Bangor University

DOCTOR OF PHILOSOPHY

The role of the basal ganglia in the selection and control of sequential action

Britain, Alfred Alexander

Award date:
1996

Awarding institution:
University of Wales, Bangor
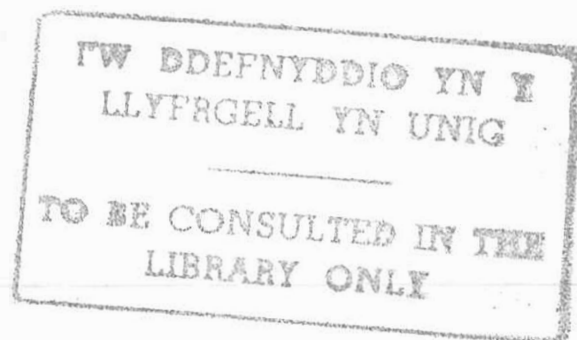
Link to publication

The Role of the Basal Ganglia in the Selection and Control of Sequential Action

Submitted for the Degree of Doctor of Philosophy

Alfred Alexander Britain

August 1996

School of Psychology

University of Wales - Bangor

Bangor, Gwynedd

Abstract

A new conceptual model of fronto-striatal interaction is described. It is suggested that frontal cortex encodes environmental cues that are used in the preparation of sequential motor programs as states of a temporally changing control signal. It is proposed that the basal ganglia use this information to facilitate desired motor patterns selectively and to inhibit competing representations. Similarity relations between states of the control signal are assumed to cause interference amongst motor patterns. In the model, reduction of dopamine causes both a degradation in the quality of the control signal and a failure to resolve competition at output. In computer simulation studies the properties and behaviour of two contrasting connectionist architectures are explored as implementations of the ideas above. The results of damaging each model were investigated to model the reduction of dopamine (as occurs in Parkinson's disease) during the performance of learned movement sequences. In one network the sequential behaviour is driven by recurrent connections from the output of a 'forward model'. The specific effects of damaging the forward model were investigated. A dynamical systems analysis of the patterns of motor interference both in the network and in Parkinson's disease is provided. In the other architecture that is explored, the control signal is composed of multiple endogenous oscillators. This model focuses on the interaction between the control signal and an explicit competitive action selection process. It is shown that degrading either the quality of the control signal or disrupting the competitive processing can yield a variety of deficits that model parkinsonian impairments. An explicit computational account of how cortical and basal ganglia systems interact to subserve both sequencing and selection functions in normal behaviour is developed and it is shown how reduction of available dopamine gives rise to the particular pattern of deficits observed in Parkinson's disease.

Contents

VI

List of Figures

X

To Daphne

Acknowledgements

There are many people whom I wish to thank for their support, assistance and advice during the completion of this research.

First of all I would like to thank my supervisor Gordon Brown for obtaining funding for this project and for his pragmatic advice and support throughout the various stages of this research. I would also like to thank my co-supervisor Mike Malloch whose ideas have influenced many aspects of this research and have had a profound impact on my own thinking in the process.

Those who deserve thanks for hours of useful discussion on the neurobiology of the basal ganglia and its possible role in motor disorders are: Ian Mitchell and Jon Brotchie of Manchester Medical School, Steve Jackson of UW-Bangor and Ann Graybiel of MIT. George Houghton of UCL, Tim Preece of UW-Bangor and Janet Vousden of Warwick University deserve acknowledgement for similar discussions on computational issues surrounding the use of neural networks for sequential processing. I also wish to thank Lew Hardy for motivating me to actually complete the thesis and for dragging me out climbing when it all got too much.

I would like to thank the technical and administrative staff of the psychology department of UW-Bangor. In particular, thanks go to Phyllis Williams and Neil Browning for computer support and to Susie James and Kath Chitty of the administrative office.

A special mention goes to my mother, Daphne, to whom this thesis is dedicated for her invaluable support especially during the bad times and to my friends, Fran and Jack, for managing to live with me during the final stages of writing up.

Finally I would like to thank the numerous family and friends whom I haven't mentioned by name, but whom I've made suffer all the same - I promise you'll never have to hear about it again!

CHAPTER ONE -

Introduction

1.1 Aims and assumptions of the thesis

The primary aim of this thesis is to develop a computational approach to understand the role of the basal ganglia and associated neural circuits in the control of action. A further aim is to increase understanding of disruption to processing mechanisms in motor disorders associated with basal ganglia dysfunction, in particular Parkinson's disease (PD). A supplementary aim is to contribute to existing theories of the control of normal movement.

Much of the behavioural research on PD demonstrates that the disorder primarily affects the initiation and feedforward control of goal-directed movements. Additionally, there is substantial evidence to show that there is a particular deficit in the planning and control of sequential actions.These features of the disorder have led theorists to suggest that the processes involved in motor programming are dysfunctional in PD (see Alexander, Delong & Crutcher, 1992). This in turn has led to a general view that the role of the basal ganglia in normal action control is the selection and organisation of sequential 'motor programs' (e.g. Marsden, 1982; Mink & Thach, 1993).

Although the notion of a motor program has served motor control theory as a useful metaphor for many years, the framework which it represents is based on a computer analogy of human information processing and cognitive representation that has been largely superceded in other areas of psychology. This analogy has failed to make contact with known features of biological organisation. In addition it offers few computational level insights into the *nature* of motor representations and processing mechanisms, or *how* they are disrupted when the system is damaged as it is in movement disorders.

A different approach which has recently gained much support in the motor control field is the 'action systems' approach (Kugler et al., 1982; Saltzman and Kelso, 1987).

This more ecological approach seeks to characterise the rules governing biological movements in terms of the structure and organisation of the environment. In particular these researchers have drawn on principles from dynamical systems mathematics to suggest that motor patterns reflect the dynamics of interacting sets of neural oscillators that fall into stable cyclic patterns. Such stable patterns, or *attractors*, are thought to be the central representation of the action as it is used in cognitive approaches, and are referred to as *coordinative structures*. The precise nature of the attractor depends on the task requirements and the structure of the environment.

Whilst this is an interesting and potentially fruitful approach, which avoids many of the problems with assuming an executive internal controller that plague standard information processing approaches, and can be reconciled with principles of biological organisation, there has been no computational modelling work to suggest how these dynamic representations or control structures may be affected in motor disorders.

Despite the many differences between the two perspectives on motor control theory that have been outlined above, a common goal of both is to characterise the nature of control structures which permit the internal initiation and feedforward execution of complex goal-directed movements. An alternative approach lies in connectionist modelling. A connectionist model can be viewed as an information processing account of a system which is specified in enough detail to explore explicit, constrained assumptions about the nature of input and output representations and to impose constraints on the process of transformation. The internal representations in a connectionist network are developed autonomously by the model. This allows the observer to examine general properties of the system's solution to a given information processing problem. In order to develop connectionist networks that are relevant to modelling movement control or sequential behaviour, a basic requirement is that the network should be capable of generating behaviour that changes over time. There are two broad categories of networks which display this property: recurrent networks and networks with endogenous dynamics. The properties of both of these types of network are explored in this thesis.

The remainder of this chapter serves as an introduction to the key issues that provide the background to the development of this thesis.

1.2 The brain in behaviour

The general notion that brain systems have a functional role in controlling behaviour is not in the least contentious, but precisely how a given neural system is involved in the control of a particular behavioural function is, by contrast, often highly controversial and difficult to determine. In general, the more peripheral structures such as the spinal cord and the retina, which have direct sensory and/or motor functions, are the easiest to define and consequently are the best studied systems. Similarly, some cortical regions which demonstrate consistent responses to electrical stimulation and which are closely linked with peripheral organs have been assigned functional roles with some confidence (e.g. primary motor and visual cortical regions). But because neural systems are not independent, and because many central brain structures form part of more than one system, strict localisation of function is not possible for the majority of brain structures.

In these cases inferential methods have to be used. In cognitive neuropsychology, patients suffering from localised brain damage are studied to assess the behavioural deficits which have been incurred. This method allows the inference to be made that the damaged brain structure is involved in a particular task, although it gives very little indication of *how* the structure participates in normal functioning. For example, the study of H.M. (an amnesia patient suffering from hippocampal damage) revealed the link between the hippocampus and memory, and provided clues about the nature of memory systems. But this type of research gives little information about *how* the hippocampus does its job of storing or retrieving memories in the normal case (see Dudai, 1989).

The information provided by neuropsychology can be enhanced by the development of animal models of human disorders in order to examine the range and nature of deficits produced when the site and extent of lesions are varied. Recently, more sophisticated and specific inferential methods have been developed. For example, regional cerebral blood flow studies (rCBF), using PET or MRI technology illuminate the brain regions that are

active when a behavioural task is being performed. Single cell recordings help to associate the activity of individual cells with overt behavioural activity.

In spite of these advances, comparatively little progress has been made in understanding the behavioural function of central systems of the brain. Central systems have been defined as "cells and circuits that mediate functions necessary for the coordinated behaviour of the whole organism", (Shepherd, 1988, p.488). Included in this definition are systems in which sensory and motor functions overlap, that have no direct connection to peripheral sensory or motor organs, that participate in multiple and complex behavioural functions or in modulating the activity of other brain areas and that are defined by the neurotransmitters that are employed. A typical example of such a system is the nigro-striatal dopamine system and associated cortical - basal ganglia circuits on which we focus here.

1.3 Dopamine and diseases of the basal ganglia

The basal ganglia comprise a group of nuclei that form part of the basal forebrain and midbrain. Basal ganglia structures are typical of a central system as described above. A simplified schematic overview of the anatomical connections of the basal ganglia is provided in figure 1.1.
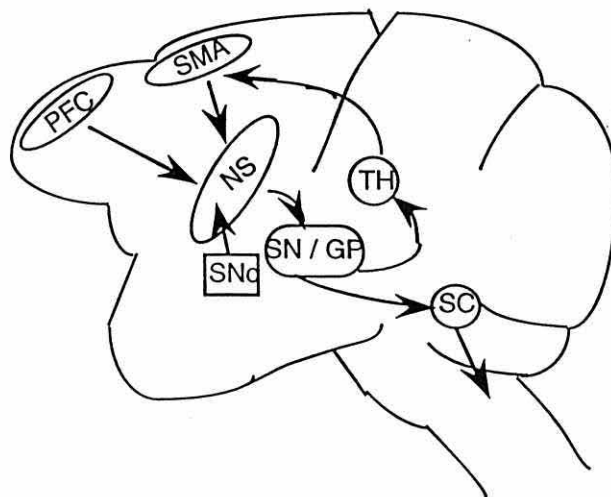
Figure 1.1. *A schematic overview of the basal ganglia and connections*

They do not receive direct input from sensory systems, nor do they project directly to peripheral motor systems. Rather, all of their input is derived from multiple cortical regions. Two such regions are the supplementary motor area (SMA) and the prefrontal cortex (PFC). Output is directed both to the brainstem, via the superior colliculus (SC), and to thalamic nuclei (TH), which in turn have diffuse connections to many cortical regions (e.g. SMA).

The primary site of cortical projections to the basal ganglia is the neostriatum (NS), although additional input is received by the subthalamic nucleus (not shown in figure 1.1). These structures project in turn to the globus pallidus (GP) and substantia nigra pars reticulata (SN). The internal segment of the globus pallidus and the SN collectively form the output structures of the basal ganglia. Additional input to the striatum is provided by the substantia nigra pars compacta (SNc). SNc cells are dopaminergic and this is the origin of the nigro-striatal dopamine system. The impression, then is of a complex system which may be implicated in multiple behavioural functions, although it is in the control of movement that the effects of basal ganglia damage are most commonly observed and most dramatic (Kolb & Whishaw, 1990).

There is a number of motor disorders that can arise from dysfunction of the basal ganglia circuitry (Mitchell 1990; Albin, Young & Penney, 1989). The most prominent idiopathic human disorders are Parkinson's disease (PD) and Huntingdon's chorea (HC). Of the two, PD is the more common, estimated to affect between 0.1% and 1% of the population, and has received more research attention. Whereas HC appears to have a genetic basis and primarily affects striatal neurons and the operation of the STN, PD results from the progressive degeneration of dopaminergic cells in the SNc and the consequent reduction of dopamine in the striatum. The cause of the cell death is unknown.

The movement deficits that occur in the two disorders contrast markedly. PD can be broadly defined as a hypokinesia and is characterised by tremor, rigidity, slowness of movement and an inability to initiate voluntary actions. Huntingdon's Chorea (HC), in contrast, is a hyperkinetic disorder involving inappropriate and excessive movements. Pharmacologically as well as behaviourally hyperkinesias can be viewed as the opposite of

hypokinesias. Whereas hyperkinetic disorders are improved by the administration of dopamine D2 receptor antagonists, and aggravated by dopamine agonists, the reverse is true for hypokinesias such as PD (Albin, Penney & Young, 1989).

From the above evidence it is clear that dopamine plays a crucial role both in maintaining the normal function of the cortico-basal ganglia circuitry and consequently in the normal control of movement. The current outstanding goal of basal ganglia research is to explain what the functional or computational role of dopamine is in the context of basal ganglia operation and what the computational role of the basal ganglia is in maintaining normal behaviour.

1.4 Hypotheses of basal ganglia function

The long association between the basal ganglia and motor disorders has tended to favour theories postulating a direct motor function for the basal ganglia (e.g. Denny-Brown, 1962). In fact, a number of non-motor and cognitive deficits are exhibited by patients with basal ganglia dysfunction, but these have been more difficult to link directly with basal ganglia structures, as opposed to frontal cortical areas which are closely linked with the striatum and become damaged as the disease progresses (Kolb & Whishaw, 1990). Some recent models have suggested roles for the basal ganglia in procedural memory (e.g. St. Cyr & Taylor, 1992) and selective attention (Jackson & Houghton, 1995).

Most of the motor hypotheses of basal ganglia function have been derived to some extent from the deficits exhibited in motor disorders, in particular PD. These have included hypotheses that the basal ganglia are preferentially involved in slow movements, in making postural adjustments, in the initiation and preparation of movement (see Mink & Thach, 1991a). For the most part, such hypotheses involve the suggestion that the basal ganglia play some part in the process of feedforward 'motor programming'.

Some investigators (e.g. Grace, 1983) have argued that the basal ganglia are involved in the "maintenance of stable posture" (cf. also Bronstein, Hood, Gresty & Panagi, 1990), or the integration of proprioceptive inputs (Martin, 1967). Swerdlow and Koob (1987) have argued that the cortico-striato-pallido-thalamic loop is involved in

"sequence switching" and in making use of appropriate contextual information. Other hypotheses have focused on the role of the basal ganglia in assembling actual motor sequences (Marsden, 1987). Precise claims vary, but the general claim that motor programs are, in some way, facilitated or suppressed by the basal ganglia is widespread.

Several researchers have, in contrast, emphasised a possible predictive role for basal ganglia circuitry. Viallet, Trouche, Beaubaton, Legallet and Khalil (1986) argue that the basal ganglia are involved in the feedforward control of movement sequences. Teuber (1976) argued that the basal ganglia are somehow involved in assessing the possible consequences of movements prior to their execution. Thus, he suggests that the output of the basal ganglia to other regions can "preset those other regions for the consequences of concurrent or impending actions.... preparing sensory systems for the kind of input that will be produced by the execution of various efferent actions." (p. 163). Nauta (1971) adopts a similar approach and uses it to account for affective dysfunction seen following some types of basal ganglia pathology. Flowers (1978) found that Parkinsonian patients were less able than controls to use predictive information in a tracking task (following a sinusoidal wave form), and concluded that the basal ganglia are involved in dynamic internal modelling of spontaneous movements thus implementing a system for internal predictive control.

Typically, hypotheses of the type reviewed above are not expressed in any more depth than to suggest that the basal ganglia are 'involved in' some broadly defined process or other. Furthermore they make different underlying assumptions about what the processes subsumed within motor programming are.

A major obstacle for theories of basal function in the motor domain is that the computational processes governing the initiation and execution of voluntary goal-directed movements remain poorly understood (Alexander, Delong & Crutcher, 1992; Jeannerod, 1988). In particular, it is far from clear that the notion of a motor program and the associated 'box and line' diagrams of the control flow of processes in neural systems constitute a fruitful way to enable theoretical development.

1.5 Connectionist models

An alternative to serial or analytic approaches to modelling motor processing in biological systems is provided by connectionist models, also known as parallel distributed processing (PDP) models or neural networks. These network models have, in recent years, been developed to provide a methodology for modelling computational processes in cognitive behaviour. The principles upon which these models rely differ in important ways from the symbolic approaches to cognitive modelling which characterise traditional motor programming theories.

Connectionist models are composed of a number of simple computational units or *nodes* which are connected together, usually in layers, to form a network architecture. Each individual unit has an activation value which is computed as a function of the input it receives from other units and the *weight* on the connection from each incoming unit. The activation level of a unit is also known as its state, and the state of all the units in a given network determine its instantaneous location in a multidimensional state space. Learning in such networks is accomplished by altering the weights on the connections, thus changing the location of the network in state space.

A variety of learning mechanisms may be used, the particular method used being determined largely by the task the network is required to solve. Hebbian associative learning is one of the simplest methods for associating a set of input patterns with a corresponding set of output patterns. The basic form of Hebb's rule is that whenever two units that are joined by a connection are concurrently active, the connection weight is altered to increase the strength of the connection. Another commonly used method, especially in self-organising networks is *competitive learning* in which the units in a response layer compete via lateral inhibitory connections to represent a particular pattern presented at the input layer. Competitive networks have proved especially useful in modelling feature detection in neural systems.

A third class of learning rule is gradient descent learning. The best studied example of this is the backpropagation rule (Rumelhart, Hinton & Williams 1986a). Backpropagation learning is used in combination with multi-layer feedforward networks

composed of units with a non-linear activation function. Learning proceeds by clamping activation values on a set of units designated as input units. Activation is then propagated forwards through the network to a set of output units. The output pattern is then compared with a target output pattern to obtain an error score for each output unit. This error is then propagated backwards through the network to obtain an error for each unit at every layer in the network. This process is repeated for each input pattern in the training set and the resulting summed squared error scores are used to change the weights on the connections to each layer. The entire procedure is repeated until the summed squared error reduces below some criterion value. In effect, over a number of iterations of learning, the weights, which are intially set to small random values, are configured such that boundaries in state space can be formed between the input-output associations that are to be discriminated. The inclusion of a layer of nonlinear *hidden* units interposed between the input and output layers is particularly important because it allows nonlinear discrimination boundaries to be formed. The patterns of activity that develop over the hidden units in a trained network may be thought of as internal representations which are useful in performing the task. This type of network has been used to model a variety of cognitive processes.

A number of general features of connectionist networks has prompted their widespread use as a tool for developing computational models of cognitive function, especially where researchers are explicitly concerned with biological constraints on computational processing that arise from the underlying organisation of neural systems.

• The internal representations that are assumed to exist in theories built on the behaviour of connectionist models are restricted to those that can be formed by the interactive collaboration of a large number of relatively simple computational elements. Even if the analogy with processing in neural circuits is highly oversimplified, as some critics have suggested, it is qualitatively different from assuming symbolic representations with implicit or 'hidden' properties.

• The use of parallel processing in such networks is much closer to brain-style processing than is achieved either by analogy with the operation of a digital computer or the methods used to control the flow of information in sequential computer programs. In

particular it avoids assuming the existence of biologically implausible executive controllers.

- The incorporation of learning mechanisms in networks allows researchers to investigate the computational demands of different tasks and to discover the rules that a network will use to generate a solution, rather than assuming that we know what those rules are a priori.

The above features of connectionist models lead to a number of properties which reflect characteristics of human information processing. These include: graceful degradation of performance in response to localised damage, the ability to generalise from exemplars to categories, and simultaneous satisfaction of multiple constraints (Rumelhart et al., 1986).

1.6 Using connectionist networks to model damage to neural systems

In view of the number of useful properties of connectionist networks in modelling functional or systems-level features of the operation of neural systems, they are natural candidates for modelling functional impairments incurred by nervous system damage or dysfunction. There are several ways in which the effects of damage in connectionist networks can be explored in a way that corresponds to damage in biological systems. For example, noise can be added to the activations of units across the network or restricted to a particular layer to simulate processing disruptions in the presence of noisy or incomplete inputs. Noise can also be added to the weights on connections to simulate retrieval impairments in memory tasks. In addition connections may be lesioned or units removed in order to model localised damage to the system. Qualitative changes in the behaviour of the model may then be examined, often with unexpected or counter-intuitive results.

1.7 Thesis summary

In this chapter we have set out the principal problems with which this thesis is concerned. In particular we have highlighted the limitations of current information processing models of motor control in providing an explicit computational level description of the processes involved in action selection and motor sequencing. We have further argued that an explicit

account will be necessary to account for the disordered motor behaviour in patients with dysfunction of the basal ganglia. The development and exploration of connectionist models that can account for the motor impairments in Parkinson's disease and also suggest explicit information processing mechanisms of the basal ganglia is the goal of this thesis. The remainder of the thesis is organised as follows:

- CHAPTER 2 contains a review of motor disorders with a particular emphasis on parkinsonian impairments in feedforward motor control, action selection and sequential behaviour. A review of neural mechanisms of the basal ganglia and cortex tailored around the same three themes follows.

- CHAPTER 3 reviews current psychological approaches to feedforward motor control and, subsequently, connectionist approaches to modelling sequential behaviour and motor programming.

- CHAPTER 4 presents a conceptual model of the involvement of basal ganglia and prefrontal cortical circuits in the control of feedforward movement. Here we put forward an argument, motivated in part by ecological considerations, that the basal ganglia could have become adapted in the first instance to control action selection and initiation. We then argue that by interaction with cortical circuits the same mechanisms can provide a basis for the control of sequential action. This chapter embodies the hypotheses which are explored in subsequent chapters.

- CHAPTER 5 presents an approach to modelling parkinsonian motor impairments using the Jordan (1990) recurrent forward model network. Simulations are reported in which a technique for simulating dopamine depletion in feedforward connectionist networks is employed in the recurrent network in order to examine how the model responds to damage and whether it can account for patterns of disruption in the control of sequential movement in Parkinson's disease.

- CHAPTER 6 addresses some of the limitations of the standard Jordan approach. A modified version of the model is implemented which incorporates McClelland's (1979) cascade mechanism in an attempt to introduce a temporal dimension to the operation of the model. We also examine alternative methods

for disrupting the operation of the model which are more in keeping with the systems-level focus employed here. A third strand of this chapter explores the properties of the attractors formed by recurrent networks in modelling the representation of sequential actions and to investigate how they are altered when the system is damaged. A dynamical systems interpretation of motor programming and patterns of impairment in Parkinson's disease is proposed in this chapter. The simulations illustrate these ideas using the cascaded Jordan sequential network.

- CHAPTER 7 contains further simulations with a simplified version of the Jordan network to examine how the model responds to damage when there are two effectors and illustrates the need for additional inhibitory mechanisms. Criticisms of the use of recurrent connections in the Jordan architecture are addressed and a non-recurrent version of the model is tested using an endogenous oscillatory control signal.

- CHAPTER 8. An entirely different computational architecture which is based on a complex oscillatory control signal is introduced. This is the oscillator based associative recall model - (OSCAR). We show by simulation that this architecture is capable of modelling sequential behaviour without the complexity inherent in Jordan networks and, we explore properties of the contextual control signal.

- CHAPTER 9. In this chapter we return to the theoretical foundations encapsulated in the conceptual model of chapter four to examine the hypotheses using the OSCAR model. We show by simulation how the contextual control signal and inhibitory "action selection" mechanisms interact to produce a variety of behavioural disruptions when the system is damaged. We discuss the results of these simulations in relation to the original hypotheses.

- CHAPTER 10. General Discussion.

CHAPTER TWO

Motor Disorders and Basal Ganglia Function

The purpose of this chapter is, first, to provide an account of the motor impairments which occur in Parkinson's disease, and second, to present an overview of relevant basal ganglia neurobiology. These strands are brought together in a discussion of constraints on computational models of basal ganglia function and dysfunction in the control of action. In each case the material is organised around the three linked themes outlined in chapter 1. These are i) feedforward control and motor planning, ii) action selection, and iii) the control of sequential behaviour.

Review of motor disorders literature

There are several motor disorders in which the primary site of dysfunction is the basal ganglia. The most common and best studied is Parkinson's disease which results from the degeneration of dopaminergic cells projecting to the striatum. In Huntingdon's disease, the other main idiopathic motor disorder, certain clusters of striatal cells degenerate. In addition, lesions to other parts of the basal ganglia in primates yield a variety of abnormal movements (Mitchell 1990).

In this chapter we focus primarily on the motor deficits that occur in Parkinson's disease (PD). However we also introduce evidence from other disorders and lesioning studies where it provides useful complementary or contrasting evidence regarding basal ganglia motor function in normal movements.

2.1 Impairments of feedforward control in PD

In this section we review evidence that is relevant to the hypothesis that the type of motor processing that is disrupted in PD is feedforward control. Some early studies suggested that the basal ganglia preferentially control slow visually guided movements (Kornhuber, 1971). In contradiction of this view Hallet and Khoshbin (1980) found that ballistic limb movements were affected to a much greater degree than slow 'ramp' movements. Consequently they suggested that PD primarily causes a deficit in feedforward control and in the programming of movements. The findings of Hallet and Khoshbin have since been corroborated by lesioning studies involving ablation of the globus pallidus, (GP). These found that ballistic (feedforward) movements were affected but visually guided movements were not (Hore et al., 1977; Mink & Thach, 1991a).

In the sub-sections below we discuss four main threads of evidence that have been used by researchers to suggest that mechanisms of feedforward motor programming are deficient in PD.

2.1.1 Increased spatial errors when visual feedback is removed

A variety of studies have reported detrimental performance in tracking and pointing tasks when visual feedback of either the hand or the target is temporarily removed (Flowers, 1976; Flowers, 1978; Viallet et al., 1985; Viallet et al., 1987). In these studies the spatial errors committed by the patient group are typically relatively small when visual feedback is available but increase dramatically when it is absent. The type of errors reported in the above studies when the feedback was removed were that the patient groups consistently made movements which would 'undershoot' towards the target. Undershooting has also found to be characteristic of saccadic eye movements in PD (Bronstein & Kennard, 1985).[1]

A variety of hypotheses have been put forward to account for the observed undershooting in PD patients when visual feedback is removed, however they largely fall into two main groups. One hypothesis is that some mechanisms crucial to the internal planning and programming of feedforward movements is disrupted (e.g. Cooke et al.,

---

[1] It should be noted, however, that it has also been found that parkinsonian movements are inherently more variable than normals (Sheridan et al., 1987; Sheridan and Flowers, 1989).

1978). For example it may be that internal representations of sensory or environmental variables affecting task performance are impaired or that the required force and amplitude parameters of the movement cannot be accurately estimated in advance. The second hypothesis is that defective proprioceptive mechanisms force the patient groups to rely more heavily on visual feedback to control their movements. Evidence favouring this position is presented by Viallet et al. (1986) and Viallet et al. (1987). Klockgether and Dichgans (1994) present a study in which they systematically test the effects of removing visual feedback of the moving limb and of the target during pointing movements. They found that the Parkinsonian group only exhibited profound movement slowing and undershooting when visual feedback of the moving hand was prevented, but not when the target location was occluded. These results prompt them to conclude that proprioceptive feedback mechanisms are impaired in PD. However, other interpretations of this data are possible. For example, it was suggested by Sheridan and Flowers (1989) that an internal predictive model of the subjects own movements may be disrupted. Hypotheses of this type would account for the same findings within a feedforward control and motor programming framework.

The issues surrounding what processes properly belong to feedforward or feedback modes of control are not at all as clear as is sometimes suggested in motor control theory (see Cruse et al. 1990). For example it is argued by Viallet et al. (1986) that feedforward control requires the integration of proprioceptive and visual information prior to movement initiation. Thus they attempt to make sense of both the over-reliance of PD patients on visual information and findings that open-loop or 'ballistic' movements are more affected in PD by suggesting that "disturbed proprioceptive information concerning the terminal hand position might .. lead to an underestimation of the target location." They further suggest that during a sequence of movements this processing deficit may account for ever increasing undershooting of the target as the sequence progresses.

Single cell recordings of pallidal neurons in primates during several different movement conditions by Mink and Thach 1991(a) provide evidence that the basal ganglia do not contribute exclusively to a single mode of control although the activation of pallidal

neurons correlated much better with ballistic movements than with closed-loop ramp movements. Further lesioning experiments (Mink & Thach, 1991b) demonstrated that pallidal inactivation primarily affected open-loop movements which depend to a much lesser extent on feedback processes. This data is difficult to accomodate in an account of parkinsonian bradykinesia based solely on non-visual feedback mechanisms.

2.1.2 Inability to prepare movements in advance of initiation

The hypothesis that the basal ganglia are involved in movement pre-programming has also been examined by assessing the extent to which patients can make use of predictive information from the environment to make anticipatory movements in both reaction time and visuomotor tracking tasks. Studies by Bloxham et al., (1984), Jahanshahi et al., (1992) and Viallet et al., (1987) examined simple reaction time (SRT) in PD patients. All found that SRT was significantly slower in the patient group than in age matched control subjects. In an SRT task all the information required to prepare a response fully is available in advance. Because SRT is slowed in PD, it is assumed that the response cannot be pre-programmed adequately. This in turn has led to the view that PD patients cannot make use of advance information to prepare their actions. Additional studies have reported an inability in PD to make use of advance information in predictive tracking tasks, notably Flowers, (1978).

Studies by Day et al. (1984) and Stelmach et al. (1986) have, however, produced contradictory results and have reported that PD patients can make use of predictive information about task requirements to speed reaction times, although the difference is less marked than in normals. Jahanshahi et al., (1992) found that given sufficient time PD's could use full or partial precueing information on both SRT and choice reaction time (CRT). The authors did find that there was an extra slowness of patients in high compatibility uncued CRT condition suggesting that some stage of processing unique to this condition is disrupted.

Whilst the question of the extent to which PD patients can or cannot make use of predictable information has not been resolved in behavioural research (see Jennings, 1995),

significant advances in relation to this question have recently been made in single cell recording studies using primates by Schultz and co-workers. They have found that a significant number of striatal neurons exhibit activity immediately prior to an expected environmental event which has become predictable through conditioning, (Apicella et al., 1992). This indicates that striatal activity is associated, at least in part, with anticipatory preparation in the context of predictive information.

2.1.3. The basal ganglia and the Supplementary Motor Area (SMA)

A further line of evidence which has been used to support the 'planning deficit' view is based on neuroanatomical and neurophysiological evidence of the connections between basal ganglia structures and the SMA. It is well established that the basal ganglia receives a large number of input fibres from the Supplementary Motor Area (SMA), [see fig 1.1]. In addition a portion of pallidal output in primates is directed to the SMA, (Schell & Strick, 1984). Goldberg, (1985) has suggested that the SMA is involved in the planning and preparation of movements prior to their initiation. One of his major sources of evidence is that electrical stimulation of SMA elicits co-ordinated sequences of muscle contractions rather than isolated contractions. Another is that SMA contains a much higher proportion of preparatory or set-related neurons which show sustained activity long before movement initiation than either premotor cortex (PMC) or primary motor cortex (MI) (Tanji & Kurata, 1985). Thus it is assumed that the SMA plays a role in storing or activating higher order motor representations. A further finding implicating the SMA-basal ganglia connection in programming of a response is that of a reduced SMA bereitsshaftpotential or 'readiness' potential in PD (Dick et al., 1989). In particular it is the early part of the evoked potential that is diminished and it is this component that correlates best with the onset of behavioural activity.

However, despite the considerable attention that has been devoted to the links between basal ganglia processing and SMA activity, the relative importance of this cortical area over other premotor areas (such as the ventral premotor area, PMv) and primary motor cortex (MI) has been diminished by recent evidence provided by a retrograde labelling

study (Hoover & Strick, 1993) that pallidal output is organised into discrete channels which project (via thalamic nuclei) to all three of the cortical areas mentioned above. These new findings contribute to the growing evidence that the basal ganglia play a complex role in the control of movement. Hoover and Strick suggest that the pallidal channel to MI may be involved in the direct control of motor output, whereas the premotor channels may be concerned with higher order control activities such as internal guidance and movement sequencing.

2.1.4. Muscular co-contraction in PD

One of the primary clinical symptoms of PD is rigidity in the limbs. The co-activation of agonist and antagonist muscles about the joint appear to be responsible for this condition. This provides some important insights into the parkinsonian deficit in feedforward control. Muscular activation about a single joint during normal voluntary movement is characterised by a reciprocal tri-phasic sequence of activation and inhibition in the flexor and extensor muscles. The duration, amplitude and relative timing of the first agonist burst (AG1) and the antagonist burst (ANT) are known to be independently variable and controlled by the central nervous system without the involvement of feedback mechanisms. This has led many motor theorists to view the triphasic pattern in simple movements as the simplest example of an adaptive motor program (Jeannerod, 1988; Schmidt, 1988).

Various studies have compared the EMG profiles of the triphasic pattern in normal elderly subjects with those of PD subjects in simple ballistic movements, (Berardelli et al., 1986; Hallet & Khoshbin, 1980; Teasdale, Phillips & Stelmach 1990). These studies have found abnormalities in both AG1 and ANT and in the temporal relationships between the two. The first agonist burst was generally small and variable in the PD groups. Instead of increasing the duration of AG1 to perform larger amplitude movements, PD patients exhibit multiple small bursts of activity. Yet, in apparent contradiction, it has also been shown that PD patients can modulate the amplitude and duration of AG activity when explicitly required to do so. Thus the failure to increase agonist activity alone cannot account for the undershooting and multiple step movements seen in PD. Benecke et al. (1987) and

Stelmach, Teasdale and Phillips (1991) have found abnormal antagonist activity either prior to or near the onset of force production. The co-activation of AG and ANT muscles at the initiation of voluntary movement may be due to a failure to inhibit antagonist activity (Stelmach & Castiello, 1992; Wickens, 1993). Lesion work reported in Mink and Thach, (1991b) supports the suggestion that motor deficits in PD could be caused by an inability to inhibit motor patterns selectively, rather than merely provide sufficient activation to the prime mover muscles. Cells in the globus pallidus normally fire constantly and inhibit their target structures in the thalamus. When this tonic inhibition was removed by pallidal ablation in monkeys, the animals exhibited a maintained state of co-contraction which interfered with the initiation of voluntary movements.

The breakdown of the normal relationship between AG1 and ANT in PD provides evidence not only of a feedforward control impairment in PD, but more specifically indicates diminished selective focus in neuromuscular activation and inhibition and a disturbance of mechanisms for maintaining normal sequential relationships. This last point is corroborated by evidence that the normal sequence of motor unit recruitment is disordered in PD (Baker et al., 1992).

The material reviewed above provides evidence of the complex and multi-faceted nature of the involvement of the basal ganglia in movement control and demonstrates the difficulty in attempting to assign a clear functional description of the breakdown in motor processing that occurs in PD. However throughout the material reviewed above there is a thread that strongly indicates that the basal ganglia are involved in motor preparation and feedforward control and that a breakdown in this respect contributes to the motor impairments in PD. Furthermore this material suggests that there is a breakdown in more than one control mechanism in PD. Some of the evidence indicates that a mechanism for initiating components of a motor program is disrupted, whilst other evidence suggests that a mechanism for extracting task-related information from the environment is disrupted.

In the following two sections we consider PD deficits in two functions which depend on intact feedforward control and use of internal representations. The first is action selection and and the second is sequential behaviour.

2.2 PD impairments in action selection

The arrangement of cortico-striato-thalamo-cortical loops which were outlined in chapter 1 show that the basal ganglia project, via thalamic nuclei, to the SMA and other cortical motor areas. This arrangement means that if, as is commonly supposed, these areas of motor cortex are responsible for storing learned motor programs and for activating representations of various body parts, then a key role for the basal ganglia in movement preparation could be to select desired actions by facilitating the appropriate programs maintained at a cortical level. In addition it could also be responsible for suppressing competing action tendencies. This hypothesis has been expressed in various forms in a number of current models of basal ganglia motor function (e.g. Albin, Young & Penney, 1989; Graybiel & Kimura, 1993; Hikosaka, 1991; Mink & Thach, 1993, Wickens, 1993), and is discussed in more depth later in this chapter.

In this section we examine evidence relating to an action selection deficit in PD's, which may arise as a result of a failure to enervate the desired cell populations in motor areas of cortex and to inhibit others sufficiently. Independent work investigating the role of attention mechanisms in action selection tasks have shown that attention serves to inhibit competing action representations during movement preparation, (Tipper, Lortie & Bayliss, 1992). Work by Jackson and Houghton (1995) has linked the basal ganglia to visual attention and has demonstrated attentional impairments in PD. This suggests that action selection impairments in PD may be linked to the involvement of the basal ganglia in attentional mechanisms that rely on inhibition.

Most examples of voluntary action in humans are not as simple as the motor tasks examined in the studies reported in the previous section. Many actions require the co-ordinated action of multiple body parts. These have been termed motor synergies (Bernstein, 1967) and are examples of complex motor programs. Horstink et al. (1990), using bimanual simultaneous motor tasks, found that the ability to share time and shift attention is impaired in PD. One task was visually cued the other was not. The authors found impaired performance in the PD group when the task was not visually-cued. It was concluded that the PD subjects were more impaired when they have to rely on internal

control to shift attention. Two other studies (Cools et al., 1984; Flowers & Robertson, 1985) have reported an impairment of shifting attentional set in PD's, where set corresponds to attentional mechanisms predisposing a subject to consistently select a required action from a number of competing alternatives.

Robertson and Flowers (1990) obtained similar results in a task where PD patients were required to learn two sequences of key presses. Although there was no difference in performance between patients and controls when tested on either sequence individually, the patients made more errors when required to shift between sequences. The authors concluded that the patients were impaired in selecting and maintaining motor set.This study suggests that there is a link between the processes of action selection and action sequencing. In the next section we review impairments in action sequencing in PD.

2.3 PD impairments in sequential control

In the previous section we concluded that PD's have difficulty in action selection, particularly when selection depends on internal cues. We have already hinted at the involvement of the basal ganglia in the programming of sequential movements due to their close involvement with the SMA. A number of studies have provided evidence of impairments in motor and cognitive sequencing in motor disorders.

One of the first studies to investigate a selective impairment in sequential tasks explicitly was that of Benecke et al. (1987). Building on the work of a previous study which found that PD's exhibit an extra impairment when required to perform simultaneous bimanual movements, (Benecke et al., 1986), the authors set out to examine whether the same was true for sequential movements. In the 1987 study patients and controls were given two movements to perform: an elbow flexion, and a squeeze with the hand. The patients were slower at performing each movement individually than controls, but when they were subsequently asked to perform the two movements rapidly in succession, e.g. squeeze then flex or vice versa, the patient group showed a significant extra slowness. The extra impairment was attributed both to an increase in movement time for each component and to an increased pause before initiating the second movement.

In a discussion of these results, the authors suggest that fluent sequential performance is achieved by integrating the two motor programs into a higher-order unit of behaviour represented by a sequential motor plan. It is the motor plan which co-ordinates the switch from one program to the next and times the onset of the second sequence element. PD, they suggest, impairs the ability to organise a motor plan and to switch from one motor program to the next. Unfortunately they are unable to make explicit suggestions about *what* the process of organising a sequential motor plan involves, *how* the switch from one element to the next is made in normal behaviour and *how* these processes might be disrupted in PD. These results would clearly provide a deeper insight into the nature of sequential impairments in PD if a more explicit theoretical framework were available within which they could be interpreted.

An alternative approach which addresses a similar issue is presented in Flash et al. (1992). The aim of this study was to identify the primitive motor elements underlying complex sequential tasks such as handwriting or drawing. They found that a unit for a complex trajectory, such as a curve or loop is formed from the superposition of temporally overlapping simple point-to-point trajectories and that it is preplanned in advance. Furthermore, they suggested that the smoothness and fluidity of a sequence of such movements results from the partial time overlap of the separate elements which are planned in parallel. The authors then tested the performance of PD and control groups on a task which required movement along a simple curved trajectory from an initial position to a target location with a specified intermediate via point. The control subjects exhibited smooth bell-shaped velocity profiles with only a slight slowing at the location of the via-point. In PD subjects, however, the velocity profiles consisted of multiple small peaks and there was a pause at the via point location. In view of these abnormalities the authors suggested that the ability to plan sequential movements in parallel and to superimpose two elemental movement plans is impaired in PD. This suggestion is supported by other studies (Teulings & Stelmach, 1992). In a target switching experiment also reported in the Flash et al. (1992) paper, PD's failed to modify movements in response to the switch. Instead the hand moved all the way to the first target, where it paused before moving to the second target. In some

cases "weird loops and returns were observed when responding to the target switch". This finding is suggestive of interference between two separately planned movements.

Another source of information regarding decomposition of complex sequential movements in PD comes from the study of reaching and grasping. It has been shown (von Hofsten, 1990, 1993) that these two behaviours are coupled at an early stage in development in humans (within the first 6 months after birth). Stelmach and Castiello (1992) report recent findings of decoupling of the two components in PD. Decoupling is inferred from the finding that the onset of the manipulation component is delayed in PD patients. Moreover the delay was significantly reduced by the introduction of a visual perturbation which enforced a reactive mode of control as opposed to a predictive feedforward organisation of the movements.

Two studies have examined how properties of a sequence such as length and complexity affect parkinsonian impairments (Stelmach, Worringham & Strand, 1986; Harrington & Haaland, 1991). Both studies showed that sequence length had less impact on RT in PD's than in normals, suggesting that the sequence is not prepared as a single unit. In a study by Jennings (1995) in which subjects learned two sequences in association with a unique stimulus cue. Cues were presented to subjects before presentation of the signal indicating which sequence was to be performed, thus allowing subjects to prepare a response sequence in advance. The predictive cue could, however, be neutral, valid or invalid thus in some circumstances causing preparation of the wrong sequence. PDs showed a similar cost in RT to controls when the first element of each sequence differed, but not when the first element was the same and the second element differed. This demonstrates that the PD group only prepared the first key press in advance and not the entire sequence.

The studies reviewed above highlight some important points about parkinsonian impairment in the control of sequences.

1. There is considerable evidence to support the suggestion that there is a problem with sequencing of actions in PD over and above those which affect single movements. Yet

even in simple single-joint movements there is a breakdown in the normal sequential relationships of muscle activation leading to coactivation of opposing muscle groups.

2. Some of the studies show that when a sequence is composed of more than one elementary unit of behaviour, PD impairs the ability to transfer from one sequence element to the next. This deficit is particularly noticeable for later components of the sequence where the locus of control is entirely internal and there is no external cue to aid selection. This deficit may be accounted for by a failure to inhibit preceding elements selectively and a concomitant failure to facilitate upcoming elements. Evidence in favour of this hypothesis is gained from the fact that errors can be attributed to intrusions from other learned motor patterns (Robertson & Flowers, 1990; Flash et al., 1992).

3. Other studies show that much of the fluency and speed of normal sequential control is achieved by the parallel activation of two or more basic units of behaviour. This ability appears to breakdown in PD and the individual components are decoupled and separately executed.

4. Whereas normal subjects in some conditions seem to construct a higher order representation controlling execution of the entire sequence, this representation is either destroyed or fails to be used in PD's. The relevant studies reviewed here suggest that in PD only the first element appears to be prepared in advance of movement initiation.

Taken together, the work reviewed in this chapter suggest that the processing functions which are disrupted in parkinsonian motor programming subserve both the selection and sequencing of actions. PD patients are more impaired when they have to rely on internal, predictive control processes than when their actions are cued by information available in the environment. In the next part of this chapter we draw on relevant aspects of biological work concerning the basal ganglia and its interaction with cortex in an attempt to gain a clearer picture of how the normal operation of this brain system is involved in the information processing functions which have been inferred from studying motor disorders.

In considering the neurobiology of the basal ganglia and other brain structures with which it is known to interact, our attention is focussed on the functional role of these systems in behaviour. In keeping with this aim we divide this review according to the three functional concerns addressed above: Internal representation and feedforward control of action; action selection and action sequencing.

Review of Basal Ganglia Neuroscience Literature

Most of the data that has been collected from motor disorders research, including that reviewed above, strongly indicates a primary role of the basal ganglia in the control of movement and that it is particularly important when the movement relies on feedforward or internal control mechanisms. However, it appears from the deficits of PD patients on cognitive tasks tapping functions such as working memory (St. Cyr & Taylor, 1992) that the basal ganglia is not simply a component of a dedicated movement control system. Because the basal ganglia is a deep central system of the brain with no direct connections to sensory or motor organs (Shepherd, 1988), specific and explicit hypotheses of its function have proved hard to formulate. Whilst this group of structures clearly does have a role in the control of action, it is probable that this role is more complex and at a more abstract level than the computation of movement parameters per se.

Hypotheses that suggest that the basal ganglia are in some way involved in the process of 'motor programming' are legion (Alexander et al., 1992; Ito, 1989). Available neuroanatomical evidence demonstrates that the basal ganglia receive inputs from almost all cortical areas, thus supplying the striatum with motor, sensory, planning and motivational information. The outputs of the basal ganglia are directed via one pathway (nigro-collicular system) to brainstem nuclei implicated in the control of eye movements and by another pathway (the pallido-thalamic system) back to multiple cortical sites, in particular motor and premotor areas. Thus the basal ganglia are ideally placed, as noted by Parent, 1990 and Parent and Hazrati, 1993 to capture and integrate multiple sources of information which have an impact on action initiation and selection. The breakdown of such a mechanism may plausibly lead to the variety of dysfunctions observed in motor disorders.

2.4 Cortico-striatal interaction and the internal feedforward control of action

2.4.1 Cortico-striatal loop circuitry

In the past the striatum, with its massive cortical input and restricted pallidal output has been viewed quite plausibly as an information 'funnel' for cortical inputs thus providing an anatomical basis for the idea that the striatum performs an integrative function. In this scheme the inputs from various cortical sites converge onto the same striatal cells.

This view has been challenged following the realisation that the striatum is both neurochemically and anatomically heterogeneous, and that the pattern of cortico-striatal inputs respects the boundaries of striatal divisions. On a gross scale the striatum is composed of three distinct parts, the caudate nucleus; the putamen and the ventral striatum. Most of the striatal input from motor cortical areas, such as supplementary motor area (SMA), ventral premotor area (PMv) and primary motor cortex (MI) is directed exclusively to the putamen (Alexander & Crutcher, 1990) whereas connections from prefrontal and other so-called associational areas of cortex project to the caudate. The ventral striatum appears not to receive inputs from neocortex but has selective links with the limbic system (Defrance, 1980, Graybiel, 1991). It is clear that the whole of the cortical inputs to the striatum are organised in a precise topographical manner such that all areas of the cerebral cortex project to the striatum, but any given striatal area receives an input from just a restricted number of cortical areas (Kunzle, 1975). Subsets of parietal, frontal and limbic cortical afferents are thus segregated into discrete longitudinal zones within the striatum, (Selemon & Goldman-Rakic, 1985). Within the category of sensorimotor projections, the longitudinal segregation is maintained according to a somatotopic organisation, in which leg, arm and face representations are segregated.

Furthermore, there is electrophysiological evidence to suggest that these discrete projection zones are maintained in the projections from the striatum to the globus pallidus (Hedreen and Delong, 1991). These observations have led to a view, that contrasts with the notion of a 'funnel', that the basal ganglia form a series of parallel loops which run from the cortex to the striatum then to the globus pallidus, before returning to specific areas of

cortex via the thalamus. Five independent cortico-striato-thalamo-cortical loops have been described (Alexander, DeLong & Strick, 1986). It has been generally assumed in the 'parallel processing' view that the loops remain segregated as they pass through the striatum to the globus pallidus and thalamus, and back to the cortex.

Each loop is classified in terms of the cortical area to which it projects. In the present chapter we will primarily be concerned with just two of the loops - the 'motor' loop, which receives projections from frontal and midline motor areas, passes through the putamen, and outputs to cortical motor areas (SMA, PMv & MI); and the pre-frontal loop, which receives projections from dorsolateral prefrontal cortex (DLPFC), passes through the dorsal portion of the caudate, and outputs back to DLPFC. The flow structure through these loops, assuming full segregation, is shown in Figure 2.1.
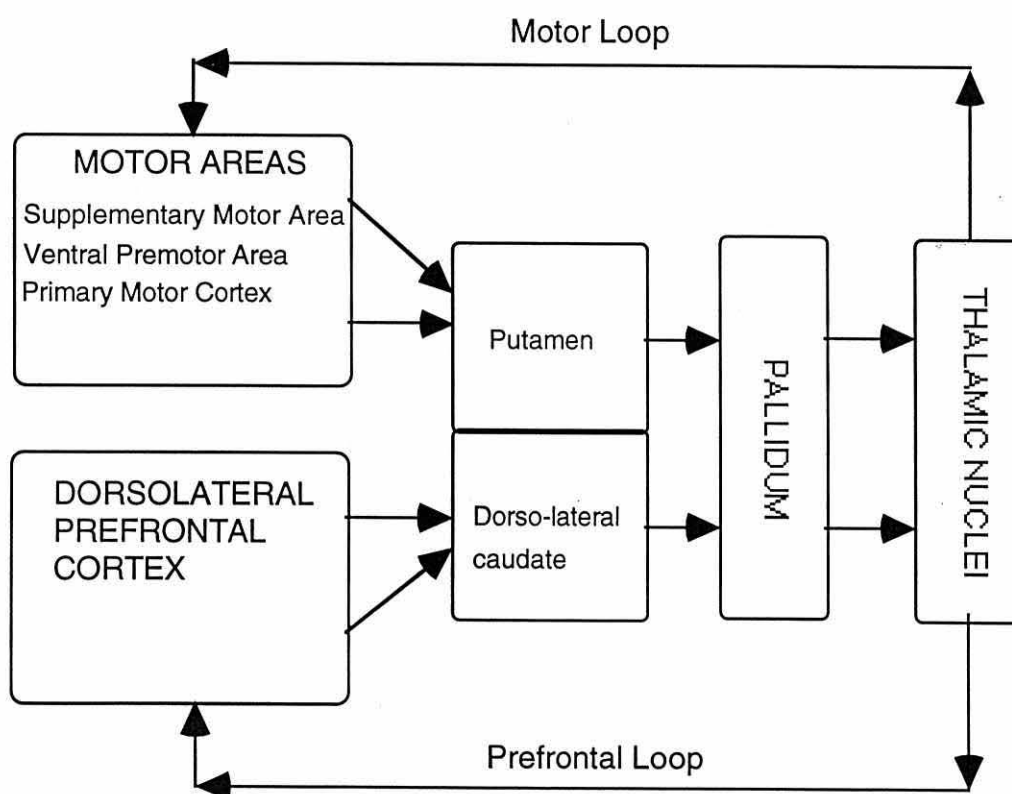
Figure 2.1. *Two of the cortico-striatal loops through the basal ganglia (Alexander et al., 1986). Top of figure: The 'motor' loop; bottom of figure: the 'prefrontal' loop.*

This basic loop architecture of basal ganglia / cortical circuits is interesting from a functional point of view for a number of reasons. The processing of the basal ganglia is entirely internal and there is no direct connection with any direct output system of the brain. It is well established that the cortical areas which form part of the motor loop are involved

in the representation and advance preparation of movements at a number of levels (Requin, 1992). The prefrontal cortex is also strongly implicated in the representation of action information, although at a higher level, and in the planning and control of actions (Goldman-Rakic, 1995).

The loop circuitry of the system provides a form of internal feedback or recurrence which may allow the information represented by cortical neural ensembles to be modified or gated on the basis of the subcortical computations performed by the basal ganglia. This feature is closely associated with the ideas of efference copy or corollary discharge which underlies most theoretical notions of internal comparator or monitor mechanisms in biology (Gray et al. 1991). In the sections below we briefly consider the lower-level structure and functional contribution of each the loops.

2.4.2 The motor loop

The motor loop is the best studied of the five loops that have been identified and is also the system which is most directly relevant to the motor processing functions of the basal ganglia. The largest cortical projections originate in frontal motor areas consisting of SMA and PMv, however there are also projections from MI and SC.

The organisation of motor representations in cortex follow some general principles which are outlined below. It is well established that the midline motor cortical areas (MI and SC) form detailed somatotopic maps of individual body parts and that micro-stimulation of points on the cortical surface reliably cause activation of the corresponding body part. MI is the cortical motor area which is connected with spinal motoneurons and thus acts as the final stage in central motor output. MI neuronal populations have also been shown to encode low-level parameters of movements, in particular direction (Georgopoulos, 1991). Traditional models of the neural motor hierarchy suggest that the more anterior cortical motor areas such as PMv and SMA are involved in representing higher level features of movements which might be termed motor program fragments. Recent electrophysiological evidence supports this notion by showing that when some premotor neurons are stimulated in primates components of complex motor acts such as grasping or pointing are activated,

(Requin, Riehle & Seal, 1992). This finding points to a 'vocabulary' of motor acts represented in frontal areas of motor cortex.

In summary the SMA contains a high proportion of preparatory or set-related neurons which are time-locked to the onset of behavioural actions and fire well in advance of movement initiation. Cells with similar properties have also been recorded in the putamen (Alexander & Crutcher, 1990). The existence of such large numbers of these cells in comparison with other motor areas supports the idea that the motor loop is involved in processes of motor preparation. Regional cerebral blood flow studies using PET technology have demonstrated activity in both SMA and striatum during preparation and execution of learned finger tapping sequences (Roland et al, 1980; Seitz et al,. 1990).

Goldberg (1985) theorises that the SMA forms part of a medial frontal system in the control of action, which functionally contrasts with a lateral system involving premotor cortex. The medial system is concerned with *internally* driven actions, whereas the lateral system is more responsible for the control of stimulus driven behaviour.

2.4.3 The contribution of prefrontal cortex

The functions of the prefrontal cortex are complex, interrelated, and as yet incompletely understood. They possess rich bilateral connections both to other areas of cortex and to the lower levels of the brain, hence their hypothesised regulatory function. The complexity and breadth of the frontal lobe connections no doubt contribute to the puzzling array of behavioural changes observed in people with lesions in this region. The nature of errors made by patients provide clues as to what is involved in frontal lobe processing. On the basis of frontal lobe lesion patients, it has generally been assumed that the frontal lobes are involved in sequencing intentional, goal-directed behaviour. In particular, they become involved in initiating, planning and guiding "non-routine" actions; those that require the default response to be overridden in some way. This view is captured in the Norman and Shallice (1986); Shallice (1988) model of control processes.

This model includes a 'contention scheduling' system, which selects between possible actions simply on the basis of current stimuli and mutual inhibition processes.

However the contention-scheduling process can be influenced by another mechanism, the Supervisory Attentional System (SAS), which can allocate attentional resources to particular schemata on the basis of high-level goals and so on. Thus the SAS does not control behaviour directly. Instead, it is assumed to modulate the contention-scheduling system by inhibiting or exciting schemata within that system. Norman and Shallice assume that the SAS is defective in frontal patients, thus frontal patients have to rely on contention scheduling. As a consequence some frontal patients seem to lack "the will to act"; Shallice argues that this is exactly what would be expected if the SAS was disconnected from the contention-scheduling system. Frith (1992) ascribes the function of contention scheduling to the loops through the basal ganglia, and suggests that action-selection problems in frontal patients reflect dysfunction in the cortico-striatal loops.

The dorsolateral prefrontal cortex, which projects to the striatum, is a portion of the frontal lobe which is traditionally linked more with motor control systems than sensory systems. Goldman-Rakic has proposed that this area of cortex is functionally specialised as a working memory system for the guidance and control of actions based on internal representations (Goldman-Rakic, 1988; 1995). Physiological recordings demonstrate that prefrontal neurons exhibit an increase in firing rate not just during the delay period of delayed response trials, but selective firing in relation to a particular to-be-remembered target (Funahashi et al., 1989). As soon as a motor action is initiated the cells return to their baseline firing rate. This establishes that prefrontal cortex contains stimulus or context specific memory fields which are used in the control of action. Goldman-Rakic argues, in common with Frith, that this distinguishes the role of prefrontal cortex from those systems which control behaviour by association or sensory mechanisms, and that the prefrontal-striatal loop only comes into play when a movement is initiated on the basis of internal representations.

2.4.4 Recent evidence of convergence in striatum

The picture we have elaborated so far emphasises the parallel, segregated nature of cortico-striato-pallidal throughput. If, as some research suggests, this information remains

compartmentalised through pallidal and thalamic stages of the circuitry as well, this raises the question of what processing is actually performed by the basal ganglia. In fact, recent evidence has shown that there is more convergence in the striatum than is assumed by the strict 'parallel loops' model discussed above (Parent and Hazrati, 1993). This convergence is, however, of a constrained and interesting nature.

Flaherty and Graybiel, (1991) have shown that where projections from somatosensory cortex are concerned there is systematic somatosensory remapping of information at the striatum. Further information regarding the functional consequences of somatosensory remapping in the striatum is provided by Brown (1992) and Brown et al. (1994). This study used 2-deoxyglucose (2-DG) methods to measure glucose uptake in rat striatum. This is a high resolution technique for visualising neural activity in awake, behaving animals. They found that when a somatosensory stimulus was applied to the trunk / hindlimb or forelimb of the animal a 2-DG map was produced that revealed different combinations of striatal unit activity and different arrangements of the body part features at different anteroposterior levels of the striatum. In further studies the authors report that the distribution and pattern of activation varies according to the context in which stimuli were applied. Brown refers to their finding as a *combinatorial map* in striatum in which different body parts are grouped together as a functional unit depending on both the nature and context of stimulation.

2.4.5 Implications for action representation and control

The evidence reviewed above suggests that both the prefrontal and motor circuits through the basal ganglia interact in the control of goal-directed actions. The prefrontal circuit acts in an executive role to provide planning and goal information which substitutes for the sensory information that is available when movements are driven by environmental events. In the motor circuit the striatum takes as input a variety of sensory and motor information which is integrated into a combinatorial map of the states of individual body parts that may be involved in an upcoming action. Brown et al. (1994) suggest that this map may form the basis for representing the complex interactions between body parts in skeletomuscular

movement. In the next section we discuss biological evidence that suggests that the intrinsic processing of the basal ganglia contributes to efficient action selection through the combined action of inhibitory and disinhibitory mechanisms.

2.5 Basal ganglia output: action selection through disinhibition

One of the notable features of basal ganglia motor disorders is that they fall into two broad categories: hypokinesias and hyperkinesias. In many ways these two categories of disorder are functional opposites. The hypokinesias, notably parkinsonism, are characterised by lack of movement, rigidity and a form of motor 'flattening' which is best described as the omission or severe reduction of important distinctive features of complex actions. Hyperkinesias, by contrast, are seen as wild or excessive movement in which apparently irrelevant or meaningless movement fragments are inserted into an otherwise meaningful train of action.

An obvious way to make sense of this clinical picture at a systems level would be to suggest that in the normal control of movement it is not enough to simply activate the body parts involved in the movement but that it is also necessary to suppress unwanted movements which might conflict with the intended action. Hyperkinesias such as Huntingdon's disease may be caused by the failure to sufficiently suppress those unwanted movements. The symptoms of Parkinson's disease, on the other hand, may either result from over-inhibition of all movements leading to a lack of spontaneous movement or co-activation of competing movements at a low level leading to rigidity and motor flattening. The scheme of striatal organisation outlined above suggests that different body parts remain segregated in order to selectively activate some representations and suppress others in a competitive framework as suggested by Wickens (1993). The striatal combinatorial map serves as a way to facilitate the parallel activation of several motor components necesssary for complex simutaneous movements or internal control of motor sequences. These represent potentially conflicting demands on striatal processing, which in an noisy system or in the presence of uncertainty about the goals of the action would require an extra filtering mechanism to establish an effective working parameter space. In this section we

review evidence concerning the architecture and operation of the intrinsic circuitry of the basal ganglia motor circuit and a possible role of dopamine in facilitating the process of action selection.

2.5.1 Intrinsic circuitry of basal ganglia

The transmitter systems comprising the internal circuitry of the basal ganglia have been identified. Some pathways are inhibitory (GABAergic) and some are excitatory (glutamatergic). The resulting circuit diagram shows that there are two separate pathways from the striatum to the output nuclei, which are the internal segment of the globus pallidus (GPi) and the substantia nigra pars reticulata (SNr). According to Alexander and Crutcher (1990), the *direct* pathway is inhibitory containing GABA and substance P. Activation of this pathway disinhibits the thalamic stage of the circuit, and thus effectively provides positive feedback to frontal motor areas. Striatal disinhibition of nigrocollicular cells is a required stimulus for the activation of eye movements (Chevalier and Deniau, 1990).

The *indirect* pathway projects initially to the external segments of the globus pallidus (GPe), and is inhibitory consisting of GABA and enkephalin; then it projects to the subthalamic nucleus (STN) via purely GABAergic projections. Finally it is linked to the output nuclei via an excitatory, probably glutamatergic projection from the STN. The GPe exerts a tonic inhibitory influence on the STN. Activation of the inhibitory GABA/ enkephalin projection from the striatum suppresses the activity of GPe neurons and thereby disinhibits the STN increasing excitatory drive to the output nuclei, increasing inhibition in the thalamus, thus providing negative feedback to the frontal motor areas .

To summarise these results, in the model of Alexander and Crutcher (see figure 2.2), the two pathways have opposite effects on the thalamic targets of the basal ganglia and consequently on the cortical regions to which they ultimately project.
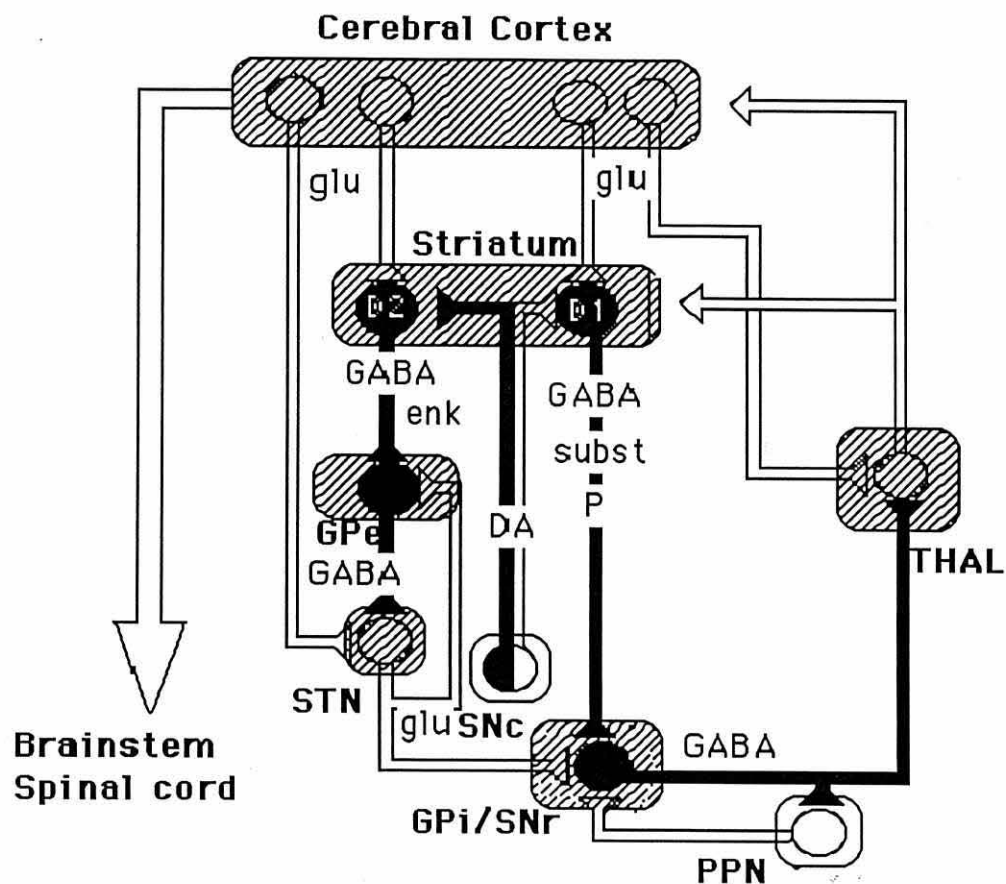
Figure 2.2. *Intrinsic circuitry of the basal ganglia. Black lines indicate inhibitory pathways, White lines indicate excitatory pathways.*

source : Alexander G.E. & Crutcher M.D.(1990)

This arrangement yields an interesting picture in which balanced activity on the two pathways must be maintained for the system to operate optimally. If the net inhibitory 'direct' pathway is overactive and the activity of the 'indirect' pathway is underactive the activity in cortical motor areas will be suppressed and vice versa. It is also known that nigro-striatal dopamine directly influences the activity of both pathways. However it has been shown that the transmission of dopamine to the striatum has opposite effects on the two pathways, (Gerfen et al., 1990). This is due in part to differential distribution of D1 and D2 dopamine receptor subtypes on the two pathways, and also due to the co-release of

different neuropeptides with GABA on the two pathways. It is thus proposed that destruction of dopamine cells by administration of the neuro-toxin MPTP causes parkinsonian symptoms by suppressing the direct pathway and causing the indirect pathway to be overactive. By contrast, hyperkinesias can be explained by the same model if the STN or striatal neurons projecting to the GPe are destroyed, (Albin, Young and Penney, 1989)

2.5.2 A possible role for the basal ganglia in the selective disinhibition of actions

The above model of intrinsic processing in the basal ganglia contributes to the prevailing view that the output of the basal ganglia acts to modulate the activity of cortical circuits responsible for the initiation and execution of motor programs on the basis of input from various areas of cortex and the limbic system, (Graybiel, 1990). Furthermore it has been explicitly proposed that the mechanism of basal ganglia action is to 'disinhibit' its target structures, (Chevaler & Deniau, 1990; Hikosaka, 1991). It is suggested that both the cortical and brainstem targets of basal ganglia output are held under tonic inhibition prior to the onset of movement and that innervation of these structures due to basal ganglia input causes them to be released from inhibition. The clearest evidence in support of this mechanism is derived from studies of the superior colliculus (the major brainstem target of SNr output) in the initiation of saccadic eye movements, (Hikosaka & Wurtz, 1985). It has subsequently been proposed that striato-thalamo-cortical circuits may operate in an analogous way in the initiation and execution of limb movements, i.e. to disinhibit *selectively* those limbs or muscle groups participating in an upcoming movement, (Graybiel, 1991; Graybiel & Kimura, 1995). Although this possibility remains speculative, the existence of GABA mediated inhibitory circuits, which are involved in voluntary movements, in primate premotor cortex, (Matsumura et al., 1991; 1992) provide support for this idea.

An interesting arrangement has been observed in the synaptic connections to and from striatal projection neurons, (Smith & Bolam, 1990). It has been found that cortical afferents terminate in symmetrical synaptic connections on the heads of the dendritic spines.

Dopaminergic afferents tend to terminate on the same cells as cortical projections, but dopamine synapses are commonly made asymmetrically on the spine neck and dendritic shaft closer to the soma than the cortical neurons. This arrangement puts dopamine input in an ideal computational position to modulate striatal output on the basis of cortical input. According to this model dopamine serves to modulate the activity of striatal output neurons to cause selective inhibition of pallidal cells in advance of movement inhibition. This selective, focussed inhibition in GPi is markedly reduced in monkeys with MPTP induced parkinsonism (Filion et al., 1989).

2.5.3 A revised model of basal ganglia processing

Recent additional evidence provided by double anterograde labelling studies adds an extra dimension of complexity to the intrinsic processing of the basal ganglia, above that provided by anatomical or physiological methods, (Parent & Hazrati, 1993; Mink & Thach, 1993). It has been revealed that the projection from the STN to the pallidum is more important than previously thought and that there are substantial direct projections from motor cortex to the STN, avoiding the neostriatum altogether (Kita, 1992). In addition it has been shown that the STN provides widespread excitation to pallidal terminals, whilst the striatal efferents have a more focussed inhibitory effect. These observations serve to enhance the schema of basal ganglia information processing already outlined: that is the selective facilitation of desired motor programs via the focussed inhibition of pallidal neurons by striatal circuits and the inhibition of competing motor programs due to the widespread excitation from the STN. Thus specific populations of pallidal neurons which are tonically active during a maintained posture are turned off prior to the onset of a voluntary movement and the activity of other surrounding neurons is increased, (see Mink & Thach, 1993). This revised model can account for the apparently paradoxical finding that bradykinetic symptoms are reproduced in MPTP-induced parkinsonism where GPi inhibition of thalamic nuclei is increased *and* following GPi ablation (which necessarily yields a reduction in activity). The model suggests that in the former case overactivity in the projection from thalamus to motor cortex would lead to coactivation of motor programs as

suggested in the introduction to this section. In the latter case, there would be an unselective over-inhibition of all motor programs.

The evidence reviewed above provides an account of how dopaminergic activity is involved in striatal processing, and how destruction of dopamine cells affects the output of the basal ganglia. However we have not yet considered the question of what causes an increase in dopaminergic activity in the substantia nigra in the first place. In the next section we consider some recent evidence which bears on this issue.

2.5.4 Dopaminergic activity and responding to environmental context

Theoretical insights into the computational role of dopamine have recently emerged as a result of a series of studies conducted by Wolfram Shultz and co-workers, (e.g. Shultz, Apicella & Ljungberg, 1993; Shultz & Romo, 1990; Romo & Shultz, 1990; Apicella, et al., 1992). These investigators have identified a number of key features concerning the conditions under which single dopamine neurons in the SNc are activated during the acquisition of behavioural tasks in primates.

One of their primary findings is that dopamine neurons do not show phasic changes in activity to movement parameters but to environmental stimuli which have behavioural significance; either primary stimuli such as food or water, or conditioned stimuli which trigger immediate behavioural reactions. In studies where a stimulus precedes a behavioural reaction, DA activity is time-locked to the stimulus rather than the onset of movement, (Shultz, 1992). Moreover dopamine cells apparently do not encode any physical properties of the stimulus itself, such as spatial location. Instead, the phasic activity merely indicates to postsynaptic structures that an event of interest has occurred.

A second important result is that dopamine activity in relation to an environmental event is highly dependent on the behavioural context. Using a "go / no-go" task Shultz and Romo, (1992) found that DA neurons initially fired in response to novel stimuli regardless of behavioural consequences, although the strength of response soon diminishes with repeated presentations. Orienting saccadic eye movements were also observed following the presentation of a novel stimulus. The authors conclude that this response occurs with a

novel stimulus because it has not yet been established whether it is neutral or not. DA neurons responded identically with phasic activation in the go and no-go task conditions. But if identical stimuli were used in a 'no-task' condition the dopaminergic response was absent. This suggests that activation of DA neurons is highly context-dependent. When the stimulus acts as a trigger for an immediate behavioural reaction, whether it be making a movement or withholding a movement, then DA activity occurs. If, however, the no-go condition is not associated with a reward, then the DA activity is extinguished. (Shultz et al., 1992; 1995).

A third finding is that the extent of learning is a key factor in the propensity of an environmental stimulus to elicit DA activity, when that stimulus is not itself a primary reward, or novel in some respect (Ljungberg, Apicella & Shultz, 1992). In this study it was found that intrinsically neutral stimuli that had gained behavioural significance by association with a primary reward were effective in activating DA neurons. As task performance became more automated with overtraining, requiring reduced attentional control, the DA response became progressively reduced. This is interesting as some researchers have suggested that there is high basal ganglia activity only in the early stages of learning a finger tapping sequence (Seitz et al., 1990).

This work suggests that dopamine plays a role in facilitating the preparation or selection of actions on the basis of contextual information available from the environment.

2.6 The role of the basal ganglia in the control of sequential behaviour

Whilst the evidence reviewed above serves to reduce the 'mystery' surrounding basal ganglia function in the control of movement (Marsden, 1982; Mink & Thach, 1993), it still does not make it clear why sequential movements are so much more profoundly affected in PD than simple movements which require internal control. In Huntingdon's disease (HD) too, performance on sequential tasks is more impaired. As one example St Cyr, Taylor and Lang (1988) report that one of the first common complaints of HD patients is that they have trouble in planning sequences of events. In both PD and HD cognitive deficits are most notable on tasks which require advance planning, working memory and some element of

sequential processing such as the Tower of London and Wisconsin Card Sorting tasks (St. Cyr and Taylor, 1992). There is not, at present, very much work available which bears directly on this question. A recent single-cell recording study by Kermadi et al. (1993) has provided evidence of context-dependent activity in caudate cells of monkeys who have been trained to learn a sequence of three targets which light up in order and then, after a delay period to push the target buttons in the same order in which they were lit. They found that many of the neurons selectively fired in response to a given button, but only if it was either preceded or followed by a specific other button. This context-sensitive activity could form the basis for striatal control of movement sequences. There is a suggestion in Mink and Thach (1993) that some of the work of Shultz and co-workers reviewed in the previous section and the Kermadi et al., study indicate that there may be a special role of frontal inputs to the anterior striatum in the control of movement sequences, although they state that it is not clear from the biological evidence how the information is used to control the final output of the basal ganglia.

