r - st_i^\rho \quad (4.1)$$

From this set of deltas we then created a set of connections for each cell in the memory network to pattern specific cells in the target decoding network. A delay d_{ij}^ρ from cell i to a target decoding cell j for pattern ρ that would project to a decoding target cell is derived from the list of deltas such that:

$$d_{ij}^\rho = \Delta_i^\rho + d_{base} \quad (4.2)$$

Where d_{base} is a base delay which will typically be the median delay in the set of connections for ρ . This is picked such that for cell refractory period R we respect

$R + \lambda/2 < d_{base}$. Here we chose $d_{base} = 20$. If we compare, for example, Δ_0^A to Δ_1^A in Figure 4.2 (the two deltas arrows at the top of the figure for spike pattern AY) we would obtain delays $d_{24j}^A = 2 + 20 = 22$ for cell 24 and obtain $d_{23j}^A = -2 + 20 = 18$ for cell 23. Hence target cells would have afferents from cell 24 with delay 22 ts and afferents from cell 23 with delay 18 ts. If this is done for all cell spike-times of pattern A then spikes would reach the target network synchronously leading to coincidence detection of A.

Here we outline the ‘DelayMapping’ algorithm which takes a list of point processes (spike times) for a limit-cycle memory pattern ρ :

$$spike_list^\rho = [[neurone_id, spike_time], [neurone_id, spike_time], \dots]$$

from which we want to determine the list of propagation delays d^ρ to decode pattern ρ and project to a set of target decoding neurones for pattern ρ which we denote as:

$$targets^\rho = [target_neurone_1, target_neurone_2, \dots]$$

Algorithm 4.2.1: DELAYMAPPING($spike_list^\rho, targets^\rho, st_o, \lambda, d_{base}$)

$ts_r \leftarrow random(st_o, st_o + \lambda)$

$deltas \leftarrow []$

for $i \leftarrow 0$ **to** $spike_list^\rho.length()$

do $\left\{ \begin{array}{l} st \leftarrow spike_list^\rho[i][1] \\ \text{if } st \geq st_o \text{ and } st \leq (st_o + \lambda) \\ \text{do } \left\{ \begin{array}{l} neurone_id \leftarrow spike_list^\rho[i][0] \\ deltas[neurone_id] \leftarrow ts_r - st \end{array} \right. \end{array} \right.$

$d \leftarrow []$

for $neurone_id \leftarrow deltas.keys()$

do $\left\{ \begin{array}{l} \text{for } target_id \leftarrow targets^\rho.keys() \\ \text{do } d[neurone_id][target_id] \leftarrow deltas[neurone_id] + d_{base} \end{array} \right.$

return (d)

Importantly, efferent connections from each memory network neurone can be created such that they target independent decoding cells. This allows a simple memory network to encode multiple input patterns such that these patterns can be recognized

by individual cells. Interestingly, this approach reflects quite closely similar forms of recurrent neural network coding where networks of cells are randomly generated and used as ‘reservoirs’ for storing arbitrary stimulus patterns. Indeed, original work by Buonomano & Merzenich on spiking networks with short-term plasticity was perhaps the first instance to demonstrate this unique method of random network generation coupled with decoding/classification procedure for complex function (Buonomano & Merzenich, 1995). This method has since been applied in other contexts including Liquid State Machines (Maass, Natschläger, & Markram, 2002) and Echo State Networks (Jaeger, 2002). What is clearly distinctive of the present approach, however, is that instead of synaptic plasticity, coincidence detection with propagation delays is maintained as the sole form for the generative spiking repertoire stored in the memory network (the ‘reservoir’) and the decoding is accomplished via delay mapping onto a decoding network (the ‘classifier’).

Importantly, however, while the delay mapping method is effective at finding each of the delay values required to decode a specific pattern, multiple patterns are likely to share neurones which have the same temporal offset with identical deltas leading to ambiguities. The most obvious form of ambiguity is the overlap of spike times as shown in Figure 4.2 by the crossed spike times. This leads to ambiguous features resulting in potential decoding errors.

4.2.2 Collision Removal

A rudimentary approach to resolving this problem of ambiguous decoding can be done by ignoring neurones with spike times that overlap for the same cells between responses by removing their connections to the decoding network. We call this method ‘collision removal’. Figure 4.2 shows spike times for the individual patterns and crossed spikes representing collision events between patterns. Collision removal reduces ambiguous decoding of the response patterns such that differentiability becomes once more possible. Unfortunately, collision removal could lead to patterns having too few distinct neurones to produce sufficient spikes for decoding. That is, if a target cell has a coincidence threshold of T but the number of distinguishable spiking cells from the memory pattern becomes lower than T due to collision removal, then the activation of this response pattern will never produce sufficient activity for it to be decoded. To remedy this, memory network neurones for which decoding efferent connections were severed can have connections corresponding to more weakly represented patterns to be ‘reinstated’. For example, from Figure 4.2 pattern BX has the fewest

remaining spikes (only three) for decoding, by reinstating connections for this pattern selectively we can enable differential decoding once more. This, however, requires additional steps of evaluation and manipulation that are impractical and possibly least likely for biological systems.

Initially, we instead opted for an even simpler solution which consisted of reinstating connections with delays with some arbitrary value within the normal delay range. This allowed for weakened memory patterns to be recognized if enough memory cells spikes matched the delay patterns. While this process did not guarantee successful response, we observed overall improvements in correct turning behaviour (Jeanson & White, 2013a). This was required here because we had not yet established a way to evaluate the amount of ambiguity between spatiotemporal pattern responses. Indeed, ambiguities can also exist for cells that do not have overlapping spikes between patterns but have some neurones with identical spike time offsets (we explore this in Section 4.4). Hence, for the first robot dynamic memory experiment, coincidence detection thresholds on decoding cells were set to be equal between response patterns. That is, the same T was used for cells decoding different patterns. Later, in Section 4.4, we establish a method to measure the amount of ambiguity between two patterns so that connections no longer need to be removed, but where the coincidence detection threshold of decoding cells for each pattern is adjusted independently.

4.3 T-maze

4.3.1 Sensorimotor Memory Task

In order to test whether sustained activity in DCD networks can lead to the temporary storage of stimuli and produce ‘decision making’ behaviours, we investigated the ability of DCD networks to provide memory and control to a Khepera robot to solve the well known T-maze task (Blodgett & McCutchan, 1947; Jeanson & White, 2013b, 2013a). In this task, an agent is initially placed at the bottom of a T shaped maze and learns to take a left or right turn when it reaches the intersection of the T depending on a cue stimulus presented before reaching the intersection. Hence the initial stimulus serves as a cue that must be memorized so that a delayed response is elicited once the agent arrives at the intersection of the maze. Figure 4.3 illustrates the conditions of this task. Interestingly, this task was successfully applied to evolved robot controllers implemented with continuous time recurrent neural networks (CTRNN) (Blynel & Floreano, 2003) as well as NEAT (Charles, Pinville, &

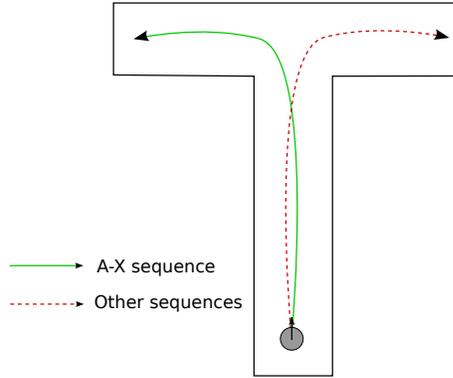


Figure 4.3: T-maze working memory paradigm.

Doncieux, 2012). However, it was noted that a confound may occur in the typical arrangement of this task. Ziemke and Thieme showed that agents may reach the correct side of the T-maze without memory when only a single stimulus is presented. This occurs when agents react immediately to the stimuli by following the wall on the side of the appropriate direction and simply continue following it at the intersection thus completing a successful turn (Ziemke & Thieme, 2002). As suggested by Charles et al. this can be remedied using a bi-stimuli paradigm based on Braver et al.’s AX-CPT working memory test (Braver, Cohen, & Servan-Schreiber, 1995). By having one stimulus presented early (the context stimulus A or B) and a second stimulus presented after a time delay (the cue stimulus X or Y) then agents can be instructed to turn in one direction, e.g., left, only for one specific combination of the two stimuli and turn in the opposite direction, e.g., right, for all other cases. To be successful, an agent must thus be able to store the context stimulus for the duration between presentation of the context and the cue.

In order for the appropriate response to take place in a DCD network a spatiotemporal pattern which will preferentially stimulate one motor effector over another must be activated within the period following the presentation of the cue. Crucially, however, this can only be achieved if a stimulus induced firing-chain is activated when the context stimulus is presented, that this firing-chain sustains, and that this firing-chain activity coordinates effectively with the cue stimulus so to activate the correct differential motor activation. To achieve this, the agent must not only store information but continue to advance up the maze in order to receive the cueing stimulus and engage in the appropriate trajectory to accomplish the task. Solving this T-maze task will thus not only demonstrate memory capacity on behalf of the agent but also an elementary form of adaptive behavioural control.

In the present implementation of the task a simple two wheeled robot with distance sensing was designed to autonomously move up the base of the T and make a left or right turn based on the presentation of a context stimulus A or B, shortly followed by a cue stimulus X or Y. When the combined stimuli AX are presented then a left turn should be taken while a right turn is expected in the three other cases. Distance sensing was used for wall avoidance. However, context and cue stimuli were directly presented to the agent’s DCD memory network. We describe the architecture of the agent and its controller in the following.

4.3.2 Sensory-Motor Network

The Enki robot simulation library was used to implement an e-puck robot with realistic noise, collision, and friction conditions (Magenat et al., 2007). The robot’s one-dimensional circular camera with a 120° field of view was divided into a left visual field (lvf) and right visual field (rvf). Each visual field was segmented in a numerical array of 30 values. Five averaged activation values per visual field were generated from 6 radially consecutive recordings. These activation values then served to dictate the spiking state of five sensory spiking cells per visual field. Sensory cells had refractory periods of 5 time steps (5ts) to increase sensitivity of the input while motor cells had refractory periods of 10ts. Spikes were triggered when the averaged normalized visual fields between 0.0 and 1.0 had values of 0.35 or less. Hence we used an inverse stimulation paradigm where the absence of obstacles stimulated the agent’s visual sensors. Spikes then propagated to the contra-lateral motor area which possessed 4 neurones for each wheel. These connections were fixed with many-to-many projections from each sensory cell to all motor cells in the contra-lateral sides using random delays between 10ts and 20ts. This relatively narrow delay range was chosen based on our prior finding that narrow delays could lead to more stable reactive behaviour as discussed in Section 3.2.4 (Jeanson & White, 2012). Contra-lateral projections were chosen based on the successful avoidance behaviour first suggested by Braitenberg (Braitenberg, 1984). Motor neurones behaved as coincidence detectors with threshold 2. Finally, activation of the four respective motor neurones were summed and total activation was limited to a maximum value of 1.5 to limit maximum motor speeds. The overall sensory-motor network (SM) can be seen on the right-hand side of Figure 4.4.

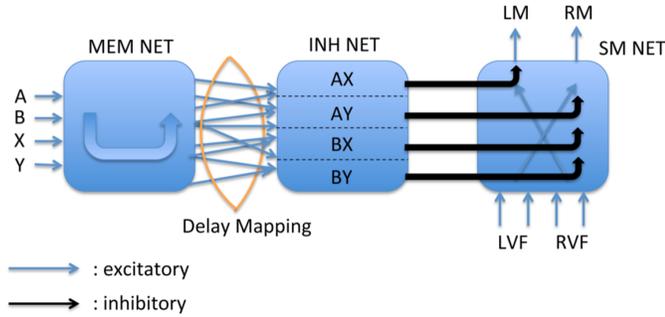


Figure 4.4: Illustrated network model with Memory, Inhibitory, and Sensorimotor networks.

4.3.3 Memory and Inhibitory Networks

The memory network (MEM) had 25 excitatory neurones which were fully connected and assigned a random uniform distribution of broad delays between 20ts and 40ts. Neurones in MEM had coincidence threshold 2 and refractory periods of 10ts. The MEM network was thus an instance of our Standard DCD network model. A context stimulus A or B was presented to the network after 450 time steps had elapsed from the onset of a trial followed by a cue stimulus X or Y after 350 time steps. Each stimulus type consisted of synchronously activating 6 individual cells for each pattern: cells 0 – 5 for A, 6 – 11 for B, 12 – 17 for X, and 18 – 23 for Y. No other control was applied to the network during trials.

Memory network activity was projected (Proj1) from its excitatory neurones to an inhibitory network (INH) which had the purpose of decoding the spatiotemporal pattern in MEM and inhibiting the appropriate cells in SM for correct maneuvering. INH contained 72 neurones with no internal connections but simply spiked when a sufficient number of coincident signals from MEM reached a cell in INH. This coincidence detection threshold was set to 8. These, in turn, projected their inhibitory signal so that 1/4 of cells in INH projected to the left motor neurones in SM while the other 3/4 projected to the right motor neurones in SM. Projections from INH to SM (Proj2) were created with a fixed many-to-many scheme and had delays set randomly in the range 10ts to 20ts. Relatively narrow delays were used so that the likelihood of coincidence on the SM motor cells would be reliable enough to disrupt the stable wall avoidance behaviour and induce the appropriate turning action. The overall network is illustrated in Figure 4.4.

Crucially, Proj1 connections were procedurally derived with respect to delay value and connectivity based on prior runs of the memory network with each of the four

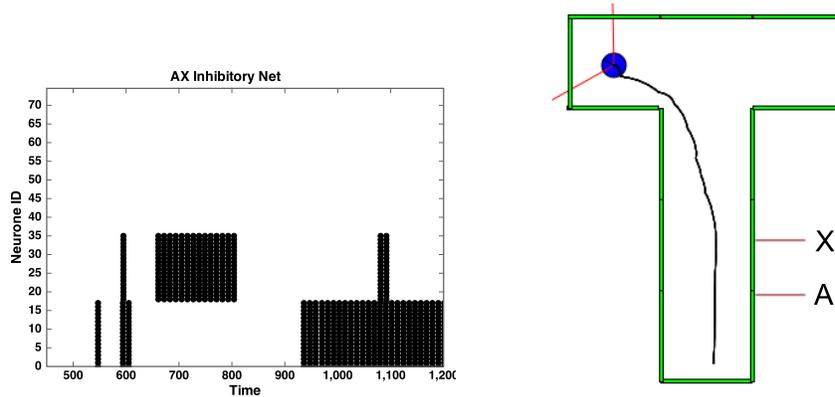


Figure 4.5: Left: spike raster plot of activity in the INH network during presentation of context stimulus A and cue stimulus X. Right: robot path in the AX stimulus condition. Lines originating from the robot mark the 120° field of view.

combined context and cue stimulus conditions (AX, AY, BX, BY). To accomplish this, we made use of the delay mapping method described by Algorithm 4.2.1 which resulted in Proj1 connections serving the role of transforming spatiotemporal memory patterns in MEM into spatial activation patterns in the INH network. By doing so, we thus manually handled the problem of having an agent learn the delay values for Proj1 connections¹.

4.3.4 Results

We generated 10 random memory networks (MEM) and derived their corresponding projection connections (Proj1) to the inhibitory network. One hundred trials were tested for a given MEM/Proj1 network pair. For each run, the initial orientation of the agent was randomly set so that it would face left or right with 45° variation on either side. We also randomized the narrow delay ranges for Proj2 and SM networks on each trial.

We first describe a sample trial for stimulus condition AX. We expected that the INH network would fire such that a subset of cells detecting pattern AX should be activated after the presentation of the context stimulus A and cue stimulus X have settled. Figure 4.5 shows in the left raster plot the activity of the INH network during a trial with context stimulus A presented at time step 450 and cue stimulus

¹While initial experiments were made to accomplish this using an evolutionary computational approach as in the reactive light-seeking task, no effective delay mapping was found using this method.

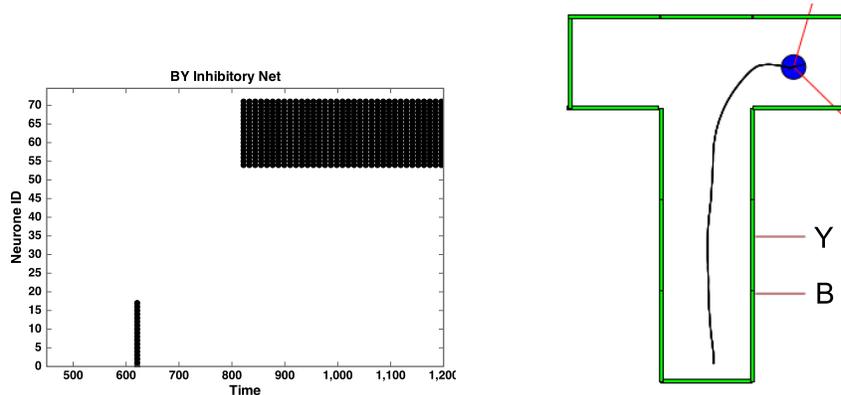


Figure 4.6: Left: spike raster plot of activity in the INH network during presentation of context stimulus B and cue stimulus Y. Right: robot path in the BY stimulus condition. Lines originating from the robot mark the 120° field of view.

X presented at time step 800. From the raster plot we notice that prior to time step 800 cells 18 – 35 corresponding to the recognition of pattern AY were temporally activated. This suggests that activation of memory A was sufficient to induce AY recognition. However, after the cue stimulus X was presented at 800ts and a short period of stabilization (approximately 130ts) cells 0 – 17 were activated which means memory AX was correctly decoded from MEM by the projections Proj1. This led to sufficient inhibition of the left motor neurones via projections Proj2 to the SM network. The graphic on the right hand side of Figure 4.5 shows the path taken by the robot when moving up the T-maze. In particular, we notice that soon after the presentation of the X stimulus a leftward turn is initiated. This led to a correct left turn reaching the top left corner of the T-maze. Similarly, correct right turns were obtained on most runs for context/cue stimuli pairs AY, BX, and BY. A sample run is shown for the BY stimulus pair in Figure 4.6.

We ran 100 trials for 10 random memory networks to determine the effectiveness of dynamic memory storage using the proposed delay mapping method. Between trials, random delays for Proj2 and SM connections were reset leading to slight variation for inhibition and sensorimotor control. For performance measurement, we independently counted the number of correct left turns when AX was presented (25% of times) and the number of correct right turns when all three other stimuli pairs were presented (75% of times). Simulation results showed that correct left turns were performed on 82% of AX trials and correct right turns were performed on 84% of AY, BX, and BY trials. Overall, the average success rate was 83.4% (SD 8.85). An agent with no

memory will have equal chance of turning left or right at the end of the maze which would lead to a 50% success rate. Similarly an agent biased towards left or right turns only will also have an overall success rate of 50%.

To reject the null-hypothesis that memory stimuli do not have an effect on correct turning behaviour, we ran experimental trials in a control condition by running 100 trials for 10 random robots with no presentation of the context or cue stimuli to the memory network. The lack of context and cue stimuli presentation led to a mean success of 50.6% (SD 3.32). The control condition simply enabled the SM network to drive the agent up the maze while avoiding walls. At the top of the maze the agent turned left or right due to sensitivity to the upper wall combined with noise. From these results we performed Welch’s t-test assuming independent variance between both conditions. We found a significant difference between the memory condition (M=83.4, SD=8.85) and the no-memory condition (M=50.6, SD=3.32) with $t(11.753) = 10.9$, $p < 0.001$. Hence, we may reject the null-hypothesis and, instead, support the finding that presentation of context and cue stimuli to the memory network was effective at enabling the agent to make correct turns.

Investigation of the individual performance of memory networks revealed that, despite collision removal and reinstating neurones, stimuli pairs could still be occasionally ambiguously decoded. This could lead to conflicting inhibition of the motor cells resulting in indecisive turning behaviour. This explained the less than perfect success rate with dynamic memory. Here we used a coincidence detection threshold of 8 for Proj1 connections for all generated memory networks. However, we believe greater memory pattern disambiguation could be achieved in future work by identifying sufficiently high individual coincidence detection thresholds for decoding neurones in INH with respect to each context and stimulus memory pattern pairs.

These results show that transmission delay coincidence detection neural networks can perform the role for dynamic memories whereby stimulus patterns are stored differentially and reliably without requiring internal structural/metabolic change. This suggests that *any* random network which possesses these characteristics could potentially be exploited by adjacent networks for dynamic memory. In addition, the short settling time into periodic attractors suggests that this mechanism could subserve fast delay response integration, rapid scene categorization, and other neural functions requiring millisecond timescale exchanges. Interestingly, this work also demonstrates that not only single stimuli conditions can be encoded but that combined stimuli (e.g., contexts and cues) can also be differentiated by a decoding network for control. Despite these positive results, we aimed to improved the method by which ambiguity

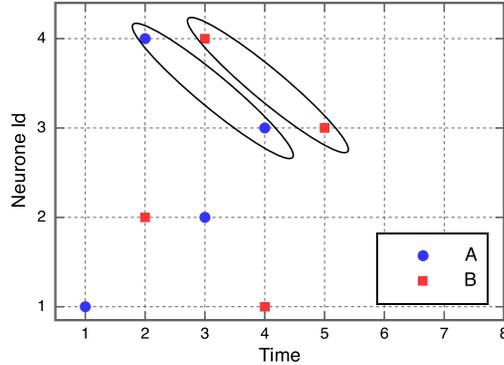


Figure 4.7: Two spike patterns with a single ambiguity between both patterns. The circled spike pairs lead to two inner-deltas that are equal for both patterns, hence the ambiguity.

resolution could be accomplished without making use of the collision removal method. Furthermore, we wished to explore the capacity for these networks to encode a greater number of patterns within a memory networks of the same size.

4.4 Ambiguity

4.4.1 Assigning Decoding Thresholds

In Section 4.2, we have shown that removing cells with colliding spike-times can be effective for reducing ambiguity. This method, however, can lead to a rapidly decreasing number of available neurones from which to decode patterns. Furthermore, this method fails to capture all ambiguities between patterns because only spikes with a temporal offset of 0 time steps are detected. While we showed that reinstating spikes could help mitigate the effects of collision removal, we aim to develop a method requiring fewer complex steps and with more biological plausibility. To do so, we propose instead to determine the exact decoding coincidence detection threshold on target cells with respect to each encoded pattern. This approach should allow neural significance to be preserved while making disambiguation possible. After describing the fundamental notion behind spike timed ambiguities, we demonstrate in the next section how to determine the minimum coincidence threshold of target cells.

While delay mapping enables the reliable detection of a spatiotemporal spike pattern via the synchronous spatial activation of target cells, two spatiotemporal patterns (say A and B) may possess a number of identical subsets of cells firing with identical

spike-time offsets. We denote the set of spike time differences between all spike pairs within the same pattern the set of ‘inner-deltas’ of that pattern. Figure 4.7 shows two spike patterns where one of the 6 ($4*3/2$) possible inner-deltas for pattern A matches an inner-delta for pattern B . This leads to a single *ambiguity* out of a potential six between these two patterns. Because only two spikes in each pattern are being considered, we denote this as an ambiguity of degree 2 or $\Lambda = 2$. However, it is also possible for a third, fourth, or indeed N spikes to form ambiguities of degree as high as $\Lambda = N$. We denote the maximum ambiguity between two response patterns as Λ_{max} . Importantly, if at least one ambiguity of degree x exists between two patterns and $x = \Lambda_{max}$, then a decoding coincidence threshold of $T = x + 1$ will have to be used. Interestingly, two ambiguities of degree 2 that share a common spike will induce the presence of a single ambiguity of degree 3, thus the maximum degree of ambiguity between both pattern is $\Lambda_{max} = 3$. However, if two ambiguities of degree 2 are independent, i.e., 4 spikes in total, then the maximum degree of ambiguity remains at $\Lambda_{max} = 2$. Hence, while in the former case a decoding threshold $T = 4$ will be necessary, a decoding threshold of $T = 3$ will be sufficient in the latter case.

4.4.2 Defining Ambiguities

For a spatiotemporal pattern ρ we can find all spike time differences between all cells that are spiking during a period λ . The spike time difference between two cells i and j , i.e., the ‘inner-delta’ between i and j , for pattern ρ is denoted as δ_{ij}^ρ . The total number of inner-deltas is equal to all edges in a complete graph or $N * (N - 1)/2$ for N neurones participating in the pattern. Two patterns, A and B , can have a number of matching δ 's. Such that $\delta_{ij}^A = \delta_{ij}^B$ for some cells i and j . We call these ambiguities $\alpha_{ij}^{A,B}$. Returning to Figure 4.7, the single ambiguity between patterns A and B can be expressed as $\alpha_{3,4}^{A,B}$. Importantly, ambiguities between large groups of cells can occur as well such that $\alpha_{i,j,\dots,n}^{AB}$ exist. As mentioned above, if two ambiguities α_{ij} and α_{jk} exist then there will necessarily be an α_{ik} present as well. Hence, two ambiguities of degree $\Lambda = 2$ that share a cell implies that an ambiguity of degree $\Lambda = 3$ exists between both patterns. This is a useful value to determine since it will indicate that a decoding threshold of at least $T = 4$ should be used to uniquely identify both patterns from each other (Jeanson & White, 2013c).

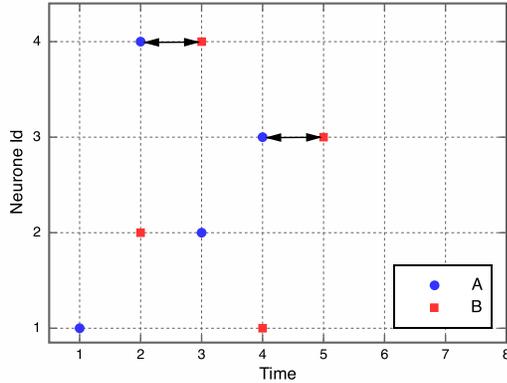


Figure 4.8: Second method for finding the maximum ambiguity between two patterns. Between pattern spike-time differences are calculated for each of the N neurones. Here, the largest set has two differences (cells 4 and 3) of time step $+1$, followed by a set with one difference (cell 2) of -1 and a set with one difference (cell 1) of $+3$. Hence a maximum degree of ambiguity of two corresponding to the size of the largest set.

4.4.3 Finding Ambiguities

Here, we present two methods to find the largest degree of ambiguity Λ_{max} between two spatiotemporal responses.

In the first method, we can determine the number of pairs that share a common neurone so to find Λ_{max} in order to derive the minimum decoding threshold by counting the number of paired inner-deltas that are ambiguous between any two patterns. To do so, we first evaluate all $\binom{N}{2}$ inner-deltas for each of the patterns. These sets are then compared between both patterns to find how many are ambiguous. Finally, we count among the ambiguous pairs, the subset which shares the greatest number of cells to determine the highest degree of ambiguity. This results in a complexity order of $O(3\binom{N}{2})$. This rather computationally costly evaluation can be substantially reduced by adopting the second method.

The second method first evaluates the spike-time difference between both response patterns for each neurone i among N as shown in Figure 4.8. We then find the largest set of equal spike-time differences Ω_{max} . This Ω_{max} corresponds to the largest degree of ambiguity in the network Λ_{max} . Because the maximum number of Ω sets is N , we can find the largest degree of ambiguity in N steps, hence a complexity of $O(N)$.

4.5 F-Maze

In Section 4.3 we showed that delay mapping can be applied effectively for the derivation of delays from dynamic memory networks in a robot task requiring agents to navigate a single intersection T-maze based on external stimulus cues (Jeanson & White, 2013b). DCD network simulations up to here, however, have hinted at a potentially rich set of reliable and differential responses with respect to independent stimuli. Hence, it would be interesting to determine how DCD networks can scale up with respect to task complexity and memory requirement while keeping a fixed size of 25 neurones for the memory network. In addition to this, we will explore in Chapter 5 their tolerance to noise and memory capacity, and in Chapter 6 mathematical tools to formalize their spiking characteristics. Here, we expand on the T-maze by adding an additional intersection to the maze, thus becoming what we denote as the ‘F-maze’ as seen in Figure 4.9.

As in the T-maze task, the agent must autonomously travel up the bottom of the first T by avoiding walls, make the first correct turn based memory network stimuli, then continue to the end of the second T to make a second turn based on a separate set of context and cue stimuli. Simulations were performed using the same platform as in the T-maze experiment described in Section 4.3. The sensorimotor network (SM) was composed of DCD neurones that enabled the agent to move up the maze corridors while avoiding walls. These motor neurones received external afferents from a population of inhibitory cells (INH). These inhibitory neurones served to inhibit left motor neurones to provoke a left-hand turn and inhibit right motor neurones to perform a right-hand turn. The inhibitory neurones decoded the spatiotemporal spiking patterns produced by DCD cells in the memory network (MEM) as depicted in Figure 4.4.

4.5.1 Procedure

For the F-maze task, eight possible stimuli pairs could be produced. After 450ts either a context stimulus A or B was presented to the agent, after 350ts a cue stimulus X or Y was presented. A correct response by the agent was noted if a left-turn was performed at the end of the first T in case AY, BX, or BY were presented, and a right-turn if only AX was presented. In the case of a right-turn (AX) the spiking activity in the agent’s memory network was cleared to prepare for the subsequent memory storage. This was done to test the independent storage of the first context/cue pair from the second context/cue pair. After this, a second context stimulus C or D was

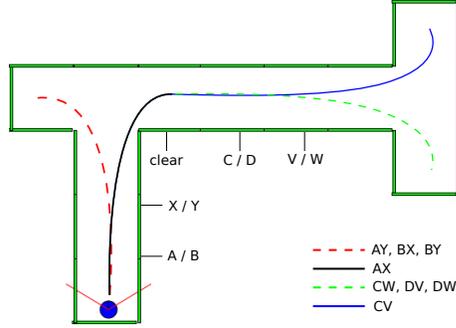


Figure 4.9: F-maze.

presented to the memory network. Once another 350ts had passed the second cue stimulus V or W was presented, thus producing CV, CW, DV, or DW spatiotemporal patterns in the memory network. If CV was presented then a left-hand turn was expected, while a right-hand turn was expected for all other combinations. Figure 4.9 illustrates the possible stimuli related paths that would lead to a successful turn. To decode these eight spatiotemporal patterns we used 144 cells in the inhibitory network amongst which 8 sets of 18 cells were divided to recognize the individual spatiotemporal patterns AX, AY, BX, BY, CV, CW, DV, and DW respectively. The probabilities of each turning scenario are 75% for a left-hand turn at the first intersection, 18.75% for making a right-hand turn at the second intersection, and 6.25% for a left-hand turn at the second intersection. These unequal probabilities serve to show that the agent responds selectively to the combined memory stimuli AX for the first part of the maze, and CV for the second part of the maze. We conducted 100 runs for each of the 10 randomly generated agent memory networks from which decoding delays were derived using the delay mapping described in Section 4.2 and decoding threshold derivation methods described in Section 4.4.

4.5.2 Results

Preliminary stimulation of the memory network for the 8 specific input combinations enabled us to apply the delay mapping procedure to determine propagation delays for connections from the memory network to the inhibitory network (Proj1 connections). After delay mapping was accomplished for each of the 10 generated memory networks, ambiguities were computed between each of the eight paired context/cue network responses to establish the highest degree of ambiguity. This resulted in the following Λ_{max} for networks 1 to 10 respectively: 11, 15, 14, 14, 12, 14, 12, 10, 9, 15. Hence,

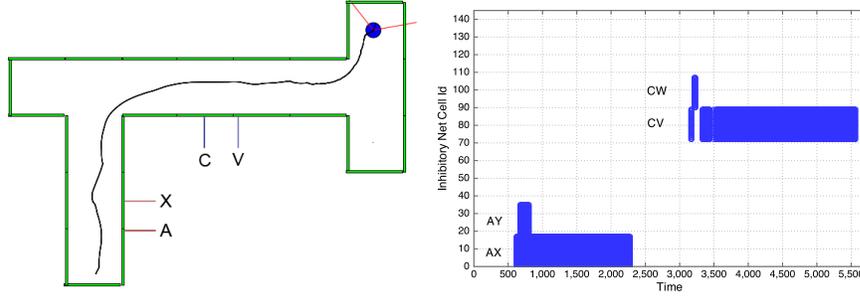


Figure 4.10: Left: Robot path on a trial in the F-maze with stimuli AX-CV. Right: Inhibitory network activity.

decoding thresholds T for inhibitory targets were set to the corresponding $\Lambda_{max} + 1$.