2.7 Summary of reviews

In this chapter we have reviewed evidence concerning the function of the basal ganglia in motor behaviour and the control of action. We have tailored this review around three main themes reflecting aspects of movement control in which the basal ganglia are involved: feedforward control / motor programming, action selection, and motor sequencing.

In the first instance we reviewed experimental evidence from the motor disorders literature with an emphasis on Parkinson's disease. The neural control of feedforward movement involves not only the basal ganglia but also the various cortical motor and prefrontal areas with which it interacts via the pallido - thalamic circuits and importantly the cerebellar-thalamo-cortical system. These two major ascending systems are presumed to act co-operatively to control the eventual motor output (Houk, 1992) and consequently it has proved difficult to establish the link between any specific 'programming' function and the neural system which is responsible.

The evidence we have reviewed above suggests that the basal ganglia system is more involved in action selection and movement initiation. The emphasis on movement initiation is mainly supported by the evidence which shows prolonged reaction times in PD and the existence of a high number of preparatory or set related cells in SMA and striatum, (Alexander & Crutcher, 1990). In contrast Mink and Thach have recently challenged this assumption based on the finding of late timing of pallidal timing, (Mink & Thach, 1991b). They suggest that the set related activity is used to inhibit competing motor programs and to concurrently release the desired motor program from inhibition. They use the analogy of the use of a handbrake when performing a hill-start. This suggestion is consistent with a primary role in action selection for the basal ganglia. Interestingly the PD data reviewed above revealed comparatively few examples of selection deficits and those that have been found use rather indirect measures. Conversely there are numerous examples of selective deficits in sequence control in PD. Sense may be made of this pattern if we assume that the processes involved in motor sequencing are in some way a generalisation or extension of the processes of action selection and initiation. However, a coherent computational level description of how this might occur in the context of basal ganglia motor processing has not yet been achieved.

In the latter part of the chapter we reviewed evidence from neurobiological studies which suggest functional roles for the basal ganglia and cortical circuitry in the control of movement. The themes which emerge from this evidence are:

1. The basal ganglia are closely linked to several regions of neocortex which are thought to play some role in feedforward movement control and motor planning. One area of particular significance is the prefrontal cortex. It has been suggested that this area plays a role in working memory and the internal representation of contextual information which is used for motor planning.

2. The basal ganglia have been shown to play a functional role in disinhibition of eye movements. A similar role has been suggested in the selective disinhibition of limb movements.

3. Dopamine plays a complex role in maintaining the normal balance between direct and indirect outputs of the striatum. In addition recent evidence has shown that it may play a role in associating motor responses with environmental cues.

In sum these threads of evidence suggest that the combined frontal / basal ganglia system plays at least two independent but interacting roles in the feedforward control of movement. The first is a predictive role which involves extracting contextual task-related information from the environment so that movement components can be prepared independently of or in anticipation of environmental cues. The second is a mechanism which controls the normal relationship between posture and movement. Normally during a held posture cortical action shemata are held under inhibition. When an event that predicts a response occurs this mechanism serves to provide selective facilitation or pre-activation to cortical schemata which are involved in an upcoming movement, whilst further inhibiting competing alternative actions. Dopamine is crucial to the operation of both systems.

This simplified functional model can potentially account for a number of features of the Parkinson's disease literature. Failure of the first sub-system can account for the unusually high dependence on environmental information in PD and the inability or slowness to utilise predictive cues to speed response preparation. It would also explain why the disease most notably affects voluntary movements. Failure of the second sub-system may account for many of the clinical symptoms of bradykinesia and the co-activation of antagonistic muscle groups during motor execution.

Importantly, this model also yields an account of why sequential movements are especially affected in PD. Sequential movements have a higher dependence on internal control than simple movements in the sense that environmental information may only be available to prepare the first sequence element. Thus if only poor quality predictive or contextual control is available to the system in preparing the later components of the sequence, there will be greater pressure on the mechanism for resolving competition between candidate actions. Since this mechanism too is dysfunctional in PD there may not be sufficient contrast to uniquely select and prepare a single action schema.

CHAPTER THREE

Psychological and Computational Models of Movement Sequencing

In chapters one and two we introduced the general motor programming model which has dominated psychological approaches to the control of movement for much of the last 25 years. Furthermore we outlined a number of ways in which theorists have attributed the motor impairments in PD to a breakdown in the processes of preparing and executing feedforward motor programs. In this chapter we review the motor programming approach in more depth and also compare it with other psychological approaches to human motor control, notably the action systems approach. In an effort to make the inadequacies of these approaches clear we construct a computational formulation of the problem of sequential action control and then we review computational models which have been applied to problems of movement control and sequencing.

3.1 Psychological approaches to movement

3.1.1 Feedback and feedforward control

As a general distinction motor control processes can be divided into two types: *feedback* and *feedforward*. Feedback mechanisms are comparatively well understood and methods derived from engineering control systems theory serve sufficiently well to model the major characteristics of biological feedback control. Feedforward mechanisms, by contrast, have proved much harder to characterise. The fundamental problem with the notion of feedforward control is that a reasonably accurate internal model of the desired movement is required, before the movement is initiated. Unlike feedback control which can be achieved

via on-line error correction, feedforward control presupposes the internal representation of at least certain high level motor primitives.

Early theories of motor control were dominated by Sherrington's notions of the reflex and response chaining. It was considered that feedback signals generated by a reflex action provided the signals for the next action and so on until the task was completed. The alternate motion of legs when walking was considered to be a characteristic example of chained reflex control.

The dominant view that these principles could explain all movement control was overthrown by two developments. The first was Lashley's demonstrations that many movements could be adequately controlled without peripheral feedback (Lashley, 1917) and his call for a motor theory which could account for an account of serial order in sequential behaviour without chaining (Lashley, 1951). The second was the finding that many motor patterns (e.g. locomotion) are generated by autonomous oscillating circuits in the spinal cord and are not affected by deafferentation (Von Holst, 1939 / 1973). These circuits became known as central pattern generators (CPG's). Many other elementary action units can be elicited by stimulation of spinal preparations, indicating that spinal circuits store hard wired 'programs for movement' (Giszter, 1992).

These findings led naturally to the idea that higher order movement control may be accomplished by CNS storage of adaptive motor programs. These programs, unlike the fixed action patterns, would be constructed through practice and learning, would be goal-directed, and would enable the motor system to organise and sequence elementary action-units in such a way as to effect smooth, efficient movements which achieve the task goal without recourse to ongoing feedback as a source of input information. Any theory that involves the construction of motor programs necessarily invokes the concept of stored representation of actions, a system for the organisation of action representations, and a set of processes which manipulate representations in preparation for movement execution. These ideas are reviewed in the sections below.

3.1.2 Action representation

The task of identifying the nature of the representations used by the brain for the feedforward control of movements and how they are learned is a persistent and ongoing problem for motor theorists (Georgopoulos, 1991; Requin, 1992; Rosenbaum, 1991). It is also the subject of much controversy. Some researchers claim that only abstract representations of the action goal are required and that only simple parameters such as muscle stiffness need to be computed, the rest being taken care of by the innate properties of the muscles and skeleton (Bizzi et al., 1991; Bizzi et al., 1992). At the other extreme the most explicit representational schemes are derived from robotics and engineering research in which movement planning requires the calculation of joint torques from estimates of position, velocity and acceleration. (Hollerbach & Atkeson, 1987; Kawato et al., 1987).

At the centre of this debate lies the concept of the motor program. This term was coined at a time when computers were beginning to be used as models or at least metaphors for human information processing, (Keele, 1968). Whether or not Keele's original suggestion that movements are stored in memory as a sequence of instructions to the musculature that can be carried out uninfluenced by peripheral feedback was intended to be taken literally, the resulting program concept has pervaded almost every aspect of motor control research since.

Part of the reason for the popularity of the idea of a dynamic program as opposed to a static memory representation for encoding movements, lies in the inherently sequential nature of movement itself. Even the simplest single-joint movement involves a sequence of muscle activations as recorded by electromyograph (EMG) (Jeannerod, 1988). Since Lashley noted the inability of S_R chaining theory to account for the preservation of serial order in movements under feedforward control, it seemed clear that serial order information must be held in memory, and the motor program offers an intuitive and appealing framework to accommodate this need.

An early example of this type of notion was the memory drum proposed by Henry and Rodgers (1960). According to this model, programming a movement would involve organising an appropriate sequence of basic action units or sub-programs. How this may be

achieved is not made explicit by the model. A prediction of the model is that the time to prepare a movement will depend on the complexity of the motor task. This may refer to the length of a sequence or the number of sub-programs involved.

Appealing as it is, the motor program concept has a number of drawbacks and has received substantial criticism over the years. Early criticisms were founded on the problems of storage and novelty. It was presumed that each different movement requires a different motor program and given that we have an almost unlimited capacity for performing different movements, storage is a serious question. The generation of novel movements is a problem if it is presumed that the program defines all aspects of the movement. Schmidt (1975, 1982) attempted to meet these criticisms by proposing a generalised motor program based on schema theory.

According to Schmidt one motor schema represents a movement category as defined by constant relative timing of the constituent sequence despite variation in the time for the whole movement to complete. At a lower level the parameters of the schema which define any instance of a movement, such as duration, amplitude, effector are computed just prior to execution (see Schmidt, 1988). This variation of program theory is undoubtedly more powerful than the original, and can account for a substantial amount of motor learning data, although some of Schmidt's more specific claims have been countered by contradictory evidence (Gentner, 1987 on proportional duration; Wright et al., 1990 on effector independence)

Perhaps a more fundamental weakness in the motor program approach has been the lack of explicit and specific models of how the serial order of the instruction sequence is maintained in memory, particularly when complex movements involving more than one limb, or recruiting several motor programs simultaneously, must be performed.

Another version of the schema hypothesis was presented by Arbib, (1990), who clearly recognised this problem. He suggests that the neural systems involved in motor control compute a coordinated control program which consists of perceptual and motor schemas. Arbib notes that objects may be grasped in a number of different ways, each of which may be regarded as a separate grasping schema. Perceptual schemas encode

information about the properties of an object to be grasped and then serve to gate the various motor schemas, such that an appropriate one is selected. Arbib's contribution is important because he recognises the need to be explicit about the control system which operates over motor programs, moreover he attempts to be explicit both about the neural processes he is concerned with and the computational details of his proposal.

3.1.3 Hierarchical systems of motor programming

An underlying conceptual framework implicit in the motor program view is that the systems governing movement control are embedded in a hierarchy, which allows complex movements are built out of simple elementary units. This framework reflects the influence of language production models in which "word" level units are super-ordinate to component "letter" or "phoneme"- level units. It was seen in the previous section that motor programs can be described at several levels of complexity and the common notion is that higher-level or more abstract schemas exercise selective control over lower-level sub-schemas (see Gallistel, 1980).

This framework is commonly represented as a three layer hierarchy, (Requin, 1992; Shaffer, 1992)

1. SEMANTIC / "WORD-LEVEL" The highest level of programming structure is also the most abstract. This is a representation of the goal of the action and is non-motoric.

2. LEXICAL / "LETTER-LEVEL" The intermediate level involves sequencing and providing context-dependent parameters (e.g. serial order, timing) to basic action units.

3. PHONEMIC / "PHONEME-LEVEL" The lowest level represents the pattern of neuromuscular activations which cause the desired movement to be executed.

This scheme has been widely-adopted, but is it well founded? There is considerable evidence to suggest that in both speech and limb movements the kinematic structure of the movement and force and timing parameters are sensitive to the context in which the movement is performed. In speech the well known phenomenon of co-articulation

demonstrates that the articulatory features for the production will achieve different configurations depending on the phoneme which is to follow. A similar type of "look-ahead" processing has been reported in a study of the kinematics of typing (Flanders & Soechting, 1992). Marteniuk (1987) found that kinematic profiles differed for a grasping task when the subjects either had to subsequently throw the grasped object into a large container or fit it into a small container. Furthermore, analysis of speech errors in which later elements of a sequence are incorrectly inserted too early in the sequence indicates that the sequence is planned as a whole prior to initiation of the first element (Shaffer, 1992). It has also been found in typing studies e.g. (Gentner, Grudin and Conway, 1980) that the onset times for the keystrokes of later letters can precede the onset time for earlier letters.

These studies indicate not only the existence of a hierarchy of action representations but also show that the preparatory organisation of lower level elements of a sequence occurs in parallel, rather than being stimulated one after another as would be predicted by a strict chaining model, of the type criticised by Lashley (1951). The inherent parallelism of sequential movements is one of the characteristic defining features of skilled actions and accounts for the fluidity of over-learned movements as well as the various serial order errors which can occur in their performance. These consist primarily of insertions, omissions and transpositions or exchange errors. It is generally thought that any model of human sequential behaviour should be able to account for the existence of such errors during performance (Houghton, 1990).

One of the basic assumptions of this hierarchical scheme is that the imposition of constraints is a purely top-down process, such that higher levels impose constraints on the lower levels. In contradiction of this assumption there is substantial evidence from the speech processing literature to show that lower level features can affect processing at the higher levels (e.g. Dell and Reich, 1981; Harley, 1984). This data suggests either that the apparent hierarchy is illusory and the characteristics of movement sequences emerge from a parallel system with multiple interacting constraints at and between all-levels, or, as suggested by Gallistel (1980) the hierarchy is *flexible* such that elementary units can exert an influence over high-level units.

3.1.4 Processes involved in motor preparation

Motor programming refers not only to the representations which are used in the feedforward control of action but also to the processes that are involved in preparing an action in advance of movement initiation. It is a general feature of human motor control that any task-specific information that is available in the environment is used to speed the preparation of the movement. For example, if all the information is available in a stimulus to uniquely specify a required response then preparation is shortened to a minimum. Expanding on this principle the Hick-Hyman law (Hick, 1952) demonstrates that the amount of time used for preparation increases in a log-linear relationship with the number of choices available. In terms of information processing this means that as there is less predictive information available from the environment the load on internal response selection and response preparation mechanisms is increased.

Motor response preparation is only part of a serial stage model developed by information processing theorists interested in what happens in the duration between a stimulus or environmental event and a response.
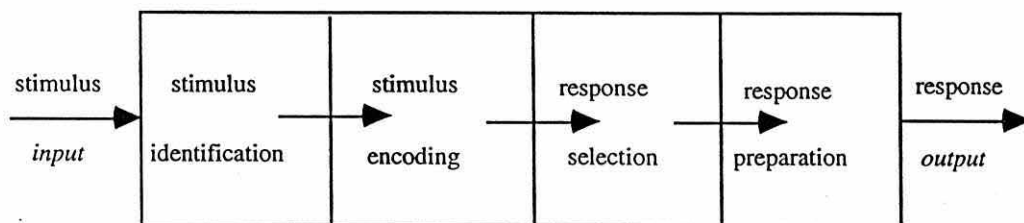
Figure 3.1. *A discrete stage information processing model*

It has been classically assumed that processing is both serial and discrete (Donders, 1869; Sternberg, 1969). That means that processing of a response proceeds from stage to stage and that all processing of one stage must be complete before processing of the next stage begins (see Figure 3.1). These assumptions are strongly supported by Sternberg's additive factors method (AFM), which is consistent with a wide range of the available RT data given the above constraints. Recently, however, evidence has accumulated that processing may

be continuous and that there may be partial output from a stage before processing is complete, (see Mulder, et al., 1992). Moreover, connectionist simulations by McClelland (1979, 1992) showed that the results of applying the AFM is not incompatible with some forms of continuous processing.

These questions are particularly relevant to the current focus on the processes involved in motor preparation and motor disorders of the basal ganglia. The emerging picture of the neural organisation of motor control is increasingly at odds with an ordered serial stage approach to processing and classical conceptions of a hierarchy (Requin, Riehle & Seal, 1992; Requin, 1992). This picture instead reveals an organisation consists of multiple, parallel interconnecting networks operating at overlapping time scales. This has prompted increased enthusiasm for connectionist approaches to modelling motor control, and in particular a recent article (Alexander, Delong & Crutcher, 1992) on basal ganglia involvement in motor programming that called for a rejection of serial models in favour of a parallel approach. Despite some apparent misconceptions over what is actually offered by connectionist approaches, the nature of the argument is clear. More crucially, perhaps, the increased RT's of PD patients in movement preparation warrant a close look at factors affecting the time course of response processing. Even if we accept that there are functional serial stages of information processing intervening between perception and action, it would be unwise in any theoretical account of action / response preparation to treat this stage in isolation from earlier stages.

There is considerable evidence of direct associational links between representations of objects or events and representations of actions. This is sometimes known as 'object affordance'. Following Arbib's schema theory, we can postulate separate perceptual representations of a cup and a thimble. Both representations 'afford' motor schemas for grasping, but a cup would stimulate a different grasping schema from a thimble, (one involving thumb and all fingers, the other thumb and forefinger only). The mere presence of an object does not usually result in execution of this motor schema. But damage to prefrontal cortex can result in utilisation behaviour, a failure to inhibit actions. This

indicates both the existence of associational links between objects and actions, and the existence of neural mechanisms to inhibit actions.

Negative priming experiments have shown that selectively attending to one object, involves concurrent inhibition of other objects competing for attention. Specifically however, it is not the object or the location which is inhibited but the action associated with the object (Tipper, Lortie & Baylis, 1992). This research is important not only because it demonstrates that the processes of motor preparation involve competition between action representations and shows the importance of inhibitory mechanisms to suppress unwanted actions, but also because it suggests a close link between action selection and motor sequencing.

3.1.5 Critique of motor program models

The problem with the information processing paradigm and motor programming perspective that we have discussed above is not that it is clearly wrong. There is a wealth of data across the fields of skill acquisition, motor learning and control that can and has been usefully addressed using this approach. However, there is also an increasing quantity of anomalous data which cannot easily be accounted for by existing models. Abernethy and Sparrow, (1992) point to the failure of Schmidt's (1979) impulse variability model to account for complex multi-joint movements as one example. Another example is provided by Gentner's (1987) critique of the proportional duration hypothesis which is central to Schmidt's motor schema theory. The proportional duration hypothesis states that a single generalised motor program can be identified if the relative timing of its component movements remain constant despite variations in overall movement time. This suggests the existence of a multiplicative rate parameter which allows the motor program to be run off at different speeds. This hypothesis cannot account for phase transitions and other dynamical features of motor co-ordination, (Kelso & Ding, 1993; Mpsitos & Soinila, 1993).

Another weakness of the motor programming approach is that it does not provide a satisfactory account of how the degrees of freedom problem is solved. A characteristic of biological, in particular human, effector systems is that they have excess degrees of

freedom available for the performance of any task which is not completely specified by environmental constraints. For example, there are many possible limb configurations and trajectories which will allow an object in peripersonal space to be picked up. How is a feedforward control program constructed so that it can not only choose among the multiple possibilities, but takes advantage of the excess degrees of freedom to optimise the movement? To date, no satisfactory answer to this question has been forwarded within the motor programming framework.

This last point indicates a strong link between the processes that are involved in the selection of a single action representation or motor program from amongst competing alternatives and the generation of a higher order sequential action in which a number of future actions must be selected, or at least facilitated in advance of initiation. Thus it would seem that the problem of action selection is subsumed within any serious attempt to deal with the problem of the generation of sequential actions.

Even if we are to assume as a simplification of the problem that there are identifiable basic action units represented in memory, independent of external environmental factors, theoretical accounts based on the motor programming perspective are, at present, able to say little about the constraints or rules of interaction which operate between representations. Progress in this area is essential for the development of a coherent account of the programming abnormalities that cause motor impairments in PD.

As a further critical point it is far from clear that the assumption of independence from environmental factors is a sensible and valid simplification. A substantial and growing body of evidence suggests that the representation of actions (and not just constraints on their activation) are tightly bound up with the structure and dynamics of the environment. In the next section we review an approach which is formulated on these principles and which is more concerned with the way that internal and external variables interact to influence the motor pattern that is produced.

3.2 Dynamical approaches to the control of action

An emerging paradigm which contrasts in many ways with the motor programming approach is the *action systems* approach (Kugler, Kelso & Turvey, 1982; Saltzman & Kelso, 1987). According to this view the characteristics of an action arise from the collective behaviour of sets of dynamic control variables related to individual muscles or limbs. These have been termed *coordinative structures* (Turvey, 1977).

A coordinative structure may be regarded as an internal control structure for the autonomous generation of motor acts in the same sense as the central motor programs discussed in the previous section (Cruse et al., 1990), although it is formulated on different fundamental principles. The roots of the approach lie in the ecological psychology of J.J. Gibson (Gibson, 1979) and the mathematics of dynamical systems. Modelling work using this approach commonly involves examining the behaviour of coupled systems of non-linear differential equations describing the motion of limbs during the performance of a specific motor task. This approach places far less emphasis on the importance of central representations of motor actions and considerably more on the regulation of global dynamical parameters such as stiffness and damping.

In the sections below we describe some of the main features of this approach as they are relevant to this thesis.

3.2.1 The equilibrium point hypothesis

The equilibrium point (EQ) hypothesis for the generation of goal-directed limb movements is consistent with the ideas of the action systems approach, and underpins much of the associated work. The EQ hypothesis arose from the use of the metaphor of a tunable spring to describe muscle properties, whereby the spring (muscle) has a resting equilibrium length which is a function of its stiffness. The stiffness parameter and thus the equilibrium length is altered by the neural input to the muscle (Feldman, 1966).

Subsequent development of the hypothesis by Bizzi and co-workers (review in Bizzi et al., 1992) led to the suggestion that feedforward control of point-to-point limb trajectories may be achieved on the basis of a CNS specification of a temporal sequence of

'virtual' equilibrium points. This establishes a virtual trajectory from the current endpoint location to the desired endpoint location. According to this view all that needs to be encoded to compute a trajectory is a final hand position, thus avoiding the need to perform complex inverse kinematic and dynamic computations as is required in other approaches (e.g. Hollerbach & Atkeson, 1987). This hypothesis both deals effectively with the degrees of freedom problem and reduces the computational burden of trajectory formation to the setting of a stiffness parameter. Considerable support for the EQ hypothesis has been derived from experiments on simple trajectories although the model has yet to be extended to account for complex movements involving the co-ordination of more than one limb.

3.2.2 Oscillators and CPG's

The co-ordinated behaviour of control variables in co-ordinative structures has been described and modelled as the interactions of weakly-coupled oscillators (Saltzman & Kelso, 1987). The first to suggest that coordinated action patterns arose from the coupled interaction of internal oscillators was Von Holst (1939 / 1973). Oscillatory mechanisms such as those postulated by Holst are now known to operate as spinal neural circuits and are known as central pattern generators. CPG's autonomously control the rhythmic outputs of the motor neurons during such activities as locomotion. CPG's are thus internal control structures that could equally be described as co-ordinative structures or central motor programs.

3.2.3 Attractor dynamics in biological motor systems.

A key concept of the action systems approach is the notion of an *attractor*. An attractor is a stable state of a dynamical system. There are broadly two types of attractor: *Point* attractors exist where the system moves towards a single point in the space of possible states, where it subsequently remains at equilibrium. *Periodic* attractors exist where the system moves towards a repeating sequence of states. The class of periodic attractors may further be subdivided into *limit-cycle* attractors in which the same sequence of states repeat exactly on each cycle and *chaotic* attractors which exhibit more complex dynamics. Saltzman and

Kelso (1987) suggest that point-attractors may provide a useful way to represent the desired endpoint location of a discrete movement, whereas periodic attractors may instead characterise rhythmic or repeating sequential actions. Furthermore Kelso and Ding (1993) argue that chaotic attractors possess the necessary properties to capture variability and flexibility in co-operative movement dynamics. Mpsitos and Soinila (1993) share this enthusiasm for using dynamical systems concepts to model motor behaviour in biological organisms, this time using sea slugs as a model neural system. These authors echo a point made by Thompson and Stewart (1986) - that a single dynamical system may contain multiple competing attractors. These concepts hold promise for exploring the effects on coherent motor patterns when competition between attractors is not resolved.

3.2.4 Critique of the action systems approach

One of the chief limitations of the action systems approach from a modelling point of view is that it is based on ecological principles which are, at root, anti-representational. It may consequently be argued that in taking this stance the proponents are not addressing a central problem: The need to specify the nature of control structures for feedforward control. Most modelling efforts to date which have been based on this approach have explored the behaviour of coupled dynamical equations of motion at a more abstract level than is applicable to the level of problems we are concerned with here. Also this approach has so far only been applied in detail to restricted classes of movement, notably repetitive or cyclical movements, and has not been extended to account for the control of heterogeneous sequences. Nevertheless the approach has many appealing features and the dynamical systems concepts employed by some researchers, in particular the notion of periodic attractors have recently begun to be employed in some neural network models. Thus elements of this approach provide important constraints for computational models of movement control.

3.3 Connectionist models of movement and sequential behaviour

Much of the appeal of connectionist models in psychology has been based on properties which simulate the mechanisms of human memory and perceptual organisation. In essence

these models associate a set of input patterns with a set of desired output patterns through the use of a learning algorithm. In most cases the learning algorithm causes adaptation of the weights on the connections between input and output units, such that eventually an approximation to the desired output for any input pattern is generated on a set of output units of the network.

Most neural network architectures are inherently static with respect to time, both in terms of processing and output (Mozer, 1993). That is, for a given state of learning in the network the same single input pattern will always give rise to the same single output pattern regardless of what preceded that pattern, or what might follow. This is clearly a fundamental drawback when attempting to model any form of sequential or other dynamic behaviour. A number of schemes have been proposed which are capable of producing a sequence of outputs from a single input pattern. Some of these are specifically oriented towards the motor domain, whilst others deal with the serial order problem from a more abstract perspective. These models will be discussed below in a critical examination of their properties with respect to human motor behaviour. First, however, we consider in more detail the computational problems associated with modelling human sequential behaviour.

3.3.1 Computational models and the serial order problem

Let us, for the moment, disregard the broader problem of movement and concentrate simply on the problem of producing a specified sequence of outputs from a single input pattern.

In the sections above dealing with the motor programming approach, we highlighted the conceptual foundations of this approach. These are: 1. Representations of elementary movement units, be they keystrokes when typing, phonemes in speech, or straight-line segments in trajectory formation. 2. Some serial ordering mechanism which allows recombination of these basic units into different higher-level units which form temporally extended behaviour, e.g. a word, sentence or goal-directed action. 3. A hierarchical scheme of action representation in which the correct serial order of the lower level units of a learned action is recalled by accessing the higher level representation of the whole action.

The reason advanced by Houghton (1994) for the deficits of many psychological models based·on these premises is that hierarchically ordered sets containing copies of passive or inert symbols do not account for empirical evidence regarding human sequential behaviour from a variety of domains. These include speech (Dell, 1986) writing (Caramazza et al., 1987) and typing (Rumelhart & Norman, 1982) (Houghton, 1994; p.126). Moreover it is clear from the material reviewed in chapter two that the nature of the breakdown of sequential action in motor disorders cannot be accounted for by a process model which simply manipulates hierarchical sets of ordered symbols.

At a computational level some important gaps in the key assumptions of psychological theories of motor programming have been revealed by attempts to implement computer models of human serial behaviour using these conceptual primitives. Two main topics are discussed below: the representation of temporal information in connectionist networks, and mechanisms for producing serial behaviour.

3.3.2 Representation of temporal information in networks

Since the most popular neural network architectures, notably the multi-layer perceptron, perform a spatial discrimination operation on the input set, a simple and common approach to representing time in these models is to recast the temporal information as a spatial pattern across the input units. Thus, to represent n different states of the network, the input pattern vector is divided into n blocks, each representing the state of the network at a successive time slice. The entire pattern vector is then processed in parallel by the model. Examples of models using this approach are the verb tense learning model (Rumelhart & McClelland, 1986) and the word recognition model (Seidenberg & McClelland, 1989). Other models which explicitly represent temporal information in this way cited by Elman (1990) are (Cottrell, Munro & Zipser, 1987; Elman & Zipser, 1988; Hanson & Kegl, 1987). Elman points to three major drawbacks with representing time as a spatial pattern in this way:

1. This system crucially presupposes a buffer mechanism which intervenes between the coded information and the world. If the system's task is to effect a simple response to sequentially structured information from the environment, then a buffer is required on

the input side such that all the information can be presented simultaneously. Similarly, if the system is to produce a sequential response as in motor control, then a motor output buffer is presumably required. Both psychological data and biological organisation argue against the existence of such buffer mechanisms.

2. The number of sequence elements or past states is limited by the dimensionality of the input and output vectors. This again is hardwired by the programmer rather than learned by the network and is consequently not a flexible property of the system.

3. This approach does not make it easy for the system to distinguish the same sequential pattern at different temporal positions. In other words the system is sensitive only to the absolute position of a sequence in a temporal stream of information but not relative position.

An alternative approach is to use delay lines on the links between input and hidden units and/or the links between hidden and output units. Although this approach means that the dimensionality of the input vector does not directly code for the different states required at output, the number of delay lines must still be adjusted by external means to the length of the sequence to be produced. Both of these approaches have been found to be unsatisfactory for modelling the properties of human sequential behaviour.

3.3.3 Serial processing mechanisms in connectionist models

The limitations of approaches in which temporal information is explicitly represented in the network have encouraged the development of models which have their own intrinsic dynamics.

Simple Chaining Models

A chaining model is one in which the activation of each consecutive element in the sequence is directly caused by the preceding one. The idea of chaining is derived from early associationist theories of S - R chaining and Sherrington's reflex chaining of responses. The appeal of chaining is that it is the simplest computational scheme by which serial behaviour may be produced without resorting to biologically implausible notions of serial buffers which are assumed in those approaches discussed above.

There are unfortunately a number of serious problems with simple chaining accounts. First they have profound difficulty with repeated elements in a sequence. Take, for example, the letter sequence "E, V, E, R, Y". Here the first instance of "E" is associated with "V", but the second instance of "E" is associated with "R". Thus upon activation of the letter "E", the system has no way to discriminate between the two possible subsequent responses.

One approach that has been used to correct this problem is to introduce *tokens* for different occurrences of identical *types* thus forming a context-sensitive chain. The most widely adopted example of this approach is the Wickelphone representation system (Wickelgren, 1965). A Wickelphone is a tripartite representation of each sequence element such that it is flanked by the preceding and subsequent elements on either side. The Wickelphone representation would then be: "$_E_V$, $_EV_E$, $_vE_R$, $_ER_Y$, $_RY_$". Approaches based on tokens have been heavily criticised (e.g. Houghton & Hartley, 1995) for failing to capture any similarity relationship between the different instances of the same type. Because of this defect they will also fail to reflect interesting behaviour resulting from the interaction and interference of similar sequences or sub-sequences represented in the same associative control structure.

A further count on which chaining systems have been criticised is that they will not exhibit human patterns of error, (such as transitions between elements). As soon as one element is not produced or recalled correctly the system will effectively halt. This behaviour is consistent with evidence that there is hierarchic control of motor sequences (Sternberg et al., 1990).

Activation Gradient Models

A different class of models which overcomes some of the problems with chaining models assumes at least a two-tier hierarchic coding of motor representations. A higher order representation or motor plan activates a set of elementary response units in parallel. Assuming that only one response can be executed at any one time, the system will execute the response with the highest activation value. The activated response units participate in a form of response competition at each time step to yield one response that has a higher

activation value than the others. Once a response has been executed, it inhibits itself and is eliminated from the competition. This basic theoretical scheme owes much to the notion of "contention scheduling", proposed by Norman and Shallice (1986). Various computational schemes have been put forward for implementing this mechanism in which the competition effectively yields a gradient of activation across sequence elements such that they are activated in parallel but executed in the correct serial order (Burgess & Hitch, 1992; Estes, 1972; Grossberg, 1978; Houghton, 1990; Rumelhart and Norman, 1982). Advantages of the activation gradient approach are that the parallel activation of sequence elements gives a natural account of co-articulatory effects, whilst noise in the system during the response competition produces a variety of errors which correspond to those produced in human serial performance in linguistic tasks.

In the sections below, we discuss in detail three specific examples of current models of human sequence production, these are the Rumelhart and Norman (1982) model of typing; the Houghton (1990) competitive queueing architecture and the Jordan (1986a, 1990) models of sequential control of movement.

3.3.4 The Rumelhart and Norman (1982) Model of Typing

Rumelhart and Norman (1982) (hereafter R&N) developed a model which was designed to account for a variety of data available from studies of typing control. This is an interesting task, since it is quite clear from the speeds achieved by expert typists that it is an example of a motor control task with a high component of feedforward control, and it is a prime example of sequential action control.

In addition to accounting for the speed of expert typing the model is intended to account for the high degree of parallelism in typing, similar to co-articulation effects in speech. An example of this is that fingers will move towards the keys for later elements in a sequence before preceding keypresses have been executed (Flanders & Soechting, 1990). Contextual effects are also evident in the timing of individual keystrokes, whereby the stroke time depends on both preceding and succeeding letters. The model also seeks to model error data in typing, in particular transposition, doubling, and alternation reversal

errors. This model is one example of the activation gradient class of models discussed above.
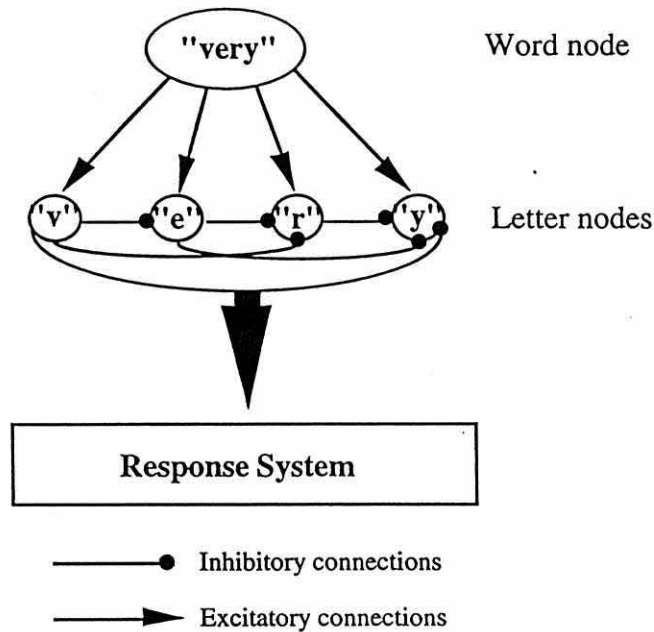
Figure 3.2. *Architecture of the Rumelhart and Norman (1982) typing model*

The R&N model, a simplified illustration of which is shown in figure 3.2, assumes a hierarchical model of motor programming containing words and letter level representations. The letter representations are mapped onto a response system consisting of two-dimensional coordinates for the movements of the hands and fingers in a configuration corresponding to the layout of a standard "QWERTY" keyboard. When a schema for a word to be typed is accessed, the letters which make up the word are all activated in parallel, by topdown excitatory connections from word to letter nodes. Lateral inhibitory connections between the letter nodes then ensure that each successive letter in the word receives more inhibition than the one that precedes it. In this way an activation gradient is set up such that the first letter has the highest initial activation and the final letter has the lowest activation. Some noise is added to the activation values at this stage.

The target positions are fed downwards from the letter nodes and the current locations of the fingers are fed back to the letter nodes. The model adjusts the position of the fingers to minimise the difference between these values in parallel with a force that is

proportional to the activation of the relevant letter node. Whenever the difference for the most highly activated letter node reaches a criterion value a keypress is launched. Following that the letter node corresponding to the keypress is deactivated thus releasing the succeeding letter nodes from inhibition.

This model is a step forward from those reviewed above because it takes seriously the need to account for coarticulatory phenomena and parallel response competition in the production of sequential actions. The serial ordering mechanism forms part of the intrinsic dynamics of the model which also provides the time course of sequence production.

In spite of this there are a number of problems and limitations with the R&N model. First the assumption that only types and not tokens for instances of the same letter are represented in the model means that a special doubling operator must be invoked in order to type words such as "book". This is not in itself a serious problem, but it means that the network has severe difficulties with words which contain two instances of the same letter that are separated by intervening letters e.g. "perception". They get around this problem by assuming that such a sequence would be broken by an unspecified parser into two sections, namely "perc" and "eption". However evidence exists to show that parallel activation effects occur across sequences containing more than one instance of the same letter (Shaffer 1975). Another problem mentioned by Houghton (1994) is that the lateral inhibitory scheme for producing serial order means that words containing the same letters in different orders such as "rat", "tar" and "art" will interfere catastrophically with each other.

A final limitation with the model is that several components involved in its operation are assumed rather than made explicit. For example, the interaction between the letter representations and the response system is unspecified except in general terms. Moreover no account of learning is provided since a learning algorithm is not included and the model must be hand-tuned.

3.3.5 The Houghton (1990) CQ Network

Houghton (1990) has produced a different mechanism known as *competitive queuing* (CQ) which shares some similarities with the Rumelhart and Norman model, whilst correcting a

number of its deficiencies. In particular the CQ approach incorporates a learning mechanism and is thus forced to be computationally explicit about how the activation gradient over letter / item nodes is achieved through learning.
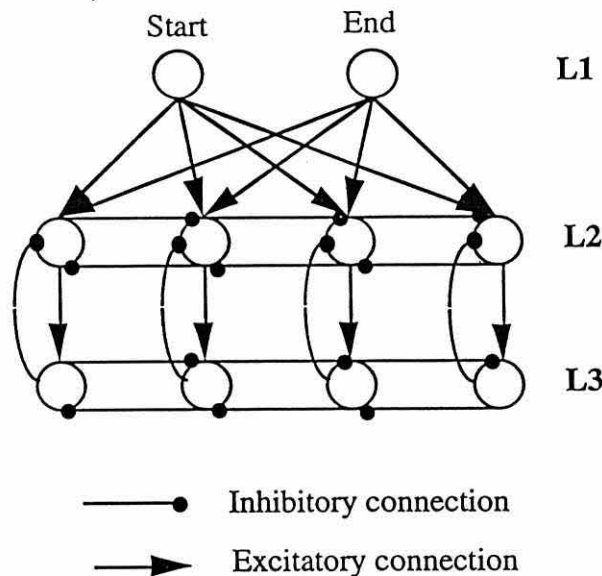
Figure 3.3. *Basic structure of a competitive queuing architecture*

The overall model has a three layer architecture in which layers one and two (L1 and L2) correspond to the word and letter levels of the R&N model. The first major difference is that here the 'word' node is replaced by a pair of nodes denoted "start" and "end" nodes, these are connected via weighted links to a set of item nodes. Time varying activity over the node pair in L1 during learning and recall allows a serially ordered activation gradient to appear over the set of item nodes. The second difference is that although there are lateral inhibitory connections at L2, this is not the primary sequencing mechanism as it is in the R&N model. The lateral connections are bi-directional and in addition there are self excitatory connections on each of the units. This pattern of connections implements an "on-centre, off-surround" field in which each unit feeds back positively onto itself and inhibits all the others. Typically only mild inhibition is used at L2, thus allowing several units to be active simultaneously. L3 is identical to L2. Feedforward connections from L2 to L3 are one to one and excitatory, thus if the weights are set to unity, L3 simply represents a copy of the activations at L2. At L3 however the competition is more selective due to stronger

inhibition in the lateral connections. This competition leads to only a single winner where one node is fully active and all the others are fully inhibited. The final important feature is that there are one to one inhibitory feedback connections from L3 to L2 which mean that when a unit has won the competition at L3 its corresponding unit at L2 is inhibited. Houghton refers to this function of L3 as a *competitive filter*.

During initial learning of a sequence (a word for example), the start node is fully activated when a word is presented to the network, the activation then falls gradually to a minimum at the end of presentation at which point the end node is activated. Thus at each time step of presentation each phoneme in L2 is associated with a different state of activation of the start node. The activity of L2 nodes that have been active gradually decays away, and L2 nodes that have yet to be activated have an activation of 0.

The weights from the start node to L2 are changed using a simple Hebbian rule which only allows increases in the weights. This means that by the end of presentation there is a gradient on the weights form the start node such that the first phoneme in the list has the strongest weight and the last phoneme the weakest. However, when the end node is turned on, the actual pattern of activation over L2 is the inverse activation gradient from the order of presentation (i.e. the last phoneme has the highest activation and the first one the lowest). The end node weights are clamped with a copy of this activation pattern. In sum at the end of a successful learning trial the weight vectors from each of the word nodes encode opposing gradients in the strength of connections across the complete phoneme list.[2]

During recall the start node is fully activated at the start of the sequence and gradually decays as before. End node activity is defined to be the inverse of the activation of the start node for each time step. This time varying control signal causes partial activation of several elements at L2 at each time step, in particular it causes anticipatory pre-activation of upcoming sequence elements. This forms the basis of coarticulation in the network. The competitive filter in L3 ensures that only the most highly activated L2 element is selected for output at any time step, and also suppresses the activity of that element following selection thus maintaining serial order of the output.

---

[2]Houghton (1990) actually describes an additional supervised learning regime for learning more difficult sequences. For simplicity, this is not discussed here as the end result should be the same.

The network described above is the simplest version of a variety of recently proposed models which use these computational principles (Burgess & Hitch, 1992; Houghton, 1994). These models have been shown to provide good accounts of a variety of phenomena associated with human serial behaviour such as coarticulation (Houghton, 1990); serial order errors in spelling (Houghton, Shallice & Glasspool, 1994); memory span, word length effects and similarity in serial short term memory (Burgess and Hitch, 1992).

The central features of the model which are of interest for present purposes is that sequencing is controlled by the joint action of two processes: a dynamical control signal provided by the paired sequence nodes, and inhibitory processing in the competitive filter.

3.3.6 The Jordan (1986) Sequential Network

Jordan (1986a) presented an alternative model of sequence production which, like the models reviewed above, relies on internal time-varying activity in the network to produce the sequential behaviour rather than an explicit representation of past events or hardwired delay connections. However, this model belongs to a class of recurrent architectures which use copies of the past states of the network to drive the sequential behaviour (see also, Elman, 1990). This model too is intended to account for the high degree of parallel activation in different motor components during sequencing, but is based on entirely different computational principles.

The basic Jordan model is an adaptation of a standard three-layer backpropagation network to allow recurrent links from the output units to a set of 'state' or 'context' units which are fully connected to the hidden units and serve as additional inputs to the network at each time slice of the learned operation of the network. The input units are referred to as 'plan' units and remain constant throughout presentation of the entire sequence to the network. The state units, on the other hand, change their value with each sequence element due to the recurrent connections from the output units. A variation of this network also intended to model aspects of sequential behaviour was presented by Elman (1990). In

Elman's model the recurrent connections to the context units originate from the hidden and not the output layer. The basic Jordan architecture is illustrated in figure 3.4 below:
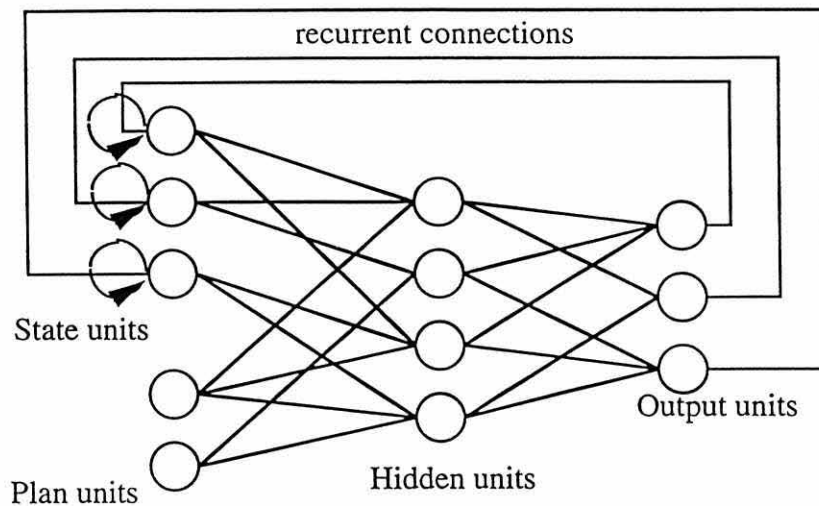
Figure 3.4. The *basic Jordan sequential network architecture (not all feedforward connections shown).*

In terms of a motor programming model, the pattern of activation on the plan units may be thought of as a high-level motor plan which encodes a distributed representation of a sequence, such as a word, to be produced. The output pattern vector encodes each of the sequence elements (e.g. letters if a typing task is used) in turn as the network cycles through the sequence of states. The state units trace the different states of the network.

Learning in the network is exactly the same as backpropagation learning in a standard feedforward network. For each epoch the model is presented with an input vector consisting of the plan + state unit activations (at the first time slice the state unit activations will be zero) and targets for the desired output for each successive time slice in the sequence. At each time slice activation is passed forwards through the network and the output of the network is compared to the target output for that sequence element. Delta values are computed at each layer as the error at the output units is propagated back through the network. Finally the values of the output units are copied back onto the state units. The delta values are accumulated for each element of the sequence and the weights on the internal connections of the network are changed accordingly[3].

---

[3] During training the target output activations are copied back to the state units rather than the actual ouputs. This scheme is implemented by Jordan to speed training time. During testing of the learned

For the sake of clarity figure 3.5 below depicts output and state unit activations at each time step for a network that has learned a single sequence in which each of four output units is turned on in succession and all the others are turned off. Note that if only one sequence is learned, no plan units are required.

Time step	Output units	State units
1	1 0 0 0 (A)	0 0 0 0 (Start state)
2	0 1 0 0 (B)	1 0 0 0 (A)
3	0 0 1 0 (C)	0 1 0 0 (B)
4	0 0 0 1 (D)	0 0 1 0 (C)

Figure 3.5 *Hypothetical output and state vectors for a learned simple Jordan network. The letters in brackets indicate the sequence element.*

In its most basic form as described above the processing of the sequential network strongly resembles the simple chaining models discussed in section 3.3.3. If we assume that only one sequence is being learned and no plan units are included, then the only input to the network comes from the state units. The associations learned by the network are as follows: A is associated with the initial state vector (0 0 0 0), B is associated with A, C is associated with B ... and so on.

The network in this form also suffers from the same criticisms that were levelled at chaining models. Some researchers (e.g. Brown, Preece & Hulme, 1996; Houghton, 1990; Houghton & Hartley, 1995) have suggested that the Jordan model can indeed be viewed as a simple chaining model. This position has, however, never been explicitly tested. There are a number of appealing computational features of the Jordan model and these are discussed below:

<u>Parallel Distributed Processing</u>

A major difference between the Jordan approach to producing sequential behaviour and the activation gradient models we have considered is that it is that it is embedded in a parallel

---

network, however, actual activations are used.

processing system that uses distributed representations. The basic properties of PDP or connectionist systems were reviewed in chapter 1. These include the capability to allow generalisation and similarity amongst stored patterns because of the general property that similar inputs tend to lead to similar outputs. In his theory of serial order (Jordan, 1986a) Jordan assumes that actions are produced in a temporal context composed of actions nearby in time. Knowing the context makes it possible to specify the appropriate action to be produced. He further assumes that a continuity property holds in the evaluation of the next-state function such that similar states are nearby in time. Thus the state vector will evolve continuously over time and is a natural way to represent temporal context in the model . By associating states with desired actions, the generalisations made by the network are to spread actions in time and as learning proceeds there is a tendency towards parallel execution of nearby actions.

Storing the history of past states

We now turn to the mechanism by which continuity in the state representation over time is achieved. In order that the state vector provides a continuous temporal context, it is defined as a function of not just the previous output but a number of previous outputs. If it were merely defined as a function of the previous **n** outputs, this would cause problems with distinguishing repeated subsequences of length **n**. Jordan corrects this problem by defining the state as an exponential function of all past outputs, where the strength decreases with distance in time. This is achieved by providing each state unit with a recurrent connection back to itself with a weight determined by the parameter μ. The value of μ may vary between 0 and 1. With μ set to 0 the system behaves like a simple chaining model in which only the previous output is stored. If the value of μ is large enough the state vector provides a strong temporal context which is potentially capable of allowing recall of a sequence element even if the previous element is recalled incorrectly.

Constraints between output activations

In most feedforward connectionist models a desired value is specified for each output unit. However, this is undesirable in the current model since only a subset of the output units are

likely to have targets at any time slice. For example, if we assume that in a typing task each of the output units represents a finger, only one finger is "active" at any one time. If values have to be specified for the other non-active fingers then the system will be strictly sequential and no overlap of activations will result. If, however, parallelism is an inherent property of the system we may simply allow non-active units to find their own values.

Jordan achieves this by imposing constraints on the output units rather than by explicitly setting target values. If, during learning, an output unit meets the constraints imposed on it then no error score is back-propagated from that unit. Thus in the above example no error would be propagated from the inactive fingers. This simple approach effectively amounts to a 'don't care' constraint on the inactive fingers. Because nearby states in time are similar the non-active fingers will be partially activated to an extent which reflects their position within the sequence and will tend to move towards their target activations at future time slices. The action which immediately follows the current action will move fastest and the final action in the sequence will move most slowly. This means that each keypress is almost fully prepared at the time step preceding execution. In fact, this scheme may be viewed as an alternative way of setting up an activation gradient as seen in the R&N and CQ models, except based on different processing principles. More constraints may be added in to the learning procedure. In general, the tighter the constraints, the less parallelism is exhibited by the system and vice versa.

Attractor dynamics

The final property of Jordan networks that we discuss is the dynamics of learned behaviour in the network. In particular in his 1986a and 1986b papers Jordan shows how the output and next state equations of the network effectively yield a dynamical system in which learned sequences act as attractors in state space. That is, if the network is started in a region of state space near a learned sequence, the network will be drawn towards the learned sequence. Unlike other attractor networks, e.g. Hopfield networks (Hopfield 1982), or Boltzmann machines, (Hinton and Sejnowski, 1986), which converge on a stable point in state space, the Jordan networks with recurrence converge on a stable cycle or limit-cycle corresponding to the learned sequence. This yields the interesting possibility that

is, if several sequences are encoded by the weight matrix of one network then complex patterns of generalisation and interference may occur between sequences.

3.3.7 The Jordan (1990) Model of Motor Learning and Control

Jordan (1988, 1990) set out to address the problem of how a control system might learn the feedforward commands necessary to control a multi-jointed articulated limb with excess degrees of freedom through one or more sequences of goal-directed movements and exhibit desirable properties of human movement such as smoothness, economy of movement and so on. A representation of a two-dimensional task space and a two-jointed limb of the type modelled by Jordan is illustrated in figure 3.6 below.
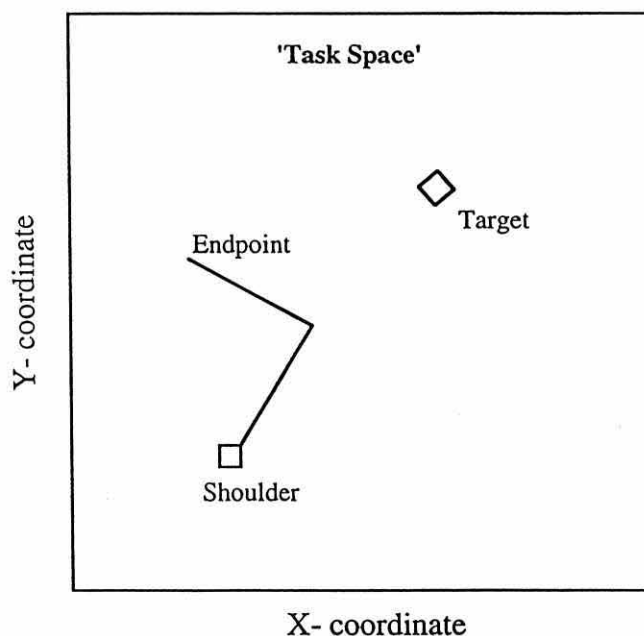
Figure 3.6. *A task space and simulated kinematic arm.*

His approach was to break the computational problem into two stages:

1. Learning to provide appropriate control signals to the limb such that an externally defined movement goal may be achieved, for example touching a point in an extrinsically defined task space with the tip of the limb (note the target in figure 3.6).

2. Learning to translate a higher-order motor plan into a set of control signals which accomplish the sequential task and optimise certain desirable quantities such as smoothness in task space.

Much of the control structure needed to achieve the second phase is already available in the sequential network just described, assuming that the plan units encode a sequential motor plan and the output units encode the kinematic variables which control the movement of the limb.

Solving the first problem is not straightforward because when the system is attempting to learn even a single goal-directed movement, the error information is provided in task space quantities, (e.g. the distance of the tip of a finger from a target) but not in control space coordinates (such as changes in joint angles). Yet an error in control space coordinates is needed to change the weights in the sequential network appropriately.

Moreover the problem is intensified by the fact that there are, in general, fewer degrees of freedom needed to specify a desired goal state, than there are available to control the limb. This means that there is a many-to-one mapping from control space to task space. In other words there are many possible limb configurations which may achieve the same task goal. The problem that arises is one of indeterminacy when the system is required to perform the corresponding inverse mapping from task space to control space which is one-to-many. This problem is a constrained statement of the degrees of freedom problem first recognised by Bernstein (1967).

As an example consider that the task is to reverse an articulated lorry into a loading bay. In this case, the task space error is the distance between the back of the lorry and the end of the loading bay. The control coordinates are the changes that need to be made to the steering wheel. The first problem to solve is to learn how to translate errors in the position of the back of the lorry, into small changes in the angle of the steering wheel.

Jordan's solution to this problem is to augment the sequential controller with a structure he terms a *forward model* (see also Jordan and Rumelhart, 1992). The input to the forward model is provided by the control variables that are the output of the sequential controller and its output is an estimate of the task space location of the tip of the articulator.

The estimate supplied by the forward model may be compared with the desired location of the tip of the effector at that time step which provides a suitable error term which may be propagated back through the forward model thus providing an error term at the control variables which are the output of the sequential controller.

The forward model is implemented as a standard three-layer backpropagation network with logistic hidden units intervening between the input and output layers. Learning of the forward model takes place in a separate initial phase which involves supplying the network with random values for the control variables, comparing the forward model estimate of endpoint location with the observed endpoint location and propagating the resulting error back through the network to change the weights accordingly. This phase of learning may be likened to a baby's developmental phase in which it waves its arms and legs about, apparently at random, in an effort to gauge the effects of its movements on the environment and thus establish a basis for future control. In the lorry example, a forward model may be learned by initially making random small changes to the steering wheel and examining the effect they have on the position of the back of the lorry. This knowledge may then be used to complete the task by running the forward model in reverse (i.e. using the desired and current location of the back of the lorry to determine accurately what changes need to be made to the steering wheel).
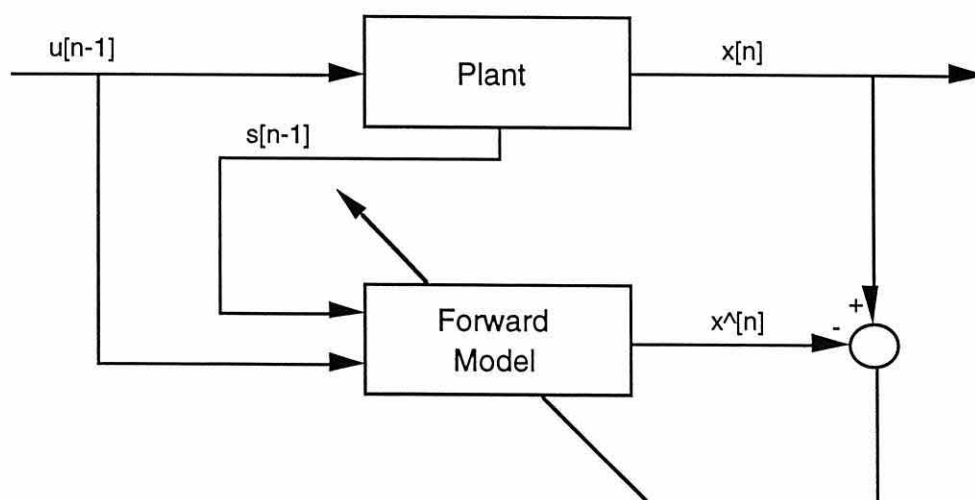
Figure 3.7 *Learning the forward model*

The process of learning a forward model is shown in the control flow diagram of figure 3.7. The plant represents the system to be controlled (e.g. the lorry in the above example). The input to the plant at time n - 1 (u[n-1]) is copied to the forward model as is the current state of the plant (s[n -1]). The forward model prediction of the output (x^[n]) is compared with the actual output (x[n]) and the resulting error is used to adapt the forward model.

Once the forward model has been, at least partially, learned the weights are held fixed and the network may be adjoined to the sequential controller network by identifying the output units of the controller with the inputs to the forward model.
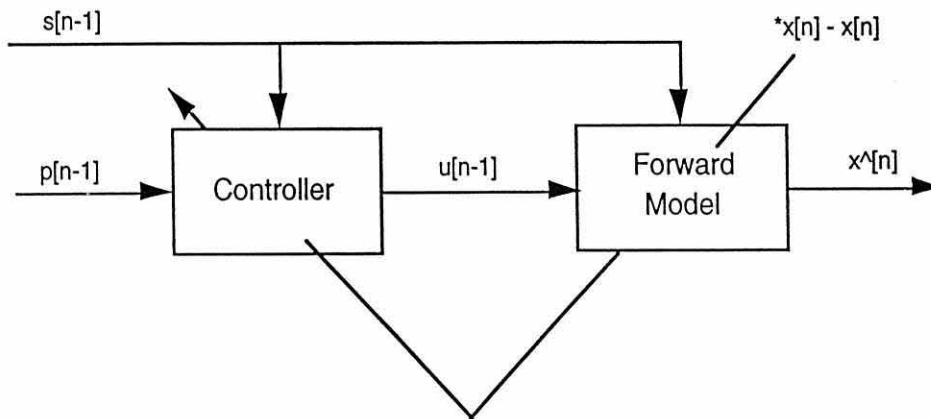
Figure 3.8. *Learning the controller using the performance error (*x[n] - x[n])*

Goal-directed sequential actions may then be learned by forward propagation of activation through the composite network to arrive at an error at the task units. Although the predicted performance error (*x[n] - x^[n]) could be used, if this were the case then the controller could only learn an inverse of the forward model. If the forward model is inaccurate, then the controller will also be inaccurate. The performance error is then propagated back through the fixed weights of the forward model to the articulatory units. The error in control space then flows backwards through the controller and the state units are updated. The state units which feed into the controller subnetwork may receive recurrent connections from either the articulatory units and the task units, or just the task units. At the end of each sequence the weights are changed.

The forward pass of activation through the forward model to provide a task space estimate is essential during the learning process in order to apply one of the many possible sets of control values to the articulatory units so that the inverse mapping can be performed during the backwards pass.

The forward modelling approach provides a flexible and effective solution to the problematic inverse transformation when using supervised learning with excess degrees of freedom. It is superior to other previously tried approaches, notably direct inverse modelling (see Jordan and Rumelhart, 1992 for a fuller discussion). One of the major advantages of this approach is that the forward model does not need to be accurate across the whole task space in order for accurate goal-directed movements to be learned. Additional constraints on the learning process which optimise desired quantities (such as smoothness in task space) can be incorporated into network processing as extra error terms at either the task units or the articulatory units.

Jordan has extended this approach in the domain of movement control in a variety of ways to produce a number of more complex systems than the one presented here. These include: Co-ordination of more than one articulator (Jordan, 1990); forward modelling of limb dynamics which incorporates use of backpropagation through time (Jordan and Rumelhart, 1992); the inclusion of multiple task and internal constraints (Jordan, 1992); and combining feedforward and feedback control in a single system (Jordan, 1990).

3.4 Chapter Summary

In the first part of this chapter we reviewed theoretical approaches to feedforward movement control from psychology. This review focussed on the various motor programming theories and examined the conceptual foundations on which they are based. It was concluded that, whilst this approach has provided a coherent and structured description of the problems faced by an autonomous system in the internal generation and control of feedforward actions, it stops short of providing an explanation in the form of explicit models of processing mechanisms. Part of the problem is that several of the mechanisms

presumed by the approach are built on biologically implausible computer primitives - for example serial-buffer mechanisms with intrinsic ordering properties, or static symbolic representations of action. Various computational models have been built based on these properties and have been criticised for failing to reproduce the characteristics of human sequential behaviour. One major problem is that they cannot account for interaction and competition between movement representations, or the internal 'soft' constraints on motor output which affect the overall shape and form of movements in a system with excess degrees of freedom.

A contrasting approach which grew out of a dissatisfaction with a profusion of 'box-and-line' models and the apparent 'loan on intelligence' (Kelso, 1982) they entailed is the action systems approach. Alternatively known as the dynamical systems or task dynamics approach, this theoretical position encompasses ideas from ecological psychology, dynamical systems mathematics and biological considerations. The emphasis here is on the impact of task properties (i.e. the structure and dynamics of the environment) on the motor patterns produced by biological organisms. Few computational models have, to date, been constructed to test this theoretical position. Some models, like the activation gradient models reviewed above, are based on well established associative learning mechanisms and processes commonly found in biological systems (e.g. lateral inhibition). However, a limitation of many of these models is that the representational primitives are overly simplified localist representations which are identified with a discrete unit of linguistic behaviour. Other recent models have employed dynamical systems concepts such as attractors, stability and bifurcations to explore the complex interactions which give rise to the output behaviour of movement control systems (e.g. Mpsitos & Soinila, 1993).

The most diverse and comprehensive computational modelling approach reviewed above is the Jordan combined sequential network and forward model architecture. One of the major strengths of this model is that it combines elements of motor programming theory by addressing the need to specify adaptive internal control structures for feedforward control and taking seriously issues of the complexity of motor representations, with elements of task dynamics and concerns about system-environment interaction. Another

attractive feature of this model is that it has been developed to control the goal-directed movements of a simulated kinematic limb within a simplified environment (or task space). Thus any changes in the processing of the model can be observed directly in qualitative changes in the output behaviour of the limb. This feature counteracts one of the criticisms of the model, which is that it is embedded within a complex and powerful general purpose connectionist architecture (the multi-layer perceptron with backpropagation learning algorithm). Systematic analysis of the state space of large networks of this type is notoriously difficult, thus the use of a simulated limb provides an intuitive way of visualising and evaluating the performance of the model.

CHAPTER FOUR

A Conceptual Model of Cortico-Striatal Motor Programming

In this chapter we develop a conceptual model which relates the cortical and subcortical mechanisms in motor processing to the selective impairments found in motor sequencing in Parkinson's disease. The framework incorporates elements of existing models of basal ganglia operation and computational models of sequencing. In constructing the model we first consider the motor control problems that biological organisms have to solve, and how the basal ganglia may have become adapted through evolution to fulfil the hypothesised function. The model is based on the combined evidence reviewed in the previous two chapters and its aim is to produce a constrained set of computational level hypotheses that act as a theoretical foundation for the modelling work in subsequent chapters.

4.1 Motor control in a hypothetical 'simple' organism

Let us for the moment consider a hypothetical 'simple' animal, for present purposes imagine a crayfish with much of the complexity removed. The animal is simple in the sense that it possesses the minimum brain mechanisms to control a few hardwired motor programs or fixed action patterns necessary for survival. Let us further suppose that the stimulating conditions for the initiation of one of these programs are provided by the environment. We will restrict our discussion to only two programs: program A - a mating response and program B - an escape response. The simplified perceptual system of our imaginary crayfish, which can only detect either mates or predators, constantly scans the immediate environment for the presence of either. The presence of a mate provides the 'key

stimulus' for the activation of the mating program and similarly the presence of a predator in the locality stimulates the activation of the escape response program. [4]

This simplest possible scheme will not work even when we only consider the two motor programs described above. Let us assume, for instance, that the water is cloudy and that the perceptual input is somewhat degraded as a result. Also it is possible that the mate and predator are of similar size. In such a situation there is potential conflict between the two different responses and it is not clear how the competition would be resolved in the system as it has been described. Furthermore, the effects of activating the wrong program are potentially disastrous.

A slightly more complex scheme would have both programs held under inhibition in normal circumstances. One might then presuppose an intermediate mechanism between the perceptual and motor systems which serves to *disinhibit* the appropriate motor pattern only when sufficient information is available about the nature of the key stimulus. This alternative scheme confers a number of advantages on the behaviour of the organism.

1. It reduces the likelihood of producing an inappropriate response because of the existence of general inhibitory control.

2. On account of the mechanism for selective disinhibition, the system is less prone to 'freeze up' due to parallel activation of conflicting response tendencies.

3. Both programs can be maintained in a full state of readiness up until the moment when one or the other is released.

Work reported by ethological and physiological psychologists has demonstrated that the principle of general inhibition - selective disinhibition is ubiquitous in the motor control of lower animals (see Gallistel, 1980; Lorenz, 1977). Lorenz states that:

"In many lower animals the most important function of the highest centres of the nervous system is to exercise a permanent inhibiting effect on the various endogenous-automatic behaviour patterns of the organism, and on the basis of external information to 'set the pattern off' at the appropriate moment". (Lorenz, 1977 p.59)

---

[4] This example draws on work reported by Lorenz (Lorenz, 1977). He provides an example of young Cichlid fish which react optically to both their mother which they have to follow and to that of a predator of the same size which they have to escape from. The observation was recorded by Kuenzer (1968).

He provides the examples of the earthworm which continues to crawl even when its supra-oesophageal ganglion has been removed and the crab which, when lesioned in a similar way, continues to eat as long as food is available. An interesting comparison may be made with human patients suffering from prefrontal lesions in whom comparative perseveration and utilisation behaviours are common.

The hypothesised conceptual link between the primitive inhibition and innate releasing mechanisms discussed above and basal ganglia function is clearly illustrated by the work of Hikosaka (Hikosaka, 1991) and Chevalier and Deniau (1990) on the role of the striato-nigro-collicular system of the basal ganglia in the control of saccadic eye movements which was discussed above. It has not been shown that the more recent striato-pallido-thalamic system plays an analogous role in the control of limb movements, but in developing the conceptual model we shall postulate that this is the case.

4.2 Motor control in higher animals

The motor control problem for higher animals and humans differs in two main ways from the 'simple' scheme we have outlined above. First, rather than having a restricted set of innate genetically-coded motor patterns to select between, higher animals, especially man, have a vast repertoire of stored motor programs and motor program fragments from which flexible, adaptive motor synergies can be formed to accommodate a wide range of environmental conditions and task constraints. This puts a far greater computational burden on the disinhibition mechanism.

Second, in complex motor tasks, commonly not all the information needed to perform the task is available from the environment. This means that action selection must take place at least in part on the basis of internally coded contextual information accrued from previous experience 'to fill in the gaps', thus allowing shortening of reaction times and anticipation of the appropriate response. In the most extreme case, human voluntary action, no information is either necessarily available from the environment nor utilised. Here the internal contextual representation provides a substitute for the key eliciting stimulus. It has been demonstrated that an arbitrary environmental event can become a

substitute for an eliciting stimulus through reinforcement learning. It is plausible that, assuming the brain machinery exists, the stimulus may also be associated with a substitute internal event. A strong candidate for such a role would be a dynamic pattern of activity in prefrontal cortex (see section 2.4.3).

4.3 Developing the conceptual model for the selection and initiation of action

We are now in a position to forge the links between the above features of higher order motor control and the function of the prefrontal and basal ganglia circuits. The descending nigro-collicular system which is involved in the disinhibition of eye movements is phylogenetically older than the complementary ascending system and it is reasonable to suppose that an analogous system would have played a similar role prior to the evolutionary development of the cortical mantle. In this case the input to the disinhibiting mechanism would be relayed directly from the perceptual apparatus.
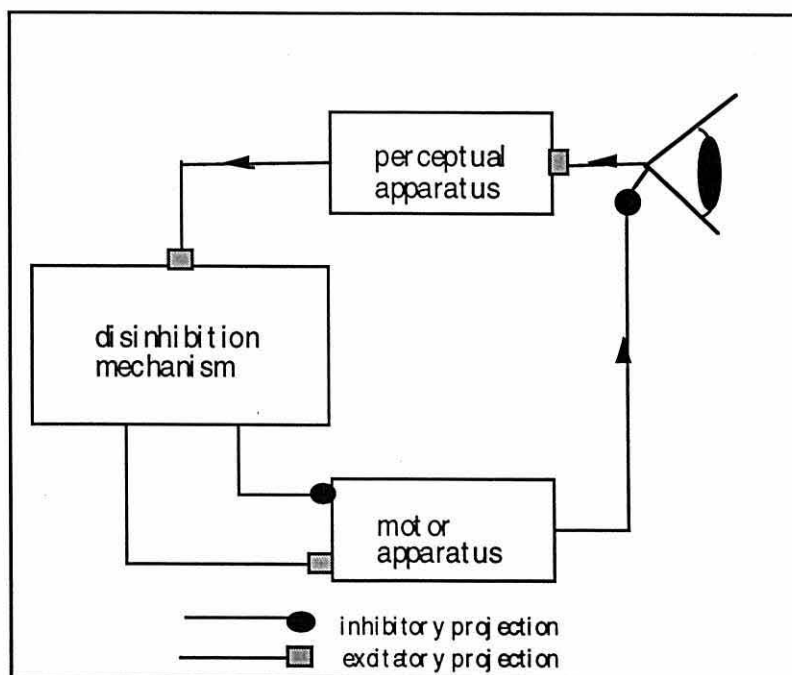
Figure 4.1. *A schematic illustration of a simple disinhibiton mechanism analogous to the basal ganglia in the release of saccadic eye movements.*

A schematic illustration of this idea is presented in figure 4.1. In the simplified scheme depicted in figure 4.1 gaze is held at a static location by the double excitatory pathway from perceptual apparatus to the motor apparatus which holds the muscles controlling eye movement under tonic inhibition. However, when the perceptual apparatus detects a change in target location, the inhibitory pathway from the disinhibition mechanism is activated which releases the eye movement controller from tonic inhibition and thus allows the eye to move to the new location.

However a major difference in the modern mammalian brain is that the striatum receives input not directly from perceptual systems but via multiple cortical stations including prefrontal cortex. Moreover its output is directed back to motor areas of cortex rather than spinal motor circuits. We suggest that these internal loops which interact with neocortex serves to handle much of the complexity that arises in higher order motor control (motor programming), but that the intrinsic circuitry of the basal ganglia still perform the same basic function of selective disinhibition.
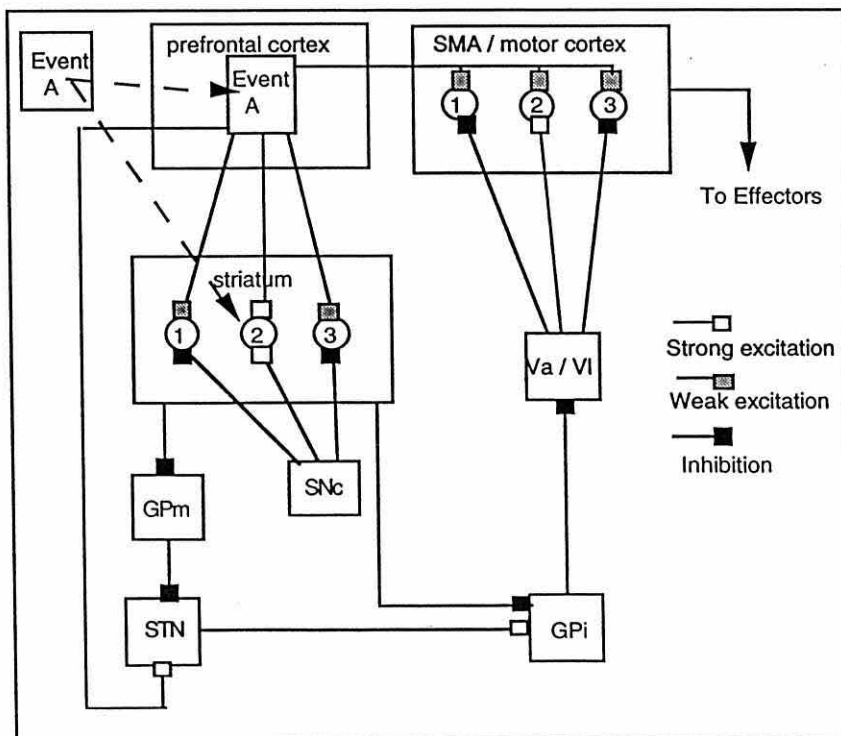
Figure 4.2. *A schematic diagram of hypothesised basal ganglia / prefrontal control in the selection and initiation of "action 2", which is predicted by "event A" in the environment*

Figure 4.2 above provides a simplified schematic illustration of the main features of the conceptual model we are proposing. To make the discussion more concrete we have assumed that event "A" is an event in the environment which predicts or is associated with desired action schema "2". For example, suppose that event "A" is the release of a ball thrown towards us and action schema "2" is a catching schema, (action "1" may be a throwing schema and action "3" a kicking schema).

When event A occurs, a number of internal changes are presumed to co-occur. For simplicity we consider here only a few of them:

• The occurrence of event A is detected by visual attention networks (Jackson et al., 1994) and is relayed via cortico-cortical connections to dorsolateral prefrontal cortex where it is associated with an instantaneous state of the dynamic contextual signal represented by prefrontal cellular assemblies.[5]

• The detection of event "A" also activates a set of cortically represented motor schemata associated with a ball throwing event. This is not in itself sufficient to prepare an appropriate response but may be thought of as orienting the system towards responding to a ball-throwing event. Simultaneously, information is transmitted via excitatory cortico-striatal projections to striatal medium spiny cells (Alexander et al., 1992).

• Projections from SMA and PMA (not shown in figure 4.2 for simplicity) innervate motor program fragments associated with the stimulated schemas in these cortical areas in a motor analogue of the combinatorial map revealed by the studies by Brown et al. (1994) and Flaherty and Graybiel (1993). Consequently actions schemas 1, 2 and 3 are depicted at striatal level in the diagram.

As a consequence we suggest that a pattern of activation is set up over the striatal schematic map which is associated with the state of the prefrontal context signal corresponding to the occurrence of event "A". This is the basic role of the internal context signal on motor output. Action schema "2" is shown as receiving stronger excitatory signal than either actions "1" or "3" at the striatal complex due to multiple sources of information which have

---

[5] The idea of a dynamic contextual signal can, in many ways, be thought of as a computational restatement of the idea of working memory - a role which has been assigned to DLPFC (Goldman-Rakic, 1995).

an impact on striatal neurons, including sensory or attentional information caused by the occurrence of event "A".

Also impinging on the striatal combinatorial map are dopaminergic projections originating in the SNc. These provide the basis for the selective disinhibition function of the basal ganglia. Shultz and colleagues have provided evidence that dopamine cells can become sensitive to events in the environment which predict behavioural responses, thus dopamine cells also fire when event "A" occurs (see chapter two). The complex action of dopamine on the indirect and direct output pathways of the striatum is depicted crudely in the diagram as providing excitation to action "2" and inhibiting competing actions "1" and "3". Following Wickens (1993a) amongst others, this is intended to indicate a model of dopamine action in which there is competitive processing over the striatal combinatorial map. In this way dopamine performs the contrast enhancement function suggested by Servan-Shreiber et al. (1990) and yields a focussed disinhibition of the appropriate cortical motor schema.

After a period of learning, however, we suggest that the contextual representation alone is sufficient to cause the correct pattern of activation over the striatal cell assemblies. We suggest that cortico-cortical associative networks also exist such that a given state of the context signal serves to provide a more discriminative facilitation of the appropriate cortical schema directly. The subcortical portion of the loop is still required, however. An excitatory projection from neocortex to STN maintains a tonic inhibitory influence over the cortical action schemata even though action "2" has attained a state of readiness. Release of action "2" is made possible by focussed dopaminergic controlled inhibition of thalamic targets pertaining to action "2".

A key novel feature of the conceptual model is encapsulated in the notion that dopamine cells in the SNc may also be sensitive to the internal representation of the context of event "A" encoded by the dynamic signal in DLPFC. Shultz and co-workers have already found that dopamine cells come to respond to neutral contextual stimuli in the environment associated with a key stimulus through learning. This finding was in primates performing simple tasks. It is conceivable, then, that in higher primates and man a similar

link may exist with the hypothesised contextual representation provided by DLPFC activity to subserve the generation of voluntary, goal-directed movements when no external cues are provided. Even though the model is speculative, it can potentially account for a number of currently unresolved issues concerning basal ganglia motor processing and motor disorders. Dopamine is hypothesised to aid feedforward motor programming in the following ways:

1. During early stages of learning it serves to improve the quality of the pattern of activation which is distributed over the striatal combinatorial map and thus the correct selection of cortical action schemata, due to its sensitivity to salient environmental events. It is possible that this function could also help to time the initiation of movements.

2. Dopamine facilitates the transference of the initiation cue from an external source (the event) to an internal source (the contextual representation) as a result of its secondary association with the state of the context signal.

So far we have only applied the model to the processes involved in the selection and mobilisation of a single action schema. In the section below we describe how the model suggests that the processes underlying motor sequencing are performed by precisely the same computational machinery that is already in place for action selection.

4.4 Accounting for cortico-striatal involvement in sequence control

According to the computational models reviewed in chapter three, sequences of movements can be adequately produced by either of two methods: competitive processing using inhibition, or a dynamic control signal which has a continuity property that maintains similarity relations between successive temporal contexts.

In terms of the conceptual model we have described how both systems may be implemented by cortical - striatal loops and we view both as playing complementary roles in the control of sequential output. The contextual control signal is provided by dorso-lateral prefrontal cortex (DLPFC). This enables a desired sequence of movement outcomes to be encoded as different states of the contextual signal. After learning the sequence a pattern of

activation is set up over the combinatorial map in striatum for each successive state of the context signal. The intrinsic processing of the basal ganglia serves a complementary function of facilitating the desired sequence elements at each salient state of the context signal and inhibiting competing representations.

4.5. Predictions of the conceptual model

The model proposed above makes a number of predictions:

The capacity for autonomous generation and control of voluntary movements has developed out of a system which is adapted to provide fast and accurate action selection in response to salient environmental variables. This can account for the finding of Mink and Thach (1991a) that the basal ganglia do not show a preferential role in self-generated or externally-guided movements. Equally it suggests that the reason that voluntary movements are most profoundly affected in PD is due to the fact that action selection takes place solely based on the internal contextual representation which more susceptible to disruption when available dopamine is reduced.

Similarly sequential actions are dependent on the internal contextual signal for the generation of later elements of the sequence rather than external sources of information. Thus it makes sense that PD patients should show greater impairments in overtly sequential actions.

The model predicts that changes take place in the system during learning of skilled movements. In the early stages of learning the internal context signal will inevitably be relatively impoverished, thus emphasis is placed on inhibitory competitive processing on the basis of external cues. This mode of control may be expected to yield the performance characteristics typical of early learning: high muscular stiffness, poor coarticulation and low anticipation. As learning progresses, and a better internal contextual representation is established there is a reduced dependence on external information and inhibitory processing. This in turn leads to reduced stiffness, greater smoothness in the transition between sequence elements and a higher degree of anticipatory movements. To provide an

-example, one may recall the changes in the performance of the clutch, gearshift, accelarator sequence as one learns to drive a car.

The model suggests that there is a complex inter-dependent relationship between the cortical and sub-cortical loop in the interaction of basal ganglia and prefrontal cortex. Efficient feedforward control of action sequences is equally dependent on a good quality internal contextual representation and the integrity of the nigro-striatal dopamine system. A reduction of the available dopamine in the system as occurs in Parkinson's disease may be expected to yield deficits at a number of different levels: Interference between sequence elements, interference between different sequences (motor plans), interference between different groups of effectors.

4.6 Chapter summary

We do not claim that the conceptual model put forward in this chapter is in any way a complete description of the function and operation of the frontal and motor loops through the basal ganglia. Much biological detail has been eliminated and other elements have been highly simplified. Rather, it represents an attempt to draw together some of the many disparate current hypotheses with evidence from neurobiology, movement control, and parkinsonian studies into a coherent framework within which computational level questions can be explored. The main hypotheses are reiterated below:

1. Frontal cortical areas (projecting to the basal ganglia) can be viewed as providing a dynamic internal contextual signal that drives the sequential flow of action. This contextual signal effectively substitutes for direct environmental input to subcortical structures.

2. The basal ganglia are involved in competitive action selection – inhibiting undesired responses and facilitating intended responses on the basis of the partial action specifications that are provided by the "frontal" contextual control signal. The action-selection processes are similar to those that occur when the environment, rather than frontal cortex, provides the main input to the basal ganglia.

3. There is a two-way influence of action selection processes on contextual control processes and vice-versa which is encouraged by the loop structure of basal ganglia-cortical interaction

4. There is a trade-off between these processes: there is a greater need for efficient action selection when the dynamic contextual control signal is degraded or when an action sequence is only partially-learned

5. Deficits in either frontal/contextual control and sub-cortical "action selection" can interact in unpredictable ways.

In the remainder of this thesis we explicitly explore these hypotheses by computer simulation.

CHAPTER FIVE

The Effects of Lesioning a Jordan Sequential Network

In the previous chapter we have applied the linked concepts of 'internal context' and 'selective disinhibition' to the control of sequential motor actions and highlighted their potential importance in accounting for the motor deficits in diseases of basal ganglia dysfunction, notably Parkinson's disease. Also we have emphasised the need for a *dynamical* approach when developing models of the control of movement.

In this chapter we explore the theoretical and computational issues that emerge from this conceptual foundation. These issues concern the role of the basal ganglia in normal movement control and, in particular, the effects of basal ganglia dysfunction on sequential motor behaviour. More specifically we examine the role of a dynamic internal representation of context and a *predictive* forward model in controlling the sequential behaviour of a simulated robot arm.

5.1 Modelling patterns of deficits in parkinsonian sequential actions

A popular method within connectionism is to select or develop a model that reflects human performance and which is based on certain key theoretical and computational primitives. Connectionist lesioning techniques are then applied to the model and the pattern of breakdown of behaviour can be examined and compared with human data from neuropsychological studies.

Numerous examples of this general approach are available in the literature. McClelland and Rumelhart (1986) modelled anterograde and retrograde memory deficits in amnesics by degrading the weights on connections during learning. Cohen and Servan-

Shreiber (1989) modelled aspects of abnormal use of context in schizophrenics by altering the sensitivity of selected units in a way that related to the effect of abnormal levels of excitatory neurotransmitters on specific groups of neurons. Levine (1986) addressed the joint problems of perseveration and utilisation behaviour in frontal lobe patients by disrupting interactions between pools of sensory and reinforcement units in an ART network (see also: Bapi & Levine, 1990; Changeux & Dehaene, 1994).[6] Finally considerable attention has been paid to investigating the behaviour of damaged connectionist networks in relation to the behaviour of patients with acquired dyslexia (Hinton & Shallice, 1991; Mozer & Behrmann, 1990; Plaut, 1991). Perhaps the most interesting of these approaches was that of Plaut (1991) who used networks that build attractors around representations in the network's state space. He analysed the effects of damage on the behaviour of the network in terms of the distortion and interference created in the properties and layout of attractors.

To construct a connectionist approach to motor disorders in a similar fashion to those described above, three preliminary decisions need to be made:

1. What pattern of impairments are to be modelled.

2. What model is to be used in simulating normal behaviour.

3. What damage is to be incurred in the model.

These issues are dealt with in the following sections.

5.1.1 Parkinsonian impairments in sequential control

In most of the connectionist approaches reviewed above, the performance of the model in question is assessed over a corpus of data which has been derived, and usually simplified, from empirical work. The performance of the model may then be directly compared with the human data gathered using the same task criteria. An indicator often used to evaluate the model in such simulation studies is the match produced between a graph of human performance and model performance.

---

[6] A description of ART networks may be found in Carpenter and Grossberg (1987).

Although this methodology is appealing on account of its apparent simplicity and objectivity, a few critical comments need to be made. First, the task of setting up the input and output representations of a model to reflect the information content of the human dataset is usually not straightforward and the inevitable simplifications can lead to an inbuilt bias towards achieving the desired effect, not as a result of properties of the model, but due to properties of the system of representation. An example of this is the verb tense learning model of Rumelhart and McClelland (1986) which was subsequently criticised on account of the "T.R.I.C.S" (The Representations It Crucially Supposes) used in the Wickelphone system of input representation. Another potential problem is that many models have a number of system parameters unrelated to the particular modelling problem that may be 'tuned' to provide a better match between the graphs. This form of atheoretical parameter tweaking is more easily avoided if some other performance criteria are used.

There are several further reasons why this approach is inappropriate for modelling many movement control phenomena and, in particular, the movement deficits in PD. First, in motor experiments with PD patients there is no easily definable analogue of a corpus of input data that is commonly found in domains with a linguistic basis such as speech, reading, spelling etc. Similarly on the output side the results of experiments with patients are not analysed with respect to robust effects in normal behaviour such as, for example, the recency effects and serial position curves characteristic of human short-term memory performance. There is considerable variability in the designs and techniques used to investigate PD movements, not to mention variability in the results obtained which again makes it difficult to model an effect which has implications outside the particular experimental approach used as a basis for the model. Much of the present inconsistency in empirical research in PD is attributable to the complex and multi-faceted nature of the parkinsonian deficit. It also reflects the relative youth of the research field in comparison with other more easily accessible areas of cognitive science and the concomitant lack of a firm theoretical framework of normal motor control.

In view of these problems, however, a necessary prelude to developing a modelling approach to PD impairments is to attempt to identify qualitative characteristics of

parkinsonian behaviour. A unique recent study by Brentari and Poizner (1994) offers some useful insights. They present an analysis of the sign language errors produced by a deaf parkinsonian signer. This approach has the methodological advantages of a linguistic-level analysis in that the information content of movements can be identified, whilst allowing physical observation the nature of the deficit in the movement of the limbs.

They categorise the errors of their patient (R.H.) into two main groups: spatial errors, which include a variety of sign distortions all commonly referred to as *reductions* and *timing errors*, which primarily consist of abnormalities in the timing of transition between handshape postures. Their findings are summarised below:

Reductions

This group of errors reflects the tendency of PD patients to produce abnormally small movements across a variety of domains: writing, ballistic movements, walking, speech, eye movements. A classic finding in the literature is that if patients are required to move to a spatial target, their movement will consistently fall short of the target. In the current study R.H. exhibited several types of movement reduction, as follows:

- Distalisation. Here the path of the movement is maintained but is transferred from a proximal joint such as the shoulder or elbow to a more distal joint such as the wrist or knuckle. This has the effect that the resulting movement is of diminished amplitude.

- Shadowing. This distortion occurs in gestures using two hands, where one hand remains static and acts as a base for the other active hand. In R.H. the base hand followed or shadowed the active hands movements.

- Contact Deletion. Many signs require the contact of the signing hand with some part of the body or the other hand. In the PD signer the movement terminated, almost without exception, short of making contact.

- Handshape Reductions. An important group of reductions refer to the formation of handshape postures themselves. For many signs there are a set of active fingers which are more active during sign execution and a set of passive fingers which remain relatively static during execution. The passive fingers can be either at full extension 'open' or fully flexed 'closed'. Typically the passive (and active) fingers in the patient

signer would be neither open nor closed, but would assume an intermediate 'lax' posture between the two.

An important point is that all of the above errors are imprecise executions of correct targets rather than crisp executions of incorrect targets.

<u>Timing disruptions</u>

Reports of timing abnormalities abound in the experimental literature on PD. Unusually slow movements are characteristic of the disease as are abnormally long latencies between the components of compound movements such as reaching and grasping and extended movement sequences. Various other time estimation and reproduction disturbances have been reported during the performance of repetitive movements. The primary timing disturbance analysed in patient R.H. was a prolongation of handshape change during the transitional movement of compound signs (i.e. those that are composed of two or more separate handshapes). Typically the patient would used a much greater percentage of the transitional component to execute the handshape change than did normals.

There are two notable features of all the errors exhibited by R.H. First, the errors all map onto known abnormalities in Parkinsonian movement generally, which suggesting that the error classification scheme produced here reflects generalisable characteristics of parkinsonian motor deficits. Second, all of R.H.'s errors favour ease of articulation at the expense of distinctiveness. These disturbances provide a useful and systematic framework for analysing the performance of a computational model.

5.1.2 Modelling normal sequential motor control

The connectionist architecture that we use as a basis for the computational work contained in this chapter is the combined sequential / forward model network (Jordan, 1988, 1990) which was reviewed as a model of motor learning and control in chapter three. We have already briefly discussed the main properties of this type of recurrent connectionist network and the motivations for its use in modelling movement control.

Our primary motivation for focusing on this model in connection with modelling motor disorders of basal ganglia dysfunction is the central importance of the recurrent links

to the state units in allowing the network to produce sequential behaviour. The state representation acts as a form of *internal context signal* which is closely related to our hypotheses of cortico-basal ganglia interaction, outlined in chapter four. Although there is no explicit inhibitory - disinhibitory mechanism in the model, the mapping from state representation to motor output is continuous, and consequently the values encoded on the state units have an influence on the level of competition or interference between the learned action and other possible actions.

As we have pointed out above, one of the difficulties with attempting to model the abnormal characteristics of parkinsonian movement is that there is no corpus of data which can be directly compared with the output of a connectionist models as is often done with more strictly cognitive phenomena such as memory (McCloskey & Cohen, 1989) or dyslexia (Plaut, 1991). An alternative approach is to compare the qualitative behaviour of the model when it is damaged in different ways.

One of the major advantages of using the Jordan scheme is that the output of the model controls the movement of a simulated robot arm in a two-dimensional space, by using a forward model that encodes the kinematic control variables in a separate phase of learning. This feature means that any computational disruption will be reflected in changes in the movement of the arm in the task space, and can be visualised using a representation of the changes in arm position over time. There are a number of characteristics of parkinsonian movement that can be assessed using this model following the categorisation scheme outlined above.

5.1.3 Methods for damaging the model

In the course of the work reported in this chapter we use connectionist lesioning techniques to examine how the sequential behaviour of the network is affected by disruptions to the state unit activations during learned performance. In psychological terms, this is intended to model the effects of disrupting the internal contextual representation. In order to produce computational disruptions that are consistent with our hypotheses of basal ganglia dysfunction caused by loss of nigro-striatal dopamine we use a method for simulating

dopamine depletion in connectionist networks derived from Servan-Schreiber, Printz and Cohen (1990), and Cohen and Servan-Schreiber (1992). This method and a detailed description of the network architecture and tasks are described in detail in a later section.

In the section below, we outline the computational features of the Jordan model that we explore in this chapter and illustrate how these features relate to the conceptual model of basal ganglia motor processing described in chapter four.

5.2 Overview of the modelling approach and hypotheses

Figure 5.1 below provides a schematic illustration of the key components of the model architecture used in this chapter.
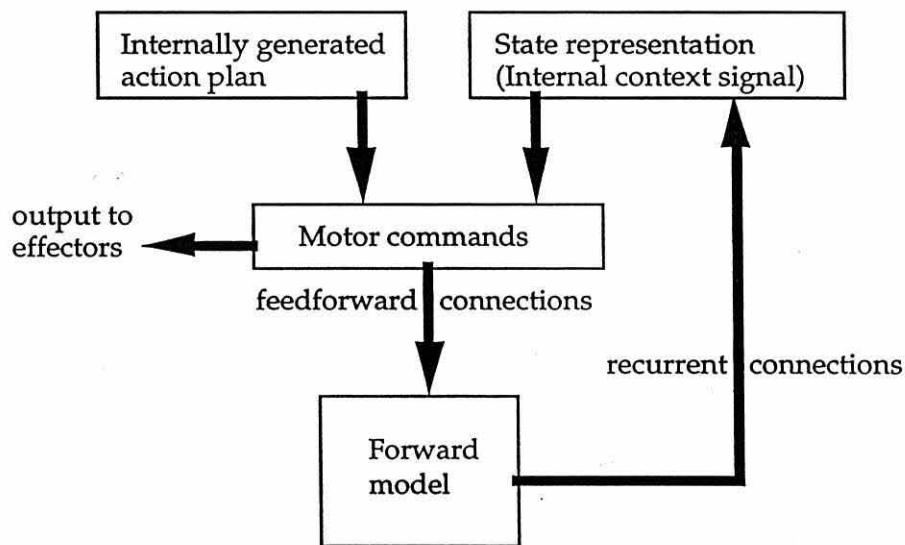
Figure 5.1. *Global structure and key components of the combined Jordan architecture*

We have described the individual components of the model in some detail in chapter three. In the sections below we outline those aspects of the model that are of specific interest in this chapter and, we show how they relate to the ideas expressed in chapter four.

The state units

As we described in chapter three, the state units can be thought of as a representation of temporal context. In effect they act as a dynamical control signal of the type that (we have suggested in chapter four) relate to the role of fronto-striatal circuits in the internal control of sequential action. In this chapter we aim to explore the computational properties of the

state unit representation and examine the effects of damage which disrupts the dynamics of the control signal.

Interacting subsystems

The Jordan model has a complex architecture in that it is composed of two separate but interacting subnetworks, the sequential controller and the forward model (FM). The two subnetworks have different functional roles in the operation of the whole system. This contrasts with simpler models in which there is no functional distinction between different parts of the network[7]. Issues of network complexity and functional modularity have been addressed in a number of other contexts (e.g. Jacobs et al., 1991). This feature is of interest in relation to basal ganglia-thalamo-cortical circuitry since this loop circuit is complex (i.e. composed of several interacting subsystems). Moreover, it is possible that the complex nature of parkinsonian motor deficits can, in part, be attributed to the partial interaction of many different subsystems with the basal ganglia in the production of voluntary movement, (e.g. frontal cortex, M1, SMA, cerebellum). In the context of the model we can compare the case in which all of the input to the state units arises from one source (i.e. the output of the forward model, as depicted above), to the case in which it has a composite source (i.e. the output of both the controller and the forward model).

Recurrence

The sequential behaviour exhibited by the network is driven by recurrent connections from either the forward model as depicted in figure 5.1, the articulatory (output) units, or both. The use of recurrence in connectionist models is an interesting topic from a purely computational perspective. However, in the current model the fact that the system has access to the output of the forward model during sequence processing is of particular interest.

The forward model.

We have already discussed the motivation for a forward model (FM) during learning. This study is more concerned with the 'forward' role of the FM in a learned network via its provision of input to the state units, (see figure 5.1). In this sense the FM may be thought

---

[7] This does not include the distinction between input/output and hidden units in connectionist networks.

of as providing a predictive internal estimate of the expected sensory or environmental consequences of performing each individual sequence element. The notion of a predictive forward model which is acquired through learning has a number of links with the conceptual model outlined in chapter four, although at a slightly different level of interpretation. It amounts to a computational hypothesis of how the subsystem implementing the contextual control signal (cortico-striatal circuit) and the subsystem performing action selection (pallido-thalamic circuit) mutually influence each other during sequential processing. As illustrated in chapter two, a number of researchers have emphasised a predictive role for the basal ganglia similar to that implemented by the FM in the current model. In the simulations below we explore how disruptions to the FM influence the sequential behaviour of the whole system.

The hypothesis developed in this chapter is that the function of the basal ganglia loop circuitry encompassing the striato-pallidal complex, thalamic nuclei and multiple areas of cerebral cortex may be viewed as the disinhibition of cortical action schemata or motor programs on the basis of a high-level motor plan, and the dynamic context signal provided by the state units. In this case the input to the state units is provided by the output of the forward model which is an internal estimate of the outcome of performing that action. Thus, the system as a whole may be seen as selecting desirable actions and concurrently inhibiting competing responses, through the computation of a FM. Furthermore it is argued that disruption to an FM which provides input to the internal contextual signal could cause impairments of sequential movement control which are similar to those that are observed in PD. In the section below we describe in more detail the architecture and computational function of the combined controller and forward model networks.

5.3 Description of the connectionist architecture and tasks

5.3.1 The basic operation of the model

As was discussed in chapter three, the function of the connectionist architecture that forms the basis of the Jordan model is to learn to control a simulated two-joint kinematic arm with a movable shoulder in a two dimensional task space. The model simulates goal-directed motor behaviour by learning to touch a specified target in task space with the endpoint of the arm. Sequential motor tasks are simulated by requiring the endpoint to touch a sequence of different targets in a specified serial order.

During the first stage of learning the forward model is supplied with random values for the control parameters and learns the forward mapping to the endpoint location, which is represented by a set of 'task units'. Following that the feedforward controller can be trained to guide the endpoint to a desired target location in task space by observing an error between the desired location of the endpoint and the actual endpoint location represented by the task units and then in effect running the forward model backwards to obtain a control space error which can be used to adapt the feedforward controller. If, as is the case here, the controller is a sequential network, then the system can learn to produce motor sequences.

It should be noted that the FM output does not give the actual position of the endpoint of the arm, but a predictive *estimate* of the position of the endpoint. The actual endpoint position is encoded by the kinematic variables in control space and must be calculated separately in the computer simulations. It has been shown that the FM estimation will only be accurate if the FM can be learned to zero error for all locations in task space (Jordan & Rumelhart, 1992). This is clearly an unfeasible assumption in a noisy, complex system. However, accurate goal-directed movement does not require an accurate FM. In the system studied here it was found that the tolerance to inaccuracy in the FM under normal conditions was quite high. This has the benefit of making the system robust in the face of noise during the learning process, (Jordan, 1992).

However, even after learning, the estimate provided by the FM should have an impact on sequential performance because it provides an internal contextual representation

of the subsequent action as input to the control system, via the recurrent links to the state representation. This is the *predictive* role of the FM that we referred to above. Although, in the model learning and performance are two discrete and separate processes, we assume that in human motor systems these are continuous and concurrent processes, and that the FM participates in both the learning and prediction roles continuously.

5.3.2 The task space and simulated arm

The task space is a two dimensional space varying between 0 and 1 on both the X and Y coordinates. Targets may be expressed either as points or regions within this space. The arm has two segments, each of length 0.25 and two joints which can vary between 0 and 180 degrees. The location of the shoulder is expressed in Cartesian coordinates. The control parameters of the arm thus have four degrees of freedom: two joint angles, (theta and phi) and two translational displacements of the shoulder, (X(s) and Y(s), and the task parameters have two degrees of freedom, (the Cartesian coordinates of the target or endpoint of the arm in task space). The location of the endpoint is calculated as follows:

$$X \text{ (end)} = X(s) + \sin(\text{theta}) + \sin(\text{phi}) \tag{5.1}$$
$$Y \text{ (end)} = Y(s) + \cos(\text{theta}) + \cos(\text{phi})$$

5.3.3 Network architecture and training regime

The network architecture we used closely follows that used by Jordan (Jordan, 1990; Jordan and Rumelhart, 1992), which was described in chapter three. The network can be divided into two separate subnetworks: the sequential controller and the FM. The combined network architecture is illustrated in figure 5.2.
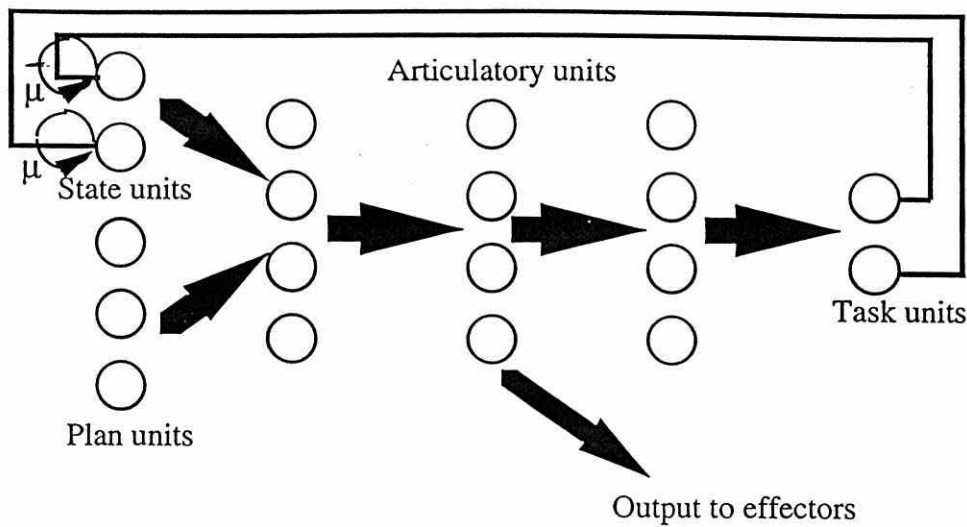
Figure 5.2. *The network architecture*

The network has two sources of input: the plan units, and the state units. The plan units are initialised independently at the beginning of a sequence and there are no particular constraints on the representations except that they are different for different sequences. The state units receive recurrent connections from the output units of the FM and/or the sequential subnetwork. The recurrent connections are linear and have a constant weight of one.

A layer of four hidden units intervenes between the input and output units of the sequential subnetwork to allow arbitrary, nonlinear associations to be formed. The number of hidden units is the same as reported by Jordan (1990). The effects of varying the number of hidden units was examined in pilot simulations but no differences were found on the tasks used. The four output units of the sequential subnetwork are the articulatory units. These encode the kinematic control variables of the arm indicated in the previous section. The articulatory units provide the input to the FM subnetwork, which simply constitutes a standard three layer network in which the two output units encode the Cartesian coordinates of the estimated endpoint location in task space. In order to explore issues of partial system interaction, in some simulations the only recurrent connections were from the FM output units and in others there were connections from the articulatory units as well.

Finally there are recurrent connections from the state units onto themselves, which were included by Jordan, (1986a) to allow the network to learn extended sequences with repeated elements. These connections have a variable weight governed by the parameter μ, which acts as a form of short term memory over the state units.

All of the hidden and output units in the network are standard logistic units, (Rumelhart, Hinton & Williams, 1986a) with activations varying between 0 and 1. Supervised learning by gradient descent is implemented in the model using the backpropagation of error algorithm with generalised delta rule described in Rumelhart et al., (1986).

We now turn to the first set of simulations. In section 5.4 we describe the preparatory work involved in learning the sequences which were subsequently used in the lesioning studies and in section 5.5 the lesioning studies are reported.

5.4 Preliminary Simulations

The aim of the first set of simulations reported here was simply to learn the sequences which will be subsequently used in the lesion studies and to replicate some of the basic properties of the architecture reported in previous accounts (Jordan, 1990; Jordan & Rumelhart, 1992).

5.4.1 Learning the forward model

The FM was learned in a separate initial phase using gradient descent learning with the backpropagation algorithm. The learning rate was set to 0.1 and the momentum term to 0.9. The normal method for training a network of this type is to iterate through a restricted training set. In order to simulate the selection of random inputs, a different method was used here. On each learning trial, a configuration of the arm was generated by randomly selecting values for each of the kinematic variables on the input units. An error score was derived by comparing the network's estimate of the endpoint location in task space, represented by the values on the task units, with the actual endpoint location calculated independently. This error was backpropagated through the network using the generalised

delta rule, and the weights were changed after each iteration. Learning was terminated when the¨ sum squared error averaged across fifty consecutive iterations was below a criterion value of 0.01.

Jordan (1990) reports that the FM in that study was given approximately 2,000 trials of learning. In the current implementation approximately 6,000 trials of learning were provided to ensure sufficient learning. Figure 5.3 below depicts the error reduction in the FM during learning. It is clear from the figure that this method of learning does not yield a smooth reduction in error as normally occurs when a standardised input set is used as training data. However, the method that was used is more in keeping with the notion of system identification through random exploration of the parameter space. Moreover, it does not matter if the results of learning are not completely reliable, since Jordan has shown that only a partially learned is sufficient for accurate goal directed performance after the second stage of learning (Jordan, 1988, 1990; Jordan & Rumelhart, 1992).

No. Iterations / 50
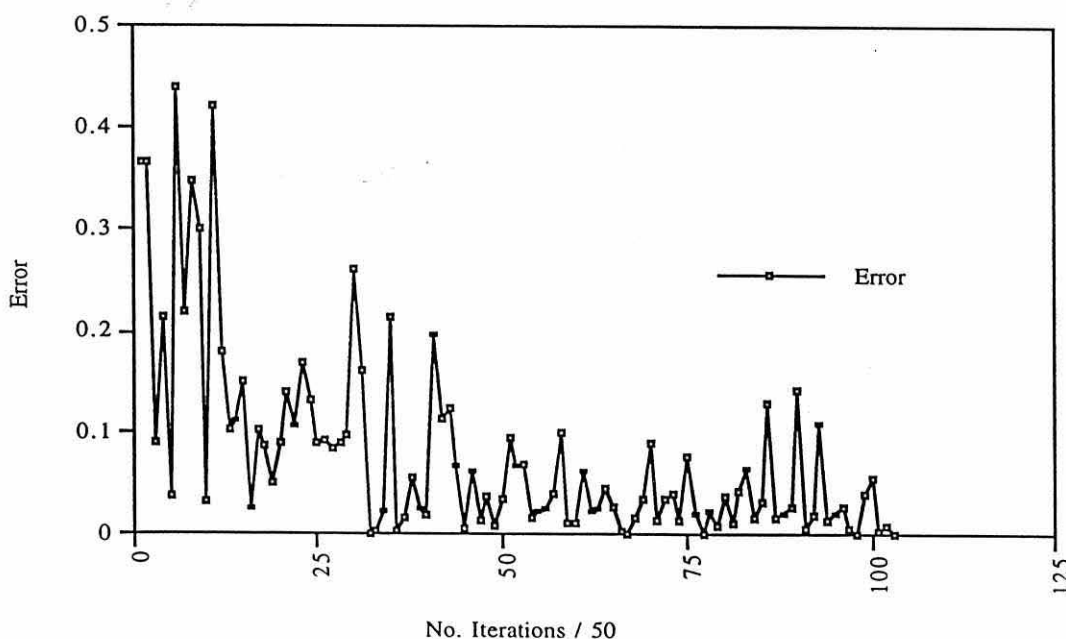
Figure 5.3. *Learning the forward model*

5.4.2 Learning a sequence

Figure 5.4 depicts the basic form of the task the network was required to learn for the following studies. The task is to perform a four element sequence touching each of the targets, in the labelled order, with the endpoint of the arm. In most cases the transition from

position 4 to position 1 is also learned, thus completing the cycle. This task is analogous to a common laboratory task in motor control studies in which a subject is seated at a table and directed to touch a sequence of targets displayed on the table in succession. The basic structure of the task was only varied minimally in all the simulations reported in this chapter. Our reasoning for this was to provide a standardised test for comparison of performance under different conditions. We investigate the effects obtained using variations of the basic sequence in later studies.
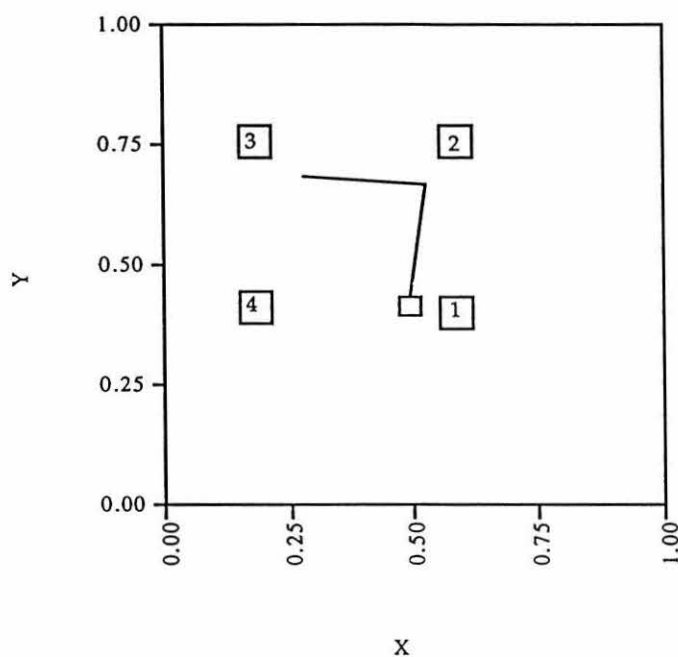
Figure 5.4. *Task structure and initial configuration of the arm*

5.4.3 Simulation 5.1: Learning a 4-element sequence

The aim of simulation 5.1 was simply to learn a four element sequence with the combined forward model and controller network.

Method

In simulation 5.1 the combined network was trained on the sequence depicted in figure 5.4. The weights in the sequential controller were initialised to small random values between 0.3 and -0.3, the learning rate was 0.1 and momentum was 0.9. The learned weights in the forward model were held constant during learning. After each forward pass

of activation through the sequential network and forward model the output of the task units was compared with the target location and the resulting error score was backpropagated.

Results

The targets and final configurations of the arm are depicted in figure 5.5.
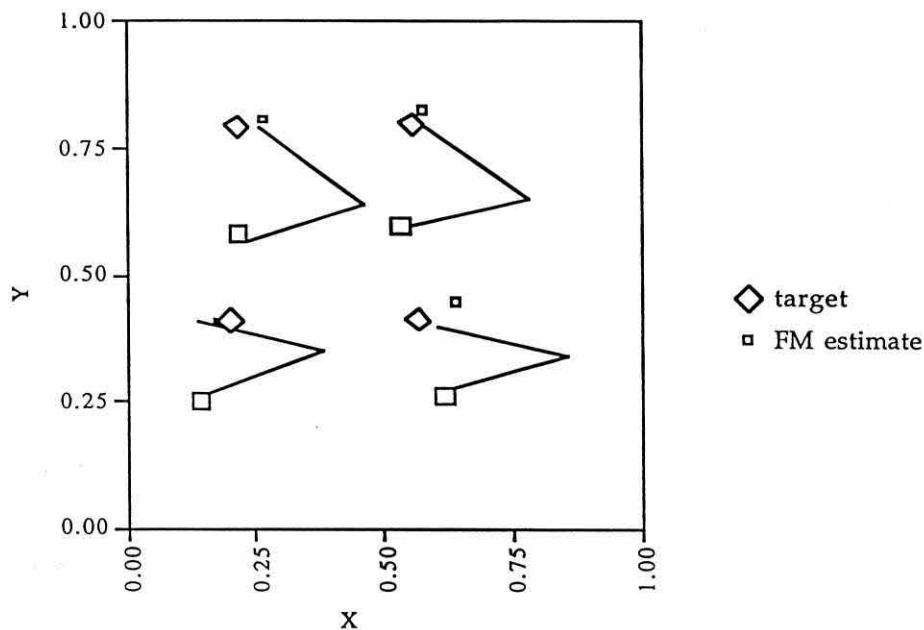
Figure 5.5. *Learned performance of the combined network on a sequential task*

It can be seen from the figure that the network has adequately learned the sequential task, in that endpoint of the arm is, in each case, close to the location of the target. However, the network has clearly followed a rather 'non-biological' optimisation strategy. This involves keeping the arm in a static position allowing the weights in the controller subnetwork to encode only the translational displacements of the shoulder during sequential performance.

Discussion

Although the strategy adopted by the network in this simulation seems to follow an entirely reasonable optimisation principle from a computational point of view, the solution is very different to the changes in kinematic configuration depicted in the figures of Jordan (1990). In these all of the degrees of freedom of the arm are used equally. One reason for the

discrepancy may be that Jordan had implemented a variety of internal constraints on the behaviour of the network.[8]

In the context of the current simulations it is plausible to impose a 'rest configuration' constraint of the type suggested in Jordan (1990) on the movement of the shoulder only. A rest constraint acts to penalise large changes in activation in a subset of units from their initial activation levels. This strategy would force the network to use the degrees of freedom available in the articulated arm itself. This is reasonable since, in the real system, the shoulder is attached to the body which has considerably greater mass than the arm segments which would make it extremely inefficient for humans to adopt the strategy used by the network. The rest configuration constraint effectively simulates this energy cost. Thus a further simulation was conducted in which a rest constraint was applied to the movement of the shoulder.

Method

The rest constraint is simply computed as an extra error term at the articulatory units, which is defined as the square of the difference between the current shoulder location and the initial shoulder location. In order that this constraint does not compete with the task constraints it is multiplied by a weighted parameter (lambda) that is proportional to the total sum-squared error on the task units. This term is added to the unit delta value during computation of the error at the articulatory units.

The previous simulation was rerun with this additional error term included.

Results

The results of imposing a medium strength rest configuration constraint (lambda = 0.6) on the movement of the shoulder are depicted in figure 5.6. It can be seen that the network now has found a solution in which the shoulder moves very little from its original location, and that most of the work is done by changing the joint angles. Although the results presented here are derived from a single case, they are typical of the results obtained across

---

[8] In fact, we later discovered that the most probable reason for this solution is that the network was given only a single task and thus is not subject to competing constraints on the articulators. Even without the imposition of internal constraints the network will produce more realistic solutions if more than one sequence is learned. Jordan, (1990) however does not indicate anywhere that more than the test sequence was learned in any task.

a number of cases in which the initial configuration of the weights was varied using different random seeds and learning variations on the test sequence.

By comparing the FM estimates in figures 5.5 and 5.6, it can be seen that the imposition of the rest constraint alone causes an increase in the inaccuracy of the FM prediction, whilst not affecting the overall accuracy of the solution. This provides a useful indication of the tolerance of the network to FM inaccuracy.
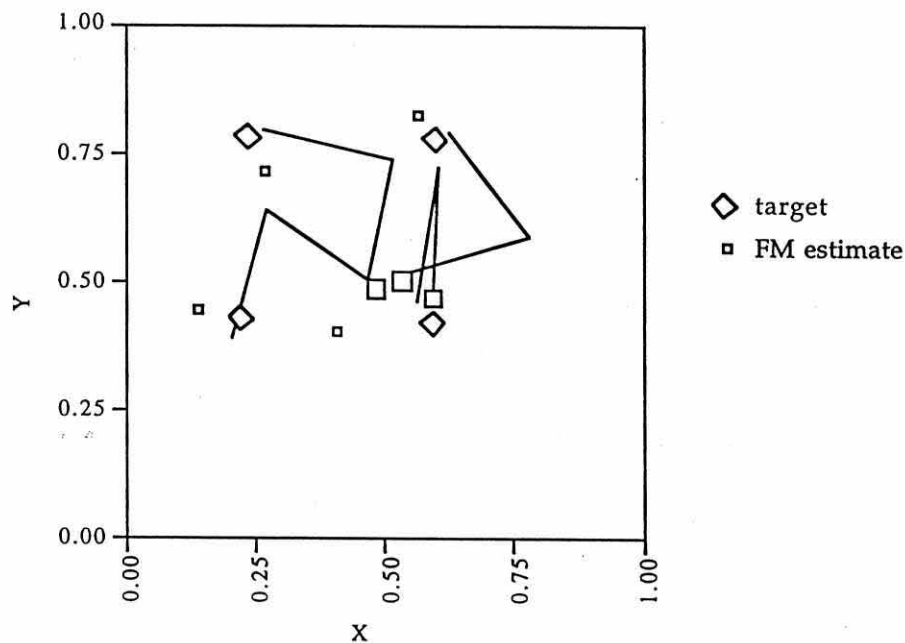
Figure 5.6. *Learned network performance using a rest constraint on shoulder movement*

Jordan (1992) states that the imposition of internal constraints in general makes sequences easier to learn, due to the reduced number of possible solutions to the inverse mapping problem. In contrast, in a number of pilot trials we found that in many cases the addition of rest constraints at the articulatory units actually hindered learning. Often the error score oscillates during learning and problems with local minima were experienced. There are many possible explanations for these differences which could be accounted for by variations in the setup of the network or the task, including task difficulty, the number of degrees of freedom in the arm or the range of mobility of the joint angles.

These issues were not fully explored as a number of sequences could be effectively learned using the set up used above. Consequently, the learned performance to the sequence as depicted in figure 5.6 was used as the prototype learned state of the system for the subsequent lesion studies in this chapter.

5.5 Simulating the effects of basal ganglia damage by lesioning the learned network

The aim of the series of simulations reported in this section was to investigate the effects of disrupting the processing of those parts of the system assumed to reflect the functional role of the subcortical parts of the cortico-basal ganglia motor loop (i.e. the forward model and recurrent connections to the state units). The specific intention was to examine whether disruptions which simulate the loss of striatal dopamine would cause motor deficits which are qualitatively similar to those which occur in either experimental or idiopathic parkinsonian motor disorders.

The basic methodology we adopt in the following simulations is to test the network with the learned solution depicted in figure 5.6 following a variety of different types of damage. These are explained in the following section.

5.5.1 Lesioning techniques and locus of disruption

The fundamental comparisons we make in the lesioning studies are between, (a) a normal control condition represented by the performance of the intact network, (b) damage that simulates dopamine depletion, and (c) non-specific damage which is caused by techniques commonly used in standard lesioning studies. The methods used to inflict damage on the network are discussed below.

This methodology is uncommonly strict in a connectionist lesioning study of this type. A more common approach is to simply add noise to the unit activations or the weights of the network and compare performance with the undamaged network. Our aim here, however, is to examine whether damage that simulates dopamine depletion has a qualitatively different effect on behaviour than other types of damage.

The method we used to simulate nigro-striatal dopamine loss is that developed by Servan-Schreiber, Printz and Cohen (1990) and subsequently used by Cohen and Servan-Schreiber (1992). They draw upon a variety of neurobiological evidence to suggest that the modulatory action of dopamine is essentially to increase the signal - noise ratio in the response to an input to the system - thus enhancing the strength of output, should an input be detected (Chiodo & Berger, 1986; Foote, Freedman & Oliver, 1975; Bloom, Shulmann & Koob, 1989). This suggestion strongly parallels an emerging view based on recent basal ganglia work, namely the view that dopamine reacts to a contextual input in such a way as to enhance outputs which are consistent with the context and to suppress competing alternatives. This notion is also consistent with the finding that when striatal dopamine is depleted, the pallidal neurons which are targets of striatal projections show excessive and *unselective* activity (Filion et al., 1989).

This model of dopamine action finds a natural interpretation in connectionist networks which incorporate units with a logistic activation function through modulating the gain of the logistic function. The standard logistic function with a variable parameter representing gain is shown below in equation 5.2:

$$\text{activation} = \frac{1}{1 + e^{-(\text{gain}*\text{netI})}} \tag{5.2}$$

In this equation 'netI' is the net input to a unit and 'gain' is a positive real number. The gain parameter alters the slope of the function, so that with a reduced value of gain the slope becomes flatter. Servan-Schreiber, Printz and Cohen (1990) demonstrated that across a chain of units the value of gain could have a significant impact on whether a response was correctly or incorrectly made. This method was subsequently used in a feedforward network to model schizophrenic performance on the stroop task (Cohen and Servan-Shreiber, 1992). In this model potentiating effects of dopamine were simulated by increasing the value of gain in the relevant units in the model. Our rationale here is similar: to model a decrease in available dopamine, we reduce the value of gain in the FM.

Two different levels of damage were inflicted. A 40% reduction of gain was used to simulate mild parkinsonism, and 70% gain reduction to simulate advanced parkinsonism. The effects of damage were compared with a control condition using the undamaged performance of the network and, with two different forms of 'unmotivated' control damage. The first type of control damage was to add random, normally distributed noise to the FM weights. This provides a comparison with the hypothesised effects of gain reduction, whilst still affecting only the processing of the FM and associated state units. The other type of control damage was to add random noise directly to the state units. This will affect sequencing but not the FM prediction. The level of disruption caused by each manipulation was matched on the basis of the increase in error it caused for each element of the learned sequence. The effects of gain reduction in the FM were compared with the effects of equal levels of gain reduction in the controller subnetwork. In the current version of the network there are only two state units and these are connected to the FM output units. This has the consequence that reducing the gain in the controller will not preferentially damage the sequencing over individual responses.

Figure 5.7 illustrates the effects of both a 40% and 70% reduction of gain on the slope of the logistic activation function.
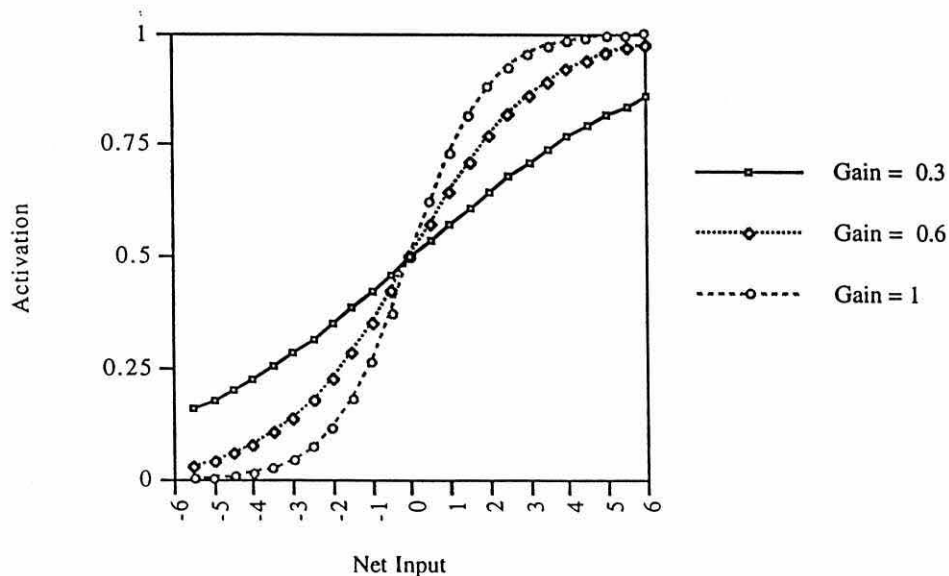
Figure 5.7. *The effects of reduced gain on the slope of the logistic activation function*

In addition to a progressive decrease in the angle of the slope as gain is reduced, there is increasing linearity of the function such that when gain is severely reduced the function is almost linear.

5.5.2 Assessment of deficits

Since the current version of the model does not have a temporal competence (Morasso & Sanguineti, 1992) (i.e. each sequence element is performed at each discrete time slice) there is no satisfactory way of assessing response time or movement time in performance. This is quite a severe limitation of the model since much parkinsonian impairment in sequential performance is exhibited as prolonged inter-response times (see chapter 2 for review). However, deficits that may be investigated using the model correspond to those kinematic or spatial deficits discussed in section 5.1:

Reductions

This classification covers a number of the characteristic deficits associated with hypokinesia. Reductions would be exhibited by the model if movements undershoot to targets or there is a tendency to avoid extreme changes in joint angles/ shoulder movement.

Deletions

Any deficit in the model's performance that leads to the omission of a sequence element or failure to complete a sequence may be regarded as indicative of a parkinsonian impairment. In contrast serial order errors such as transposition errors, which are not a feature of PD, will be not be regarded as simulating basal ganglia dysfunction.

Selective sequential deficits

Any deficit that is indicative of a selective impairment to motor sequencing should increase in severity as the sequence progresses. If the impairment has equal severity across all sequence elements then a selective sequential deficit is not indicated.

Learning impairments

Jordan (1992) has demonstrated that once a sequence has been learned, the network is able to generalise to a shifted version of the same sequence within only a few learning trials. This provides a useful task for examining learning impairments in the experimental

conditions tested. This is related to evidence demonstrating a failure in PD to encode high level structural information in complex motor programming (evidence reviewed in chapter 2). Thus we might expect relearning to take considerably longer in the experimental condition on a time scale similar to learning a sequence from scratch.

5.5.3 Simulation 5.2: A comparison of mild gain reduction in the forward model and controller subnetworks

The aim of this simulation was to compare the effects of selectively damaging either the forward model subnetwork or the controller subnetwork using the gain reduction method to simulate the effects of reduced dopaminergic activity in that subsystem. The level of gain reduction in this simulation was set at the 'mild' 40% level. Damaging the forward model subnetwork is assumed to be analogous to damage to a subsystem that is involved in sequencing, but not directly involved in computing the parameters of motor programs. By contrast, damage to the controller subnetwork is included as a 'dissociation' control which reflects impaired computation of motor parameters, but not sequential information. Reducing the gain in the forward model is thus intended to simulate impaired basal ganglia operation in motor processing, whereas reduced gain in the controller may be regarded as damage to some other neural structures involved in motor computation.

Method

The weights of the learned network were held constant, and the network was run with learning turned off, for each of the conditions described above. In condition A the forward model was damaged and in condition B the control subnetwork was damaged. In each case the reduced gain was applied to both the hidden and output units of the appropriate subnetwork.

Results

The results of both manipulations are depicted in figures 5.8 (a) and (b) overleaf.

(a) Gain reduced by 40% in the forward model subnetwork

(b) Gain reduced by 40% in the controller subnetwork
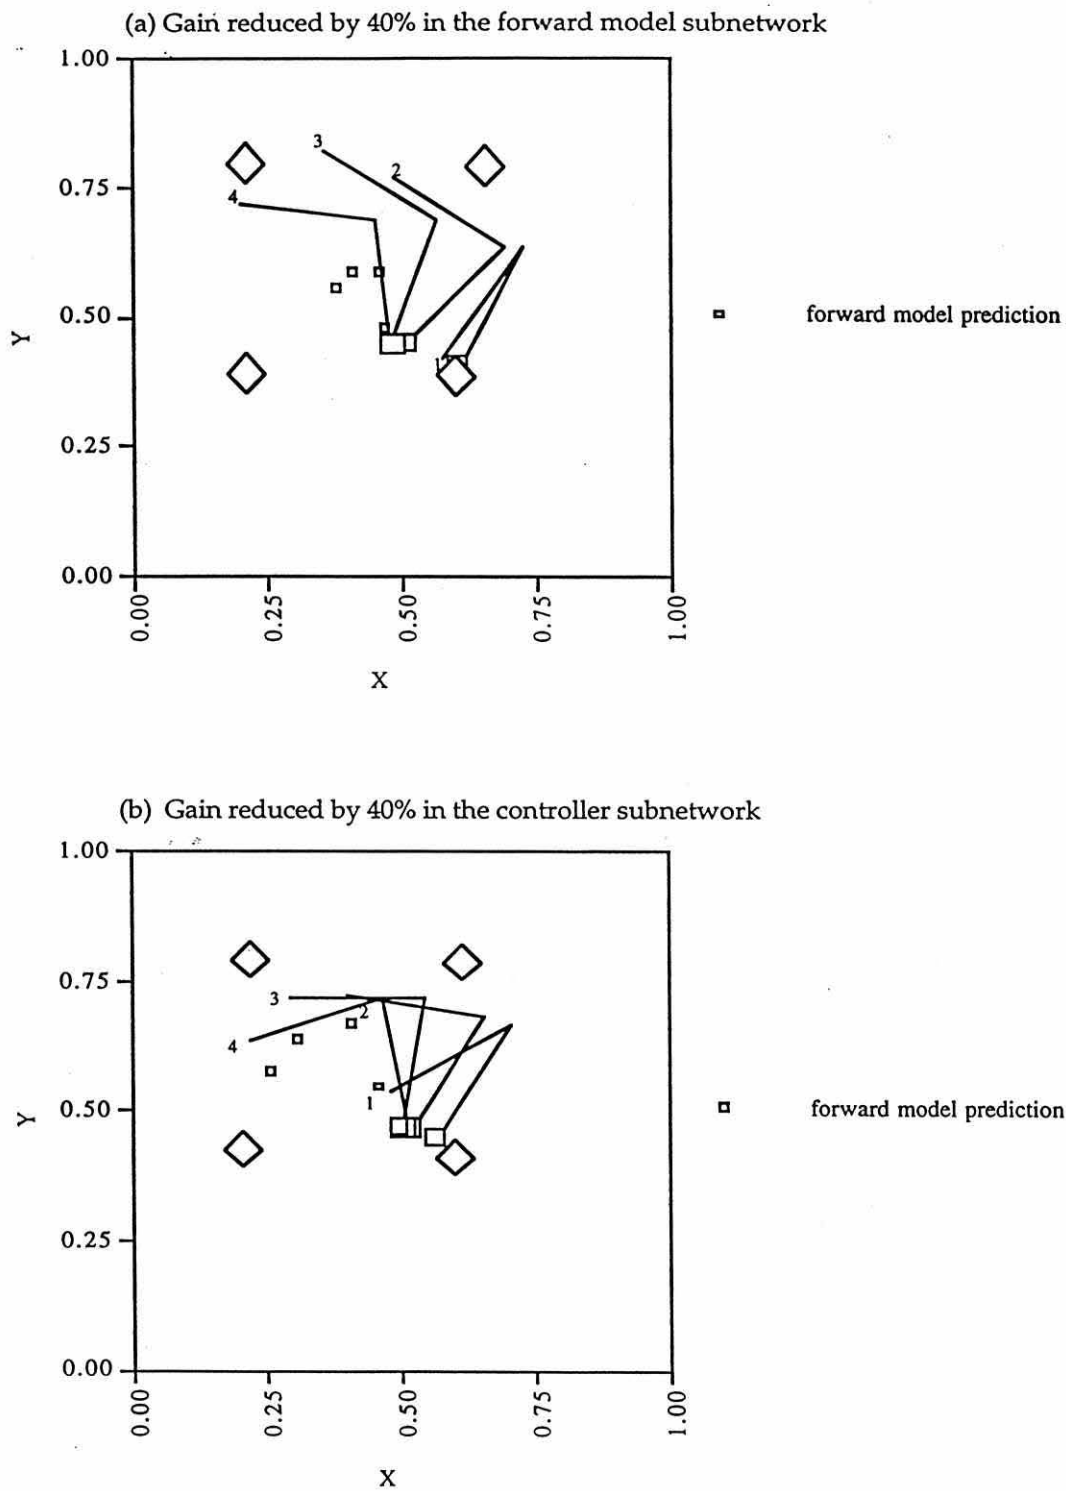
Figure 5.8. The effects of 40% reduced gain in conditions A (FM damage), and B (controller damage)

As before each figure depicts the state of the arm at each successive time slice and also the locations of each of the four targets, indicated by clear diamonds. The learned sequence is to touch each target in an anti-clockwise direction, starting at the lower right-hand location.

The endpoint location corresponding to each of the four timesteps/targets is indicated numerically on the figure.

In both conditions there are similar reductions in area covered by the sequential action as a whole. Other simulations which were conducted using different initial weight configurations led to similar results. The decrease in range of movement is primarily in the proximal joint which is the one that does most of the work in the normal performance depicted in figure 5.6. However the two manipulations do differ in important ways:

1. In condition A (FM damage), movement to the first target is unaffected, whereas in condition B (controller damage) that too is reduced.

2. The FM estimate is much more inaccurate when the FM has been damaged than when it is intact.

3. When the FM is damaged there is an approximate increase in error across sequence length, but this property does not hold in condition B (element 2 is as far from its target as element 4, and element 3 is the most accurate).

Discussion

Points 1 and 2 above are direct consequences of the nature of the manipulation. Due to the forward propagation of activation in the network, the first endpoint location is calculated without any dependence on the FM output, thus the first element will be unaffected by damage to the FM. This result indicates a potential flaw in the logic of the above comparison. If the first element of the sequence is taken to be the starting point of the network, which is reasonable because the initial values on the state units encode this position, then the disruption to the controller does not even allow the network to compute the start position. But if the initial position forms part of the feedforward motor program, and requires computation, then the FM will be involved in the computation. This problem does not however affect comparison of the remaining elements of the sequence.

Point 3 is a consequence of the fact that the FM only is connected to the state units. As the distance between successive FM outputs is reduced by the decrease in gain, thus making each successive estimate of the state more similar to each other. This reflects the processing dysfunction which we have hypothesised leads to parkinsonian symptoms.

In accordance with the hypothesis, it can be seen from figure 5.8(a) that there is a reduction in movement amplitude, such that the network makes transitions which are consistently too small and of approximately constant amplitude for elements 2, 3, and 4 of the sequence. Yet, at each time step the network is attempting to make the transition from one sequence element to the next.

Superficially this deficit cannot be distinguished from the processing deficit in the controller. However, the deficit produced by damage to the controller does not preferentially affect sequential performance, since it does not affect the state unit representations. This point is illustrated by comparing the size of the error for each condition across sequence elements during task performance (figure 5.9).
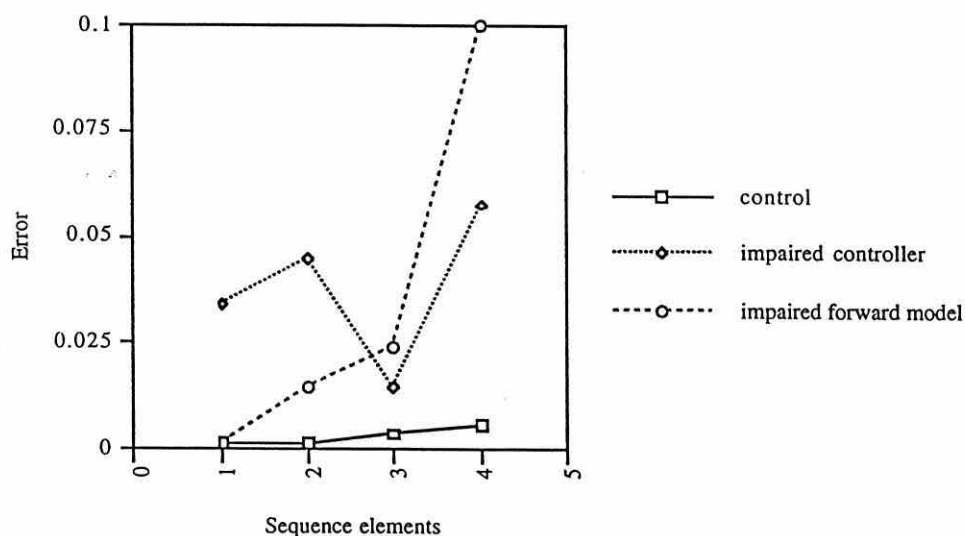
Figure 5.9. *Error scores across sequence elements (gain reduced by 40%)*

It should be noted that these error values are derived only from the simulation presented here as an illustration. If averaged scores were used, then the error for the impaired controller would be approximately the same for all sequence elements.

In sum the results of this first simulation indicate that the effects of a 40% reduction in gain are to cause reductions in movement amplitude leading to undershoots to targets in both conditions. However, only damage to the FM caused selective sequential deficits indicative of a 'parkinsonian' disruption.

5.5.4 Simulation 5.3: A comparison of severe reduction in gain in the forward model and controller subnetworks.

There is a well established difference in the motor disorders literature between 'mild' parkinsonian impairments, which are normally observed, either in younger patients or during 'ON' periods of dopamine-replacement medication and severe deficits which occur in older subjects and notably during 'OFF' periods. It is under these conditions that the akinetic disturbances, such as freezing, are most evident.

Thus, we conducted a further simulation in which the gain reduction was increased to 70% to simulate the effects of severe dopamine reduction and consequently a much more severe disruption to processing. All other details were the same as in previous simulations.

Results

The effects on network behaviour resulting from this manipulation are depicted in figures 5.10 (a) and (b) below. In both conditions the effects of the disruption are much more severe, as would be expected. In both cases it can be seen that the FM estimates are very tightly clustered and are shifted towards the initial position of the network. However, in condition A, the network makes the transition from sequence element 1 (the initial position in this condition), to element 2 in two time steps. Interestingly, this indicates that even though an undershoot has occurred on the first time step, the network still moves towards the target for that time step, rather than moving on to the next target. This interpretation is warranted since, in moving towards target 2 at the third time step, the endpoint is moving *away* from target 3. At time step 4 the network remains stuck at target 2 and fails to complete the rest of the sequence.

In contrast, condition B (damage to the control subnetwork) merely exhibits the failure of the network to move away from its initial position, although the direction of the little movement there is indicates that the endpoint is being drawn towards the nearest target to its initial position.

(a) Gain reduced by 70% in the forward model subnetwork

(b) Gain reduced by 70% in the controller subnetwork
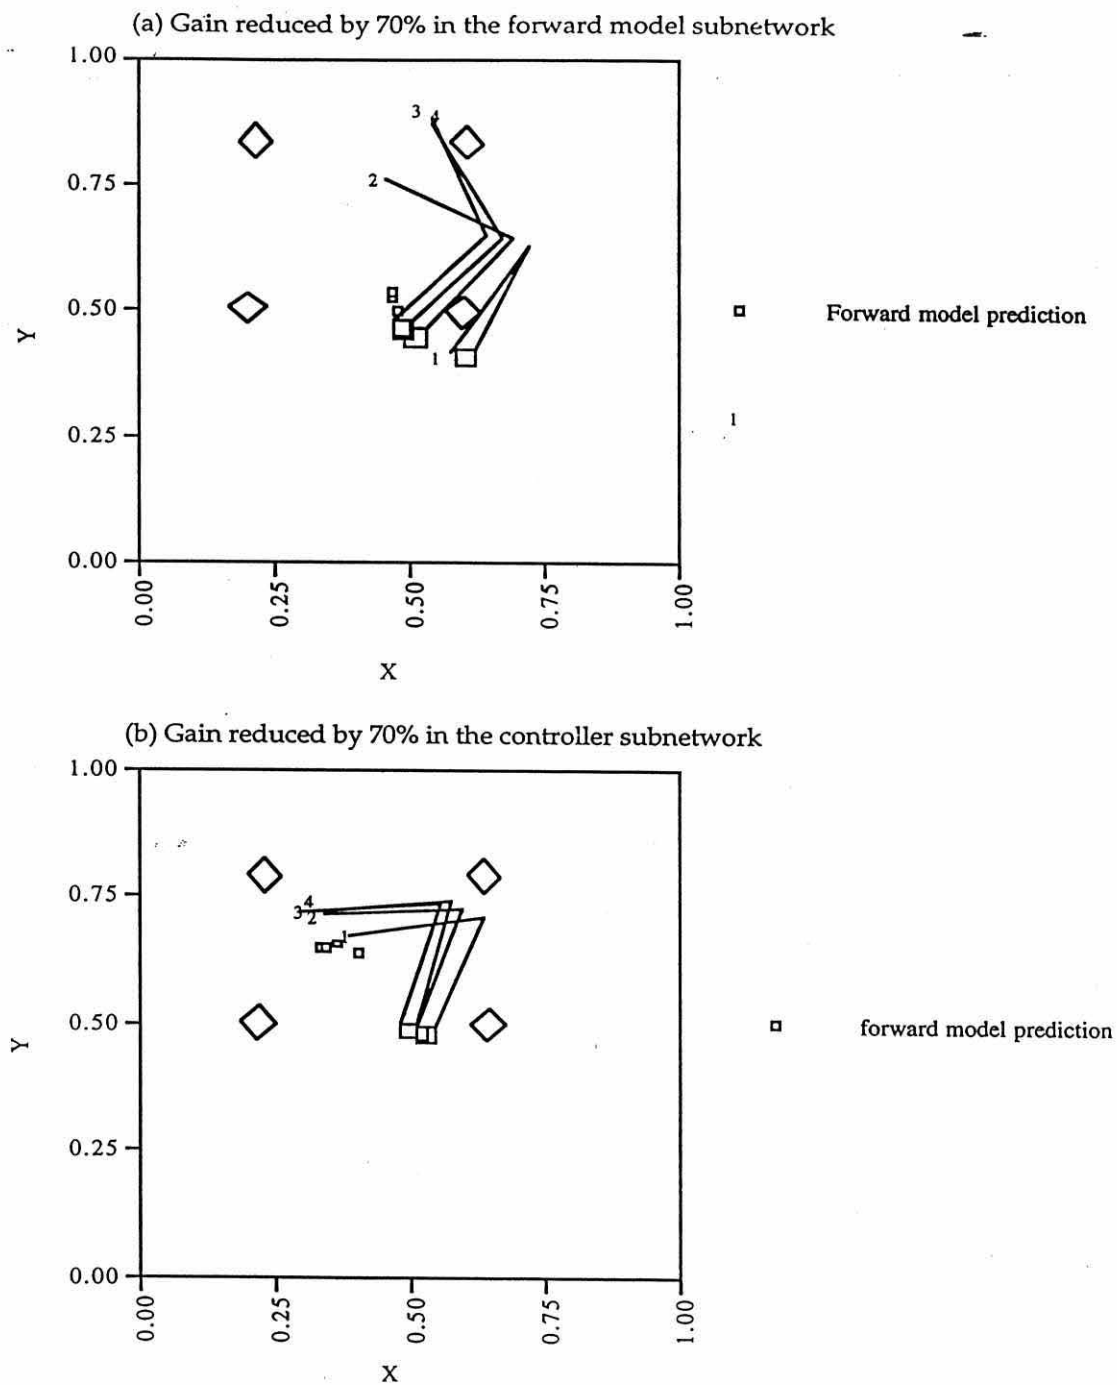
Figure 5.10. *The effects of reducing the gain by 70% in condition A (top panel) and condition B*

(bottom panel)

Discussion

The performance impairments induced by severe gain reduction in both conditions are more pronounced than those caused in the previous simulation, leading not only to reductions in

the size of movement indicative of parkinsonian hypokinesia, but also failure to complete the sequence. This can be related to work reported by Berridge (1989) and Berridge and Whishaw (1992) in which mice with striatal lesions fail to complete innate grooming sequences. It can also be related to the profound rigidity and freezing that occurs in PD patients when they are 'OFF' medication .

There are again substantial differences between the effects of the disruption in the controller subnetwork and the forward model subnetwork, which support the functional hypothesis that the basal ganglia are involved in a separate function to the direct computation of movement parameters, which has more to do with sequencing control at a higher level and in particular the potentiation or selective disinhibition of each individual action in the sequence based on an internal temporal context signal (the state vector).

The qualitative nature of the deficits produced in this simulation have their roots in an interaction between the manipulation (gain reduction) and both the nature of the representations learned by the network and the computational basis of sequential processing in the Jordan network. A discussion of the implications of these computational features for modelling motor disorders is provided in section 5.6.

5.5.5 Simulation 5.4: Damaging the forward model with lesions not motivated by hypotheses of dopamine depletion

The previous simulation reproduced some of the basic parkinsonian deficits that would be expected if the nature of the disruption reflects striatal dopaminergic reduction. The aim of this simulation is simply to examine the deficits that are produced in control conditions where the damage is 'unmotivated' and where the effects of damage would not be expected to yield parkinsonian symptoms. The nature of the control disruptions has been described in the general methodology above. Both disruptions involve the addition of random noise to those parts of the network affected by the disruption in condition A in the previous simulation. The first of these involves adding noise to the state unit activations themselves, the second to the weights of the FM. In both cases the noise has a gaussian distribution and the power of the noise is calculated as a percentage of the normal value of the weight or

unit activation. In both cases the noise is rather large and is approximately matched to the total error produced by a 70% reduction in gain on the FM. The primary difference between the two disruptions is that noise applied directly to the state units would be expected to produce more severe deficits in output behaviour, since the previous states have already been incorporated. Alternatively, when the FM is disrupted in this way, the effects will likely be dampened by the subsequent computation of state unit history.