A sample run with stimulus combination AX-CV is shown in Figure 4.10. From this figure we see how the agent initially moved up the base of the maze while avoiding walls. The context stimulus A was presented at time step 450. After a transient period we notice the decoding from cells 0-17 dedicated to AX began in INH at time step 618. This suggests that pattern A was sufficient to induce an AX response in INH. This was briefly followed by simultaneous decoding of AY for cells 18-35 in INH. This induced the agent to swerve slightly to the left as a result of co-activation of the patterns. This ambiguous decoding can be expected if enough features in A are sufficient to be recognized from the decoding delays for AX and AY. However we noticed that as soon as stimulus X was presented at 800ts only cells 0-17 in INH effectively decoded pattern AX in MEM. This enabled the agent to inhibit right motor neurones and perform a correct right turn at the first intersection. After passing the first right hand turn the MEM network had activity reset to accommodate for the subsequent stimulus pair. This explains the lack of decoding between time steps 2278 and 3329. After presenting context stimulus C at time step 2879, CV was correctly not decoded by INH which shows that C alone does not contain sufficient spatiotemporal features to produce a turning behaviour. However, after cue stimulus V was presented at 3229ts CV was recognized after a 100 time step transient period and became recognizable to INH at time step 3329. This initiated a correct left-hand turn leading to successful behaviour at the second intersection of the maze. While CV and CW were momentarily decoded due to the second transient period response when V was presented, this did not affect the correct left turn behaviour of the robot at the second intersection.

Similarly a sample run with stimulus combination AX-DW is shown in Figure 4.11. The behaviour of the robot is nearly identical to the previous sample run with

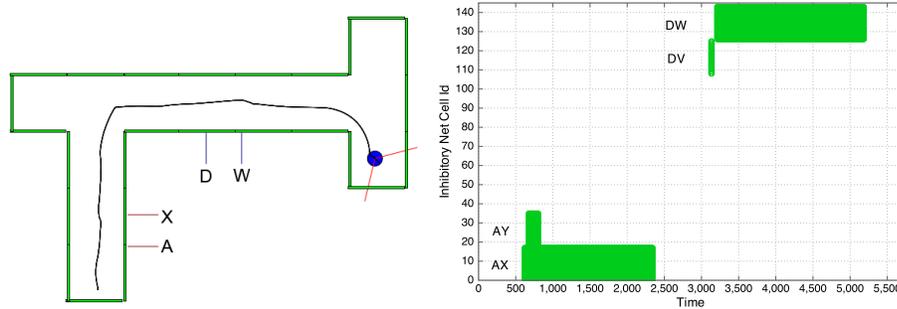


Figure 4.11: Left: Robot path on a trial in the F-maze with stimuli AX-DW. Right: Inhibitory network activity.

respect to the first intersection. Context stimulus D was presented at 2887ts. After a transient period during which DV was momentarily decoded, the onset of the cue stimulus W at 3237ts pattern DW was consistently decoded to perform a correct right hand turn by inhibiting right motor neurones of the SM network. While a number of these preliminary experiments demonstrated successful behaviour, we undertook a more statistically meaningful analysis by running 100 experiments with memory and without memory control.

To evaluate behavioural control performance, tested 10 random agent three-network controllers for which delay mapping and decoding thresholds for the stimuli pairs were resolved. Each agent was tested on 100 trials. In these trials context and cue stimuli were presented to the memory network. Individual scores for each turn were 88% for FL, 92% for SL, and 80% for SR leading to an average success rate of 86.6% (SD 3.94). To ensure that memory had no effect on correct turns, we conducted control runs during which no turning stimuli were presented to the memory network. Again, 10 agents were tested 100 times with randomly generated networks. If memory had no effect then the agent should have reached the first left-hand corner (FL) 50% of the time, the second left-hand corner (SL) 25% of the time, and the second right hand corner (SR) 25% of the time, resulting in an overall mean success rate of 50%. Results from the control in fact revealed an overall mean success of 42.9% (SD 4.38). This seemingly lower than chance result is due to some agents failing to reach any of the corners at all.

To confirm that we may reject the null-hypothesis that memory had no effect on effectively controlling the robot, we performed Welch's t-test assuming independent variance between both memory and no-memory conditions. We found a significant difference between the memory condition (M=86.6, SD=3.94) and the no-memory condition (M=42.9, SD=4.38) with $t(17) = 22.86$, $p < 0.001$. Hence, we reject the

null-hypothesis and confirm that presentation of context and cue stimuli to the memory network had a significant effect at enabling the agents to make correct turns in the F-maze.

Errors occurred due to noise interference with the balanced inhibition of the sensorimotor network to perform adequate turns. Despite this, these results support the view that this dynamic memory network model can be effective at controlling agents in a problem of greater scale than the T-maze task alone. Furthermore, this was accomplished without collision removal but by identifying a sufficiently high coincidence detection threshold for decoding cells in INH.

4.6 Summary

The F-maze results confirm that dynamic memories in DCD networks can effectively scale up to harder memory tasks. In particular, this was accomplished without having to increase the size of the memory network. These results demonstrate that DCD networks can perform the role for dynamic memories whereby stimulus patterns are stored differentially and reliably without requiring internal structural/metabolic change. Importantly, this suggests that any random network which possesses these characteristics could potentially be exploited by adjacent networks for dynamic memory purposes. In addition, the short average transient periods over a few hundred time steps (in the order of 10's of milliseconds) between memory stimulation and stabilization into a periodic attractor suggests that this mechanism could account for fast delay response integration, rapid scene categorization, and other memory functions requiring millisecond timescale exchanges. Crucially, the dynamic memory mechanism presented here requires that an effective delay mapping exists in order to decode the spatiotemporal memories into either other spatiotemporal patterns, temporal patterns, or as we have shown here, spatial patterns. While we offer a procedural method to determine these, biological mechanisms such as evolution or development may be effective at selecting adequate projection connections and delays to make use of the stored memory patterns. Synaptic pruning of inadequate connections could promote adequately tuned delayed connections, degrees of myelination could vary during development to adjust transmission speeds, or even STDP could lead to the reinforcement of connections with adequate delays. However, we do not propose or imply biological plausibility here. Such investigations would form part of future research. In general, we believe dynamic memories of the sort presented here should be especially well suited for providing a large yet fixed set of stimulus

response primitives which can be coordinated rapidly for further processing.

A weakness to the present approach can be found in the presentation of the context and cue stimuli at fixed temporal intervals. When fixed spatial positions were used for the memory stimuli we found that the firing-chains resulting from the convolution could often lead to non-recognizable memory patterns by the decoding projections for some offset to the period λ . While preliminary studies have shown statistically significant performance with fixed spatial stimuli, we believe larger networks should enable greater statistical options for the decoding by the projection network. Alternatively, a biological mechanism such as a phase carrying spike waves could justify the presentation of sensory stimuli to the memory network such that presentation occurs at an optimum timing for decoding. Sensory coding via phase carrying waves have indeed been theoretically described and observed biologically (Gerstner & Kistler, 2002; Nadasdy, 2012). Again, such investigations could form future research into more advance DCD networks.

Overall, the experiments presented here demonstrated that unique firing-chain responses are produced for independent stimuli, and that the presentation of subsequent stimuli enabled the convolution of the joint inputs in a dynamic fashion. While we demonstrated that 8 distinct patterns based on 8 primitive inputs could be stored in a network of 25 neurones, we have yet to establish a precise estimate of the total number of distinct firing-chains which can be produced by networks of a given size and their tolerance to noise. Such an estimate could provide a much better appreciation for the coding capacity and noise tolerance of these networks. This memory capacity and tolerance to noise are investigated in Chapter 5. Beyond memory, however, we anticipate that delay coincidence detection networks can likely be organized to produce a wide range of additional dynamics under various connectivity and parametric conditions with potentially important consequences for our understanding of cognition and behavioural adaptation via neural processes.

Chapter 5

Dynamic Memory: Noise Tolerance and Memory Capacity

In the previous chapter we were able to show that DCD networks could serve the function of dynamic memories in robot maze navigation tasks. In particular, we showed that three networks of DCD neurones could be coordinated so as to enable a robot to perform wall avoidance, stimulus storage, and decoding for elementary decision-making turning behaviours. While the T-maze task served to exhibit the fundamental capacity of DCD networks within a three-network architecture, we also showed that the identical memory network architecture could also be used in a scaled-up version of the task with four additional stimulus cues in the F-maze task. These experiments, however, did not test the extent to which stimuli could be encoded while delivering reliable and differentiable responses. To gain a better appreciation for the limits in terms of useful encoding that these networks can support, we explore in this chapter their ability to cope with interference with respect to spiking but also their capacity limit to encode effective memories. By doing so, we should gain a better appreciation for the practical application of these networks in neural coding for cognition.

5.1 Noise Tolerance

Noise within a system occurs when this system is subjected to events that are independent of its normal functioning. Here, we investigate the degree to which noise affects the spatiotemporal characteristics of the Standard DCD network model initially described in Section 3.3.1. First, we review how using ambiguities (c.f. Section

4.4) can be found in spatiotemporal patterns of spikes. We then introduce the notion of degrees-of-freedom with respect to ambiguities. Subsequently, we use this degrees-of-freedom notion to establish lower and upper bound estimates for the number of noisy cells that can be present in spatiotemporal patterns of spiking neurones with ambiguities. We then explore the condition where decoding does not require all available neurones, but only the subset sufficient to provoke coincidence detection on decoding neurones. However, because none of the estimates take into account the correlated nature of firing within these networks, we then investigate empirically the effect of noise via spike removal and misfiring on DCDNs to quantitatively determine how many neurones can undergo interference before decodability becomes impossible.

5.1.1 Degrees of Freedom

A critical feature for any memory system is the ability to provide sufficient differentiation between encoded patterns. Before estimating of how DCDNs behave under noisy conditions we isolate the spatiotemporal features based on a subset of spiking neurones that can be used to uniquely identify patterns without having to take into account the remaining noisy or ambiguous neurones. To do so we establish the upper and lower-bounds for a network’s tolerance to ambiguities, which we described in Section 4.4, by computing what we call the maximum and minimum *degrees-of-freedom* of a network – respectively df_{max} and df_{min} based on the number of common cells shared between ambiguous relations. This degree-of-freedom corresponds to all *paired relative spike times* (i.e., inner-deltas) that are not ambiguous between any two patterns. In other words, the greater the degree-of-freedom the greater the potential for differentiation between patterns. An optimistic bound exists because three ambiguities found could correspond to no more than 3 cells. For instance, if two cells c_1 and c_3 form two ambiguities with a third cell c_2 , then they are ambiguous with each other. However, in the pessimistic case, three ambiguities could exist among 3 independent pairs of cells.

To determine df_{max} we take the smallest set of h neurones that can hold A ambiguities of degree 2. We call this set a ‘hyper-node’ since the reference times of each spike for cells in a hyper-node are correlated between patterns involved in the ambiguity. For example, when recalling Figure 4.7, pattern ρ_B has a single ambiguity formed by cells 3 and 4. Cells 3 and 4 can thus be effectively replaced by a single hyper-node of size $h = 2$ with a unique spike time difference (inner-delta) with other neurones in the network. More specifically, the inner-deltas between cells 1 and 3

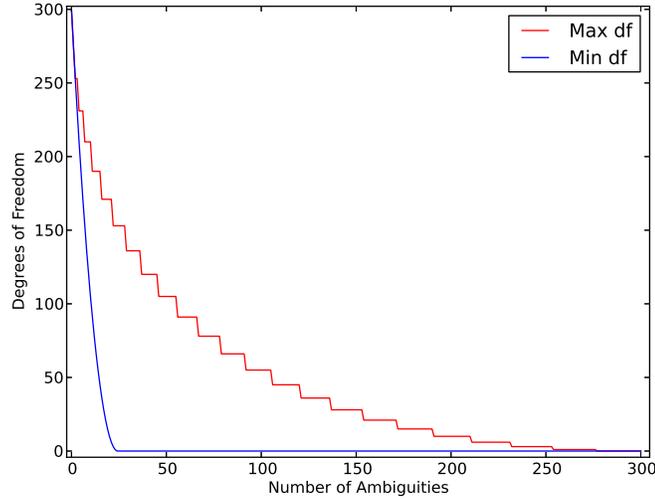


Figure 5.1: Maximum (optimistic) and minimum (pessimistic) bounds for the degrees-of-freedom of a DCD network of 25 neurones with a single spike per cell for period λ . As the number of ambiguities of degree 2 increases the degree-of-freedom (df) decreases in a log like fashion.

(δ_{13}) and between 1 and 4 (δ_{14}) provide the same information since 3 and 4 have an ambiguous inner-delta. Thus pattern ρ_A could be distinguished from pattern ρ_B using δ_{13} or δ_{14} , however these two deltas do not provide independently new information due to the ambiguity in which cells 3 and 4 participate. Hence all cells (N) not participating in an ambiguity (h) along with this hyper-node (1) form the new set of size $N_h = N - h + 1$ of uniquely detectable spike-times to decode, where N is the total number of neurones in the network. In pattern ρ_A and ρ_B , the hyper-node contains 2 neurones (3 and 4) and the remaining 2 neurones (1 and 2) together with this hyper-node give $N_h = 3$. To determine the maximum degrees-of-freedom given N and A , we need the smallest set of cells h that form a connected graph with A edges using the following equation:

$$h = \left\lceil \frac{1 + \sqrt{1 + 8A}}{2} \right\rceil \quad (5.1)$$

Where h counts the number of cells participating in a hyper-node. Knowing h we thus derive $N_h = N - h + 1$ (the number of effective neurones remaining plus the hyper-node) to evaluate df_{max} according to:

$$df_{max} = \frac{N_h(N_h - 1)}{2} \quad (5.2)$$

Determining df_{max} thus gives us the most optimistic decoding efficiency of the network when the number of ambiguities of degree 2 present in the network is known. It is optimistic, because it tells us the fewest possible number of ambiguities (and therefore neurones) that should be ignored in order to distinguish between patterns. However, the worst case scenario arises when a maximum number of cells have correlated spike times, thus increasing the maximum degree of ambiguities in the network Λ_{max} . This leads to a minimum degree-of-freedom df_{min} . Knowing the number of cells N and the number A of ambiguities of degree 2, we know that at worst an ambiguity exists between $N - A$ cells, hence at worst $N_h = N - A$. Thus we determine df_{min} with the following:

$$df_{min} = \frac{(N - A)(N - A - 1)}{2} \quad (5.3)$$

Using these upper and lower bounds for degrees-of-freedom of a network we can plot this range for our network of 25 cells knowing that the maximum limit for ambiguities of degree 2 is $N(N - 1)/2 = 300$. This df range is shown in Figure 5.1. From this Figure we can determine the region between df_{max} and df_{min} in which decoding will be possible given a number of found ambiguities. For instance, in the optimistic case, 24 ambiguities (of degree 2) will lead to a df_{max} of 153. This means 153 sets of paired spike times are available to discriminate patterns.

At first glance, degrees-of-freedom could be used to find how many noisy neurones can be tolerated by a network when we know the number of ambiguities that are present among different response patterns for different inputs. For instance, if 24 ambiguities (of degree 2) are found between patterns then a df_{max} of 153 is available to a network which means that 18 neurones are available to distinguish between patterns. This leaves a total of 7 neurones that could behave erratically. However, this approach assumes that all 18 neurones are necessary to distinguish between patterns. In reality, a subset of spikes could be used to distinguish between patterns if the decoding neurones have sufficiently low coincidence detection thresholds.

Taking this into account suggests that additional noisy cells could be tolerated by the network while preserving decodability. We develop this more sophisticated view of sub-pattern decoding in the following section.

5.1.2 Noise Tolerance Estimation

If a spatiotemporal pattern is only expected to be decoded by low coincidence threshold cells (where T is the coincidence detection threshold) then many sub-patterns of

size T can be selected so as to uniquely identify the initial input. Indeed, assuming spiking independence and no ambiguities between network responses, $N - T$ neurones can be noisy with a single subset of nodes of size T being sufficient for differentiation. This suggests that while all other cells may be noisy, decoding via delay mapping is possible from those few spike-times. The case becomes less straightforward when spike times are correlated over connectivity and time because noisy activity will likely propagate to other cells. We first present a solution to the noise estimate when only ambiguities between patterns are considered and explore correlated firing in the next section via empirical simulation experiments.

5.1.2.1 Noise with Ambiguity

If X neurones are noisy in a network of N cells then the remaining cells of the pattern $N - X$ that are stable can be used to decode the pattern. Of course, for this to be possible the coincidence threshold of target decoding cells will have to be less or equal to $N - X$. However, if this pattern must be differentiated from other possible input response patterns for the same network then ambiguities may arise. Here we present a method to estimate a priori the necessary number of non-noisy cells that must be available to decode a pattern when the number of ambiguities between this pattern and others is known.

We denote S as the number of differentiable sub-patterns in pattern ρ that are detectable given the number of cells N , the number of ambiguities A of degree 2 found, and the decoding threshold T . Here the decoding threshold T dictates the minimum size of the sub-pattern. Hence a value of $T = 2$ implies that we are interested in sub-patterns of 2 spike times in two different cells. The maximum and minimum number of sub-patterns S with at least two available spike times for decoding can be obtained by summing the number of combinations $\binom{N}{t}$ (i.e., N choose t) of spike times for sets of size t such that:

$$S = \sum_{t=T}^{N_h} \binom{N_h}{t} \quad (5.4)$$

Where $N_h = N - h + 1$ when computing the maximum number of sub-patterns S_{max} (optimistic) and $N_h = N - A$ in the case of minimum sub-patterns S_{min} (pessimistic). Hence, similarly to the degree-of-freedom, we obtain an upper and lower bound estimate for the number sub-patterns S in a network as shown in Figure 5.2. From this Figure we notice that while over 10^7 sub-patterns are available from net-

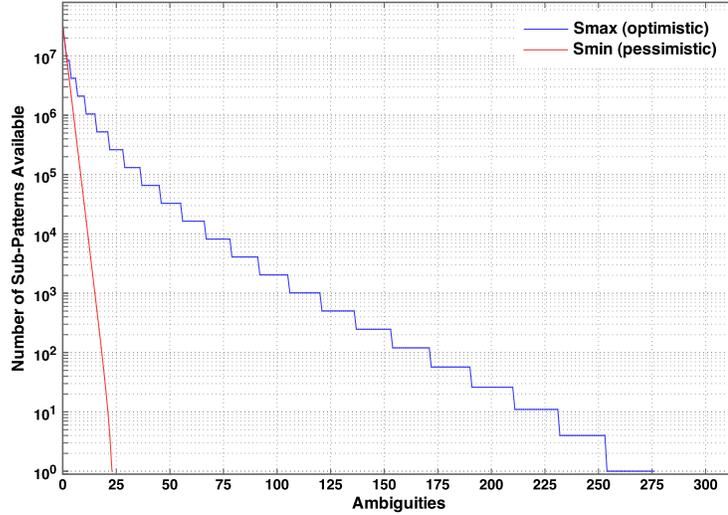


Figure 5.2: Optimistic and pessimistic number of available sub-patterns (S) with at least two spike times in a network of 25 cells as ambiguities increase.

work responses without ambiguity in both optimistic and pessimistic conditions, this availability decreases in a logarithmic stepwise fashion as the number of ambiguities (of degree 2) increases linearly. In the pessimistic case we find that availability of sub-patterns reaches 0 as soon as 24 ambiguities are found.

Applying this to noise, S can be used to determine the number of neurones X that can be disrupted such that at least one sub-pattern remains available to uniquely identify the complete pattern from others. Assuming the number of ambiguities (of degree 2) decreases linearly as the number of cells N is reduced, we want to find the number of spike-times that will allow for at least one sub-pattern given the proportionally reduced number of ambiguities $A_{partial}$. To do so in the optimistic scenario (most noise tolerance) we must find the subset of cells $N_{partial}$ such that $A_{partial} = \frac{N_{partial}}{(N/A)}$ which satisfies $S = 1$, where S is computed with $N_{partial}$, $A_{partial}$, and T . From this we obtain either the maximum or minimum noise tolerance – depending on how h is computed – such that the number of tolerable noisy cells $X = N - N_{partial}$. Figure 5.3 illustrates the optimistic noise tolerance for a network with 25 neurones and with a required decoding threshold $T = 2$. It is important to note, however, that noisy spike-times should not coincide with spike-times of alternate patterns, or ambiguous decoding will arise.

Conveniently, this sub-pattern counting method is applicable to any discrete spatiotemporal events for which we wish to determine upper and lower-bound noise tolerance estimates and, of course, for which we have point process data. While

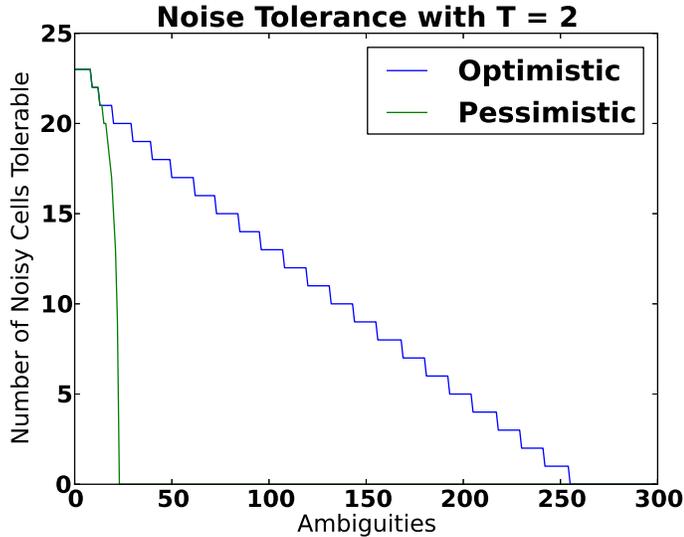


Figure 5.3: Optimistic and pessimistic noise tolerance in a network of 25 cells with fixed decoding threshold $T = 2$ as ambiguities increase.

computing the number of sub-patterns was accomplished here for spiking events that occur on regular intervals, it would not be difficult to extend this method to other spatiotemporal events with limit-cycle characteristics. What this method lacks, however, is the ability to provide specific estimates with respect to the noise tolerance of patterns when events are not strictly independent. In particular, the neural and network characteristics of DCDNs with high degrees of connectivity, fixed delay distributions (e.g., 20 to 40 time steps) and homogeneous decoding thresholds (e.g., $T = 2$), suggest that sustained spiking responses are driven by a high degree of correlated spikes in the network. In relation to this, Derrida and colleagues have shown that a degree of incident connections in attractor neural networks of at least $\log N$ per cell will lead to spike correlation (Derrida et al., 1987; Nützel et al., 1994). However, their network model was not concerned with adequate timing for coincidence detection. Hence, in DCDNs if I incidents are sufficient to produce a post-synaptic spike only I/T are in fact effective, hence correlation in DCDNs begins to occur for $I/T \geq \log N$. In a DCDN of 25 cells that are fully connected and a coincidence detection threshold of 2 suggests that a substantial degree of correlation is present since $24/2 > \log 25$. While we have argued that a sufficiently high degree of incidence is necessary for sustained firing, we must determine the impact of noise on this correlated firing for a more formal understanding of noise tolerance by DCD networks.

5.1.3 Noise in Correlated Spiking Activity

Noise during neural activity is to be expected for reasons of global or local interference. In spatiotemporal spiking dynamics, noise can take the form of additive spiking, subtractive spiking, and misfiring where a neurone reaches its threshold at an uncharacteristic time for a given pattern. To determine how tolerant to noise DCDNs are we explore empirically both subtractive spiking (by the removal of spikes) and misfiring. We later explain why spike addition is not relevant to the present model context.

To measure noise for a given network we investigated the amount of ambiguity which exists between repeated trials under identical input conditions (synchronous input to 6 cells) but where spike-removal or misfiring was applied to an increasing number of cells after the pattern settled in a stable limit-cycle attractor. Interestingly, we can make use of ambiguity analysis to quantify the level of noise tolerable by determining the amount of ambiguity between a noisy and a non-noisy response pattern for the same network under identical initial input. Indeed, if a response pattern is compared to itself then the maximum number of ambiguities will be reached and its *maximum degree* of ambiguity Λ_{max} will be equal to the number of neurones in the pattern. However, if noise results in one neurone changing its spike timing (shifted or removed) then the maximum degree of ambiguity between the noisy and non-noisy pattern will be $\Lambda_{max} = N - 1$, signifying that $N-1$ neurones can still be used to decode the spatiotemporal pattern via coincidence detection.

In the following, we empirically investigate noise tolerance via spike-removal and misfiring in both prolonged noisy conditions and short term, temporary noisy conditions in Standard DCD networks (c.f. Section 3.3.1). For each of these conditions we begin with a preliminary investigation of sample runs and finally quantify noise tolerance via population sampling in multiple networks.

5.1.3.1 Noise via Prolonged Spike-Removal

Noise via spike-removal can either take place as prolonged inhibition of a subset of neurones or as temporary inhibition of the subset. We first present DCDN tolerance to prolonged inhibition followed by temporary inhibition.

To investigate prolonged inhibition of spiking (spike-removal), thirty different networks were generated. Each network was presented with an identical single pulse input to neurones 0 to 5 synchronously (6 total) at time step 0. An input to at least 24% of neurones (here 6 out of 25) within the delay range period (here 20 time steps) was found to be the minimal condition to obtain sustained firing response in more than

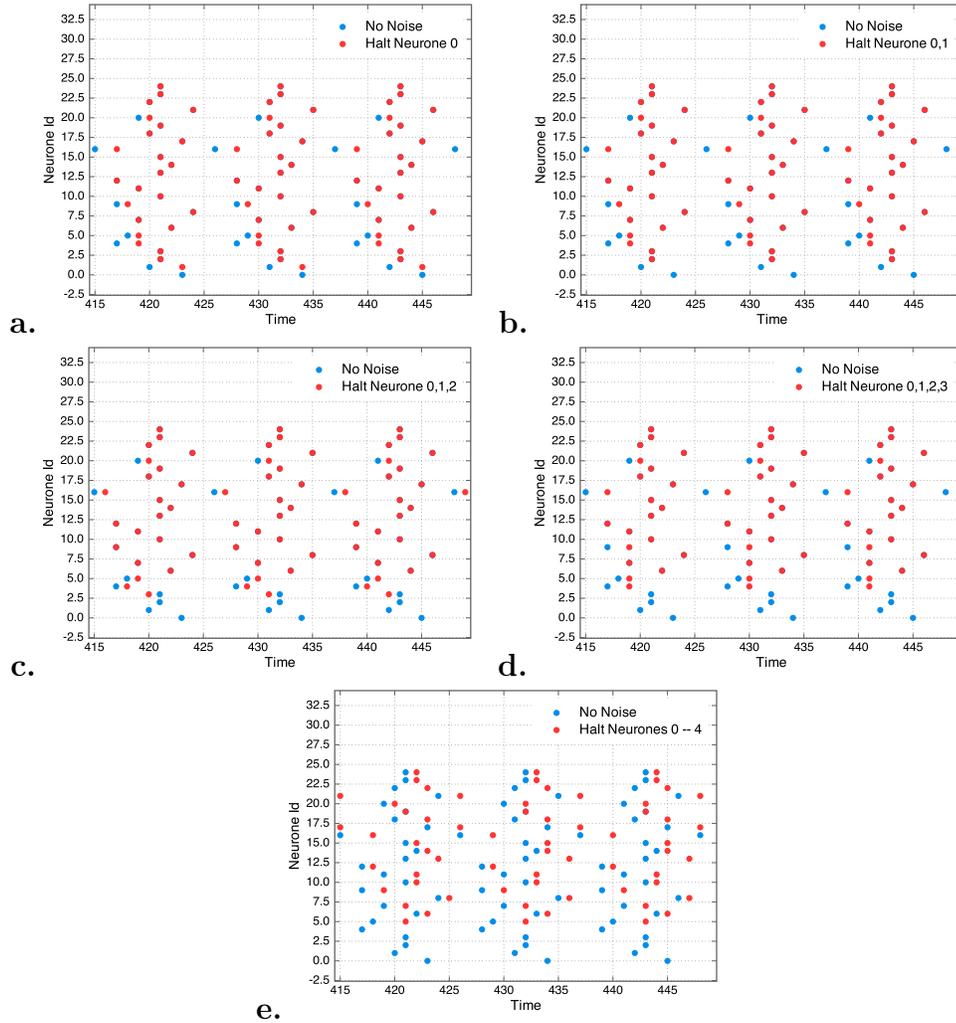


Figure 5.4: Sample Spike removal raster plots. No-noise raster (blue) is superimposed by noisy raster (red) when spike-times coincide. **a.** Neurone 0 is prevented from spiking – this led to 7 total misfiring neurones (including cell 0). **b.** Neurone 0 and 1 are prevented – led to no additional neurone misfiring. **c.** Neurones 0 to 2 are prevented – led to 8 neurones misfiring. **d.** Neurones 0 to 3 are prevented – led to 9 neurones misfiring. **e.** Neurones 0 to 4 are prevented – led to 23 neurones misfiring.

99% of runs for these networks (c.f. Section 3.3.1 for more details). Input was applied synchronously in order to simplify the setup. A total of 25 runs per network were made by increasing the number of inhibited neurones after the 800th time step was reached until each neurone was perturbed. This 800 time step initialization period was found to be sufficient for a stable attractor to be reached. Figure 5.4 illustrates the activity for a sample network after increasingly preventing the number of neurones firing. The no-noise raster (in blue) is superimposed by the noise raster (red). If cells are not affected by noise their spike times (in red) will be identical to non-noisy spike times (in blue) thus hiding non-noisy spikes. However, the appearance of non-noisy spike times (blue) indicates that the corresponding cell has been affected by noise, or in this case the inhibition of the subpopulation. Overall, we found that spike removal led spiking activity to settle into new spatiotemporal configurations. More interestingly, from the sample run illustrated in Figure 5.4 we notice that when a small number of neurones (4 or less) are prevented from spiking then only up to 9 neurones in total no longer spiked in step with the non-noisy response. However, when only a single additional neurone was prevented from spiking (see Figure 5.4.e.) then a total of 23 neurones spiked inconsistently with the non-noisy response. We investigate this interesting nonlinear effect in greater detail in the following section.

Stability In order to more easily measure the effect of noise on the spike times of correlated neurones in DCD networks, we developed a measure of spiking stability to quantify the change in timing when different amounts of noise are applied. As a measure of stability, we compute the spike time difference s for each neurone i to find how close their period converges to a target period λ over multiple renewal processes using the stability equation (3.3). Rather than computing a single stability estimate for each cell over the entire network simulation time, we repeatedly measure the stability s across small windows of size $n = 2$. This method results in an ‘instantaneous’ spiking stability measure that provides an indication of how stability changes over time during network simulation. Figure 5.5 illustrates this *averaged instantaneous spike time difference* (AISTD) for one network under 5 different conditions of spike inhibition. An AISTD of 0.25, for instance, indicates that the corresponding neurones shifted on average from their stable period by $2.5/10^{th}$ of λ , i.e., 2.75 time steps. Here, when no inhibition was presented to the network (Figure 5.5.a.) initialization instability quickly converged (0.0 spike time difference) after approximately 400 time steps and remained in this stable pattern during the entire network simulation. When four neurones were prevented from spiking starting from the 800th time

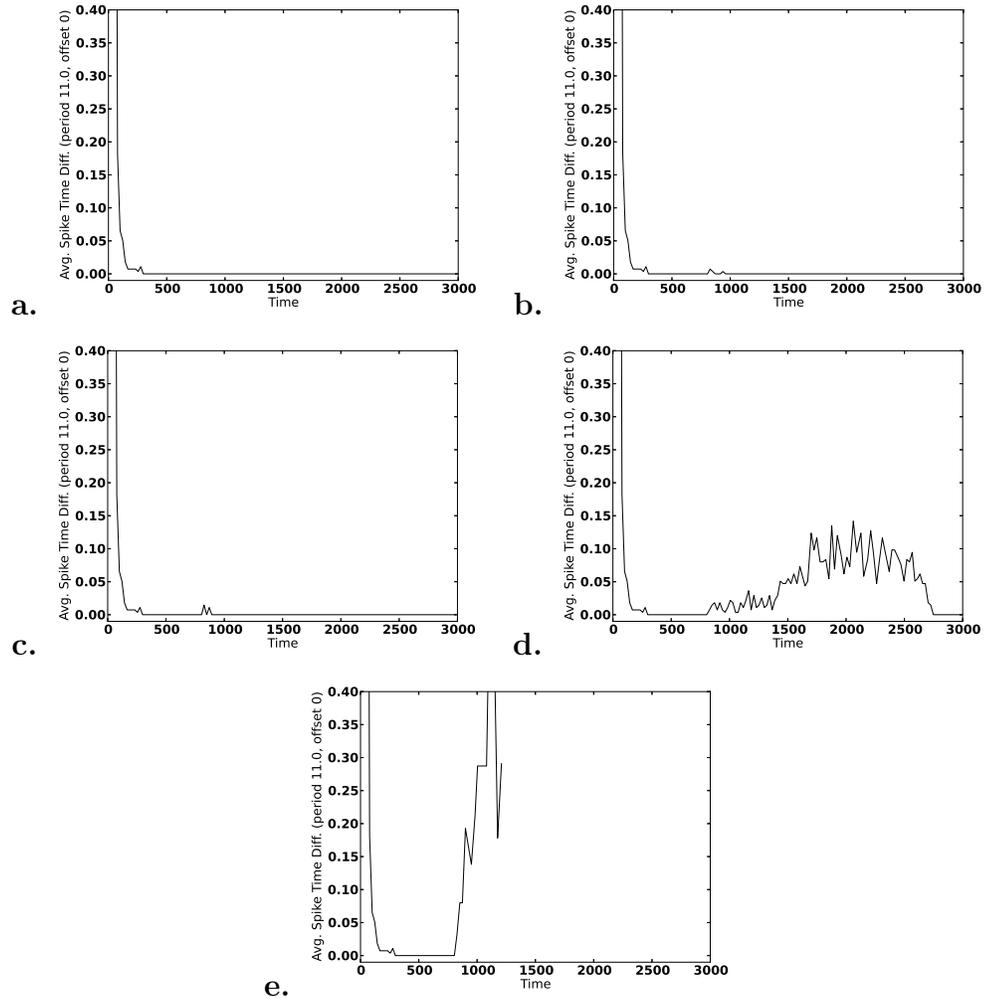


Figure 5.5: Spiking stability in prolonged spike-removal via *averaged instantaneous spike time differences* for a single network with identical input between noise scenarios. **a.**: no noise removal. **b.**: neurones 0–3 are inhibited from time step 800 onwards. **c.**: neurones 0–7 are inhibited. **d.**: neurones 0–11 are inhibited. **e.**: neurones 0–15 are inhibited.

step (Figure 5.5.b.), a small jump of 0.004 in instability took place but with a quick return to a stable pattern after approximately the 1000th time step. A similar short period of instability took place when eight neurones were inhibited (Figure 5.5.c.). This suggests that only brief and temporary instability occurs with relatively small perturbation via spike removal. When twelve neurones were inhibited, however, a much more pronounced and longer period of instability took place with an AISTD of up to 0.14 (Figure 5.5.d.). Hence, by preventing twelve neurones from firing, many of the remaining cells had spike times altered. However, a return to stable firing after approximately the 2750th time step suggests that the network was able to find a stable self-sustaining pattern provided the 13 remaining neurones. When sixteen neurones were prevented from firing after the 800th time step the average instantaneous spike time difference measure indicates a substantial amount of instability of up to 0.72 that eventually led to cessation of all activity after approximately 1250 time steps (Figure e.). This can be explained due to the much reduced number of neurones remaining (only 9) that were not prevented from spiking. Indeed, such a small number of neurones make it difficult for self-sustained activity to occur.

It is important to note, however, that when spiking does return to a stable regime it does not imply that the firing returns to the originally stable spiking pattern. The instantaneous spike difference only indicates regions of time when variability in spike timing occurs – not the characteristics of the patterns themselves. To examine the degree to which patterns preserve original spatiotemporal characteristics after perturbation we make use of ambiguity analysis as discussed above. Indeed, two identical patterns will have a maximum degree of ambiguity Λ_{max} equal to the number of cells in the network, however, comparing a non-noisy pattern with a noisy counterpart will have a reduced degree of ambiguity.

Criticality From the raster plots in Figure 5.4 we notice that in most cases removing a few neurones had a low to moderate impact on the number of other cells changing in spike-timing. As seen from the top left plot, preventing cells 0 to 3 from firing led to a total of 9 neurones misfiring: the 4 manipulated neurones plus 5 correlated neurones (Figure 5.4.d). However, by additionally preventing neurone 4 from spiking 24 of the 25 cells misfired as a result (Figure 5.4.e). This suggests that a critical tolerance to spike removal was reached by the spike pattern after which important deviation from the initially stable response pattern occurred. Formally criticality can be defined as:

Definition. *A DCD network undergoes a critical transition in activity when its maxi-*

maximum degree of self-ambiguity changes by some $\epsilon \geq \frac{N}{3}$ between two noise input amounts – where N is the number of neurones in the network.

This critical reduction in ambiguity is made more obvious by looking at the change in the maximum ambiguity as noise increases. The individual ambiguity results are shown for all 30 generated networks in Figure 5.6. In accordance with our definition we chose an $\epsilon = 8$ provided the 25 neurones in the network and rounding down to the nearest integer. In Figure 5.6 network responses to noise which led to a critical reduction in maximum degree of ambiguity when only one additional cell was removed are circled. For example, the top left plot shows how ambiguity dropped by 16 degrees of ambiguity when 4 or more neurones were prevented from spiking. By inspecting this behaviour for each random network in Figure 5.6, we notice that while critical reduction in ambiguity is not systematic for every network, this phenomenon does occur in more than half of the networks. The presence of criticality suggests that some cells play a greater supporting role on the spatiotemporal response of DCD networks than others.

Interestingly, other cells can, in some cases, negatively impact the network response in the context of noise. Indeed, an increase in self-ambiguity, as seen in the fourth network from the top the middle column took place after the 6th neurone was prevented from spiking. Hence, while a critical reduction was reached when 5 neurones were prevented from spiking, an increase in self-ambiguity took place after preventing the 6th cell suggesting that, in the absence of other cells, some neurone play a suppressive role in network tolerance to noise via prolonged removal of spikes.

Average Noise Tolerance While the sample runs above suggest that noise via spike removal can indeed be handled by DCDNs, an accurate measurement of the effect of spike-removal can be made by averaging the maximum degree of ambiguity between the non-noisy and noisy patterns. For statistical soundness spike removal was performed over 30 randomly generated networks of fully connected neurones with random delay assignments between 20ts and 40ts. To determine a sufficiently high number of samples to validate noise tolerance we measured the difference in ambiguity variation as additional networks were averaged. More specifically, we measured this change for each level of noise and set a change in average Λ_{max} below 0.5 to be satisfactory. We found that with up to 14 networks the largest difference in average Λ_{max} was above 0.529 but below 0.5 with 15 networks or more for each level of noise. For increased confidence we performed averages across a total of 30 networks leading

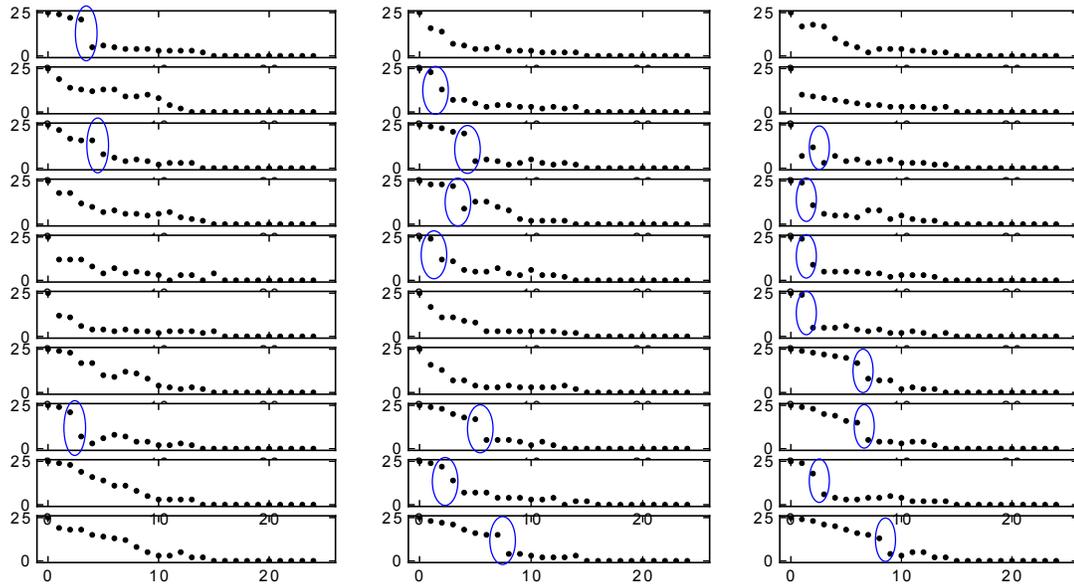


Figure 5.6: Degree of ambiguity (y-axis) as the number of neurones prevented from spiking (spike removal) increases (x-axis) amongst the 30 different networks. Critical reductions in ambiguity above 8 are circled.

to a change in mean close to 0.4 – this sampling size was used throughout the different noise conditions in these experiments.

Figure 5.7 illustrates both the average ambiguity as the number of spiking cells prevented from firing increases, as well as the normalized average. This measure compares the spatiotemporal ambiguity between the non-noisy pattern response and the noisy pattern responses in the time interval 1400 to 1411 time steps. The normalized average is a more accurate indicator because, as neurones are increasingly prevented from firing, the degree of ambiguity that can be found also decreases, hence the number of ambiguities found is divided by the maximum possible with respect to neurones available. From this Figure we noticed that on average up to 15 neurones (60%) can be removed such that at least one pair of stable spiking cells remains available for decoding. When removing 3 cells the maximum degree of ambiguity is approximately 14 suggesting that a coincidence threshold of 14 or less incident spikes on target cells will make it possible for a decoding network to identify the pattern despite noise on any three neurones in the memory network.

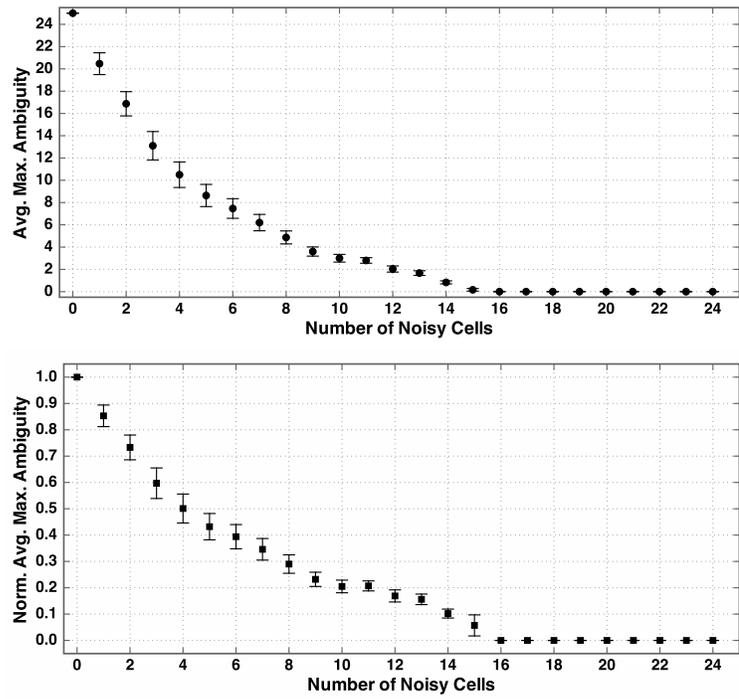


Figure 5.7: Average maximum degree of ambiguities (AMA) and normalized AMA as the number of spiking neurones prevented from spiking is increased in the prolonged condition. Higher ambiguity denotes greater noise tolerance. Error bars display standard error across 30 trials.

Scaling Finally, we aimed to generalize this empirical measure of noise via spike removal for networks of different sizes. To do so, we compared the noise response for the 25 neurone case ($N = 25$) above with networks of size $N = 100$ and $N = 1000$ neurones. When simulating larger networks, we proportionally increased the number of input cells, preserving it at 24% input, and reduced the number of connections so that the average number of afferents per cell remained at 24 ($I = 24$). The delay distribution, however, remained the same. By computing the average over 30 networks for each of these sizes and normalizing not only the degree of ambiguity in the scale 0 to 1, but also the number of neurones inhibited in the scale 0 to 1 we can compare how networks of different sizes behave with respect to an increase in noisy cells. Figure 5.8 shows the compared changed in average ambiguity between all three networks ¹. Overall, we noticed general non linear reduction in noise tolerance with noticeable slowing down typically after 30% of cells where prevented from spiking. Interestingly, this reduction in ‘self-ambiguity’ due to noise was more pronounced in the larger networks than in smaller networks after approximately 20% of cells were perturbed. For the 25 neurone networks, the standard error remained above 0.01 up until 64% of the neurones were prevented from spiking, at which point no more activity could be sustained. However, the standard error remained above 0.01 before 48% of neurones were inhibited for the 100 neurone networks, and above 0.01 before 36% of the neurones were inhibited for the 1000 neurone networks. Hence larger networks displayed poorer tolerance to noise than smaller networks. However, larger networks responded with less variability suggesting better conditions for noise tolerance prediction. Hence, from the general reduction in self-ambiguity after preventing more than 30% to 40% of neurones from firing we conclude that DCD networks exhibit a non negligible capacity for noise resilience with decreased noise tolerance but increased reliability in larger networks. Overall, this indicates that despite this severe noise condition where subsets of neurones are completely unable to participate, some features of the original pattern are preserved and usable for decoding.

5.1.3.2 Noise via Temporary Spike-Removal

While prolonged inhibition of neurones provides an indication of network performance under severe perturbation conditions, it is important to gain an appreciation for the

¹While noise measure were recorded by increments of 1 noisy cell in the $N = 25$ scenario, increments were increased to every additional 4 noisy neurones in the $N = 100$ condition, and every 40 neurones in the $N = 1000$ condition for a total of 25 points in each of the three network size conditions.

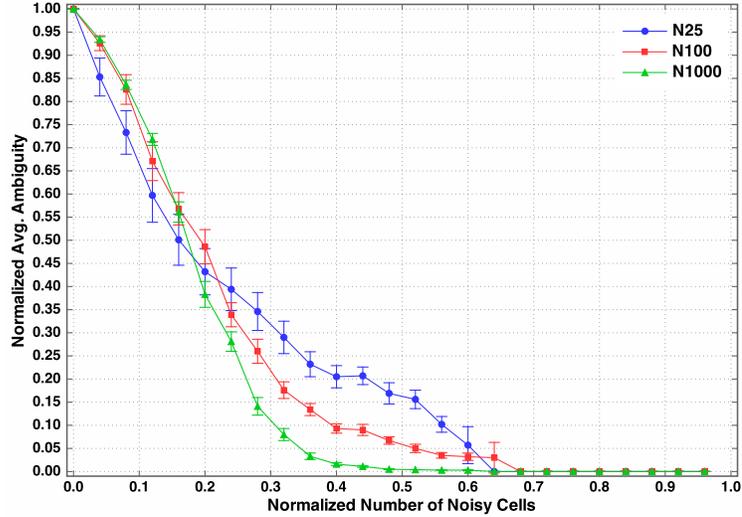


Figure 5.8: Normalized average ambiguity in noise via spike removal for a prolonged period. A comparison is made between DCD networks of 25 neurones, 100 neurones, and 1000 neurones. Error bars show standard error across the average 30 networks.

dynamics that arise in these networks under very subtle noise conditions. To test this other extreme we applied temporary inhibition to spiking cells and progressively increased the number of cells affected by this momentary inhibition. In particular, we want to determine whether this behaviour will be different at all, and if so, if it will be in respect of qualitative and/or quantitative change from prolonged inhibition.

Temporary inhibition was applied so as to disable a neurone or set of neurones for a single spike that occurred between time steps 800 and 811. After this, cells could resume spiking provided sufficient afferent coincidence detection. Similar to the prolonged condition, we measured noise by comparing the response of the non-noisy condition with the noisy-condition after this latter returned to a stable attractor. Here, we aim to investigate empirically how close to the original spiking pattern a network returns after temporary removal spikes occur for a subpopulation of cells.

Stability Thirty networks were generated on which single spike removal in the time interval 800 and 811 time steps was applied for an increasing number of neurones. Figure 5.9 illustrates how temporary inhibition affects the averaged instantaneous spike time difference (AISTD) of the network responses with respect to 4 of the 25 conditions for one sample network. From this Figure we notice that after initial input stimulation and a transient period the network reaches a stable limit-cycle attractor after approximately 600 time steps. By temporarily inhibiting 8 neurones in the time interval 800 to 811 (Figure 5.9.b) a noticeable perturbation with a peak at 0.1

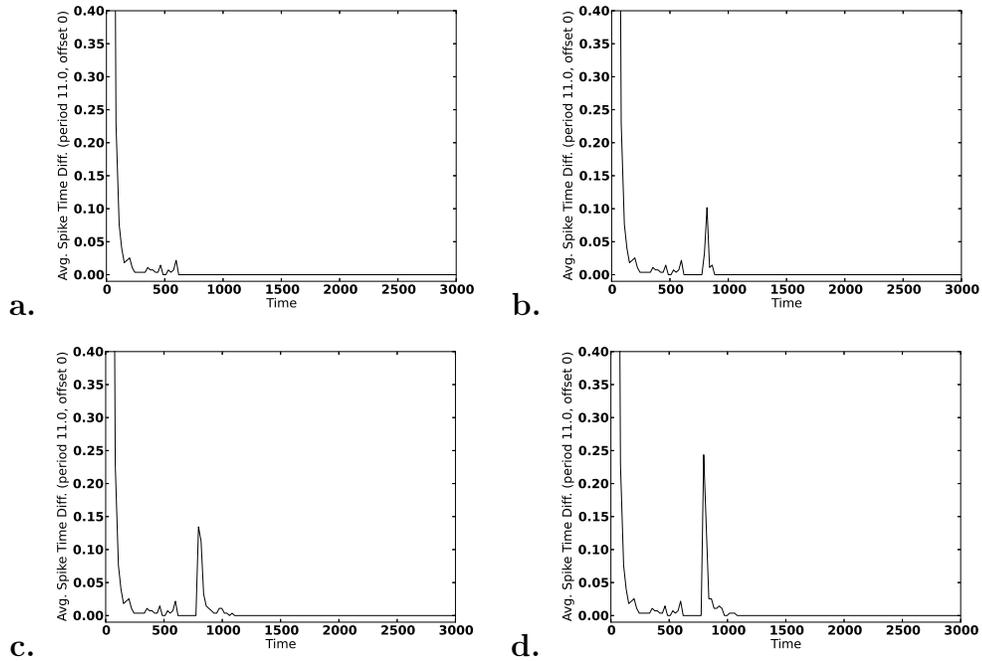


Figure 5.9: Spiking stability in temporary spike-removal via averaged instantaneous spike time differences for a single network. **a.**: no noise removal. **b.**: neurones 0–7 are inhibited only once in the time step interval 800 and 811. **c.**: neurones 0–15 are inhibited. **d.**: neurones 0–24 are inhibited.

occurred with continued instability for approximately 100 time steps after which the network returned to a stable attractor. This perturbation increased to a peak of 0.13 (or $1.3/10^{th}$ of a spike time) when 16 neurones were induced with noise. Finally, when all neurones had spike times removed temporarily a spike time shifting peak of 0.24 was reached. Overall, the network returned to a stable attractor in each condition within approximately 300 time steps from the onset of spike removal.

Criticality Similarly to the prolonged removal of spikes condition, ambiguity between the non-noisy condition and each of the 24 noise conditions was computed for each of the 30 different networks. This was done for the time interval 1400 to 1411 time steps. Despite the temporary nature of the perturbation effect, computing the average ambiguity across all 30 networks as this perturbation increased displayed a critical change in 11 of the 30 networks when assuming a criticality detection $\epsilon = 8$. Figure 5.10 illustrates the individual network responses. Overall, however, this critically was not as prominent as in the prolonged inhibition condition. For a more general perspective we computed the average ambiguity across all networks which we describe in the following.

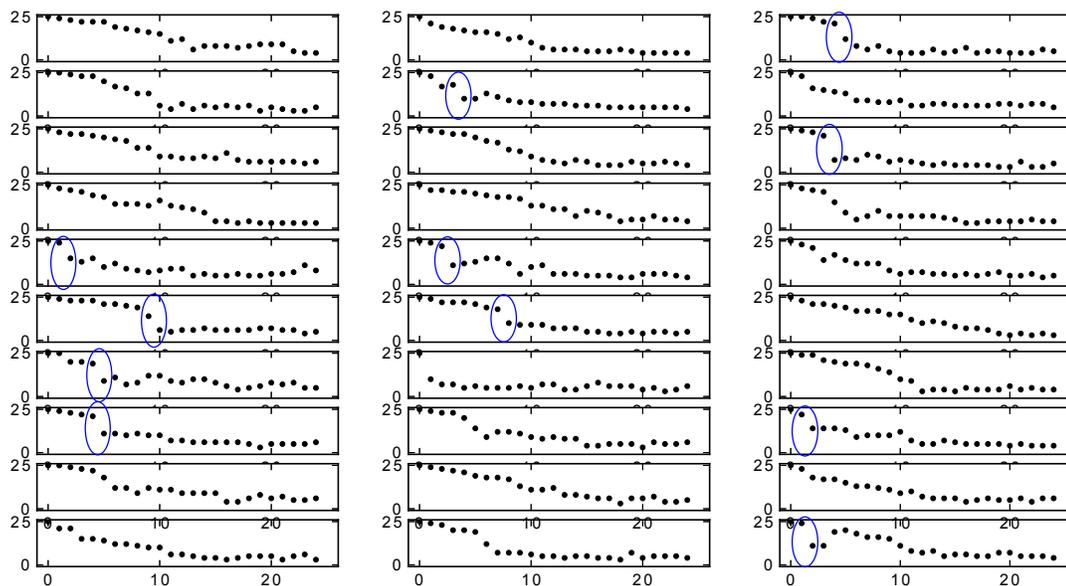


Figure 5.10: Degree of ambiguity (y-axis) as the number of neurones temporarily prevented from spiking increases (x-axis) amongst the 30 different networks. Critical reductions in ambiguity above 8 are circled.

Average Noise Tolerance To gain a general appreciation for the behaviour of these small DCD networks when inhibition to single spikes per cell is applied we computed the normalized average maximum self-ambiguity Λ_{max} between all thirty networks. Figure 5.11 shows a progressive decline in self-ambiguity as the number of temporarily inhibited neurones increases. Because all neurones resume firing by time step 1400, the normalized average ambiguity plot does not differ from the simple average ambiguity plot other than in scale. In contrast to the prolonged noise condition, temporary inhibition of neurones had less of a destabilizing effect on the originally encoded pattern. When contrasting these results with the prolonged removal condition from Figure 5.7 we found that when 24% of neurones (6 cells in the $N=25$ condition) were prevented from spiking the normalized average highest degree of ambiguity between noisy and non-noisy patterns was 0.394 (39.4% of the maximum Λ). However, when 24% of neurones were prevented from spiking only once in the temporary removal condition, the normalized average highest degree of ambiguity was 0.545 (54.5%). Interestingly, in the temporary spike removal condition, applying noise to all neurones did not lead to complete long term disruption of the pattern. In the temporary condition, however, we noticed that at least 16% of neurones (4

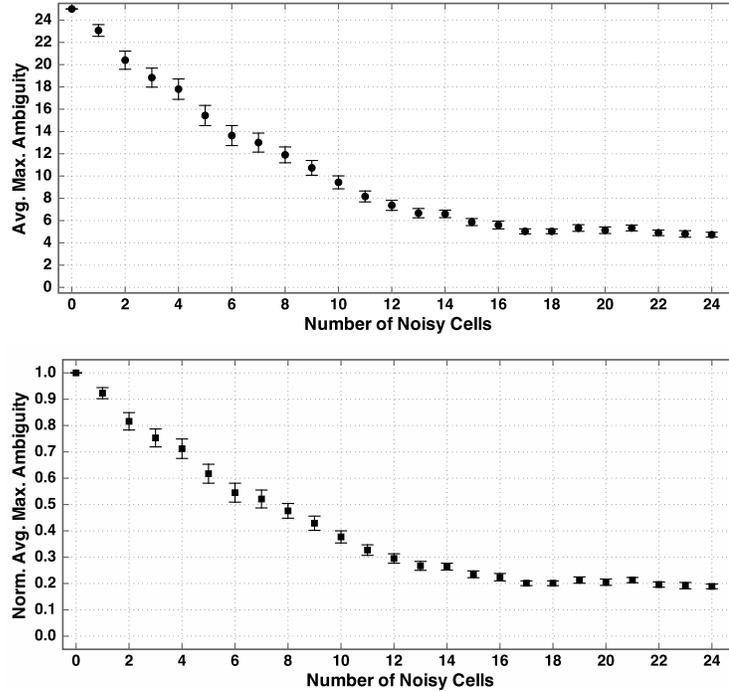


Figure 5.11: Average and normalized maximum degree of ambiguities as the number of spiking neurones prevented from spiking is increased in the temporary condition. Higher ambiguity denotes greater noise tolerance. Error bars display standard error across 30 trials.

cells) could still be used to decode the spatiotemporal patterns encoded in memory. This suggests that DCDNs could present interesting stability characteristics when a widespread, yet short, burst of noise is presented to the network.

Scaling When comparing between the average ambiguity for networks of 25 neurones with networks of 100 and 1000 neurones, we found a similar trend in reduced ambiguity as the number of spikes prevented from spiking in the time step window 800 to 811 increased. Figure 5.12 shows the nonlinear decrease in ambiguity with increased temporary noise when averaged over 30 random DCD networks. For the 25 neurone network we found that when 48% of neurones (12 cells) are temporarily inhibited, the normalized maximum Λ covered 29.5% of cells with a standard error of 0.018. In the 100 neurone network condition, noise to 48% of cells led to a normalized maximum Λ of 31.7% with a standard error 0.015. Finally, in the 1000 neurone network condition, we found that noise to 48% of cells led to a normalized maximum Λ of 30.5% with standard error of 0.004. Similar to the prolonged noise condition above, standard errors were noticeably lower in larger networks suggesting that temporary

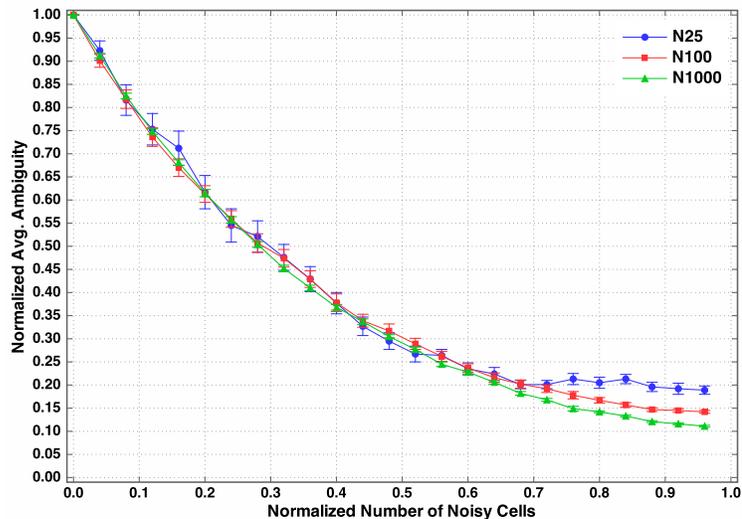


Figure 5.12: Normalized average ambiguity in noise via single spike removal (temporary). A comparison is made between DCD networks of 25 neurones, 100 neurones, and 1000 neurones. Error bars show standard error across the average 30 networks.

removal in larger networks should be more predictable than in smaller networks. In addition, comparison across all three network sizes shows a high degree of similarity in ambiguity behaviour up until approximately 2/3 of cells are subjected to noise. This suggests a scale-invariant behaviour in noise tolerance via single spike removal for noise applied to less than two-thirds of cells.

Comparing Prolonged and Temporary Noise Overall we found that noise in the form of absence of spikes has a declining influence on the identification between patterns as this noise increases. If sufficiently low decodable thresholds can be used on target decoding neurones then over half of the neurones can be perturbed before losing all distinct features. While in the more extreme, prolonged perturbation case, up to approximately 60% of neurones could be perturbed across network sizes. Subtle interference from single spike removals led to important spike pattern changes making increasingly fewer neurones available for decoding in the temporary noise condition. In this latter case, however, decoding remained possible even after all neurones were interfered with exactly once.

Interestingly, both prolonged and temporary spike removal scenarios displayed greater self-ambiguity in the smaller networks after a sufficient level of noise was applied. In particular, this is noticeable after approximately 20% of cells had noise applied in both the prolonged and temporary noise conditions. In the temporary spike removal condition, for instance, a little under 20% of neuronal spikes could still

be used to decode the original pattern in the 25 neurone networks when all neurones were affected. By contrast, a little less than 15% were available in the 100 neurone networks and approximately 11% were available in the 1000 neurone networks. At first glance, this suggests that smaller networks allow for proportionally greater tolerance to noise by making more features available for decoding under high levels of noise than do larger networks. The fact that features remained in all networks despite noise being applied to all cells suggests that noise does not systematically lead to a complete change in the spatiotemporal pattern and that partially identifiable spike timing patterns remain. The conditions for this feature preservation to take place, however, are much less likely in the larger network conditions because of the larger number of possible spike timing permutations. Indeed, while up to 5 cells remain available for pattern identification in the 25 neurone network, up to 110 cells remain available in the 1000 neurone network.

Despite these informative results, a number of other conditions could be explored with respect to spike removal in DCDNs. For instance, we could apply temporary noise more than once after transient periods of stability have been reached to observe the compound effect of spike inhibition on decodability. Another avenue of exploration would be to extend the period of temporary inhibition and its effect on decodability. In particular, we suspect that spike removal that prolongs for as long as the longest transmission delay (the implicit spiking memory of the network) then inhibition effects will be very similar if not identical to the prolonged inhibition condition found here. With respect to other quantitative alternatives, we believe the results found here exhibit the two most extreme cases of noise for DCD networks with high degree of connectivity and broad delay distributions. Because of this, we expect the stability, criticality and ambiguity profiles of quantitatively different noise conditions to lie in between the results presented here. Hence, we leave these avenues of exploration as open future questions and now turn to an alternative mode of noise: *misfiring*.

5.1.3.3 Noise via Prolonged Misfiring

A qualitatively different form of noise in spatiotemporal patterns occurs when the events at a particular spatial location vary in their temporal occurrence independently from the other events in the process while not disappearing. For instance, neurones can be made to fire pre-emptively with some given probability at each time step. Unlike spike removal which can only prevent target spiking from taking place, an offset in firing – or *misfiring* – may affect target cells by either preventing targets from

spiking or provoking new targets to spike. Importantly, misfiring does not prevent the noisy neurone from continuing to participate in the spatiotemporal pattern. Hence, in a prolonged condition of misfiring cells that are targets to the neurones made to misfire may shift in spike timing but the overall activity should not reduce.

Similar to the noise removal experiments 30 networks were generated and tested in all 25 cases corresponding to the number of neurones set to misfire. To investigate the effect of misfiring on stable spatiotemporal patterns we forced spiking on a subset of neurones in the population with a probability of $1/\lambda$ (one over the period of the repeating pattern – here $\lambda = 11$) at every time step. This probability forced the cells to fire with the same likelihood as non noisy cells but at different times when compared to their normal non-noisy counterpart. Hence, irrespective of whether a neurone was in a refractory period or not it was forced to spike preemptively before sufficient afferents led to threshold crossing about once on every limit cycle period λ . Similar to the spike removal scenario, we applied this misfiring in a prolonged fashion such that a subset of neurones were set to misfire from the 800th time step until the end of the simulation.

Stability Subsets of neurones were set to misfire stochastically after time step 800. In a similar way to the previous spike removal conditions we expected that this spike time difference would increase as the number of neurones set to misfire increased. Indeed, when looking at a sample network response we found that after time step 800 the average instantaneous spike time difference increased substantially as the number of neurones increased. Furthermore, due to the continued nature of this induced noise, networks perturbation was prolonged for the duration of dynamic memory simulation. Figure 5.13 illustrates the case for the different numbers of noisy cells. When no misfiring occurred (Figure 5.13.a) the network reached a stable attractor quickly after the onset of input within 250 time steps. When misfiring was applied to 8 neurones at time step 800 (Figure 5.13.b), perturbation reached AISTD peaks as high as 0.28. When misfiring was applied to 16 neurones (Figure 5.13.c) shifts in spike timing reached peaks up to 0.39. Finally, when misfiring was applied to all neurones (Figure 5.13.d) perturbation reached peaks as high as 0.49 (or $4.9/10^{th}$ of a spike time). While not as pronounced as prolonged spike removal, this amount of perturbation was more substantial than in the temporary spike removal condition. In contrast to other noise conditions the high probability of misfiring combined with its continued application led to network responses that never returned to a settling state of stable firing. In the temporary spike removal condition, for example, we found that

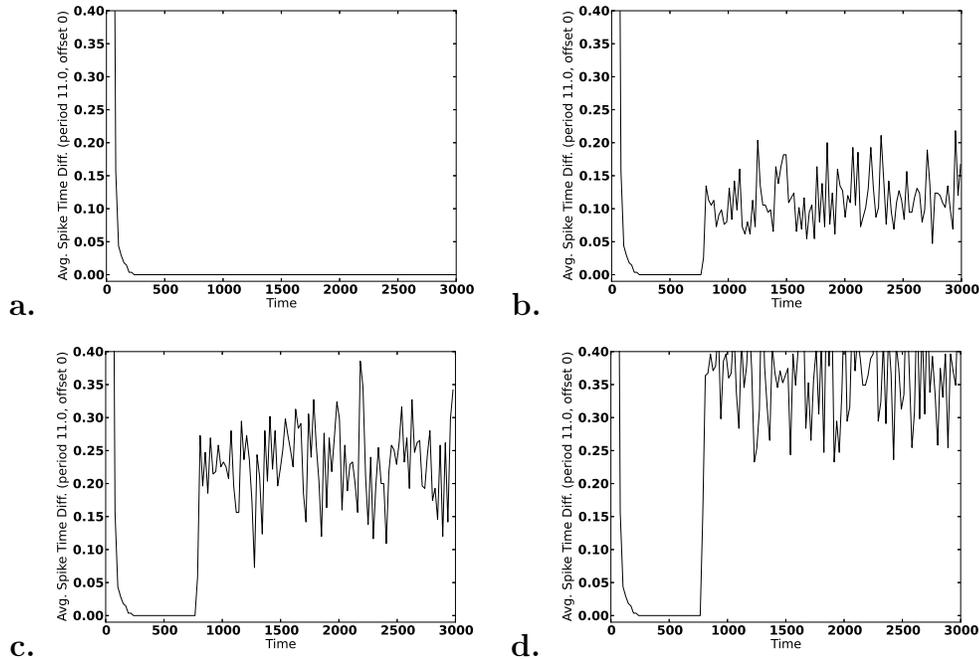


Figure 5.13: Spiking stability during prolonged misfiring via averaged instantaneous spike time differences for a single network. **a.** : no misfiring. **b.** : neurones 0–7 are set to misfire from time step 800 onwards. **c.** : neurones 0–15 set to misfire. **d.** : neurones 0–24 are set to misfire.

a return to stable activity was reached within a few hundred time steps after spike removal. Here, however, the high probability of misfiring prevented the network from returning to a stable state within a few hundred time steps because the misfiring rate was greater than the rate of return to a stable attractor. However, a lower probability of misfiring could potentially enable the network to return to a stable limit cycle if this probability were reduced to induce misfiring at a rate below the rate of settling. This could be examined by decreasing the probability of misfiring by raising λ by multiple exponents such as $1/\lambda^2$, $1/\lambda^3$ or less. We leave this as an opportunity for future work.