Results

The effects of the two control disruptions are depicted in figures 5.11 and 5.12 below.
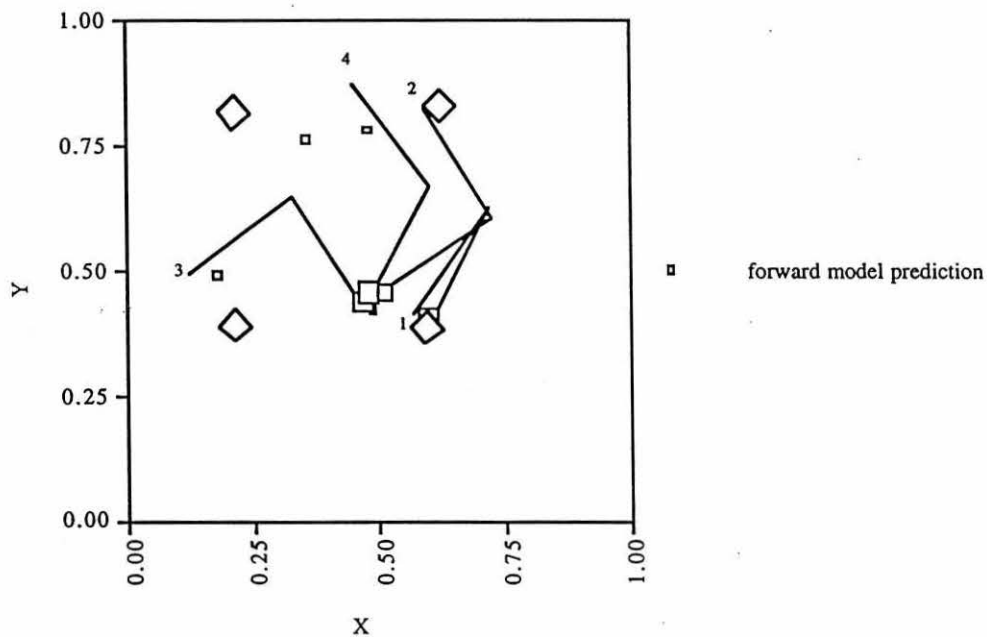
Figure 5.11. *The effects of random noise on the state units, (overall error matched to 70% gain reduction)*
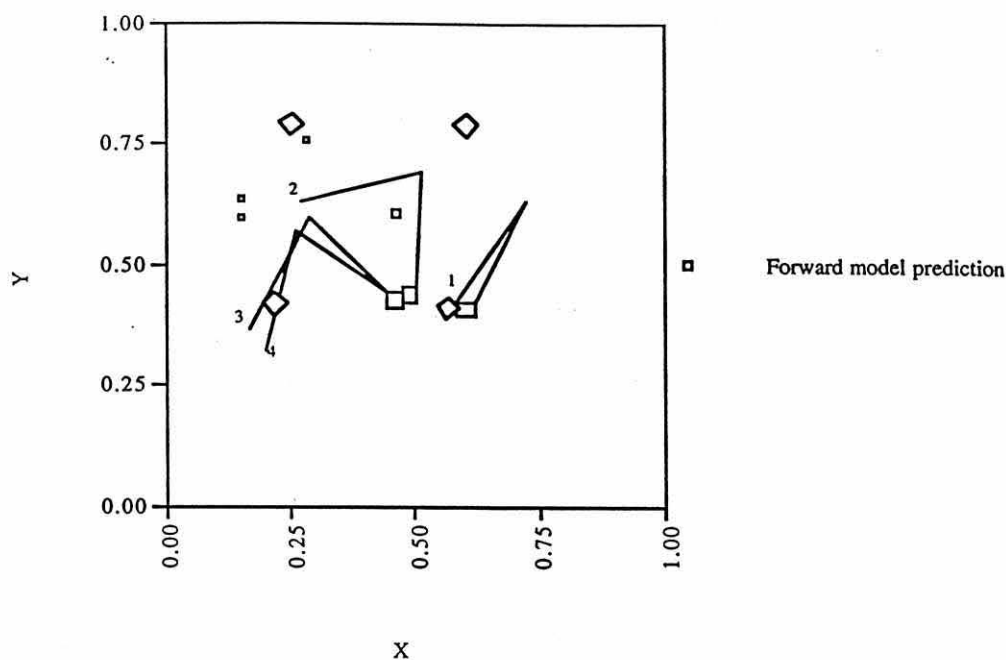
Figure 5.12. *The effects of random noise applied to the FM weights, (power = 50%)*

It can be seen that in neither case are parkinsonian effects produced. In the first case (figure 5.11) the transitions between elements three and four are excessive and there is a breakdown of serial order, (elements three and four are transposed). This is because there is excessive and errorful difference between successive state representations and the property of temporal similarity is destroyed.This is not a feature observed in PD and is not consistent with theoretical ideas about dopamine reduction.

In the second case the serial order of the original sequence is maintained. Whilst the intact temporal averaging process is likely to ensure this for lower values of noise, it is unlikely, due to the random element, that it will hold in all cases, especially with very large amounts of noise. As would be expected the close clustering of FM estimates exhibited when the gain was reduced is not produced.

<u>Discussion</u>

These results indicate that it was the experimental manipulation that produced the parkinsonian nature of the deficits in the previous simulations, and not some more general property of the learned task. These deficits are not reproduced by the addition of random noise.

5.6 Interpreting the behaviour of the model in terms of motor control and movement disorders

In the learned network each of the target locations corresponds to the output of the network when an individual 'motor program' represented by the current articulatory configuration is instantiated. Which of the four learned motor programs is produced at any one time step is determined by the pattern of activity on the state units. The state units encode a decaying activation trace of past outputs of the FM, the depth of which is determined by the value of the parameter μ. The pattern that exists on the state units may be thought of as a temporal context that facilitates or potentiates the production of one of the learned motor programs over the other three alternatives.

In order to develop an account of the disrupted behaviour of the model in relation to motor disorders, it is first necessary to discuss explicitly two fundamental properties of the model. The first is the nature of the representations formed in the learned network, and the second concerns the effect of the disruptions on the movement of the state vector over time.

5.6.1 Attractor dynamics and sequence representation in Jordan networks

In considering the representations formed by the network, it is useful to separate the notion of representing each target state (sequence element) individually, from the representation of the sequence as a whole. During learning the network effectively builds a higher order unit of representation, which encodes the sequence that is learned. Jordan (1986b) has shown that the learned sequence acts as a limit-cycle attractor, such that if the network is started near the learned sequence it will move towards the trajectory encoded as the sequence. Each target location can also be thought of as an attractor, although in this case it is a point-attractor represented by the location of the target.

An easy, although strictly-speaking inaccurate, way to imagine the state space of the learned network for this sequence is to imagine a three-dimensional surface much like an egg tray for four eggs, where the target locations lie at the bottom of each of the four egg cups. Each egg cup represents a basin of attraction for each sequence element. Now one may imagine a shallow groove that cuts a path through the ridges separating each basin.

This groove represents the 'virtual' trajectory encoded by the network. The qualitative properties of attractors, such as the depth and layout of basins have important consequences for the behaviour of the network when it is damaged.

5.6.2 The effects of reduced gain in the forward model vs. the controller subnetworks

The four learned states of the network are not simply discrete alternatives. Jordan (1992) has shown that there is continuity in the mapping between states and outputs. This property means that when the state unit representation is degraded, the level of implicit competition and potential interference from those alternative actions, which are partially facilitated by the degraded contextual representation, is increased.

The simulations conducted here demonstrate an interaction of this 'continuity' property with the reduction in gain, such that, (in both conditions), when gain is severely reduced the network will move slowly towards the cyclic attractor[9]. However the differences between condition A (FM damage) and condition B (controller damage) suggest that this behaviour also interacts with the recurrent connections to the state units, when the gain reduction affects the sequence information encoded by the state units.

Jordan ascribes the attractor dynamics properties in his simulations, in part, to the μ parameter which allows the state units to encode the recent past history of state unit activations to an arbitrary length (the length depending on the value of μ). In brief, the existence of the recurrent connections of the state units onto themselves means that nearby states in time are more similar to each other, and since the state units encode the position of the endpoint in task space this confers the useful property of smoothness in task space on network performance. The size of μ dictates how much sequential structure is embedded in the state unit representations. Thus if $\mu = 0$ the state units encode only the activation on the output units at the previous time slice, but if $\mu = 0.5$, for example, then the state units encode an exponentially weighted average of half each previous activation into the arbitrarily distant past. This means that if the distance between each successive movement is reasonably large, with μ greater than 0, each successive state is more distinct from the

---

[9]This assumes that the network has been started near the attracting sequence.

previous state, than it would be if μ was 0, but if the movement amplitude is small then $\mu >$ 0 acts to make each successive state, and consequently, each successive movement, more similar to the previous one.

This interpretation can provide an explanation for the difference between conditions A and B in this simulation. In condition A, the network correctly computes the articulatory unit activations at time step 1. Due to the disruptive manipulation, the FM estimate is inaccurate but nevertheless constitutes a sizeable difference in the state unit activations, when it is fed back, since they were previously 0. Now, at time step 2, even though the FM estimate hardly changes at all due to the gain reduction, there is enough 'energy' in the system on account of the previous large change in state unit activations to propel the system into the basin of attraction for target 2. Now, at time step 3, the initial large change in state unit activation has far less weight than the more recent change which was minimal, although it is sufficient to move the endpoint very close to target 2.[10] However, there is now, very little 'energy' in the system due to the unchanged FM estimate, and the network is trapped at that point.

In contrast, in condition B, in which the controller is damaged, the network fails to compute the articulatory activations correctly at the first time step due to the intervention, and consequently the network is started at an unvisited initial position. However, the FM estimate is quite consistent with the actual endpoint position, since the FM is undamaged. Although the difference in the state unit activations is as large as in condition A at this first time step, the actual distance travelled by the arm at the following time step is reduced due to the gain intervention in calculating endpoint position, and so on throughout the sequence.

The existence of attractors at the target locations make a critical difference to the behaviour of the network as gain is reduced in the FM. By comparing figures 5.8 (a) and 5.10 (a), it can be seen that, although the endpoint location at time 2 is similar in both cases, when the gain reduction is mild, the system is driven on towards target 3, because there is less interference from element 2 in the state unit representation, than when the damage is more severe. However, it then appears to fail to make the transition to element 4,

---

[10]Note that the existence of an attractor exerts its influence at this time step, such that the endpoint continues towards target 2, rather than moving towards element 3.

although there is insufficient information to tell whether it would make this transition given sufficient time. It is important to note that target attractors are not a property of the environment itself, but of the combined environment / task representation encoded by the weights in the control subnetwork. This has considerable bearing on the application of computational principles to fundamental theoretical issues in motor control.

5.6.3 How does the model bear on issues of motor program representation and sequential control?

Condition A damages only the internal aspects of processing which affect FM estimation and consequently the internal contextual representation which would lead to a particular deficit in sequential behaviour. It does not directly affect the computation of movement parameters or the actual endpoint of the arm. This is consistent with the hypothesis that the basal ganglia are involved in action selection, rather than motor computation. Condition B, on the other hand, does affect parameter computation and does not have any impact on sequential processing. This is more consistent with what might be expected in the case of primary motor cortex or cerebellar damage.

Both conditions exhibit parkinsonian-like reductions and deletions in output behaviour, but only those in condition A can be accounted for in terms of a failure to shift between successive actions, due to a dysfunction in the internal representation of context.

The model demonstrates some interesting properties regarding the nature of the representation of internal context, and the interaction with damage. However, there are also some fundamental limitations of the model which give way to concerns about the 'predictive' role of FM estimation and its effect on behaviour.

5.7 Generalisation impairments following damage

A variety of evidence indicates that the cortico-striatal loop involving prefrontal cortex plays some role in sequence learning (evidence reviewed in Curran, 1995). In particular there have been a number of recent suggestions that implicit sequence learning is impaired in Parkinson's disease. Precise claims regarding the root of the reported deficits vary.

Suggestions include attentional dysfunction, a failure of procedural learning mechanisms, and difficulty with learning higher-order sequential representations. However, conclusive evidence that pinpoints the nature of basal ganglia contribution to motor sequence learning is still lacking (Curran, 1995).

According to the conceptual model outlined in chapter four, impaired sequence learning in PD could be attributed to a failure to acquire a strong internal contextual control signal, through a breakdown of interaction between prefrontal cortex and striatum. This dysfunction may be hypothesised to lead to increased interference from established motor patterns with similar contextual cues.

This hypothesis is consistent with behavioural data indicating a failure in PD to adapt old or well-learned procedures to accommodate new task constraints (St. Cyr et al., 1988; St-Cyr and Taylor, 1992). It also makes sense of reports that PD's have difficulty in switching motor set, Robertson and Flowers (1990) and that they have trouble superimposing two different motor plans, (Benecke et al., 1987; Flash et al., 1992). Although only some of the available studies have found that primary sources of error are intrusions from other learned sequences, it is reasonable to assume that few have explicitly tested the possibility.

Jordan (1992, p. 420) reports a simulation in which he demonstrates that a network of the type used here can generalise to a shifted version of the original sequence in only a small number of learning trials (he reports that the network is close to a solution on the second sequence after only five trials, compared with several hundred or thousands of trials for initial learning). This task provides a more interesting case for examination of the learning impairments after lesions than merely giving the network a new sequence to learn from scratch, since it is possible to explore the extent to which previous learning interferes with generalisation after lesions of different types and at different locations in the network.

5.7.1 Simulation 5.5: The effects of lesions on generalisation performance

In this simulation we therefore examine the learning impairments caused by the various disruptions employed in the previous simulations. In the context of the model and task, a

moderate learning impairment would be indicated by an increased number of iterations to reach a stable solution, and increased error in the final solution. A severe learning deficit would be indicated by failure to reach a stable solution, or failure to decrease error below that achieved on the first trial.

Method

The task in this simulation required the network, initialised with the learned weights used in the previous simulations, to learn an identical sequence in which all targets were shifted downwards in task space. The lesions that were used to damage the network were identical to those in the previous section.

Results

In the analysis of results it was assumed that generalisation is only evident if the error reduces to a satisfactory level over a relatively short period (e.g. tens of iterations). Error reduction over a long period (thousands of iterations) indicates that the network is undertaking a major reorganisation of the weights, and thus losing the solution to the original sequence. Therefore, in the graphs below we only consider changes over the first 50 trials to represent generalisation. The learning criterion was set at 0.1 error for the sequence.

The first two figures below, 5.13 and 5.14 depict the error reduction and final configuration of the arm respectively, after 90 trials of learning on the new sequence.
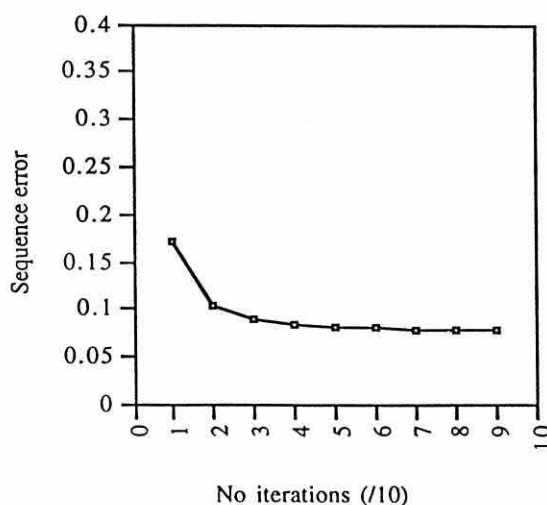
Figure 5.13.*Generalisation performance in the undamaged network*

It can be seen from the above figure that the network achieves a stable error score below criterion, (<0.1), after the first 30 trials. Although this is more than the five trials reported by Jordan, this is of an order comparable to that reported in Jordan (1992). However, it can be seen from figure 5.14 that most of this error is accounted for by the first element which has failed to generalise from its learned position in the original sequence. Overall it is clear that generalisation is better on those sequence elements where the targets lie within the boundaries of the original sequence.
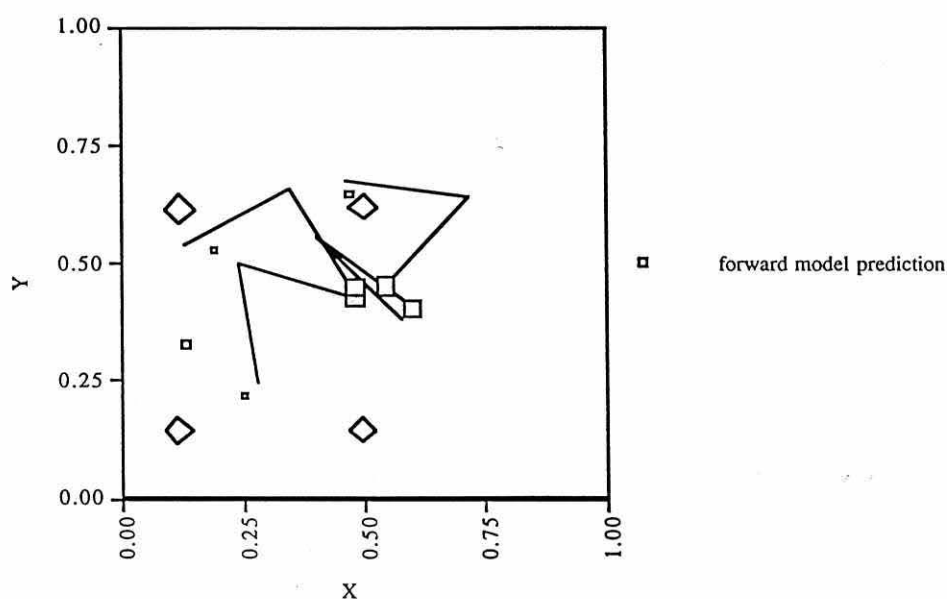
Figure 5.14. *Solution to the shifted version of the sequence after 30 learning trials in the undamaged network*

Further learning trials following the achievement of this stable state led, after approx. 1000 trials, to accurate performance on all elements. However, this solution was preceded by a substantial period of increase and oscillation in the error score, which indicates that the original solution was lost in the process. The apparent difficulty in achieving accurate generalisation even in the undamaged network may be attributed to a number of factors. Jordan (1992) states that generalisation performance in particular cases will depend on the relationship between sequences, but also the target degrees of freedom and the degrees of freedom in the output representation. In this case there is evidence of interference between

the two sequences and this factor must be considered when assessing the effects of disruptions.
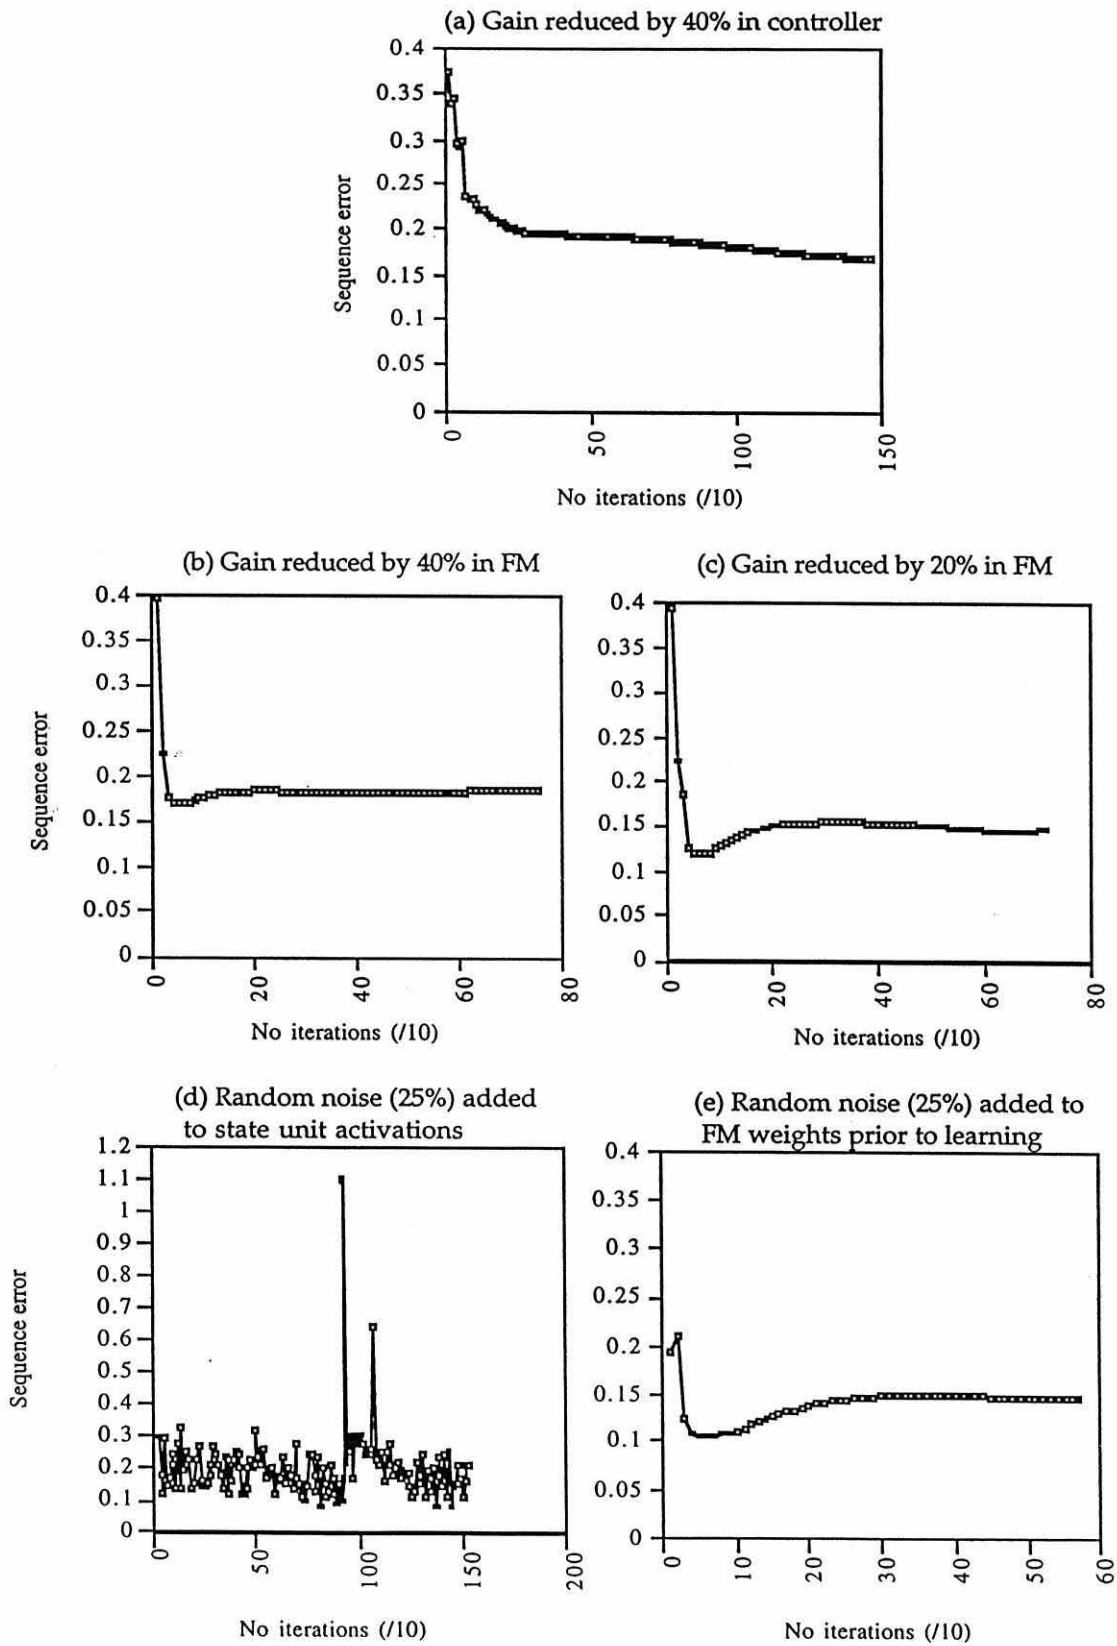
Figure 5.15. *Generalisation performance of the network following damage*

The effects of disruptions on error reduction are depicted in figures 5.15 (a - e) below. Only the mild form of gain reduction (40%) was used to compare differences between damage to the controller and FM subnetworks (a & b). An even milder (20%) reduction in FM gain was used to examine any more subtle effects that might arise in this condition, (c). Finally this may be compared with the effects of the two control disruptions used in the previous simulation, (d & e).

In all of the above cases except figure 5.15 (d), learning asymptotes at a similar error score (between 0.13 and 0.18). The more severe damage caused by 40% gain reduction in figures 5.15 (a) and (b) result in the slightly higher value. However, the profile caused by damage to the controller, (a) differs from the other three, (b, c and e). In the latter three cases there is evidence of some generalisation in the reduction of error below the final asymptotic value in the first fifty trials. The subsequent rebound indicates convergence towards a local minimum caused by interference between sequences. When the controller is damaged (a), the network shows no indication of generalisation and simply converges slowly toward the local minimum state.

The effects of adding noise to the state units (d) clearly prevent the network from learning. Due to the use of gradient descent learning, it is not surprising that this manipulation prevented the algorithm from working. However it would be expected that noise of this magnitude applied at each learning trial would have a similar effect regardless of the learning method.

The existence of a local minimum at the asymptotic error score of (approx. 0.15) is further revealed by the similarity in arm configurations when the effects of gain reduction are compared with the effects of noise on the FM weights in figures 5.16 (a & b) below.

(a) Gain in the forward model reduced by 20%

(b) 25% noise added to the forward model weights
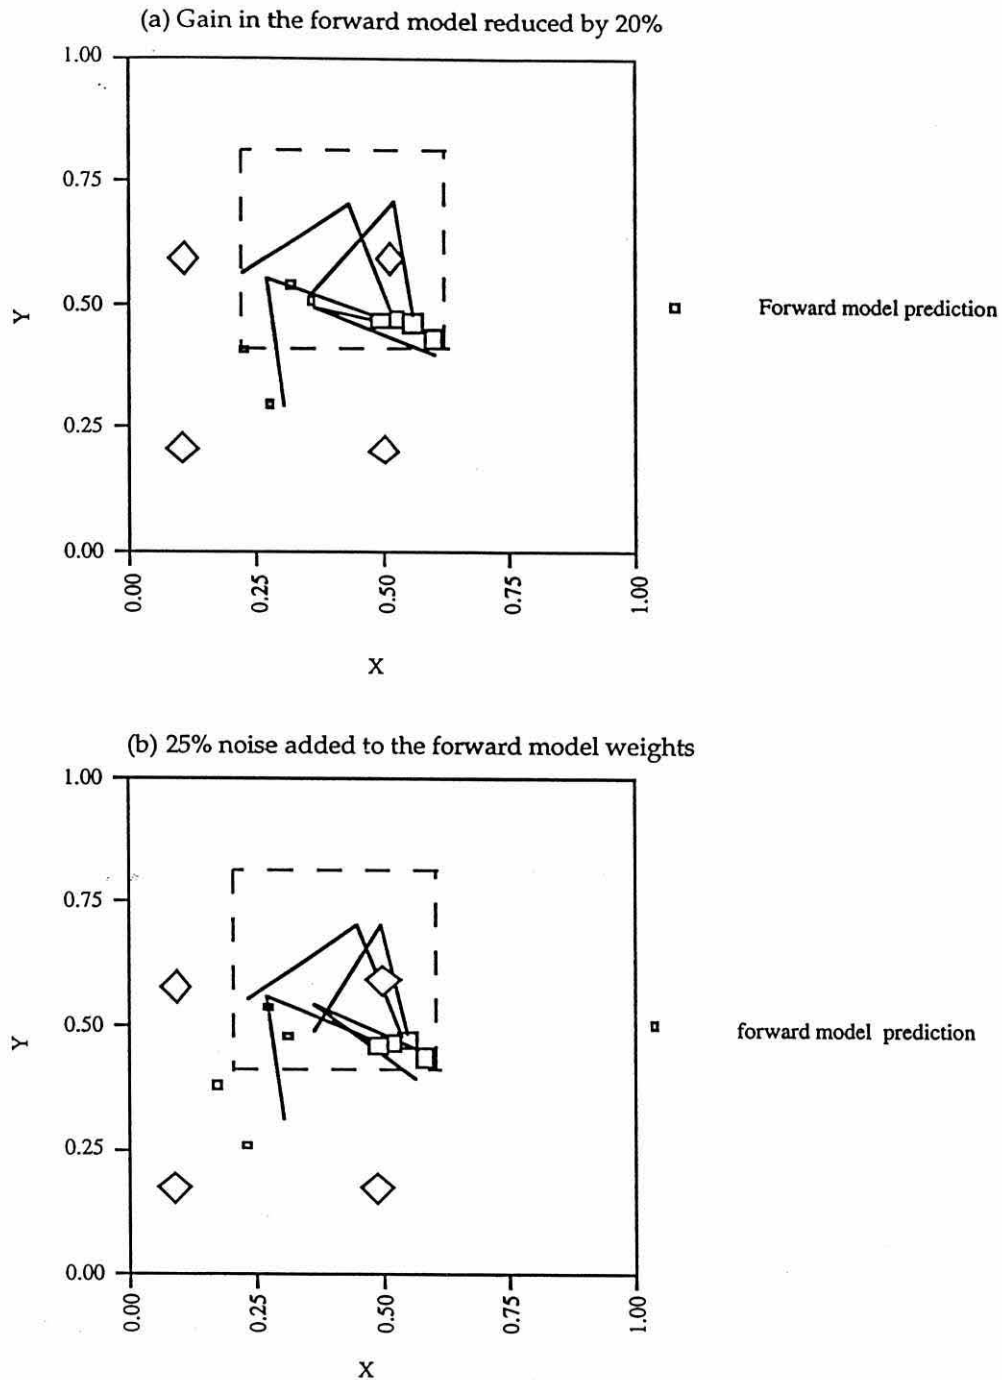
Figure 5.16. *The effects of gain reduction versus noise in the FM weights on network performance (the trajectory of the previously learned sequence is shown by the dashed line)*

Discussion

It is clear that the results of this simulation were heavily influenced by the difficulty in achieving an accurate generalisation even in normal conditions. Nevertheless some of the features of the damaged conditions are worthy of comment. Generalisation performance in

the model is a function of both task difficulty and the extent and site of damage caused by gain reduction. Under more favourable conditions both levels of gain reduction modelled here may be expected to allow generalisation to occur, although in approximately double the time taken under normal conditions. A more severe level of FM damage (70%) was also examined, although not reported above. This condition exhibited behaviour very similar to the controller damage reported. In the case of damage to the controller, no evidence of generalisation was found. This suggests that damaging the controller in this way, rather than the FM, is less likely to lead to generalisation under any conditions.

The extreme similarity between the 'solutions' in task space when gain was reduced as opposed to adding noise to the FM (see figures 5.16 (a) and (b)) is interesting because, in the previous simulation, these manipulations had quite different effects on the output behaviour of the network. It will be noticed that in both cases the distribution of endpoint locations lie within a region of task space where the two sequences intersect. This finding yields the interesting possibility that the local minimum which impeded generalisation (even in the unlesioned network) is, in fact, a spurious attractor state in the network's task representation caused by the previously learned task. Re-examination of figure 5.14 (b) reveals that this behaviour is also evident in the generalised solution arrived at by the undamaged network.

The wider distribution of endpoint locations and greater accuracy in the unlesioned condition suggest either that the damage in the other conditions caused the attractor to assume a different size and/or shape, or that the reduction in system 'energy', which we referred to earlier, makes the system less able to resist or avoid the spurious attractor state formed at the intersection. The effect of the attractor state is powerful and appears to outweigh the previous differences in network behaviour caused by the different types of damage. To see this, reconsider figure 4.21 (d). It can be seen that the mean error score, despite large fluctuations, lies in approximately the range (0.15 - 0.2) which indicates the attractor state.

The interaction between attractor states and measures of system energy are well known and have been extensively studied in other types of recurrent network, in particular

Hopfield nets (Hopfield, 1982) and the Boltzmann machine (Hinton and Sejnowski, 1986)[11]. Thus, it is not unduly surprising to find similar general properties apply in this network. However, the attractors produced by these networks are static (point attractors), whereas the attractors formed here are dynamic (cyclic attractors). Notice that the endpoint locations in the above diagrams do not converge to a point in either case. But the cycle of states is reduced in amplitude in the damaged condition.

These findings are of considerable interest both in analysing the response of these networks to damage and in the computational study of dynamical action representation and motor disorders.

5.8 Chapter Summary and Discussion

At the start of this chapter we presented a computationally motivated hypothesis that the function of the cortical - basal ganglia loops can be characterised as computing a FM. The general form of the hypothesis is consistent with both neurobiological and behavioural data concerning the basal ganglia and motor disorders. It is also consistent with a number of previous theoretical suggestions regarding the role of basal ganglia motor circuits and with our own theoretical approach. The computational model we used to explore the hypothesis draws on two previous sources of computational work: the sequential network of Jordan, (1986a, 1990) for modelling the control of movement, and the disruption of a gain parameter applied to the logistic activation function for simulating dopaminergic reduction (Cohen and Servan-Schreiber, 1992). The effects of damaging the network indicate a number of limitations in the model, but also a number of interesting properties which require further investigation. Both of these are briefly discussed below.

5.8.1 Simulating PD motor deficits

The gain reduction disruption of the FM reproduces a variety of postural, kinematic and sequential deficits observed in PD. These impairments are not reproduced by the other

---

[11] The Boltzmann machine utilises a global parameter called Temperature, which reflects system energy during the learning process. It is interesting to note that the temperature parameter is applied to the logistic activation function and acts , in that network, in an analogous way to the gain parameter in this network.

types or sites of disruption that were examined. One of the major limitations of the model in its present form is that there is no way to investigate temporal disturbances in network behaviour. This is due to the discrete representation of time in the network, and the identification of each successive action with each successive time step in the operation of the model. Temporal disruptions such as prolonged onset / reaction times and increased inter-response intervals in sequential behaviour are cardinal features of the parkinsonian disorder and account for much of the available behavioural data. Moreover, they provide a great deal of information about the underlying processing deficit. In the next chapter we consider one method of introducing a temporal dimension to network processing, and examine how it affects both processing and output behaviour of the network.

5.8.2 A predictive forward model in basal ganglia motor processing?

Although the disruptions to the FM produced deficits that were qualitatively different from those arising from damage to the controller subnetwork, it would be unwarranted to ascribe those differences simply to a 'predictive' role of the FM estimate of endpoint position. This is easily demonstrated by consideration of the time course of events in network processing. At time step 1 the values of the state units are set to zero, and activation is passed forward through the controller subnetwork. The associated pattern of activity emerges on the articulatory units, which specifies the computed position of the endpoint of the arm. Activation is then passed forwards through the FM subnetwork and the estimate of endpoint location is obtained. The values on the FM output units are then transferred to the state units and time step 1 ends. Thus, it can be seen that the FM estimate is functionally inert in processes affecting the behaviour of the network at time step 1. The FM does have an impact on the behaviour of the network at time step 2, due to the recurrent connections, but this is not a predictive role since the values that are imposed on the state units are estimates of current events, not future events. A number of points concerning the role of the FM in the learned network can be made:

- The forward model is essential for learning in the current framework using supervised learning (Jordan & Rumelhart, 1992). During learning the action of the forward

model is predictive, although for this feature to become apparent back-propagation-through-time would be required (see Jordan & Jacobs, 1990).

• However, the FM does not exert its influence on the sequential behaviour of the network in the current model through prediction. The nature of the behavioural impairments produced when it is disrupted are due to the fact that it is a separate, but interacting subsystem which is not involved in the computation of control parameters directly, but influences the relationship between sequential actions due to the recurrent connections to the state units.

• Thus the forward modelling role is less crucial to the behaviour of the model than the fact that the damage occurs in a subsystem which is involved in computing information about the relationships between actions, rather than in computing the parameters of the actions themselves. It is this functional role which best reflects the theoretical approach put forward in chapter four.

• Any disruption to the FM in this model exerts its effect on the structure of the information encoded by the state units. For example, the effects of gain reduction could potentially be reproduced by any manipulation which makes successive state vectors more similar to each other. Thus a more principled and systematic approach to exploring the properties of the model could be achieved by manipulating the state representation itself, rather than invoking an otherwise redundant system to manipulate it indirectly.

5.8.3 Attractor dynamics and impairments in sequential processing

One of the most interesting features of the simulations reported in this chapter is the formation of dynamical attractors in task space. The interaction of dynamical task representations with disruptions to the representation of internal context provides a useful tool for analysing network behaviour following damage, and, in combination with the introduction of temporal processing, a potentially fruitful way to characterise the computational dysfunction in motor disorders. These issues provide the focus for the next chapter.

CHAPTER SIX

System Dynamics and Temporal Variability: Implications for Models of Motor Dysfunction

The primary goal of the work reported in this chapter is to address two computational issues that were raised in the previous chapter.

- The nature of the dynamic attractors that characterise the model's representation of sequential tasks.

- The temporal limitations of the model in the simulation of motor disorders.

Not only are both of the above issues central to the function of the connectionist network which forms the basis of the current model, but they both have important implications for theories of motor control generally, and for models of basal ganglia dysfunction in particular. Our approach in this chapter is to offer a dynamical systems interpretation of the motor deficits in Parkinson's disease. We examine the properties of the attractors built in a modified, fully-recurrent version of the neural network used in the previous chapter which permits greater temporal variability in performance. Finally we examine the response of the system to damage.

The notion of task dynamics is a central theme of the action systems approach to motor control (e.g. Kugler, Kelso & Turvey, 1982; Saltzman & Kelso, 1987; Schmidt, Shaw & Turvey, 1993) that is challenging traditional motor programming approaches. Temporal variability is a characteristic feature of normal motor control in biological organisms and is considered by many to be essential for the development of adaptive, skilled motor behaviour (e.g. Wing, 1992). Moreover, researchers in neurobiology are

becoming increasingly aware of the importance of temporal dynamics in the modulation of responses in neural systems (Hounsgaard & Mitgaard, 1989).

A number of lower level computational and mathematical approaches have focussed on dynamical analyses, e.g. dopamine transmission (King et al., 1981, Wickens, Alexander, & Miller, 1991) and tremor in PD (Beuter et al., 1991). In chapter one we introduced theoretical ideas which stress the importance of both dynamical representations and temporal aspects of processing at a systems / functional level.

The work presented in this chapter also addresses a number of sub-goals. First, it demonstrates a close inter-relationship between the two themes, outlined above, in the computational model and shows how sharpening the temporal resolution of processing in this type of network can help understand both the network's behaviour and the layout of dynamical attractors in a complex system. Second, it provides a deeper analysis of the role of the contextual signal encoded by the state vector in the sequential control of the system. The effects of directly manipulating the content of the state vector are explored as a means of studying the relationship between the dynamics of a contextual control signal and the motor representations encoded by the weights. Third, the effects of disrupting the timing relations between the two sub-systems of the network are examined as a follow up to the limitations discussed in the previous chapter.

In the first part of this chapter, we describe the computational mechanism used to introduce temporal processing to the model and apply it to the basic sequential control task studied in the previous chapter. We will refer to this as the cascade mechanism (McClelland, 1979).

6.1 The cascade mechanism

The cascade mechanism was originally developed by McClelland (1979) to model the time course of processing in connectionist networks. The basic principle of cascade processing is that it allows activation to build up gradually on units in each layer of the network, before eventually reaching a stable, asymptotic activation state. This contrasts with the usual procedure for updating unit activations in feedforward networks, whereby the complete

activation update for each unit in each successive layer of the network is computed in a single discrete time step.

6.1.1 Description of the cascade equations

The computational basis of the cascade is encapsulated in the following equation which is substituted for the normal equation governing the calculation of unit activations.

$$a_i(t) = tau \sum w_{ij} a_j(t) + (1 - tau) a_i(t - 1) \qquad (6.1)$$

The above equation essentially computes the new activation a unit based on a running average of the net input to that unit. The parameter *tau* is a rate constant that governs the speed of processing. Speed of processing is indicated by the number of cascade cycles required for a unit to reach its asymptotic level of activation.

The original cascade model in which the above equation was used, comprised only linear units. This equation is therefore not suitable for use in models utilising a nonlinear activation function. However a variation of the above scheme which is compatible with nonlinear units has since been suggested, (McClelland & Rumelhart, 1988). In this case, equation 6.1 is modified to calculate the net input to a unit, as in equation 6.2 below.

$$net_i(t) = tau \sum w_{ij} a_j(t) + (1 - tau) net_i(t - 1) \qquad (6.2)$$

The unit activation is then computed in the normal way, in this case using the logistic function.

It has been reported that a feedforward system incorporating the cascade mechanism exhibits interesting temporal properties during the course of processing (McClelland & Rumelhart, 1988). Most often, cascade models have been used to account for a variety of reaction time data (Cohen, Dunbar & McClelland, 1990; McClelland, 1979). Indeed, one of the most appealing aspects of the cascade mechanism is that it introduces a temporal capability to models of those psychological processes for which time is a critical component

of the experimental analysis. The processing involved in movement preparation or motor programming is a prime example. Reaction time (RT) paradigms have formed the basis for much of the standard research on normal motor programming and have also provided a considerable quantity of the available information on abnormal motor processing in PD (e.g. Jennings, 1995).

6.1.2 Using the cascade in a model of motor processing

A basic feature of human motor processing is the temporal variability in the organisation of both simple and sequential actions under different environmental or internal conditions, and it has been argued that this provides an important tool in the investigation of movement processing (Wing, 1992). Consequently, it may also be argued that a *temporal competence* is a requisite feature of any computational model of movement control and in particular those seeking to model processes affecting the internal control of complex / sequential actions.

A fundamental criticism of the type of system we studied in chapter four is that it lacks a temporal competence of this nature. Morasso and Sanguineti (1992) point out that in these networks the problem of learning timing structure is not addressed because the learned structure is embedded in the training data. In the previous chapter we saw how this limitation imposes an artificial restriction on our ability to observe the behaviour of the network between target states and to examine the effects of task space attractors on behaviour at a finer time-scale.

A number of variations on the current scheme have been put forward which do address the problems of learning the temporal dynamics of goal-directed movements, (Flash, Jordan & Arnon, 1994; Kawato, Furukawa & Susuki, 1987). These approaches have had some success in modelling force control and trajectory formation in motor planning problems. Both of these approaches use variations of backpropagation-through-time (BPTT) (Rumelhart, Hinton et al., 1986b). This has the advantage that the resulting temporal structure of the output is an emergent property of the network computation. Although BPTT is a sophisticated and powerful learning process, the downside is that it is

complicated and is not neurobiologically plausible. This, as we have already discussed, can actually make understanding the relationship between network computation and biological processing function harder rather than easier.

The use of the cascade equations provides a model with the capacity to demonstrate temporal variability in the processing of different input patterns, or under different conditions, without any real increase in the complexity of network processing. Learning is carrried out in the normal way and the cascade is only introduced during testing of the network.[1] Thus the introduction of the cascade mechanism does not have any impact on temporal structure or biases inherent in the training data. What it does allow, however, is a window onto temporal properties of network processing during the production of learned responses, which would not otherwise be possible.

This facility is particularly useful in analysing the disturbances of sequential processing caused by disruptions in the current model and the interaction between network lesions and the task space attractors identified in the previous chapter.

6.1.3 Using the cascade in recurrent networks

A novel feature of the use of the cascade mechanism in the current model is its incorporation into a recurrent network. This is potentially interesting in the following respects:

If the cascade mechanism is used, as it has often been to date, in a purely feedforward network the effects of introducing temporal dynamics are limited to the single network response to a particular input pattern. In a recurrent network, however, the effects of altering the temporal flow of information through the network may have extended consequences across the set of responses which define the sequence to be produced.

Part of the problem in Jordan sequential networks implicated by the criticisms of Morasso and Sanguineti is not just that the temporal structure of the sequence is embodied in the training data, but that the processing of each subsequent response in a sequence is

---

[1] It has been shown that the network will produce the same asymptotic unit activations over a number of cycles with the cascade as it will in one pass through the network without the cascade (Cohen, Dunbar & McClelland, 1990)

initiated externally as part of the computer program that implements the model. Thus the decision about when to transfer from element 1 to element 2 of a sequence does not form part of the function of the model. The cascade mechanism can allow independent transfer between sequence elements to occur, simply by letting processing on a subsequent element to begin once the state units have reached asymptote. Progression through a given sequence then occurs autonomously as a function of the temporal dynamics of network processing.

One of the interesting features of the current model discussed in chapter five is that the complete network is composed of two semi-independent subsystems. Using the cascade mechanism, the flow of information between systems can be simulated with greater resolution and the effects of altering the timing relations between interacting subsystems can be investigated. This is of particular relevance in connection with modulatory neurotransmission mechanisms such as the dopamine system since the effects of damage are unlikely to affect merely the instantaneous enervation of target structures but also the time course of activation gradients. It has been suggested that the time course of the action of striatal dopamine is an integral part of the function of the system (Graybiel & Kimura, 1995).

In sum, two primary benefits of using the cascade mechanism in computational models of sequential motor programming can be identified:

- The introduction of temporal variability in the initiation of successive sequence elements and simulation of the time course of processing of the sequence, in both normal and lesioned conditions.

- A visualisation tool to assist in the analysis of system properties caused by the interaction of task-space attractors and disruptive manipulations to network processing.

6.1.4 Two modes of cascade processing sequential networks

It will become apparent in the sections that follow that there are two qualitatively different methods for incorporating the cascade mechanism into a recurrent network of the type considered here, which give rise to fundamentally different interpretations of its function. The network, before introduction of the cascade, is only partially recurrent since processing

of each sequence element stops at the state units. The first option with the cascade is to reproduce this assumption, with the exception that processing on the next sequence element only starts when the state units have reached asymptote. The asymptotic values on the state units are then held constant during processing of the subsequent sequence element. The second option is to use the cascade mechanism to simulate a fully recurrent network. To achieve this, it is necessary to allow activation to cycle continuously round the network. This is made possible by cascading activation through the state units as well.

In preliminary simulations it was demonstrated that in both cases, when the cascade rate (tau) is set to 1, the network behaves precisely as it would without the cascade. That is to say, when the network is presented with an appropriate pattern of activity on the plan units, the response of the network is to retrieve the associated sequence whereby each target location is retrieved at each successive time step. In the first case, when tau is then set to an intermediate value (between 1 and 0), activation builds up gradually on the articulatory units until eventually they reach an asymptotic level which corresponds to the learned activation values. The state units are then clamped with those values and processing on the next sequence element begins. This procedure yields an activation gradient for each of the articulatory units during processing of that pattern, and consequently can be used to calculate a 'virtual' trajectory for the endpoint of the arm during processing of that sequence element. Clearly, the optimal path of this virtual trajectory between two spatial points in task space, in the absence of any other influences, is a straight line. The behaviour of the network indicated by the path traced by the endpoint of the arm during activation build-up, in particular deviations from an optimal straight-line path, provide a means of visualising the influence of competing attractor states on the network under various conditions.

The use of the term 'virtual trajectory' above is deliberately intended to coincide with the same term developed in the context of the equilibrium point (EQ) hypothesis of trajectory formation and motor planning (Bizzi et al., 1992; Feldman, 1966) which is gaining increasing empirical support. Although the current model does not adhere precisely to the claims of the EQ hypothesis there are a number of points of contact between the underlying theoretical perspective which has motivated the EQ hypothesis and those which

control the operation of the model. In brief, the EQ hypothesis suggests that the motor system does not need to compute kinematic or dynamic derivatives in planning a trajectory as robotic controllers do (see Hollerbach & Atkeson, 1987; Kawato et al., 1987), but instead utilises a representation of the desired task space location of the endpoint of the active limb which in combination with the natural elastic properties of muscles acting about a joint establishes a virtual equilibrium point which causes the endpoint to be at rest at the desired location. Multiple such equilibrium points may be defined between the initial and final locations of the endpoint thus establishing a neurally encoded virtual trajectory. The activations of the muscles controlling the limb are then altered so that the endpoint traces a path following the virtual trajectory. Masson and Pailhous (1992) have suggested that the basal ganglia may perform the function of stiffness regulation using an EQ- based model.

The relationship between this theoretical position and the current model lies in the fact that the attractor states at target locations, in dynamical systems terms, are equilibrium points. The gradual change in activation of the articulatory units corresponds to the synergistic pattern of change in activity required in the muscles controlling the limb to trace the virtual trajectory to the desired endpoint location.

It is not clear in advance what effects will result from reducing tau below 1 in the case where the state units are cascaded as well. However, this manipulation is more likely to be interpreted as a model of dysfunctional processing in the network. This is due to the changing state representation at each cascade cycle. These issues are examined in the following section.

6.2 Using cascade processing in the analysis of attractor dynamics and trajectory formation

In the series of simulations presented in this section, the cascade mechanism was applied to the network used in the previous chapter. The aim of the studies was to examine the effects of both methods of using the cascade in a sequential network which have been discussed above. In particular we were interested in whether the addition of the cascade could provide

useful information in the visualisation and analysis of attractor dynamics in the behaviour of the network in both normal and lesioned conditions.

6.2.1 Simulation 6.1: Cascade processing in a sequential network

<u>Method</u>

In this first simulation we reproduced the task described in simulation 5.1. The network was initialised with the learned weights obtained in the condition with the rest configuration constraint depicted in figure 5.6. It will be recalled that in this network there are two state units which receive connections from the task units only. The cascade equations are only applied to the hidden and output units of the controller and forward model subnetworks. In this case the value on the state units remains static during the processing of each sequence element. Processing of each sequence element terminates when the forward model output units (i.e. the task units) reach asymptote. The cascade rate was arbitrarily set to an intermediate value of 0.5.

<u>Results</u>

The results of this manipulation are depicted in figure 6.1 below.
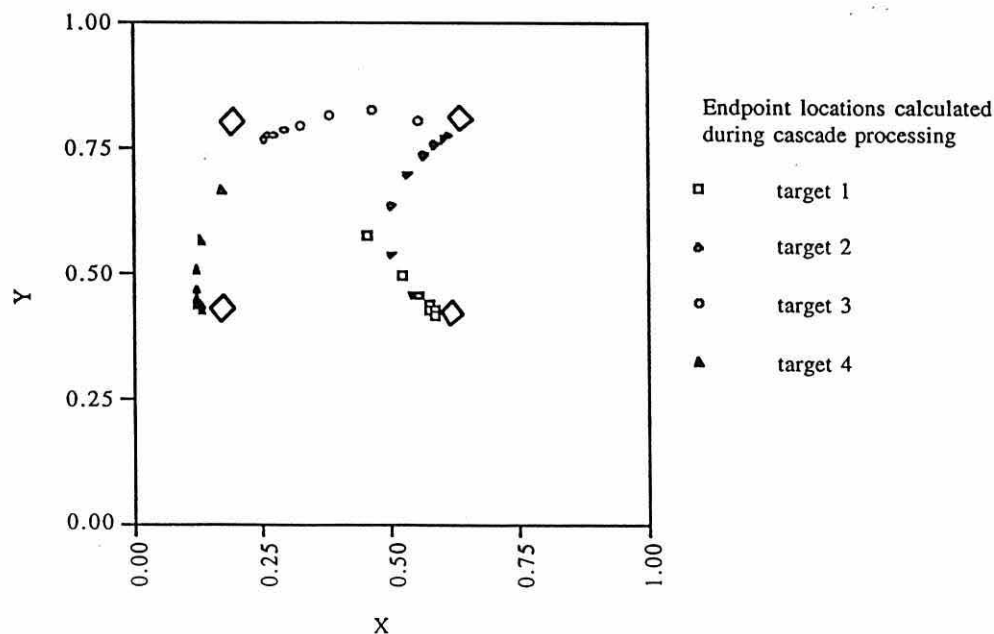
Figure 6.1. *Cascade processing in a learned 4-element sequence*

For simplicity, the representation of the kinematic configuration of the arm has been omitted in this, and all subsequent figures in this chapter. The data points represent the endpoint of the arm after each cascade cycle, calculated from the activations of the articulatory units. The forward model estimates of endpoint location are also omitted.

It can be seen that each sequence element transition takes approximately eight cascade cycles to complete. Comparison of figure 6.1 with figure 5.6 shows that the asymptotic activations of the units are approximately equivalent to those achieved without using the cascade. It is apparent that the gradual build-up of activation on the articulatory units caused by the cascade process forms a trajectory towards the target location. For transitions 1, 3 and 4 of the sequence, this trajectory does approximately follow the optimal straight-line path between the start and target locations. Transition 2, on the other hand, follows a curved path.

Discussion

This simulation demonstrates that the cascade process causes a change of activation over time which establishes a virtual trajectory of the limb endpoint from initial to final locations. This implies that the movement representations learned by the network are not confined to the target locations but are dynamic representations of the trajectories between locations. This feature of network behaviour is consistent with the basic ideas behind the equilibrium point hypothesis.

The reason for the curved path that occurs in transition 2 is not entirely clear, although the nature of the curve suggests that the activation profile is influenced by the initial position of the endpoint, which is the square data point at approximately 0.48, 0.58.

6.2.2 Simulation 6.2: Cascaded activations passed on to the state units

Method

In this simulation the network and task remained the same as in the previous simulation, except that activation at each cascade cycle was passed on to the state units. This method of processing means that the activation on the state units also changes at every cascade cycle.

In this regime a pattern of activity is simply clamped onto the plan units and activation is passed around the network until the output units reach an asymptotic level of activation.

When the cascade rate is set to 1, there is no change from the normal behaviour of the network, with activation at the output units eventually stabilising at the final position of the sequence. The aim of this simulation is to examine the effects of reducing tau below 1, in this condition. The network was thus run with four different cascade rates 1, 0.8, 0.4 and 0.1.

Results

The results are shown in figures 6.2(a - d).
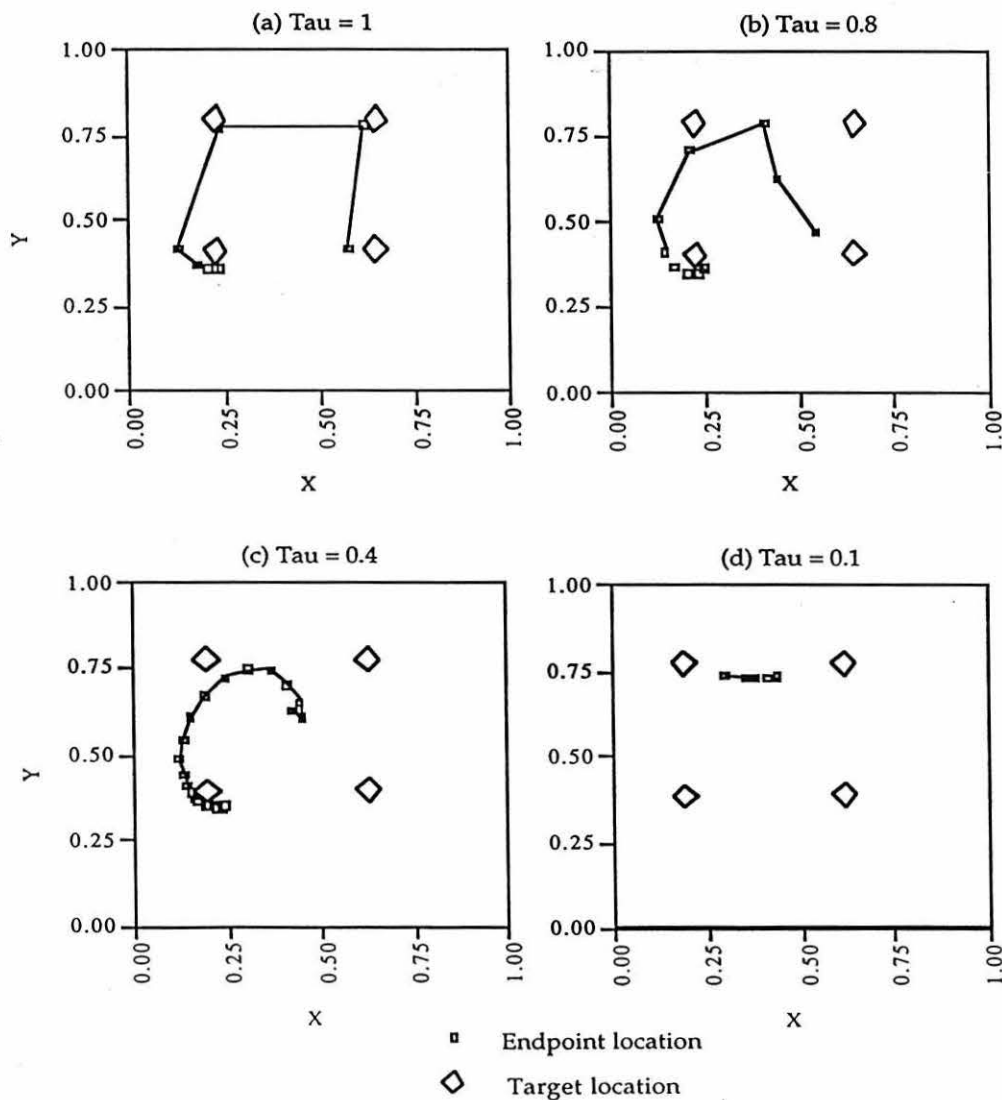
Figure 6.2. *Reducing tau when using cascade processing with dynamic state units*

Figure 6.2(a) illustrates the resulting endpoint trajectory when tau = 1. Because the state units receive their full activation values at the end of each cascade cycle, the endpoint is driven round the full sequence in four cycles, terminating at the final location. In figure 6.2(b), however, it can be seen that even a small reduction in the cascade rate causes disruption to the learned sequence. Intuitively, it may be expected that the disruption would be cumulative, causing increasing spatial errors as the sequence progresses. The general form of the trajectory, however, is consistent with the learned sequence and the location of the arm at the termination of processing is at the correct final location. The disruption is more severe when tau is reduced to 0.4, as depicted in figure 6.2(c). One feature of cascade processing is that the activations of the articulatory units do not move as far from their initial activations at each cycle, as the cascade rate is reduced. This means that after the first time slice in figure 6.2(c) the endpoint location is closest to the part of the learned trajectory between sequence elements 2 and 3. It is a consequence of the fact that the learned sequence acts as an attractor, that the network moves towards the closest point of the learned trajectory and then completes the sequence from there, completely omitting the first two elements. In each of the conditions above, the amount of time (no. of cycles) that the network is engaged in processing increases as tau is reduced, and the general form of the sequential trajectory is maintained to the final location. When tau is reduced below some critical value as shown in figure 6.2(d) with tau = 0.1, performance is catastrophically affected. In this case, processing terminates after only a few cycles. The reason for this is probably that at such a low cascade rate, the unit activations only change very slightly and artificially asymptote at some arbitrary point in the cycle. It will also be noticed that in this case the network is traversing a portion of the learned trajectory in reverse, (left to right).

Discussion

The results of this simulation indicate that when the cascade mechanism is used in conjunction with a recurrent network as described in this section a different interpretation is required. When the cascade rate (tau) is set to any value below 1, abnormal or incomplete production of the learned sequence results. In the previous chapter it was argued that a more global, systems level, manipulation was required to simulate dysfunction of the basal

ganglia in this computational model, in particular those which affect the processing function of the state units. A reduced cascade rate is a potential candidate for this role.

The cascade rate is a single global parameter which is usually assumed to reflect the time course of processing a response. In a purely feedforward network the cascade process does not directly affect the input-output relations of the model. In this model, however, each intermediate activation state at the task units causes a change in part of the input to the network at the following time slice. The pattern of activity on the plan and state units is then passed forwards to compute a new endpoint location. If the network is functioning normally this new position should match the target location in task space at that time slice. In the introduction to this section we argued that this process is akin to establishing a neurally specified equilibrium position for the endpoint of the arm. In this simulation it was shown that when tau = 1 the equilibrium positions lie at the target locations, but when tau is reduced the equilibrium positions fall short of target locations leading to multiple step movements of decreasing amplitude as tau is decreased. Multiple step movements and reduction of amplitude are symptoms of parkinsonian bradykinesia that were outlined as modelling criteria in chapter five.

Other features of the network behaviour when tau is reduced are a reduction in the overall shape of the trajectory (see figure 6.2(b)), and the omission of effortful sequence elements. This is seen in figure 6.2(c) where the network preferentially moves to the part of the learned trajectory closest to its initial position. These aspects of network behaviour are due to the interaction of impaired state unit function with the attractor dynamics of the learned trajectory. The presentation of a reduced cascade rate as a suitable model of motor dysfunction in PD remains only a tentative suggestion at this stage, but it does give rise to further motivations for studying the interaction of attractors formed by these networks and how they interact with state unit encoding of context.

6.2.3   *Attractor states and cascade processing in recurrent networks*

In the simulations described in this section, we examined two methods of incorporating the cascade mechanism into Jordan sequential networks. In the first method the state unit

representation remains static during processing of each sequence element. In the other method the state units vary at each cascade cycle as a function of the build up of activation on the task units. In either case, regardless of the method of updating the state units, an attractor state is regarded as one towards which the network will converge and at which the network will eventually asymptote.

In the first case where the input is constant, each target element of the sequence is an attracting equilibrium state at which the network asymptotes, after describing a trajectory towards that point in task space. In the introduction to this section we pointed to the close relationship between this form of processing and the equilibrium point hypothesis in motor control theory. In the second case, there is only one fixed-point attractor which lies at the termination point of the sequence, this being the point at which the articulatory unit activity asymptotes. Yet the learned trajectory towards this point is also attracting since the network will preferentially converge on the final endpoint location via the learned path, even when processing is disrupted by reduction of the cascade rate.

The differences between these two methods of processing hold important implications for the internal representation and control of action sequences. In the first case each individual state is explicitly represented as an equilibrium point, and a separate process is required to inhibit or reset the state representation to a new value which will drive the network on to the next state. In the second case, each intermediate target location is represented implicitly only as part of the learned attracting trajectory towards the final location, at which the network naturally terminates without the need for extra processes.

In the next section we consider a different task, which is an interesting special case both in terms of network dynamics and in motor control theory. This is the case in which the learned sequence forms a closed cycle of states.

6.3 Attractor dynamics and motor representation in connectionist networks

Jordan has highlighted the property that when a learned sequence is a cycle in task space, the resulting attractor is a *limit cycle* attractor (Jordan, 1986b; Jordan, 1990). This means that, instead of converging on a fixed point in task space as in the previous sequence we

studied, the network will continue to move round the learned sequence of states until terminated by an external condition. Strictly speaking, if a limit cycle attractor is present, the system will move towards the learned trajectory and then cycle, repeating each of the attracting states exactly on each cycle. If the system is subsequently perturbed slightly, it will return to the learned trajectory and continue as before.

Whereas the concept of fixed-point attractors has become important in computational theories of memory and representation (Hinton & Shallice, 1991; Plaut, 1991) the notion of periodic attractors and dynamical representations holds a special place in motor control theory (Kelso & Ding, 1993; Mpsitos, 1993; Saltzman & Kelso, 1987). In this section we report a series of simulations that further examine the nature of attractors in the current model. The results of these simulations show that the dynamics exhibited by this network have more complex properties than have previously been reported.

Until now we have given only an intuitive treatment of the concept of attractors in a dynamical system. To fully appreciate the results reported in this section, it is necessary first to underline in more detail some basic elements of dynamical systems theory. In the following sections we discuss both how it is relevant to the network under consideration, and how it is applied in both motor control theory and neurobiology.

6.3.1 Dynamical systems and recurrent neural networks

A dynamical system is a mathematical model which characterises how the output of a given system changes over time. In theory any process in which change is observed can be modelled as a dynamical system, although clearly some processes are more difficult to model than others. A useful representation of a dynamical system is a state-space model, which consists of a set of input variables, a set of state variables and a set of output variables. A dynamical system allows the modeller to determine how the output of the system changes over time as a function of the previous states of the system and the input to the system.

In general, any dynamical system can be described by a pair of equations which express how the state changes as a function of the current state s[n] and the input p[n], where the index n refers to a discrete time slice:

$$s[n+1] = f(s[n], p[n]) \qquad (6.3)$$

and an output equation:

$$x[n] = g(s[n]) \qquad (6.4)$$

which shows that the output of the system is a function of the current state of the system. Furthermore, by combining these two equations a composite equation may be obtained which maps the state and input onto the output at the next time step:

$$x[n+1] = h(s[n], p[n]) \qquad (6.5)$$

where h is the composition of f and g.

In order to develop a dynamical system, it is clear that the functions f and g must exist, that is to say that the state of the system at time [n +1] must not be completely independent of the state at time[n]. If it is possible to derive expressions for f, g and h, then it may be said that a solution to the dynamical system has been found. As in most mathematical modelling, the resulting dynamical system is much easier to analyse if the relationship between state, input and output is linear than if it is nonlinear. The state-space representation is useful, particularly in nonlinear dynamical systems where a general solution cannot be found, because it allows the modeller to plot the behaviour of the system over time as a trajectory in state space.

The sequential neural network which forms the basis of the model studied here can be characterised as a dynamical system. Unlike a purely feedforward network in which each state of the network is only related to the pattern of activity on the input units and is independent of any previous state of the network; the current state, and consequently the output of the sequential network is a function of previous states, due to the recurrent connections to the state (context) units[2]. Because all the hidden and output units have nonlinear activation functions, the dynamical system as a whole is nonlinear.

---

[2] Although it is easy to see why Jordan has used the term state units to describe the function of these units, it may at this point become a source of some confusion. The state unit activations do not constitute the state of the network. The state of the network is encoded by the pattern of activity of all units in the network at any time slice. We henceforth follow Elman (1990) who uses the term context units, which in

The order of the dynamical system depends on the value of the parameter μ which is the weighting parameter on the recurrent connections from the context units onto themselves. If $\mu = 0$, then the state at time [n] is a function of the state at only the previous time slice [n - 1] and the dynamical system is first order. If, however, $\mu > 0$ then the current state is a function of the n previous states and the dynamical system is of order n. A number of interesting properties arise from using the system with $\mu > 0$ and these are discussed in more detail in connection with the simulations reported below.

The state space model constructed here to describe the behaviour model uses the task space i.e. the space of possible endpoint locations of the arm to represent the space, and the calculated endpoint location of the arm at any time slice to represent the state.

A final point which is worth noting here relates to the introduction of the cascade mechanism. When the cascade is used each individual hidden and output unit can also be characterised as a dynamical system, since its activation state at each cascade cycle is dependent on its net input and its previous state. It would thus be possible, if so desired, to construct a state space representation of the activation trajectory of each of the hidden units in the network. This approach may yield interesting information about the dynamics of individual units, but has not been pursued here. A similar approach has recently been used in conjunction with extremely simple recurrent networks (Wiles & Elman, 1995).

6.3.2 Stability and bifurcation in a dynamical system

One of the goals of dynamical systems modelling is to determine the long term behaviour of systems, under various conditions. Two topics that are central to the study of this problem are stability and bifurcation theory. It is not our intention to discuss in depth the mathematics supporting these concepts. An excellent introduction to these subjects is provided by Sandefur, (1990). Our aim is merely to give sufficient information to enable us to discuss the behaviour of the network in the subsequent simulation studies.

The study of stability is an integral part of determining how the dynamics of a system unfolds over time. The state space of a linear dynamical system may contain one or

this model better describes their function and psychological interpretation.

more equilibrium positions or fixed points. Fixed points may be attracting or repelling. If the system is started close to an attracting fixed point it will move progressively closer to the equilibrium position and stabilise there. In contrast a repelling fixed point will push the system further away from the equilibrium position and the solution is said to be unstable.

A different form of behaviour often found in nonlinear systems, is when the solution goes, not to a fixed point in state space, but cycles repeatedly round a sequence of states. This type of solution is what we mean by a limit cycle attractor. Figure 6.3 reproduces graphs from Jordan (1986b) which demonstrates stable limit cycle behaviour in a basic sequential network.

In figure 6.3 the square represents the learned trajectory. When the network is started at a point either inside the learned trajectory (left panel) or outside (right panel), the network spirals towards the square, eventually reaching a limit at the square where it cycles repeatedly round the four points of the original sequence. It should be noted that in a simple limit-cycle such as this the system repeats the same four states with each cycle.
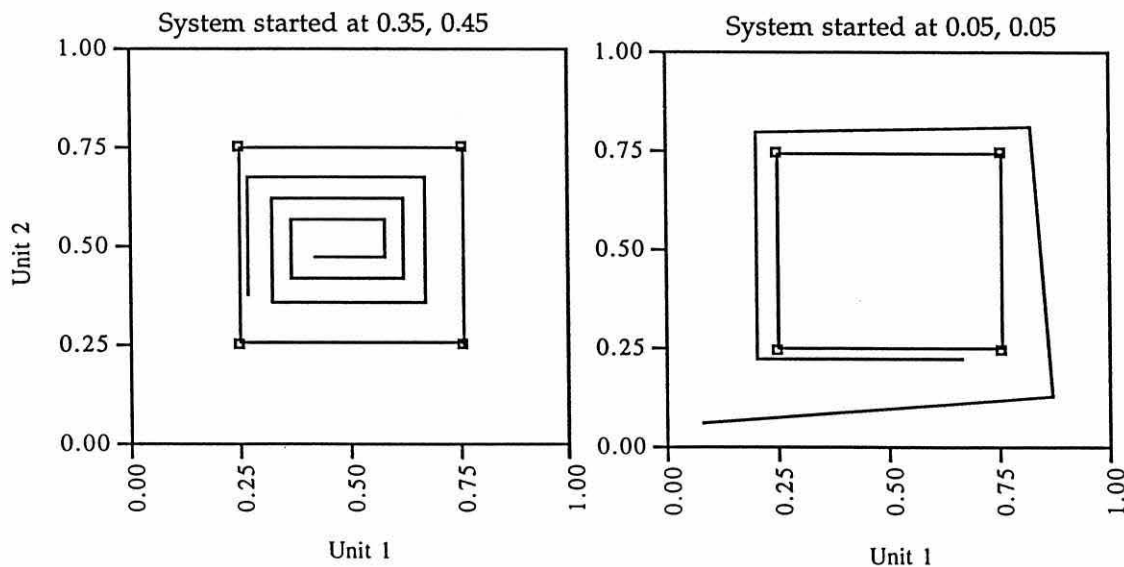
Figure 6.3. *A learned cyclic sequence acts as a limit-cycle attractor in the state space of the two output units of the network*

Nonlinear systems of several equations are capable of exhibiting more complex behaviour than that discussed so far. Complex behaviour may result from the existence of multiple limit-cycle and fixed point attractors in a system's state space or it may be governed by the

presence of another class of attractor - that is so-called chaotic attractors. The study of chaotic attractors is a rapidly expanding field and they are still not well understood. Here we restrict ourselves to discussing the behaviour of systems governed by an *attracting set*, (Sandefur, 1990). In short, the existence of an attracting set means that although we do know that the system will move towards the state-space trajectory defining the attracting set, unlike the simple limit-cycle attractor, it is not possible to predict which points on the trajectory will be intercepted at any particular time step. This depends on system variables including where in state space the system starts (dependence on initial conditions). The notion of an attracting set is best illustrated by an example. The same example will also allow us to introduce the idea of bifurcation in dynamical systems.

Let us assume a nonlinear dynamical system with the two dependent variables: a(n) and b(n). Let us further assume that the qualitative behaviour of this system for successive values of n is as depicted in figure 6.4(a)
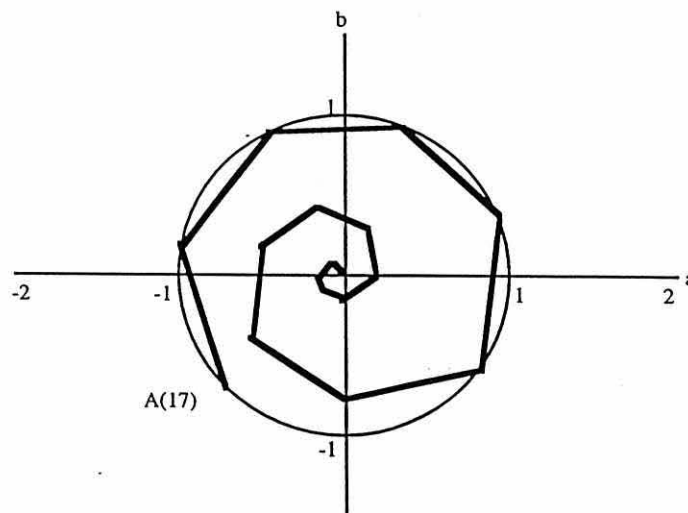
Figure 6.4(a). *Behaviour of a dynamical system for which the unit circle is an attracting set. The system is started at the repelling fixed point 0,0.(Adapted from Sandefur, 1990. p.400)*

Figure 6.4(a) shows the behaviour of the dynamical system for the first 17 time steps. The data points at each time step have been joined by straight lines. The system spirals away from its initial location (0,0) and approaches the unit circle given by: $a^2 + b^2 = 1$. Once it

intercepts the circle, the system stabilises and all subsequent points lie on the circle. This behaviour indicates that the unit circle is an attractor for the dynamical system.

In figure 6.4(b), the same system is shown after it has been left to cycle for many time steps and the line representing the unit circle has been removed. Because the system never repeats the same set of states exactly on different cycles, the data points after many cycles round the attracting set provides a trajectory map which characterises the underlying shape of the attracting set, (in this case the unit circle shown in figure 6.4(a)).