Criticality To examine the degree of variability between randomly generated networks we plotted the independent ambiguity analysis plots to show how network spatiotemporal characteristics change as the number of misfiring cells increased. Figure 5.14 shows the self-ambiguity for individual networks as the number misfiring cells increases. From this figure we notice that noise sharply reduced the number of original response pattern features when over a third of the cells were misfiring. However, when less than a third of the noisy cells were perturbed critical change in

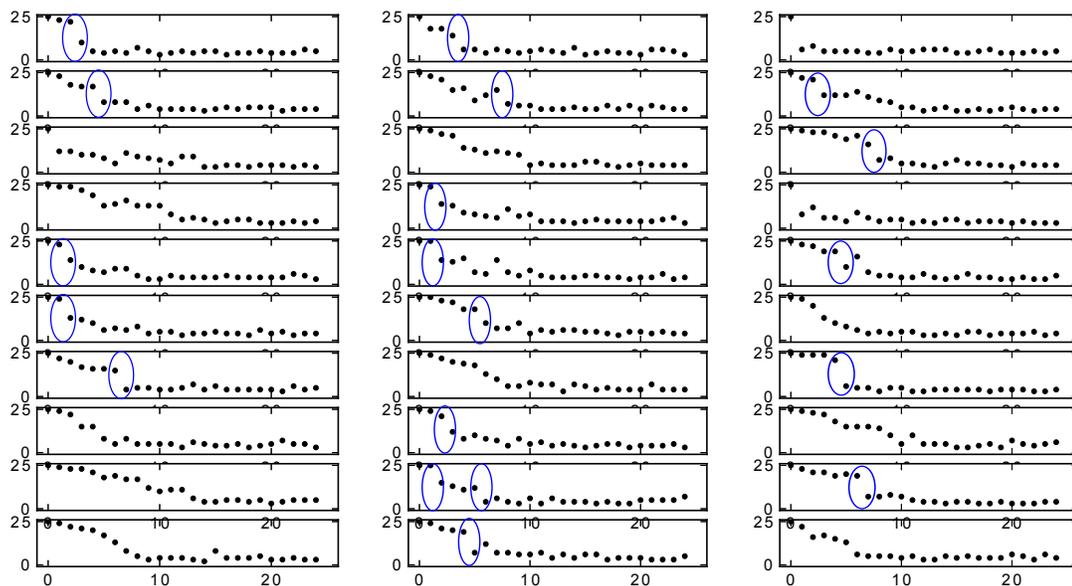


Figure 5.14: Degree of ambiguity (y-axis) as the number of neurones misfiring continuously increases (x-axis) amongst the 30 different networks. Critical reductions in ambiguity above 8 are circled.

self-ambiguity in 18 of 30 networks was found for a criticality threshold of $\epsilon = 8$. Interestingly, multiple critical changes was found in one network (9th network in the middle column) which further suggests that individual cells will variably exercise a supportive or suppressive role of the original spatiotemporal pattern. Although criticality does not appear to be guaranteed for any given network from the plots in Figure 5.14, the variable significance of individual cells suggests that the spatiotemporal response to noise could also be dependent on the specific order in which cells are selected to misfire. Here, because networks were generated with random connectivity, noise was applied in order from cell 0 to cell 24. This is sufficient to demonstrate average network behaviour. However, one could investigate in individual networks with respect to their behaviour when alternate orders in which cells are selected is tested. We leave this an opportunity for future work. Overall, the observation that individual cells will have varying effect on the network dynamics when confronted with noise is made possible by viewing the individual activity of networks. Taken together however, a smooth non-linear degradation of feature availability occurs as we show in the following.

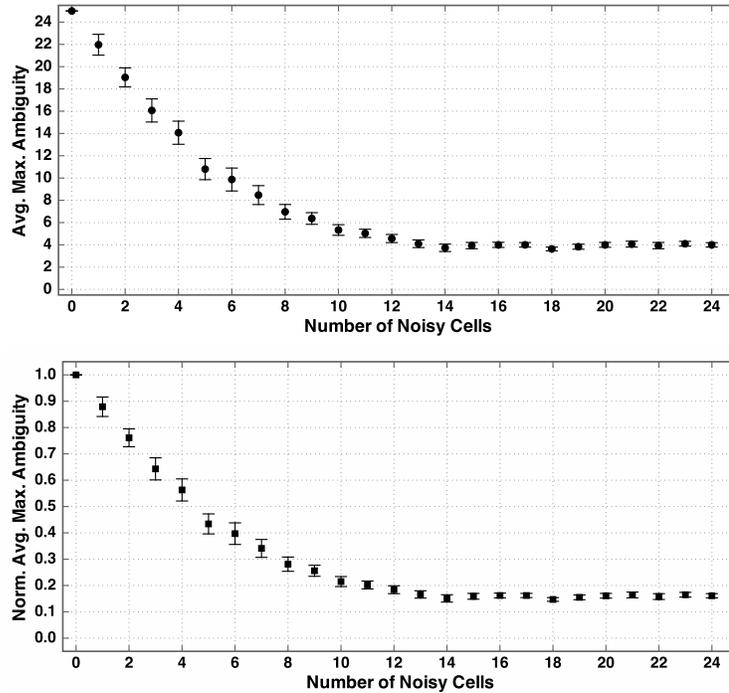


Figure 5.15: Average and normalized maximum degree of ambiguities as the number of spiking neurones set to misfire is increased in the prolonged condition. Higher ambiguity denotes greater noise tolerance. Error bars display standard error across 30 trials.

Average Noise Tolerance Similar to the spike removal scenarios, we computed the maximum degree of ambiguity between the non-noisy pattern and noisy patterns to quantify the number of neurones which can be noisy while permitting the pattern to be identifiable. We did so in the time interval 1400 to 1411 which lies within the disrupted portion of the pattern as seen in Figure 5.13. Figure 5.15 shows how the average maximum degree of ambiguity across networks decreased as the number of misfiring neurones increased. This shows that the networks became less tolerant to noise as the number of misfiring cells increased up until approximately 50% of the neurones were affected by noise, after which a plateau was reached leaving about 20% of cells available for pattern decoding irrespective of the increasing noise. Hence ambiguity was not completely lost after a large number of cells misfired because some original spike times could retain identical spatiotemporal characteristics when compared to the non noisy pattern. For instance, when all 25 cells were set to misfire an average maximum degree of ambiguity of 4.13 (16.5%) was found suggesting that decoding neurones with coincidence threshold of 4 or less would still be capable of decoding original patterns despite the stochastic noise on all neurones.

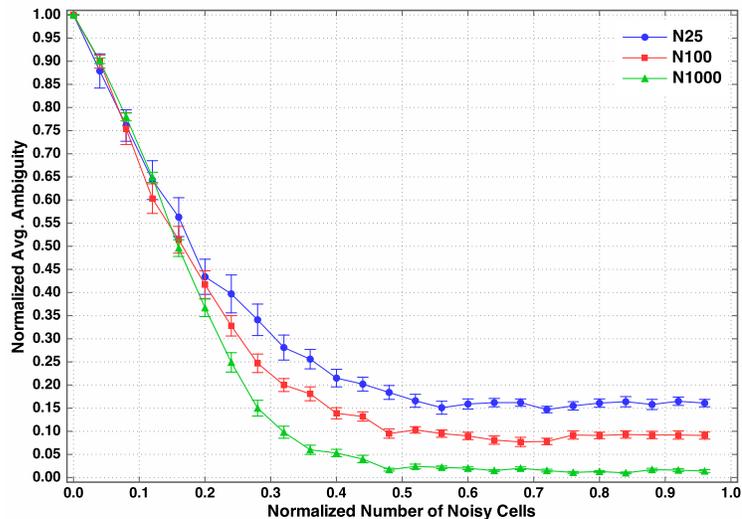


Figure 5.16: Normalized average ambiguity in noise via prolonged spike misfiring. A comparison is made between DCD networks of 25 neurones, 100 neurones, and 1000 neurones. Error bars show standard error across the average 30 networks.

Scaling In order to obtain a better understanding of how this behaviour scales to larger networks, we ran the same continued misfiring experiment on networks of 100 and 1000 neurones. Figure 5.16 shows the normalized average ambiguity for a normalized number of misfiring cells as the number of neurones set to misfire increased. From this figure we first notice that all three network sizes led to a similar non-linear decrease in ambiguity until a plateau was reached when about 50% of neurones were affected by noise. Despite this similarity, pattern preservation was greater in the smaller networks than in larger ones as seen in the previous spike removal cases. Approximately 15% of neurones were available in the 25 neurone networks after the plateau was reached, while only approximately 9% were available in the 100 neurone networks and approximately 2% in the 1000 neurone networks. This further suggests that despite proportional connectivity, input and noise, smaller DCD networks are better suited for noise tolerance when studied independently from other input conditions.

Variation in ambiguity was slightly more important in the smaller 25 neurone networks with an average standard error of 0.015 when 48% of neurones were affected by noise versus a standard error of 0.01 for the 100 neurone networks and 0.003 for the 1000 neurone networks with 48% noise. As in the spike removal conditions, this suggests that despite proportional connectivity and input stimulation, larger networks are more stable in the face of noise via prolonged misfiring.

Noise via Temporary Misfiring

In a manner identical to spike removal, we analyzed the effects of temporary misfiring and contrast it with long term perturbation to observe their effects on DCD networks. Here we show results as the number of misfiring cells increased and where cells were set to misfire exactly once at time step 800 during the entire simulation.

Stability A sample analysis of stability of the network over time via average instantaneous spike time difference (AISTD) revealed that, similar to the temporary spike removal condition, instability occurred momentarily at the onset of noise induction and returned to a stable attractor after a short period of less than 500 time steps. As the number of misfiring neurones increased the amplitude of the AISTD increased. As seen in Figure 5.17, instability on the initialization of the network with input quickly converged to a stable attractor within the first 500 time steps as seen in Figure 5.17.a. When noise was introduced at time step 800, an instability peak occurred with increasing amplitude as the number of misfiring neurones increased. A peak of 0.1 was reached at time step 800 when 8 neurones were induced with noise as seen in Figure 5.17.b. This means that spike times shifted by $1/10^{th}$ of a time step in that short time period. When forced temporary misfiring was applied to 16 neurones, a shift of 0.178 time steps (or $1.78/10^{th}$) occurred overall (Figure 5.17.c). When misfiring was applied to all neurones, as seen in plot d, this spike time difference peak reached 0.27 (or $2.7/10^{th}$ of a time step). By simply taking the difference between these peaks we can already identify a non-linear decrease in change as the number of cells to which noise is applied increases linearly. While a correlation between perturbation behaviour and how dynamic memory is susceptible to noise is likely to exist, stability alone is not sufficient to account for how memory depreciates with noise because the network possesses spatiotemporal characteristics that could be independent of perturbation change alone. this is especially made obvious in this case, as in the temporary spike removal case, where the network returns to a table attractor shortly after noise induction. After examining criticality, we instead identify this change more accurately by analyzing the degree of ambiguity between the network's behaviour under different amounts of noise.

Criticality We have seen in the previous conditions how a high degree of variability in tolerance to noise exists when approximately the first 30% to 40% of neurones are noisy. In the case of temporary misfiring variability is also relatively high when perturbation is applied to the first third of the population; however, this degradation

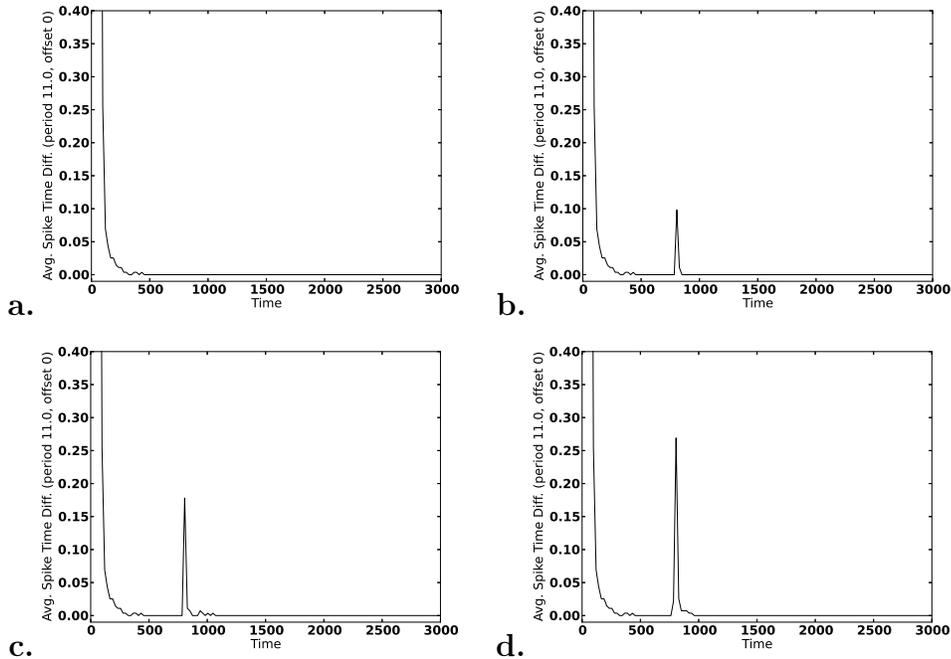


Figure 5.17: Spiking stability in temporary misfiring via averaged instantaneous spike time differences for a single network. **a.**: no misfiring. **b.**: neurones 0–7 are set to misfire within time step 800 to 811 with probability $1/\lambda$. **c.**: neurones 0–15 set to misfire. **d.**: neurones 0–24 are set to misfire.

in tolerance is overall more gradual than in the previous cases. Figure 5.18 shows this behaviour for each of the 30 networks tested. Among all tested networks, 9 of 30 led to a critical reduction in maximum degree of ambiguity for a given amount of applied temporary noise. In comparison to the previous forms of noise, this relatively subtle effect from temporary misfiring on other affected cells suggests that more spatiotemporal features will remain similar to the original non-noisy pattern.

Average Noise Tolerance In order to quantify the amount of noisy neurones the network could tolerate so that the stored input pattern remained decodable we again computed the maximum degree of ambiguity between the non-noisy condition and the non-noisy conditions and averaged this across 30 generated networks. Figure 5.19 illustrates the progressive decrease in ambiguity between non-noisy and noisy patterns. This is synonymous with the number of neurones which can be noisy while decodability via coincidence detectors remains possible – provided the coincidence threshold of target cells is correspondingly reduced. As could be expected from looking at the individual plots in Figure 5.18 the decline in ambiguity is nearly linear suggesting that perturbation to neurones has less impact on target cells due to correlation than

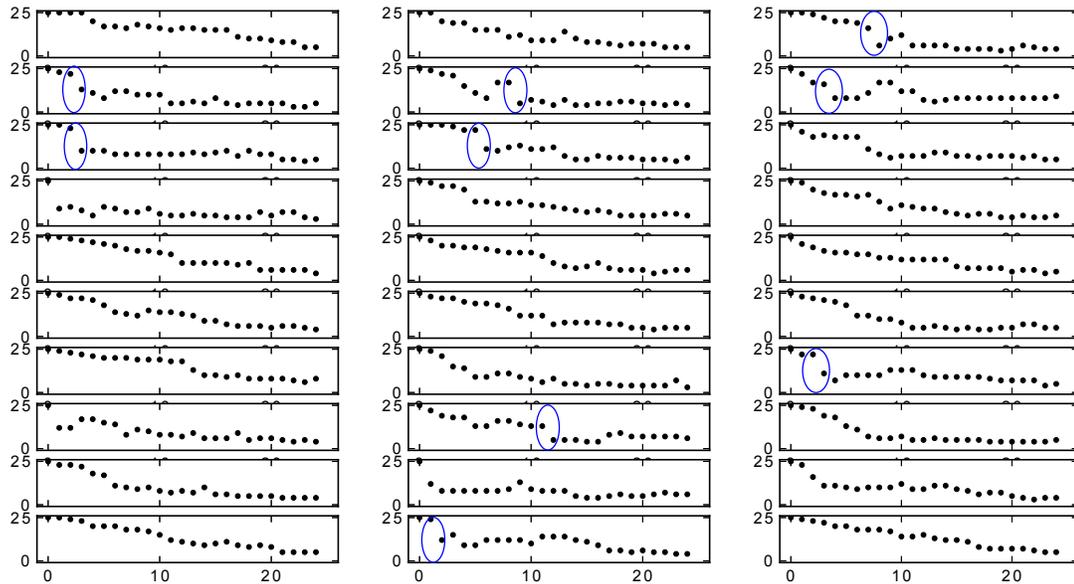


Figure 5.18: Degree of ambiguity (y-axis) as the number of neurones misfiring only once increases (x-axis) amongst the 30 different networks. Critical reductions in ambiguity above 8 are circled.

does prolonged misfiring or spike removal. Instead, preservation of spatiotemporal characteristics enabled a linear degradation of available features. This is indeed the least severe noise condition because not only neurones remained active even after perturbation, but this perturbation was only applied once. Still, when all neurones are affected via single misfiring, no more than 5 neurones (or 20%) of the ambiguities were available for pattern decoding.

As a note of interest, this noise condition is reminiscent of applying novel input stimuli which trigger target neurones to spike at an offset time on a pre-existing spatiotemporal pattern. This form of input ‘convolution’ was the case as seen in the robot memory tasks (Sections 4.3, 4.5) where after the presentation of a context stimulus to 6 neurones, a cue stimulus of neurones was presented at a later time. Delay decoding required that the coincidence detection of decoding cells had to be relatively high despite only having to differentiate between four patterns (AX, AY, BX, and BY). This was because the cue stimulus had the same impact as a noise stimulus to 6 neurones would have here. As seen from Figure 5.19 approximately 50% of ambiguity with the original pattern remains after noise to 6 neurones (24% of the population) is applied to the DCD networks under investigation. By analogy,

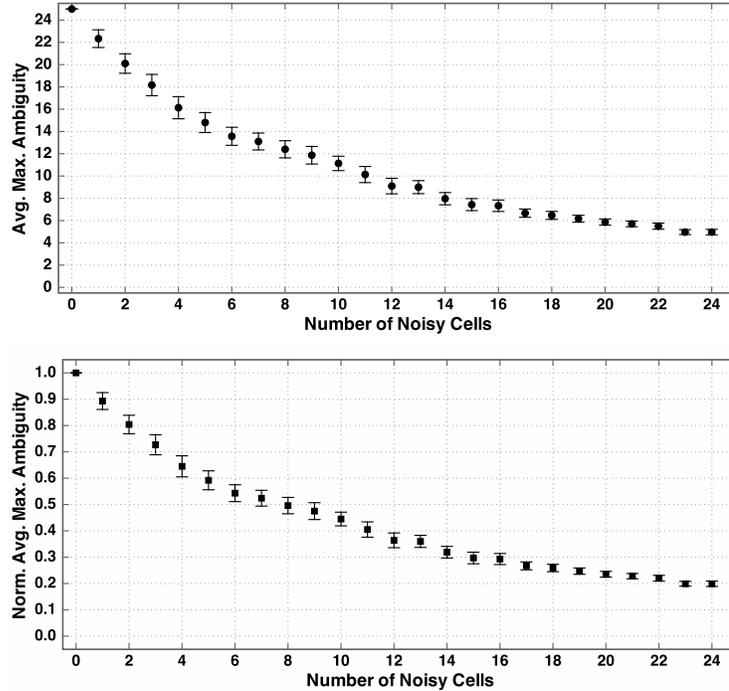


Figure 5.19: Average and normalized maximum degree of ambiguities as the number of spiking neurones set to misfire is increased in the temporary condition. Higher ambiguity denotes greater noise tolerance. Error bars display standard error across 30 trials.

this means that to decode a novel pattern convolved with the original a decoding threshold above $N/2$ will be necessary to capture the new pattern features – note, however, that this does not take into consideration the requirement of differentiation from other patterns to decode. In the case of noise the opposite scenario is of interest where the decoding network should ignore misfiring cells by reducing the coincidence threshold of decoding cells. Delay mapping then acts as a filter to only allow the detection of properly timed (non-noisy) spikes. Indeed, this hints at the relation between convolved pattern storage, disambiguation and noise which we explore in greater depth in Section 5.2.3.

Scaling Similar to the previous experiments we simulated network behaviour for 30 networks of size 100 neurones and 30 networks of size 1000 neurones to obtain a comparison between the smaller network of size 25 and these larger counterparts. Overall, we found that behaviour was largely comparable between network sizes as seen in Figure 5.20. All three network sizes showed a near linear decline in ambiguity at the level of noise increased linearly. When 48% of neurones were set to misfire once

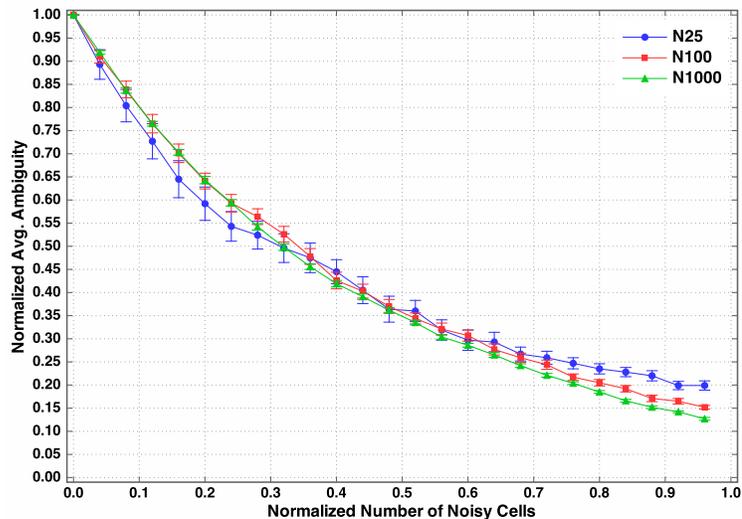


Figure 5.20: Normalized average ambiguity in noise via single spike misfiring. A comparison is made between DCD networks of 25 neurones, 100 neurones, and 1000 neurones. Error bars show standard error across the average 30 networks.

in the 25 neurone networks, self-ambiguity remained present in 36.4% of neurones with a standard error of 0.028. In the 100 neurone scenario, when 48% of neurones misfired self-ambiguity of 37% with a standard error of 0.015 was reached. Degradation was smoothest in the 1000 network case, with 36.2% ambiguity remaining and a standard error of 0.005 after 48% of cells misfired. While variability in ambiguity was clearly lower in larger networks feature preservation remained very similar with only slight divergence in performance after approximately 60% of cells had noise applied. As in the previous noise cases, smaller networks preserved greater ambiguity with the original non-noisy response in comparison to larger networks by the time all cells had noise applied. Despite these nuances scaling networks in size did not have a substantial impact on their response to noise via temporary misfiring. Just like the smaller 25 neurone network, larger networks exhibit more linear reduction of feature availability in this noise condition than when facing prolonged misfiring. This comparison is similar to that between the temporary spike removal and prolonged spike removal where in the temporary case the decrease in ambiguity was also less sudden but more gradual. This can be expected due to the more subtle interference from single spike perturbation versus the persistent interference of neurones in the prolonged case.

Further Considerations The third type of noise that should be considered in spiking networks is spike-addition. Here additional spikes could occur within the assembly of cells due to independent external input stimulus or other response factors.

However, in our present network model, the high degree of connectivity would lead to firing with rates as high as the inverse of the refractory period. This suggests that any additional spike will not in fact be able to influence cells which tend to spike as soon as they return to their resting state. Hence, we did not conduct analysis of spike-addition for the present network model.

Overall, we found that misfiring does not eliminate the possibility for decoding even when noise of this type affects all neurones in the network. On the other hand, spike removal, despite performing well in the temporary noise condition, cannot cope with noise to more than 60% of neurones when inhibition is applied for a prolonged period. Despite this, one could say these networks react to noise with ‘graceful degradation’ in virtue of their decelerating decrease in decodability as the number of noisy cells increases.

5.1.3.4 Discussion

We first introduced the notion of degree-of-freedom of firing-chain patterns. Related to this notion, we showed that noise estimates can be made a-priori by computing the number necessary cells required to form sub-patterns that are sufficiently large to produce a coincident spike on targets yet sufficiently small so as to not be affected by potentially noisy cells while ambiguities with other patterns exist. This was accomplished by evaluating the size of these sub-patterns as the number of ambiguities increases. This estimate, however, does not provide an indication with respect to the sensitivity of spikes which, in these networks, are highly correlated to previous spikes in the population. We thus conducted empirical evaluations of the number of neurones that can be used for decoding in the cases of spike-removal and misfiring. While ambiguity estimation was used here to compare spatiotemporal characteristics of the same pattern in the non-noisy and noisy scenarios, no consideration with respect to potential ambiguity between these patterns and other possible patterns in the same network was made. Indeed, for all empirical conditions investigated here in the 25 neurone network cases, except in prolonged spike removal, 15% to 20% of neurones preserved decodable features that could be used to recognize the original patterns by using coincidence detection thresholds of equal or lower size on decoding neurones. Delay mapping from the memory network to the decoding network would then serve as a filter by preventing, with some stochastic probability, the decoding of noisy features. In a more broadly applied context, however, these preserved features despite high amounts of noise could coincide identically with other memory patterns

under the same noise conditions. This would result in completely unusable features when high rates of noise are present. A better appreciation for this memory to noise tradeoff will be explored in Section 5.2 where we examine the relation between pattern ambiguity and noise tolerance. This will revisit the notions of noise with ambiguity as it relates to unpredictable firing.

5.2 Effective Memory of Delay Coincidence Detection Networks

Delay coincidence detection (DCD) networks will display a broad set of dynamics depending on their amount and distribution of connectivity, delay characteristics, and neuronal properties. Here we continue to take under consideration the stable limit-cycle attractor of fully connected networks of 25 neurones with coincidence detection threshold of 2.

Here, we aim to evaluate the informational characteristics of these networks when acting as dynamic memories. As seen on the robot T-maze navigation task, fixed network topologies that fit DCD network characteristics offer a physical ‘platform’ for distinct firing-chains to occur provided distinct input was presented to the network. What we wish to establish presently, is the actual memory capacity of these networks in terms of the number of distinct patterns they can store.

We first evaluate self sustained dynamics in networks of 25 neurones with the synchronous stimulation of 6 neurones as input. We quantify the number of patterns that can be stored for a given coincidence detection threshold T on decoding cells (the decoding threshold for short) by counting the number of patterns that can be disambiguated for that decoding threshold T . For example, if 10 distinct patterns across 25 neurones requires at least 8 neurones to be detected coincidentally to disambiguate the patterns, then a decoding threshold of 7 or less is insufficient because it will imply that at least two patterns will be confused. Evaluating network capacity thus requires that the decoding threshold be found for a set of responses of a given network. We will return to this with a more formal definition of effective memory. In addition, we evaluate the capacity of a larger network of 50 neurones to demonstrate the scalability of DCD networks as potentially significant mechanisms of dynamic memory in nervous systems.

In the following, we begin by investigating how the degree of correlation between input patterns, in terms of the number of identical cells they share, affects the degree

of ambiguity between their corresponding network responses. Indeed, from a memory perspective it is important to know whether or not the network properties conserve input pattern feature similarity when spiking responses are produced. We then present our results on the empirical evaluation of the number of patterns that can be stored for a given decoding threshold. We then complement this analysis with empirical evaluations of the likelihood and 'weight' for a given degree of ambiguity and its relation to noise.

5.2.1 Input Correlation

Often we expect systems which undergo similar configuration changes to exhibit similar outset behaviours. Typical neural networks that are trained to encode input pattern by adjusting connection weights can converge towards similar attractors (response patterns) when new inputs only share a subset of features with the original trained input. This suggests that inputs that share some degree of initial correlation will lead to outputs with some, if not greater, correlation. This is not necessarily the case, however, with all neural networks. As we show in the following, DCD networks possess highly diverging response characteristics reminiscent of chaotic attractors despite similarity on the input for both spatial and temporal features. In Subsection 3.3.1 we have seen that when a sufficiently high degree of connectivity is present in these networks, spike coincidence correlation leads to stable attractors. However, the initialization period before the sustained periodic regime is reached will chaotically drive DCD networks into unique periodic patterns. In the following we show how similarity on the input does not lead to similarity or convergence on the output in DCD networks.

5.2.1.1 Preliminary Investigation

To gain an intuitive sense as to the degree of correlation to expect between network responses for which similar inputs were provided, we stimulated a network on two separate occasions with a pulse to 6 neurones simultaneously where 5 of these were identical between both occasions. Figure 5.21 illustrates the spiking response of the network over the first 110 time steps (110ts) for both input conditions. We found that despite 83% overlap of the input pulses, all spike timing overlap was lost between the network responses by time step 100. Hence, while many cells had overlapping spikes in the first 100 time steps, this overlap progressively reduced as time progressed. This preliminary investigation thus provides anecdotal evidence for the hypothesis that a

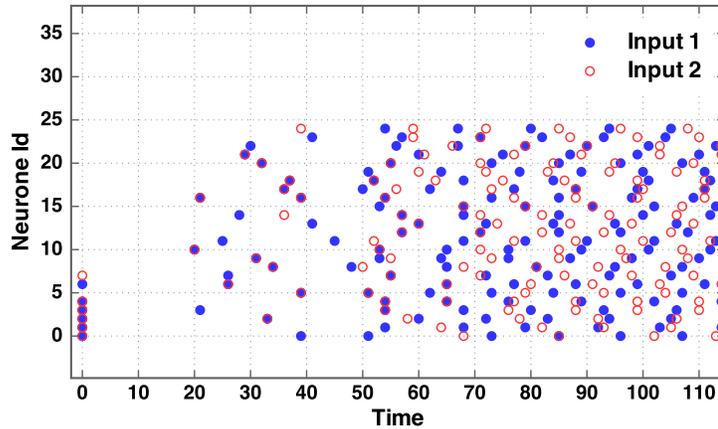


Figure 5.21: Sample raster showing divergence of network response despite similarity of the input. Both inputs are run on the same DCD network. Input to six neurones only differs by one cell as seen at time step 0.

high degree of divergence takes place within a relatively short period of time in DCD networks.

From a decoding perspective, however, the lack of overlapping spike times does not prevent the occurrence of ambiguity. As seen in Section 4.4, responses should not simply be distinct in spike timing, but in their inter-spike (inner-delta) offsets. Indeed, two responses that are offset by some number of time steps will remain fully ambiguous. Hence the ambiguity measure, which takes into account inner-deltas, is more useful in providing an accurate indicator of the amount of overlap between network responses.

5.2.1.2 Input Ambiguity

To establish the amount of ambiguity between related inputs, we generated all possible (177,100) pattern responses for each possible synchronous input combination to 6 neurones in random DCD networks of 25 cells. We repeatedly selected two random input conditions and performed ambiguity analysis between their responses. For each of the 10 randomly generated networks, we averaged the highest degree of ambiguity Λ found (c.f. Section 4.4) across paired comparisons by classifying them with respect to the number of shared input neurones. The final mean ambiguity was then averaged over the 10 randomly generated networks which was sufficient to obtain a standard deviation below 0.05 for all but the 4 and 5 input conditions. This lower standard deviation was due to the much lower likelihood that any two random input patterns would share 4 or more common on the input.

In total, there are 6 classes: from 0 common input neurones to 5 common input neurones. Because there are only 20 inputs with exactly 5 common cells for any given input pattern, the probability of comparing two patterns sharing 5 neurones was much lower than comparing two patterns with fewer common neurones². Accuracy is thus unequal between overlapping conditions. Figure 5.22 illustrates the results from, on average, 14,165 comparisons for each network.