The properties exhibited by this nonlinear system are clearly different from stable limit cycle behaviour, in which the same point would be repeated on each cycle once the circle has been attained, and indicate a more complex form of attractor. However, they also provide a potentially valuable technique for visualising the characteristics of attractors in complex systems whose behaviour would be difficult or impossible to analyse by formal methods. A method for examining the properties of attractors in the sequential neural network model is developed in the simulations presented later in this chapter.
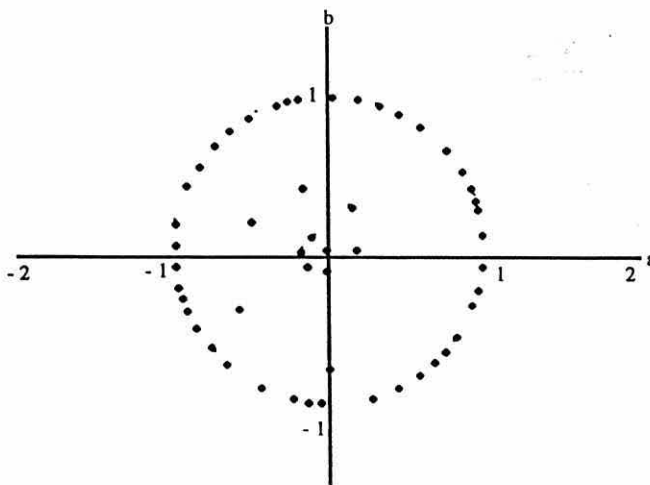
Figure 6.4(b). *Stable behaviour of the system in the limit. (Adapted from Sandefur, 1990. p.401)*

Let us now assume that the dynamical system involves an unknown parameter - x. A *bifurcation* occurs if, when the value of x is changed beyond a certain point, the qualitative behaviour of the system changes. There are many different types of bifurcation which can

occur, but we shall concentrate on the *Hopf bifurcation* which is common in nonlinear systems of several equations and is of particular interest here. The mechanics of the Hopf bifurcation may be briefly described thus: In the hypothetical system of figure 6.4 there is a constant c such that for x $\leq$ c, the fixed point at (0,0) is attracting, but for x > c the fixed point is repelling and an attracting set (the circle) bifurcates from the fixed point, and the qualitative behaviour described above is observed. However, the closer x is to c the smaller the radius of the circle and as x decreases to c, the radius goes to zero. A schematic depiction of the Hopf bifurcation is illustrated in figure 6.5 below:
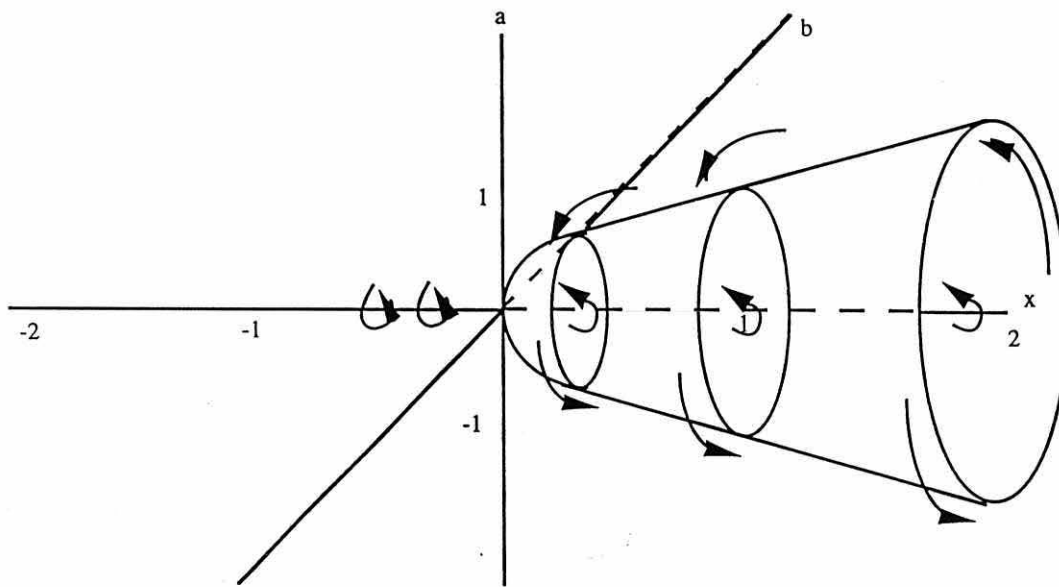
Figure 6.5.*The Hopf bifurcation. As x is increased above 0, the system spirals to solutions on a plane of the parabolic cone. (Adapted from Sandefur, 1990. p. 406. The parabolic cone is approximated in the figure)*

A full analysis of a nonlinear dynamical system would require that we observe the qualitative behaviour of the system for all values of x (within a definable range) and for all possible initial positions in state space. A more practical alternative in the first instance is to discover the behaviour for sampled values of x across the defined range, for a selected initial condition. This is the approach pursued in the following simulations.

A constant set of initial conditions are naturally provided by configuration of the simulated arm at the first position of the learned sequence. In the model there are two independent parameters which it is of interest to manipulate in place of the dummy variable

x: the cascade rate - tau, and μ which is the weight on the recurrent connections of the state units to themselves. Our reasons for examining the effects of varying these parameters on the overall qualitative behaviour of the network are that both are directly involved in the sequential processing function of the network and thus reflect our basic hypotheses concerning the processing deficit underlying motor disorders in basal ganglia dysfunction.

6.3.3 Dynamical systems and the neural control of movement

To date, few researchers have explicitly considered the application of dynamical systems to the study of biological movement control. There is, however, a growing body of research literature in this area which is closely in accord with the concepts involved. This includes the action systems approach to human movement control and, as mentioned previously, the equilibrium point hypothesis. Here, we aim to indicate the potential afforded by dynamical systems concepts in redefining current views of motor programming and in drawing together these various different approaches.

Jordan's work, which we have reviewed extensively, has linked the basic dynamical systems equations to the function of the connectionist networks he has developed in the feedforward control of movement. Mpsitos and co-workers, on the other hand have focussed on the possible development of chaotic attractors in the neural systems of simple invertebrates, which serve as the basis for representation of motor patterns. Saltzmann and Kelso (1987) have pointed to the development of limit-cycle behaviour in ensembles of oscillators in providing a representation of repetitive movements.

There are various issues raised in the work of these authors which lend interesting interpretations to the formation of attractors in the current model. Firstly, whilst a fixed point attractor may be useful for representing a static item of information such as an object name (see Plaut, 1991), dynamic attractors may be more useful for encoding movement information, which necessarily implies a trajectory through different states.

Moreover, it seems reasonable to suggest that the nature of the representation, reflects properties of the motor pattern to be produced or the state of the motor system. Thus, an attracting fixed point may represent a static state of the motor system, such as a

held posture or desired end point state of a movement. This is consistent with the basic ideas behind the equilibrium point hypothesis, and also makes sense of the observed motor equivalence in many goal oriented limb movements, i.e. it is irrelevant how the limb approaches its final state, the representation merely encodes the final state.

A stable limit-cycle attractor is suited to represent repetitive movements or sustained motor patterns. Repetitive movements are the type most often studied using the action systems approach, leading to suggestions that motor pattens may be encoded as limit-cycles (e.g. Kugler, Kelso & Turvey, 1982). In order to model the flexible, inherently variable adaptive motor programs characteristic of skilled human behaviour, it may be argued that a more complex and variable form of representation is required. In fact Mpsitos (1993) argues that this may even be a requirement of modelling the motor patterns of simple organisms. In this context, we may consider the formation of attracting sets as discussed in the previous section as a basic requirement of the system. The attracting set in this interpretation bears considerable resemblance to the psychological notion of a motor schema or generalised motor program (Schmidt, 1975; 1988). The general form of the pattern is retained, but the values of the dependent variables are variable. These sorts of attractors are intrinsically less stable and more variable than either fixed-point or limit-cycle attractors. This may have certain advantages in a behavioural sense, such as making it easier to blend two or more existing motor representations to achieve a new pattern of behaviour if a novel environmental situation or context is encountered.

The existence and behaviour of complex attractors in connectionist networks that produce movements or pseudo-movement sequences is consequently extremely interesting. A further line of support for the approach concerns bifurcations and the modelling of both normal and abnormal motor behaviour. There has been a long standing debate which has polarised motor researchers and is as yet unresolved over how motor programs deal with changes in independent variables such as the required speed of a movement. Schmidt, (1982) has argued that all components of a motor program change according to a proportional duration constant such that the relative timing amongst all components is maintained at different speeds. This argument is under many conditions well supported.

There is, however, a variety of contradictory evidence, reviewed by Gentner (1987) which demonstrates variation in the relative timing of components which is not consistent with a simple multiplicative rate parameter. Moreover many studies (e.g. Tuller & Kelso, 1990) have shown that motor patterns exhibit phase transitions or qualitative shifts in the overall pattern with changes in an independent variable. This has been shown in limb and speech movements. Both of these lines of evidence are consistent with the idea of a Hopf bifurcation in the attractor states which occurs when the independent parameter is sufficiently changed. The evidence cited by Schmidt would nevertheless agree with results over the range in which the motor pattern is structurally stable and no bifurcation occurs. We therefore argue that the Hopf bifurcation is a model which can account for the empirical findings supporting both positions in this controversy, and in particular can account for the phase transitions which occur in normal movements.

The same line of reasoning can be used in modelling the abnormal movements exhibited in PD and other basal ganglia motor disorders. In chapter five we have already argued in favour of modelling dopamine reduction in PD as a change in a system parameter which makes it more difficult for the system to select between actions, by reducing the contextual cues favouring one action over an alternative action. We have suggested that abnormal behaviour may be modelled as a bifurcation in the attractor states of the system, leading to production of an inappropriate response. For example, a number of parkinsonian deficits such as postural rigidity, lack of spontaneous movement and the failure to sustain motor activity in repetitive or sequential movement may be the result of the system collapsing to a fixed point stable state.

Figure 6.6 provides an illustration of how normal and parkinsonian motor activity may be modelled using the dynamical system depicted in figure 6.5. To make this example more concrete, we assume that the motor pattern is a cyclic motion of a limb which involves alternating patterns of activity amongst the agonist and antagonist muscles controlling the limb. The dependent variables a and b have been denoted as the agonist and antagonist respectively. The independent variable - x has been replaced by a global measure of system 'energy' or 'activation'.

In the top panel, normal motor activity is represented as sustained cyclic patterns of activity in the plane of the parabola of figure 6.5. The shape of the pattern is determined by the dynamical interaction of the opposing muscle groups. The size of the movement is determined by the value of the 'activation' parameter. Reduction of activation below zero (the intersection) yields a static posture.
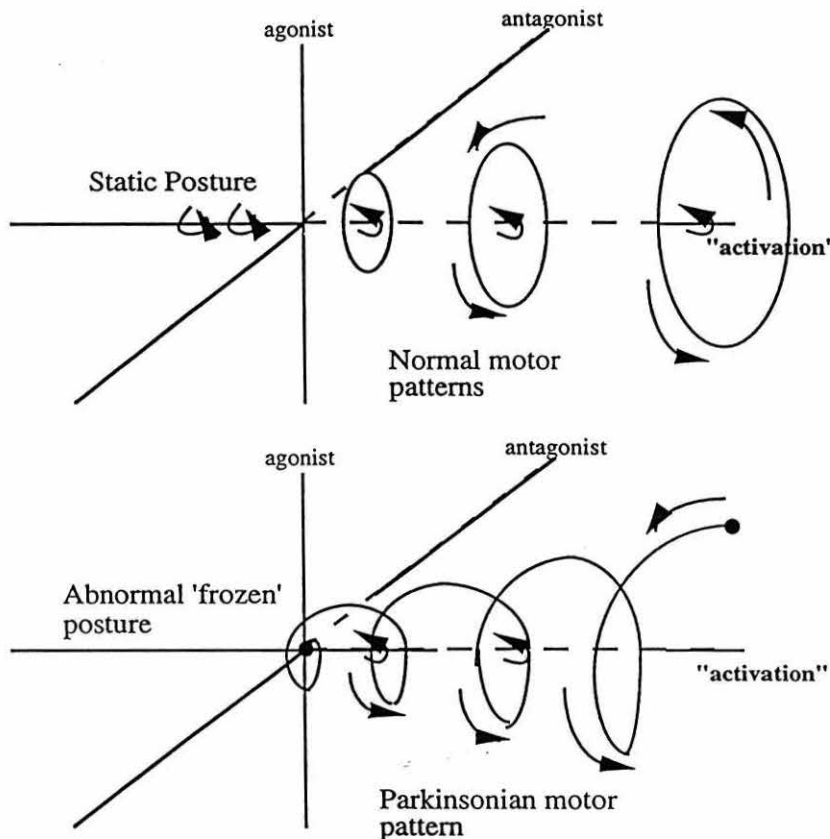
Figure 6.6. *Dynamical systems interpretation of normal and parkinsonian motor behaviour*

In the bottom panel a parkinsonian motor pattern is depicted. Here the system is started near the attracting set representing the large amplitude movement at the right of the top panel. Here, however, due to insufficient 'activation' the motor activity is not sustained and rapidly dwindles to the fixed point at the intersection. Here there is co-activation of the muscles which causes a rigid abnormal posture at this point. The use of agonist and antagonist should not be taken too literally here as the dependent variables could easily reflect some internal system variable such as reverberating circuits or coupled oscillators for coding motor variables. Similarly the uni-dimensional 'activation' parameter is

oversimplistic in terms of modelling basal ganglia dysfunction. The main aim here is to convey the main points of the general approach. The same framework has the potential to account for other abnormal movement states in a similar fashion, for example under different circumstances the behaviour of the system may include additional inappropriate elements, or may reduce to a stereotyped limit-cycle.

6.3.4 Aims and hypotheses of the simulations

The theoretical basis for the simulations presented in the remainder of this chapter is that the attractor built by the network over the learned sequential trajectory corresponds to a neural representation of the global properties of the sequential action, similar to a motor schema. As such it builds on conceptual ideas already in circulation (Borrett et al., 1992; Mpsitos, 1993; Wickens, 1993b). We assume that the properties of the attractor are linked to certain overall characteristics of the movement or state of the effector system. For example it makes intuitive sense that a fixed point attractor should reflect a held posture, or static end-state of the effector, whereas a dynamic attractor is required to reflect sustained activity of the effector, as in a sequential action.

Furthermore, we suggest that alterations of the properties of the attractor would be reflected in altered properties of the motor output. For example reduction in the size of an attracting set, may cause a reduction in the overall amplitude, force of each component of a complex action. In this way dynamic attractors provide a compact encoding of higher level properties of an action, but their nonlinearity means that the suggestion is not as simplistic as, say, a linear scalar such as multiplicative rate parameter, (Schmidt, 1982). Such schemes have been criticised (Gentner, 1987) because they do not account for nonlinearities in motor output such as phase transitions which are characteristic features of biological movement. Phase transitions (bifurcations) are, however, a natural feature of attractors in non-linear dynamical systems.

Lastly, we suggest that critical changes in the properties of an attractor may take place when key system variables are manipulated. In a complex representational space containing a number of attractors, changes in system parameters may lead to interference

which could cause severe and unpredictable effects on the behaviour of the system. We relate these computational ideas to the disorders of movement in basal ganglia dysfunction. In accord with the conceptual model of basal ganglia operation in chapter four, key parameters in the sequential neural network are those which affect the contextual representation encoded on the state units. We examine the behaviour of the system when these parameters are manipulated.

6.4 Examining the properties of attractors in cyclic sequences

In this section we report a series of simulations in which cyclic sequences are learned by the network.

6.4.1 Tasks and architecture

For consistency the same basic task is used when testing the network in all the simulations below. This task will be referred to from now on as "sequence 1".
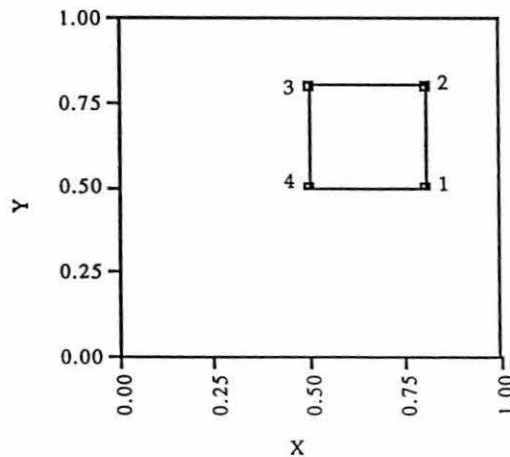
Figure 6.7. *Sequence 1 (condition A)*

This task is essentially the same as that used in the previous simulations, except that the transition from the final sequence element back to the first element is also learned, thus forming a closed cycle in task space. Figure 6.7 illustrates the location of targets and the desired trajectory of the limb endpoint for sequence 1.

In order to examine the properties of the attractor built by the network on this sequence and to test the hypotheses presented in the section above, three conditions were devised:

Condition A. In this condition only sequence 1 was learned by the network before testing. This condition is illustrated by figure 6.7

Condition B. In this condition two sequences are learned, one of which is sequence 1, the test sequence. The other sequence is qualitatively similar to sequence 1, but covers a different area of the task space. We refer to this sequence as the 'dummy' sequence because it serves as an exemplar of a class of other motor patterns which may affect performance of the test sequence - 'sequence 1'. In this condition, there is no point of intersection between sequence 1 and the dummy sequence. Figure 6.8 illustrates the layout of sequences in condition B.
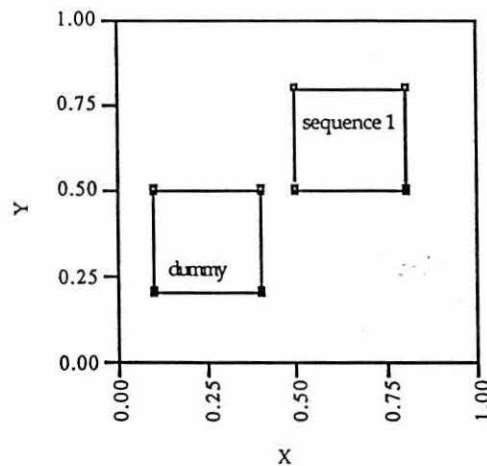
Figure 6.8. *Condition B*

Condition C. This condition is the same as condition B, except that the dummy sequence is shifted in task space so that there is considerable overlap between the two sequences. This condition is illustrated in figure 6.9 overleaf.
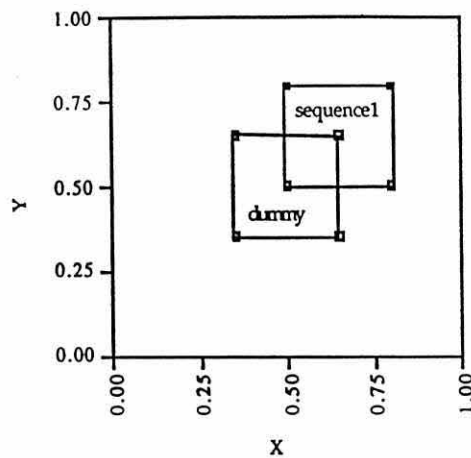
Figure 6.9. *Condition C*

The rationale behind this design is that it allows us to examine both the properties of the attractor built on sequence 1 and, later, the effects of network lesions on performance in three different test cases. Each test case reflects a different level of task complexity, with the level of potential interference depending on the existence and location of the dummy sequence.

In all cases the task was learned by a version of the network architecture in which there are recurrent connections from both the task units and the articulatory units to the state units. There are thus six state units instead of two. In all other respects the network remains the same as in previous simulations.

6.4.2 Methodology and interpretation of simulations

Our focus in this chapter is on the representation of sequential movements as whole actions. In chapter five we were primarily interested in the step by step transitions in performing one complete sequence. Here, our interest is at a more abstract level, and is concerned with global nature of the sequential action as indicated by the properties of the attractor which forms the network's representation of the action.

For these reasons the basic methodology in all of the simulation conditions reported below conforms to the following steps:

1. Learning. The task is presented to the combined controller and forward model network. The forward model weights have already been learned and are held fixed. Each sequence to be learned is identified by a unique pattern of activity across the three plan units of the controller subnetwork. Sequence 1 is, in each condition, denoted by the pattern 0, 1, 0; the dummy sequence is denoted by the orthogonal plan 1, 0, 1. Each learning epoch consists of presenting the network first with the plan for sequence 1, then, if appropriate the plan for the dummy sequence. Error derivatives and weight changes are calculated for each element of each sequence. As before, the network parameters during learning are: learning rate = 0.1 $\mu = 0.5$ and momentum = 0.9.

2. Testing. The weights in both the controller and the forward model subnetworks are held fixed while testing takes place. In order to test the network, recurrent cascade processing is switched on with tau = 1. The asymptote criterion for the cascade process is defined to be when the largest change in activation of any of the articulatory units is less than 0.01.

The use of the cascade in this way means that the network autonomously moves towards the attracting set thus defining the nature of the attractor without any external intervention. If the network moves towards a single equilibrium point, processing in the network will self-terminate as the cascade process asymptotes. This indicates a fixed point attractor. If the attractor is a limit cycle processing will not terminate as the network continues to cycle endlessly round the sequence of states defining the limit cycle. Furthermore, by plotting the location of the endpoint of the arm at each cascade cycle a trajectory map indicating the shape and location of the attractor can be produced.

6.4.3 Simulation 6.3: Cyclic sequences exhibit complex dynamics

In this first simulation using the above methodology we tested the network behaviour in condition A, in which only sequence 1 was learned.

Results

Figure 6.10 depicts the trajectory map of the network after 12 cascade cycles, that is three complete cycles round the four element sequence. Each point represents the location of the

endpoint of the limb and successive endpoint locations have been joined by a line to indicate the progression of network behaviour.
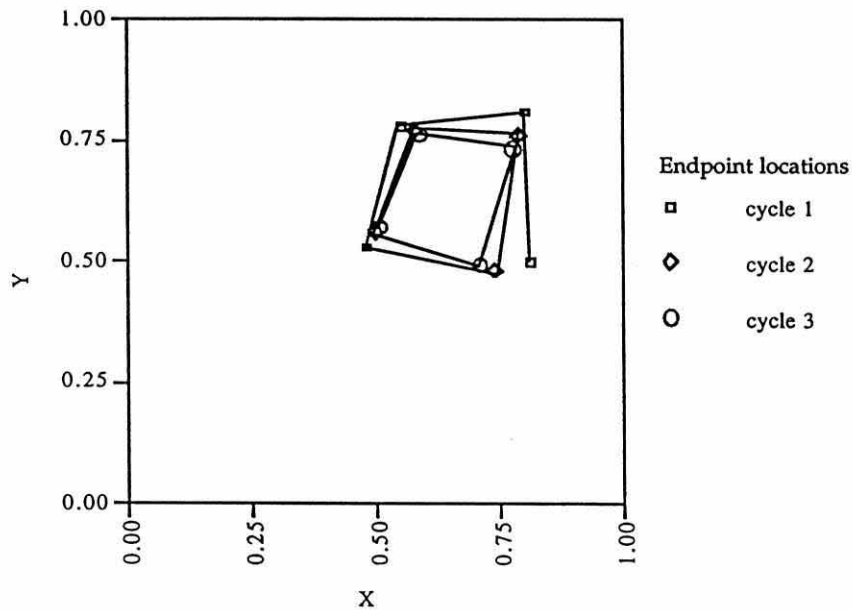
Figure 6.10. *State of the network after three complete cycles of the sequence*

Figure 6.10 shows that the network continues to cycle round the sequence, but that the behaviour is more complex than a simple limit-cycle on the learned sequence. The sequence as a whole rotates in a clockwise direction and reduce in size as the network continues to cycle. This behaviour can be seen more clearly in figure 6.11 which shows the progression over 36 cascade cycles, although only the first six cycles are joined up for the sake of clarity.
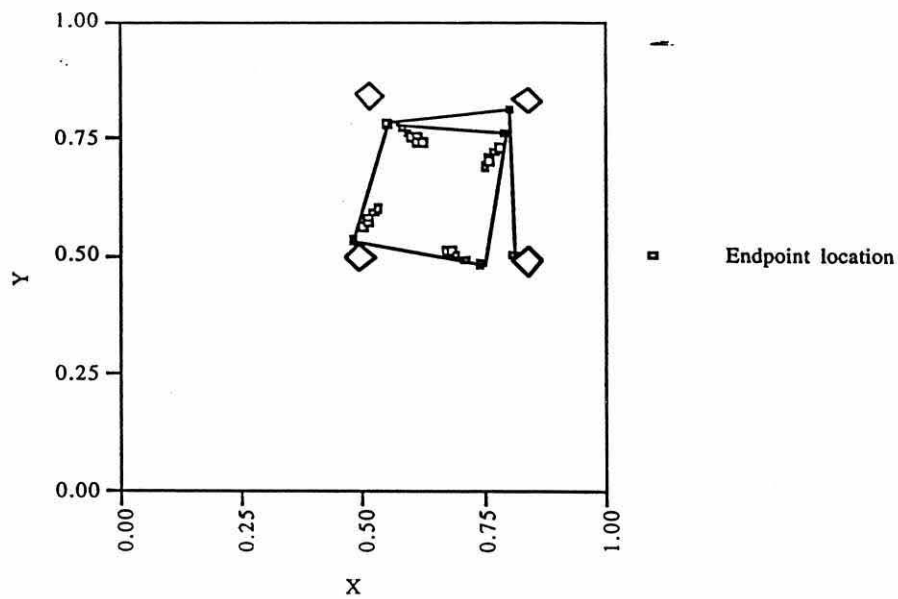
Figure 6.11. *Behaviour of the system over nine complete cycles of the sequence*

Figure 6.12 below illustrates the behaviour of the network in the limit. This figure indicates the attracting set formed by the learned network representation. Targets are not shown.
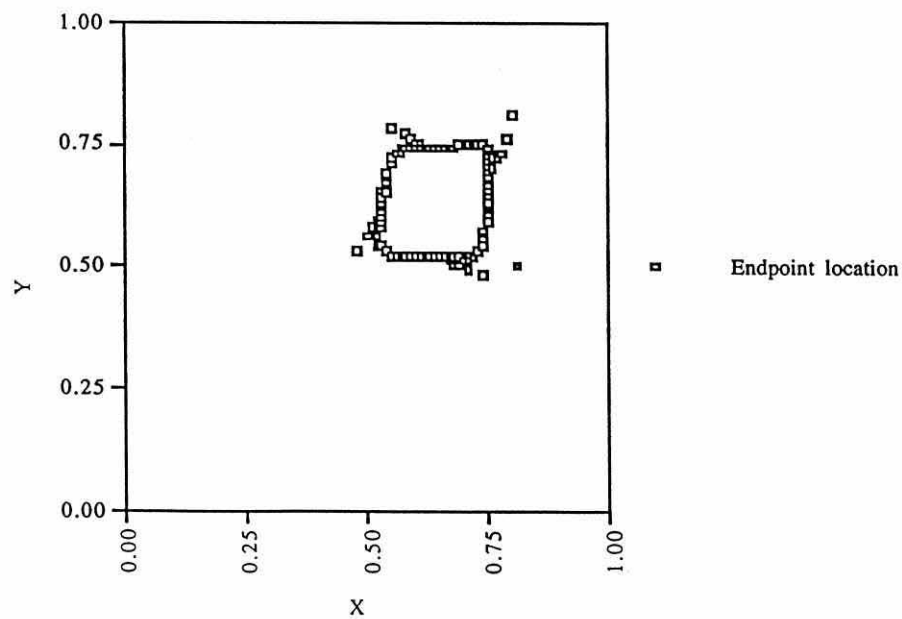
Figure 6.12. *The attracting set for sequence 1*

Figure 6.13 depicts the activations of the context units during testing of the network.
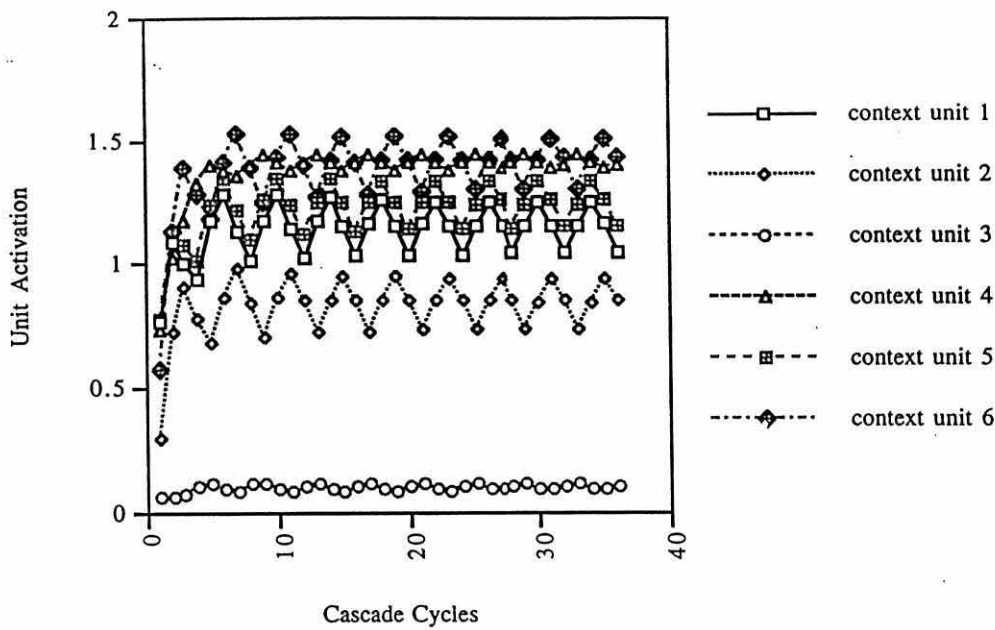
Figure 6.13. *Context unit activations in condition A*

Discussion

Although the behaviour of the network is clearly cyclic, the network in this instance did not follow the simple limit cycle trajectory reported by Jordan. The reason for this lies in the value of the parameter μ. In Jordan's simulation the value of μ was zero (Jordan, 1986b) presumably to provide the simplest possible demonstration of limit-cycle dynamics in the network, whereas in the current simulation μ was set to 0.5. This explains the rotational behaviour because the context vector does not merely represent the output vector, but encodes a more holistic representation of the whole sequence. Thus residual activation on the context units causes the activations of the articulatory units to be altered slightly each time a specific target location is revisited.

Close inspection of figure 6.13 reveals that although each individual context unit exhibits a four-cycle sequence of states, which corresponds to the learned trajectory, if any two cycles are compared for the same unit it can be seen that the corresponding elements have a slightly shifted relationship. Moreover, each of the units has a characteristic amplitude and phase.

These results hold some important implications for the analysis of the network behaviour as it allows visualisation of the entire attractor rather than a few selected points at the target locations. Examination of the context unit activations allows visualisation of the effect of manipulations to the contextual representation on the output behaviour of the network.

6.4.4 Simulation 6.4 Network behaviour when two sequences have been learned

In this simulation we tested the network in conditions B and C. In both conditions an additional sequence has been learned. The aim of this simulation was first to examine whether the existence of an additional 'action' affects the shape of the attractor formed over sequence 1. Secondly, it was to test whether proximity of the second sequence is a factor in determining the shape of the attractor for sequence 1.

The procedure followed in simulation 6.3 was repeated using the learned weights for condition B and then for condition C.

Results

Cyclic behaviour similar to that reported in the previous simulation for condition A was observed in both of the conditions tested. There were, however, differences in both the shape and size of the attractor for sequence 1 caused by the introduction and most notably the proximity of a second sequence. These effects are illustrated in figure 6.14 below.
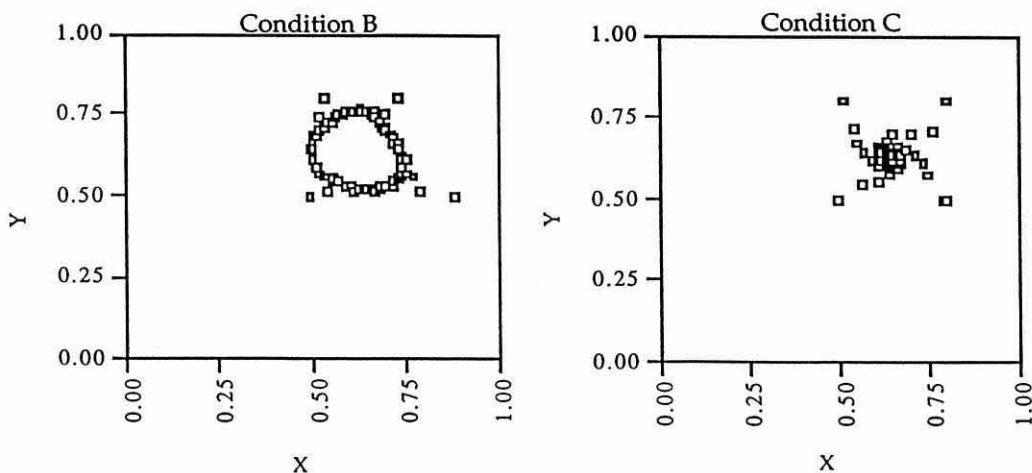
Figure 6.14. *Proximity effects of a second learned sequence*

In the representation of the endpoint trajectory for condition B, it can be seen that the presence of a second sequence in the weight space of the network causes some noise in the shape of the attractor for sequence 1. The previously square shaped trajectory, with distinctive corners at the locations of the targets and straight lines in-between, has been 'squashed' at the corners into less distinct circular form. The size of the attractor is much smaller in condition C where the two sequences overlap. These results demonstrate a clear proximity effect when a second similar sequence is also represented by the network.

Discussion

It must be assumed that in biological systems, several motor patterns are encoded simultaneously by the same representational system (modelled here as the learned weights in the control subnetwork). On the one hand this leads to useful generalisation properties, but at the same time can cause interference which is detrimental to performance. In many areas of cognitive performance including perception, memory and action, it has been shown that greater similarity leads to greater interference. Here, we have modelled this basic property using proximity as our measure of similarity.

Jordan, (1986b) showed that when one cyclic sequence was learned with $\mu = 0$, the network formed an attractor which corresponded precisely to the learned sequence. We have shown here that with $\mu = 0.5$, this result holds and that the endpoint trajectory between target states is a straight line, which is consistent with behavioural data showing that the optimal straight line endpoint path between two points is used in human limb movements (Flash & Henis, 1991). However, when a second sequence is also represented, the properties of the attractor for the test sequence (sequence 1) are altered, degrading the representation. The interference is increased as the similarity between the two sequences is increased.

6.4.5 Simulation 6.5: The effects of degrading the plan representation

In the previous simulation we examined interference in the output of the network, when the input unambiguously specified the target sequence. The input, in this model, is provided by the plan vector, where the name reflects the intuition that it represents an internally

generated motor plan. In this simulation we investigate the behaviour of the network when the plan is ambiguous.

Method

The network was again initialised first with the weights from condition B (non-overlapping sequences) and then the weights from condition C (overlapping sequences). This time, however, the plan units were initialised with values equidistant between the plans for each of the two learned sequences, (the plan vectors were orthogonal). Thus each of the plan units was initialised with a value of 0.5.

Results

The results are illustrated in figure 6.15 below.
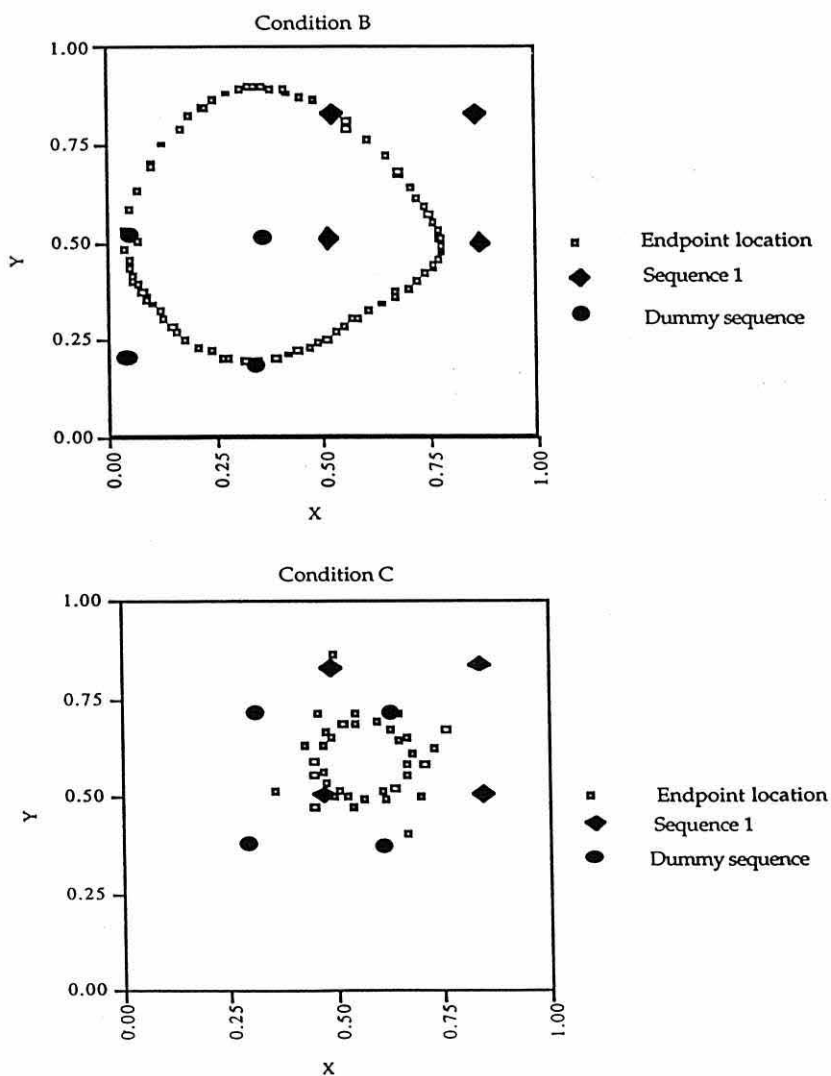
Figure 6.15. *Network behaviour with a plan vector which is equidistant between the two learned plans*

In the trajectory map for condition B, at the top of the figure, the rim of the attracting set has expanded to accommodate approximately equal portions of the two sequences. In condition C, by contrast, the attractor is restricted to the area defined by the intersection of the two sequences.

Discussion

The results presented above demonstrate two important properties of the network, by analysis of the attractor topology. Firstly, they show that there is continuity in the mapping between input or plan states and the output states for whole actions. In both cases the network generalises to produce an attracting set which covers an area, which is in some sense, midway between the two learned sequences. Secondly, there is a qualitative difference between the overlapping sequence (condition C) and the non-overlapping sequence (condition B). In condition C, the attractor does not include half the area of each of the learned sequences, but covers the intersection area instead.

6.4.6 Implications for motor representation and action selection

In terms of motor planning and the selection of action, both of the above results give good computational grounds for proposing that an extra mechanism exists to suppress competing representations and enhance the desired representation during response preparation.

The proposed basal ganglia inhibitory / disinhibitory mechanism does not form an explicit component of the model, as discussed in chapter five, but forms an integral part of the operation of the network due to the use of the logistic activation function in the processing of each unit in the network. A limitation of this network model in terms of the conceptual model of chapter four is that the notion of contextual control cannot be treated separately from the mechanisms of inhibition and competition that we are assuming are implemented by the basal ganglia. The only meaningful disruption that can be explored in this model is disruption to the control signal itself. These problems are addressed in chapter nine.

When the model is functioning normally as in simulation 6.4, the effect of a second sequence on the form of the attractor does not affect reasonably accurate retrieval of the

target sequence, when only one cycle round the sequence is considered, but does affect the overall robustness of the representation. Simulation 6.5 shows that the network exhibits useful generalisation properties in the structure of the plan to output mapping, which suggests that features of already learned actions can be readily extracted and combined in the context of a novel plan. However, it also suggests that if the plan is degraded in some way, spurious attractors may develop, leading to abnormal production of the intended action.

In total these results suggest that if there is a system level dysfunction which affects the topology of the attractor for a given sequence, the existence of nearby attractors will yield unpredictable interference effects which impair the sequential behaviour of the network. This issue is considered in the next section.

6.5 Examining the effects of system disruptions on network behaviour and attractor dynamics

In chapter five we considered various methods of simulating the effects of dopamine depletion on the output behaviour of the model. In particular we examined a gain reduction model derived from the work of Cohen and Servan-Schreiber. Reducing the gain of the output function of a single unit impairs the ability of the unit to discriminate between different patterns. Following those simulations we concluded, for various reasons, that this manipulation is at too low a level to be meaningful in the current model and that a more global systems level parameter, which directly affects the contextual encoding, on the state / context units should be considered. Our reasons for this are twofold and are rooted in the relationship between the connectionist model used here and the conceptual model of basal ganglia operation discussed in chapter four. First, the intact operation of the context vector is crucial to the normal sequential operation of the network, and any disruption of the forward model only has an effect on network behaviour in so far as the context vector is disrupted. Second, that the internal representation of context is precisely what we have suggested is disrupted in fronto-striatal interaction.

In this chapter we have highlighted two system parameters which directly affect the processing of internal context in the model: the decaying activation trace - μ, and the rate parameter - tau. We consider each of these parameters in turn.

The parameter μ establishes a form of short-term memory of past states of the network, which is what allows us to think of the state units as encoding a contextual representation (see also Elman, 1990). The size of μ affects the "depth" of context that is encoded. If $\mu = 0$ the system is sensitive only to the previous state of the network. As μ is increased the system becomes more sensitive to temporally remote contexts. This ability is important for learning extended sequences, especially those containing repeated elements (Jordan, 1986a).

Tau, as we have previously stated, is a rate parameter which in a sequential network of this type has two different interpretations. If the context units are not cascaded, it determines the speed of unit update in the network. This version has no direct impact on the context unit activations. Alternatively, if the context units are cascaded, as in the simulations above, then a reduction in tau, does affect the context representation. In this case, a reduction in tau below 1 reflects a reduction in system "energy" which causes successive states of the context vector to become closer to each other. This as we saw in earlier simulations can impair sequential performance in the network.

In the following series of simulations we investigate the effects of changing both μ and tau under a variety of different conditions.

6.5.1 Simulation 6.6: The effects of altering the value of μ

Method

The parameter μ, which controls the depth of sequential information encoded by the context units has a range of meaningful values between 0 and 1. If it is set to 0 only the output at the previous time slice is encoded. Alternatively if it is set to 1 then the whole of the activation at each time slice is included. In this simulation we examined the effects both of altering the value of μ through the full range of values within its natural range. This operation was first performed for condition A only in the first instance.

Results

Figure 6.16 shows trajectory maps of the endpoint locations after 27 cascade cycles and in the limit respectively.
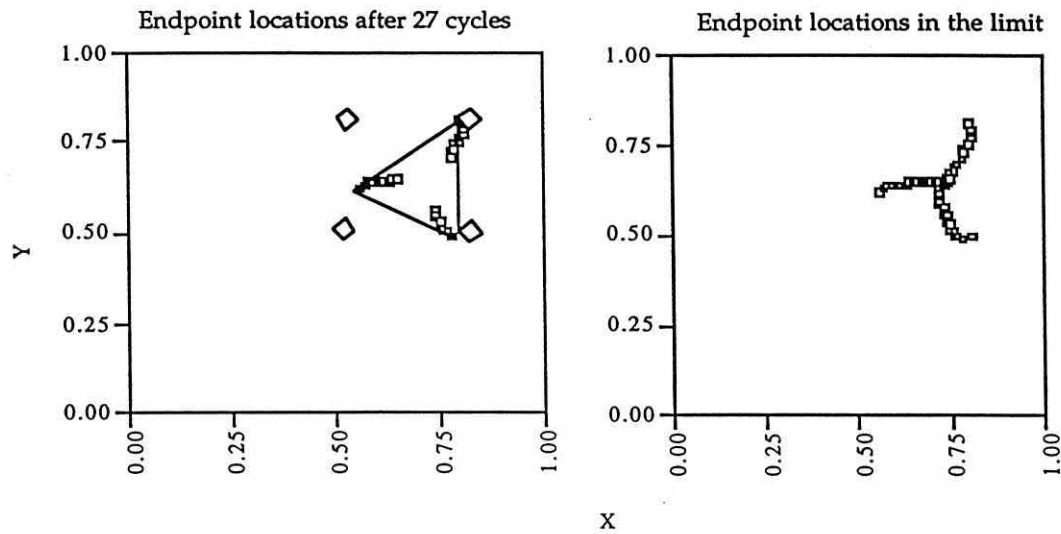
Figure 6.16. *Network behaviour with* $\mu = 0$ *in condition A*

The change in network behaviour can clearly be seen in both figures. The rotation is predictably now not in evidence, because there is no residual activation on the state units. It can also be seen that, in the limit, the system collapses to a fixed point attractor at the centre of the region, previously occupied by the cyclic attractor. This indicates a bifurcation which occurs at some point as μ is reduced to zero, such that a fixed point at the centre of the target sequence becomes attracting. Interestingly, there is also a phase shift from the four element pattern of the original sequence, to a three element pattern in which the second element is interposed between the original third and fourth elements.
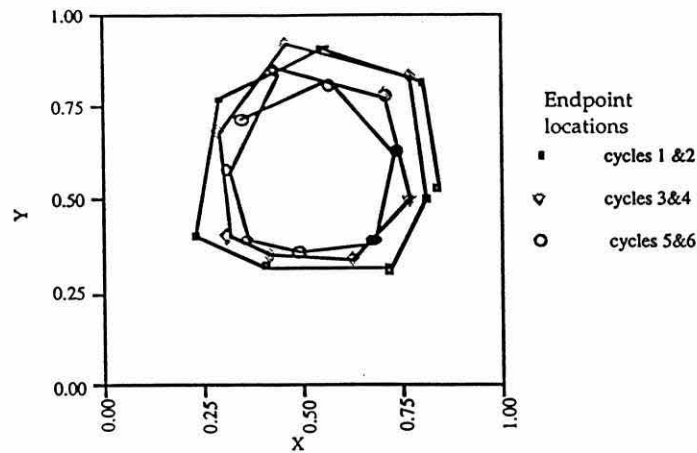
Figure 6.17. *Endpoint locations after 6 cycles with* $\mu = 1$

Figure 6.17 illustrates the behaviour that results when μ is increased to one.
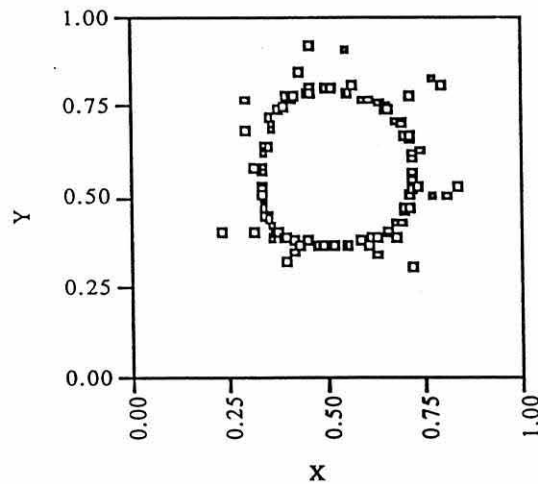
Figure 6.18. *Network behaviour in the limit with* $\mu = 1$.

In figure 6.17 successive endpoints are linked to show the first six cycles around the learned sequence. The qualitative behaviour of the network is the same as in the original case, i.e. a cyclic sequence of states. Here, however, the sequence is a seven-cycle in task space, as opposed to a four cycle, and the size of the rotation has increased.

Figure 6.18 shows that in the limit the network converges on an attracting set indicated by the circular pattern formed by the endpoint locations after the first few cycles. This attracting set is similar to the original set depicted in figure 6.11, although it has a

greater area and its centre is shifted on the x-axis. The greater area of this attractor than that produced with $\mu = 0.5$ is indicative of a roughly cone shaped attracting set which increases in diameter with an increase in μ. This progression is confirmed when we examine the changes in attractor shape across a range of values of μ. These changes are illustrated in figure 6.19.
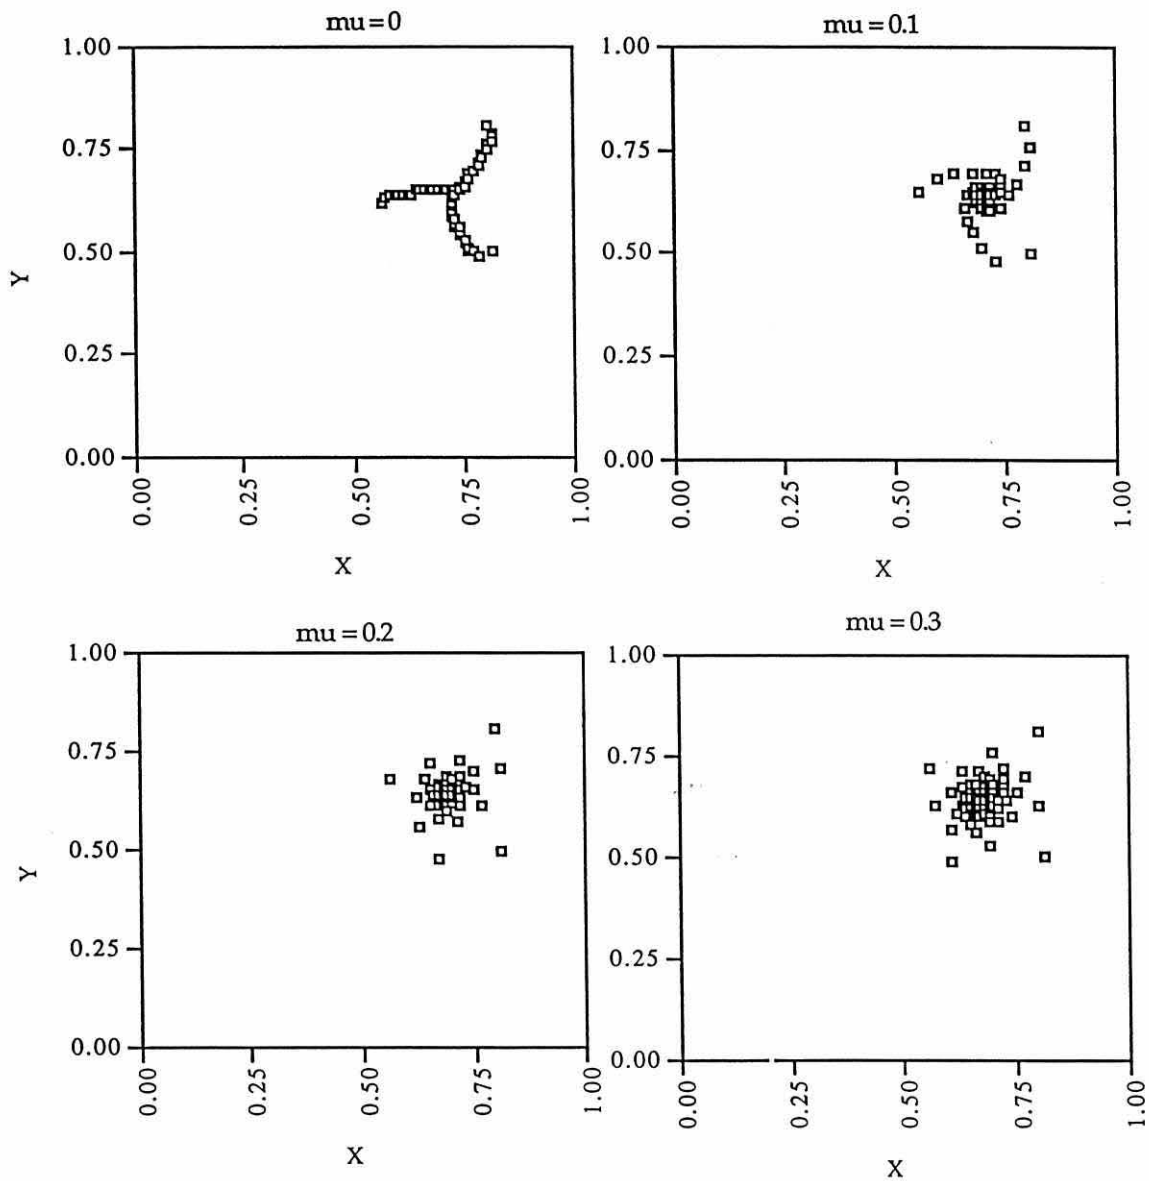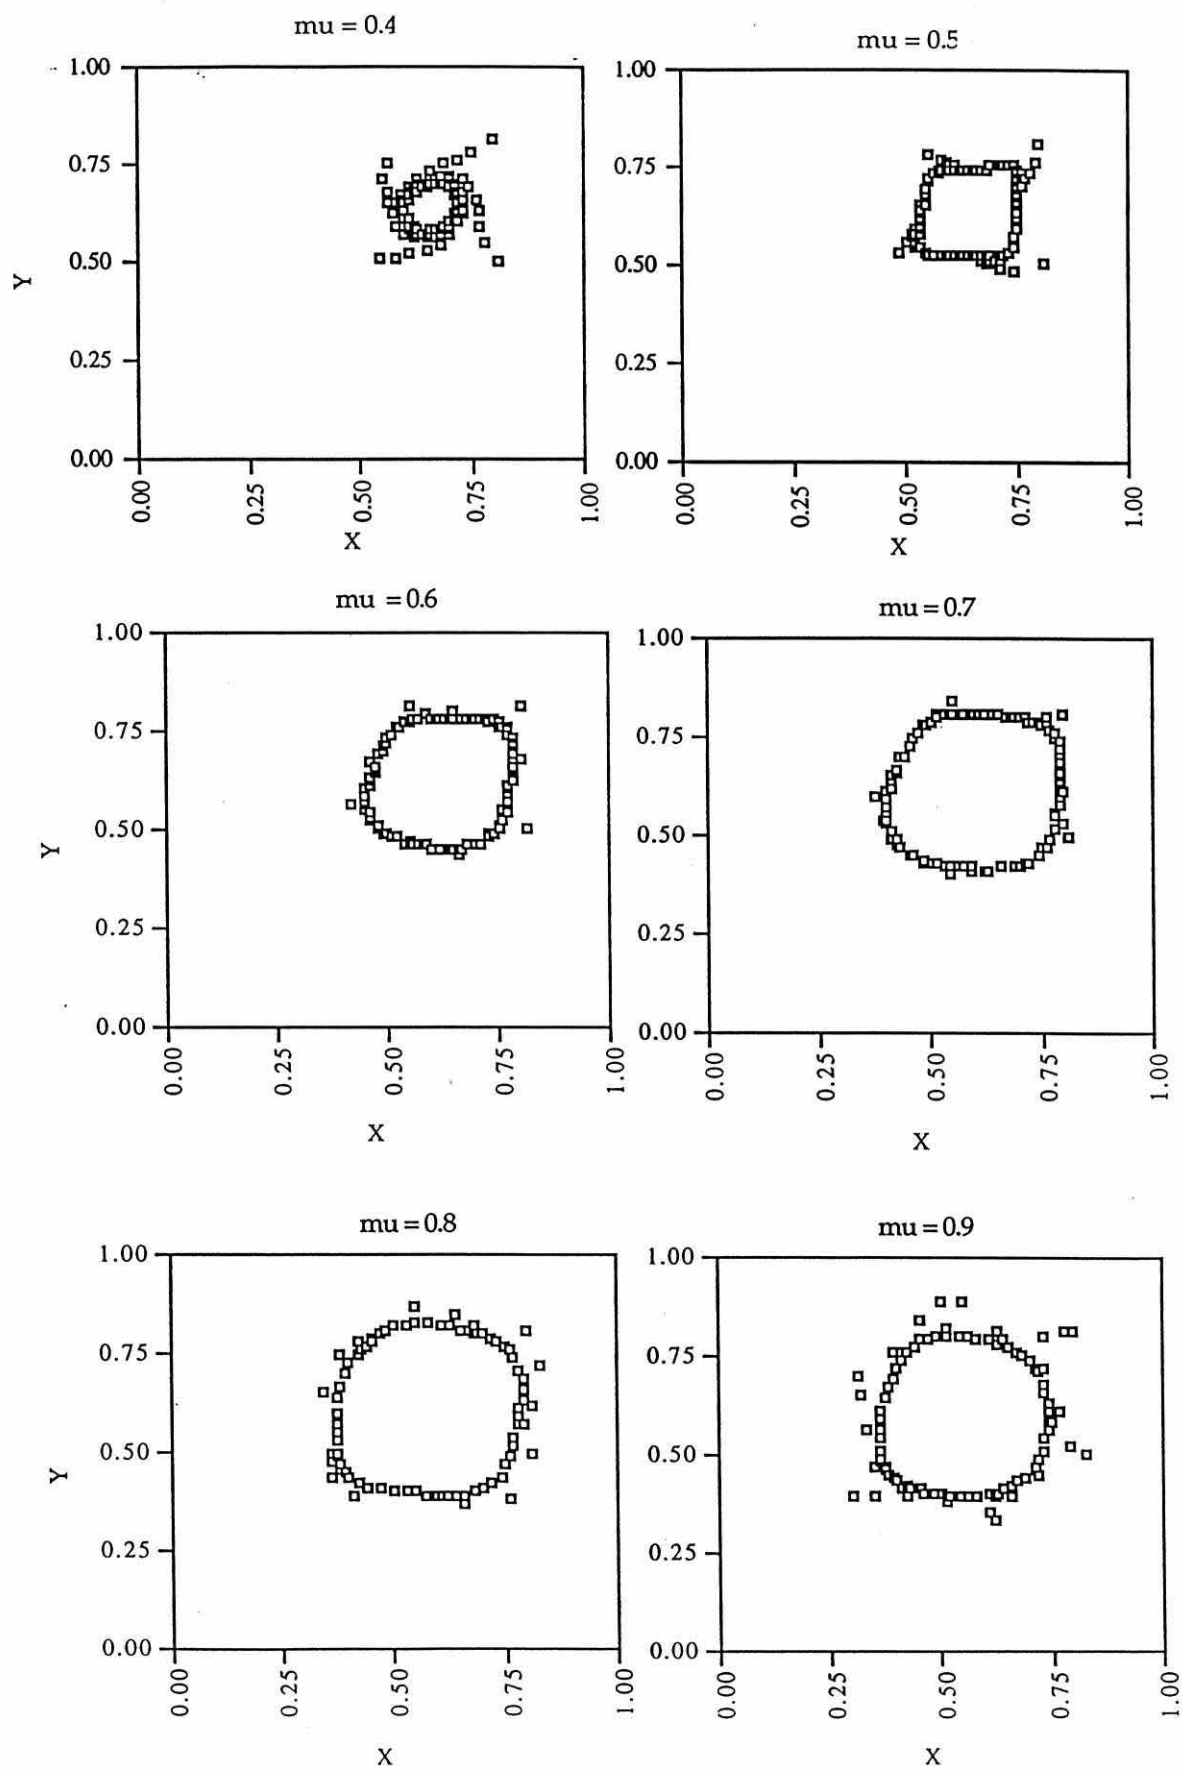
Figure 6.19. *Varying μ throughout its range in condition A*

The series of trajectory maps depicted in figure 6.19 reveals a number of features of the change in system behaviour as μ is altered throughout its range of values. These are enumerated below:

1. Bifurcation of the attractor from a fixed-point to a cycle at some critical point between μ = 0.4 and μ = 0.

2. Increase in diameter of the rim of the attractor as μ increases.

3. Discrete phase transitions in the number of cycles to perform one complete sequence for different values of μ. Compare μ = 0 and μ = 0.5, for example.

4. Reversal in direction of rotation (clockwise / anti-clockwise) as the system passes through critical phases. To see this compare μ = 0.1, μ = 0.4 and μ = 0.5.

5. This last point entails that some value, or range of values exist between the values for which a change in direction is observed, at which stable behaviour of the attractor is evident. For example between μ = 0.4 and μ = 0.5 there should exist a value at which the attractor is a stable 4-cycle. Further simulation showed that μ = 0.48 is such a value. This is illustrated in figure 6.20 below:
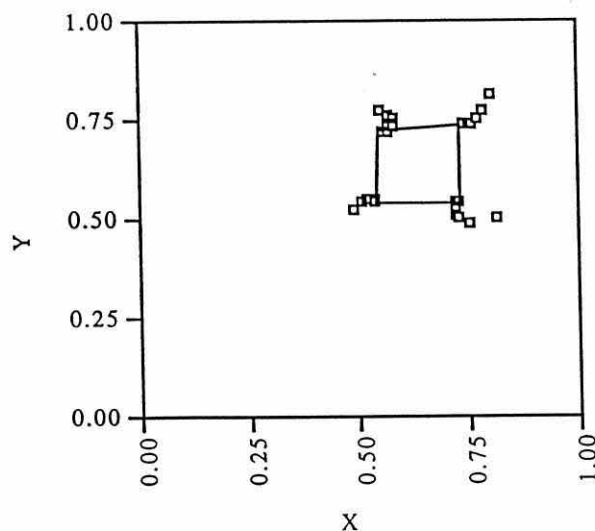
Figure 6.20. *Stable limit cycle behaviour with μ = 0.48.*

Discussion

These results do not constitute an exhaustive analysis of the system, but they do indicate that the Hopf bifurcation presented above describes the behaviour of this system as the μ parameter is reduced to zero sufficiently well to warrant its use as a working model. Whether or not this is entirely accurate, there is a strong indication that when a cyclic sequence is learned with μ greater than zero, the attractor for the system is a more complex attracting set based on the trajectory of the sequence rather than a simple limit-cycle. This property considerably extends the potential of this type of sequential network in modelling action representations.

The finding of phase transitions in the period of the cycle is also interesting in terms of dynamical approaches to motor control generally (Tuller & Kelso, 1990), although it is not clear how to interpret this finding in the current model. The rotational component of each cycle round the sequence which increases in size as μ increases is a direct consequence of the value of the parameter. Thus we do not attempt to offer any parallel interpretation in terms of biological systems, although it is a system characteristic which should be noted and accounted for when considering the properties of this type of contextual representation.

In sum the results of this simulation suggest that when only one cyclic sequence is encoded by the network weights, and μ is altered through the range of its values the behaviour of the attractor formed by the system appears to conform to the basic Hopf bifurcation model presented earlier in this chapter. The implications of this behaviour for motor control and motor disorders have been discussed in an earlier section. Thus it is possible to regard μ as a system parameter which can be varied to modulate the output characteristics of a learned motor pattern. At some critical value of μ between 0.3 and 0.4 the system bifurcates to a fixed point attractor which is assumed to represent a fixed posture. In the next simulation we examine how the behaviour of the network changes as μ is altered in the more realistic situation that multiple motor patterns are co-represented in the state space of the system. In particular we are interested in the patterns of interference that occur when a second sequence is non-overlapping (condition B), and when it is overlapping (condition C).

6.5.2 Simulation 6.7 The effects of altering μ when more than one sequence has been learned

Method

Here the previous simulation was replicated for conditions B and C.

Results

Figure 6.21, on the following two pages, depicts trajectory maps for varying values of μ in conditions B and C respectively. The figure depicts condition B in the left hand panels and condition C in the right hand panels. The appropriate value of μ heads each pair of panels. The black circles in the figure indicate the location of targets for the dummy sequence in each condition. In some cases lines have been included to clarify the trajectory of the endpoint.

In both conditions B and C depicted in figure 6.21, the change of the topography of the attractor for sequence 1 as μ is varied exhibits differences from condition A. The information gained from this series of trajectory maps regarding the interference effects of a second attractor in the state space could not have been gained from looking at the learned value of μ (0.5) alone. In particular, qualitative differences between conditions B and C can clearly be seen. In condition B, with values of μ less than 0.5 the system is driven very quickly to a fixed point at approximately 0.75, 0.75. In condition C, by contrast, at low values of μ (<0.5), a cyclic tendency is evident even at $\mu = 0$, and at $\mu = 0.3$ a sizeable attracting set has developed. However, at values of 0.5 and above an inverse pattern emerges, whereby the attractor condition B expands dramatically, whereas in condition C, the system is driven to a fixed point near 0.8, 0.5.
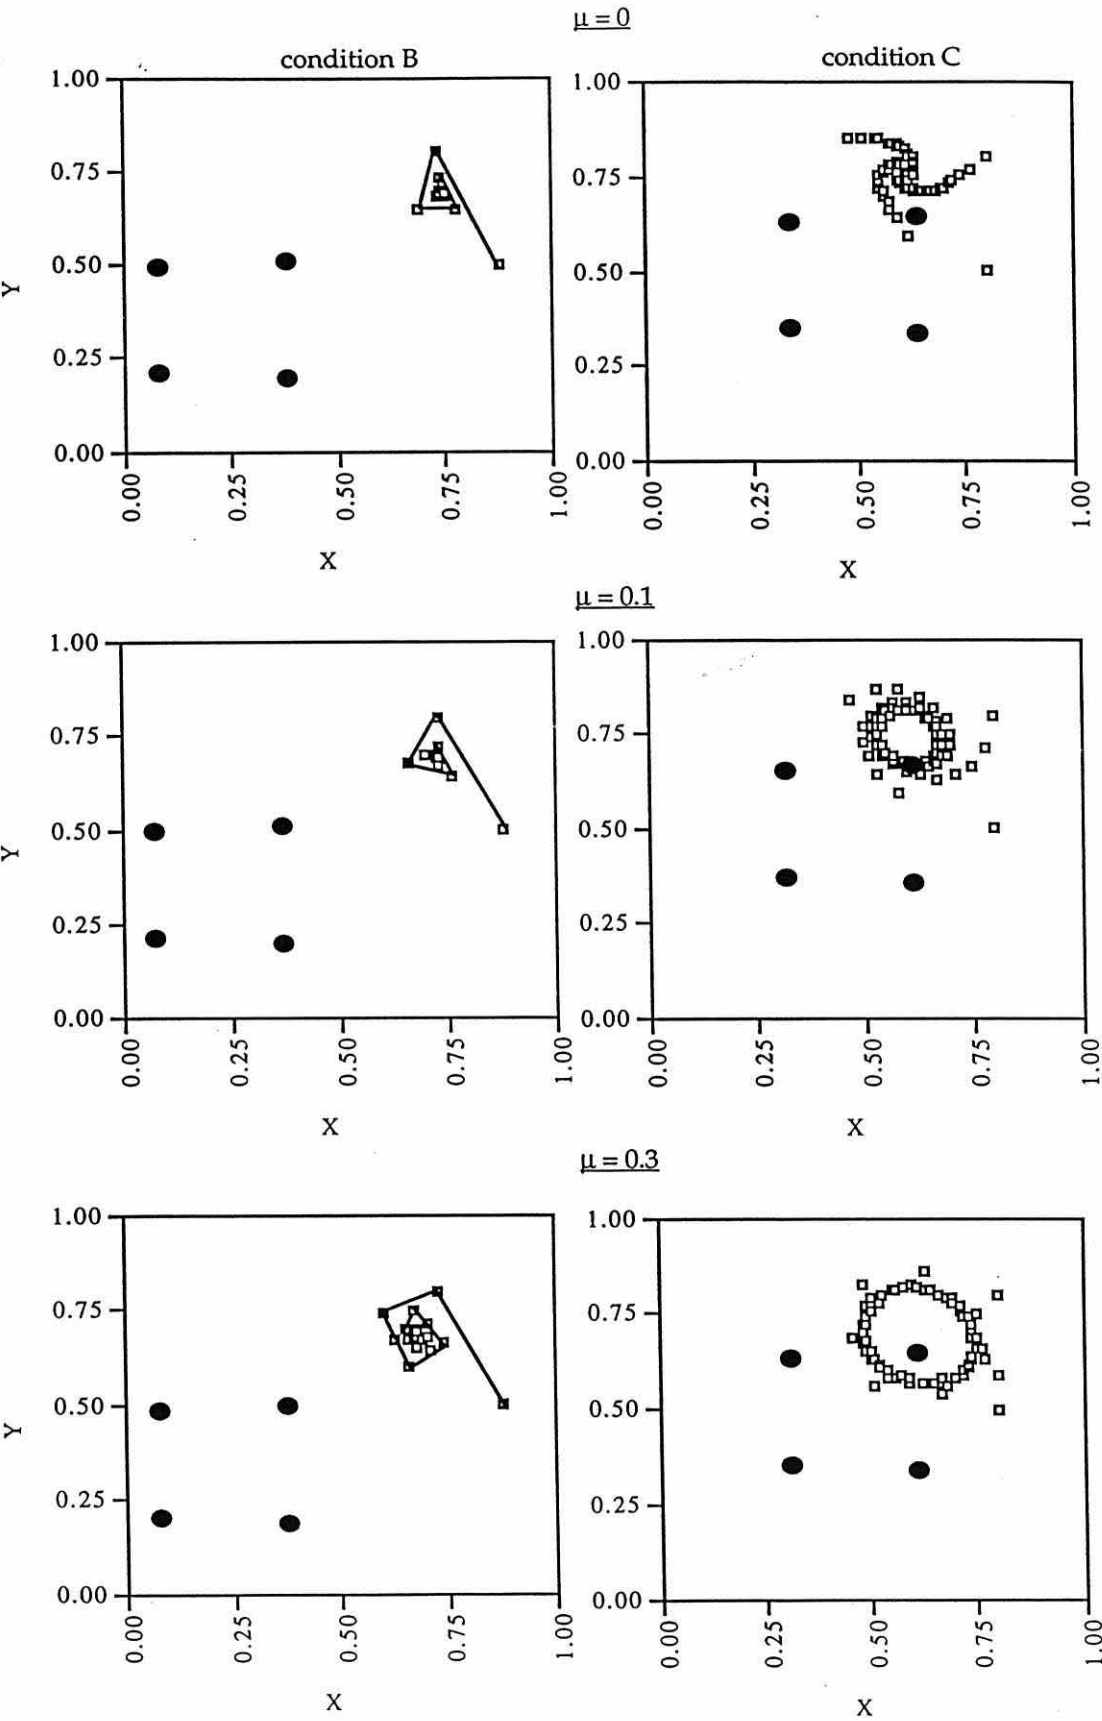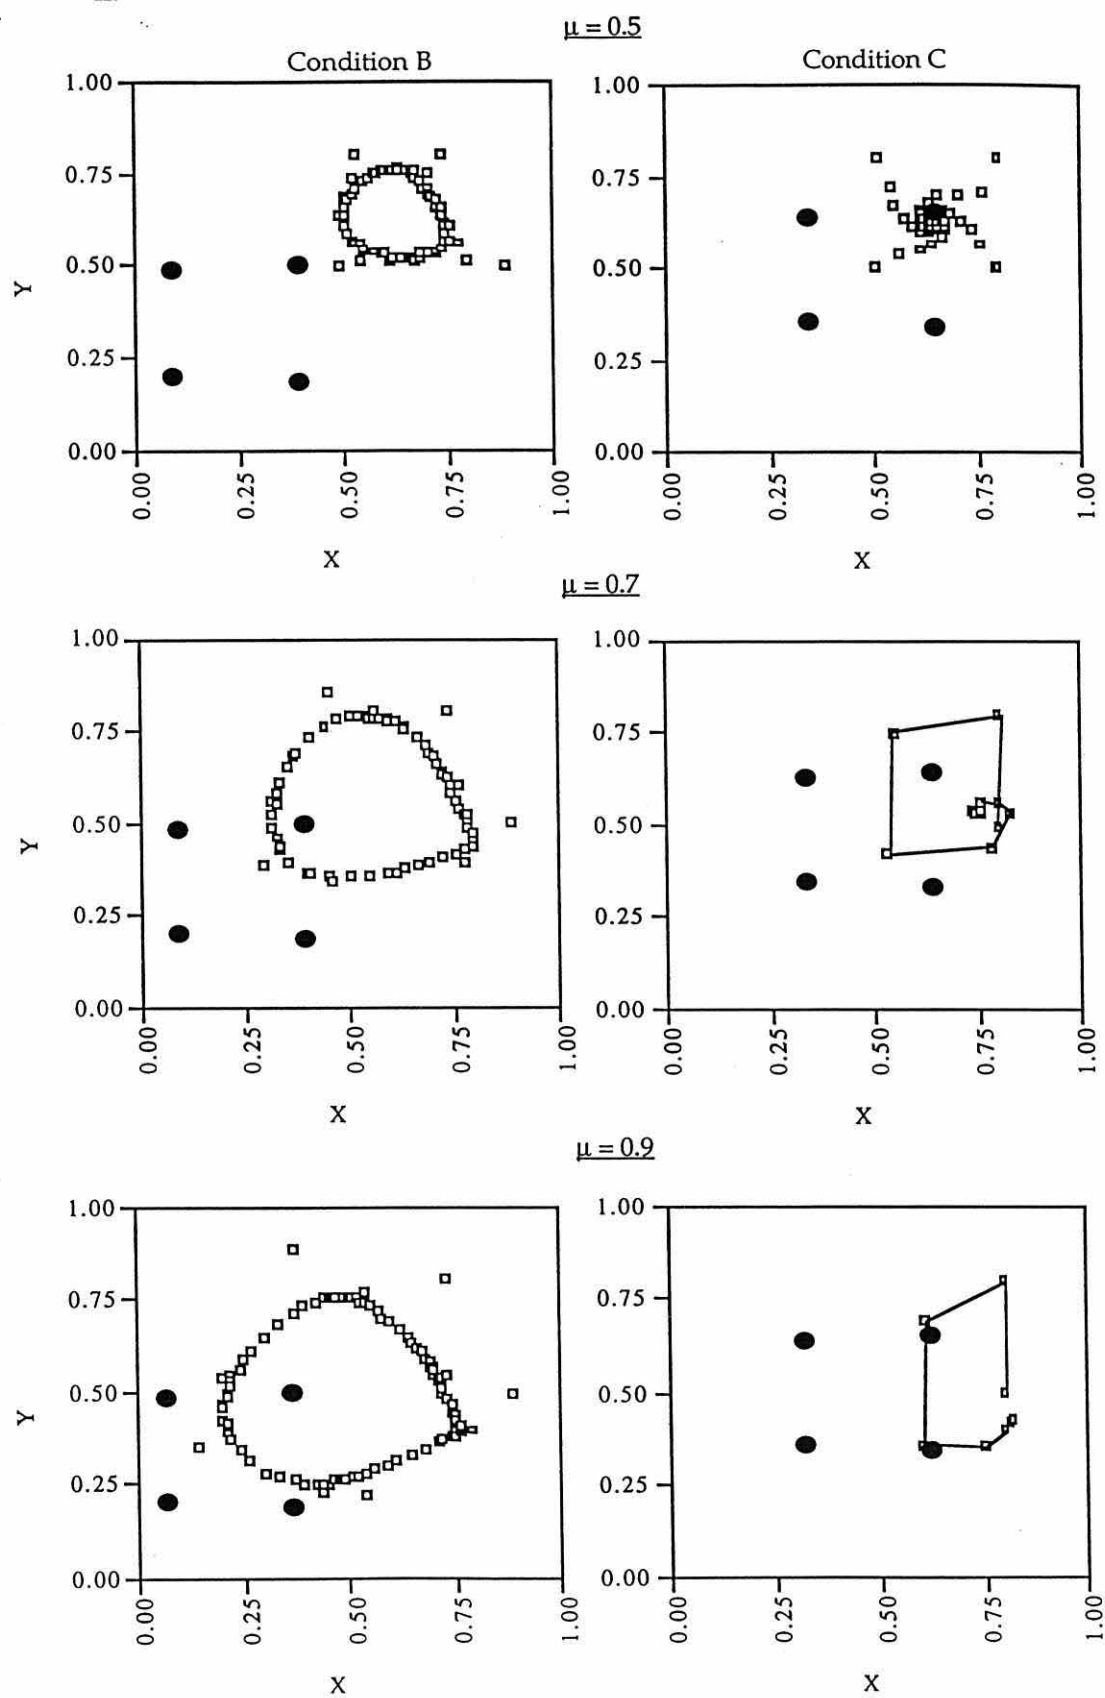
Figure 6.21. *Varying μ in conditions B and C*

Discussion

A dynamical systems interpretation of these results suggests that there is an interaction between the value of μ relative to the learned value of 0.5 and the relationship between the attractors for the dummy and test sequences. With μ below 0.5, the overlapping sequence has an attracting effect on the attractor for sequence 1, whereas the non-overlapping sequence has a repelling effect. With μ greater than 0.5 the inverse true. In other words, the non overlapping sequence has an attracting effect, indicated by the fact that the attractor expands to integrate elements of the dummy sequence in the set as μ is increased to 0.9. The overlapping sequence, however, has a repelling effect, driving the system away from the elements of that sequence and towards a fixed point near the lower right hand side of the task space.

These simulations indicate that the attractor dynamics in Jordan sequential networks are more complex than has previously been reported and indicates that interference problems may be significant in these networks if many sequences are learned. A broader discussion of the role of the parameter μ in a learned network is deferred to the end of the chapter. Here it is sufficient to note that the manipulation of this parameter in the sequential network holds a variety of computational implications for the use of attractors to represent patterns of movement in a distributed system. These simulations do not, however, find a clear interpretation in the modelling of movement disorders. We now turn to an investigation of a system level disruption which is motivated by the hypothesised effects of basal ganglia damage due to dopamine depletion.

6.5.3 Simulation 6.8: The effects of reducing the cascade rate

In the introduction to this section, we hypothesised that reducing the value of tau in the fully recurrent version of the network is a systems level manipulation which is sufficient to model the notion that dopamine performs a sort of contrast enhancement operation by increasing the signal to noise ratio in activating selected populations of pallidal neurons. In chapter five we attempted to model this function using reduced gain on the forward model units, following Cohen and Servan-Schreiber, (1992). The systems level effect of this

manipulation was to cause successive states of the context vector to be closer to each other, thus reducing the discrimination capability of the network. Here we suggest that reducing tau should have a similar effect on the processing of the network as a whole, but that it is a more principled manipulation, because it is designed to affect *only* this aspect of network processing, whereas there are several potentially confounding factors associated with altering the unit gain in a complex network.

A reduction in the value of tau in the current model should make it more difficult for the network to move on to the next element of the sequence because each subsequent state of the context vector is more similar to the previous one thus failing to drive the system forwards, thus leading to a slowing of performance measured in cascade cycles. However it is also predicted that there will be further abnormalities in the behaviour of the network due to the fact that the sequence representation is a complex dynamical attractor in task space, as shown in the previous section. These may be visualised using the same techniques as in previous simulations. Furthermore, it is possible to examine the interference effects of other task space attractors reported in the previous section on network behaviour when tau is reduced.

In order to assess the effects of this disruption under different conditions in terms of motor disorders the behaviour of the network is interpreted with reference to the categorisation of parkinsonian deficits outlined in chapter five.

Method

The general methodology of the simulations is similar to that described in the previous section. The same conditions A (test sequence only), B (test sequence and a non-overlapping 'dummy' sequence), and C (test sequence and an overlapping 'dummy' sequence) were evaluated at different values for tau. In each case the location of the endpoint of the arm was calculated at each cascade cycle. To confirm that the manipulation does, in fact, affect the states of the context vector as described above, each of the context unit activations were recorded at each cascade cycle.

Results

Figure 6.22 below depicts the activations of the individual context units over cascade cycles for two values of tau in condition A. The top graph is a slightly reduced value of tau (0.8) and the bottom graph is a much reduced value of tau (0.2).
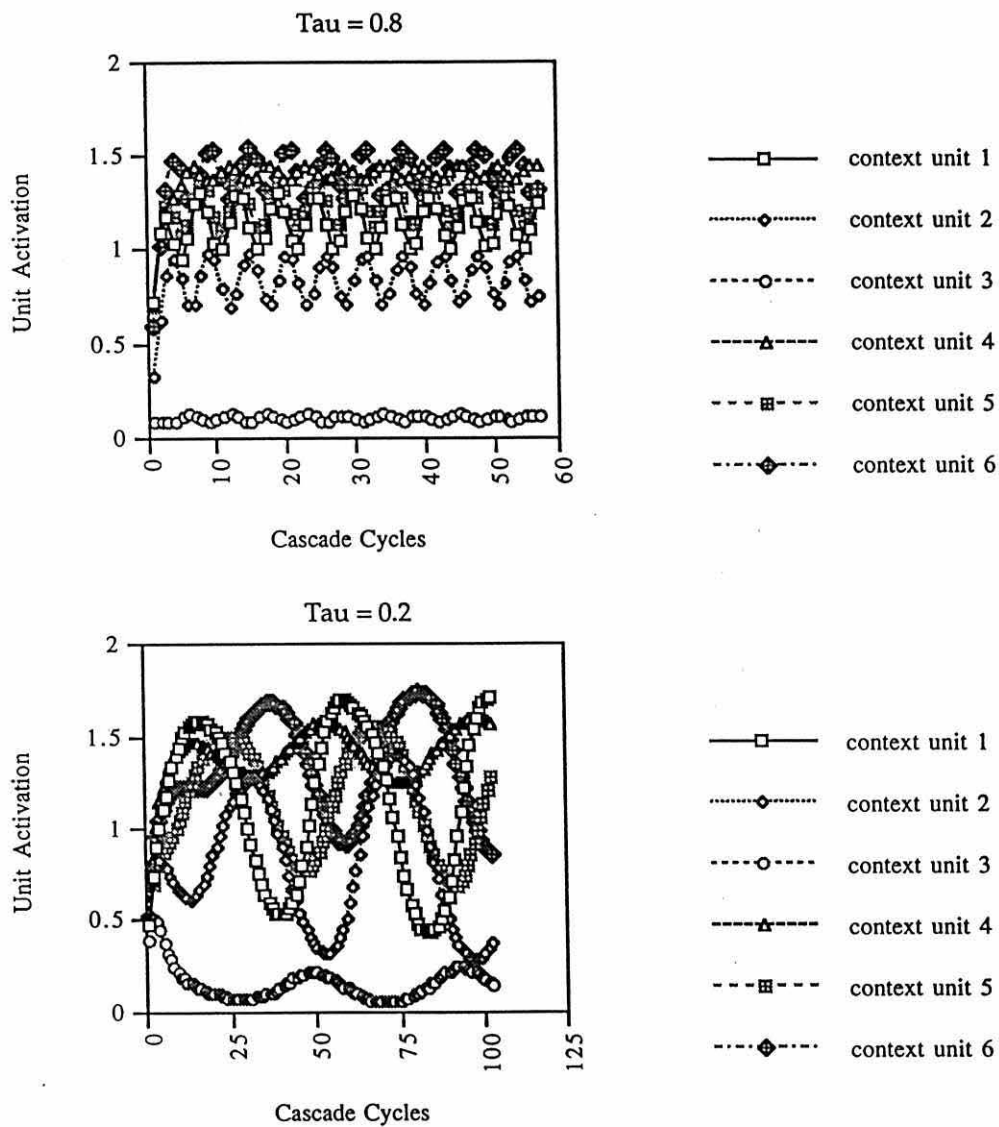
Figure 6.22. *Comparison of context unit activations for two different values of tau*

By comparing the upper and lower graphs of figure 6.22, it can be seen that two main changes are associated with a decrease in tau. Firstly, as predicted there is a reduced change in activation with each cascade cycle. Because this is true for every context unit, consecutive states of the context vector are more similar to each other, than with a higher

value of tau. The second difference is that the amplitude of each oscillation of each unit, where one oscillation corresponds to one complete cycle of the endpoint round the sequence, is substantially increased. It can also be seen that as a result of these two factors, the period of oscillation is also increased.

Condition A (test sequence only)

Trajectory maps of the endpoint for decreasing values of tau in condition A, in which only the test sequence was learned, are presented in figure 6.23 overleaf. There are two basic changes in the behaviour of the system when tau is reduced in condition A. First, the step size accomplished with each cascade cycle reduces and the time to complete one cycle of the sequence increases. Second, the diameter of the cyclical trajectory of the endpoint is increased.

There is, however, a critical point at which the change in activation between cascade cycles is so small that the asymptote condition on cascade processing causes the system to halt. This can be seen when tau is reduced to 0.1. The point at which this occurs is dependent on the setting of the asymptote condition and does not reflect any emergent behaviour in the network. The increase in size of the attracting set appears to reach a limit at tau = 0.3, further reduction of tau does not cause any further increase. It can also be seen that as tau is reduced, the network starts ever further away from the first point of the sequence, this is merely due to the fact that in cascade processing the activations of all units start at zero, and increase as activation is cascaded through the network. Zero activation on all the articulatory units presumably encodes an endpoint location near the top left hand corner of task space, and since the change in activation is reduced as tau is reduced, so the activations after one cascade cycle encode a location nearer that point.
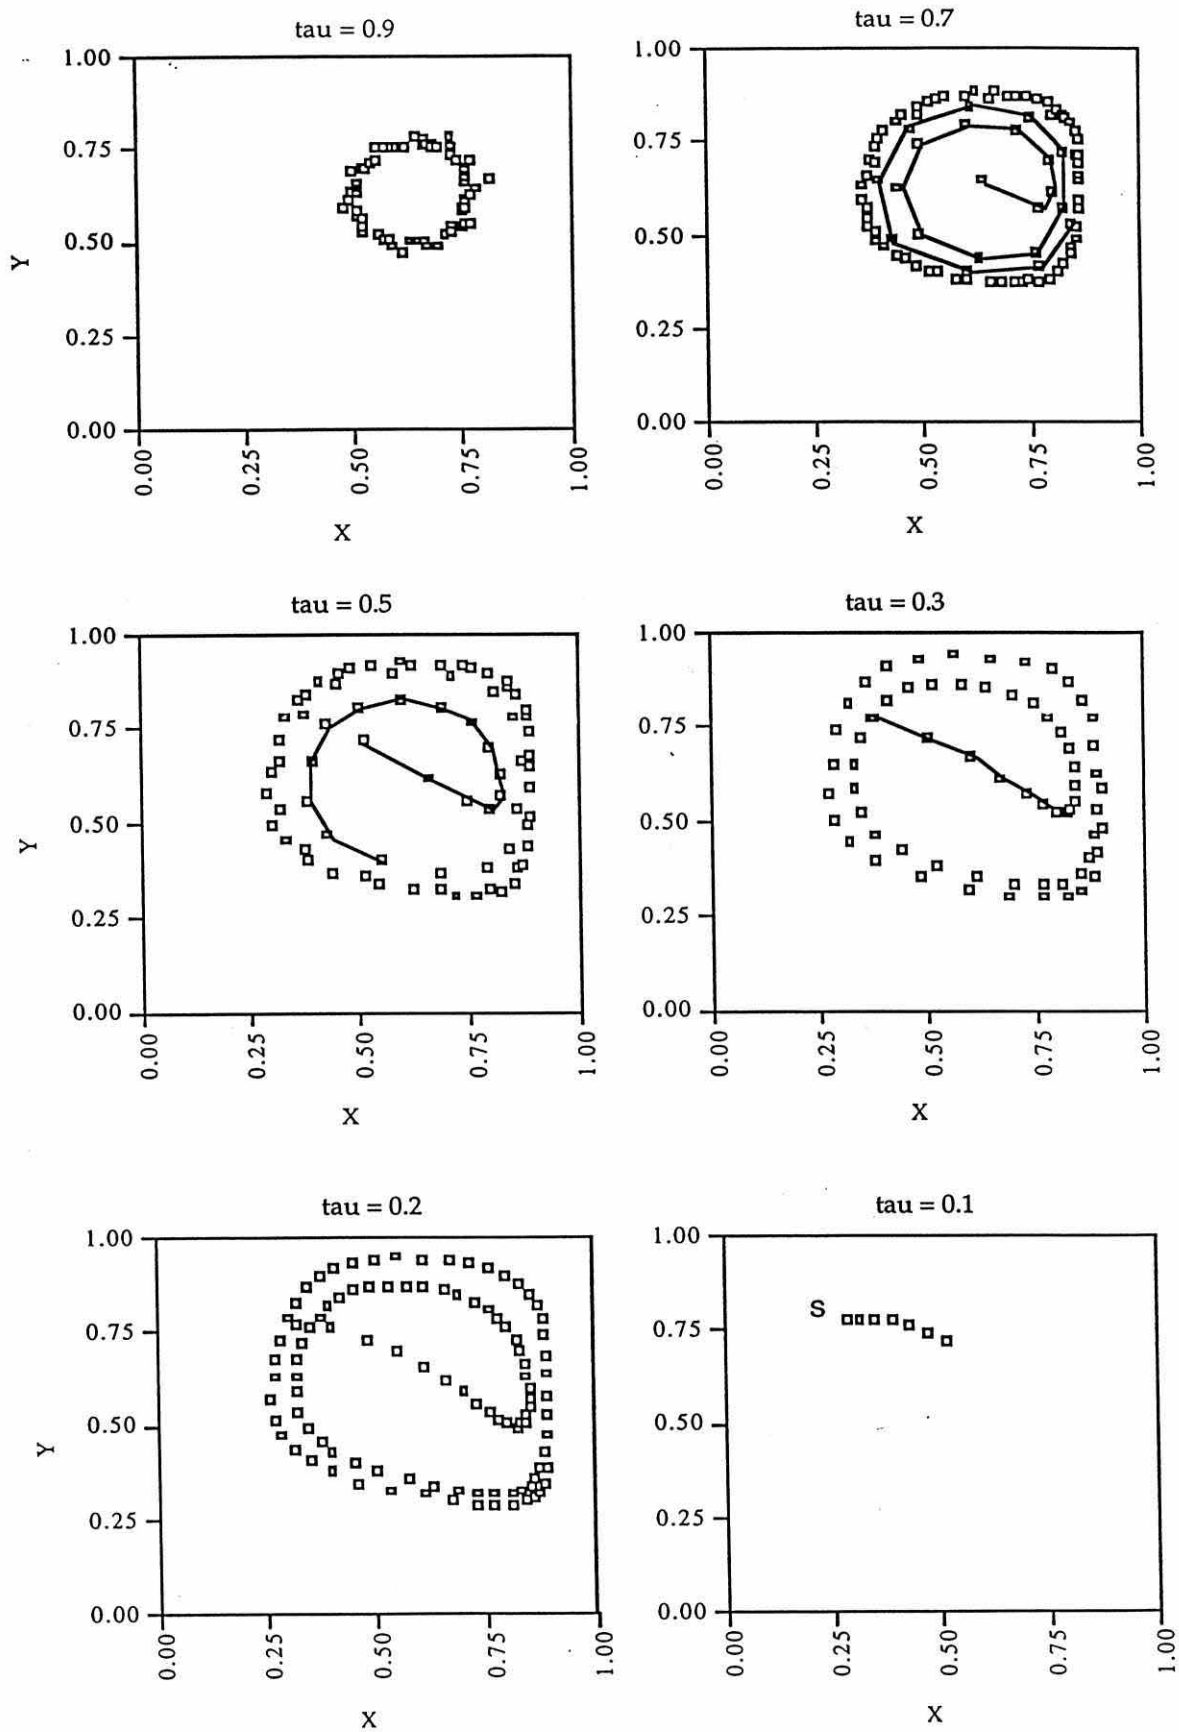
Figure 6.22. *Decreasing values of tau in condition A*

Conditions B and C (second sequence also learned)

Figure 6.24 below depicts the results of reducing tau in conditions B (non-overlapping) and C (overlapping). Two values of tau are shown for each condition: 0.9 and 0.5.
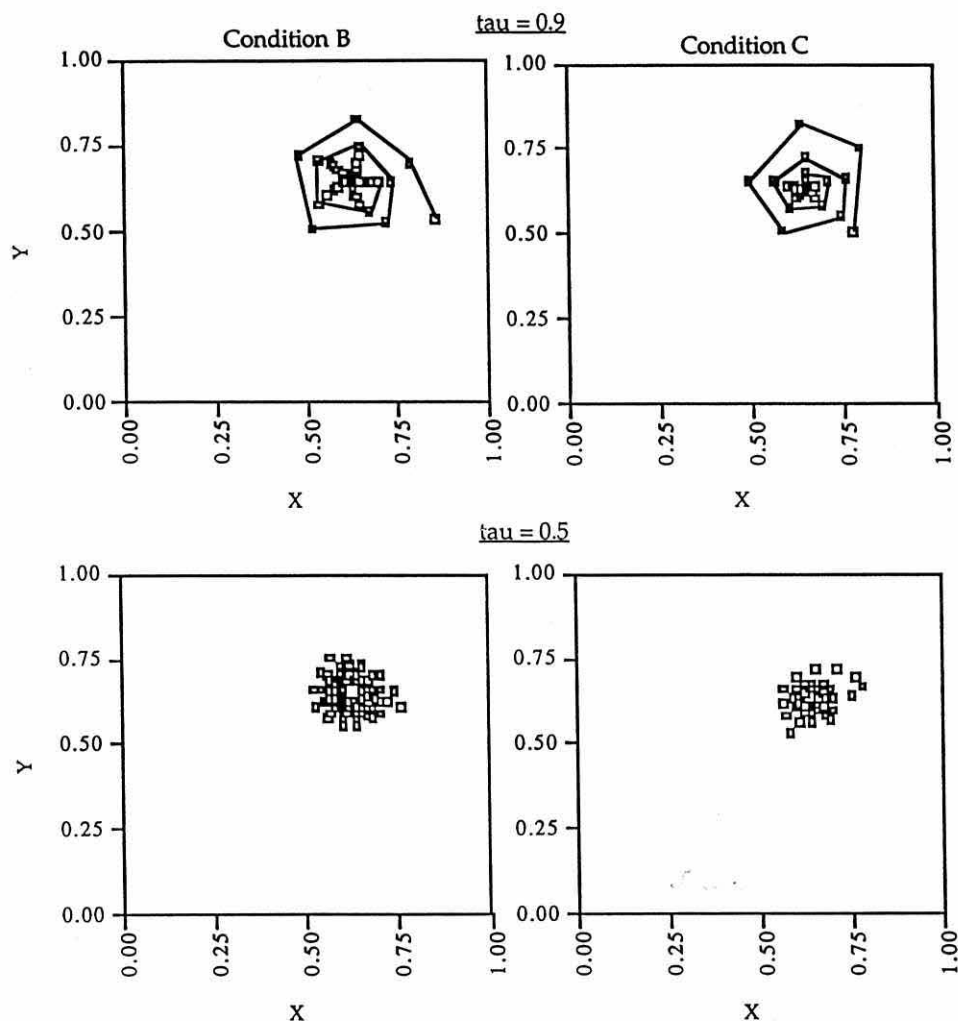
Figure 6.24. *Reduction of tau in conditions B and C*

The pattern of behaviour produced in both conditions B and C is qualitatively different from that produced in condition A. Instead of cycling outwards in task space when there is a second sequence, the network cycles inwards to a fixed point. This behaviour is produced with even the very small reduction in tau seen in the top pair of panels, (tau = 0.9). When tau is reduced to 0.5, the same behaviour is evident, but as in condition A, the amplitude of each movement is decreased, so it takes much longer to reach the fixed point. There is a small quantitative effect of the greater interference caused by the overlapping sequence in

condition C, in that the network moves towards the fixed point more quickly at both values of tau, than in condition B.

Discussion

In summary, the pattern of behaviour produced by this manipulation fails to fulfil the criteria we put forward for modelling sequential motor deficits in Parkinson's disease with this network, in chapters two and five, but is partially consistent with the hypotheses presented at the beginning of this chapter. To recap briefly, in order to model parkinsonian impairments the network should produce, at least, the following qualitative deficits:

1. A reduction in the amplitude of each movement to model the general property of hypokinesia.

2. A progressive reduction in amplitude as the sequence progresses to model a special deficit in sequential behaviour.

3. Bifurcation of the attracting set to a fixed point in the presence of competing actions.

4. 'Omissions' of target locations caused by smoothing of the sequential trajectory.

In all conditions A, B, and C there is a reduction in movement amplitude as tau is reduced, but in none of the conditions does the impairment worsen as the sequence progresses. Thus there is no special deficit in sequential behaviour.

In condition A, the system does not bifurcate to a fixed point attractor, instead the system expands to cycle over a greater area. There are two possible computational explanations for this behaviour. Either, using the Hopf bifurcation model, reduction in tau causes the attracting set to shift to a plane further away from the point of bifurcation. Or alternatively, reduction in tau causes a weakening of the attractor on the learned sequence thus allowing the system to drift farther away from the learned sequence. This makes sense because the strength of the attractor may be linked to the size of the activation change at each time step. When the system energy is reduced, so the attractor is weakened. In the absence of any other influences in the state space, the outward movement of the endpoint is not constrained.

However, in conditions B and C, the system does bifurcate to a fixed point with even a very slight reduction in tau. This supports the notion that in the presence of more

than one attractor the pressure on the selection mechanism is increased, and the outcome is a qualitative change in the behaviour of the system. In this respect, the behaviour of the system was consistent with the Hopf bifurcation model of motor disorders, put forward in section 6.3.2.

6.5.4 Simulation 6.9: Selective reduction of tau in the forward model

One reason why the network may have failed to produce a selective sequential deficit in the previous simulation is that tau was reduced in both the controller and forward model subnetworks. In chapter five, we put forward several independent motivations for selectively disrupting the forward model subnetwork, including the fact that it is not directly involved in the computation of movement parameters. Several of the simulations in chapter five showed that selective disruption of the forward model subnetwork provides a much better characterisation of parkinsonian deficits. In this simulation we selectively reduce tau in the forward model subnetwork. This allows us not only to view the effects of reducing the cascade rate in the forward model alone, but also to investigate what happens when the relative timing between the two subsystems is disrupted. We have previously suggested that this may be one feature of the computational dysfunction in communication between cortical and subcortical areas (Brown, Britain, Elvevag & Mitchell, 1994).

Method

In order to reduce selectively tau in the forward model subnetwork, the program was simply rewritten with two variables coding for tau, one for the control subnetwork, the other for the forward model. When both variables were set to 1, the network behaved normally as before. However, the variable for tau in the FM could be selectively manipulated. Again this manipulation was performed in conditions A, B and C.

Results

Condition A: (test sequence only). The results of reducing forward model tau in condition A are displayed in figure 6.25.
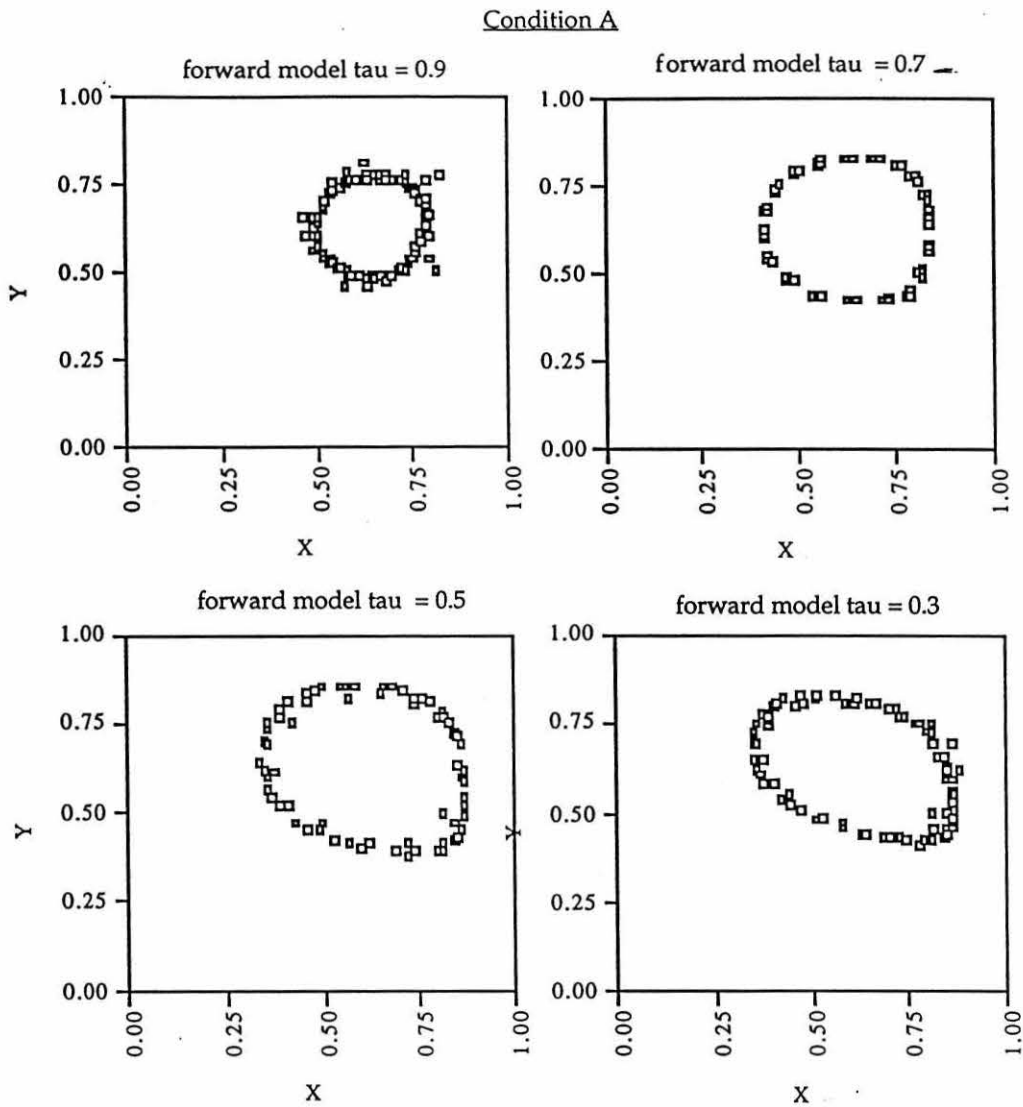
Condition A

Figure 6.25.*Reducing the forward model cascade rate in condition A*

When the discrepancy is relatively small, as in the top pair of panels a similar pattern of results to those when tau is reduced in both subnetworks is seen. The step size is diminished but there is no qualitative change to the shape of the attractor. However, the more extreme discrepancies shown in the bottom two panels depict progressive squashing of the attractor into an elliptical shape as the difference in relative cascade rates between the two subsystems is increased. Although it is not clearly visible in the figure, there is no selective sequential deficit in this condition.

Figures 6.26 and 6.27 below display the results for conditions B and C respectively.
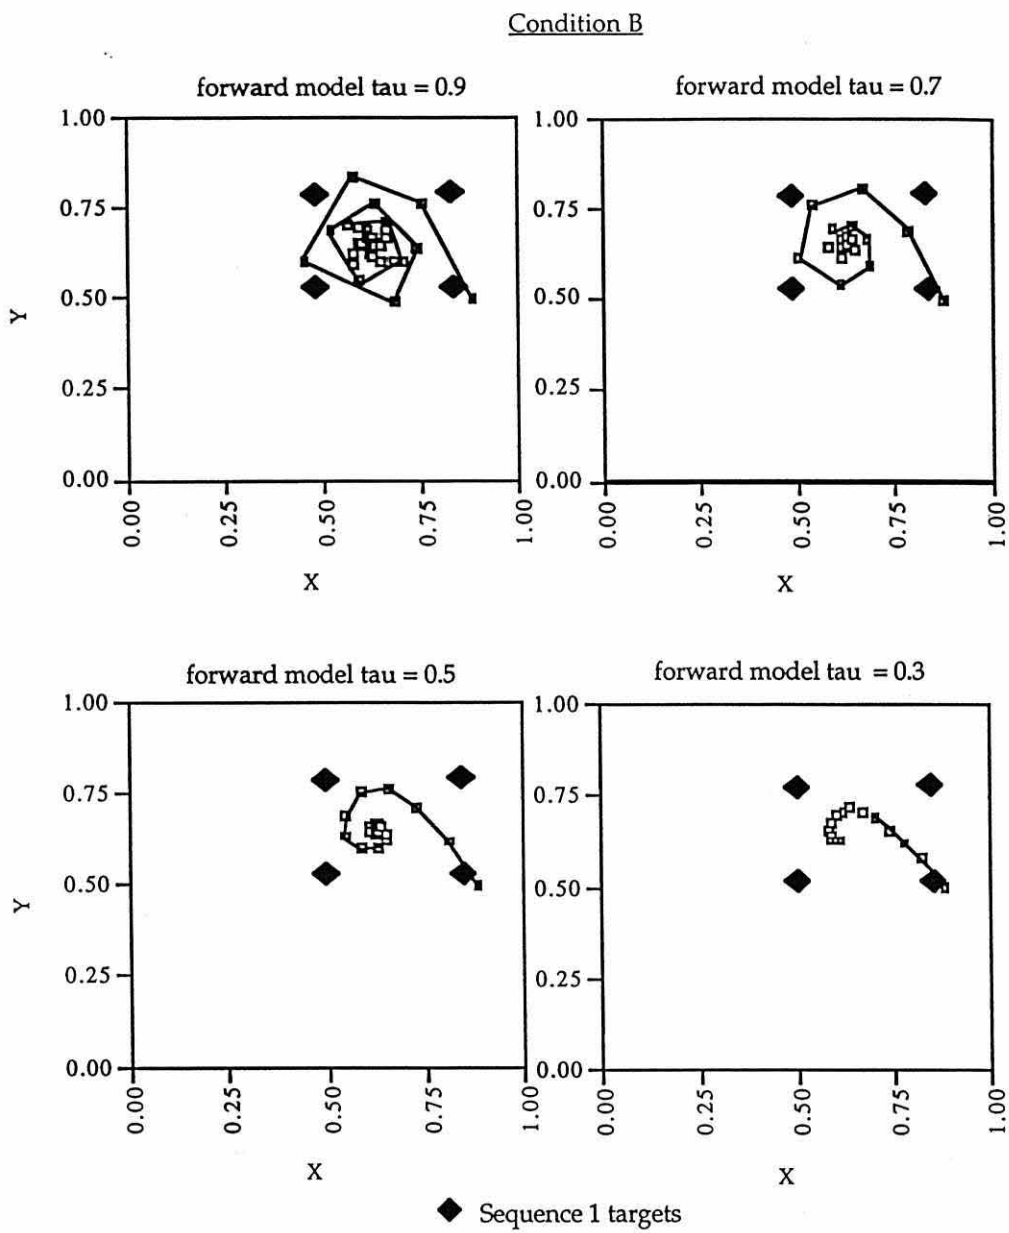
Figure 6.26. *Reducing the forward model cascade rate in condition B*

Condition C

forward model tau = 0.9

forward model tau = 0.7

forward model tau = 0.5

forward model tau = 0.3

◆ Sequence 1 targets
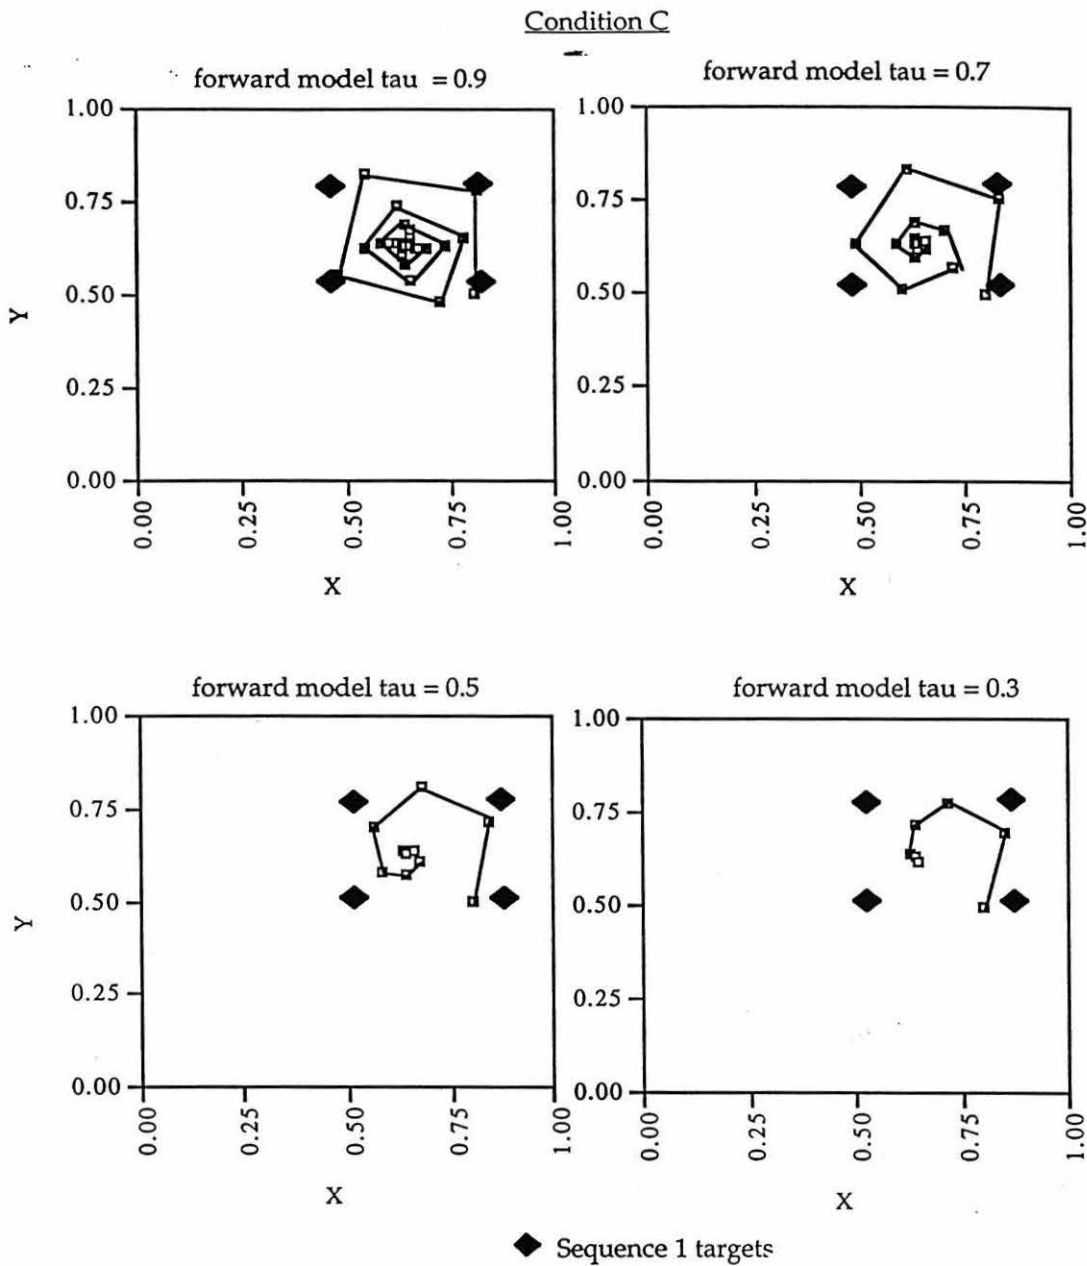
Figure 6.27. *Reducing the forward model cascade rate in condition C*

It can be seen from both of the above series of trajectory plots, that when there is an additional sequence encoded by the weights, regardless of location, there is a qualitative difference in the behaviour of the system from condition A, where the target sequence alone is encoded. Specifically, in both conditions, for all reduced values of tau in the forward model, the cyclic attractor bifurcates to a fixed point at approximately 0.6, 0.6. In all cases, the system moves towards the fixed point in fewer cascade cycles as tau is reduced and the trajectory describes a smaller arc in task space, indicating a 'reduction' in the shape of the

whole action. At more severe levels of disruption (tau = 0.3 in condition B) there are 'omissions' caused by smoothing of the trajectory and failure to complete the sequence before converging on the fixed point. In all cases there is decreasing amplitude in the movement of the endpoint as the sequence progresses, indicating a sequential deficit which was not evident in condition A. Finally, at all levels of tau, the system converges on the fixed point more rapidly in condition C, where there is greater interference, than in condition B.

Discussion

The results produced in this simulation for conditions B and C provide the best model of parkinsonian motor deficits, according to the set of criteria set out in chapter five, achieved with any of the disruptive manipulations performed so far. The deficits are in many ways comparable with the effects of reduced forward model gain produced in chapter 4. The advantage here is that the computational basis for the deficits is clear and is consistent with the biological notions put forward in the conceptual framework of interference in parkinsonian action selection and initiation. Moreover the visualisation and analysis in terms of the attractor topology and dynamics developed throughout this chapter provides the basis for a deeper level of computational explanation of the behaviour of the network under disruption than would otherwise have been possible. The key computational features underlying the impaired behaviour of the network in modelling parkinsonian deficits are summarised below:

1. Disruption of a subsystem, not explicitly concerned with the computation of movement parameters (the forward model).

2. Disruption of the forward model by reducing the cascade rate (tau) causes successive states of the context vector to be more similar to each other.

3. When this occurs there is interference from other nearby attractors in state space.

4. This interference leads to progressive reductions of movement amplitude and slowness, eventually leading to 'freezing' of the system at a fixed point.

5. The more the cascade rate is reduced, the more pronounced the deficits, and the sooner freezing occurs.

This pattern of behaviour is consistent with the failure of Parkinson's patients to sustain motor activity and ean be attributed to over-inhibition of action representations due to the breakdown of a selective disinhibition mechanism. In the final section of this chapter, we turn to a broader discussion of the issues addressed in this chapter, and we highlight some of the limitations and shortcomings of this modelling approach.

6.6 Chapter summary and discussion

6.6.1 Chapter summary

In the first part of this chapter we set out three main computational aims. These were: 1. To introduce an independent temporal component to processing in the combined sequential network. 2. To analyse in greater depth the nature and dynamics of the attractors formed by the network following learning. 3. Investigate the response of the learned network to disruptions which are motivated by the conceptual model of basal ganglia damage presented in chapter four.

These computational aims were linked to a set of hypotheses which relate network processing viewed as a nonlinear dynamical system to motor processing in biological systems. The main theoretical and computational contributions resulting from the work in this chapter are set out below.

1. A distributed global action representation or motor schema may be learned on the basis of a 'high-level' motor plan and a set of internal contextual variables, which drive sequential responses.

2. The action representation may be characterised as an attractor in a multi-dimensional neural state space.

3. This representation maps onto a virtual location of the endpoint of a limb in task space.

4. A fixed-point attractor in state space or virtual equilibrium point in task space may serve to represent a stable state of the endpoint of a limb. This may encode a held posture if it is the same as the current location of the endpoint, or a desired final endpoint location of a simple goal directed movement during motor planning.

5. If the neurally encoded equilibrium point moves through task space over time, a virtual trajectory is established, which provides the control parameters for the effector system, according to the principles of the EQ hypothesis. The location of the equilibrium points establishes the path of the endpoint, where the path between two points is a straight line. The distance between equilibrium points reflects the required force of the movement. Thus if they are close together, only a small force is required to cover the distance in one time slice, but if they are further apart a greater force is required to cover the distance in the same time period.

6. The same principles allow the representation of more complex repetitive or sequential movements as a more complex attracting set, the simplest case of which is a limit cycle.

7. Nearby attractors in state space interact with one another. This is a computational feature of representation in a distributed system, and is desirable in potentially allowing adaptation through integration of existing motor patterns. However, there is also potential for debilitating interference between motor patterns.

8. Normal feedforward control of sequential motor behaviour requires an intact internal contextual control signal. Disruption of the context signal either by changing the depth of information encoded or increasing the similarity between successive states, causes interference from nearby attractors to profoundly affect motor behaviour.

9. Motor disorders and the disinhibitory function of the basal ganglia may be understood in a dynamical systems framework, using the concepts of stability and bifurcations in the behaviour of the system. If the Hopf bifurcation is used as a model of system behaviour, reduction of a parameter reflecting dopamine function causes a reduction in the size of the attracting set representing the desired motor pattern, and at a critical point causes a fixed point to become attracting, leading to freezing of the system at that point. However, in normal behaviour the ideal state in which the system should be maintained is reasonably close to the critical point, such that selective inhibition / disinhibition requires only small changes in the key parameters. Maintenance of this balance may reflect the role of dopamine in the normal operation of the basal ganglia, although this function is not illustrated by the model.

10. Other problems may be associated with meta-stability in the system and points far from the bifurcation point. Meta-stable solutions such as simple limit-cycles in which each state is repeated exactly, may reflect stereotyped behaviours, in which a cycle of actions is continually repeated without modulation according to contextual influences. Stereotyped behaviour patterns has been linked to abnormalities of basal ganglia function, for example in Tourettes syndrome. Furthermore stereotypy can be thought of as the opposite of PD in terms of system dysfunction (Ian Mitchell, personal communication). This suggests that the approach put forward here has the potential to account for other abnormal states caused by basal ganglia dysfunction.

11. The model presented in this chapter combines the notions of attractor representation of motor patterns in a distributed dynamical system, internal contextual control of sequential behaviour and bifurcations with changes in key parameters to simulate features of normal and disordered control of limb movements in a two-dimensional space.

6.6.2 Review of the modelling approach

The modelling approach developed in this chapter using the modified 'cascaded' Jordan network represents a considerable improvement on the work presented in chapter five. The use of the cascade mechanism permits a more comprehensive analysis of the behaviour of the network than was achieved in chapter five. The inclusion of a second learned sequence and the examination of the effects on performance of the target sequence when the system is disrupted provides insights into the problems of cross-talk in a parallel distributed system. This issue holds important implications for striatal processing as discussed above. In chapter four we suggested that one of the multiple sources of information provided to the striatum is a contextual control signal that facilitates internal control of sequential actions. In this model recurrent connections drive the control signal and we have examined the effects of disrupting the control signal. This manipulation produced deficits which are analogous to some aspects of parkinsonism. We have argued in this chapter that by making successive

states of the control signal more similar to each other, competition between motor patterns impairs output behaviour.

However, in chapter four we also proposed that the basal ganglia implement a supplementary mechanism for resolving competition via general inhibition and selective disinhibition. This has not received explicit treatment in the modelling work so far. We have simply assumed that there is inherent competition to the distributed nature of the network and have concentrated primarily on the properties of the control signal. The question of inhibitory mechanisms and their role in processing still needs to be addressed.

Another limitation of the current model is that movements of only a single limb have been considered. The explicit role of inhibitory mechanisms for maintaining a held posture in a subset of limbs whilst selectively disinhbiting another set may become apparent if more than one limb is modelled. In addition problems of bi-manual coordination and simultaneous movement in different effectors can only be examined if additional effectors are included. In the next chapter we address these and other limitations of the current framework.

CHAPTER SEVEN

Limitations of the model and further issues in modelling motor control

Despite a number of interesting results obtained with the model in the previous chapter, there are various problems and limitations that remain to be addressed. In this chapter we outline the potential shortcomings and present extra simulations which were performed to illustrate our arguments. We also address here some theoretical issues related to design aspects of the sequential connectionist network itself. The primary issues discussed below are:

1. The role of inhibitory mechanisms in a multi-effector system.

2. The role of the parameter μ in the learned network.

3. Contextual representation without dependence on recurrent connections.

7.1 Coarticulation and the need for inhibitory mechanisms in a multi-effector system

A limitation of the current model is that only one effector is represented, and is used to perform all of the tasks. One of the central arguments in favour of the basal ganglia performing a selective disinhibition function in facilitating programmed limb movements is that some limbs (e.g. the prime mover) will need to be disinhibited, whilst others will need to be concurrently inhibited.

We have argued throughout that either unselective overinhibition of all body parts or partial co-activation of conflicting body parts is what gives rise to the motor symptoms in PD. Although we have addressed this issue by modelling one multi-jointed effector, it may be argued that a much fuller range of deficits would be observed if the model were

expanded to include two or more effectors. For example, failure to sufficiently inhibit the 'unselected' hand during a one handed signing movement may lead to the shadowing effect observed in the deaf parkinsonian signer, studied by Brentari and Poizner (1994).

Generalisation to a multi-effector system is possible within the current framework. Jordan (1990) reported a number of simulations in which the network encoded two effectors, similar to the one encoded in our simulations. In this scheme, targets were only assigned to one of the effectors at any one time slice, and the units coding for the other effector were free to take on any activation values at all. Jordan showed that when an additional internal 'smoothness' constraint is applied to network processing, the network exhibits coarticulatory behaviour in the unselected effector which effectively smooths the transition between sequence elements, thus simulating a form of motor 'synergy' which is typical of skilled human movements. If such synergy formation is a natural part of the motor learning process, then this argues in favour of the need for additional inhibitory mechanisms for preventing synergetic movement of two limbs when it conflicts with task demands. This point is illustrated by a simple simulation below.

7.1.1 Simulation 7.1: The need for inhibitory mechanisms in a dual-effector system

This simulation and all further simulations carried out in this chapter were performed using a simplified version of the network used in the previous simulations, in which the forward model is collapsed and the activation values of the endpoint locations for each effector are computed directly by the controller subnetwork. An architectural diagram of the system is presented in figure 7.1.

The activations of output units 1 and 2 code directly for the X and Y coordinates of the left hand endpoint in task space and similarly output units 3 and 4 code for the right hand endpoint location. This simplification reduces to an assumption that the kinematic configuration of the effectors is unconstrained at this level of motor computation, and all that is required is a specification of endpoint location in task space. Although this may be psychologically naive, it is a reasonable assumption for the purpose of the current simulations. In particular it allows us to examine properties of the model that are due solely

to the properties of the sequential control network, without the added complexity of the forward model.
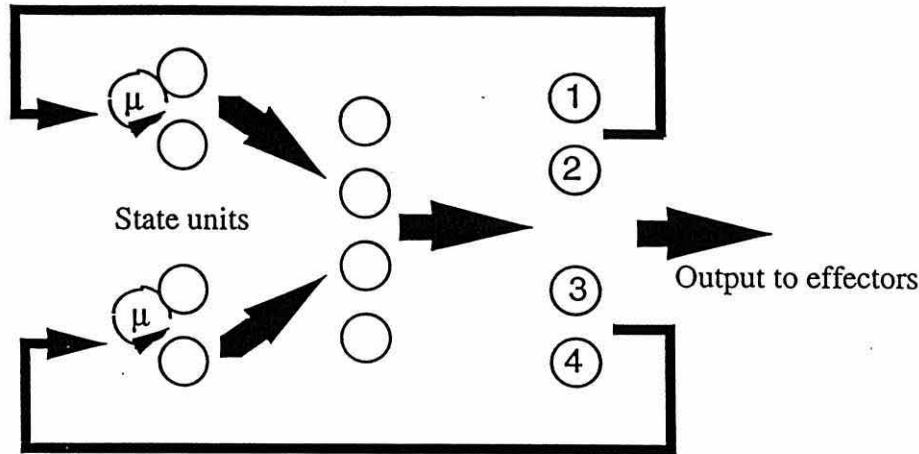
Figure 7.1. *Network architecture. Connection pattern has been schematised and plan units are not shown*

In this simulation we examine the effects of 'don't care' constraints on the movement of an unselected effector in two task conditions. In one there is no action encoded on the unselected effector, whilst in the other condition a sequence has been encoded for the right hand effector.

Method

The basic task given to the network was to learn a cyclic four element sequence in a two dimensional task space, similar to those used in chapter six. The basic task was divided into two different conditions. In the first condition (condition 1) the network learned to perform the sequence with the left effector with 'don't care' constraints on the right effector. The don't care constraint is implemented by simply not backpropagating an error from the output units coding for the unselected effector, during learning. This allows those units to take on any values within the range 0,1. In the second condition, (condition 2), the sequence is learned with the left effector as in condition 1, but a mirror version of the sequence is also learned with the right effector (in the context of an orthogonal plan vector). In both cases learning was accomplished with a learning rate of 0.1, momentum = 0.9 and $\mu = 0.5$.

The network was then tested by assuming a task in which the learned sequence is to be performed by the left hand effector only. Thus the left effector is referred to as the selected effector. Testing took place by presenting the network in each condition with the plan vector which was associated with the sequence learned by the left hand. The movement of each of the effector endpoints was plotted in the familiar way.

Results

The results are depicted in figure 7.2 below.

Condition 1

Endpoint locations for the selected effector (LH)   Endpoint locations for the unselected effector (RH)

Condition 2

Endpoint locations for selected effector sequence 1   Endpoint locations for the unselected effector (RH)
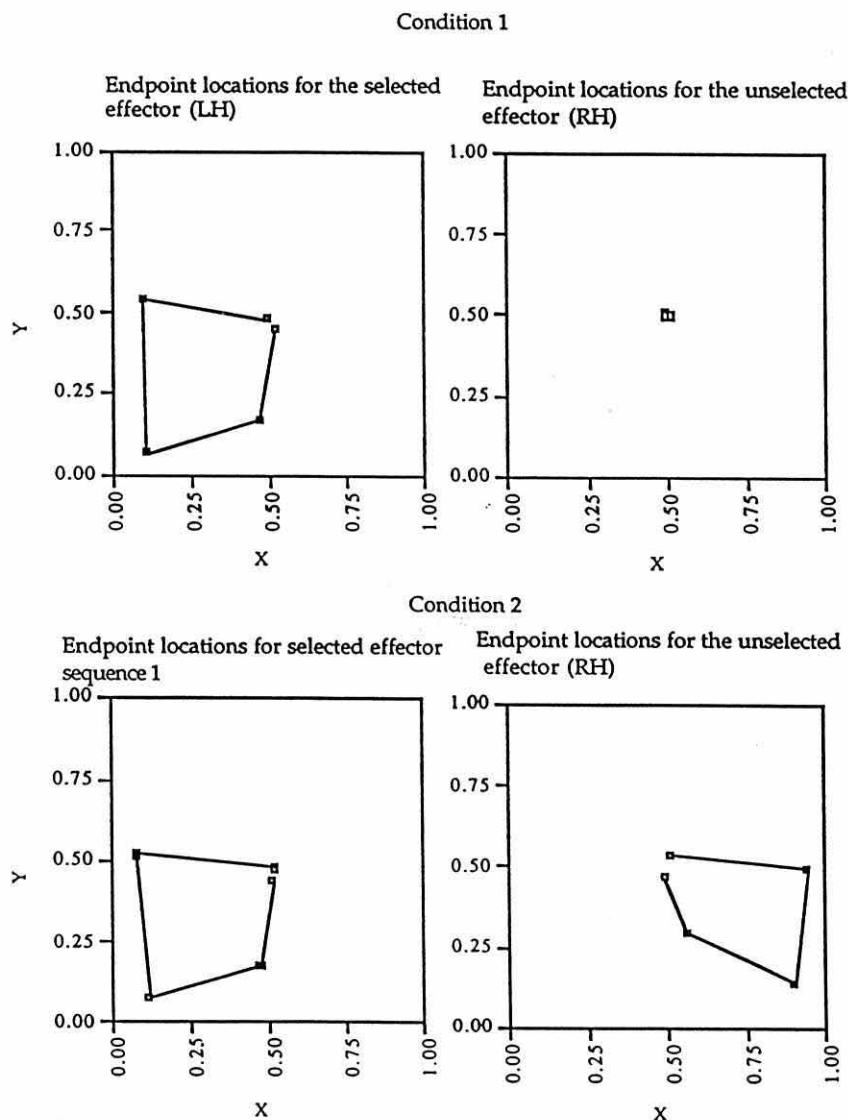
Figure 7.2. *The effects of don't care constraints on the behaviour of an unselected effector when either 1 or 2 actions are encoded*

It can be seen in the top pair of panels in figure 7.2 that when only the target action is encoded by the network, the unselected effector remains static throughout the movement.

However, when a similar action is encoded for the right hand (RH) effector also, the effect of the 'don't care' constraints on the unselected effector is to cause it to trace out a degraded form of its learned 'mirror' version of the target sequence (see the bottom right panel of figure 7.2).

Discussion

This simulation provides a simple demonstration of both the value and the limitations of the Jordan approach from the current perspective. The introduction of don't care constraints on the clusters of output units representing each effector causes a form of motor synergy to emerge, which is in keeping with current knowledge about the organisation of motor behaviour[1]. Yet, it is also clear that if this synergetic behaviour is to be suppressed, extra inhibitory mechanisms, which do not form part of the current architecture, would need to be employed.

Naturally, it could be argued that this 'toy' simulation with only two mirror actions is too simplistic, and that if many actions were encoded the multiple constraints on the unselected effector would abolish the 'synergetic' behaviour. The stronger counter-argument, however, is that on scaling up, the variable multiple constraints (i.e. interacting attractors as shown in chapter five) on the unselected effector would be likely to result in unpredictable and noisy activations in the control of that effector. This would, indeed, abolish the apparent synergy, but would argue even more strongly in favour of a supplementary inhibitory system to suppress the unwanted movements.

The need for extra inhibitory control demonstrated in this simulation is an important consideration from the point of view of the computational architecture. In the previous chapter we suggested that explicit inhibitory mechanisms are not required because gradient descent learning itself prevents competition between similar actions occuring (unless the network is disrupted). However, when don't care constraints are incorporated, as they must be to achieve co-articulation in the multi-effector system, the discrimination achieved

---

[1] This behaviour only occurs in the presence of some internal constraint which promotes coarticulatory smoothness. This is provided here by the parameter μ. The role of μ in this context is discussed further in the following section.

by learning alone is not sufficient to prevent undesirable co-activation of the unselected effectors. Thus, in this case, extra mechanisms must be assumed to exist.

7.2 The Role of μ in sequential behaviour and encoding contextual information

Throughout this and the previous two chapters we have conducted the simulations with recurrent connections of the context units onto themselves and with an intermediate value of the parameter μ, ($\mu = 0.5$). There are a number of psychological and system level motivations for including the self-recurrent connections on the state units, including the notion that they allow the state representation to be thought of as a representation of temporal context. This notion is central to our use of the Jordan sequential network in modelling motor disorders. In order to explicitly examine the properties of the state representation when set up in this way, we include here a further discussion of the role of μ, supported by simulation results using the architecture presented above.

Jordan's primary motivation for including μ in the network operation is that it makes sequences with repeated elements easier to learn. This is the case because the state units will hold different representations for the repeated element depending on its position in the sequence, which is what allows us to think of them as encoding sequential context, and thus refer to them here as *context units*.

There are, however, additional motivations underlying this property. If the auto-recurrent connections are omitted, or μ is set to 0, then the Jordan network implements a simple chaining architecture, in which the network is only sensitive to the output at the previous time slice in determining the subsequent action. Arguments in favour of chaining architectures in human control of serial order and motor programming are notoriously difficult to maintain in the face of the weight of contradictory arguments originated by Lashley (1951).

Indeed, the obvious similarity between the Jordan network and a chaining architecture is so striking that it has caused some current critics to oppose its use as a model of human sequential behaviour regardless of the properties of a state representation which

encodes extended context (Houghton, 1990; Houghton & Hartley, 1995). Consequently some care is required in evaluating whether or not these criticisms are valid. One of the main arguments used by these critics against chaining architectures is that they cannot show true co-articulatory behaviour in the sense that upcoming sequence elements are facilitated ahead of time, such that a discrete series of movements are performed as one fluid action. We have already discussed at length why this must be a fundamental property of human sequential motor programming and why also it must be regarded as central to the breakdown of sequential movements in motor disorders. Furthermore, we have, at length considered the contribution of the contextual representation encoded on the state or context units (with μ) in the current network to the behaviour of the system in that respect. However, we have not as yet explicitly demonstrated this property.

Jordan has used an independent internal 'smoothness' constraint to promote co-articulation in the network, (e.g. Jordan, 1990; Jordan, 1992). This is simply an extra error term which serves to minimise the difference between the activation of a unit at the current time slice and its activation at the previous time slice. Thus the network with excess degrees of freedom in the articulators finds a solution in which the transition between targets is smooth in task space. However, Jordan (1986a) suggests that the parameter μ exhibits similar properties. This we showed to be the case in chapter six using the cascade mechanism to visualise the behaviour of the network between target locations (see figure 6.1). Here the system, with no additional smoothness constraint, preferably exhibits a virtual trajectory which approximates a straight line in task space. This is consistent with data on human movement planning and provided additional psychological motivation for using $\mu > 0$.

Here we report simulations that explicitly compare the properties of μ and the smoothness constraint used by Jordan in facilitating upcoming movements and also compare system behaviour with and without μ. The general form of the task on which the network is tested is illustrated in figure 7.3 below. The sequence of targets to be touched is numbered 1-4 in the figure. However, the task is divided between the two effectors such that the left hand effector is assigned only to targets in the left hand half of the task space (1

and 3) and the right hand effector is similarly assigned to the right hand targets 2 and 4. There is considerable scope for anticipatory behaviour in the structure of the task. For example, at time step 2, the RH effector should be touching target 2, but the extent of anticipatory behaviour may be seen in the position of the unselected LH effector at that time step. If it remains near target 1, then there is little anticipation, but if it is nearer target 3, then the network is anticipating the required location of the LH effector at time step 3.
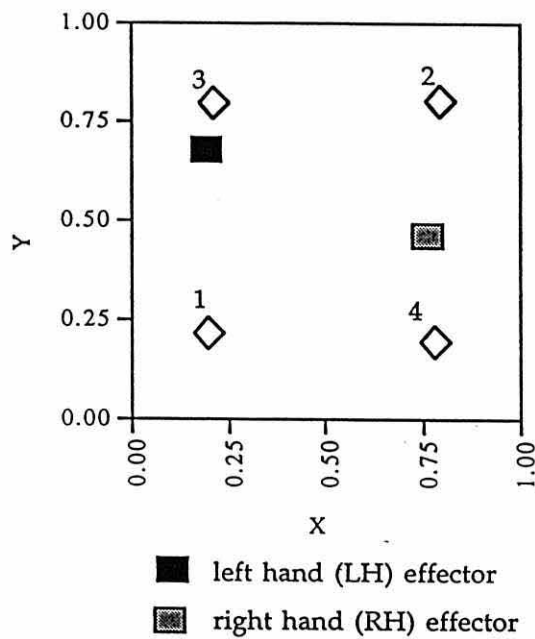
Figure 7.3. *The dual effector sequence task.*

7.2.1 *Simulation 7.2: Comparing the effects of μ and a smoothness constraint on anticipatory behaviour in sequential processing*

Method

The task described above was presented to a sequential network with an architecture and processing parameters identical to those described in the previous section. Two versions of the network were used in the simulation. In the first version the network was implemented as before, with a value of 0.5 for μ. In the second version an additional smoothness constraint was implemented on the output units of the network during calculation of the backpropagated error term. Thus instead of computing the error at each output unit **i** simply as:

$$error[i] = target[t] - activation[i]; \qquad (7.1)$$

the following term was added:

$$smoothness[i] = gamma(activation[i] - prev\_activation[i]); \qquad (7.2)$$

where prev_activation is the stored value of the units activation at the previous time slice, and gamma is a weighting parameter which is proportional to the size of the sum squared error. This is included to prevent the smoothness constraint ultimately competing with the task demands. This smoothness constraint was constructed from information provided in Jordan (1990). Thus the final form of the error term for each output unit is:

$$error[i] = (target[t] - activation[i]) + smoothness[i]; \qquad (7.3)$$

In the first version of the network, without the smoothness constraint, no error is backpropagated for the unselected effector in any one time step. When the smoothness constraint is incorporated, the smoothness term only provides the error signal.

Results

The results of comparing the two versions of the network are illustrated below in figure 7.4. It can be seen that when μ alone is used (black squares), the network exhibits good anticipatory behaviour in the unselected effector. At time step 2 the LH effector is very close to its next target, (target 3) and at time step 3 the RH effector is midway between its previous target, (target 2), and the subsequent target (target 4).

By comparison with the condition in which the smoothness constraint is applied (stippled squares), it can be seen that there is very little difference between the two systems.
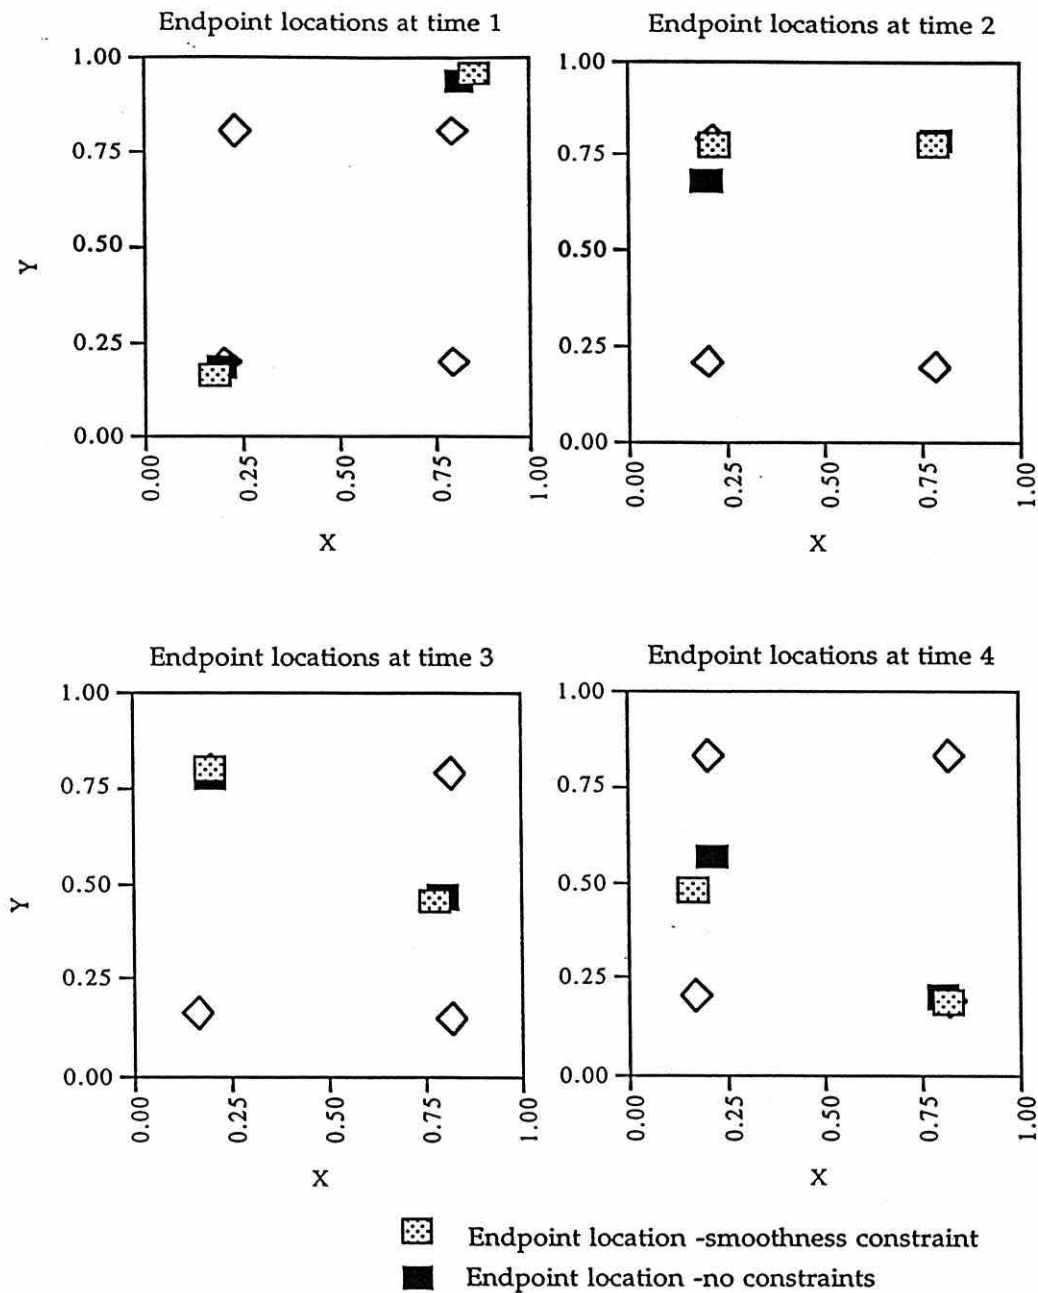
Figure 7.4. *Performance of the network with smoothness constraint (stippled squares) and without*

smoothness constraint (black squares)

Discussion

This simulation shows that in this task there is very little extra anticipatory behaviour to be gained by deliberately incorporating an additional internal constraint. By increasing the range of contextual information encoded on the state / context units, the μ parameter alone is sufficient to allow anticipatory behaviour.

7.2.2 Simulation 7.3: Comparing the effects of $\mu = 0.5$ and $\mu = 0$ on anticipatory behaviour in sequential processing

Method

In this simulation a further condition is tested in which the network is learned with $\mu = 0$ for comparison with the anticipatory behaviour produced with $\mu = 0.5$
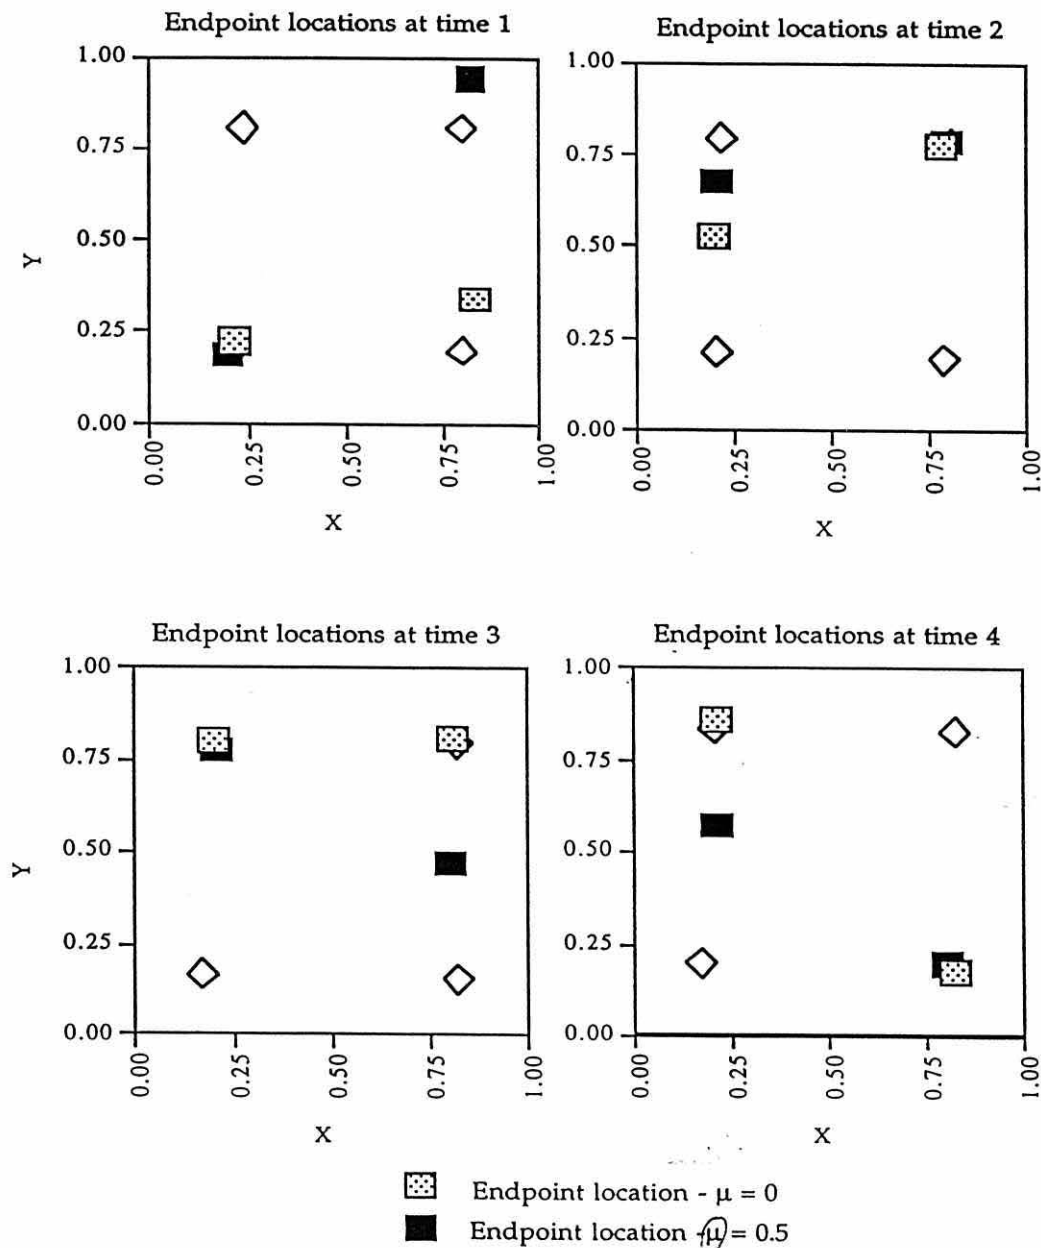
Figure 7.5. *Comparison of network performance with $\mu = 0$ (stippled squares) and $\mu = 0.5$ (black squares)*

Results

The results are depicted in figure 7.5. Here the black squares indicate effector locations at each time step with $\mu = 0.5$ and the stippled squares indicate effector locations with $\mu = 0$. The figure shows that although some anticipation is evident at time step 2 in the LH effector with $\mu = 0$, this behaviour is abolished at time steps 3 and 4, where the unselected effector remains at the previous target location.

Discussion

In sum the two simulations presented in this section demonstrate the importance of the contextual properties invested in the state vector by using μ with a value greater than zero in providing suitable activations for an unselected effector during sequential performance. This behaviour is not enhanced, in this case, by incorporating additional internal constraints.

For completeness, we tested a further condition in which the smoothness constraints are used, but μ is set to zero, in order to test the contribution of the smoothness constraint alone. This test produced a somewhat surprising result in which the network failed to learn and the effectors converged on the fixed points shown below in figure 7.6, for all time steps.
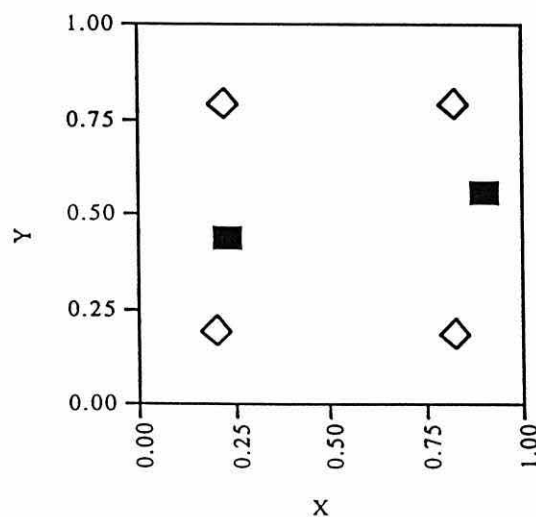
Figure 7.6. *'Frozen' state of the system when the smoothness constraint is combined with $\mu = 0$*

The static state of the network depicted in figure 7.6 occured, presumably as a result of the smoothness constraint conflicting with the task demands. Since the value of gamma was identical to simulation 7.3 , it may be concluded that without the remote contextual encoding afforded by μ, each effector was forced to settle at a stable intermediate state.