From these results we found that an average degree of ambiguity between 4 and 5 was found for each input correlation condition. When 5 input neurones among the 6 were identical the degree of ambiguity between response patterns tended to increase. This increase is not significant, however, due to the low number of comparisons possible with 5 correlated input cells (on average 10 across all networks) and only 25 neurones in the population. Indeed, only 20 responses exist for any set of 5 identical input cells for a total of 6 inputs. These results suggest that DCD networks lead to fully diverging responses despite similarity on the input. Hence, it would appear that stimuli that share many features can still be clearly differentiated from the output network response. While this could be advantageously exploited for the detection of subtle nuances in the input, it also prevents feature similarity to be preserved in the network. This suggests a distinct functional role for DCD networks in contrast to other networks that offer feature similarity preservation. Instead, DCD networks could effectively produce and maintain distinct responses for otherwise highly similar stimuli; furthermore, they could regroup arbitrarily distinct stimuli into associated wholes by increasing ambiguity via reduced decoding thresholds. This latter mechanism could enable a rudimentary form of classification via ambiguity by considering patterns with common spatiotemporal features without regard for input similarity. These features of DCD networks have motivated the research reported in this thesis and are supported by the demonstration of DCD network utility in the robot task.

Having gained greater insight into the transformational properties of DCD networks with respect to input correlation, we can now evaluate how patterns produced by a single network for different inputs result in a change in degrees of ambiguity as the number of compared patterns increases. From the study on input correlation we now know that effective memory will not be affected by the spatial origins of the response patterns encoded in a given network.

²To find the number of inputs with c common cells we compute the binomial coefficient $\binom{25-c}{6-c}$. For example, each input pattern shares 4 common cells with $\binom{25-4}{6-4} - 1 = 210 - 1 = 209$ other inputs.

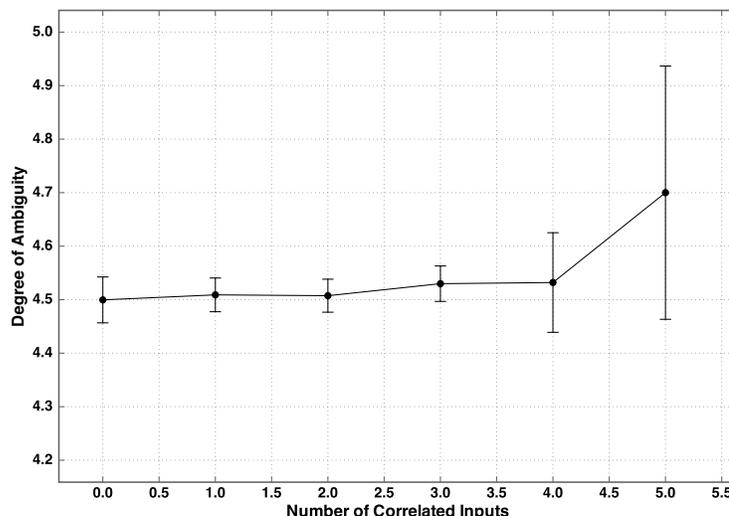


Figure 5.22: Response correlation with respect to input correlation averaged across 10 random DCD networks. This reveals the degree of ambiguity between any two response patterns for inputs that have no identical input cells up to 5 identical input cells. Bars indicate standard deviations across the 10 trials. The number of comparisons with 5 identical cells is much less likely than with fewer common input cells.

5.2.2 Effective Memory

Up to this point, we have seen how DCD networks demonstrate the ability to respond diversely in a way that remains independent of the amount of similarity between input patterns. However, as the number of patterns elicited from a network increases, the amount of ambiguity between any two patterns increases. As seen in Section 4.4, a decoding network should have an increased decoding threshold in order to maintain the ability to distinguish between patterns. What is not clear, however, is how this degree of ambiguity increases as the number of patterns to be decoded increases. We investigate this in the following.

5.2.2.1 Spatiotemporal States

DCD networks which behave with no more than a single spike per cell over a period λ can occupy very large state spaces. The total number of possible spiking states is $(\lambda - 1)^N$ (in the discrete condition), where $\lambda \in \mathbb{N}$ for a set of N spiking neurones with limit-cycle period λ during which a single spike per cell occurs. This leads to $1 * 10^{25}$ states with a λ period of 11 and 25 neurones. However, one should note that for each of these combinations there will be $\lambda - 1$ patterns that exist that are

merely offsets of one another. Thus, despite being distinct with respect to spiking onset, decoding of the patterns will be fully ambiguous. Hence the upper bound number of distinctly decodable spatiotemporal patterns is $(\lambda - 1)^{N-1}$ (here $1 * 10^{24}$). This, however, expresses the maximum number of uniquely identifiable spiking states for a network of size N without taking into account limitations set by connectivity, delay, thresholds, input conditions, and other factors by which real world networks are constrained. In particular, the present investigation continues to explore the case where 6 neurones are stimulated synchronously for input to a population of 25 neurones. To find the total number of possible input patterns we compute the binomial coefficient $\binom{25}{6}$ which leads to 177,100 combinations of input patterns that can result in self sustained activity for this input condition. Within this reduced set, we anticipate that a number of overlapping inner-deltas between neurones as well as the limits imposed by propagation delay ranges and refractory periods will further reduce the memory capacity of DCD networks. We evaluate this effective capacity in the following.

5.2.2.2 Empirical Evaluation of Effective Memory

Definition. *The effective memory of a network is the maximum number of spatiotemporal spike patterns that it can produce that are differentiable.*

Importantly, the maximum number of patterns that are differentiable in DCD networks has an upper bound set by the maximum degree of ambiguity Λ_{max} when Λ_{max} reaches the number of neurones in a network minus one, i.e., $\Lambda_{max} = N - 1$. Beyond this, it cannot be guaranteed that any two patterns can be differentiated. This definition of effective memory is broad in scope and fails to distinguish between the types of network response patterns including dissipative responses, limit cycles, and chaotic attractors. However, dissipative responses present poor dynamic memory qualities while the possibility of decodability of chaotic attractors remains unknown. Hence we restrain the task of estimating effective memory to network and input conditions which lead to sustained limit cycle attractors.

Despite this, an exhaustive test of the effective memory of a DCD network would require that we apply all possible inputs which can be stored across all N cells and which span a time period as long as the longest propagation delay δ_{max} in the network. This, however, would result in testing $2^{N \times \delta_{max}}$ possible input conditions. Considering that networks smaller than 25 neurones were found to be insufficiently large to produce sustained limit-cycles of complex firing (c.f. Section 3.2), an exhaustive test of all

significant inputs remains computationally infeasible.

Instead, to gain an appreciation of the memory capacity of these networks, we can take advantage of the fact presented in Section 5.2.1 that a restrained input set does not, in turn, restrain the qualities of the output response patterns. Hence, comparing responses from a DCD network for inputs to only 6 neurones which receive synchronous pulses, should lead to no more correlation than the comparison of any two other patterns within the $2^{N \times \delta_{max}}$ input space. Investigating the differentiability of response patterns for this limited set of input combinations should thus be representative of the effective memory of DCD networks. For this reason, to obtain an average estimate of the effective memory of DCD networks, we computed network responses for all possible input combinations to 6 neurones in the form of a synchronous input pulse, i.e., $\binom{25}{6} = 177,100$. We then incrementally compared a new pattern (randomly selected from the pool of responses) to each of the previously compared patterns – i.e., comparison without replacement. The maximum degree of ambiguity Λ_{max} found between this pattern and the pool is then recorded. Due to the high computational cost of ambiguity analysis for such high numbers of responses, we limited the final number of patterns compared to 88,550 (half the total number of responses). By using this incremental process we can reveal how likely it is for new patterns to conflict by an ambiguity of a given degree with other patterns. Alternatively, one could simply re-sample from all possible spatiotemporal responses so as to form increasingly large sets of patterns and compute the maximum degree of ambiguity found – i.e., comparison with replacement. This approach, however, would only reveal the worst case conditions of ambiguity since only the two most ambiguous patterns would count. While this is a stricter measure of effective memory, we opt for the former relaxed version which reveals the natural capacity of DCD networks to encode input patterns as the number of stimuli required to differentiate increases.

For statistical soundness, we ran this incremental ambiguity analysis as the number of patterns compared increased over a total of 30 random networks for which we each obtained the 177,100 input responses of 3000 time steps. This population size was determined by computing the average of the maximum degree of ambiguity found for an increasing number of networks. Convergence was considered to have occurred when this average varied by less than 0.1 of a degree of ambiguity. This was obtained when at least 23 networks were analyzed. For increased confidence in the result statistics we simulated a total of 30 random networks.

Figure 5.23 shows the average decoding threshold required to uniquely identify each pattern within a subset as the number of patterns per subset increases. Each

point represents the averaged degree of ambiguity among the given number of patterns compared over 30 random DCD networks for which responses were simulated. This is plotted as a log-linear relationship showing how the average maximum Λ increases in a logarithmic fashion with respect to an increase in the number of patterns to decode. Because only sets of up to 88,550 patterns were compared, the average Λ_{max} only reached an average of approximately 13. This suggests that substantially larger sets of patterns could be retrieved uniquely. To estimate this effective memory limit, we performed a regression to project the maximum memory capacity when $\Lambda = N - 1$ is reached, i.e., when only a single neurone remains available to disambiguate between patterns. To perform this regression we used the following base function model:

$$\Lambda_{max} = a * \log_e(x * b) + c \quad (5.5)$$

Where x is the number of patterns to differentiate. Identified values for a , b and c were 0.684, 2.240, and 4.655 respectively. Figure 5.23 shows the fitted curve as a straight line in a log-linear coordinate mapping. Importantly, we found that an average Λ_{max} of 24 would be reached when approximately 10^{12} patterns are retrieved from the networks; in which case each would remain uniquely identifiable under average conditions. While this represents much less than the maximum theoretical limit of 10^{24} patterns, this memory capacity remains rather large given the relatively small number of neurones. Furthermore, this projected effective memory limit is much greater than the total of $1.771 * 10^5$ responses that can be uniquely elicited by all combinations of synchronous input to 6 neurones among 25 cells. This suggests that DCD networks expand input patterns within their structure in a manner that efficiently exploits both spatial and temporal dimensions for response generation and produce less overlap than what is present in the input. Indeed, this was previously reflected by the independence between input responses despite high levels of input correlation seen in Section 5.2.1.2.

The large projected effective memory based on regression, however, ignores the increasing variability found as the number of patterns compared increased. This increase in variability is also reflected in Figure 5.23 by the broadening of the averaged Λ_{max} as the number of patterns increases about the regression curve. This finding is a direct result from the incremental approach to ambiguity analysis which we accomplished by comparing new patterns to the previously analyzed pool of patterns – i.e., comparison without replacement. If the re-sampling approach was adopted, however, convergence towards higher degrees of ambiguity would be recorded. To capture the

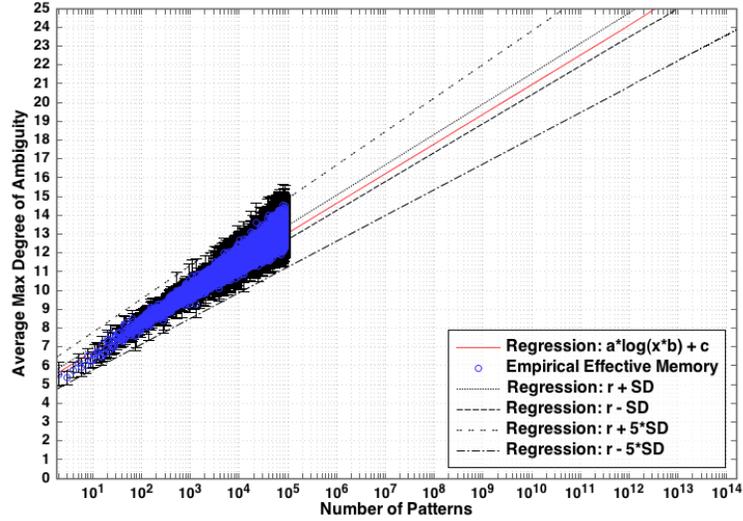


Figure 5.23: Average maximum degree of ambiguity over 30 random networks of 25 neurones as the number of patterns retrieved increases up to a total of 88,550 patterns. Vertical error bars are standard deviations of the ambiguities. The solid line corresponds to the natural logarithmic regression with coefficients a , b and c . Other lines correspond to the regression of the ambiguities plus or minus multiples of the standard deviations.

amount of variability and gain more accurate estimates of both pessimistic and optimistic estimates of effective memory, we computed regressions over the ambiguity points to which was added or subtracted multiples of the standard deviations. Figure 5.23 shows this for the average maximum degree of ambiguity plus one standard deviation ($r + SD$: pessimistic), plus 5 standard deviations ($r + 5 * SD$: very pessimistic), minus one standard deviation ($r - SD$: optimistic), and minus 5 standard deviations ($r - 5 * SD$: very optimistic). The very pessimistic regression led to an estimated pattern differentiation capacity just above 10^{10} , while a capacity of over 10^{14} differentiable patterns is expected to be reachable in the very optimistic condition. Interestingly, effective memory for networks possessing different characteristics could be more systematically compared with respect to both their estimated effective memory from regression but also in terms of their optimistic and pessimistic capacity from regression on multiples of their standard deviations. This could help reveal important distinctive features with respect to capacity and reliability under different cellular, network, or input conditions.

To further convince ourselves of the substantial dynamic memory capacity of these networks we conducted a scaled experiment by simulating DCD networks of 50 neurones. Similar to the previous experiment, a total of 30 networks were randomly

generated and tested on a large number of distinct input conditions. The connection probability was reduced to 50% to preserve the average of 24 afferent connections per cell. Network input stimuli were doubled to 12 neurones stimulated synchronously. However, due to the much larger number of possible input conditions ($\binom{50}{12} > 10^{11}$) we recorded only 88,550 responses from independent input patterns. Figure 5.24 shows a memory capacity above 10^{16} patterns that could potentially be encoded as predicted from the logarithmic regression when $\Lambda_{max} = 49$. Regression was computed using the same regression function (5.5) but with coefficient values $a = 1.128$, $b = 1.095$, $c = 6.717$. Hence doubling the population of cells from 25 to 50 led to an increase by 4 degrees of magnitude in memory capacity. However, by examining the regressions resulting from multiples of the standard deviation added or subtracted from the average maximum degrees of ambiguity, we found that the projected optimistic and pessimistic estimates varied more markedly in these larger networks. In particular, by computing regression in a very pessimistic condition ($r + 5 * SD$) we found that memory capacity was reduced to 10^{11} resulting in a decrease of capacity by 5 degrees of magnitude. However, for a very optimistic regression estimate ($r - 5 * SD$) the projected capacity reached over 10^{40} differentiable patterns. This represents an increased capacity of over 24 degrees of magnitude.

The noticeable increase in variability for this population of larger networks could stem from different sources. Naturally, the small number of inputs on which the networks of 50 neurones were tested in proportion to the possible input combinations could have led to less accurate regression estimates. However, the small number of response trials cannot justify a relatively earlier onset of greater variability than in the 25 neurone network evaluation. It is also possible that the ratio of connections to the delay range was much greater in 50 neurone networks resulting in greater ambiguity in spiking among afferent connections. Indeed, in the 50 neurone networks, a total of 1225 connections were established while only 600 connections existed in the networks of 25 cells. However, networks of both sizes had transmission delays spanning over 20 time steps. This ratio could compete with the effect of a larger number of potential spatiotemporal patterns for larger networks, which alone would lead to a lower chance of ambiguity. The combined effect could thus lead to networks responding at times in substantially different ways, and on other occasions rather similarly. We leave the opportunity of confirming this for future work.

By examining their possible behaviour under non-noisy conditions DCD networks thus offer a very high capacity for input differentiation and retrieval via delay mapping. The results here, however, are averaged across multiple runs which does not fully

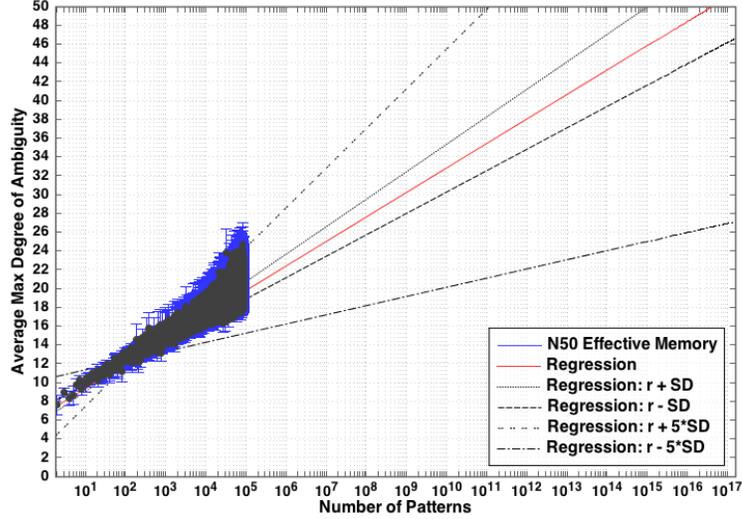


Figure 5.24: Average maximum degree of ambiguity over 30 random networks of 50 neurones as the number of patterns retrieved increases up to a total of 88,550 patterns. Vertical bars show standard deviations. The solid line corresponds to the natural logarithmic regression with coefficients a , b and c . Other lines correspond to the regression of the ambiguities plus or minus multiples of the standard deviations.

reveal the features of these networks when investigated under individual conditions. Furthermore, it would be beneficial to gain insight with respect to the probabilities that storing a novel pattern will cause successful or failure of retrieval. We explore these in the following.

5.2.2.3 Probability of Failure

Definition. *Failure in differentiability* occurs when two separate inputs result in identical spatiotemporal responses in DCD networks such that no sub-pattern features exist in order to distinguish between patterns.

What the previous averaging over a population of networks fails to reveal is that occasionally two differing spatiotemporal patterns can lead to identical responses such that $\Lambda = N$. This can occur for a number of different reasons. First, if a few input neurones possess less effective delay connectivity this can limit the coincidence onto target cells when paired with other input cells. However, if both the remaining subsets of effective input cells are identical then an identical response will occur. Naturally, failure to differentiate between dissimilar input responses can also occur if delay and connectivity from both input pattern cells are identical. For a network where neurones have a mean number of afferents I per cell and a delay range D_{range}

time steps then the probability for a neurone n to have any given afferent delay configuration is $P(n_{conn}) = \frac{1}{D_{range}}$. In the networks under investigation each neurone projects to all other 24 cells with uniform random delays ranging between 20 and 40 time steps thus leading to the probability of $\frac{1}{20^{24}}$ that a neurone's projections have a given delay structure. While this represents a very low probability, it is, in fact, not necessary that all afferent connections possess complementary delay features because only a sufficient subset of coincidentally effective spikes will matter. More specifically, when the coincidence detection threshold of cells in the network is low, the chances that two distinct neurones produce similar spatiotemporal responses increases because only a subset of their connectivity needs to match for this to take place.

To determine the likelihood of complete failure in the differentiation between patterns we counted the number of times new patterns caused a degree of ambiguity $\Lambda = 25$ for every 10,000 patterns explored up to the limit of 88,550 patterns. We then averaged these counts across all 30 networks to gain a more accurate estimate of this likelihood. Figure 5.25 shows an increase in the probability of failure starting with only a 0.05% chance of having two identical responses for different inputs when up to 10,000 patterns are stored. This increases progressively up to 0.92% chance when a total of 88,550 stimuli are stored. Hence, after half of the possible input patterns have been compared, less than 815 of them will be non-differentiable. This provides for a rather good performance when the criteria for failure is set to the highest degree of ambiguity $\Lambda = N$. As seen from the standard deviations for each bracket, variability increases by one order of magnitude by the time half the possible input patterns are stored. Albeit small, this suggests that the confidence of successful decoding decreases when very large numbers of patterns are elicited by peripheral networks.

As the probability of failure suggests, recording Λ_{max} alone does not reveal the entire potential of DCD networks and may be insufficient as a practical measure of memory performance. While it does show how the differentiation between any two patterns becomes increasingly difficult as the number of patterns to be recalled increases, it do not reveal the number of patterns that are potentially ambiguous with eachother. To do so we introduce the notion of the *weight of the degree of ambiguity* (weight of Λ) in the following.

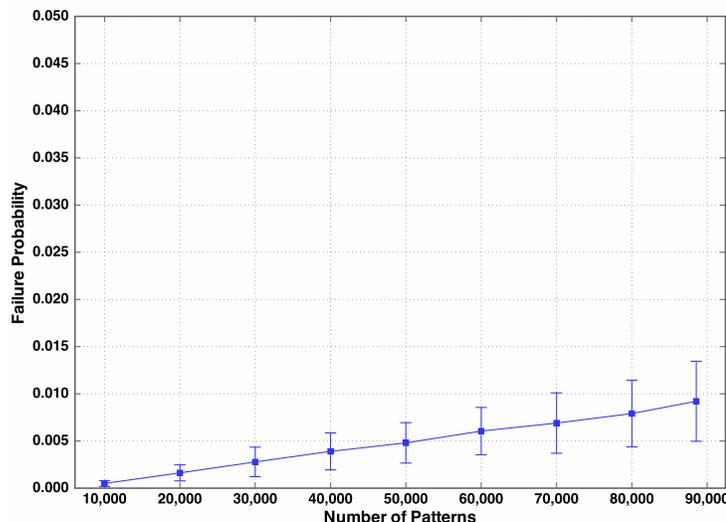


Figure 5.25: Failure probability (0 to 1) averaged over the 30 networks. This measure shows an increase in the probability that two separate input patterns produce fully ambiguous results, i.e., $\Lambda = 25$. Vertical bars indicate standard deviations.

5.2.2.4 Weight of Degree of Ambiguity

Unfortunately, counting the highest degree of ambiguity between any two patterns and the likelihood measure of failure shown above does not distinguish between large sets and small sets of ambiguity relations. Because the highest degree of ambiguity measure is based on the comparison between a newly added pattern and one from the pool, only two patterns are used to expose the worst case of ambiguity in a growing set of elicited responses. However, worst case faulty recall between two patterns may not be as detrimental to a memory system as partially faulty recall between a larger number of patterns. Hence, the condition for failure could be relaxed with respect to the degree of ambiguity but enhanced with respect to the minimum number of patterns that will produce ambiguous responses. We denote this instead as failure for a given ‘*pattern weight*’.

Definition. The *weight of ambiguity* for a given degree Λ , denoted $w_p(\Lambda)$ for a network, is the number p of response patterns, from some set, that share a degree of ambiguity of Λ for any given subset of cells.

If a particular degree of ambiguity Λ is present across x response patterns for the same neurones then the weight of that degree of ambiguity is x for that particular network and set of cells. For example, if out of 26 patterns A to Z (i.e., $p = 26$), only patterns A, B, and C share an ambiguity of degree 5 for cells $\{3, 11, 5, 8, 17\}$ then

$w_{26}(5) = 3$. However, if patterns A and B share an ambiguity of degree 5 on cells $\{3, 11, 5, 8, 17\}$ and patterns C and D share an ambiguity of degree 5 on cells $\{2, 12, 6, 4, 9\}$ then overall $w_{26}(5) = 2$, despite a degree of ambiguity of 5 being present across a total of 4 patterns, those patterns do not share the same ambiguous set of cells. To account for all weights of ambiguity of degree 5, irrespective of the set of cells involved, we can sum all patterns involved to provide a quantity which we call the *total weight of ambiguity* denoted $W_p(\Lambda)$. In our example with A and B sharing a degree of ambiguity of 5, and C and D sharing a degree of ambiguity of 5, we thus have a $W_{26}(5) = 4$ for all four patterns A, B, C, and D.

To empirically evaluate the weight of ambiguity in our Standard DCD network model, we computed the total weight of ambiguity W_p for three response sets of size $p = 10$, $p = 100$, and $p = 1000$. We repeated this measure across 30 random Standard DCD network models for which 10,000 response patterns were collect each. Importantly, cells that participated in higher degrees of ambiguity were not counted again for lower degrees of ambiguity. That is, if cells 5, 2, 8 formed a degree of ambiguity of 3 amongst p patterns, then degrees of ambiguity of 2 with cell pairs 5, 2, 5, 8 or 2, 8 were not counted.

Figure 5.26 shows the average total weights of ambiguities for each of these pattern subset sizes. Figure 5.26.a shows that an average $W_{10}(2)$ of 308.6 (SD 25.33); this quickly decreases for a degree of ambiguity of 3 to $W_{10}(3) = 163$ (SD 15.47); and $W_{10}(4) = 71.3$ (SD 9.06) with an average total weight of ambiguity below 30 for degrees greater than 4. When 100 response patterns are compared for weight of ambiguity analysis as seen in Figure 5.26.b, a degree of 2 led to an average total weight of ambiguity of $W_{100}(2) = 17,871.23$ (SD 383.89); for a degree of 3 we found a proportionally smaller reduction in the 100 responses case with $W_{100}(3) = 16,434.3$ (SD 416.188); for a degree of 4 we obtained $W_{100}(4) = 7,815.63$ (SD 254.61) and average total weights of ambiguity below 3,200 for degrees 5 and above. Finally, when 1000 response patterns were compared as seen in Figure 5.26.c, the average total weight of ambiguity for degree 2 was $W_{1000}(2) = 290,113.93$ (SD 1,307.42); interestingly for a degree of 3 this weight was greater than for degrees of 2 with $W_{1000}(3) = 988,044.17$ (SD 12,893.99); this weight then began to diminish for degree 4 with $W_{1000}(4) = 704,909.83$ (SD 15,859.05); weights then continued to diminish for higher degrees of ambiguity.

The greater degree of ambiguity found for $W_{1000}(3)$ than for $W_{1000}(2)$ can be explained by the fact that when greater numbers of response patterns are compared the likelihood for patterns of greater degrees of ambiguity to be found across patterns

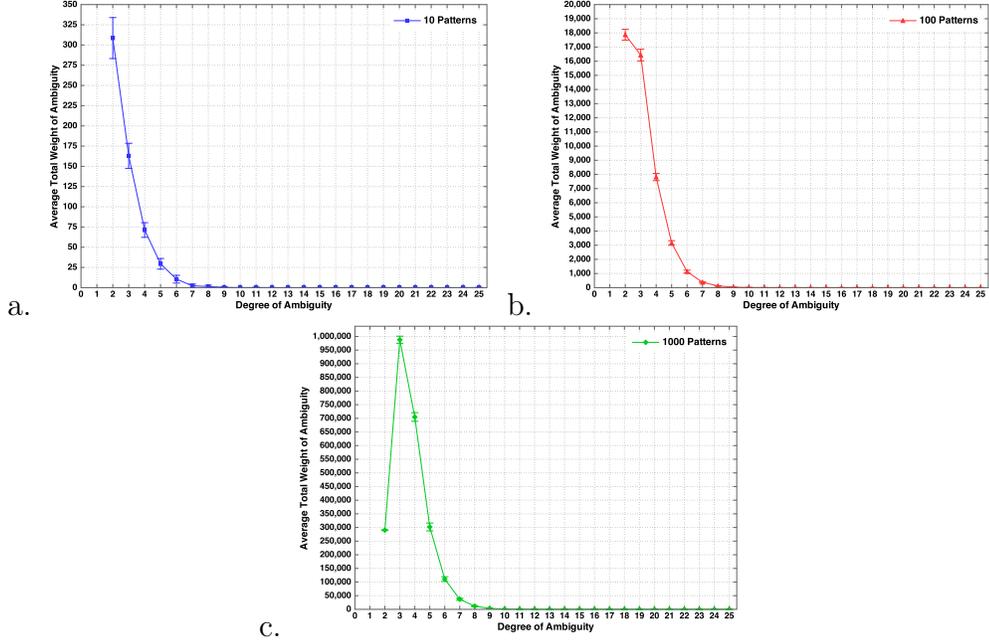


Figure 5.26: Averaged total proportional weights of degree of ambiguity amongst three different sizes of response subsets: 10, 100, and 1000 patterns.

is also greater, yet because cells that participate in ambiguities of higher degree are not counted for lower degrees of ambiguity, the predominant degree of ambiguity will be exposed when the total weight of ambiguity is measured. In this case, the weight $W_{1000}(3)$ was predominant over $W_{1000}(2)$ but also over degrees of ambiguity 4 and above. When looking strictly at the weight of ambiguity (and not the total) for the set of cells which spanned the largest number of response patterns, we found that $\max(w_{1000}(2)) = 982.93$ meaning that almost all but approximately 18 patterns on average had an ambiguity of degree 2 for a specific pair of cells, while $\max(w_{1000}(3)) = 535.07$ meaning that only approximately half the response patterns had an ambiguity for a specific triplet of cells. It is important to note, however, that $w_p(\Lambda)$ fails to capture all the subsets of cells that form ambiguous groups which is why we opt to use the total weight of ambiguity $W_p(\Lambda)$ as the better medium to quantify the predominance some given amount of ambiguity Λ .

Overall, a sub-linear relationship between degree of ambiguity and total weight of ambiguity appears to be the trend with respect DCD network response dynamics in our Standard model. However, we also found that when comparisons are made across larger polls of network responses, a total weights of ambiguity for larger degree can be more predominant than for smaller degrees of ambiguity when we prevent constitutive subsets of cells to count as ambiguities of those subset sizes. From a functional

viewpoint, this total weight of ambiguity measure can serve as an indicator for the pattern sizes that are most likely to be decodable from a network when this network is sollicitated for p number of patterns. In particular, when $W_p(\Lambda)$ is predominant for that p then patterns of size greater than Λ should be used for information coding. Stricter measures may require even greater pattern sizes to be used is the risk for ambiguous decoding is to be minimized at the cost of requiring higher coincidence detection thresholds and/or reduced synaptic strengths.

5.2.3 Effective Memory and Noise Interaction

5.2.3.1 Estimation of Effective Memory with Noise

As seen in Section 5.1 on noise tolerance, the interference of a single neurone in a DCD network of only 25 cells can lead to the loss of just two or three neurones when that cell misfires to more than six neurones when that cell is subjected to prolonged spike removal. This previous analysis revealed that stochastic effects on a few cells can propagate to other neurones in the network leading to higher degrees of variation than would be initially measured from noise sources alone. Presently, we consider how decoding remains feasible when a subset of neurones behave stochastically in combination with increasing pattern storage demand from having to differentiate between multiple responses.

In the most direct approach, we could measure the impact of noise on the effective memory of a network by reproducing noise simulations in combination with input differentiation tests. Measuring the limits of capacity via ambiguity analysis, however, could not be performed simply by directly measuring the increase in the degree of ambiguity between responses from alternate input patterns. This is because as noise increases, the degree of ambiguity between different patterns would not increase but in fact decrease due to the alternate spiking regime reached by the DCD network. Since increasing the number of patterns to decode promotes the likelihood of ambiguity, noise and effective memory result in competing behaviours. Thus, measuring the impact of noise on effective memory would have to be done in comparison to the non noisy memory load condition so that a measurement of the variation between both noisy and non noisy conditions could be made.

To avoid the computational and analytical challenges imposed by this approach, we instead propose to measure the effect of noise on memory capacity by taking the maximum degree of ambiguity obtained in both noise and effective memory scenarios as a unit of reference. We can readily apply this method using the independent

ambiguity results for both noise and effective memory conditions already obtained. If the assumption that maximum degree of ambiguity is a sufficient marker of behaviour in both noise and effective memory then this approach should provide valuable insight as to how the noise-to-effective memory relationship impacts the performance of DCD networks. However, this approach makes the additional assumption that the impact of noise on the neural behaviour within the same network is equivalent, with respect to disruption, to the impact noise would have when comparing across separate network responses. We assume in the following that this equivalency can be made.

We should recall from the noise analysis in Section 5.1 that as the number of noisy cells increased the degree of self-ambiguity *decreased*. However, a decrease in self-ambiguity suggests that the presence of any between pattern ambiguity should also be diminished. Yet, ambiguity between patterns was measured for effective memory which increased as the requirement for storage grew. Interestingly, self-ambiguity was maximum when no cells received noisy input, similarly ambiguity between patterns is maximum when the effective memory reaches its maximum capacity, i.e., the highest degree of ambiguity is reached. This suggests that a relationship between self-ambiguity and ambiguity for memory load can be made.

To relate noise and effective memory we take the inverse of the regression function for the progress of effective memory as additional patterns are compared (see Equation 5.5). By using the inverse function we can map the effect of increasing noise for each of the four conditions (c.f. Section 5.1) on the number of response patterns that can be distinguished and therefore retrieved. This inverse function identifies the number of patterns x that are differentiable given the degree of ambiguity present in a population of noisy responses according to equation:

$$x = \frac{e^{\frac{\Lambda-c}{a}}}{b} \quad (5.6)$$

Where a , b , and c are the regression function coefficients found in Section 5.2. Hence by applying the Λ found as the number of a noisy neurones increases in the 25 neurone model from the four noise conditions considered in Section 5.1 including spike removal in the prolonged and temporary conditions as well as spike misfiring in the prolonged and temporary conditions we obtain a relationship between noise and effective memory as illustrated in Figure 5.27. By plotting the number of patterns that can be differentiated on a log scale we obtain a clear mapping from noise to effective memory. From this figure we notice that when no noise is applied the effective memory is expected to reach its maximum capacity of a little over 10^{12} differentiable

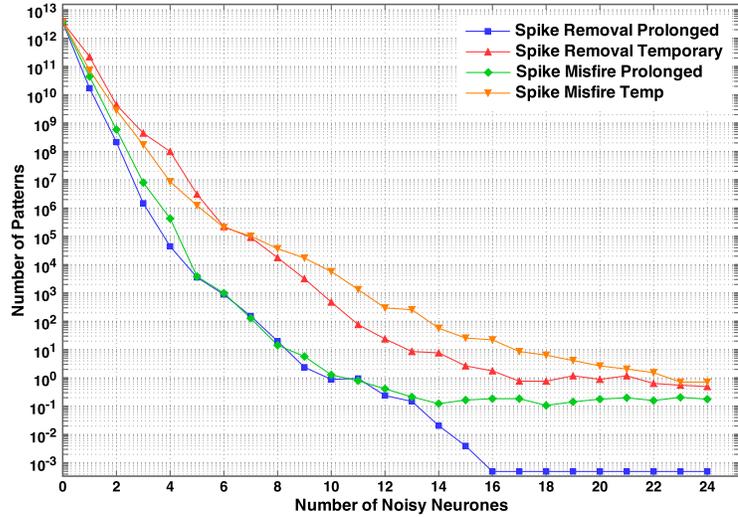


Figure 5.27: Relationship between noise and number of identifiable patterns in DCD networks of 25 neurones based on independent analysis of noise and effective memory.

patterns. However, as the number of noisy neurones increases a rapid depreciation in effective memory is expected to occur with worse case scenarios taking place when prolonged noise is applied. For the prolonged noise conditions (removal and misfiring) this method predicts that only a single pattern and pattern features remain available after noise is applied to 10 or more neurones. Less dramatically, this method projects that temporary noise will lead to a greater number of differentiable patterns as the amount noise increases.

Arguably, this method makes it uncertain if other factors will affect memory capacity when multiple noisy responses are compared. If a substantial interaction would, in fact, occur then the current method will produce an incorrect evaluation of how noise impacts effective memory capacity. As mentioned earlier, determining the potential of this impact could be assayed by conducting simulations which include noise while testing effective memory a developing a measure of comparison.

While this approach might provide valuable insight in the relationship between memory capacity and noise in DCD networks, it is important to realize the competing effect that these two conditions have on the performance of DCD networks. We explore this in greater detail in the following.

5.2.3.2 The Decoding Threshold Tradeoff

We saw in Section 5.1 that by decreasing the coincidence detections threshold of decoding neurones (i.e., the decoding threshold) patterns which were previously am-

biguous could be differentiated once more. Hence a decrease in decoding threshold provides greater tolerance to noise. Conversely, as a larger number of patterns need to be differentiated the degree of ambiguity increases leading to the requirement for higher decoding threshold $T = \Lambda_{max} + 1$ in order to maintain differentiability of the spatiotemporal responses. If DCD network dynamics are indeed found to be a significant coding mechanism for dynamic memory in the brain, then this tradeoff between noise tolerance versus encoding capacity may be dynamically regulated by some adaptive, possibly homeostatic, control mechanism. Such a mechanism could manipulate decoding thresholds by increasing or decreasing them towards a sufficient tradeoff based on contextual need. When a tradeoff target is reached then the largest payoff between the number of patterns that can be identified uniquely and noise tolerance is reached. While a number of mechanisms could possibly implement this, this remains material of investigation for future work which we discuss in Section 7.2.

5.3 Summary

In this chapter we first investigated the effect of noise on a Standard DCD network model by exploring analytically the upper and lower bound noise tolerance when ambiguities are present. This method, however, fails to describe how DCD networks cope with noise when taking into account the correlated effects that missing spikes of misfiring cells have on the overall response dynamics. To take account of this, we performed empirical evaluations of the behaviour of these networks under prolonged and temporary spike removal conditions as well as prolonged and temporary misfiring of cells. By looking at temporary versus prolonged conditions for both removal and misfiring we could reveal the minimum effect of noise (temporary scenario) and maximum effect of noise (prolonged scenario). For each of these four conditions, we tested the response as the number of noisy cells increased in order to observe the change in tolerance. In all cases, a sublinear effect on self-ambiguity took place as the number of noisy cells increased, although this was more pronounced in the prolonged scenarios. Importantly, we found that in all conditions, decodability with a sufficiently low decoding threshold T could be achieved when up to approximately 60% of cells were perturbed. In the cases of temporary misfiring, however, some of the original pattern features remained available after all neurones were perturbed only once suggesting that limit cycle attractors display robustness in the face of noise. Furthermore, we observed stability conditions of the response patterns with respect to the type and duration of noise, and found that critical changes in noise tolerance often took place

which was especially common in prolonged noise conditions. We finally compared the behaviour of networks of sizes 25, 100, and 1000 neurones in all four noise conditions to evaluate how this tolerance scales with network size. While scaling revealed that larger networks were proportionally less tolerant to noise, their responses varied less across trials suggesting greater reliability when coping with noise.

The second Section of this chapter investigated the memory capacity of Standard DCD networks empirically. To do so, we began by verifying how closely responses pattern resembled each other when input patterns shared similar features. By computing ambiguity between responses in relation to the number of identical synchronous input spikes we found that DCD network response correlation remains independent of input correlation. With this result, we then continued to explore how ambiguity increased as the number of compared patterns increased. This resulted in a measure of effective memory for DCD networks. Due to the high computational effort required, this empirical evaluation was halted after 50% of all 177,100 possible input related responses were compared. Despite this input evaluation limit, computing regression on the available data revealed that a maximum degree of ambiguity Λ_{max} would reach the number of cells in the network, i.e., the limit for decodability via coincidence detection, when approximately 10^{12} patterns are encoded for networks of only 25 neurones. This significant memory capacity, however, was also accompanied by increasing variability in capacity as the number of patterns to store increases. Interestingly, while capacity increased to in networks of 50 neurones to 10^{16} possible patterns, variability increased even further than in the 25 neurone network. While effective memory measures the absolute capacity of the network such that any two patterns will remain differentiable on average, it fails to capture the prevalence in responses for ambiguities of a particular degree. We define this as the weight of degree of ambiguity. First, this evaluation provided a direct measure of how common an ambiguity of particular degree is in a set of response patterns of a given size. Furthermore, this measure indicated that as the number of patterns increased the prevalence of lower degrees of ambiguity reduced while the prevalence of higher degrees of ambiguity increased. While this should be expected, this weight measure provides insight with respect to how likely two patterns are to be ambiguous for a particular Λ .

Finally, this chapter investigated the relationship between effective memory and noise. To do so, we performed a direct mapping from the noise tolerance estimated results in each of the four conditions found in Section 5.1 to the effective memory regression from Equation 5.5 using the ambiguity value. While an empirical evaluation of the effect of noise on effective memory capacity could reveal greater insights with

respect to this relationship, the proposed mapping revealed an encoding capacity proportional to noise when applied to self-similarity ambiguity evaluation.

Chapter 6

Formal Investigations

In previous chapters we have shown that DCD networks possess the capacity to not only produce diverse dynamics under various connectivity and neuronal parameters settings such as synchronous firing, chaotic firing and stable limit cycle dynamics, but we have also shown functional applications of these networks for fundamental cognitive processes including reactive behavioural control and dynamic memory. Furthermore, we have established measurements of important informational features such as effective memory, weight of ambiguity and tolerance to noise. These contributions, however remain primarily seated in empirical simulation results. While scaling analyses for both noise and dynamic memory enable some predictive perspective of DCD network performance on yet unforeseen conditions, greater insight into the true potential of these networks of coincidence detectors with transmission delays should be achievable through more in depth characterization of structural and functional attributes of these networks. Even further, functionally effective features of DCD networks, such as spiking regimes or the number of firing-chains, could be deduced using analytical methods that only require structural characteristics. While we have shown that simple knowledge about the delay range in connectivity can lead to differing spiking dynamics (see Section 3.2.3), we aim here to further characterize the amount of activity in networks and their number of states analytically by using structural properties and as few dynamical properties as possible. By doing so, a greater understanding of the overall behaviour and functional efficacy of DCD networks for cognition should be gained.

To accomplish this, we build on previous work that has formally investigated the roles of connectivity, firing probability and noise in neural networks and other network systems in general. We first present a method for estimating the amount of spiking activity in the network at any given time step in discrete DCD networks

provided knowledge of the input and model characteristics. We then develop formal components which we combine to form a general probabilistic model for the number of spiking states in the network. From this, we then introduce constructive steps towards estimating the number of possible firing-chains within DCD networks.

6.1 Estimating Network Activity

To begin our formal investigation we present a method to estimate the global activity A_t of a network, where A_t is the number of neurones firing at time step t . We aim to do so provided the initial amount of activity A_0 , the average degree of incidence to neurones I (the number of afferent connections to any given cell), the coincidence detection threshold T (the number of simultaneous spikes required to reach a target cell at time t and induce a spike) and the number of neurones N in a DCD network.

We first look at the probability that sufficient source cells to a target cell c have fired. We also begin by assuming that at least T source cells are active and that propagation delays are all equal to 1 time step. Given this, there is a $\frac{1}{\sum_{i=T}^I \binom{N}{i}}$ chance that a set of at least T connections among a maximum of I incident connections from any possible N source cells will be effective at driving c , where $\binom{N}{i}$ is N choose i. However, if there are more than one active cells then there will be a greater chance for the target to spike. In particular, there are $\binom{A}{T}$ combinations of exactly T cells firing among A active cells that will be effective at triggering a target to spike. However, a cell can receive exactly I afferents, such that $I \geq T$. Hence, the total number of combinations of potentially effective active cells is $\sum_{i=T}^I \binom{A}{i}$. Furthermore, there can be fewer than I cells active such that $A < I$, in this case the number of effective active cell combination is $\sum_{i=T}^A \binom{A}{i}$. Hence, we define $K = I$ if $A \geq I$, and $K = A$ if $I > A$. Using K , we then define the number of effective active cell combination as $\sum_{i=T}^K \binom{A}{i}$. We can now define the probability function of spiking based on the number of active neurones as:

$$\gamma(N, I, T, A) = \frac{\sum_{i=T}^K \binom{A}{i}}{\sum_{i=T}^I \binom{N}{i}} \quad (6.1)$$

This equation can be used to estimate the probability of firing of a target if all input is synchronous and delays on the afferent connections are equal. For DCD

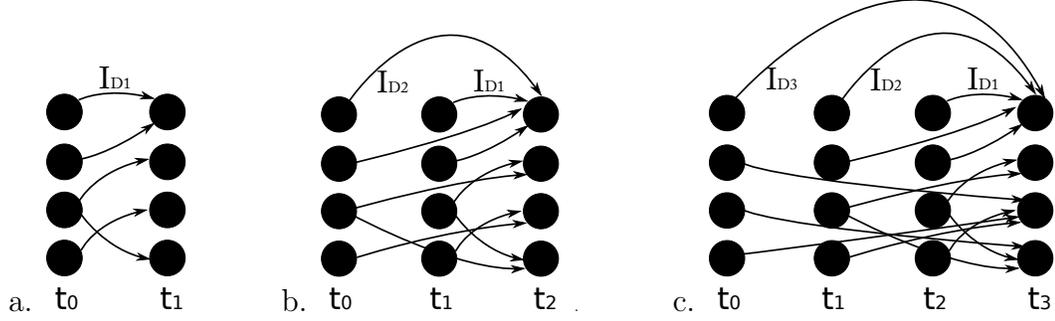


Figure 6.1: Illustrated delay of incidence for 1, 2, and 3 time steps. Illustration a. shows that a_1 (Eq. 6.3) can be computed when the proportional activity a_0 is known at time step 0 and taking into account delays of one time step I_{D1} . Illustration b. shows how the proportional activity a_2 (Eq. 6.4) is dependent on delays of incidence I_{D1} and I_{D2} . Illustration c. shows how activity from incident connections delays with 1, 2, and 3 time steps away are used to compute a_3 (Eq. 6.5). Note: the number of incident connections is arbitrarily illustrated here as we assume low variability between neurones in the present activity estimation.

networks, however, this is not the case. Instead, we must take into account two other important temporal factors: propagation delay and refractory periods of cells. We begin by integrating propagation delays with the following observation: the number of incidents I with a particular delay D will vary. We define the number of incident connections with a delay of a single time step as I_{D1} , with delays of two time steps as I_{D2} , etc.

Assuming that all cells have an identical coincidence threshold T we can estimate the proportion of active cells a_n (where $a = \frac{A}{N}$) for time step n by iteratively summing the probabilities of active cells from previous time steps in the manner illustrated below:

$$a_0 = \frac{A_0}{N} \tag{6.2}$$

$$a_1 = \gamma(N, I_{D1}, T, a_0 \times N) \tag{6.3}$$

$$a_2 \times 2 = \gamma(N, I_{D2}, T, a_0 \times N) + \gamma(N, I_{D1}, T, a_1 \times N) \tag{6.4}$$

$$a_3 \times 3 = \gamma(N, I_{D3}, T, a_0 \times N) + \gamma(N, I_{D2}, T, a_1 \times N) + \gamma(N, I_{D1}, T, a_2 \times N) \tag{6.5}$$

$$a_n \times n = \gamma(N, I_{Dn}, T, a_0 \times N) + \dots + \gamma(N, I_{D1}, T, a_{n-1} \times N) \tag{6.6}$$

Here we have the initial proportional activity described by Equation 6.2 and assume no spikes are currently propagating towards target cells in the network (no

historical spiking). Furthermore, by dividing the right hand side of this sum by the number of steps evaluated n then we capture the final proportion of activity for that time step n . This summation is possible because we assume independence of the activity of the events over time. However, summing the activity should stop once the longest delay in the network has been reached, i.e., when $n = D_{max}$. In order to measure activity at a time step n beyond the maximum delay (where we designate D_{max} as m for short), we can express the proportion of active cells as:

$$a_n \times m = \sum_{i=n-m}^n \gamma(N, I_{D_{n-i+1}}, T, a_{i-1} \times N), \text{ for } m < n \quad (6.7)$$

Equation 6.7 takes into account connectivity (degree of incidence), number of cells, coincidence threshold, and transmission delays. By taking into account m as the maximum delay, we recognize that there is a period after which summing past activity will no longer change the proportion of activity of the population at time n . However, this equation still ignores a more subtle feature of biological spiking cells: their refractory period. As mentioned earlier, the refractory period of cells limits the actual amount of effective incident input. Indeed, this prevents neurones from spiking for a time period after an action potential was just emitted thereby reducing the proportion of cells being active given recent activity in the network.

Because this activation estimate seeks to establish a global activation value without taking into account individual neurone states, we cannot account for the individual refractory period, R , of cells. Instead, our probability of activation estimate can discount current activity by narrowing the window of past activity integration. As an approximation, we can do so by integrating over a recent time window of activity and ignore older activity that would have been suppressed by some original activity. We choose this time window between an offset passed time step o and $n - 1$, such that $o = m - R \times a_{n-m}$, and only if $o > 0$. Hence, if no cells were active at D_{max} time steps in the past, then $o = m$ meaning that all previous time steps of activity propagating on connections are taken into account, but if $a_{n-m} > 0$ then only the most recent activity, if any, will be counted. Thus we can adjust Equation 6.7 to take into account this general refractory period of cells by the following expression:

$$a_n \times m = \sum_{i=n-m}^{n-o} \gamma(N, I_{D_{n-i+1}}, T, a_{i+o-1} \times N), \text{ for } m < n \quad (6.8)$$

Unfortunately, implementing refractoriness in this way enables the model to only take into account refraction after D_{max} time steps have passed. This condition, how-

ever, is less problematic as our general scope is to understand long term dynamics. Despite this, the initial network activity will not reflect realistic activation values because neurones will not be prevented from firing after an action potential was emitted. Furthermore, the accuracy of the refractory effect remains coarse. In particular, it could over compensate for high levels of past activity and prevent maintenance of some base level activation caused by subsets of cells that previously would not have fired. Also, the discrete nature of the offset time step o may accumulate errors for estimations over long periods of time.

Despite these limitations, this model captures the essential components necessary to accommodate for some general effects of connectivity, delays, and refraction in DCD networks with the ability to produce some expectation of the amount of activity over time when actual network simulation is not possible. To show this, we performed numerical evaluation of Equation 6.8 under two familiar conditions: narrow delays and broad delays. We set parameters to reflect the Standard DCD network model as closely as possible such that relevant activity estimates were produced. These parameters were set to 25 neurones ($N = 25$); incidence of 24 cells ($I = 24$); coincidence detection threshold of 2 ($T = 2$); initial activity on 6 cells ($a_0 = 0.24$); refractory period of 40 time steps ($R = 40$); and a minimum delay of 20 time steps. Finally, in the narrow delay condition, the maximum delay was set to 25 time steps, while a maximum delay of 40 was used in the broad delay condition. Hence, all parameters except for the longer refractory period were set to be identical to our original simulation results in Chapter 3 Section 3.2.1. Figure 6.2 shows two plots depicting the analytically estimated activity of a DCD network in both the narrow delay condition (Figure 6.2.a) and broad delay condition (Figure 6.2.b). When proportional activation reached $a = 1.0$ then all network cells spiked. These plots show how the present analytical method expects highly synchronous and regular spiking in the narrow delay condition and more complex and continued firing in the broad delay condition. At a macro behavioural level, these estimates reflect broad simulation dynamics presented in Section 3.2.1.

While this activity estimate method is not as accurate as a network simulation, these results show that some evaluation of general network behaviour properties can be made. However, to gain a more accurate estimation of the spatiotemporal features of DCD networks, we aim to provide a method to estimate analytically the number of distinct firing-chains these networks can produce. In the following section, we develop a distinct set of state probability estimates in order to achieve this as closely as possible.

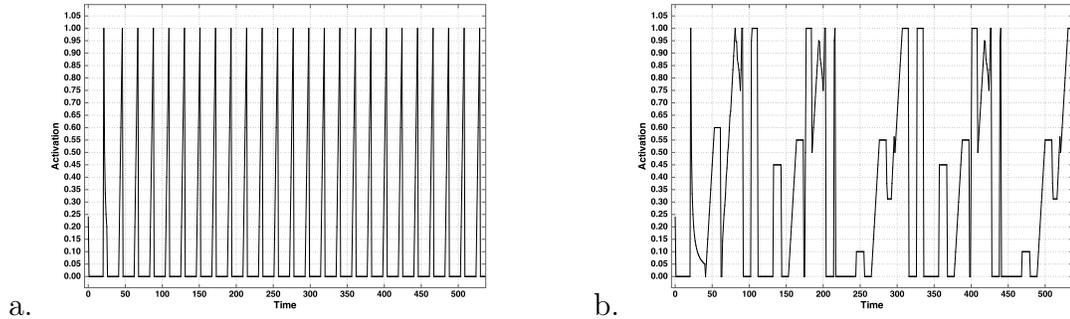


Figure 6.2: Analytical estimate of network activity following Equation 6.8 in the narrow delay condition (plot a.) and the broad delay condition (plot b.).

6.2 Probability of Spiking and the Number of Firing-chains

As discussed in earlier chapters, and in particular in Section 3.3, a characteristic attribute of DCDNs is their ability to produce temporally and spatially sensitive spiking patterns, which we call firing-chains. If we can achieve an accurate estimate of the number of firing-chains possible within a network, then we can estimate, a priori, the number of distinguishable patterns in these networks provided some input condition. If we can further determine how many inputs lead to sustained activity then estimating the number of decodable memories also becomes possible.

Here we determine numerical expressions that can be combined so that a probabilistic model of network states can be established based on network features. Ideally, an analytic solution to the number of patterns a network can express would be described succinctly with sole recourse to structural properties of the network. For DCD networks, these would be: number of neurones, degrees of connectivity, the number of coinciding spikes to reach threshold, membrane and channel properties leading to refractoriness (refractory periods), and distances (propagation delays). For simplification and annotation purposes, we will refer to Θ , in the following, as the set of *average* structural properties of the network and neurones such that $\Theta = \{N, I, T, R, D\}$. Hence if we refer to $I = 24$, for example, we imply that on average and with uniform random distribution each neurone receives 24 incident input connections. Fortunately, all properties except for the state of refraction of cells are independent of the activity in the network. Hence, we begin by introducing probability terms with respect to connectivity and effective spike timing. We then attempt to approximate, eventually independently of activity, the effect of spike refraction in order to estimate the

probability that a neurone will spike. Following this, a general expression is developed, based on this spiking probability estimate, for the number of states that can be reached by a DCD network. Finally, initial contributions are made to determine the bounds for the number of firing-chains available in the network and correlation conditions which could serve as a mapping from state to firing-chain number.

6.2.1 Effective Connectivity

At the outset, we wish to estimate the spiking probability of cells in DCD networks, knowing structural conditions alone. To do so, we begin by estimating the number of spatial input patterns a coincidence detection neurone can detect. By assuming that all afferent connections have no propagation delay, we want to find how many subsets of connections will be large enough to cause a post-synaptic potential, i.e., that connectivity is effective. We denote this number of subsets as the number of spatial input patterns L a cell can detect.

In this case, a neurone requiring T synchronous incident spikes with a total of I afferent connections will be capable of detecting L spatial input patterns on a single input step, such that:

$$L(I, T) = \sum_{i=T}^I \frac{I!}{i! \times (I - i)!} \quad (6.9)$$

Where $L(I, T)$ is the sum of all input combinations between T and I coincident inputs. Hence, as long as one of these spatial input patterns is such that afferents are loaded with synchronous spikes, then the target cell will in turn emit a spike.

As a note of interest, if connection weights are all equal and spikes contribute equal post-synaptic depolarization, the cell's firing threshold becomes linearly proportional to the minimum number of incident spikes T . Alternatively, an integrator neurone could fire given a fraction T/f of the number of incident spikes over up to f time steps making it sensitive to a greater number of possible input conditions but at the cost of occupying f additional time steps. Hence, while an integrator would detect $L(I, T/f)$ input patterns over f time steps, a coincidence detector could detect $L(I, T)^f$ input patterns over the same period of time, making it exponentially more sensitive to rapid variation in the spatial input arrangement of spikes. Interestingly, by applying Equation 6.9 we can conceptually transform the network system of the target cell with I afferents into a system of a single target and a single incident connection depicted

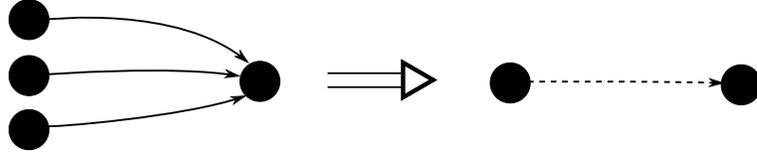


Figure 6.3: The sum of combinations allows to conceptually reduce the number of afferent combinations to a single afferent.

in Figure 6.3.

By making use of Equation 6.9 we can derive the probability s_{Θ} that a target cell will spike based on the probability p_{Θ} that sufficient source spikes are present (and synchronous) on the afferent connections, but without regard for refractory periods. We can do so by applying the cumulative binomial distribution:

$$s_{\Theta}(p_{\Theta}) = \sum_{i=T}^I \binom{I}{i} p_{\Theta}^i (1 - p_{\Theta})^{I-i} \quad (6.10)$$

Here, successful arrival of spikes coincidentally on the target will take place with probability p_{Θ} . Crucially, however Equation 6.10 ensures that coincident spikes count as a success if at least T and up to I afferents reach the target synchronously. Indeed, without this cumulative probability distribution, coincident spikes from more than T afferents would discount the probability of spiking, leading to an erroneous spiking probability. This result differs from the spike estimate suggested by Gerstner and Kistler in Section 2.2.7 (see Equation. 2.5) in that the probability of spiking is not contingent on the previous state of the cell but instead on the probability p_{Θ} that spikes on the afferents are properly timed (i.e., effective). However, this would require that p_{Θ} can be determined without reference to network activity. In addition, our approach takes into account local properties of cell connectivity without the requirement for knowing the total number of cells in a network. While we may at first regard the inclusion of the cell's previous state in Equation 2.5 as a step towards including refractory periods of cells, Gerstner and Kistler's method does not clearly incorporate refraction in their approach (Gerstner & Kistler, 2002, p.307). Before taking into account refractory periods, however, we first account for the probability p_{Θ} that spikes reaching a target are adequately timed to induce coincidence detection, thus constituting what we denote as *effective afferent spikes*.

6.2.2 Temporally Effective Afferent Spikes

To further complete the spiking probability estimate s_Θ from Equation 6.10, we must determine the probability p that spikes are suitably temporally arranged to lead to coincidence detection, i.e., the likelihood of being in the condition under which there is synchronous arrival from T afferents on a target neurone. We presently propose a set of estimates for the probability p_Θ , for any given neurone, that afferent spikes will be temporally effective.

Across multiple afferents, only spikes that are within temporal range to each other will be effective at producing a post-synaptic spike. This range D_{range} is determined by the minimum delay D_{min} and maximum delay D_{max} such that $D_{range} = D_{max} - D_{min}$. Hence two afferents carrying spikes from cells that fired more than D_{range} time steps apart cannot lead to synchronous arrival on a target because the delay difference between the two connections is too small to compensate for the offset of the spike times. Hence our spike distribution analysis is concerned with the scenarios where all afferent connections have delays that have fewer than D_{range} time step difference from one another.

6.2.2.1 Temporal Breadth

We first assume that a uniform distribution of delays reach a target with a total of I incident connections to this target such that $I = D_{range} + 1$. Hence, each afferent has a unique delay among the $D_{range} + 1$ possible values. For example, Figure 6.4 illustrates the case where in a network with $D_{range} = 3$ then, assuming uniform random distribution, four afferents would each possess a different transmission delay separated by a single time step. Furthermore, we assume for the time being that source neurones spike only once in the time period D_{range} (i.e., the output frequency of source cells is below $\frac{1}{D_{range}}$ or that $D_{range} < R$). Under these conditions, a coincidence of spikes is not guaranteed because the order in which spikes occur on the afferents of different delays may reach the target asynchronously. Indeed, the *order* in which spikes are emitted from the source neurones will determine whether a post-synaptic spike occurs. For example, from the illustrated model in Figure 6.4, the target cell has 4 source cells ($I = 4$) a, b, c, and d. Each project to that target with only a single time step difference such that $I = D_{range} + 1$, e.g., $D_a = 40$, $D_b = 39$, $D_c = 38$, and $D_d = 37$. If a coincidence detection between two cells ($T = 2$) is enough to produce a post-synaptic spike, then we are interested in all possible output pairs on the four cells such that spikes occur exactly one time step from each other in the

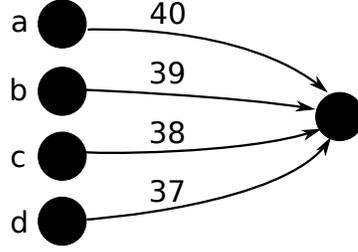


Figure 6.4: Four afferents to a target. With a delay range $D_{range} = 3$ and assuming uniform distribution of delays, then each projection are one delay time step apart.

correct pair order. For each spiking sequence among the four cells, there are 3 paired sequences of spiking. For example, if the cells fire in order (c, a, b, d) then the pair (c, a) the pair (a, b) and the pair (b, d) are potential firing sequences that could trigger a post-synaptic spike. Generally, among the $I!$ possible orders of I cells firing, the total number of paired subsets (subsets of two cells) that will fire in any given order is:

$$I! \times (I - 1) \quad (6.11)$$

This leads to a total of 72 possible pairs in our 4 afferent example. However, only a subset of these will be in the correct order (i.e., a before b, b before c, and c before d). The number of *effective subsets* is found with:

$$(I - 1)! \times (I - 1) \quad (6.12)$$

Hence with 4 afferents there are 24 spiking sequences among which there are 72 ordered paired inputs among which 18 ordered pairs will be effective by reaching the target coincidentally. In other words, $\frac{1}{7}$ pairs are effective over this temporal distribution. To determine the number of triples, quadruples, or any effective subset of size σ provided $I \geq \sigma$, we can use the equation:

$$(I - \sigma + 1)! \times (I - \sigma + 1), \text{ for } I \geq \sigma \quad (6.13)$$

Hence, the number of triples ($\sigma = 3$) will be $(4 - 2)! \times (4 - 2) = 4$ effective triples for $I = 4$. This is convenient if we are searching for the number of effective ordered sets that will lead to a coincidence threshold T by setting $\sigma = T$. Yet, the number of all possible subsets of size σ is:

$$I! \times (I - \sigma + 1), \text{ for } I \geq \sigma \quad (6.14)$$

Generally, out of all possible subsets of size σ , by taking the ratio of Equation 6.13 over Equation 6.14 the probability of effective subsets that will trigger a target to spike is thus:

$$p_{orders} = \frac{(I - \sigma + 1)!}{I!} \quad (6.15)$$

Where we p is assumed to be applied in the context of network parameters in Θ . This brings us closer to a probability of an effective input by considering one spike per afferent connection and exactly $I = D_{range} + 1$ connections. But what happens when one of more afferent neurone does not fire within D_{range} time steps from the others – i.e., when there is a missing spike?

6.2.2.2 Missing Afferent Spikes

In the case where one or more afferents do not possess a spike within a D_{range} time interval of other spikes, then the previously found probability of effective sequences of spikes should be discounted. We cannot simply subtract the missing neurone spikes from the number of incidents I and compute Equation 6.13 to determine the discounted probability, since neurones that fire near the middle of the order (with average delay projections) are more likely to play as an effective role as neurones at the beginning or end of the order (maximum or minimum delay projections). More specifically, neurones that are both preceded and followed by spikes from other neurones in the same pool of source cells will have twice the impact of neurones which are either preceded or followed only. Hence, when a spike is missing from a tail-end connection with delay D_{max} or D_{min} (e.g., from neurones a or d in Figure 6.4) delays will reduce the number of effective subsets of length σ by $(I - \sigma + 1)!$. Thus subtracting this from Equation 6.13 we obtain a total number of effective orders when ‘tail’ connection spike are missing to be:

$$(I - \sigma + 1)! \times (I - \sigma), \text{ for } I \geq \sigma \quad (6.16)$$

Yet, we know from Equation 6.14 the total number of subsets of size σ . We thus obtain the probability that subsets of cells will be effective when a spike is missing from ‘tail’ connections with either delays D_{max} or D_{min} to be:

$$p_{missing_tails} = \frac{(I - \sigma)! \times (I - \sigma)}{I!}, \text{ for } I \geq \sigma \quad (6.17)$$

However, when spikes are missing from neurones for which delays lie between two other propagation delays (e.g., from neurones b or c in Figure 6.4), then twice the number of initial effective patterns will be lost, hence leading to the probability of effective subsets to be:

$$p_{missing_between} = \frac{(I - \sigma)! \times (I - \sigma - 1)}{I!}, \text{ for } I \geq \sigma \quad (6.18)$$

These observations thus help build a more accurate estimate of the probability that any input sequence separated by one time step on the source cells will be effective on a target neurone. In particular, the case where neurones are misfiring or $I < D_{range} + 1$ can now be accounted for. This complements the basic, yet less likely, condition developed in Section 6.2.2.1 where $I = D_{range} + 1$. However, these estimates still lack the ability to account for connection topologies where the number of incident connections I is greater than $D_{range} + 1$. Indeed, neurones will likely have a large number of afferents such that $I > D_{range} + 1$, even if this number of incident connections respects the condition that delays remain within the same range. Hence, we need to take into account afferents which have duplicate delays to other connections leading to what we denote ‘spatial breadth’.

6.2.2.3 Spatial Breadth

Finally, to complete our insight into the probability p_{Θ} that incident connections will be temporally effective at producing a spike on the target, we need to consider the impact on this probability when additional connections with delays from D_{min} to D_{max} project to the target. This will necessarily be the case when $I > D_{range} + 1$. Indeed, additional connections from other source cells, if excitatory, should increase this probability of post-synaptic spiking. Adding new connections in this way thus constitutes an exploration of connectivity from a ‘spatial breadth’ view point, since additional source neurones with identical delays to existing source cells will affect the likelihood of post-synaptic spiking over the spatial population of cells rather than over increasingly extended delay periods.