This finding further supports the conclusions reached in chapter six about the importance of internal contextual information in preventing freezing of the system due to multiple conflicting internal constraints, as well as in promoting fluid sequential behaviour.

7.3 Contextual representation without recurrent connections

One of the central arguments underlying the criticisms of the Jordan architecture as a chaining system lies in the recurrent connections from the output units to the state units, thus providing the input at the next time step. It is this assumption that makes the apparent similarity between the sequential network and a chaining system so great.

However, if the state units were to encode an independent moving context vector, and it could be shown that similar behavioural properties still hold, the chaining criticisms would need to be questioned.

Jordan (1992) reports a version of the network in which the state units have internal oscillatory behaviour entirely independent of the recurrent input from the output units. In the following simulation we implement a similar idea in which the state units act as simple oscillators and the recurrent connections from the output units are removed. The recurrent connections of the state units onto themselves are left intact with $\mu > 0$. This yields a system in which the sequential behaviour is entirely determined by an independent dynamical context vector.

7.3.1 Simulation 7.4: Examining the dependence of the network on recurrent connections

in producing anticipatory behaviour

Method

The only difference between the oscillator version of the sequential network used here and the normal architecture lies in the set up of the state unit interconnections. The scheme we have used is illustrated in the diagram of figure 7.7 below:
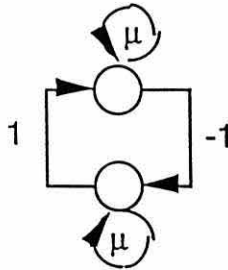
Figure 7.7. *Interconnections between one pair of oscillatory state units.*

The four state units are divided into two pairs. As shown in the diagram, each pair has interconnections between the units with weights of 1 and -1. At each time step the state unit activations are updated according to the following scheme:

$$\text{activation}[i] = (\text{activation}[j] \times \text{weight}[i,j]) + \mu(\text{activation}[i]) \qquad (7.4)$$

In each pair, one unit is initialised with an activation of 1 and the other with an activation of 0, to start the oscillatory behaviour. The value of μ in this simulation was 0.5.

Results

The results of this simulation are compared with previous results on this task with recurrent connections and $\mu = 0.5$. The comparable effects on the behaviour of the network are illustrated in figure 7.8. The figure shows that although the performance of the network with independent state units is quite acceptable on this task, the behaviour of the unselected effector is not as smooth as in the condition with recurrent connections.

In an effort to improve the performance of the network we conducted a second simulation with the value of μ increased to 0.9. The results of this simulation are depicted

in figure 7.9. The results indicate that in this case reliable anticipatory behaviour in the unselected limb is produced.
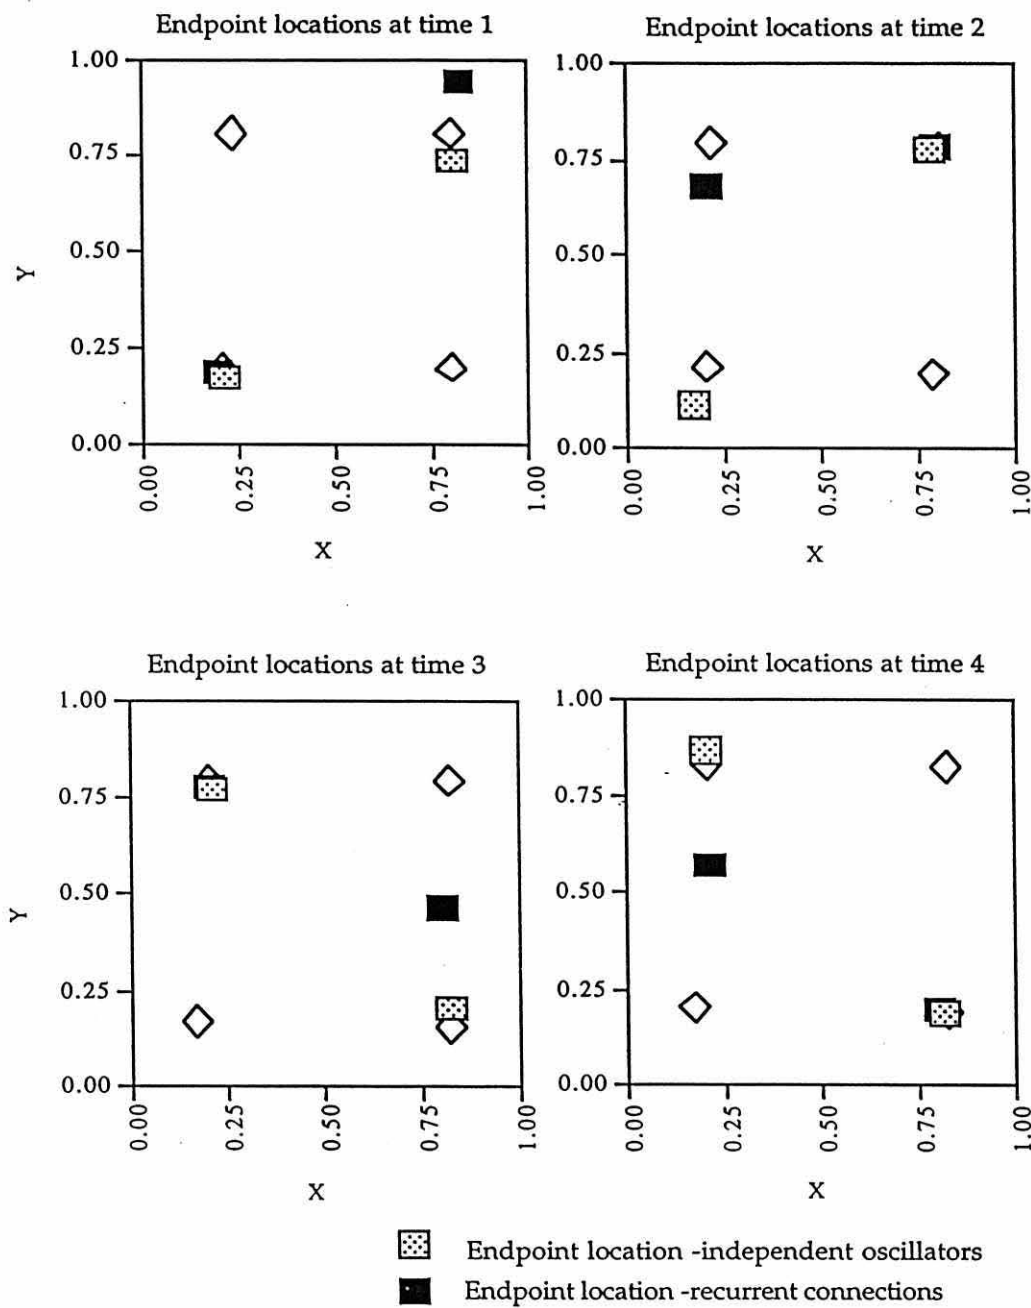
Figure 7.8. *Comparison of network performance with the state units functioning as independent oscillators (stippled squares) and using recurrent connections (black squares)*
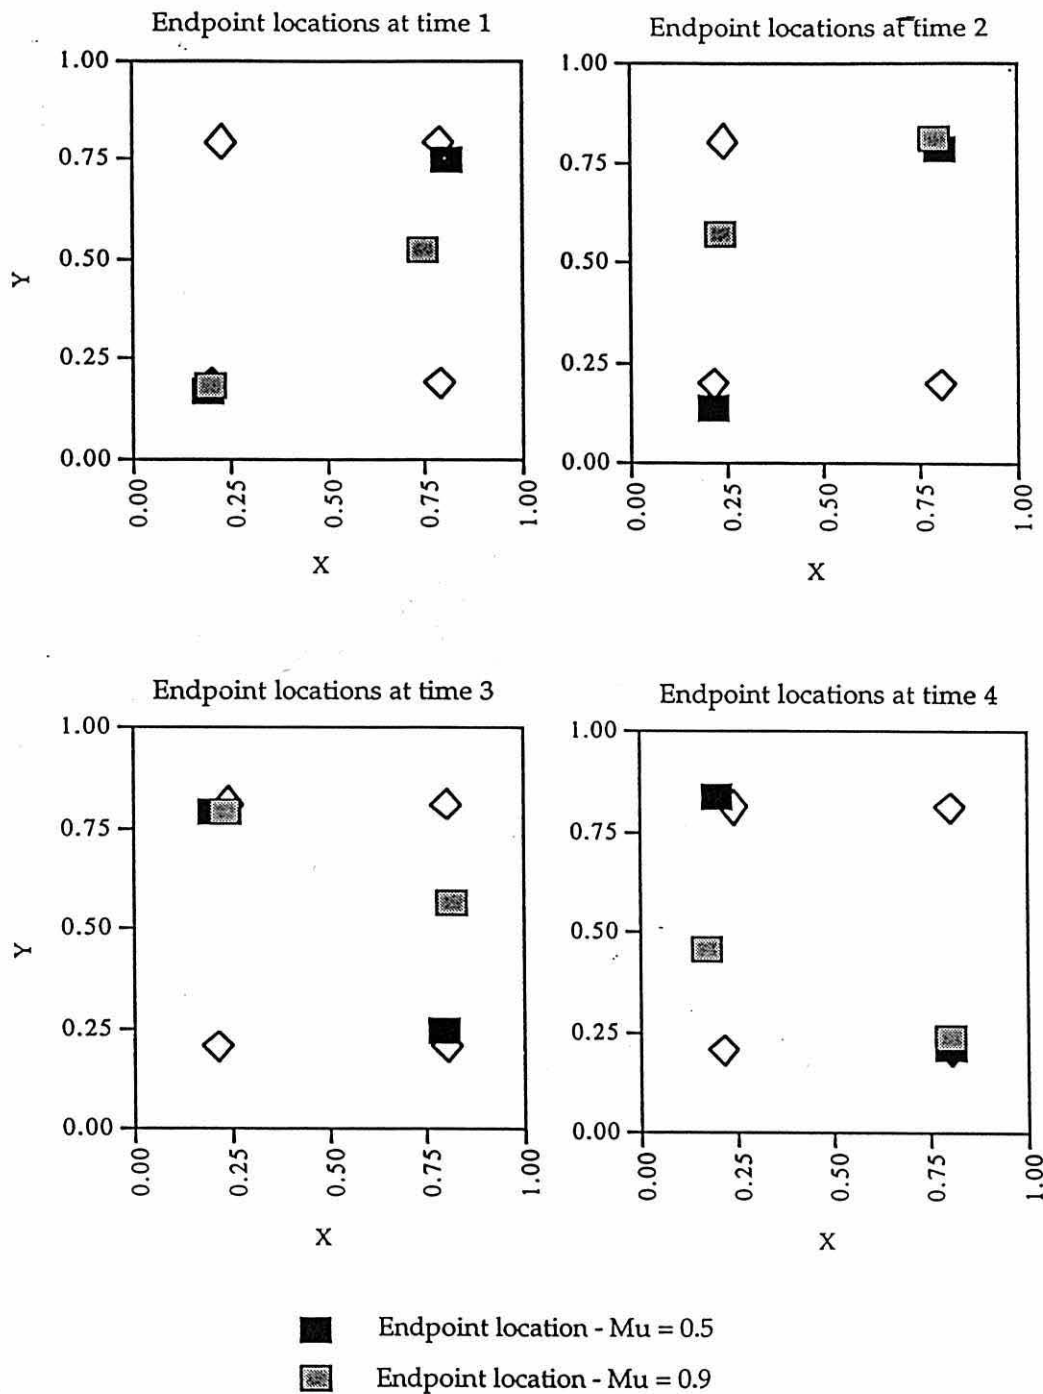
Figure 7.9. *Comparison of network performance using oscillatory state units with* $\mu = 0.5$ *and* $\mu = 0.9$

In a final comparison between the two versions of the sequential architecture implemented here, we recorded the learning profiles of both versions for the different values of μ used in the simulations. These data are illustrated in figure 7.10. Comparison of the two graphs with $\mu = 0.5$ reveals that learning is faster with recurrent connections (1000 iterations as opposed to 2000), although the learning curve is smoother with oscillators. When μ is

increased to 0.9 in the second case, however, the learning time is reduced dramatically, indicating a further benefit of a higher degree of contextual encoding.

recurrent connections to context units

oscillatory context units
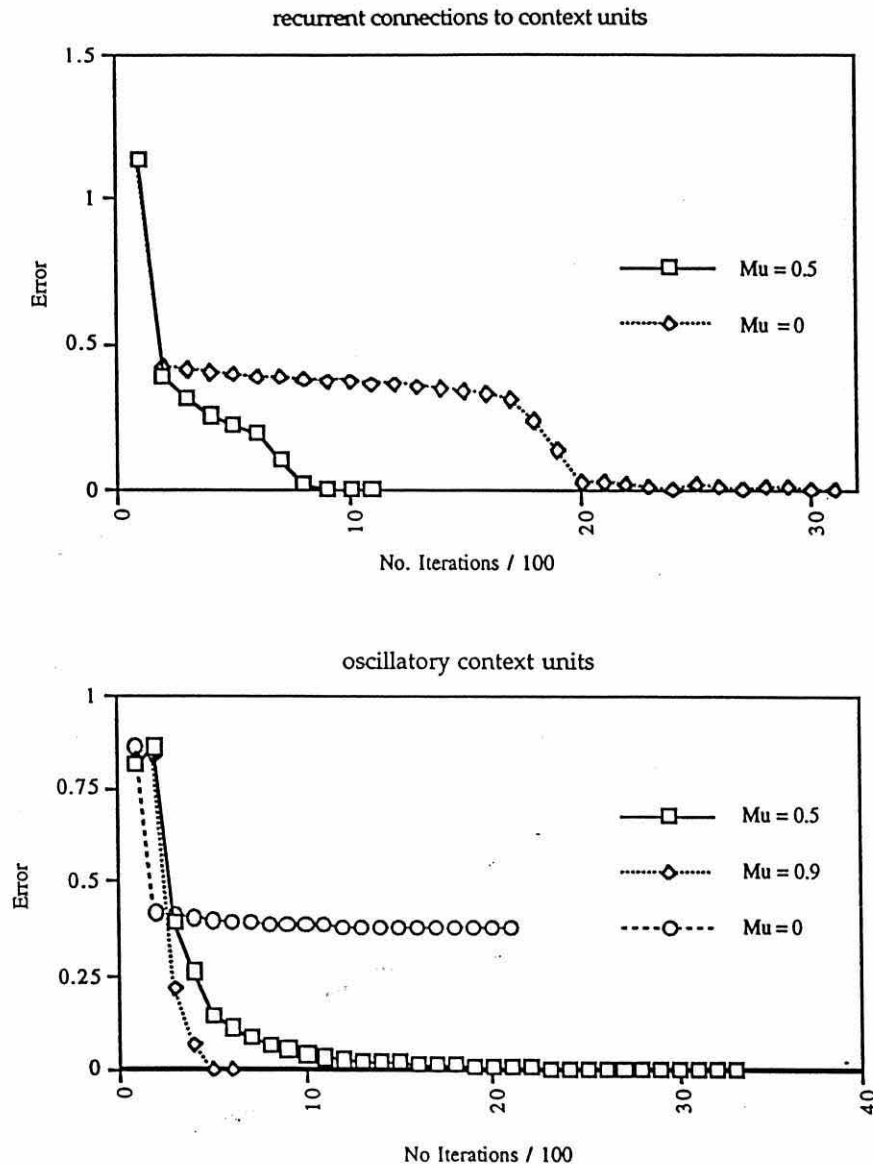
Figure 7.10. *Learning in the network with and without recurrent connections for different values of* μ

Finally, it is notable that with $\mu = 0$ in the case where the state units are oscillators, the network fails to learn and the error score asymptotes at approximately 0.4. The likely reason for this is that with $\mu = 0$ each state unit pair operates as a simple symmetric anti-phase oscillator, which repeats itself every second time step, and as a consequence valuable discriminating information is reduced. This suggests that if oscillators are to be used for encoding contextual information, it is essential that repetition is avoided and that a more

complex vector will have better discrimination properties over long time periods. The points raised by this and the previous simulations are discussed next.

7.4 Chapter summary and discussion

The work reported in this chapter was conducted primarily to supplement and clarify the results reported in chapters five and six, regarding the operation of Jordan sequential networks and their use as a basis for modelling an action control system. A number of simulations were conducted using two sets of output units to represent the 'virtual' endpoints of two effectors in a simplified sequential network architecture in which the forward model was omitted. This simplification was justified on the grounds that the points of interest related to the sequential subnetwork and were not dependent on the forward model.

7.4.1 Dual effector interference: The need for inhibitory mechanisms

In simulation 7.1 'don't care' constraints were applied to the unselected effector in two versions of a sequential positioning task. In one task the only sequence was learned by the selected effector during testing, but in the other the unselected effector had learned a similar sequence. The effects of the don't care constraint, which is central to Jordan's account of dual effector coarticulation in this network, was to trace its target sequence during testing. This result suggests that extra inhibitory mechanisms would need to be included in the model to prevent unwanted synergetic movement of limbs during task performance. Several features of the breakdown of motor behaviour in basal ganglia function are indicative of a failure of inhibitory / disinhibitory mechanisms (see chapter 2).

The results of this simulation suggest that the tasks to which this model has been applied, both in Jordan's own papers and in the simulations of chapters four and five, do not adequately capture the need for inhibitory mechanisms in human serial processing. Furthermore the style of processing in the Jordan model is not suited to addressing this issue. Various other architectures, notably the CQ class of models discussed in chapter three take the problem of parallel response competition as a theoretical primitive built into

the design of the model and would be better equipped to address this problem in the context of basal ganglia control. One advantage of the Jordan architecture over the computationally much simpler CQ models is that it does serve to demonstrate, as we showed in chapter five, the sort of interference problems that can arise in a fully connected massively parallel and distributed architecture.

7.4.2 The state units and representation of temporal context

In previous simulation chapters we have explored the role of the state units as a contextual control signal in both the original Jordan network and the cascaded version presented in chapter five. In doing so we have examined the response of the network to various types of disruption which affect the sequential control afforded by the state units. The key properties of the state vector in these simulations were shown to be related to the use of the parameter μ. It was argued, after Jordan (1986a) that μ could promote coarticulation in the model in a similar fashion to an explicit smoothness constraint applied during learning.

In simulation 7.2 we compared the coarticulatory effects produced by using μ alone and μ with an additional smoothness constraint in the dual effector model. It was found that the smoothness constraint added very little to the performance of the network. In contrast, if μ is reduced to 0 as in simulation 7.3 the coarticulatory effects are banished. Furthermore when the smoothness constraint alone was used, the network failed to learn the sequence due to a failure to resolve conflicting task and smoothness criteria and the network lapsed into a stable intermediate state between the two sets of spatial targets. These results support our view that the use of μ to represent extended contextual information over the state units is an essential part of the psychologically interesting behaviour of the network.

Throughout the course of this thesis we have considered that an important component of a theory of basal ganglia control of action selection and sequencing is the internal representation of temporal context as a means of linking internal control processes to the structure and dynamics of the environment when the information is either not available from the environment or when the movement must be prepared in advance. Furthermore we have put forward the suggestion that portions of the 'frontal' loop through

the basal ganglia may be involved in constructing and maintaining a contextual representation. The simulations performed in this and previous chapters suggest that the theme of internal representation of temporal context is worth pursuing in further modelling studies.

7.4.3 The use of recurrent links and chaining criticisms

A central feature of the Jordan network are the recurrent links from the output units to the state units. Recent critics have argued that the use of recurrent links to produce serial behaviour in the model render it susceptible to the limitations of a chaining system (Brown, Preece & Hulme, 1996; Houghton & Hartley, 1995). Jordan (1985) suggested an alternative to using recurrent links which involved independent oscillatory dynamics over the state units. Jordan (1992) used both oscillatory dynamics and recurrent links to improve the learning performance of the network.

In simulation 7.4 we set up an alternative version of the network used in the previous simulations which has no recurrent links and in which the state units are driven by independent oscillatory dynamics. We compared the performance of the two versions of the network to investigate whether the coarticulatory effects in the unselected effector were still present without recurrent connections. The simulation demonstrated that the same effects were present, although to a reduced degree. In a further simulation we increased the value of μ to examine whether stronger contextual input could improve performance. This was shown to be the case (see figure 7.9). Finally we examined learning performance in both versions of the network and found that, with comparable values of μ ($\mu = 0.5$), the network with recurrent connections learned faster. However, learning performance could be improved in both versions by increasing the value of μ.

This simulation does not address the criticisms directed at the model on the basis that it is unlikely to produce the sorts of errors that are found in human linguistic serial behaviour. It does, however demonstrate that the production of sequences showing coarticulatory effects are not dependent on the existence of recurrent connections, thus weakening the chaining argument against the architecture. The source of dissatisfaction

with Jordan networks extend beyond the use of recurrent links alone. These limitations are discussed in the following section.

7.4.4 Further limitations of the Jordan architecture

The criticisms to which we refer are rooted in the use of the multi-layer perceptron architecture to implement the Jordan model and the backpropagation learning algorithm for forming associations. Together these two components form an extremely powerful and complex computational system which is in part why it has been used so extensively to model different aspects of cognitive function. There are various ways in which using this system to model sequential behaviour and motor control are unsatisfactory.

First, many simple sequences, such as keypresses or pointing to targets, can be learned in a single trial. It is impossible to mimic this sort of learning using gradient descent methods, which typically take several hundred or thousands of iterations before a satisfactory error score is achieved. Furthermore if additional sequences are to be learned performance on the already learned sequences is disrupted as the knowledge encoded in the weights of the network has to be reorganised to accommodate the new sequence. This difference suggests that the way the model encodes and stores information about the sequences it learns is fundamentally different from the way simple sequences are encoded and stored in human performance of serial tasks. Whereas Jordan has expanded on the basic architecture in subsequent papers to produce yet more complex systems with the aim of accounting for more aspects of motor behaviour, the implication of the above criticism is that a simpler associative mechanism, which can form associations in a single trial, may yield more information because the system as a whole is easier to analyse. We develop an alternative approach to meet this criticism in the following chapters.

7.4.5 Summary

In summary, we have demonstrated the crucial importance of the contextual representation over the state units in controlling sequential behaviour in different versions of the Jordan network. Also we have shown that disruptive manipulations that weaken the contextual

representation produce a variety of interference effects both amongst sequence elements and amongst different stored sequences. In this chapter we have additionally demonstrated impairments in dual effector performance. The impairments which are produced in the behaviour of the model simulate a number of features associated with parkinsonism such as reductions, deletions and de-coupling of sequence elements.

There are a number of desirable features of a model of cortico-basal ganglia interaction which are not present in the networks explored thus far. These are:

1. An internal representation of temporal context which is independent of previous outputs of the network and which has sufficiently extended endogenous dynamics to represent context into the arbitrarily distant past to avoid problems with repetitions.

2. Inhibitory/ disinhibitory mechanisms which resolve response competition whilst maintaining coarticulatory behaviour.

3. A simple associative mechanism that is capable of single-trial learning.

These issues are addressed in the following chapters.

CHAPTER EIGHT

Endogenous Contextual Control of Sequential Behaviour

8.1 Introduction

In the material presented in this chapter our aim is to capture and illustrate some high-level computational principles involved in the generation of sequential action using a complex contextual signal with endogenous dynamics, and to show how these computational considerations support the notion of complementary roles for an interacting frontal system and basal ganglia complex presented in chapter 4.

The generic architecture which we use as a basis for the work reported in this chapter is the OSCAR (Oscillator based associative recall) model which has been developed by Brown, Preece and Hulme (1996) to simulate various aspects of serial short term memory. This architecture belongs to the class of item-to-context models reviewed in chapter three and bears certain similarities to the CQ architecture which was described in some detail. The model provides a suitable contrast to the Jordan model we have explored in previous chapters. In particular, the model uses simple Hebbian association to perform the association between sequence elements and states of the context signal. This strips away the in-built complexity introduced by using hidden units and gradient descent learning. A second advantage of this approach is that the interaction between contextual control of sequencing and competitive inhibitory processes may be explicitly explored. Although the basic OSCAR architecture does not implement a competitive component, it is a relatively simple task to introduce competitive processing to the model.

The use of OSCAR has advantages over CQ approaches which are entirely dependent on inhibitory processing due to the simple nature of the control signal used in

these models. This means that the trade-off between inhibitory and contextual modes of control cannot be effectively explored using a CQ architecture. Finally, the OSCAR model makes use of distributed representations both of "action" schemata and states of the dynamic signal. This has the potential to yield more interesting behaviour of the model in response to disruption than the use of simple local representations as in the CQ models.

In the first part of this chapter, we use ecological considerations to motivate the construction of an appropriate endogenous dynamic contextual signal. It is shown that a more complex, multi-dimensional dynamic context signal can be constructed from a range of oscillators that cover a wide range of frequencies, and such a signal receives independent motivation from studies of timing behaviour in humans and animals.

8.1.1 Oscillators and representation of temporal context

In this section we consider the question of how an independent internal, contextual signal to control internally generated sequential behaviour could plausibly be constructed. In doing so we shall briefly reiterate some of the points that have emerged in the course of this thesis.

In chapter four we developed an argument that the behaviour of simple organisms and higher order motor control differ in important respects. It is easy to see how an organism can achieve sequential behaviour if it is merely responding and selecting actions in response to events in the external world. The dynamic structure of event sequences in the world will drive the sequential behaviour of the organism directly - the organism needs only a disinhibition mechanism of the type discussed in chapter four and in Lorenz (1977) in order to respond sequentially to a succession of perceived events, without the need for any internal dynamic representation.

However if the internal generation of sequential action is to be achieved, the organism must have access to some form of internal dynamics that can be used to drive internally generated sequential action. As we have suggested, it is likely that the form of such an internal dynamics will have evolved to serve the ability of the organism to become

attuned to temporal regularities in the environment (such as, for example, the periodic availability of food sources).

There is thus a crucial connection between the ability to generate *sequential* behaviour, and the *internal* contextual control of action. This aspect of internal control is not satisfactorily explained by motor programming models as discussed in chapter 3. Yet, if we are to specify the nature of internal control structures within the framework of alternative dynamical approaches to movement control this is an important step. To make use of the ecological principles of the dynamical systems approach we need to explain how temporal regularities in the world can be represented internally. Our next aim, therefore, is to consider how an appropriate internal dynamic contextual representation could be provided.

An alternative approach to the recurrent networks that we considered in previous chapters is to generate an endogenous dynamic signal using oscillators. We illustrate the intuitions underlying this approach with an example. Many simple organisms, such as wasps, are able to arrive at the same time of day at a known food location, and the distribution of arrival times has a peak at the actual time of availability of a food source (see e.g. Gallistel, 1990). One obvious way to build a system which is capable of this behaviour would be based on an internal oscillator, which repeated its output every 24-hours. Thus, the time of food availability at a particular location can be associated with a unique state of the internal oscillator, which subsequently acts as a representation that can guide the wasps foraging behaviour (cf. Gallistel, 1990).

Sequential behaviour could emerge from an oscillator controlled system of this type by assuming that successive actions in a sequence are associated with successive states of the oscillator. Suppose for example that another energy source is readily available at a different location at say midday, and a third food source is available elsewhere in the early evening. If appropriate associations are formed between representations of the food locations and corresponding states of the internal oscillator, the wasp will display basic sequential behaviour as it flies from one location to the next, cued by the internal oscillator.

8.1.2 Oscillators and the control of sequential behaviour

The above example shows how an internal oscillator could provide a basis for the control of sequential action. With regard to the internal generation of motor behaviour in particular, oscillators can be seen as an 'evolutionary primitive', in that the most primitive forms of swimming and locomotion rely on simple oscillatory control (e.g. Kopell & Ermentrout, 1990). Von Holst (1939 / 1973) suggested that co-ordinated action patterns arise from the coupled interaction of internal oscillators. The oscillatory mechanisms Holst postulated are now known to operate as spinal neural circuits and are known as central pattern generators (CPGs). These control the rhythmic outputs of the motor neurons. The CPGs themselves are under the influence of descending projections from higher motor centres. Activity at the level of the CPG is further modulated by sensory inputs from proprioceptors or the environment.

Since the work of von Holst oscillatory mechanisms have been discovered at higher levels of the motor system and applied to many motor functions. Weakly coupled oscillators form a central part of Bernstein's (1967) concept of motor *synergies* which he puts forward as the prime way the motor system is constrained so that disparate limbs act in concert to achieve an external goal, thus reducing the excess degrees of freedom. Building on Bernsteins ideas, proponents of the action systems approach have proposed that complex motor patterns may be represented by systems of interacting oscillators, (see chapters 3 and 6). Our work in this chapter is related to these ideas but our interest is not directed so much at the representations of the motor patterns themselves. Rather we are concerned with the method used by the brain for extracting and coding the salient information from the environment to which the motor patterns are bound.

8.1.3 Computational considerations for oscillator based control of sequencing

We have discussed at an intuitive level ways in which a single internal oscillator might provide the basis for the learning and production of simple sequential behaviour. We now discuss limitations of computational systems based on a simple oscillator and describe how a more complex model, which is based on an array of many oscillators, can be constructed.

An obvious limitation of a system based on a single oscillator such as the one we described earlier is that the time scale over which the oscillator is able to represent event sequences is limited by the period of the oscillator. In the example of the wasp described above, if the wasp possesses only an oscillator of frequency 24 hours, then it will not be able to represent sequences that span a longer period. It could not, for example, represent the fact that a source of food was available at the same time every second or third day. The constraint here is known as the "non-repetition requirement" (Brown, Preece & Hulme, 1996). This states that the oscillator-based internal dynamic contextual signal, to which external event sequences may become associated, must not repeat itself within the time period over which regularities may need to be represented. It is clear that such limitations are found in nature - Wahl (1932), for example, found that bees could not adjust to a 48-hour schedule of sugar water availability, in that the bees flew to the site just as frequently on test days one and three, as on test days two and four.

A further limitation is that while a single oscillator will in one sense provide too "fast-changing" a context signal, and hence will repeat itself as described above, a single oscillator system will also be too slow-moving, in that successive states of it will be too similar to one another to allow events that occur nearby in time to be adequately discriminated. For example, the state of the oscillator at 8.30 am will be very similar to its state at 8.45 am, because the oscillator will have only moved through a small fraction of its total cycle in the intervening time period. A further complication, if the output of the oscillator is assumed to be sinusoidal in nature, is that the output of the oscillator will change at varying rates at different points in the oscillator's cycle. This leads to a second requirement, which is the "temporal discrimination requirement". This constraint states that nearby points in time must be able to be represented by states of the internal oscillator-based system that are sufficiently distinct from one another to enable different action representations to be associated with different states of the dynamic internal context. We have already become familiar with the need for adequate discrimination between states of the context signal in our simulations with the Jordan networks in previous chapters. We

have examined in detail the effects of poor temporal discrimination on the behaviour of models which used recurrent connections to form the context signal.

In the Jordan network with oscillator based state units used in chapter seven a scheme was adopted using two identical oscillators starting at opposite positions in phase space. Houghton's competitive queuing model is very similar, and again uses a two-dimensional control signal which is composed of so-called start and end nodes, which was discussed in chapter three. At the start of a sequence the start node starts to decay from a maximum value (e.g. 1) towards zero, meanwhile the end node increases its value at a rate identical to the decrease in activation of the start node.

Both of the above examples are instantiations of the simplest functionally useful dynamic control signal. As such, there are considerable limitations on the ability of a network driven by a simple signal to correctly discriminate the serial order of sequence elements. The Jordan network requires addition of self -recurrent connections on the state units to add higher dimensionality to the control signal; the CQ model requires additional competitive processing to effectively control serial behaviour. In both cases the action of the oscillators alone is not sufficient to meet both requirements.

Intuitively, one solution to this problem would be to adopt a system which incorporates both fast and slow oscillators. The fast-moving (or high frequency) oscillators will change their state substantially over short periods of time, and so the overall state of the system will change sufficiently quickly for the temporal discrimination requirement. The slow-moving (or low frequency) oscillators can be assigned a sufficiently low frequency such that they do not repeat themselves over any arbitrary time interval. This will then satisfy the non-repetition requirement.

A recent modelling approach which uses this method was produced by Church and Broadbent (1990). Although there is not space to review this model in detail here, it is sufficient to note that the model has been criticised on at least two counts which are relevant to the current discussion. First, as Brown, Preece & Hulme (1996) have commented, the model does not meet the 'continuity' requirement to which we have previously referred when discussing the Jordan model. This is the basic requirement that states close together

in time are more similar than states far apart in time. This flaw provides part of the motivation for OSCAR in which it is possible to correct this problem by adopting a distributed approach, in which each element of a vector representing the internal state of the organism is determined by many different oscillators. A second criticism forwarded by Weardon, (1994) is that each oscillator is crucial to the performance of the whole model, thus the model is fragile in the face of damage or disruptions to individual oscillators.

8.2 A description of OSCAR

A central assumption of OSCAR is that an array of oscillators, running at different frequencies and phases, combine to make up the dynamic contextual signal to which items in a sequence become attached. It is assumed that oscillators are neurally implemented by means of time varying activity either at the level of synaptic circuits within individual neurons or neural circuits consisting of ensembles of neurons.

8.2.1 Composition of the context vector

Each element of the context vector is made up of the output of several different oscillators, with different periodicities.
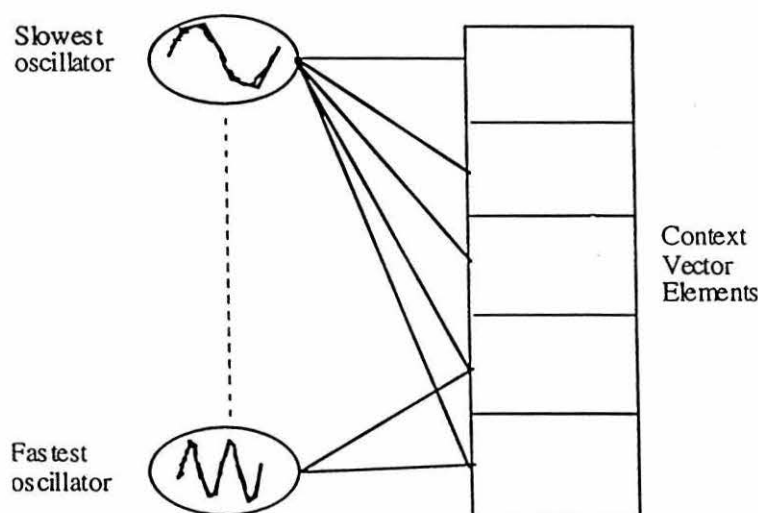
Figure 8.1. *Connections between an array of oscillators and elements of a context vector*

A simplifying assumption of the model is that the output of each oscillator can be represented in terms of a sinusoidally time-varying output. In the version of the model that we describe here there are 15 different oscillators which combine to make up a 16-element context signal. An array of such oscillators, connected up to elements of a context vector, is illustrated on the left-hand side of Figure 8.1. One of the criticisms of the Church and Broadbent (1990) model is that all the oscillators are assigned equal weight in constructing the control signal. This is undesirable because the high frequency oscillators tend to repeat themselves. If they are assigned equal weight to the slow oscillators then repeating patterns have too much influence in the construction of the signal as a whole. This can lead to bias in the behaviour of the model.

$$\text{element}(1) = \cos(\text{th}(1)) * \cos(\text{th}(2)) * \cos(\text{th}(3)) * \cos(\text{th}(4));$$

$$\text{element}(2) = \cos(\text{th}(1)) * \cos(\text{th}(2)) * \cos(\text{th}(3)) * \sin(\text{th}(4));$$

$$\text{element}(3) = \cos(\text{th}(1)) * \cos(\text{th}(2)) * \sin(\text{th}(3)) * \cos(\text{th}(5));$$

$$\text{element}(14) = \sin(\text{th}(1)) * \sin(\text{th}(9)) * \cos(\text{th}(13)) * \sin(\text{th}(14));$$

$$\text{element}(15) = \sin(\text{th}(1)) * \sin(\text{th}(9)) * \sin(\text{th}(13)) * \cos(\text{th}(15));$$

$$\text{element}(16) = \sin(\text{th}(1)) * \sin(\text{th}(9)) * \sin(\text{th}(13)) * \sin(\text{th}(15));$$

Figure 8.2. *The construction of some of the context signal elements. from combinations of different oscillators*

The approach used in OSCAR avoids this problem by assuming that slow moving oscillators can contribute to many context vector elements, whilst the faster oscillators contribute to fewer elements. To illustrate this point, Figure 8.2 depicts the construction of

the first 3 and last 3 elements of a 16-element control vector. The various "thetas" are assumed to reflect the outputs of different contributing oscillators; these rotate through different angles in each time slice. They are initialised to different random angles in each replication of a simulation.

It can be seen that th(1) contributes to every element of the context signal, whereas th(15) contributes only to elements 15 and 16. th(2) would contribute to half of the 16 elements, and so on. Thus the oscillator whose output is represented by th(15) therefore contributes much less to the overall value of the signal than the oscillator whose output is represented by th(1).

An important feature of the construction of the signal is that the distribution of sine and cosine components of the signal is chosen so that the resulting control signal is always normalised (i.e the dot product of the context vector with itself is always one).[2] This means that the context signal remains constant in overall "magnitude" but has both slow-moving and fast-moving underlying components. This is exactly what is required to meet the constraints identified above. The fast-moving underlying components serve to ensure adequate separation between states of the signal that occur nearby in time. If these were the only components, then states of the overall control signal would repeat themselves over short time intervals, and the non-repetition requirement not be met. However the use of a sufficient number of low-frequency oscillators, leading to the presence of slow-moving underlying components, means that states of the context signal never repeat themselves provided that the slowest moving oscillators have sufficiently long periods, and contribute sufficiently to the overall context vector.

8.2.2 Similarity relations in the contextual control signal

The similarity relations of such a control signal are illustrated in Figure 8.3. Figure 8.3(a) shows the 3-dimensional surface representing the cosine of each of the 32 states of the context signal that is made up of the outputs of different oscillators with the same frequency but different phases, and Figure 8.4(b) shows an end-on view of the same surface. (Note

---

[2]This scheme is used so that the cosine of the angle between vectors can be used as a single measure of similarity because all vectors are constrained to be of the same length.

that the dot-product and cosine between pairs of vectors are equivalent measures when dealing with normalised vectors). Figure 8.4 shows that vectors have a dot product of 1.0 with themselves, a dot product nearly as high with context vectors that occur near to them in time, but are unrelated to context vectors more than about 28 states, or time-steps, away. This satisfies the "similarity requirement" outlined above, although this particular signal would repeat itself when plotted over a wider range[3]. It can also be seen that the function plotted on the graph is continuous thus satisfying the "continuity requirement".
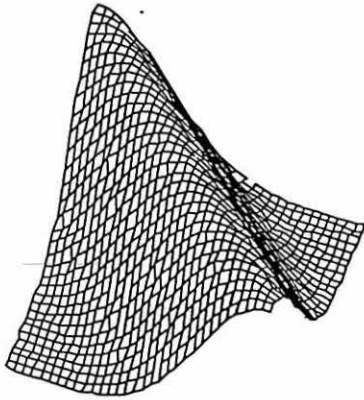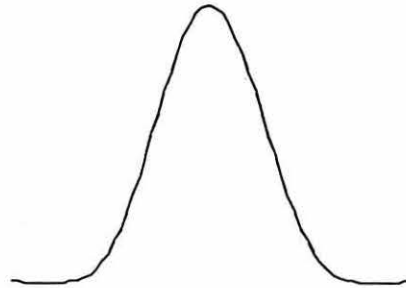
Figure 8.3(a)   Figure 8.3(b)

Figure 8.3. *Similarity relations between 32 successive states of a simple oscillator-based context vector*

A crucial parameter of the context signal is the "sharpness" at the top of the surface, because this represents the distinctiveness of neighbouring context signals. The sharper the peak of the surface, the more distinctive are context signals near to each other in time, and hence the more distinct will be the context-to-action associations for successive items. The sharpness of this part of the surface can be increased by increasing the importance of high-frequency oscillators to low-frequency oscillators in determining the context signal vector.

---

[3] Brown, Preece and Hulme, (1996) have demonstrated sets of contexts that do not repeat themselves.

We illustrate this in Figure 8.5, which shows the equivalent similarity relations for a context signal generated in this way. It will be remembered that the 16-D control signal is made up of the outputs of 15 different oscillators, with oscillator 1 contributing to most elements of the context vector, and oscillator 15 contributing to fewest. In order to capture a reasonable range of periodicities, we now assume that if the period of oscillator 1 is Tau then the period of oscillator n is sqrt(n)*Tau, and that thus the periodicity of oscillator 12 is sqrt(12) * Tau = 3.46*Tau, etc. Here Tau represents the angle through which the sinusoidal output of each oscillator progresses within each time cycle. Thus, for example, oscillator 1 might progress through 2 degrees in each time cycle, in which case oscillator 12 would progress though 6.9 degrees in each time cycle.
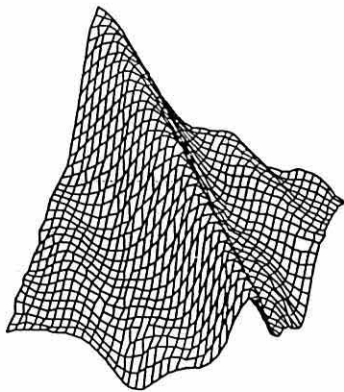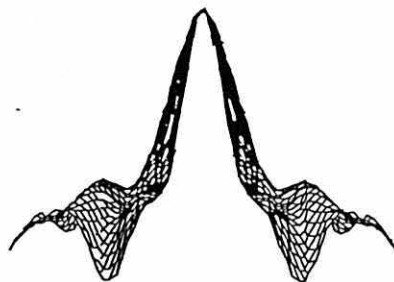
Figure 8.4(a)   Figure 8.4(b)

Figure 8.4. *Similarity relations between 32 successive states of a complex oscillator-based context vector*

A context signal generated in the way we have described above has a number of independent motivations. First of all, the use of oscillators as a fundamental mechanism for driving the sequential behaviour can be motivated by a consideration of the behaviour of relatively simple organisms. There is a vast range of evidence that many different animal behaviours are governed by endogenous oscillators and that the oscillator that becomes

entrained to dynamic aspects of the environment (e.g. circadian rhythms) is particularly important. Secondly, there is considerable evidence for the use of oscillators in the control of a variety of motor patterns in a number of species. In addition some action systems models of motor control are based on systems of coupled oscillators for the formation of motor synergies (e.g. Saltzman and Kelso, 1987).

Thirdly, the fact that each oscillator used in construction of the dynamic control signal has a frequency twice as high as that of its next slowest neighbour can be seen as independently motivated, in the sense that a similar scheme was developed independently to account for the completely different task of interval estimation (Church & Broadbent, 1990). However, it is interesting to note that some studies have found that time estimation and reproduction is abnormal in PD (Pastor, Artieda, Jahanshahi & Obeso, 1992). Although this link is not explored in the current study, it lends indirect support to the central notion that a dynamic signal composed of oscillators may contribute to the role of basal ganglia in motor control and suggests that the framework may be extended to account for a suggested role in timing the onset of movement. This will be the focus of a future study.

We now turn to a more detailed discussion of how this dynamic internal representation can be used to learn to reproduce sequential behaviour.

8.2.3 Sequence learning and production using OSCAR

The basic principle underlying the architecture is that initial sequence learning involves the formation of associations between successive sequence elements and successive states of the time varying context signal. It can be assumed that the dynamic signal effectively represents the context of learning for each sequence element. In other words, each sequence element becomes associated to the state of the context signal that exists when that sequence element is encountered.

For retrieval to take place it is assumed that the time varying context signal can be reproduced. Each successive state of the context signal can be used as a probe to recall successive elements of the learned sequence. These can be thought of as potentially representing either actions in a movement sequence; letters in a word to be spelled or typed,

or phonemes in speech production. The basic methodology involves associating successive items in a sequence with successive states of the control signal. This is illustrated in Figure 8.5 overleaf. In the present model we use simple Hebbian association.[4] Although it is equally possible to use some other method for single-trial learning such as convolution, Hebbian association is the simplest method available. This reduces complications such as those encountered with the Jordan architecture, whereby the behaviour generated by the model could be attributed, at least in part, to properties of the learning procedure.

c1 c2 c3 c4 c5 c6 c7 c8 c9 c10 c11

   | | | | | |

  e1   e2   e3   e4   e5   e6
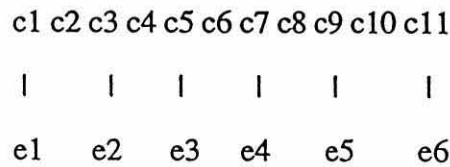
Figure 8.5. *Associations formed between successive items in a six-element sequence and every second state of the context signal*

The contextual separation between each action-context association may be altered by varying the number of time steps (i.e. the number of steps through the temporally-changing context signal). Thus while Figure 8.5 illustrates the formation of associations between action vectors and every second context state, this interval can be varied at will. The size of the interval between contextual states that are used in the association can be thought of as a *distinctiveness* parameter. Following Brown, Preece and Hulme (1996) we henceforth denote this parameter as **D**. The value of **D** determines the distinctiveness or 'quality' of the context signal .

8.3 Contextual Control of Sequential behaviour

In the remainder of this chapter we present a series of simulations which demonstrate the basic properties of OSCAR in the control of sequential behaviour. In the first simulation we aim to show how an appropriately constructed context vector can effectively control a

---

[4] The weight change matrix is calculated as the outer product of the vector elements to be associated.

sequence of action schemata using only simple Hebbian learning to perform the associations. In the second simulation we examine the effects of reducing the quality of the context signal on the ability of the system to discriminate between sequence elements. Finally we examine the effects of introducing a small amount of noise into the processing of the system.

8.3.1 Simulation 8.1: Sequential behaviour in the model

<u>Introduction</u>

The aim of this first simulation is simply to illustrate the ability of the dynamic, time-varying contextual control signal to allow the activations of different candidate actions to vary over time. Here, therefore, we adopt a simple approach, in which the model is required to learn one sequence of six uncorrelated actions, and in which it is assumed that the system is noiseless such that the context of learning can be perfectly reinstated after learning. Although these assumptions undoubtedly represent simplifications, a straightforward example provides the best method of understanding the basic behaviour of the model.

<u>Method</u>

The version of OSCAR described above, in Section 8.2, was used, with a 'medium speed' context with D set to 4. Each state of the dynamic context signal becomes associated, via simple Hebbian association, to successive actions in a six-element sequence. Each "action" is represented as a normalised vector of 16 elements. For simplicity, the action vector has the same number of elements as the context vector.

Thus the simulation proceeds through the following phases. First, six successive states of the contextual control signal are generated, each being a 16-element vector. Because of the D value of 4 successive states of the context signal with have an intermediate level of similarity to each other. The six "action vectors" will on average be uncorrelated with one another. Importantly a context signal such as that illustrated in figure 8.5 exhibits the desired continuity property that states of the context signal that are nearby in time are more similar to each other to each other than states further apart in time. This

has the consequence that successive context vectors represent states of a complex dynamic signal which evolves continuously through time.

When the context vectors and action vectors have been generated, they are associated with one another using simple Hebbian association. The result of this is a matrix in which are stored associations between Action 1 and Context State 1, through Action n and Context State n, to Action 6 and Context State 6. During learned performance of the sequence, it is assumed that successive states of the context can be reproduced accurately, given the first one. Thus to retrieve the first action vector in the sequence, Context State 1 is used as a probe; to retrieve the second action vector, Context State 2 is used as a probe, and so on.

Simulations were run using this procedure, with a different, randomly generated set of action vectors and also a different, newly-generated set of context vectors for each simulation. Because Context State 1 is similar to Context State 2, the action vector that is retrieved when Context State 1 is used as a probe will (if the model is functioning as expected) be highly similar to Action 1 (the first target) but also somewhat similar to Action 2, this property leads to the desired property of parallelism in learned sequence performance.

We examined these relations by computing the similarity (dot product) between the action vector that is retrieved at each step of the sequence retrieval process, and comparing it with the target (i.e. correct) action vector and also with the other 5 action vectors in the six-action sequence.

Results

The results are illustrated in Figure 8.6, overleaf. Each of the six panels in the figure represents a successive time step during sequence production, with one panel for each action-retrieval.
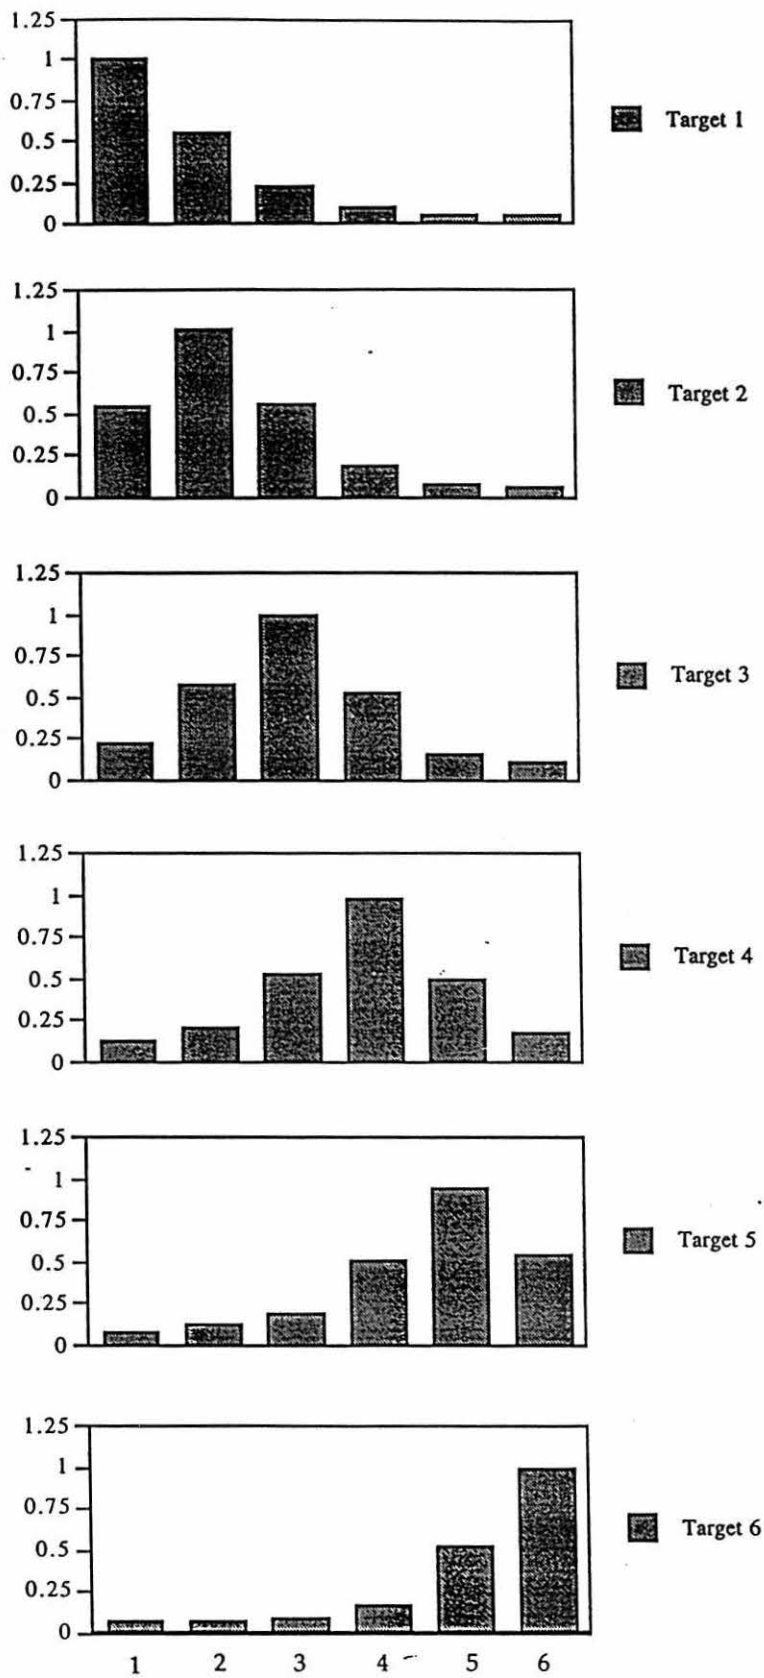
Figure 8.6. *Retrieval of a sequence of actions using a medium quality context signal*

Thus, for example, the top panel (labelled "Target 1") indicates that the first retrieved action vector was most similar to the target, i.e. Action 1, having a dot product of 1. It was next most similar to the neighbouring Action 2 (about 0.55). It was less similar to Actions 3 and 4 (about 0.25 and 0.15 respectively) and almost uncorrelated with the most distant Actions 5 and 6. In psychological terms, this may be considered as a state of the dynamic contextual control signal leading to the partial activation of several competing actions - with Action 1 being the most highly activated, Action 2 the next most highly activated, and so on.

The second panel shows the similarity relations of the second retrieved target vector. Here we see that the retrieved vector is most similar to Action 2 (the "correct" output), and next most similar to neighbouring Actions 1 and 3. Thus again in this case several potential actions have become activated, with Action 2 being the most highly activated.

An examination of the remaining panels indicates that for all six states of the context signal, the "correct" action is most highly activated with a dot product of 1, although there is a general tendency for actions that are associated with similar states of the dynamic contextual control signal to become partially activated also.

Discussion

This simulation illustrates a fundamental property of the model: that if a sequence of action representations are associated to successive states of a time-varying, oscillator-based contextual signal, it is possible to reproduce the sequence by "replaying" the dynamic context of learning. As each successive state of the context signal is replayed during retrieval, several action representations will be partially activated, with actions that are nearby in the temporal sequence being more highly activated. In general, of course, this is a desirable property: similar actions are likely to be appropriate for similar contexts, and appropriate generalisation is much more likely to be achieved in a system with this kind of property.

This type of architecture also complies with two of the fundamental requirements for modelling sequential motor preparation originally put forward by Lashley, (1951) and since supported by a number of studies (see Houghton & Hartley, 1995 for review):

1. parallel pre-activation of sequence elements.

2. potential competition between similar actions.

In the next simulation we examine the effects of varying the distinctiveness or 'quality' of the dynamic contextual signal.

8.3.2 Simulation 8.2 : Varying the quality of the context signal

Introduction

The purpose of simulation 8.2 was to examine the effects of varying the quality of the contextual control signal. More specifically, we wished to illustrate the property that a contextual control signal that changes more quickly over time, and successive states of which are therefore more distinct from one another, will lead to reduced activation of competing actions. In other words, a faster moving context signal should lead to the target action having a high activation relative to its competitors, whilst a slower-moving signal will be less distinct and may lead to a high activation of the competitors relative to the target action. In relation to the conceptual model of chapter four, this manipulation is analogous to varying the value of tau in the cascaded Jordan network studied in the previous chapter.

As illustrated above, it is a simple matter to increase the frequency of all the oscillators whose output combines to make up the dynamic contextual control signal. This results in a reduced level of similarity between successive states of the control signal, and therefore each can serve as a more discriminatory cue for a particular element of the sequence.

Method

We examined the success of the model in reproducing just the third element of the sequence that was examined in simulation 1. Performance was examined as a function of the size of the jump between contextual states that are associated with each successive sequence element. Thus, if $D = 1$, then each successive contextual state is used in the association. If $D = 2$, every other contextual state is used, and so on. In order to examine the effect of varying distinctiveness on the performance of the model in reproducing the third element of the sequence examined in simulation 1 above, we simply repeated simulation 1 for different

values of **D**. All other details were as reported in simulation 1 (in which the value of **D** was 4).
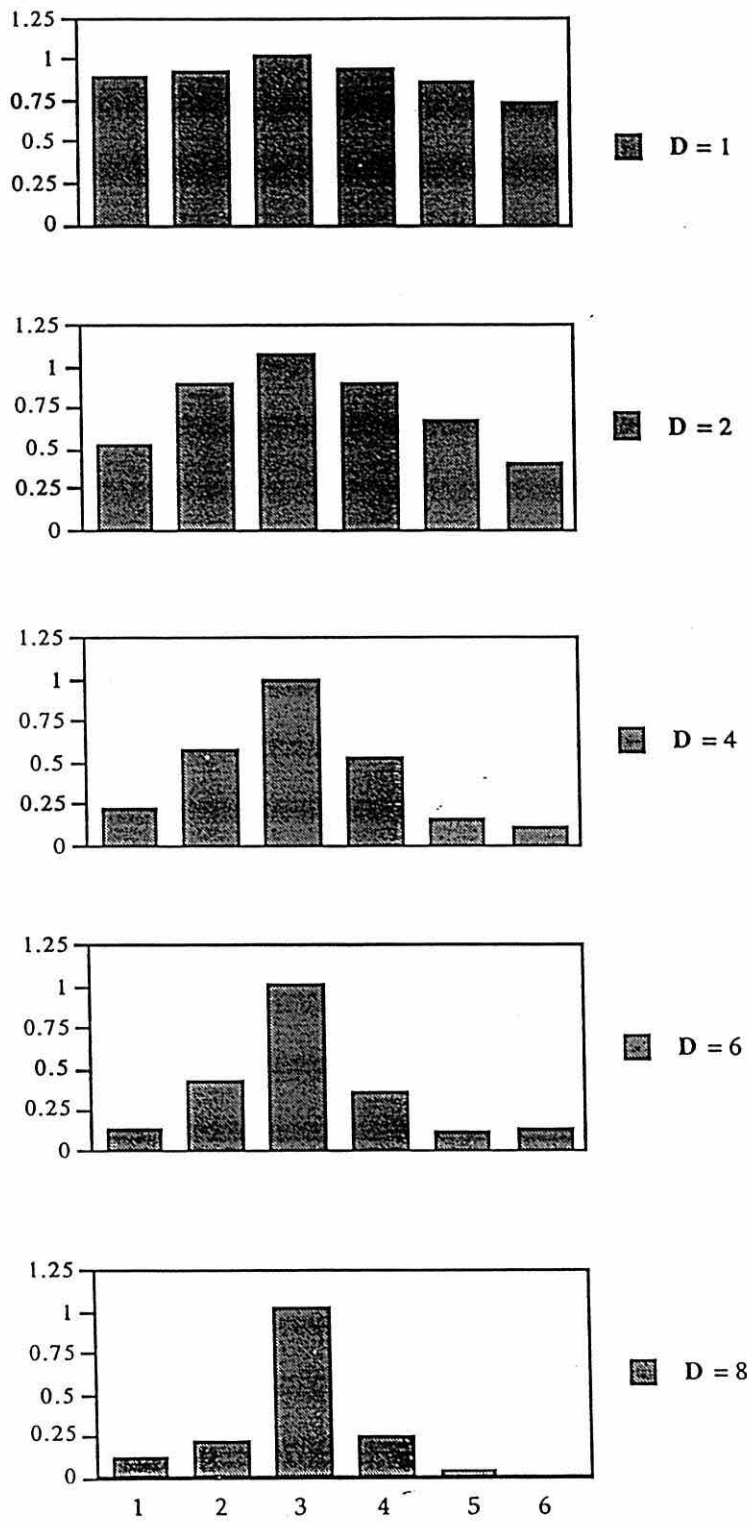
Figure 8.7. *The effects of increasing the distinctiveness of the contextual signal*

Results

The results are illustrated in figure 8.7. Each of the panels in figure 8.7 represents the activation of six possible actions, when the third state of the control signal is used as a probe with which to retrieve the third element of the sequence. Thus it can be seen that the third panel of figure 8.7 is identical to the third panel of figure 8.6 because in both cases the model is attempting to reproduce the third element of the sequence with $D = 4$.

The effects of varying the rate of change of the control signal are clear. For example, in the top panel of figure 8.7, where $D = 1$ and therefore successive states of the context signal are highly similar to one another there is a high degree of activation of elements close in the sequence to the target. More specifically the target element (action 3) is most highly activated (again with a dot product of 1), but elements nearby in the sequence are also very highly activated. In fact all sequence elements have activations above 0.75. In psychological terms, the separation between the target action and competing actions is relatively small and the potential for erroneous responses or lengthened reaction times (RT's) due to contextual interference is very high.

Compare this with the lowest panel in figure 8.7 where $D = 8$. This results in a much greater separation between the target and competing actions, and indeed the two nearest competitors actions 2 and 4 each have an activation value less than 0.25.

Discussion

This simulation serves to illustrate the point that by associating sequence elements with a highly distinct control signal, the level of similarity between nearby actions and the state of the context vector associated with the target action is reduced. By contrast, a less distinct signal increases the similarity relations between actions and the state of the context vector. This property of the model has the following psychological and computational interpretations:

The amount of inherent parallelism in the model during sequence production is variable on a continuum and is determined by the overall speed of the control signal. When the signal is slow the amount of parallelism is increased which will lead to greater fluency and speed during sequence production, but also a greater probability of errors and / or

conflict between successive actions. As the speed of the control signal is increased output becomes more strictly sequential but temporal discrimination is improved.

A comparison may be made at this point with the Jordan networks which were studied in previous chapters. In that system, a similar continuum from high to low parallelism is achieved by the particular combinations of task and internal constraints applied to the network. This scheme is, however, somewhat complicated and artificial since a decision has to be made in advance to what extent individual output units 'care' or 'don't care' about their targets at a particular time step. Moreover, these constraints are unrelated to the control signal provided by the state units. In the current scheme this property is simply achieved by varying a single parameter - the speed or distinctiveness of the control signal.

One unrealistic feature of the model made evident by this simulation is that there is no cost attached to always using a very slow control signal. There is high parallelism indicating a tendency towards speed and fluency of output, yet none of the risk of the associated disadvantages to which we have referred. This is caused by the assumption that there is no noise during processing and that the control signal can be reinstated perfectly. In the next simulation we examine the effects of introducing some noise into processing.

8.3.3 Simulation 8.3: The effects of noise during learning

Introduction

In the previous two simulations a simplifying assumption was made that learning occurs in a noiseless environment. This assumption is rather unrealistic since biological systems are noisy. Also, from a modelling point of view, it has been well established that the presence of some random noise during processing can have a significant impact on the results produced by the model.

Furthermore, in the simulations above, there are two unrealistic features which are consistently present in all the results and which may be due to the 'no-noise' assumption. First of all, there was always 100% similarity between the retrieved action and the target action. Second, there was a perfect monotonically decreasing relationship between distance

from the target element and activation. In this simulation the effects of introducing noise during the association process are examined.

Method

To examine this issue we ran a modified version of the previous simulation in which **D** values of 2, 6 and 10 were selected. In each of these conditions the task was to retrieve the target action (action 3) as before. This time, however, we introduced a varying level of normally distributed random noise in each condition. The amount of noise was calculated as a proportion of the size of the weight change during learning. Thus a noise level of 1 represented a large amount of noise, whereas a noise level of 0.2 represented a small amount of noise.

Results

The results for **D**: 2, 6 and 10 are depicted in figures 8.8, 8.9 and 8.10 respectively.
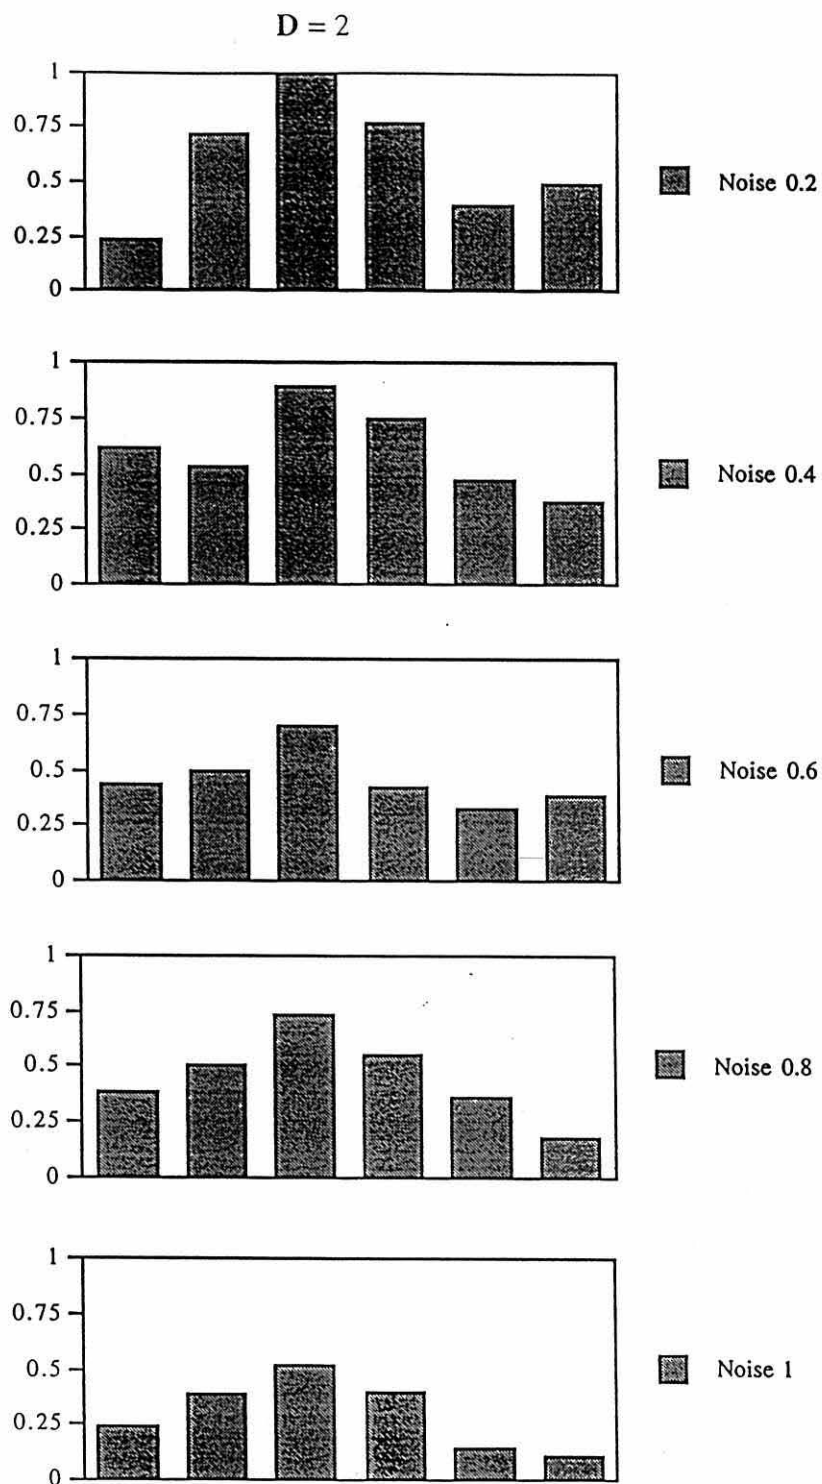
Figure 8.8. *Varying the level of noise with a slow moving contextual signal.*
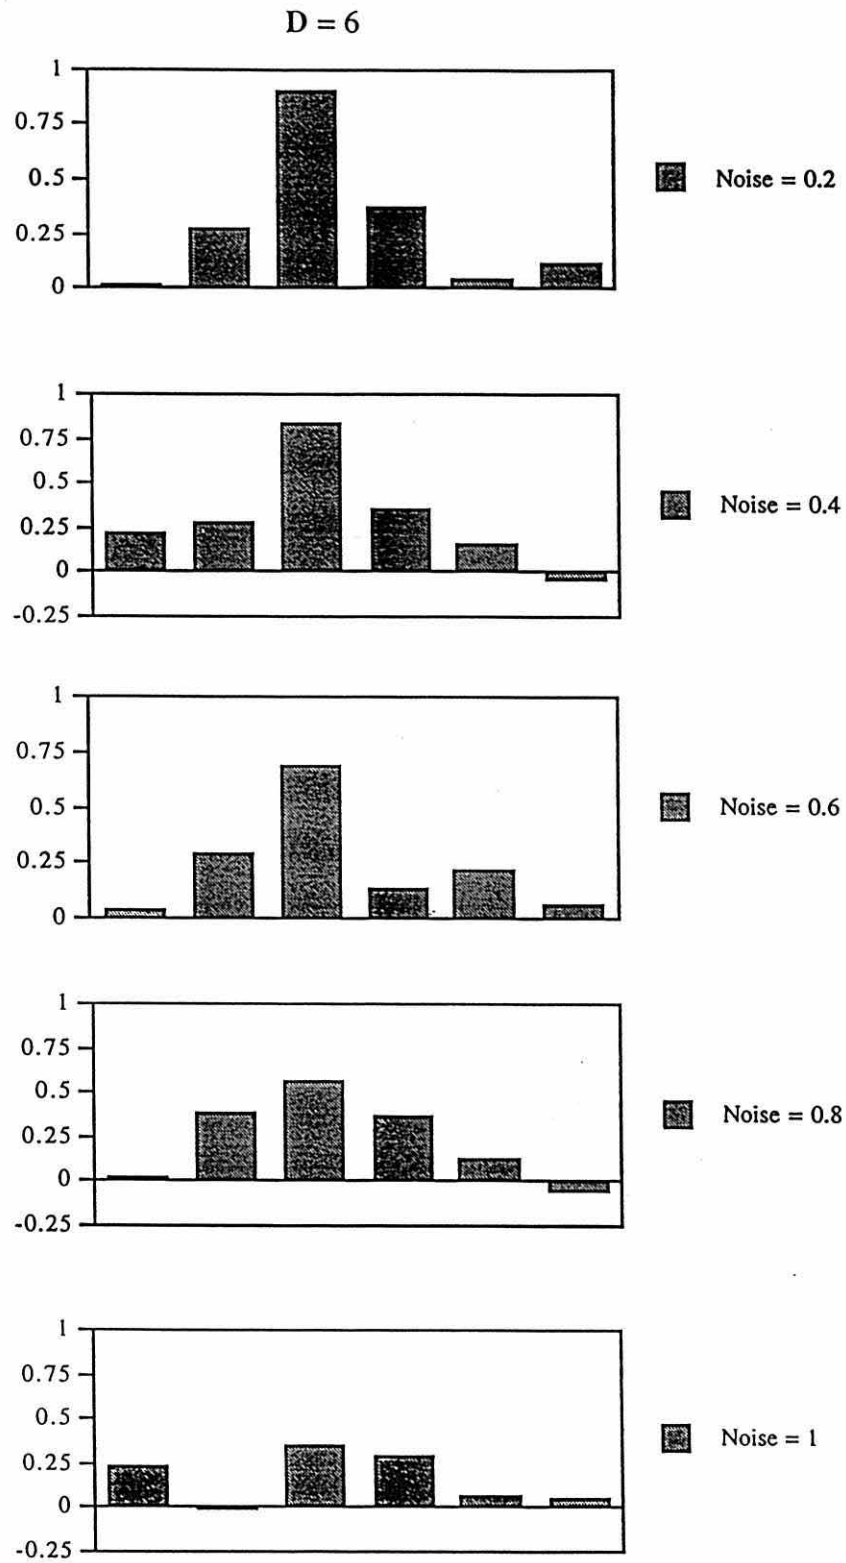
Figure 8.9. *Varying the level of noise with a medium rate contextual signal*
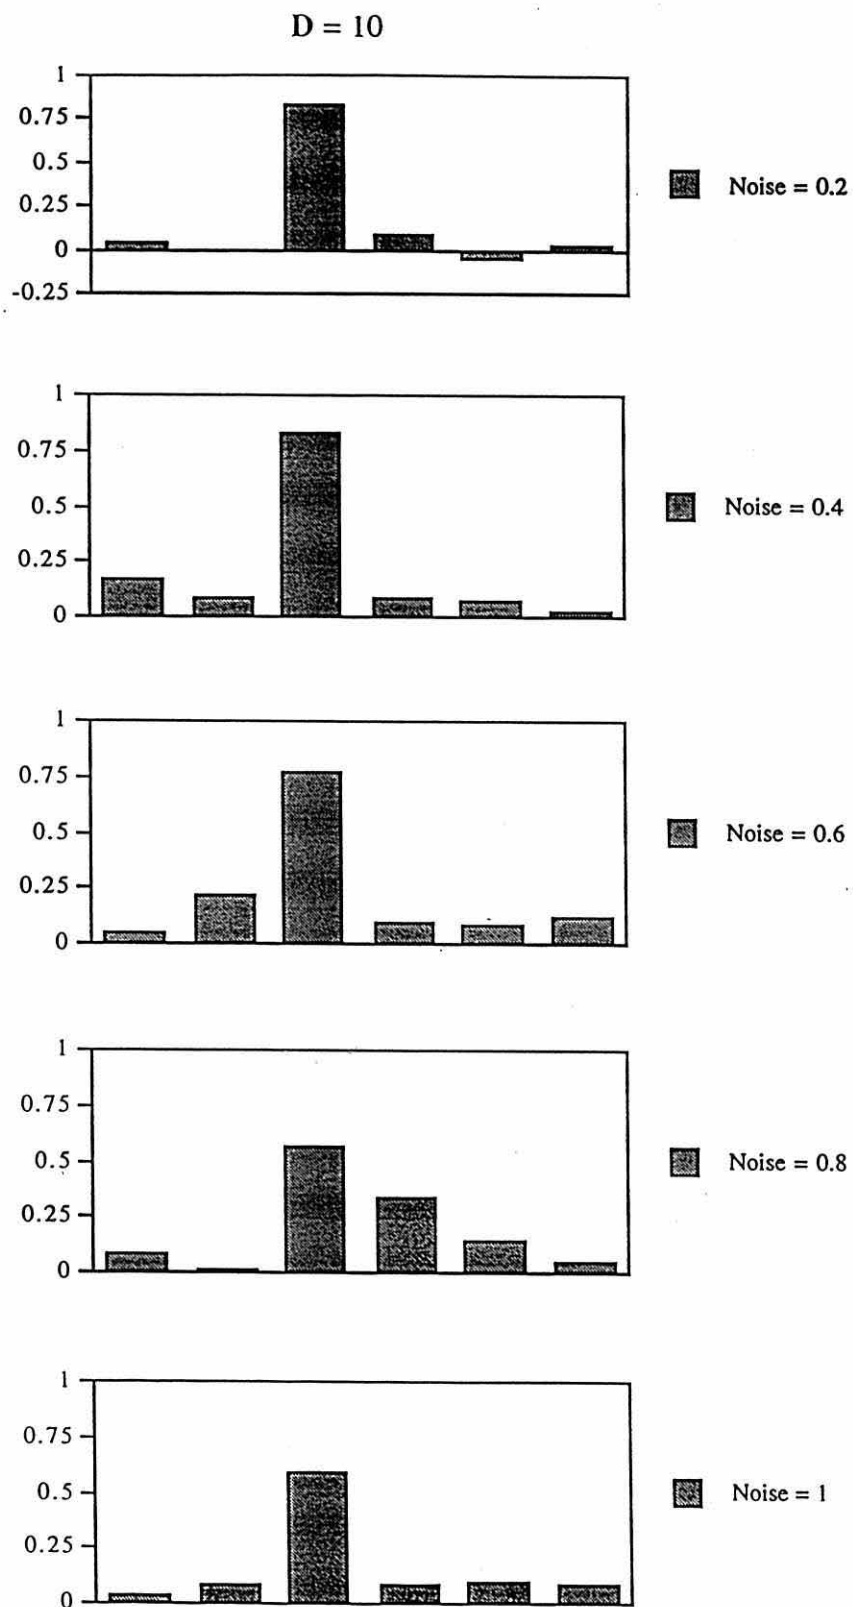
Figure 8.10. *Varying the level of noise with a fast moving contextual signal*

In each condition tested the effects of introducing noise, even at low levels was to disrupt the monotonic distance-activation relationship observed in the previous simulations. Moreover, as the level of noise is increased there is a progressive reduction in the similarity between the target and the retrieved action. In all conditions, when the noise has a power of one the activation has reduced to between 0.4 and 0.6.