Figure 6.5 illustrates the case when an additional afferent cell exists (cell e) with a propagation delay which coincides with other cells (here cells b and e have identical projection delays). This suggests that whenever neurones with identical delays fire synchronously they will reach the target synchronously. Hence specific sequences of firing will substantially favour coincidence detection on target cells due to these additional spatially significant cells. We thus seek to express how additional cells

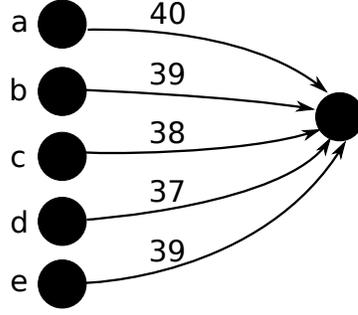


Figure 6.5: Four temporally distinct afferents and one spatially distinct afferent to a target with a delay range $D_{range} = 3$. Afferents a,b,c,d have distinct delays but afferent e has identical delay to afferent cell b.

impact the overall probability of effective input on a target cell.

Figure 6.6 illustrates the partial set of permutations for possible orders of firing that would be effective in causing the target cell to spike. With an additional spatially contributing cell that has equal propagation delay to another afferent cell, there is an increase in the proportion of effective subsets of cells (the circled pairs in Figure 6.6). The total number of possible pairs while taking into account an additional spatially contributing cell (e.g., cell e) is found to be $(I - 1)^2 \times (I - 1)!$. Hence for any temporal subset of size σ this amounts to the total number of possible subsets of size σ to be:

$$(I - \sigma + 1)^2 \times (I - \sigma + 1)!, \text{ for } I \geq \sigma \quad (6.19)$$

Now to determine the number of effective pairs when including the spatially contributing cells, we begin by counting the number of effective temporal pairs which, from Equation 6.12, we know is $(I - 2)! \times (I - 2)$ (we subtract 2 instead of 1 because, in our example, only 4 cells are temporally related yet there are a total of 5 incidents to the target). We then multiply this by the number of times the ordered permutations are repeated due to the spatial incident, hence by $(I - 1)$. Finally, we add the number of effective spatial pairs which is $(I - 1)!$. This results in the total number of effective spatiotemporal pairs S_{pairs} to be:

$$\begin{aligned} S_{pairs} &= (I - 2)(I - 2)! \times (I - 1) + (I - 1)! \\ &= (I - 1) \times (I - 1)! \end{aligned} \quad (6.20)$$

Importantly, this is equivalent to the number of effective temporal pairs only. To

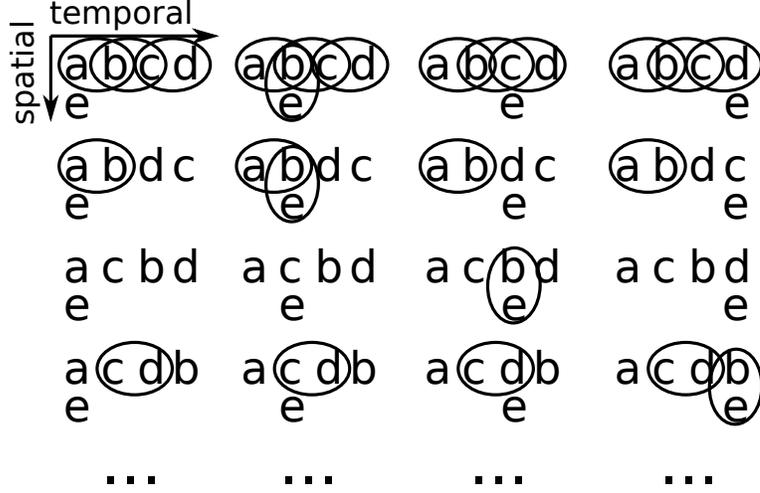


Figure 6.6: Partial set of permutations of the temporal order of the four cells a, b, c, d, and the spatial cell e. Effective pairs are circled. Pairs of cells firing one step apart and in order with a single propagation delay time step difference are circled because they will lead to a coincidence of 2. Also spatial cells that have equal propagation delays and fire synchronously are circled (here only b and e since both have propagation delays of 39).

generalize to the number of effective spatiotemporal subsets of size σ we obtain:

$$S_\sigma = (I - \sigma + 1) \times (I - \sigma + 1)!, \text{ for } I \geq \sigma \quad (6.21)$$

From this we deduce that of all possible sequences of spikes separated by a single time step expressed by Equation 6.19 only S_σ will be effective. Hence the probability of effective subsets of size σ provided one additional spatial cell is:

$$p_{spatial} = \frac{1}{I - \sigma + 1}, \text{ for } I \geq \sigma \quad (6.22)$$

This estimate thus allows to incorporate the condition where a subset of afferent fire synchronously yet possess identical propagation delays to others. There are limitations, however, to the estimates of the probability p_Θ that afferent spiking activity will be effective. We explore these limitations in the following.

6.2.2.4 Application and Limitations

At first glance these estimates appear limiting in that we only considered them in the narrow condition where afferent neurones spike exactly one step apart from each other (temporal breadth), when spikes are missing, or when equal propagation delays exist (spatial breadth). However, other network and spiking configurations should be

accounted for.

First, there is the spiking condition where afferent cells do not all spike one after the other but where some spike together despite having non-identical outgoing propagation delays; that is, despite having $I = D_{range} + 1$, simultaneous spiking can occur. Indeed, out of all $(D_{range} + 1)^I$ possible spiking combinations among source cells, assuming a single spike per afferent, only $I!$ combinations will be asynchronous across all cells. However, subsets comprising between 2 and I cells could also be firing synchronously. Hence, the probability that an initial combination of spikes will be effective will have to be calculated based on the known number of synchronous subsets. Importantly, however, this condition can be accounted for using our existing estimates by treating subsets of synchronous spiking cells with different propagation delays as a single effective spiking cell and all other cells part of that subset as missing spikes as seen in Section 6.2.2.2. Indeed, the number of effective spikes will be reduced equivalently since two source cells spiking together that have different propagation delays implies that only one of the two source cells can still play an effective role with other neurones. In the case where two afferent cells spike synchronously but also have identical propagation delays, then we return to the condition where effective spikes are computed as spatially additive as seen in Section 6.2.2.3.

Second, the spatial breadth condition estimate can also apply to networks where propagation delays are not uniformly distributed across the delay range if the number of the overlapping delays are known (which should be available from the probability distribution function). This condition is in fact guaranteed when $I > D_{range} + 1$.

Finally, there is also the possible scenario in non-uniform delay distributions where ‘delay gaps’ exist; that is, where no source cell exist which has some particular propagation delay between D_{min} and D_{max} . In this case, missing propagation delays can be treated as missing spikes by using the effective spiking probability estimates established in Section 6.2.2.2.

Hence, both estimations of missing afferents and spatial breadth can account for all initial spiking states of afferent cells to a target and also various distributions of connections within a delay range of interest.

The greatest limitation of these proposed probability estimates stems from the unclear ability of these methods to account for more than a single missing spiking cell, or more than a single additional spatial cell. Furthermore, combinations of missing afferent spikes, synchronous firing and non-uniform delay distributions are likely and important to consider for biologically realistic models. We leave these questions as areas of future investigation, which would lead to more exact estimations

of the probability p_{Θ} that afferents will produce effective spikes. While additional work could adapt these methods to scale to larger sets of missing source spikes or additional cells, and combinations of these, we believe that the methods presented in this subsection provide a solid foundation for the analytic description of probability that spikes in DCD networks will be temporally effective.

6.2.3 Spiking Probability with Refractory Periods

By determining the probability for effective connectivity leading to the spike probability s_{Θ} and the probability that those connections will be carrying effective coincident spikes p_{Θ} we now have a probability estimate that a target neurone will spike in the optimistic probabilistic condition where activity is assumed to be maximum. However, this estimate ignores the effect of refractory periods on cells which is dependent on the previous activity of neurones. Here we develop a method which attempts to capture the firing constraint imposed by refractory neurones on the overall probability that any given neurone fires.

The total number of spikes that can be present on an afferent connection is limited by the ratio of the transmission delay over the refractory period D/R . For example, on an afferent connection with a delay of 40 time steps and a source cell with $R = 10$ means that a maximum of 4 spikes could be loaded on the afferent connection such that spikes are 10 time steps apart from each other. Hence the total number of discrete spike states a connection can be in for a single sequence of intervals R is $2^{\lfloor D/R \rfloor}$. However, these permutations exist for all the R possible offsets along the connection. Hence a total of $R \times 2^{\lfloor D/R \rfloor}$ states are possible when taking into account refractory periods. Yet, if refraction is ignored, the total number of states a connection can be in is 2^D . Hence the probability r_{Θ} for a connection to be in a given state carrying spikes while taking into account refractoriness of the source neurone is:

$$r_{\Theta} = \frac{R \times 2^{\lfloor D/R \rfloor}}{2^D} \quad (6.23)$$

Importantly, we can use r_{Θ} as an overall estimate for the probability that refractoriness is affecting any given cell in the network. Under this interpretation, the probability estimate r_{Θ} depicts the pessimistic condition where activity of all source neurones is maximal, while 2^D states are possible when activity is minimal. Hence, we can use r_{Θ} to construct a scaling factor f such that $f = 1.0$ when the percentage of activity in the network is minimal ($a = 0.0$ where $a = \frac{A}{N}$), and $f = r$ when the percentage of activity in the network is maximal ($a = 1.0$). We define f as the linear

equation:

$$f(a) = (1 - r_{\Theta}) - a(1 - r_{\Theta}) + r_{\Theta} \quad (6.24)$$

We can now use this refractory scaling factor to adjust the probability of spiking under refractory conditions denoted as z_{Θ} . We define z_{Θ} as:

$$z_{\Theta} = f(a) \times s_{\Theta}(p_{\Theta} * a) \quad (6.25)$$

Here, the probability p_{Θ} for effective spike synchrony is also scaled with respect to the amount of activity in the network, since p_{Θ} is only maximally effective if all afferent neurones are providing a spike. By setting the number of effective subsets $\sigma = T$ we obtain the condition where sufficient neurones with adequately timed spikes will be available to trigger a post-synaptic potential. In combination with our estimate for effective connectivity we can compute s_{Θ} . Finally, by applying the refractory scaling factor f we obtain an overall spiking probability z_{Θ} for any neurone without recourse to individual cell states. Figure 6.7 illustrates the spiking probability distribution as a function of the overall percentage of activity in the network A_P for the different effective temporal spiking probability conditions: p_{orders} , $p_{missing_tail}$, $p_{missing_middle}$, and $p_{spatial}$. These probabilities were computed for parameters used in the Standard DCD network model, such as a degree of incidence $I = 24$, a coincidence detection threshold $T = 2$, a refractory period $R = 10$, and an average delay between 20 and 40 time steps of $D = 30$. From this figure we notice that the probability of spiking initially increased as the activity in the network increased. In the case of applying p_{orders} , a spiking probability reached 0.0104 for an activity level of 17%, and reached a maximum probability of 0.0482 when an activity level of 62% was present. As activity continued to increase, however, the probability of spiking was increasingly dampened by the refractory scaling factor f resulting in a probability of spiking that converges to 0 as the activity level nears 100%. It is important to note, however, that these results reflect the narrow condition where the I afferents possess propagation delays that are one time step apart from each other. Other possible model probabilities p_{Θ} for effective spike timing are shown in Figure 6.7. As expected, in the case of additional spatial cells, there was an increase in the probability of spiking, while missing spikes decreased this probability.

Hence, by regulating the probability of spiking s_{Θ} using the refractory scaling factor f we obtain a general expression to capture the likelihood that a DCD neurone will spike provided a particular level of activity in the network. Naturally, this remains an approximation as individual refractory periods responses of cells are not taken into

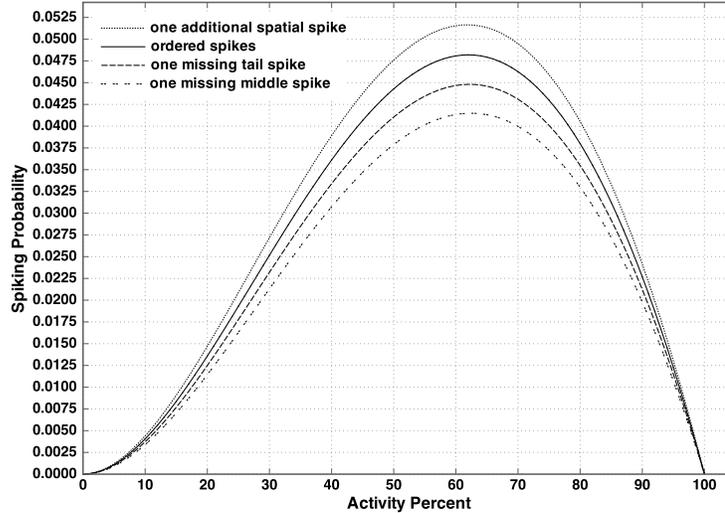


Figure 6.7: Spiking probability distribution z as a function of the overall percentage of activity in the network a for the different effective temporal spiking probability conditions: p_{orders} , $p_{missing_tail}$, $p_{missing_middle}$, and $p_{spatial}$, for parameters $I = 24$, $T = 2$, $R = 10$, and $D = 30$. By taking into account the refractory scaling factor f , probability of spiking is reduced despite increasing activity.

account and homogeneous temporally effective spiking probabilities are assumed, as well as other network parameters such as I , T , and D . This approach, however, does not prevent the fine tuning of estimates such as the probability that spikes will be temporally effective p_{Θ} to incorporate specific details about network connectivity and spiking developed in Section 6.2.2.

6.2.4 Probability of Firing States

With the probability of spiking z_{Θ} now available, we can develop an estimate for the number of states in which a DCD network can be between time step t to $t + n$. By doing so, we hope to get closer to an estimate for the number of firing-chains that are available in the network. Under ideal conditions, this would be estimated independently of the degree of activity inputted into the network. However, as seen in Section 6.2.3, the effect of refraction is dependent on this degree of activity. To accommodate for this activity dependency, we could either describe the number of spiking states for the network based on input and passed activity, or make use of an average of the spiking probability z_{Θ} such that it remains independent of the activity state of the network. Here, we opt the latter approach.

To obtain an average number of states available independently of the amount of

activity in the network, we propose to identify the average spiking probability z_{avg} by classifying the network activity under a number of k different activity domains where each occupy an α_k percentage of time during the time evolution of the network activity. For example, we can consider 2 domains, where domain 1 captures activity from 0 – 50% in the network, and domain 2 which captures 50 – 100% activity in the network such that each domain is occupied for 50% of the time, i.e., $\alpha_1 = 0.5$, and $\alpha_2 = 0.5$. To compute the average spiking probability we compute the areas under the two curves and set them to be equal. The vertical line that separates the two domains so that the areas of the curves are equal defines the the average spiking probability z_{avg} . This value may be computed using numerical analysis using, for example, a binary search technique. We leave the problem of solving the average spiking probability with more than two domains ($k \geq 2$) for future work.

Using z_{avg} , we can express the number of states in which the network can be found at any given time step as:

$$h_{\Theta}(z_{avg}) = \frac{N!}{(\lfloor z_{avg} \times N \rfloor)! \times (\lfloor (1 - z_{avg}) \times N \rfloor)!} \quad (6.26)$$

Consequently, we can expand this to express the total number of possible spiking states for n time steps as:

$$E_n = 2^N \times \prod_{t=1}^n h_{\Theta}(z_{avg}) \quad (6.27)$$

At time step $t = 0$ the network of N cells can be in any of the upper bound of 2^N possible spiking states. For time steps $t \geq 1$, however, we use the binomial distribution to express the number of possible states in which the network can be by considering the reduced lower bound number of possible cells firing $\lfloor z_{avg} \times N \rfloor$. By computing the product of these possible spiking states over n time steps we obtain a total number of combined spiking states over that period of time.

This total combined state estimate E_n , however, merely provides an upper bound for the number of possible firing-chains within a DCD network. Interestingly, if any possible subsequent state from $h_{\Theta}(z_{avg})$ could be reached at time step $t + 1$ from any of the states from previous time steps t to $t - m$, then E_n would, in fact, correspond directly to the number of firing-chains available in the network. We must account for the m time steps in the past because delays carry effective activity beyond a single past time step. As a lower bound, however, if each state of the network at time step $t + 1$ is exactly dependent on a single state over the previous time steps

then there are $h_{\Theta}(z_{avg})$ firing-chains in the network. Interestingly, this lower-bound is, in fact, an accurate estimate for the number of firing-chains in the network when spiking is sufficiently uncorrelated, i.e., that no two neurones spiking at time t have identical ancestor spikes. As mentioned in Section 2.2.3, this condition is met in standard integrate and fire networks when $I < \log_e N$ (Nützel et al., 1994). However, as discussed in Section 5.1.2.1, in the case of DCD neurones the number of effective incident connections depends on the coincidence detection threshold T suggesting that sufficient uncorrelated activity will take place when $\frac{I}{T} < \log_e N$ is satisfied. In our Standard DCD network model, however, $\frac{I}{T} = 12$ and $\log_e N = 3.219$ suggesting that correlated firing is expected in networks of the type investigated up to this point. Hence the number of firing-chains in the network will be between these lower and upper bound limits.

Interestingly, this notion of spike correlation based on ancestry may be brought to bear on the problem of estimating the number of firing-chains that can be instantiated in a DCD network. In particular, minimum correlation is achieved when $T \rightarrow I$ (T tends to I) such that $\frac{I}{T} \rightarrow 1$. In this case, uncorrelated activity will be take place as long as $N \geq 3$. However, if this state of uncorrelated activity corresponds to the lower bound number of possible firing chains of $h_{\Theta}(z_{avg})$ mentioned above, then we can map the simple network features of degree of incidence I and coincidence threshold T to the lower bound number of firing chains. Furthermore, the maximum degree of correlated activity is reached when $T \rightarrow 1$ and $I \rightarrow N$. Similarly, if this state of correlated activity has a direct correspondence with the upper bound number of possible firing chains E_n established earlier, then we can again map the maximum number of firing chains to network features I , T , and N . Hence, by fixing values for these network features, we could use their mapping to provide an approximation as to the number of firing-chains in DCD networks.

While we do not provide for a definite method of mapping network properties to the number of firing-chains in a network, the bounds for this estimate can serve as a starting point for this development. Unfortunately, the accumulation of averaged network properties such as delay mean, incidence mean, etc., suggest that accurate estimates will be mitigated unless precise evaluations of the spiking probability are made. This will have to be based on a detailed mapping of the delays, connectivity densities in the network, and refractory period of cells. Instead, we have offered overarching methods so as to circumscribe the estimates of firing probabilities, number of states for DCD networks, and limits on the number of available firing-chains within analytically reasonable bounds. By doing so, we hope to have provided a preliminary

set of analytical tools to help reveal and motivate the subtle complexity of delay coincidence detection networks and their unique neural coding features via functional connectivity.

6.3 Summary

Because our capacity to understand the behaviour and function of complex systems is limited by our ability to sufficiently describe them, in this chapter, we have gone beyond the empirical testing of modelled delay coincidence detection networks and attempted to develop approximate analytical tools to capture features that are not only critical to their structure but also function. First, we began by describing the key components of DCD networks that must be taken into consideration when seeking to develop a general model for the overall activity of a network based on initial activity conditions. In particular, this model succeeds in predicting general behaviours of the networks previously studied empirically. Despite this, we argued that some important considerations need to be addressed, such as the individual effect of cell refractory periods, in order to gain a more realistic prediction of behaviour. While we proposed a method to approximate this effect in our activity estimation in Section 6.1, this approach simply relies on reducing the temporal window of past activity proportionally to the past activity of the network and the average refractory period of cells.

Because, estimating general activity alone is not sufficient as an indicator of the functional coding capacity of these networks, we then proceeded to establish a probabilistic estimate of the number of spiking states a network could reach provided both spatial and temporal constraint parameters but independently of activity. We did so in four steps: (1) determining the probability s_{Θ} that a spatial arrangement of connectivity will lead to a post-synaptic spike. A probability which depends on the proper temporal arrangement of afferent spikes, thus requiring the following; (2) determining the probability p_{Θ} that afferent cells would be spiking in an effective order to produce a post-synaptic spike. This included considerations for missing spikes and additional neurones; (3) determining the influence of refractory periods of the overall possible states of the network leading to a general term for the spiking probability of cells z_{Θ} . (4) determining the number of spiking state of the network h_{Θ} based on z_{Θ} over time. Finally, we have introduced preliminary directions to explore the upper and lower bounds for the number of firing-chains in the network and elements that could lead to a possible mapping to obtain an estimated number. Despite this, a

more refined approximation to estimate the number of possible firing-chains in DCD networks was not reached. We believe that further considerations with respect to the degree of correlation and will be required to achieve this.

Chapter 7

Conclusion and Future Work

7.1 Conclusion

The aim of this thesis was to gain a better understanding of neural coincidence detection with transmission delays as a potentially fundamental mechanism for aspects of reactive and adaptive neuronal function. Networks of these neurones can be tuned via the appropriate selection of delays, connectivity, and coincidence detection thresholds to accomplish complex transformations of input over space and time. While neural encoding via firing rates still remains the most recognized form of information processing in the nervous system, we have shown both empirically and theoretically the potential of coding via delay coincidence neural networks represents. Although others have proposed that spatiotemporal coding via precise spike timing could be significant for neural coding, this thesis is the first comprehensive analysis of network structure, dynamics, function and application within an embodied task to our knowledge. In line with the embodied/embedded approach to cognition, we stratified in Section 3.1 the complexity of simulated robot tasks with respect to two key categories that distinguish simple sensorimotor control, or reactive behaviour, from advanced complex control, or adaptive behaviour.

After formalizing the binary neural network model used for simulation in Section 3.2.2, we adopted a bottoms-up approach to modelling cognitive behaviour to show how particular dynamic regimes could be obtained with potentially useful functional capabilities. In Section 6.1 we explored initial dynamics in DCD networks and found that narrow delay ranges led to highly synchronized spiking while broad delay ranges led to complex firing. To further explore the capacity for control of these two dynamic features we evolved artificial network controllers in Section 3.2.5 for a

reactive light-seeking task that displayed greater performance with narrow delay range connectivity. This suggested that reactive mechanism should be suitably addressed by network with a narrow distribution of delays that produce highly synchronous activity. Instead, broad delay range connectivity lead to a more diverse set of robot behaviours. Furthermore, broad delays were shown in Section 3.3.1 to reach chaotic regimes in the case of low degrees of internal network connectivity, while high degrees of connectivity lead to stable limit cycle firing. Fortunately, this regime suited the characteristics of dynamic memory initially hypothesized in early theoretical considerations.

In order to utilize self-sustained activity of stable periodic patterns as a possible form of dynamic memory, we established a preliminary understanding of the ability for DCD networks to differentiate between input patterns and reliably produce responses. In particular, the significant divergence between responses despite similarity of the input suggested that a large number of patterns could be potentially encoded in these networks. For a system to work as a memory, however, a mechanism for decoding must be implemented. To accomplish this we developed, in Section 4.2, a ‘delay mapping’ method whereby a secondary decoding network of coincidence detection cells received input from a memory network storing spatiotemporal patterns. More specifically, by deriving the appropriate delays for projections from the memory network to the decoding network, spatiotemporal responses in the memory could deterministically provoke target decoding cells to fire. After developing this method for memory decoding, it was tested in an embodied robot task requiring an agent to navigate a T-maze which we presented in Section 4.3. In particular, it was shown that a sensorimotor network of coincidence detectors behaving as a reactive system could guide the robot up the maze effectively without colliding into walls. By inhibiting the motor neurones of this network with spikes originating from a memory decoding network, dynamically stored stimulus cues could control the robot to perform correct turning behaviours. In order to help ensure unambiguous decoding of spatiotemporal patterns, we also developed in Section 4.4 a method which identifies the minimal decoding threshold to be used for decoding neurones when a number of patterns need to be independently retrieved from memory. To gain an even greater appreciation for potential storage capacity of DCD networks under high connectivity yet broad delay distribution conditions, we tested this robot in an F-maze paradigm in which two turns had to be made; each with a different set of context/cue stimulus pairs. Successful behaviour was obtained despite the same memory network size as in the T-maze task. Arguably, however, the combined storage of context/stimuli pairs elicited greater ambiguity than the storage of independent patterns. In the F-maze

task, however, memory was reset between turns. It remains unclear whether the additional memory load from storing the combined context/cue stimulus pairs from both turns could have been successful. We discuss this further in the following section on directions for future work.

A potential threat to all information processing systems stems from the presence of noise. To better appreciate the capacity for DCD networks to tolerate noise we explored in Section 5.1 their informational properties analytically with respect to ambiguity and tested via simulation their capacity to maintain decodable features despite increasing levels of perturbation. In particular, this perturbation was applied in low and high noise conditions. In the low noise condition, both spike removal and misfiring conditions were tested where cells were perturbed no more than once during simulation. In the high noise condition, cells were perturbed for the entire duration of the simulation after noise onset. Overall, these results revealed that features remain available in all conditions when up to approximately 60% of neurones are disrupted. This suggests that DCD networks offer a relatively robust capacity for tolerance to noise. For increased noise tolerance, these results indicated that the coincidence detection threshold of decoding cells would have to be reduced in order for a decoding network to tolerate the misfiring or absence of spikes. Interestingly, this comes as a competing requirement in contrast to dynamic memory capacity which demands a greater coincidence detection threshold for greater, unambiguous, capacity. To gain a better appreciation for the capacity of DCD networks, we then explored in Section 5.2 their ability to differentiate between response patterns as the number of elicited patterns increased in networks of 25 neurones. After testing for ambiguity analysis on up to 88,550 patterns the degree of ambiguity limit – equal to the number of neurones in the network – was not reached. By computing a non-linear regression on these ambiguity results, we projected that up to 10^{12} patterns could be stored such that they would remain uniquely identifiable by a decoding network. In addition, we also introduced the notion of weight of ambiguity which described the prevalence of particular degrees of ambiguity across a number of response patterns. This estimate serves to better characterize the quality of DCD networks as dynamic memories under general load constraints.

Finally, we aimed to develop a greater formal description of both activity and state probability of DCD networks. With respect to activity estimation, a probability of spiking based on cell activity was established. Furthermore, this estimate took into account neuronal refractory periods by reducing the time window of activity to be summed. While this analytical estimate for activity was shown to exhibit

general network dynamics properties similar to actual neuronal simulation, this precision remains coarse. This is primarily due to the refractory period being estimated from global network activity rather than individual neurone activity. To move beyond activity dependent estimates and gain a greater appreciation for the number of spiking states a DCD network can have over a period of time, we developed in Section 6.2 activity independent probability approximations that capture the spiking probability based on the structural and temporal arrangement of afferents. This was then combined with a global refractory scaling factor in order to obtain a probability distribution of spiking. By assuming separate domains of activity for DCD networks, we established that an average refractory effect could be determined while avoiding the incorporation of past activation in our final state probability estimate. Finally, we developed estimates for bounding the possible number of firing-chains in DCD networks from this state probability estimate. While considerations with respect to degree of correlation based on connectivity should help shed light on isolating this estimate, further analytical investigation is warranted to reach an accurate method for finding the number of firing chains possible in the DCD network based on structural parameters alone. Indeed, if such a measure could be found, then effective memory estimates should become possible analytically without requiring network simulation or potentially even ambiguity analysis.

Overall, this thesis encompasses a natural scientific question which was motivated by a desire to develop a better understanding for not only the dynamical but also the functional properties of neurones which can behave as coincidence detectors and by taking into account realistic conditions imposed by the spatial distribution of cells leading to propagation delays. The approach adopted here was rooted in the bottom-up approach to cognition with an emphasis on validating neural network processing within realistic embodied tasks. Not only have computer simulations enabled this in-principle investigation, but analytical efforts have also consolidated a better appreciation for the properties and subtleties that DCD networks have to offer with respect to their information, dynamics, and function for cognition.

While this work led to a set of comprehensive experiments, new questions naturally arose at these findings were made. In the following section, we explore remaining questions and avenues for potential research work in the area of neural processing via delay coincidence detection.

7.2 Future Work

In the following we describe areas of research which may lead to promising directions of inquiry and discovery with respect to increased biological consolidation of DCD network function, their possible regulation with respect to noise tolerance and memory capacity, their encoding capacity and limit with respect to combined stimuli, and analytic estimation of noise tolerance.

7.2.1 Conductance Based Modelling of Pyramidal Coincidence Detection

As introduced in Chapter 1, early computer simulation work on precise conductance-based models of pyramidal cells by Bernander and colleagues, as well as in-vivo observations by Prescott and colleagues, have shown that high degrees of contextual input can lead to the functional transformation of cells from integrators to coincidence detectors (Bernander et al., 1991; Prescott et al., 2006). Not only did Bernander et al.'s work show that the membrane resistance was substantially reduced from $110M\Omega$ to $14M\Omega$, but the electronic distance of the synapses to the soma also increased by a factor of 3. It would be of interest to replicate an accurate compartmental model pyramidal neurone with corresponding neighbourhood of other pyramidal cells but also inhibitory interneurons with accurate neurotransmitter related activation equations. Furthermore, realistic spatial characteristics should be taken into account. From this network model, we could test and attempt to reproduce the functional characteristics of sustained firing, limit-cycle dynamics, decoding, and more. Doing so would help dispel any worry that lower level dynamics may interfere with the functional characteristics identified in this thesis with simple binary threshold neurones. Furthermore, expressing these networks with continuous valued delay differential equations to express the propagation and activation dynamics could also support additional analysis of DCD networks with respect to dynamical characteristics.

7.2.2 Hebbian Plasticity to Regulate Capacity vs. Noise

In Section 5.2, we found that ambiguity between possible response patterns within DCD networks increased as the number of stored inputs to be retrieved also increases. Furthermore, in Section 5.1, we found that sub-features for a given response pattern could still be used to identify a pattern if many (but not all) neurones fire at an incorrect time or do not fire at all. To decode these patterns using sub-features, however,

a lower decoding threshold for target cells is required, otherwise these targets would not fire due to inappropriately timed afferent spikes. Hence a decoding threshold tradeoff exists which suggests that DCD networks should achieve a balance between encoding capacity and noise tolerance. We hypothesize that biologically plausible mechanisms may be suitable to accomplish this so that either decoding thresholds are somehow regulated or by preventing decoding neurones from temporarily behaving as coincidence detectors. Here, we present two hypothetical mechanisms that could receive greater attention in future work.

First, a form of Hebbian plasticity such as spike time dependent plasticity (STDP) (Bi & Poo, 1998) could indirectly influence the decoding threshold of cells by increasing or decreasing the synaptic strength of afferents. Indeed, an increase in synaptic efficacy leads to a reduced requirement for the number of afferent spikes to produce a post-synaptic action potential. Hence, strengthening is equivalent to the requirement of a reduced decoding threshold while synaptic weakening is equivalent to the requirement for a greater decoding threshold. However, we have seen in Section 4.4 that a greater decoding threshold is necessary to differentiate between partially ambiguous patterns. Yet, in Section 5.1 we have seen that lowering the decoding threshold can help reduce the effect of noise. Consequently, plasticity may hypothetically play a homeostatic role by dynamically regulating this capacity versus resilience tradeoff. A priori this strategy would be most effective under anti-hebbian strengthening conditions, since non-noisy stable patterns could afford to reduce connectivity strength (increasing the decoding threshold) while noisy spiking would induce strengthening of connections (reducing the decoding threshold) such that only a stable subset becomes required for effective decoding. While synaptic plasticity is a process that is traditionally viewed as purposeful for learning alone, this hypothesis suggests that it could also play an important role in regulating the degree of noise when neurones behave as coincidence detectors. We leave this notion of dynamic regulation of the decoding threshold and its mechanism as an open question for future work.

Alternatively, the effective time constant of neurone membranes could be modified by variation of inhibitory load. As shown by Bernander et al. (Bernander et al., 1991), contextual inhibition can dynamically alter the so called time constant of modelled layer V pyramidal cells from 80ms to as little as 7ms. Hence, if a reduction in contextual inhibition took place instead, leading to a lengthening of the time constant, then decoding cells could temporarily cease to behave as coincidence detectors and instead behave as integrators leading to decreased sensitivity to synchronous afferent spikes. This would prevent erroneous decoding of noisy signals yet maintain

coincidence detection on other cells. This approach, however, fails to explain the contextual rules that would selectively choose which neurones must change behaviour. Again, we leave this alternative mechanism as an open question for future work.

7.2.3 Stimulus Convolution

In the F-maze task discussed in Section 4.5, activity in the memory network was reset after the first context/cue paired stimulus was presented. If the robot made a turn into the second corridor, the memory network activity was stopped by forcing cells into a resting state, then a different context/cue pair was presented to the same memory network. This allowed for the testing of independent memory storage. While convolution of two inputs took place after the cue stimulus was presented to the memory network while the context was stored, it remains unclear whether additional stimuli could be stored effectively such that decodability could remain possible. Indeed, ambiguity analysis between context stimulus response only and the convoluted response of both context and cue lead to maximum degrees of ambiguity Λ_{max} as high as 15. Interestingly, this is similar to a noise condition where a subset of cells are forced to misfire once as seen in the noise condition of Section 5.1.3.3 when exactly 6 neurones were set to temporarily misfire. Indeed, the noise analysis estimated an average degree of self-ambiguity of approximately 14 with the original pattern when noise was applied to 6 cells. This means that approximately 10 neurones could no longer be used for decoding¹. This corroborates the results from the combined effect of a second input stimulus to 6 cells on an existing pattern in the context/cue stimuli convolution where a mean degree of ambiguity of 12.6 was obtained for the context/cue pairs in the 10 tested networks.

Ideally, we could measure the degree of ambiguity which results from the convolution of additional stimulus of a given size until the maximum degree of ambiguity Λ_{max} is equal to the number of neurones in the network. Interestingly, the prolonged noise condition when 6 neurones were perturbed via misfiring led to a degree of self-ambiguity of approximately 10, signifying that 15 neurones were no longer available for decoding. Unfortunately, the prolonged noise conditions tested in Section 5.1 apply to the case where the same neurones are continuously disrupted. In the case of stimulus convolution with a stored pattern, however, the new input stimulus could affect a different set of cells on every occasion. This would likely result in greater

¹Recall that self-ambiguity indicates the number of neurones which fire with features corresponding to the original pattern.

negative impact for spatiotemporal pattern decodability. We leave this as a question for future exploration.

7.2.4 Noise and Efficient Coding

In Section 5.1 we presented empirical results on the effect of noise on DCD network spiking dynamics. Those results showed that while our DCD network model exhibited substantial change in spatiotemporal spiking features under noisy conditions, sufficient features remained available for dynamic memory decoding if the coincidence detection threshold of decoding cells can be reduced. These empirical estimates, however, fail to be generalizable to other possible DCD network configurations, such as narrower or broader delays, variable degrees of incident connections per cell, variable coincidence thresholds within the network, etc. Here, we present work which aims to provide a greater formal understanding of the impact of noise on DCD networks provided any network configuration.

Although a number of numerical anecdotes regarding the role of noise in spatiotemporal coding networks are reported in the literature, no standard framework to quantify its impact has been proposed to date. Fortunately, existing work in network diagnostics may be suitably adapted to the characterization of noise in DCD networks. In particular, the Code Book approach to event correlation provides a set of tools to do so (Kliger, Yemini, Yemini, Ohsie, & Stolfo, 1995). In their work, Kliger et al. were concerned with the problem of adequately correlating ‘symptoms’ to their source ‘problems’. This is especially problematic in the case of complex networks where combinations of different problems can lead to combinations of different symptoms. Analogously, we may approach the notion of problems causing symptoms from a neural perspective by interpreting problems as source cells and symptoms as target cells. Furthermore, they show that Code Book is adapted to both spurious signals, which are analog to spike additions, as well as signal loss, which are analog to spike removals. Hence, adapting their approach may be suitable for a quantitative assessment of the degree of noise that a network can tolerate.

Figure 7.1 illustrates a simple correlation table relating three problems p1, p2 and p3 with three symptoms along with a network representation of the identical relationships. Here, each correlation relationship is sufficient for the symptom to occur. For example, symptom s1 can be caused by problem p1 or problem p2. However, p1 can also cause s2, but p2 does not. So if symptoms s1 and s2 occur but not s3, then we can deduce that p1 is the originating cause. The list of outgoing connections for a given

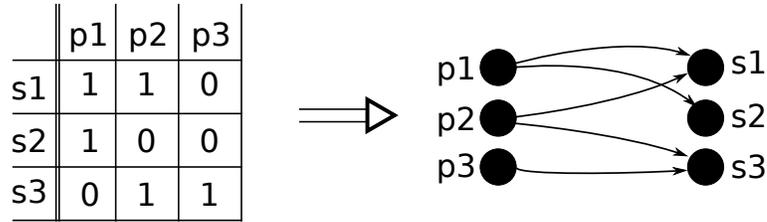


Figure 7.1: A simple correlation table relating three problems (p1, p2, p3) to three symptoms (s1, s2, s3) and the translated network representation.

problem (or columns in the table) are known as a problem’s code. In large correlation tables, when all symptoms can be uniquely identified by a subset of problems, the authors denote that set of problem codes the table’s *code book*. Kliger et al. propose that the distinction between problems should be made by measuring the Hamming distance between the problem’s codes. These correlations can be deterministic (0 or 1) or probabilistic. For the example given in Figure 7.1, the Hamming distance between p1 and p2 codes is 2, the Hamming distance between p1 and p3 is 3, and the distance between p2 and p3 is 1. This means that at least one difference between all problem to symptom correlations exists which implies that a disambiguation between problems would be possible provided any manifestations of the symptoms. This leads to the important notion of the radius of a code book. This radius is the half of the smallest Hamming distance between all problem codes. In our current example, p2 and p3 have the smallest distance of 1, which implies that this code book (list of problem codes) has a radius of 1/2.

Importantly, a radius of 0.5 suggests that a single noisy event will no longer guarantee that problems can be identified from the symptoms. A noisy event manifests when symptoms occur when they should not (spike addition), or do not occur when they should (spike removal). If the radius r of the code book is 1 or more, however, then r noisy events can take place such that the problems will remain identifiable. Furthermore, as Kliger et al. note, if the number of symptoms k is less than the radius of the codebook than $2k$ errors can be detected. To extend this insight to DCD networks, however, we need to account for three other important network factors: average coincidence detection thresholds (T), average delays (D), and average cell refractory period (R).

One potential approach to representing the ‘code book’ defined as a set of afferent neurones projecting to a set of target cells when transmission delays must be taken into account, is to regard, in the discrete case, each delay time step for a given

afferent connection as *independent potential problem sources*. By taking into account the fact that no more than $E = \lceil \frac{D}{R} \rceil$ of these steps can actually be combined with other afferent cell spike time steps, then a total of E individual spike times can be combined for each N cell to form the total number of ‘problems’ for the N target cells. Assuming that E is equal among each afferent neurone, then the total number of problems P is:

$$P = E * \binom{N}{T} \tag{7.1}$$

By solving for the number of problems P a correlation matrix between N symptoms and P problems could be derived. Using this, one could determine the radius of the code book and its tolerance to noise based on criteria set by Kliger et al. discussed above. This method, however, remains to be validated. Indeed, the added complexity of delays and correlation in firing suggests that other factors may have to be taken into account. We leave this area of investigation open for future work.

References

- Abeles, M. (1982). Role of the cortical neuron: integrator or coincidence detector? *Israel journal of medical sciences*, 18(1), 83.
- Abeles, M. (1991). *Corticonics: Neural circuits of the cerebral cortex*. Cambridge University Press.
- Abeles, M., Hayon, G., & Lehmann, D. (2004). Modeling compositionality by dynamic binding of synfire chains. *Journal of Computational Neuroscience*, 17(2), 179–201.
- Adams, F., & Aizawa, K. (2001, March). The bounds of cognition. *Philosophical Psychology*, 14(1), 43–64. Retrieved from <http://www.informaworld.com/openurl?genre=article&doi=10.1080/09515080120033571&magic=crossref||D404A21C5BB053405B1A640AFFD44AE3> doi: 10.1080/09515080120033571
- Adrian, E. D. (1928). The basis of sensation.
- Aertsen, A. M., Gerstein, G. L., Habib, M. K., & Palm, G. (1989). Dynamics of neuronal firing correlation: modulation of 'effective connectivity'. *Journal of Neurophysiology*, 61(5), 900–17.
- Amit, D. J. (1989). *Modeling brain function - the world of attractor neural networks*. Cambridge University Press.
- Amit, D. J., Gutfreund, H., & Sompolinsky, H. (1985). Storing infinite numbers of patterns in a spin-glass model of neural networks. *Physical Review Letters*, 55(14), 1530.
- Antal, A., Keri, S., Kovacs, G., Janka, Z., & Benedek, G. (2000). Early and late components of visual categorization: an event-related potential study. *Cognitive brain research*, 9(1), 117–119.
- Arik, S. (2000). Stability analysis of delayed neural networks. *IEEE Transactions on Circuits and Systems*, 47(7), 1089-1092.
- Barahona, M., & Pecora, L. M. (2002). Synchronization in small-world systems. *Physical review letters*, 89(5), 054101.

- Barlow, H. B., et al. (1972). Single units and sensation: a neuron doctrine for perceptual psychology. *Perception*, 1(4), 371–394.
- Berger, H. (1929). Über das elektrenkephalogramm des menschen. *European Archives of Psychiatry and Clinical Neuroscience*, 87(1), 527–570.
- Bernander, O., Douglas, R. J., Martin, K. a., & Koch, C. (1991, December). Synaptic background activity influences spatiotemporal integration in single pyramidal cells. *Proceedings of the National Academy of Sciences of the United States of America*, 88(24), 11569–73. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=53177&tool=pmcentrez&rendertype=abstract>
- Bi, G. Q., & Poo, M. M. (1998, December). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 18(24), 10464–72. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/9852584>
- Bienenstock, E. (1995, May). A model of neocortex. *Network: Computation in Neural Systems*, 6(2), 179–224. Retrieved from <http://www.informaworld.com/openurl?genre=article&doi=10.1088/0954-898X/6/2/004&magic=crossref||D404A21C5BB053405B1A640AFFD44AE3> doi: 10.1088/0954-898X/6/2/004
- Bienenstock, E. (1996). Composition. In *Brain theory: biological basis and computational principles*. Elsevier.
- Bienenstock, E., & Doursat, R. (1990). Spatio-temporal coding and the compositionality of cognition. In *Proceedings of the workshop on temporal correlations and temporal coding in the brain, paris, france* (pp. 42–47).
- Bienenstock, E., Geman, S., & Potter, D. (1997). Compositionality, mdl priors, and object recognition. *Advances in neural information processing systems*, 838–844.
- Blodgett, H., & McCutchan, K. (1947). Place versus response learning in the simple t-maze. *Journal of Experimental Psychology*, 37(5), 412–422.
- Blynel, J., & Floreano, D. (2003). Exploring the t-maze: Evolving learning-like robot behaviours using ctrnns. *Applications of Evolutionary Computing*, 593–604.
- Börger, C., Epstein, S., & Kopell, N. J. (2005, May). Background gamma rhythmicity and attention in cortical local circuits: a computational study. *Proceedings of the National Academy of Sciences of the United States of America*, 102(19), 7002–7. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender>

- .fcgi?artid=1100794&tool=pmcentrez&rendertype=abstract doi: 10.1073/pnas.0502366102
- Braitenberg, V. (1984). *Vehicles: Experiments in synthetic psychology*. Cambridge, MA: MIT Press.
- Braitenberg, V. (1988). Discourse on current issues in multi-neuron studies. In A. M. H. J. Aertsen, W. von Seelen, S. . G., & U. M. Leinhos (Eds.), *Organization of neural networks*. Weinheim: VCH.
- Braver, T., Cohen, J., & Servan-Schreiber, D. (1995). A computational model of prefrontal cortex function. *Neural Information Processing Systems*, 141-148.
- Brunel, N. (2000). Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Computational Neuroscience*, 8(3), 183-208.
- Buonomano, D. V., & Merzenich, M. M. (1995). Temporal information transformed into a spatial code by a neural network with realistic properties. *Science*, 1028–1028.
- Burgess, N., Barry, C., & O’Keefe, J. (2007). An oscillatory interference model of grid cell firing. *Hippocampus*, 17(9), 801–812.
- Campbell, S. A., & Kozlovskiy, I. (2012, March). Phase models and oscillators with time delayed coupling. *Discrete and Continuous Dynamical Systems*, 32(8), 2653–2673. Retrieved from <http://www.aims sciences.org/journals/displayArticlesnew.jsp?paperID=7255> doi: 10.3934/dcds.2012.32.2653
- Carr, C. E., & Konishi, M. (1990, October). A circuit for detection of interaural time differences in the brain stem of the barn owl. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 10(10), 3227–46. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/2213141>
- Cessac, B., Paugam-Moisy, H., & Viéville, T. (2010). Overview of facts and issues about neural coding by spikes. *Journal of Physiology-Paris*, 104(1), 5–18.
- Charles, O., Pinville, T., & Doncieux, S. (2012). With a little help from selection pressures: evolution of memory in robot controllers. *Proceedings of Artificial Life XIII*, 1-8.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*(58), 7-19.
- Clark, a., & Thornton, C. (1997, March). Trading spaces: computation, representation, and the limits of uninformed learning. *The Behavioral and brain sciences*, 20(1), 57–66; discussion 66–90. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/10096995>
- Cowan, J. (1972). Stochastic models of neuroelectric activity. In S. Ricce, K. Fread, & J. Light (Eds.), *Statistical mechanics* (p. 181-182). University of Chicago.

- Crook, S. M., Ermentrout, G. B., Vanier, M. C., & Bower, J. M. (1997). The role of axonal delay in the synchronization of networks of coupled cortical oscillators. *Journal of computational neuroscience*, 4(2), 161–172.
- Derrida, B., Gardner, E., & Zippelius, A. (1987). An exactly solvable asymmetric neural network model. *EPL (Europhysics Letters)*, 4(2), 167.
- Diesmann, M., Gewaltig, M.-O., & Aertsen, A. (1999). Stable propagation of synchronous spiking in cortical neural networks. *Nature*, 402(6761), 529–533.
- Douglas, R. J., Martin, K., & Whitteridge, D. (1991). An intracellular analysis of the visual responses of neurones in cat visual cortex. *The Journal of physiology*, 440(1), 659–696.
- Doursat, R. (2013). Bridging the mind-brain gap by morphogenetic neuron flocking: The dynamic self-organization of neural activity into mental shapes. In *2013 aaai fall symposium series*.
- Doursat, R., & Bienenstock, E. (2006). Neocortical self-structuration as a basis for learning. In *5th international conference on development and learning (icdl 2006)* (pp. 1–6).
- Eckhorn, R., Bauer, R., Jordan, W., Brosch, M., Kruse, W., Munk, M., & Reitboeck, H. (1988). Coherent oscillations: A mechanism of feature linking in the visual cortex? *Biological cybernetics*, 60(2), 121–130.
- Elman, J. L. (1993, July). Learning and development in neural networks: the importance of starting small. *Cognition*, 48(1), 71–99. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/8403835>
- Engel, a. K., König, P., & Singer, W. (1991, October). Direct physiological evidence for scene segmentation by temporal coding. *Proceedings of the National Academy of Sciences of the United States of America*, 88(20), 9136–40. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=52667&tool=pmcentrez&rendertype=abstract>
- Fernando, C. (2011). Symbol manipulation and rule learning in spiking neural networks. *Journal of Theoretical Biology*, 275, 29–41.
- Fujii, H., Ito, H., Aihara, K., Ichinose, N., & Tsukada, M. (1996). Dynamical cell assembly hypothesis - theoretical possibility of spatio-temporal coding in the cortex. *Neural Networks*, 9(8), 1303–1350.
- Gautrais, J., & Thorpe, S. (1998, November). Rate coding versus temporal order coding: a theoretical approach. *Biosystems*, 48(1-3), 57–65. Retrieved from <http://linkinghub.elsevier.com/retrieve/pii/S0303264798000501> doi: 10.1016/S0303-2647(98)00050-1

- Gerstner, W. (1996, March). Rapid phase locking in systems of pulse-coupled oscillators with delays. *Physical review letters*, *76*(10), 1755–1758. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/10060509>
- Gerstner, W., & Kistler, W. (2002). *Spiking neuron models: Single neurons, populations, plasticity*. Cambridge University Press.
- Gerstner, W., Kreiter, a. K., Markram, H., & Herz, a. V. (1997, November). Neural codes: firing rates and beyond. *Proceedings of the National Academy of Sciences of the United States of America*, *94*(24), 12740–1. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=34168&tool=pmcentrez&rendertype=abstract>
- Gochin, P. M., Colombo, M., Dorfman, G. A., Gerstein, G. L., & Gross, C. G. (1994). Neural ensemble coding in inferior temporal cortex. *Journal of Neurophysiology*, *71*(6), 2325–2337.
- Gould, S. J., & Lewontin, R. C. (1979). The spandrels of san marco and the panglossian paradigm: A critique of the adaptationist programme. *Proceedings of the Royal Society of London*, *205*(1161), 581-598.
- Gray, C., & Singer, W. (1987). Stimulus specific neuronal oscillations in the cat visual cortex: A cortical functional unit. In *Society of neuroscience abstracts* (Vol. 13).
- Gray, C. M., König, P., Engel, a. K., & Singer, W. (1989, March). *Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties*. (Vol. 338) (No. 6213). Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/2922061> doi: 10.1038/338334a0
- Haider, B., & McCormick, D. A. (2009). Rapid neocortical dynamics: cellular and network mechanisms. *Neuron*, *62*(2), 171–189.
- Hasenstaub, A., Shu, Y., Haider, B., Kraushaar, U., Duque, A., & McCormick, D. A. (2005). Inhibitory postsynaptic potentials carry synchronized frequency information in active cortical networks. *Neuron*, *47*(3), 423–435.
- Hasselmo, M. E. (2008, January). Grid cell mechanisms and function: contributions of entorhinal persistent spiking and phase resetting. *Hippocampus*, *18*(12), 1213–29. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2614862&tool=pmcentrez&rendertype=abstract> doi: 10.1002/hipo.20512
- Hasselmo, M. E., Schnell, E., & Barkai, E. (1995). Dynamics of learning and recall at excitatory recurrent synapses and cholinergic modulation in rat hippocampal region ca3. *The Journal of neuroscience*, *15*(7), 5249–5262.

- Hebb, D. O. (1949). *The organization of behavior: A neuropsychological approach*. John Wiley & Sons.
- Hellwig, B. (2000). A quantitative analysis of the local connectivity between pyramidal neurons in layers 2/3 of the rat visual cortex. *Biological cybernetics*, 82(2), 111–121.
- Hertz, J., & Prügel-Bennett, A. (1996). Learning synfire chains: turning noise into signal. *International Journal of Neural Systems*, 7(04), 445–450.
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8), 2554–2558.
- Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *The Journal of physiology*, 195(1), 215–243.
- Humphries, M., Gurner, M., & Prescott, T. (2006). The brainstem reticular formation is a small world, not scale-free network. *Proceedings of the Royal Society of London*, 273, 503–511.
- Ikegaya, Y., Aaron, G., Cossart, R., Aronov, D., Lampl, I., Ferster, D., & Yuste, R. (2004). Synfire chains and cortical songs: temporal modules of cortical activity. *Science*, 304(5670), 559–564.
- Izhikevich, E. (2003). Simple model of spiking neurons. *IEEE Transactions on Neural Networks*, 14(6), 1569–1572.
- Izhikevich, E. M. (2006a). Polychronization: Computation with spikes. *Neural Computation*, 18, 245–282.
- Izhikevich, E. M. (2006b, February). Polychronization: computation with spikes. *Neural computation*, 18(2), 245–82. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/16378515> doi: 10.1162/089976606775093882
- Izhikevich, E. M., Gally, J. a., & Edelman, G. M. (2004, August). Spike-timing dynamics of neuronal groups. *Cerebral cortex (New York, N.Y. : 1991)*, 14(8), 933–44. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/15142958> doi: 10.1093/cercor/bhh053
- Jaeger, H. (2002). Adaptive nonlinear system identification with echo state networks. In *Advances in neural information processing systems* (pp. 593–600).
- Jeanson, F. (2011). *Coincidence detection: Towards an alternative to synaptic plasticity* (Carleton University Cognitive Science Technical Report). Retrieved from <http://www.carleton.ca/ics/TechReports>
- Jeanson, F., & White, A. (2012). Evolving axonal delay neural networks for robot control. In T. Soule (Ed.), *Proceedings of the fourteenth international conference*

- on genetic and evolutionary computation conference (gecco '12) (p. 121-128). ACM, New York.
- Jeanson, F., & White, A. (2013b). Dynamic memory for robot control via delay neural networks. In *Proceeding of the fifteenth annual conference companion on genetic and evolutionary computation conference companion* (pp. 29–30). New York, NY, USA: ACM. Retrieved from <http://doi.acm.org/10.1145/2464576.2464590> doi: 10.1145/2464576.2464590
- Jeanson, F., & White, A. (2013c). Dynamic memory via delay coincidence detection for robot maze navigation. In *International conference on cognitive modelling (iccm)* (pp. 65–70).
- Jeanson, F., & White, T. (2013a). Dynamic memory for robot control using delay-based coincidence detection neurones. In *Artificial neural networks and machine learning–icann 2013* (pp. 280–287). Springer.
- Jeffress, L. A. (1948). A place theory of sound localization. *Journal of comparative and physiological psychology*, 41(1), 35.
- Jensen, O., Gelfand, J., Kounios, J., & Lisman, J. E. (2002). Oscillations in the alpha band (9–12 hz) increase with memory load during retention in a short-term memory task. *Cerebral Cortex*, 12(8), 877–882.
- Jirsa, V. K. (2008, March). Dispersion and time delay effects in synchronized spike-burst networks. *Cognitive neurodynamics*, 2(1), 29–38. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2289254&tool=pmcentrez&rendertype=abstract> doi: 10.1007/s11571-007-9030-0
- Judd, T. K., & Aihara, K. (1993). Pulse Propagation Networks : A Neural Network Model That Uses Temporal Coding by Action Potentials. *Neural Networks*, 6(02255107), 203–215.
- Kistler, W. M., & De Zeeuw, C. I. (2003, January). Time windows and reverberating loops: a reverse-engineering approach to cerebellar function. *Cerebellum (London, England)*, 2(1), 44–54. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/12882234> doi: 10.1080/14734220309426
- Kliger, S., Yemini, S., Yemini, Y., Ohsie, D., & Stolfo, S. (1995). A coding approach to event correlation. In *Integrated network management iv* (pp. 266–277). Springer.
- König, P., Engel, A., & Singer, W. (1996). Integrator or coincidence detector? The role of the cortical neuron revisited. *Trends in neurosciences*, 19(4), 130–137. Retrieved from <http://www.sciencedirect.com/science/article/>

- Kriener, B., Tetzlaff, T., Aertsen, A., Diesmann, M., & Rotter, S. (2008). Correlations and population dynamics in cortical networks. *Neural Computation*, *20*, 2185–2226.
- Krüger, J., & Becker, J. (1991). Recognizing the visual stimulus from neuronal discharges. *Trends in neurosciences*, *14*(7), 282–286.
- Liu, G. (2004). Local structural balance and functional interaction of excitatory and inhibitory synapses in hippocampal dendrites. *Nature neuroscience*, *7*(4), 373–379.
- Luciana, M., & Nelson, C. A. (1998). The functional emergence of prefrontally-guided working memory systems in four-to eight-year-old children. *Neuropsychologia*, *36*(3), 273–293.
- Maass, W., Natschläger, T., & Markram, H. (2002). Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural computation*, *14*(11), 2531–2560.
- Maass, W., & Schmitt, M. (1999). On the complexity of learning for spiking neurons with temporal coding. *Information and Computation*, *153*(1), 26–46.
- Magnenat, S., Weibel, M., & Bayeler, A. (2007). *Enki: The fast 2d simulator*. <http://home.gna.org/enki/> last accessed: 28-January-2013.
- Masuda, N., & Aihara, K. (2002). Bridging rate coding and temporal spike coding by effect of noise. *Physical review letters*, *88*(24), 248101.
- McCulloch, W. S., & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The Bulletin of Mathematical Biophysics*, *5*(4), 115–133.
- Michel, O. (1995). *Khepera simulator version 1.0 user manual*.
- Murakoshi, T., Guo, J.-Z., & Ichinose, T. (1993). Electrophysiological identification of horizontal synaptic connections in rat visual cortex in vitro. *Neuroscience letters*, *163*(2), 211–214.
- Murthy, V. N., & Fetz, E. E. (1992, June). Coherent 25- to 35-Hz oscillations in the sensorimotor cortex of awake behaving monkeys. *Proceedings of the National Academy of Sciences of the United States of America*, *89*(12), 5670–4. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=49354&tool=pmcentrez&rendertype=abstract>
- Nadasdy, Z. (2012). Binding by asynchrony: the neuronal phase code. *Frontiers in Neuroscience*, *4*(51), 1–11.
- Nelken, I. (1988). Analysis of the activity of single neurons in stochastic settings. *Biological cybernetics*, *59*(3), 201–215.

- Nolfi, S., & Floreano, D. (2000). *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines*. MIT Press.
- Nützel, K., Kien, J., Bauer, K., Altman, J., & Krey, U. (1994). Dynamics of diluted attractor neural networks with delays. *Biological cybernetics*, *56*(1), 553–561. Retrieved from <http://link.springer.com/article/10.1007/BF00198808>
- Panzeri, S., Brunel, N., Logothetis, N. K., & Kayser, C. (2010, March). Sensory neural codes using multiplexed temporal scales. *Trends in neurosciences*, *33*(3), 111–20. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/20045201> doi: 10.1016/j.tins.2009.12.001
- Paugam-Moisy, H., & Bohte, S. (2009). Computing with spiking neuron networks. *Handbook of Natural Computing*, 40p. Springer, Heidelberg.
- Pauluis, Q., Baker, S. N., & Olivier, E. (1999). Emergent oscillations in a realistic network: the role of inhibition and the effect of the spatiotemporal distribution of the input. *Journal of computational neuroscience*, *6*(3), 289–310. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/10406139>
- Peters, A., & Yilmaz, E. (1993). Neuronal organization in area 17 of cat visual cortex. *Cerebral Cortex*, *3*, 49–68.
- Prescott, S. a., Ratté, S., De Koninck, Y., & Sejnowski, T. J. (2006, September). Non-linear interaction between shunting and adaptation controls a switch between integration and coincidence detection in pyramidal neurons. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, *26*(36), 9084–97. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2913017&tool=pmcentrez&rendertype=abstract> doi: 10.1523/JNEUROSCI.1388-06.2006
- Prut, Y., Vaadia, E., Bergman, H., Haalman, I., Slovin, H., & Abeles, M. (1998). Spatiotemporal structure of cortical activity: properties and behavioral relevance. *Journal of Neurophysiology*, *79*(6), 2857–2874.
- Ren, M., Yoshimura, Y., Takada, N., Horibe, S., & Komatsu, Y. (2007). Specialized inhibitory synaptic actions between nearby neocortical pyramidal neurons. *Science*, *316*, 758–761.
- Robinson, D. (1975). Oculomotor control signals. *Basic mechanisms of ocular motility and their clinical implications*, 337–374.
- Rolls, E. T., & Treves, A. (1990). The relative advantages of sparse versus distributed encoding for associative neuronal networks in the brain. *Network: Computation in Neural Systems*, *1*(4), 407–421.
- Roxin, A., Rieke, H., & Solla, S. (2004, May). Self-Sustained Activity in a

- Small-World Network of Excitable Neurons. *Physical Review Letters*, 92(19), 198101. Retrieved from <http://link.aps.org/doi/10.1103/PhysRevLett.92.198101> doi: 10.1103/PhysRevLett.92.198101
- Rumelhart, D. E., & McClelland, J. L. (1986). Parallel distributed processing: explorations in the microstructure of cognition. volume 1. foundations.
- Salin, P. A., & Prince, D. A. (1996). Electrophysiological mapping of gabaa receptor-mediated inhibition in adult rat somatosensory cortex. *Journal of neurophysiology*, 75(4), 1589–1600.
- Scheier, C., Pfeifer, R., & Kuniyoshi, Y. (1998, October). Embedded neural networks: exploiting constraints. *Neural networks : the official journal of the International Neural Network Society*, 11(7-8), 1551–1569. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/12662767>
- Shadlen, M. N., & Newsome, W. T. (1994). Noise, neural codes and cortical organization. *Current opinion in neurobiology*, 4(4), 569–579.
- Shadlen, M. N., & Newsome, W. T. (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *The Journal of neuroscience*, 18(10), 3870–3896.
- Shidara, M., Kawano, K., Gomi, H., & Kawato, M. (1993). Inverse-dynamics model eye movement control by purkinje cells in the cerebellum.
- Singer, W. (1994). Putative functions of temporal correlations in neocortical processing. *Large-scale neuronal theories of the brain*, 201–237.
- Softky, W. R. (1995, April). Simple codes versus efficient codes. *Current Opinion in Neurobiology*, 5(2), 239–247. Retrieved from <http://linkinghub.elsevier.com/retrieve/pii/0959438895800328> doi: 10.1016/0959-4388(95)80032-8
- Softky, W. R., & Koch, C. (1992, September). Cortical Cells Should Fire Regularly, But Do Not. *Neural Computation*, 4(5), 643–646. Retrieved from <http://www.mitpressjournals.org/doi/abs/10.1162/neco.1992.4.5.643> doi: 10.1162/neco.1992.4.5.643
- Somers, S., & Jeanson, F. (2011). Transductionally bounded hierarchical systems. In *Proceedings of the 33rd annual confrence of the cognitive science society* (p. 3376-3380).
- Song, S., Sjöström, P. J., Reigl, M., Nelson, S., & Chklovskii, D. B. (2005). Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS biology*, 3(3), e68.
- Sotelo, C., Llinas, R., & Baker, R. (1974). Structural study of inferior olivary nucleus of the cat: morphological correlates of electrotonic coupling. *Journal of*

- Neurophysiology*, 37(3), 541–559.
- Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs*, 74, 1-29.
- Stevens, F., & Wesseling, F. (1999). Augmentation is a potentiation of the exocytotic process. *Neuron*, 22, 139-146.
- Stevens, M. C. (2009, June). The developmental cognitive neuroscience of functional connectivity. *Brain and cognition*, 70(1), 1–12. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/19185406> doi: 10.1016/j.bandc.2008.12.009
- Suemitsu, Y., & Nara, S. (2003). A note on time delayed effect in a recurrent neural network model. In *Neural computing and applications* (Vol. 11, p. 137-143).
- Tanaka, K. (1992). Inferotemporal cortex and higher visual functions. *Current Opinion in Neurobiology*, 2(4), 502–505.
- Thorpe, S., Delorme, A., & Van Rullen, R. (2001). Spike-based strategies for rapid processing. *Neural networks*, 14(6), 715–725.
- Thorpe, S. J., & Imbert, M. (1989). Biological constraints on connectionist modelling. *Connectionism in perspective*, 63–92.
- Trengrove, C. (2005). *Synfire structures and cognition: Synfire structures and cognition: A complex systems perspective*.
- Vaadia, E., Haalman, I., Abeles, M., Bergman, H., Prut, Y., Slovin, H., & Aertsen, A. (1995). Dynamics of neuronal interactions in monkey cortex in relation to behavioural events. *Nature*, 373(6514), 515–518.
- van Rossum, M. C., Turrigiano, G. G., & Nelson, S. B. (2002). Fast propagation of firing rates through layered networks of noisy neurons. *The Journal of neuroscience*, 22(5), 1956–1966.
- Van Rullen, R., & Thorpe, S. J. (2001). Rate coding versus temporal order coding: what the retinal ganglion cells tell the visual cortex. *Neural computation*, 13(6), 1255–1283.
- Vázquez Marrufo, M., Vaquero, E., Cardoso, M. J., & Gómez, C. M. (2001). Temporal evolution of α and β bands during visual spatial attention. *Cognitive Brain Research*, 12(2), 315–320.
- Vertes, P. E., & Duke, T. (2010, June). Effect of network topology on neuronal encoding based on spatiotemporal patterns of spikes. *HFSP journal*, 4(3-4), 153–63. Retrieved from <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2929633&tool=pmcentrez&rendertype=abstract> doi: 10.2976/1.3386761
- Vogels, T. P., & Abbott, L. (2009). Gating multiple signals through detailed balance

- of excitation and inhibition in spiking networks. *Nature neuroscience*, *12*(4), 483–491.
- Vogels, T. P., Rajan, K., & Abbott, L. F. (2005, January). Neural network dynamics. *Annual review of neuroscience*, *28*(c), 357–76. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/16022600> doi: 10.1146/annurev.neuro.28.061604.135637
- von der Malsburg, C. (1981). The correlation theory of brain function. (*Internal Report*): Goettingen: Max-Planck Institute for Biophysical Chemistry, *2*.
- von der Malsburg, C. (1994). The correlation theory of brain function. *Models of neural networks*, *2*, 95–119.
- von der Malsburg, C., & Bienenstock, E. (1986). Statistical coding and short-term synaptic plasticity: A scheme for knowledge representation in the brain. In *Disordered systems and biological organization* (pp. 247–272). Springer.
- Wang, J. X., & Zochowski, M. (2012, January). Interactions of excitatory and inhibitory feedback topologies in facilitating pattern separation and retrieval. *Neural computation*, *24*(1), 32–59. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/22023193> doi: 10.1162/NECO_a_00220
- Watanabe, M., & Aihara, K. (1997, November). Chaos in Neural Networks Composed of Coincidence Detector Neurons. *Neural Networks*, *10*(8), 1353–1359. Retrieved from <http://linkinghub.elsevier.com/retrieve/pii/S0893608097000373> doi: 10.1016/S0893-6080(97)00037-3
- Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of ‘small-world’ networks. *nature*, *393*(6684), 440–442.
- Ziemke, T., & Thieme, M. (2002). Neuromodulation of reactive sensorimotor mappings as a short-term memory mechanism in delayed response tasks. *Adaptive Behavior*, *10*(3), 185–199.
- Zipser, D., Kehoe, B., Littlewort, G., & Fuster, J. (1993). A spiking network model of short-term active memory. *The Journal of Neuroscience*, *12*(8), 3406–3420.