It is important to note, however, that in all cases the presence of noise has not affected the fundamental retrieval properties of the network, i.e. the pattern retrieved in the third position is consistently more similar to the target than any of the other patterns retrieved.

Discussion

The introduction of noise successfully corrected the unrealistic features of network performance reported in the introduction to this simulation and is consequently incorporated in all subsequent simulations reported in the following chapter. Although the basic retrieval properties of the network are not affected by noise in the processing, it does mean that there is a high potential cost to using a slower speed control signal indiscriminately. In general a fast-moving control signal will produce most accurate retrieval of sequence elements, although it is assumed that optimisation of performance of the system would require some situation-dependent modulation of the speed of the control signal.

For the purposes of characterising contextual control in fronto-striatal circuits in the current model we assume that medium to fast speeds of the control signal (i.e. **D** values of 4 to 10) represent the normal range of "efficient" internal contextual control. Slower speeds (**D** less than 4) represent a degraded control signal that would result from damage to the system.

8.4. Chapter Summary and Discussion

8.4.1 Summary

This first series of simulations with the OSCAR model has shown that a dynamic contextual cue can be used to drive the sequential flow of action. Furthermore, it is possible to construct an appropriate time varying context signal from simple oscillators of the type which we argued above may be responsible for driving the simplest forms of sequential

behaviour in primitive organisms. This lends initial support to the hypothesis that a multidimensional oscillatory signal alone could lead to sequential action.

In terms of the hypotheses outlined in chapter four we suggest that the model provides a dynamic notion of context which can be used to extend existing hypotheses that frontal areas of cortex are in some way involved in the contextual control of internally generated action. More specifically, we suggest that time-varying context signals which have different rates of change can be seen as computationally analogous to varying strengths of contextual control of action provided by fronto-striatal circuitry. In other words, it is possible to characterise "strong" frontal contextual control of action as a high dimensionality fast moving dynamic contextual control. Such control leads to highly specific and selective activations of candidate actions for output. The idea that frontal control of action is sometimes weaker can at a computational level be captured as the idea that alternative action candidates are more activated by a slower moving, less distinct, context. In the next chapter we explore the idea that "action selection" will be harder or easier as a function of the level of "frontal" contextual control that is provided.

8.4.2 Introducing inhibitory processing to the model

Real biological systems are noisy, and a central task facing the organism at any given point in time is to select one and only one action for output from the variety of competing actions that are likely to be activated to a greater or lesser extent simultaneously. It is plausible that some criterial degree of separation between the levels of activation of different actions must be achieved before just one unique action can be selected for output. Many of the behavioural disorders that we have considered can, we have suggested in a number of places, be viewed as an inability to achieve this selection for action efficiently.

Here we assume, in common with e.g. Houghton, (1990) that some competitive process is necessary amongst competing action representations if one and only one action is to achieve a level of activation that is sufficiently high relative to its competitors to enable it to be output. In other words, the competitive action selection process must enable the choice of one action from the variety of competing, partially activated, action

representations. We tentatively attribute this function to the basal ganglia (pallido-thalamic system). It is assumed that action representations are generally held under tonic inhibition, and that the function of activation provided by the dynamic contextual control signal is to allow disinhibition of a contextually appropriate action. These issues are pursued in the next chapter.

CHAPTER NINE

The Role of Competition and Context in Basal Ganglia Control of Action Sequences

In this chapter we return to the central themes of this thesis which are embodied in the conceptual model presented in chapter four. Our first aim is to provide a computational account of the internal control sequential actions. The account is based on two interacting subsystems: a dynamical contextual signal and an action selection system which performs competitive processing amongst candidate action schemata. The second aim is to show by simulation how the interaction between these two subsystems can account for the pattern of impairments seen in motor disorders of basal ganglia dysfunction when damage is incurred. We examine the effects of two different forms of damage in the network: degrading the control signal and altering the parameters of excitation and inhibition in the action selection subsystem. The control signal is provided by the multidimensional array of oscillators introduced in the previous chapter. In the next section we describe the simple competitive network that forms the basis of action selection in the model.

9.1 Inhibitory processes: The selection of action

This series of simulations describes our implementation of a simple competitive action selection mechanism, and will be used to illustrate the suggestion that there is a trade-off between the quality of contextual control that is available and the need for competitive processing. In other words we examined the hypothesis that there will be a reduced need for competitive inhibitory processing when there is efficient 'frontal' dynamic contextual control.

9.1.1 Simulation 9.1: An 'action selection' mechanism

Introduction

The purpose of simulation 9.1 is to illustrate how a simple competitive network can enable efficient action selection on the basis of an activation gradient across a range of candidate actions for output. There are many different ways of implementing a simple winner-take-all network of the type required for this task and their properties are well understood. For present purposes the precise choice of implementation is not crucial; many different architectures could be used to illustrate the trade off between quality of contextual control of action and the amount of competitive processing necessary. In this first simulation we simply show that a simple winner-take-all network can lead to the eventual activation of one and only one action representation.

Method

As in simulation 8.2 in the previous chapter we focused on the ability of the network to output the third element of the action sequence described in simulation 8.1. The task of the competitive network was to take as input the pattern of action activations produced by a version of the model that used a context signal with medium distinctiveness ($D = 4$, as in simulation 8.1). As its output, the competitive network was required to produce a set of activations in which one and only one action representation was maximally activated, and all other action representations were completely inhibited.

We adopted a simple mechanism in which the value of each action activation was constrained between a maximum value of +1 and a minimum value of -1. The network incorporates both an inhibitory mechanism and an excitatory mechanism. In each simulated time cycle each action representation inhibits every other action representation by an amount proportional to the positive activation of the inhibiting unit. Furthermore each action representation acts to increase its absolute level of activation; i.e. each unit has an intrinsic tendency to drive itself away from an activation level of zero, and towards an activation level of either + 1 or - 1. Each representation's activation is updated by a small amount in each of many simulated time cycles until every representation has an activation either greater than +0.9 or less than -0.9. In other words the competitive process terminates when every representation has approached either complete excitation or complete inhibition.

Here it is only necessary to remember that each representation tends to inhibit every other representation and to drive itself away from zero. There are only two parameters in the model: one that determines the strength with which each action representation inhibits every other (parameter I), and a parameter that governs the strength with which every unit tries to drive its own level of activation away from zero (this is the excitation parameter E). For this first simulation intermediate values of I and E were chosen with each being set to 0.4.

Results

The results of running this simple competitive process are illustrated in figure 9.1. The first panel of figure 9.1 shows the activations provided as input to the competitive network. This is the pattern of activations that is produced by the dynamic contextual control signal for action 3 with a medium rate context signal ($D = 4$) and a medium level noise (noise $= 0.6$). The second panel of figure 9.1 illustrates the changed levels of activation after 1 time cycle of competitive processing in the network. It can be seen that while the activation level of the target activation - action 3 remains roughly the same, the competing actions now have reduced levels of activity, with actions 5 and 6 now having negative levels of activation.

The third and fourth panels show the levels of activation after 3 and 5 time cycles respectively and the bottom panel shows the level of activities when the process terminates after 7 cycles. At this stage action 3 has won the competition with an activation level of close to 1.0 and the representations of all the competing actions are completely inhibited, with levels of activity close to -1.
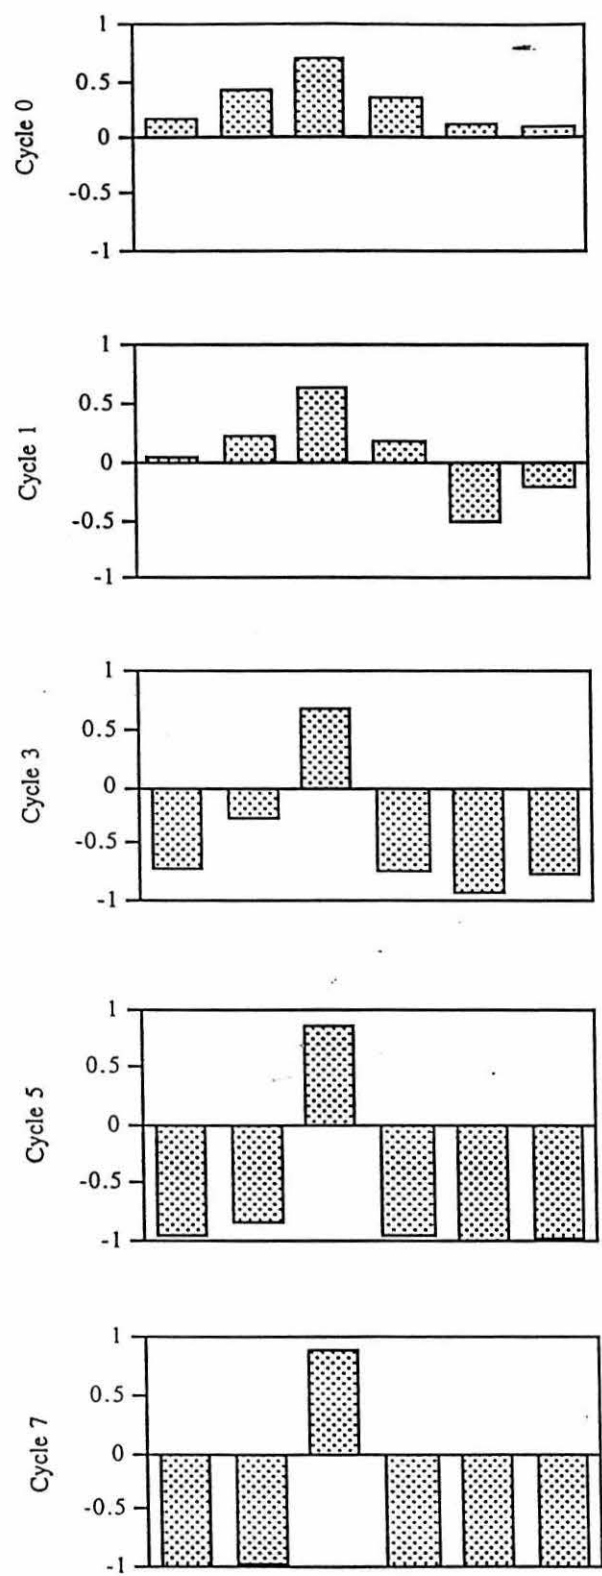
Figure 9.1. *Selection of "action 3" using competitive processing*

Discussion

The results of simulation 9.1 simply illustrate that a basic competitive network, relying primarily on lateral inhibitory connections can act as an efficient winner-take-all system, effectively performing contrast enhancement in such a way that the outcome is that just one action representation is fully activated and all competing action representations are completely inhibited. This result was achieved with a choice of intermediate values for the excitation and inhibition parameters. However it is important to note that in a network of this type successful winner-take-all behaviour will not necessarily emerge with any choice of values of the parameters that govern the strengths of the excitatory and inhibitory processing.

9.1.2 Simulation 9.2 : The effects of altering the relative strength of excitatory and inhibitory parameters on effective action selection

Introduction

The purpose of simulation 9.2 was to examine the effects of varying the choice of values for inhibition (I) and excitation (E) on the outcome of the competitive action selection process. First we report two simulations to illustrate the results of choosing extreme values for I and E and then we report a more systematic exploration of the parameter space, with a view to determining the sensitivity of final outcome to initial parameter values.

Method

All simulations were identical to simulation 9.1, except that the values of I and E were systematically varied. First of all we carried out an "over-excitation" simulation in which the value of the excitation parameter E was set to a value that was high (at 0.5) relative to the value of the inhibition parameter I (which was set at 0.1). Next, we carried out an "over-inhibition" simulation in which these values were reversed (I was set to 0.5 and E was set to 0.1).
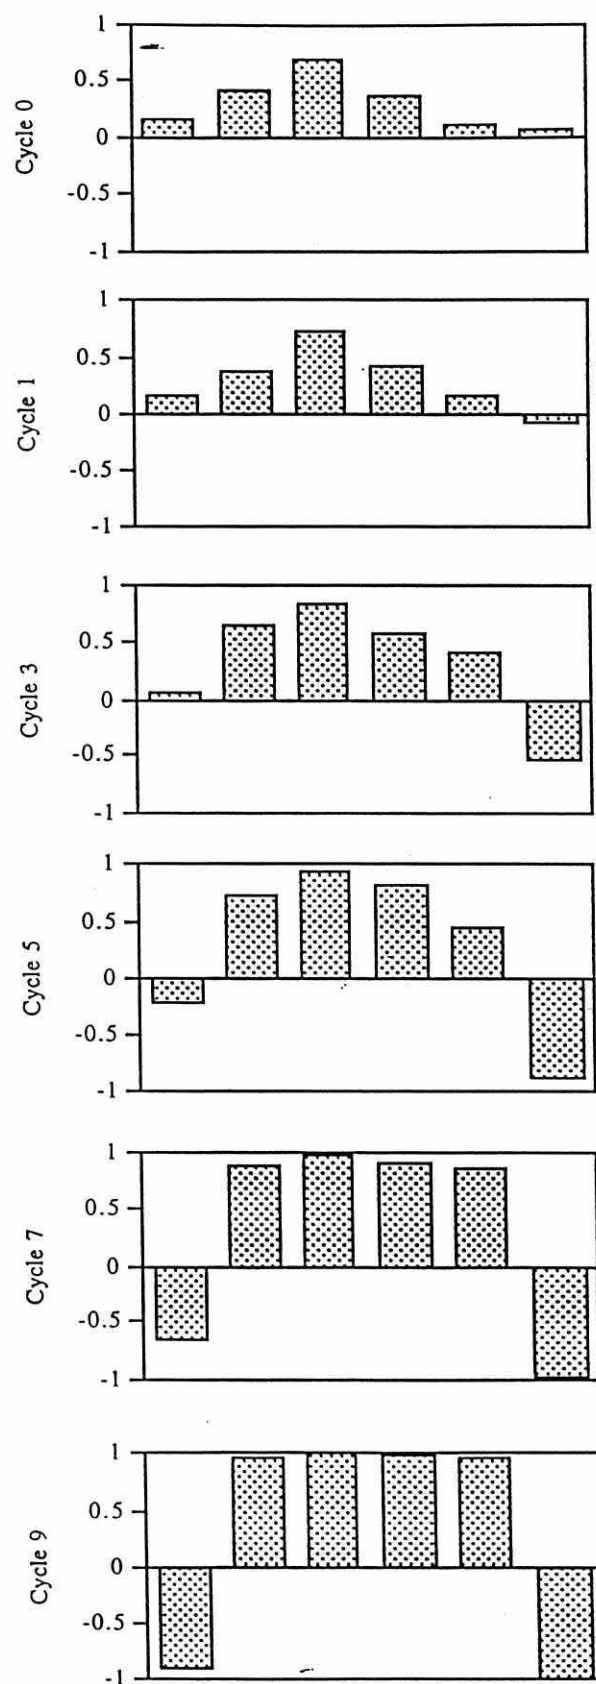
Figure 9.2. *Coactivation of several actions with insufficient inhibition*
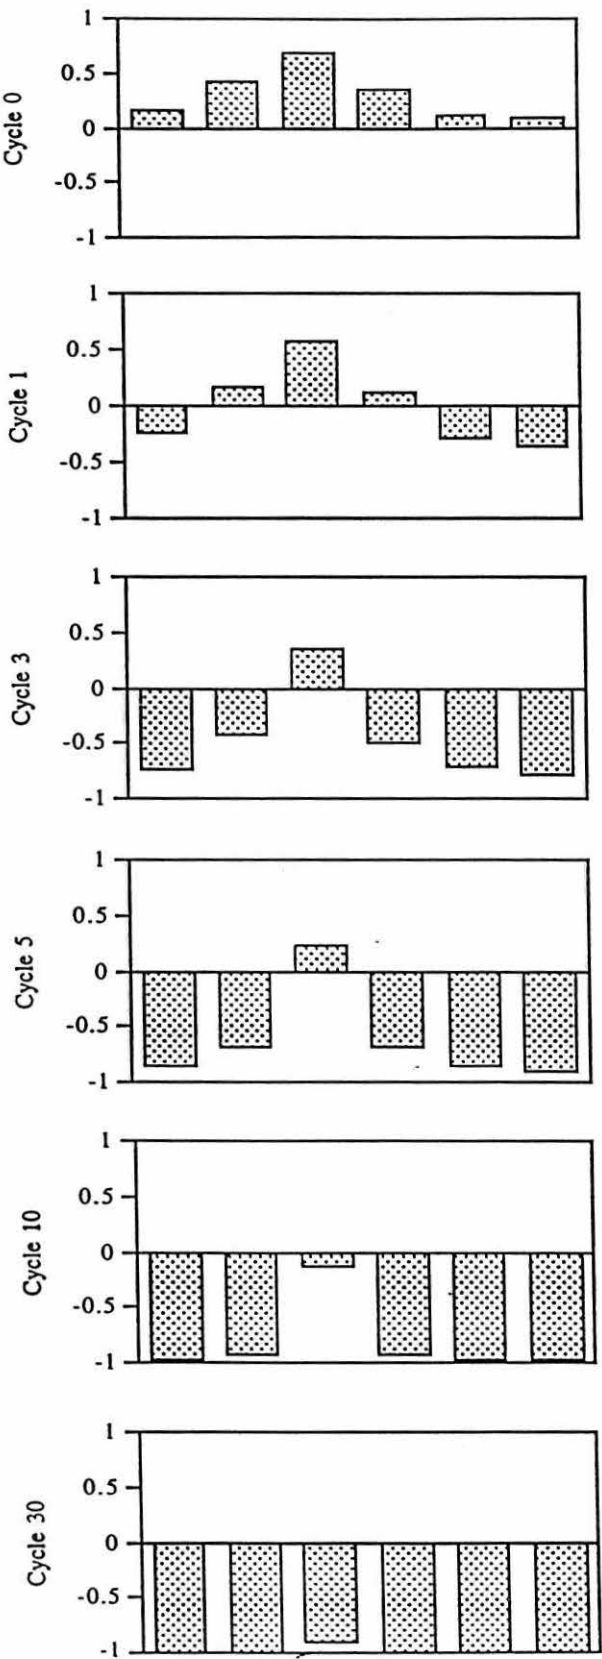
Figure 9.3. *Inhibition of all actions with insufficient excitation*

Results

Figure 9.2 shows initial, intermediate and final levels of activation, reached after 9 cycles, for the condition in which the level of excitation was high relative to the level of inhibition. It can be seen that, in contrast to the results of simulation 9.1, a stable state was reached after 9 cycles in which "coactivation" is evident. The final levels of activation for actions 2 through 5 were all at the maximum level of 1.0, while the activation levels of actions 1 and 6 were completely inhibited.

Figure 9.3, in contrast, shows the corresponding levels of activation in the condition in which the level of excitation was low relative to inhibition. In this case, the high level of inhibition eventually led to a situation in which every action representation was completely inhibited.

Discussion

These additional results illustrate that the choice of excitation and inhibition parameters is not arbitrary: only some intermediate combinations of parameter values will lead to a successful outcome in which one and only one action representation becomes completely activated.

This simulation serves to illustrate the need for modulatory processes to maintain a balance between the excitatory and inhibitory parameters which control the 'focus' of selection mechanisms in a competitive network. If the focus is too broad, several competing action representations are facilitated which will lead to overload of the final common pathway to the muscles. A number of deficits may potentially occur in this situation, such as increased response or initiation times, co-activation amongst competing muscle groups leading to increased rigidity and dysmetric movements, or it may cause the inclusion of inappropriate movements in the action which is performed. Alternatively, if the focus is too narrow all actions are inhibited and no movement will occur. Both of these conditions reflect the symptoms of motor disorders of basal ganglia dysfunction. Moreover, the notion that dopamine participates in the modulation process is consistent, at a computational level, with the theoretical ideas presented throughout this thesis and in particular the 'on-centre/

off-surround' model of dopaminergic participation in the action selection process put forward by Mink and Thach (1993) (see chapter 2).

The model also provides a computational basis for differentiating between basal ganglia type deficits in motor sequencing and normal serial order errors or 'action slips'. In this model serial order errors would be indicated if the wrong action is selected. This, as can be seen above, is not a feature of failure of the modulatory process, and it is not a characteristic of basal ganglia disorders. In the next section we pursue this theme more systematically.

9.2. General inhibition and selective disinhibition

Figure 9.4 overleaf illustrates the results of a more systematic exploration of the parameter space. Each cell in the figure represents the outcome produced by the competitive network with a particular choice of parameters. The top line of each cell represents the final level of activation of each action representation with a 1 representing a final activation level of +1.0 and 0 representing a final activation level of -1.0. Asterisks on this line indicate that the network failed to reach a stable state within the criterion of 50 time cycles. The second line of each cell represents the number of cycles required to achieve this stable state with a maximum of 50. Thus for example the top right hand cell gives the results of running the network with $E = 1.0$ and $I = 0.0$. The 111100 indicates that at the end of the competitive process the first four action representations will be maximally activated, while the 6 on the second line of the cell indicates that this state was reached after 6 cycles of competitive processing. Thus any cell in which the top line contains the pattern 001000 represents a correct response in which the third and only the third action in the sequence is activated at the end of processing.

Excitation strength

Inhibition strength	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	****** 50	****** 50	111110 40	111100 22	111101 16	111101 11	111110 9	111101 7	011111 8	011100 10	111100 6
0.1	****** 50	****** 50	011100 46	111100 22	011101 14	011110 16	111100 9	011110 21	011110 7	111110 8	011110 6
0.2	****** 50	001000 31	001000 24	011000 19	011100 12	011000 12	011100 10	011100 7	111100 7	111100 5	011100 6
0.3	****** 50	****** 50	001000 18	001000 12	001000 12	011100 13	011000 8	011000 9	011100 6	011100 6	001100 6
0.4	****** 50	****** 50	001000 29	001000 11	011000 19	001000 7	001000 7	011100 7	011000 6	011000 5	011010 8
0.5	****** 50	000000 29	000000 20	000000 15	001000 8	001000 7	001000 5	011100 9	011000 13	011000 6	011100 3
0.6	000000 35	000000 27	000000 21	001000 22	001000 11	001000 7	001000 5	001000 5	001000 5	001000 7	001000 5
0.7	000000 22	000000 18	000000 20	000000 23	001000 13	001000 18	001000 10	001000 4	001100 8	001000 5	001000 4
0.8	000000 24	000000 12	000000 13	000000 9	001000 15	000000 13	001000 8	001000 9	001100 11	001000 6	011000 4
0.9	000000 22	000000 18	000000 14	001000 14	001000 15	001000 7	001000 3	001000 5	001000 5	011000 5	011000 6
1.0	000000 18	000000 11	000000 13	000000 8	000000 11	001000 18	001000 5	001000 11	001000 8	001000 9	001000 5

It can be seen that correct outcomes are more likely to be obtained with intermediate values of both I and E (i.e. in the centre of the parameter space represented by the figure). The light shaded area of the parameter space (essentially the top right hand corner), encloses all points where coactivation of action representations has occured owing to the use of a value for the excitation parameter that is too large relative to the value of the inhibition parameter.

The dark shaded area (essentially the lower left hand triangle), indicates the area of the parameter space which led to an outcome in which all action representations were completely inhibited owing to the use of a value for the inhibition parameter that was too large relative to the excitation parameter value. Note that these results represent just one simulation with each combination of parameter values and so the precise shape of the enclosed areas must be regarded as approximate.

Inspection of the number of cycles required to reach a stable state indicates that values are generally high in the top left hand corner of the parameter space where both excitation and inhibition parameters were small. In the lower right hand corner of the parameter space, which corresponds to high values of both parameter values, stable states are typically achieved after much smaller numbers of processing cycles. In the absence of other considerations, therefore, it would be in an organism's interest always to adopt high values for both excitation and inhibition parameters. In the limiting case this would simply amount to a strategy of choosing the action with the highest activation initially for immediate output. However, because there is noise in the process such a system would be extremely unreliable. This is illustrated by the cells enclosed in the light shaded area (*lower right corner* of parameter space), corresponding to high values of both E and I. In this region we see that incorrect outcomes are achieved in which actions other than the target action may become fully active at the end of processing.

This clearly illustrates the fact that, in a noisy system, there is a cost to attempting to achieve the action selection too quickly. This is because in the early stages of the competitive process, when all action representations have similar levels of activation to one another, random noise added to the excitation and inhibition that spreads from one

representation to another can cause action representations to momentarily reverse their relative order of activation. This can lead to spurious additional winners in the winner-take-all competition.

It will also be noticed that whenever some set of activations are activated, the correct action is always included, this leads to a situation where there is co-activation of additional action representations with the correct one, *not* activation of an incorrect action alone.

Discussion

The results of these simulations illustrate several important points. First of all, it is clear that the right choice of inhibition and excitation parameters must be made if the correct outcome is to be produced by the competitive process. If excitation is too high relative to inhibition then coactivation will result in which the system effectively selects two or more competing actions for simultaneous output. If on the other hand inhibition is too high relative to excitation, the opposite outcome ensues in which all action representations become inhibited, and the system is effectively paralysed in that no action is selected for output. A further important point is that the competitive selection for action network exhibits a trade-off between accuracy of action selection and speed of reaching a stable state. If too-high values are chosen for both excitation and inhibition the system converges very rapidly but on an incorrect answer.

We discuss the psychological interpretation of these results in more detail at the end of this chapter. Here we simply note that any noisy biological system in which competitive inhibitory action selection plays a part must incorporate some mechanism for modulating the excitation and inhibition parameters to maintain the system in the appropriate region of parameter space.

The simulations we have just described have explored that parameter space of the competitive network when it is provided with input from a dynamic contextual control vector of intermediate distinctiveness (the parameter **D** was set at 4, indicating that temporally adjacent states of the control signal were neither very similar nor very dissimilar to one another). In the next set of simulations we explore the possibility that the correct choice of parameters will be less crucial when higher quality contextual control is provided.

9.3 Interactions between context and competition

9.3.1 Simulation 9.3: Examining the parameter space of effective action selection as a function of contextual quality

Introduction

The purpose of this simulation is to examine how the usable region of the parameter space of the action selection network will vary as a function of the quality of contextual control that is available to selectively activate different action representations. It might be expected that the choice of excitation and inhibition parameters will be less crucial if a high distinctiveness (fast moving) context signal is used to activate action representations. We were also interested in the possibility that the action specification provided by the dynamic contextual control could degrade to a point where there is no choice of the parameter values for the competitive network that will lead to the reliable selection of an appropriate action.

In general terms, then, simulation 9.3 explores the reciprocal relationship between contextual action specifications and competitive action selection, with a view to demonstrating that there is less need for accurate parameter selection in action selection if a good quality contextual control is available.

Method

In this simulation we again examined the ability of the competitive network to correctly select action number 3 in the sequence initially investigated in simulation 8.1. However in this case separate simulations examined the parameter space as a function of the quality of contextual specification. We therefore took the six different activation gradients for action 3 as target that were produced using versions of the context signal that used different values of **D**.

The competitive network worked in the same way as described in simulation 9.1, however in this case we ran 100 simulations using each combination of excitation and inhibition parameter values between 0.0 and 1.0 with a step size of 0.1. The dependent measure in this case was the percentage of these 100 simulations in which a "correct" action selection was achieved (i.e. in which the representation for action 3 was fully active when competitive processing terminated and every other action representation was completely

inhibited. If a stable state was not achieved within 50 processing cycles, this was classified as a failure.

Results

The results with an intermediate value of D (D = 4) can be seen in figure 9.5. Each panel of figure 9.5 shows the percentage of correct responses achieved with a different combination of excitatory (x-axis) and inhibitory (y-axis) parameters. The shaded area in the middle of the parameter space indicates that part of the space in which a correct outcome was achieved by the action selection network at least 90% of the time.

	0.0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1.0
0.0	0	0	0	0	0	0	0	0	0	0	0
0.1	0	0	0	0	0	0	0	0	0	0	0
0.2	0	100	85	25	0	0	0	0	0	0	0
0.3	0	0	100	95	80	30	20	20	0	0	0
0.4	0	0	95	100	95	95	55	50	25	25	5
0.5	0	0	50	95	95	100	90	95	45	40	15
0.6	0	0	5	80	95	90	100	80	55	40	35
0.7	0	0	5	45	80	80	95	90	90	75	75
0.8	0	0	0	25	75	75	80	90	90	65	70
0.9	0	0	0	5	50	75	75	90	75	95	60
1.0	0	0	0	0	35	30	65	85	85	100	65

Figure 9.5. *Parameter space of effective action selection with D = 4*

These results confirm the conclusions of the previous section in that most reliable performance is achieved with a roughly equal balance between excitation and inhibition, in the middle range of their values.

Figure 9.6 illustrates the results of this simulation when the contextual distinctiveness is is increased. For clarity, these results are as contour maps in which the shaded region again represents 90% correct selection or above. For values of D below 4, i.e. with a poor quality context signal there are no combinations of excitatory or inhibitory parameters which produce performance of greater than 60% correct and consequently these cases are not represented below. The first panel of figure 9.6 above is the contour map equivalent of figure 9.5. Each subsequent panel shows an increase in the area of the parameter space which will reliably produce accurate selection performance.

Region in which the correct action is selected more than 90% of the time.
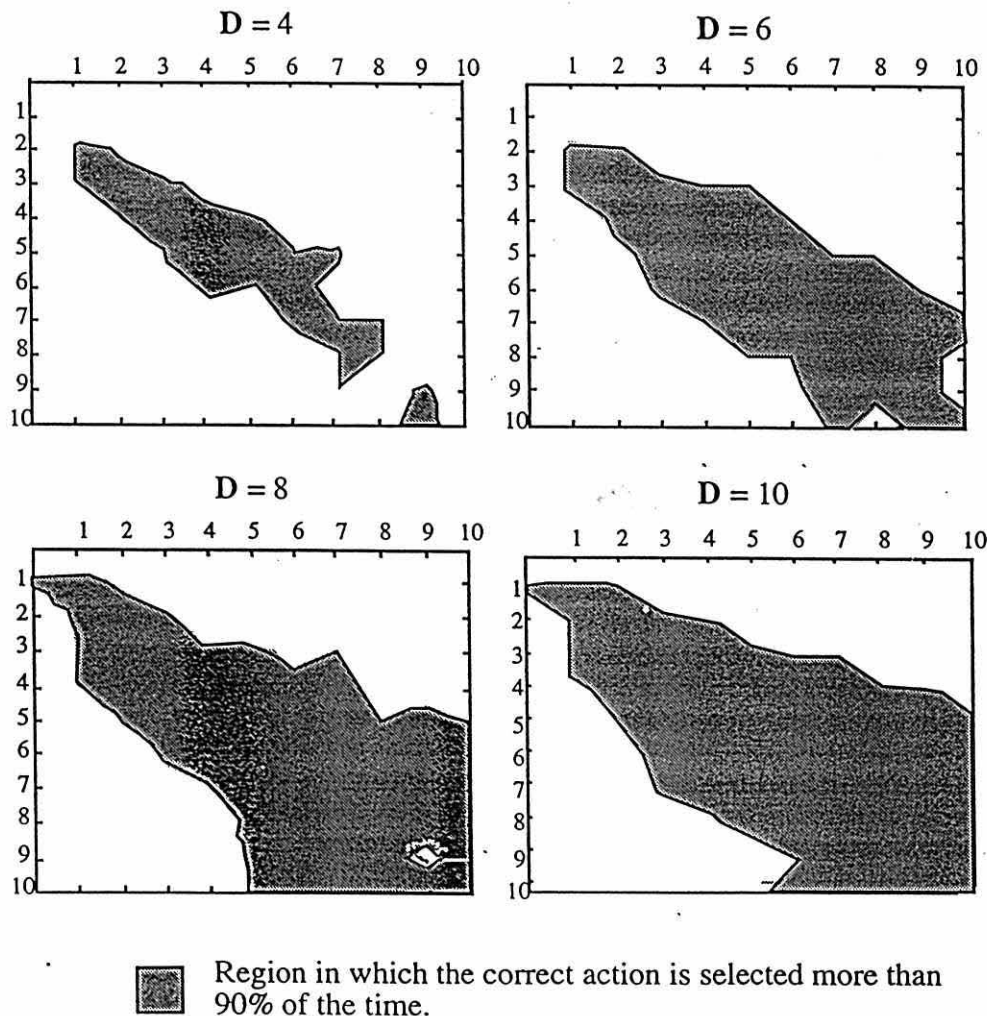
Figure 9.6. *Contour maps of the region of efficient action selection as the distinctiveness of the context signal increases*

Discussion

The results of this simulation provide a clear illustration of the hypothesis introduced in the introduction. This was that the choice of parameter settings in a competitive action selection network becomes increasingly critical as the ability of the high level contextual control signal to activate a target action selectively reduces. Indeed if contextual control is inadequate it becomes impossible to reliably select a correct action. Thus there is a clear trade off between the quality of contextual control and the need for sophisticated and precise action selection mechanisms to operate subsequently.

9.4 Response time as a function of contextual quality

In this section we consider the effects of the interaction between the quality of contextual representations and response competition during preparation of a motor action on the output behaviour of the system.

Information processing studies which provide empirical data on this point use reaction time (RT) as a standard measure of the time to prepare and execute a response. Simple reaction times (SRT's) in which there is only a single response for which all information is available in advance, i.e. it is fully precued externally, fall in a very limited range of around 150 - 200 ms. Choice reaction time studies measure the extra processing time required when there are a number of possible responses, and the choice is determined by properties of the stimulus. In this condition RT's typically lengthen as a linear function of the number of choices, according to the Hick-Hyman Law (Hick, 1952). This relationship is altered, however, if there is partial precueing of the correct response in advance of initiation. In other words, if *contextual* information is available which predicts the correct response, the level of competition is reduced because the correct response is selectively facilitated and RT is shortened as a consequence. For example, if the environmental stimuli always follow a predetermined sequence, the subject is provided with contextual information which facilitates each successive response. Studies have found that the quality of contextual information, determines the extent of precueing and thus crucially affects the extent to which RT's are shortened below uncued levels.

Now let us translate this information into the framework of the current model. We have hypothesised that frontal cortex provides internal contextual information which, through learning, substitutes for external cues when they are unavailable or inappropriate. Moreover we have suggested that a highly distinct contextual control signal is a key processing feature in the smooth feedforward control of rapid sequential actions. We have already shown that a highly distinct contextual signal facilitates accurate action selection by placing reduced demands on the competitive selection process. In figure 9.7 below we also show how a good quality contextual signal leads to reduced response times in the output of the system.

To produce figure 9.7, we simply collected selected data from the previous simulation, in which the parameters of the competitive network are in the optimal range. The parameter values used were - E = 0.5, I = 0.5. Average values were then computed over 20 trials for the number of correct selections and the number of processing cycles taken to reach a stable state. The number of processing cycles is assumed to represent a simple measure of response time. These data were recorded across the range of values for the discrimination parameter (D) examined in the previous simulation.

Figure 9.7 shows a general reduction in response time measured in processing cycles, (left hand scale), as the quality of the contextual signal improves. This is paralleled by a concomitant increase in consistency with which the correct action is selected, (right hand scale).
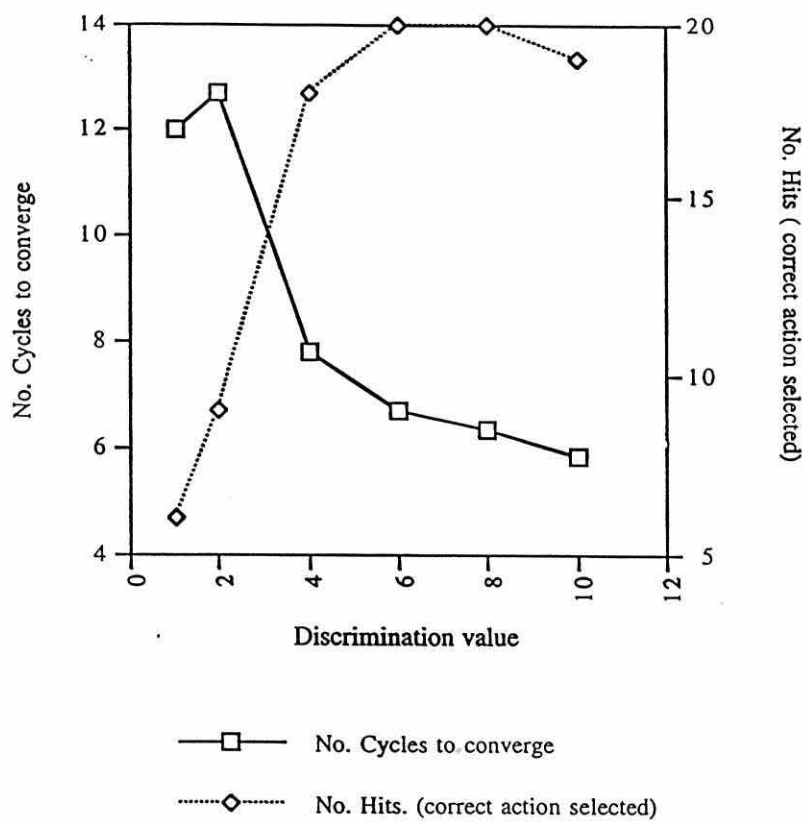
Figure 9.7. *Change in response time as a function of contextual quality*

On both measures the greatest difference in performance occurs between **D** values of 2 and 4. It will be recalled from simulation 9.3 that when the distinctiveness is less than four there is no region of the parameter space in which satisfactory (90% correct) performance is achieved. This is confirmed here, with a mean value of approx. 45% correct when **D** = 2. The substantial increase in response competition leads to the prolonged response times produced by the network when **D** is less than four.

Here we have assumed that contextual information provided directly by the environment is absent. If the system is unable to maintain a balance between excitation and inhibition, then an even higher quality of contextual input would be required to produce normal performance. If this is not possible, then greater reliance would be placed on external sources of contextual information to compensate for the deficient feedforward system and prevent performance decrements, such as those produced in the model.

9.5 Discussion

The simulations presented in this chapter serve to illustrate a number features of basal ganglia involvement in motor control and motor disorders in the context of the conceptual framework of this thesis.

9.5.1 The computational role of cortical and sub-cortical structures

The role of frontal and premotor cortico-striatal circuits through the basal ganglia in the internal generation of movement can usefully be subdivided into two interdependent computational processes:

1. An internal representation of context which substitutes for environmental stimuli in controlling the internal generation of action and which plays a special role in sequential behaviour. We suggest that this function is subserved by frontal cortex and is relayed to the basal ganglia via excitatory projections to striatum and subthalamic nucleus.

2. A system for selectively facilitating appropriate actions or sequence elements and inhibiting competing alternatives which are contextually inappropriate. We suggest that this mechanism is subserved by basal ganglia internal processing modulated by dopaminergic inputs from the substantia nigra pars compacta (SNc).

The computational interdependence between the two systems has an anatomical substrate in the recurrent loops from basal ganglia output stations (GPi) through thalamic nuclei to cortical areas, (Alexander and Crutcher, 1990).

9.5.2 A 'frontal' oscillatory context signal

The "frontal" contextual signal may be constructed from a multidimensional array of simple oscillator elements, and under controlled conditions is sufficient by itself to control sequential behaviour, (simulation 8.1). The contextual signal constructed in this way has the necessary properties of parallel activation of several sequence elements, which are required for modelling coarticulation, smoothness and serial reaction time data in human movement sequencing, (Houghton, 1990; Rumelhart and Norman, 1982). The downside of this property is that if the quality of the contextual signal is degraded, (simulation 8.2)

and/or noise is present in the system, (simulation 8.3), increased competition between responses is likely to lead to error prone recall and performance. Also, it should be noted that the internal contextual signal cannot be assumed to exist a priori, and must be learned by exposure to cues and regularities in the environment, thus a poor quality (indistinct) contextual signal will exist in early stages of learning. This provides an additional argument for the need to supplement the contextual signal with the additional disinhibitory mechanisms provided by the basal ganglia (simulations 9.1 and 9.2).

9.5.3 Basal ganglia disinhibition

In chapter 2 we reviewed neurobiological evidence that the basal ganglia play a disinhibitory role in eye movements (e.g. Hikosaka, 1991) and considered previous suggestions that this system plays a similar role in the generation of limb movements (Graybiel & Kimura, 1995). The difference between the two cases is that in eye movements a unitary disinhibitory signal will suffice, whereas in the case of limb movements, *selective* disinhibition is required. We have reviewed a variety of evidence which suggests that limbs tend to act together, either due to mechanics or representational variables (Kornblum, 1965; Bernstein, 1967), and that inhibitory mechanisms are required to prevent cocontractions in disparate effectors when it conflicts with task requirements. We have implemented a disinhibitory subsystem here as a simple competitive network, in which the competition process is modulated by a combination of excitatory and inhibitory parameters. This implementation is intended to model the "dual-pathway" action on the basal ganglia outputs, in which one pathway, (indirect) has a net excitatory effect and the other 'direct' pathway has a net inhibitory effect.

9.5.4 The computational role of dopamine

A range of neurobiological evidence, (see chapter 2) combines to suggest that dopaminergic inputs to the striatum have a modulatory effect on the two pathways, which acts to maintain a balance between them at rest, such that GPi neurons are held under tonic inhibition. Just prior to movement initiation, however, the balance is briefly altered, thus

releasing selected populations of GPi neurons from inhibition. We suggest, as do Mink and Thach, (1993), that this process facilitates the activation of the desired motor program or action representation at a cortical level, whilst concurrently inhibiting competing motor programs.

In the conceptual framework of chapter four, we assume that the desired motor program is specified to the disinhibition subsystem by either environmental variables, via sensory systems or the internal contextual signal provided by cortex. We further assume, on the basis of recent neurobiological evidence (Brown, 1992; Graybiel & Kimura, 1995), that there is a topographic mapping from striatum to pallidum, which represents either body parts or motor synergies, that allows the right selection of pallidal neurons to be targeted. In feedforward sequential control we assume that the primary source of input is provided by the internal contextual signal representing the whole sequence.

In the model, we let the retrieved activation gradient across the action vector representing the sequence provide input to the competitive network. Under normal conditions the competitive network acts to increase the activation of the target action, whilst suppressing the activations of competitors, thus enhancing the contrast between the two, (simulation 9.1).

9.5.5 Modelling motor disorders

According to the data presented by Delong (1990) amongst others, in Parkinson's disease or MPTP-induced dyskinesia, the loss of dopamine causes an imbalance in the two pathways. Specifically, the 'direct' pathway becomes over-active. This causes excessive and, importantly, *unselective* activity in the globus pallidus (GPi), (Alexander, Filion and Tremblay, 1986). A direct neurophysiological effect of this change at the level of cortical targets is difficult to determine due to the multiple and diffuse efferent projections from thalamic nuclei, which are the targets of GPi neurons, to cortex. However, Delong suggests that there is a net 'over-inhibition' of all motor programs including the desired action. This situation was modelled in simulation 9.2 by increasing the level of inhibition,

relative to the level of excitation. The manipulation caused all action representations to become inhibited.

Although, Delong's argument, appears at first glance to be the most logical suggestion in accounting for akinesia in PD, it singularly fails to account for the rigidity and muscular co-activation associated with the condition. An alternative suggestion, therefore, is that the selection process fails to select out a single action only. This causes several competing cortical motor programs to become activated, thus putting conflicting demands on the motor output system. This situation was also modelled in simulation 9.2 by increasing the excitatory parameter, relative to inhibition. These properties of the model are consistent with the revised model of basal ganglia control of action selection put forward by Mink and Thach, (1993). This model was intended to account for the apparently contradictory finding that both an increase *and* a reduction of activity in the primate GPi gives rise to parkinsonian motor deficits. As a result of the modelling work performed here, we suggest that increased activity in the GPi corresponds to increased inhibition in the competitive process leading to over-inhibition of action schemata (it will be recalled that pallido-thalamic connections are inhibitory mediated by GABA). Decreased GPi activity corresponds to increased excitation in the competitive network leading to co-activation of several schemata.

9.5.6 Basal ganglia involvement in motor learning

A number of studies have implicated the basal ganglia in motor learning. Notably, Seitz et al. (1990) using PET imaging techniques found evidence of basal ganglia activity during early stages of learning a finger tapping sequence. A number of other studies have implicated the basal ganglia in sequence learning (see Curran, 1995 for a review). This finding can be accounted for in the framework of the current model if it is assumed that the competitive action selection process plays an active role in the association of target actions with states of the context vector. This suggestion is speculative as it has not actually been modelled here. It is, however, consistent with the known loop circuitry of neuroanatomical links between frontal cortex and the basal ganglia, i.e. the basal ganglia outputs send return

projections back to frontal areas via the thalamus. It is also merely an extension of current ideas that the basal ganglia are involved in reinforcement learning of external cues which predict motor responses, to suggest that they are involved in learning internal cues for sequential control of skilled actions.

9.5.7 The nature of fronto-striatal interaction

The results of the simulations in this chapter have implications for models of fronto-striatal action, as well as for the operation of each system separately. We have suggested that frontal cortical areas provide a dynamic contextual cue that acts as input to the action-selection mechanism represented by the basal ganglia. As is illustrated by the results of simulation 9.3, there is a clear trade-off between these processes. The better the dynamic contextual specification of action representations, the less need for accurate parameter-setting in the action-selection mechanism. Thus either system may compensate for limitations in the other system, up to a point. In neuro-psychological terms, for example, impaired basal ganglia function may be compensated for by enhanced frontal contextual control. Furthermore, deficiencies in *either* frontal control *or* basal ganglia action-selection may be overcome by the provision of external stimuli that afford the target action. Furthermore, deficits in either system may lead to slowness in action selection as shown in figure 9.7.

9.5.8 Relationship with other models of sequential motor control

The OSCAR model, adapted here for use in the current framework was originally designed to model serial recall in human memory (Brown, Preece & Hulme, 1996). Its use here has shown that it has wider appeal in modelling sequential behaviour generally. In this context, OSCAR shares some features with existing models of sequencing, although there are also important differences. Existing models fall into two broad categories, defined by the mechanism used to produce sequential behaviour.

In the first category are models such as those developed by Jordan (1986a, 1990), which were used in chapters five and six. These models employ a context vector of

arbitrary dimensionality to do the sequencing, which is usually driven by recurrent connections from within the network itself. In the second category are models which primarily use inhibitory mechanisms to produce sequential behaviour. Models in this category include the Rumelhart and Norman (1982) model of typing control and Houghton's competitive queuing (CQ) approach.

In contrast, the approach developed here using OSCAR employs *both* a high dimensional contextual signal and inhibitory mechanisms to model the role of frontal - basal ganglia circuitry in the control of sequential movement. A potential criticism of this approach lies in whether the added complexity of postulating two systems can be justified when it has previously been shown that either of the two systems alone can produce sequences. This criticism can be answered on a number of grounds.

In the course of this chapter we have shown how one system can compensate for another under different conditions. On this point, we argue that the use of inhibitory processes alone is computationally slow. For example in the CQ model, the competitive process has to be run for each individual element in a sequence. Since this takes several processing cycles each time it is run, this approach has difficulty accounting for the fast production of 'overlearned' short sequences.

Using OSCAR, however, we would argue that an overlearned sequence would be represented by a high quality context signal, and thus the reliance on inhibitory mechanisms may be minimised or even eliminated entirely. This clearly reflects a property of "automatised" skilled movements, in which there is commonly highly parallel activation of limbs, which may even be involved in different tasks, without conflict or a tendency for serial order errors. This feature of the model also makes sense of the counterintuitive recent finding that the globus pallidus can be bilaterally ablated in patients with severe Parkinson's disease to dramatically liberate their movement. We have proposed in the model that the basal ganglia subserves the inhibitory subsystem. Thus it seems that having no inhibitory subsystem at all is substantially better than having one which is operating in an extreme region of the parameter space, (see simulation 9.3). On the other hand, as we have argued above there are many situations including during learning or in the presence of contextual

interference in which the influence of inhibitory processing is either desirable or even essential. Indeed in chapter four we presented a number of ecological arguments which suggest that an inhibitory mechanism must be regarded as a primitive building block of action control in biological systems.

In conclusion, we argue that the combined use of both inhibitory and contextual subsystems is justified on computational and biological grounds and that it provides a useful contribution towards future theoretical and computational accounts of action control and motor sequencing.

CHAPTER 10

General Discussion

In this chapter we discuss the work reported in this thesis in relation to the aims set out at the beginning and the issues that have arisen during the course of its development. We start with a summary of the work and its motivations in modelling aspects of basal ganglia function and motor disorders. Following that we turn to a more general discussion of the contribution of this work to theories of basal ganglia function in motor control. Finally the limitations of the current approach are discussed and suggestions are made for the future role of computational models in the development of theories of biological movement control.

10.1 Summary

The primary aim of the thesis was to develop a computational approach which could account for the finding that sequential movements are more disrupted in Parkinson's disease than simple movements. Furthermore we sought to develop an account of the underlying nature of the disruption to cortical and basal ganglia circuits in the control of sequential movements that is constrained by existing biological data.

10.1.1 The conceptual model

In chapter four we set out a conceptual model in which it is argued that the cooperative action of basal ganglia and frontal cortex subserve the internal control of sequential action using two principle mechanisms: A subsystem for the selective disinhibition of distributed cortical representations of sequence elements (motor schemata), and a subsystem for encoding contextual cueing information that is available in the environment. The dynamic contextual signal is assumed to be implemented by frontal cortex and it transmits information to the striatum via downward projections and to frontal motor areas via cortico-

cortical connections. The information encoded in the contextual signal helps to determine which action components are activated and when. This information is assumed to be most heavily relied upon when the information is not available from the environment as in self-initiated actions or sequential actions when only the first element is cued externally. The selective disinhibition mechanism is implemented by intrinsic processing in the basal ganglia and exerts an influence on cortical motor processing via the pallido-thalamo-cortical circuit which terminates in MI and SMA.

Biological evidence for the role of the basal ganglia in responding to environmental events which are relevant to behavioural actions comes primarily from the work of Shultz and colleagues which was reviewed in chapter 2. The notion of an internal context signal which extracts information from the environment is introduced to account for the Parkinson's disease data that voluntary or 'memory' guided movements are more profoundly affected than stimulus elicited movements and that sequential movements (which inherently have a high programming load) are more affected than simple movements. These features of the disorder cannot be accounted for if the system is adapted for responding only to external events. Although biological evidence which directly supports the notion of a 'frontal' contextual signal is not available there have been a number of similar suggestions regarding the role of frontal cortex in behaviour (Goldman-Rakic, 1995; Shallice, 1988).

One of the properties of a temporal context signal for the control of sequential action is that states of the signal that are nearby in time should be similar to each other, thus in a system with inherent parallelism future sequence elements should be partially pre-activated in order to account for co-articulation effects in human sequential behaviour. This property gives rise to the need for the selective disinhibition mechanism which effects a form of response competition to prevent conflict between similar action representations at output. In the conceptual model we suggest that there is a 'trade-off' between the two systems such that when a good quality context signal is available the need for an efficient action selection mechanism is reduced because there is less potential conflict between action

representations. However, if the context signal is degraded or of poor quality, there is a heavier dependency on the action selection mechanism.

In the conceptual model we suggested in common with previous hypotheses that dopamine performs a contrast enhancement function and that reduction of dopamine in PD acts to either over-inhibit all action representations or to allow undesirable 'co-activation' of more than one action representation. These computational level dysfunctions may account for a variety of the deficits exhibited in that disorder.

10.1.2 Developing a computational approach

The aim of the modelling work in the thesis was to implement a system capable of controlling sequential behaviour which is based on the computational principles put forward in the conceptual model. The response of the system to damage could then be assessed. There are two distinct approaches to modelling sequential behaviour using a contextual control signal of the type proposed above: recurrent networks in which the state representation acts as a context signal, and networks using endogenous oscillators to form a dynamic context vector. In the course of the work presented in the thesis we explored the properties of both types of network architecture and interpreted their behaviour in relation to the hypotheses embedded in the conceptual model. The main findings and limitations revealed in the simulation studies conducted with both types of network are summarised in the sections below.

10.1.3 Damaging the forward model in a Jordan sequential network

The architecture we used to implement the recurrent network approach was the Jordan (1990) architecture which utilises a forward model to learn to control a simulated articulated limb. After learning a sequence the forward model outputs an estimate of the task space location of the endpoint of the arm. This is fed back to the state (context) units via recurrent connections.

One of the main features of the Jordan model which first attracted our interest, in addition to the capacity to produce sequential behaviour, was the forward model. The

notion of a subsystem which provides a predictive internal estimate of movement outcomes appeared to be closely allied to our hypotheses of the role of frontal cortex in the internal control of action. By virtue of the recurrent connections from forward model output to the state representation of the controller subnetwork, the output of the forward model is central to the control of sequential actions, whilst not being directly involved in the computation of limb movements.

The aim of the first set of simulations conducted in chapter five was to investigate the effects of damaging the forward model by reducing the contrast enhancement properties of its hidden units. This is simply achieved by reducing the gain on the activation function. The effects of damaging the processing of the network in this way were compared to other forms of damage which were not motivated by hypotheses of dopamine depletion. The criteria for modelling parkinsonian deficits were that the model should exhibit *reductions* in the form of undershoots to targets rather than increased variability about target locations and secondly that the deficits should get worse as the sequence progresses. Although the model did comply to both of these criteria *only* when the gain was reduced in the forward model (thus providing intial support for the hypothesis), there were a number of problems with both the architecture and the modelling framework which render these simulations rather unsatisfactory.

First any attempt to model aspects of dynamic behaviour with this model are hampered by the lack of a *temporal competence*. The time course of processing in the model consists of only the discrete time steps that are identified with the desired location of the endpoint for each target in the sequence. Not only is the behaviour of the limb at intermediate time steps impossible to access, but the transition of the system from one sequence element to the next is imposed on the network by the structure of the training data rather than being a property of the intrinsic dynamics in the processing of the system. Thus, whatever damage is inflicted on the system there will never be any temporal variability in sequential performance.

Another problem with the simulations reported in this chapter was that when only one sequence was learned examination of the kinematic configuration of the limb revealed

that the network opted for a solution in which the joint angles remained static. This problem was partially solved by using a rest constraint on the motion of the shoulder, but comparatively poor learning was achieved as a result. This reduces the generalisability of the results produced. Other investigators who have used the Jordan architecture have reported difficulties in learning certain sequences (Vousden, 1996) and this appears to be a general limitation of the architecture.

10.1.4 Attractor dynamics in a cascaded Jordan architecture

In chapter 6 we introduced a modification to the standard Jordan model, by incorporating the cascade equations developed by McClelland (1979). The primary motivation for this modification was to introduce some temporal aspect to network processing and to allow network dynamics alone to determine when the transition between sequence elements occurs. This could be achieved allowing activation to cascade through the system terminating at the state units. Only when the state units have reached an asymptotic level of activation can processing of the next sequence element begin. Using this method with a sufficiently small value for the cascade rate (tau) causes the network to trace the trajectory of the limb endpoint between each target state in the sequence (figure 6.1). This scheme is interesting because the network encodes a virtual trajectory between equilibrium points at the target locations. We discuss an interpretation of this behaviour in terms of the equilibrium point hypothesis of motor control.

The bulk of the work in this chapter, however, is inspired by the finding reported by Jordan (1986a, 1986b) that sequential networks encode sequences as limit cycle attractors. The notion that movement control can be characterised as a dynamical system is a growing theme in motor control theory. We introduce a dynamical systems model of normal and parkinsonian motor control which is based on the Hopf bifurcation in a nonlinear dynamical system. In the subsequent simulations we implement an alternative version of the cascaded network in which activation is cascaded through the state units, thus allowing activation to cycle round the entire network. This produced the surprising and interesting result that the shape of the entire attractor for a given sequence is displayed.

These results also showed that the attractor is more complex than a simple limit cycle. Slight alterations of the control parameter μ can yield a stable limit cycle attractor. A study of the systematic variation of μ through a range of values demonstrated that the behaviour of this system, when a single sequence has been learned, approximates the behaviour of the hopf bifurcation model. As μ is reduced towards zero the attractor collapses to a fixed point.

This system provides the basis for a series of simulations in which the effects of a second learned sequence on the qualitative characteristics of the attractor encoding the test sequence are examined. The main finding is that the task-space proximity of a second attractor is an important factor in determining both the extent and nature of its influence on the test attractor. In particular, an overlapping second attractor produces qualitatively different effects to a non-overlapping attractor. These findings provide explicit computational demonstrations of a number of appealing ideas in motor control. The first of these is that patterns of movement may be represented as dynamic attractors. However, whilst it may turn out to be a useful property that attractors have the potential to interact with each other - for example in the generation of novel movement patterns, it is also possible that the influence of proximal attractors may cause undesirable interference effects particularly if the system is disrupted, as in motor disorders.

In a final series of simulations we investigated the effects of selectively reducing the cascade rate in the forward model as a means of simulating a degraded contextual signal provided by the state units, thus increasing the inherent level of competition between attractors. The results of this manipulation were quite striking. When only one sequence had been learned, reducing the cascade rate slightly had minimal effect on the qualitative behaviour of the system. However in both conditions where a second sequence had been learned the system spirals to a fixed point attractor near the centre of the test sequence. These results demonstrate the 'parkinsonian' behaviour that is characterised in the Hopf bifurcation model presented in this chapter.

10.1.5 The use of recurrent networks in modelling motor disorders

Whilst these simulation results exhibit general properties which are appealing in the context of the dynamical systems interpretation presented here, a limitation of this approach is that the characterisation of parkinsonian symptoms is postulated at a rather abstract level. The deficits that we have modelled here are more relevant to high-level features of PD such as the failure to sustain motor activity during repetitive tasks, thus leading to progressive reductions in movement amplitude. A typical example of this is the 'festination' exhibited in parkinsonian locomotion. These simulations are not directed towards the sorts of deficits that have been found in sequencing tasks with PD's.

One of the problems with computational modelling of movement control data as opposed to purely cognitive phenomena is that extreme simplifying assumptions have to be made about the elements of the effector system that are included in the model. This is problematic if the aim is to model data produced in behavioural studies. For example, the characterisation we have used here is an articulated kinematic limb. This makes it impossible to appeal to data that has been collected using dynamical measures or EMG profiles. Even if an attempt were made to construct a simplified representation of muscle contractions, it is hard to see how it would be sufficiently similar to allow comparison with real data.

In chapter 7 we adopted an alternative approach which was to dispense with the model of the articulated limb encoded in the forward model. Instead we used a simple Jordan architecture in which the output representations encoded the endpoints of two "virtual" limbs, in which no kinematic constraints are assumed. Using this dual-effector system we showed that in the absence of explicit inhibitory mechanisms there is undesirable 'shadowing' exhibited by the unselected effector in a model of selective single effector performance. This is a feature of Parkinson's disease reported by Brentari and Poizner (1994) in their study of a deaf parkinsonian signer. This simulation is used to argue that deficiency in an action selection subsystem may produce disrupted performance that could not be achieved by degrading the quality of the context signal alone.

A further unsatisfactory feature of the Jordan sequential network is the use of recurrent connections. It has been argued on a number of grounds that this architecture implements a chaining system. We tested an alternative version of the Jordan architecture which uses endogenous oscillators to implement the state representation and found that it exhibited similar properties to those exhibited when recurrent connections were used. This indicates that a context signal for driving sequential behaviour that is not subject to criticisms of chaining can be used instead without loss of the desirable properties demonstrated in the Jordan network. These two final criticisms in addition to other limitations of the Jordan architecture already mentioned provide part of the motivation for the second architecture we examined - the OSCAR model.

10.1.6 Endogenous oscillators and the control of sequential behaviour

There is a wide range of evidence to suggest that ensembles of oscillators are applicable to dynamical processes associated with perception and action (Turvey & Carello, 1995). In particular, there are biological and computational motivations for the suggestion that they form an integral part of processes controlling sequential action (Saltzman, 1995). One problem with the use of oscillatory dynamics to provide a context signal in the Jordan model presented in chapter 7, and other approaches such as the Houghton (1990) competitive queuing model is the low dimensionality of the control signal produced by simple oscillators. A different approach which is used in the OSCAR model is to construct a high dimensional fast moving context signal from multiple oscillators operating at different frequencies. OSCAR has already been used to account for evidence in serial short term memory (Brown, Preece & Hulme, 1996) and speech production (Vousden, 1996). We use it here as a basis for modelling the control of sequential action. In a series of simulations we show how the context signal permits parallel pre-activation of several sequence elements where those that are associated with similar contextual states are most highly activated. A further advantage of OSCAR over other models based on similar computational principles (e.g. Burgess & Hitch, 1992) is that both the context signal representation and the action representations are distributed. This is an important

characteristic if the model is to appeal to the notion of "brain-style" processing which is one of the most widely cited advantages of modelling with connectionist networks.

10.1.7 Modelling basal ganglia dysfunction in PD with OSCAR

In the final simulation chapter of this thesis we apply the OSCAR model to the conceptual model of cortico-basal ganglia interaction put forward in chapter four. The oscillator based context vector is used to simulate the 'frontal' control signal and a simple competitive process using lateral excitatory and inhibitory connections between action representations is used to simulate the selective disinhibition mechanism. By plotting the parameter space of effective action selection with a medium-speed context vector, it is shown that a balance is required in the excitatory and inhibitory parameters controlling action selection such that only one action is selected. This simulation relates to the biologically inspired notion that dopamine acts to maintain a balance between the direct and indirect pathways during the control of posture and voluntary movement. If there is a high relative level of inhibition then a static posture is maintained and no action is selected. In contrast high relative excitation can cause several actions to become concurrently active. Both cases can be interpreted as system disruptions underlying motor deficits in PD. Over-inhibition could account for the lack of spontaneous movement that is characteristic of hypokinesia, whereas co-activation of competing action representations could be the basis of hypometric movements, and co-activation of antagonistic muscle groups. Low absolute values of both excitation and inhibition could, in addition, provide an account of prolonged movement onset times in PD. These deficits interact with the quality of the context signal. A degraded contextual signal yields a reduction in the available parameter space in which effective action selection can be achieved and also causes an increase in the number of cycles required to select an action. The simulations in this chapter provide an explicit computational demonstration of the ideas expressed in chapter four.

10.2 Implications for theories of basal ganglia function

At the beginning of this thesis we criticised hypotheses of basal ganglia function that are grounded in traditional approaches to motor programming for failing to meet the demands of current knowledge about the organisation and processing capabilities of real nervous systems. We also criticised those 'black-box' models in which the hypothesised control flow between control structures is labelled, but the nature of processing assumed within structures is left unspecified. In this thesis we have attempted to improve on these methods of theorising, first by making our computational assumptions explicit and second by examining their properties in computer simulation studies.

The main thrust of our theoretical approach is based on the linked concepts of a dynamic internal representation of temporal context for the control of movement sequencing and the principle of general inhibition and selective disinhibition in the control of action selection. These theoretical primitives are motivated by both biological and computational considerations as was indicated in review chapters two and three. These two mechanisms have been combined in the current approach in an attempt to enhance current understanding of the role of basal ganglia and frontal cortex in motor programming processes and to provide an account of why sequential actions are especially affected in Parkinson's disease. This issue was most coherently addressed in the simulations using the OSCAR architecture in chapter nine and represents a considerable advance over hypotheses which have not been implemented.

Of course, the biological picture of basal ganglia and cortical interaction that we have presented here has been considerably simplified. For example, there are substantial projections from limbic cortex and the amygdala to portions of the ventral striatum (Graybiel & Kimura, 1995). These areas are thought to play an important role in motivation. This is a facet of basal ganglia function which has potentially significant consequences for hypotheses of its role in the initiation and control of action which has, of necessity, been omitted from the current approach. In addition there are numerous features of striatal organisation such as the functional distinction between patch and matrix compartments (Graybiel, 1990) which have not been considered here and would more

properly be treated by a modelling approach aimed at lower level computational features. Probably the most significant simplifying assumption that we have made in the current approach concerns the systems level action of dopamine. The notion that it serves to increase the contrast between competing response alternatives is founded on several sources of information reviewed in chapter two and the computational approaches of Cohen and Servan-Shreiber (1992) and Wickens, Alexander and Miller (1991). It is likely, however, that this assumption may need to be re-examined as better models of dopaminergic action are developed.

10.3 The future of computational models in motor control theory

A second theoretical strand which was addressed here using the cascaded Jordan architecture is the nature of representation of motor patterns and the patterns of interaction and interference which might occur in a complex distributed system. A number of recent theoretical approaches have advocated the idea that movements may be represented as attractors in a dynamical system. Our simulations in chapter six illustrate the potential for catastrophic interference if multiple dynamic attractors are present in a single system. The simulations produced here provide a platform for further research on this issue, although given the various problems associated with the Jordan architecture it is likely that other computational architectures may prove to be more useful in this respect.

One suggestion that arises directly from the work presented in this thesis is the use of interacting systems of oscillators. Here we have used oscillators to implement the contextual representation which provides input to action representations for internal control. An ecologically conservative extension of this idea is to use oscillators as a basis for representing action patterns as well. The notion of using coupled sets of oscillators to model the nonlinear dynamics of movement patterns is already well established within the perception-action approach (Schöner & Kelso, 1988). There are a number of potent issues in neural computation which remain to be addressed in systems with dynamical representations of actions, including how such a system might learn to control sequences. In such a framework, dynamical systems analysis of the behaviour of networks using the

associated concepts of stability and bifurcations are likely to become increasingly important in motor control theory. A clear application of this approach in the domain of Parkinson's disease lies in modelling abnormalities in the performance of repetitive movements and time estimation.

10.4 Concluding Remarks

In this thesis we have developed a computationally explicit model of how disruption of the basal ganglia affects the control of voluntary movements in Parkinson's disease. The model makes specific and concrete assumptions about the function of cortical and basal ganglia structures in the control of sequential actions, and about the processes that are involved in 'motor programming' of such actions. Two different computational architectures have been tested with the same underlying computational intuitions. Both architectures showed performance deficits which are consistent with features of Parkinson's disease when control processes assumed to be implemented by cortico-striatal circuitry were damaged. Furthermore some general insights into the properties of these processes in the control of motor behaviour have been obtained.

One of the current overriding goals of motor control research, and of cognitive neuroscience more generally, is to reconcile the need to explicitly characterise the nature of internal control structures and processes used in the control of voluntary movements with the principles governing the operation of biological systems. The theoretical and modelling work presented in this thesis has been pursued with that goal in mind. In so doing it forms part of a broader research enterprise in which constraints from biology, psychology and computational theory are brought together to develop a coherent approach towards fulfilling this objective.

References

Abernethy, B., & Sparrow, W.A. (1992). The rise and fall of dominant paradigms in motor behaviour research. In J.J. Summers (Ed.), *Approaches to the study of motor learning*. Amsterdam: Elsevier Science.

Albin, R.L., Young, A.B., & Penney, J.B. (1989). The functional anatomy of basal ganglia disorders. *Trends in Neuroscience, 12*, 366-375.

Alexander, G.E., & Crutcher, M.D. (1990). Functional architecture of basal ganglia circuits: neural substrates of parallel processing. *Trends In Neuroscience, 13*, 266-271.

Alexander, G.E., & Crutcher, M.D. (1989). Coding in spatial rather than joint coordinates of putamen and motor cortex preparatory activity preceding planned limb movements. In A.R. Crossman, & M.A. Sambrook (Eds.), *Neural mechanisms of movement disorders*. John Libbey & Co.

Alexander, G.E., Delong, M.R., & Crutcher, M.D. (1992). Do cortical and basal ganglionic motor areas use "motor programs" to control movement. *Behavioral and Brain Sciences, 15*, 656-665.

Alexander, G.E., Delong, M.R., & Strick, P.L. (1986). Parallel organisation of functionally segregated circuits linking basal ganglia and cortex. *Annual Review of Neuroscience, 9*, 357-381.

Apicella, P., Scarnati, E., Ljungberg, T., & Shultz, W. (1992). Neuronal activity in monkey striatum related to the expectation of predictable environmental events. *Journal of Neurophysiology, 68*, 945-960.

Arbib, M.A. (1990). Programs, schemas, and neural networks for control of hand movements: Beyond the RS framework. In M. Jeannerod (Ed.), *Attention and Performance XIII: Motor Representation and Control*, (pp. 111-138). Hillsdale, N.J: Erlbaum.

Baker, J.R., Davey, N.J., Ellaway, P.H., & Friedland, C.L. (1992). Short term synchrony of motor unit discharge during weak isometric contraction in Parkinson's disease. *Brain, 115*, 137-154.

Bapi, R.S., & Levine, D.S. (1990). Networks modeling the involvement of the frontal lobes in learning and performance of flexible movement sequences. In *Proceeding of the 12th annual conference of the cognitive science society*, (pp. 915-922). Cambridge, MA.

Benecke, R., Rothwell, J., Dick, B., & Marsden, C.D. (1986). Performance of simultaneous movements in patients with Parkinson's disease. *Brain, 109*, 739-757.

Benecke, R., Rothwell, J., Dick, B., & Marsden, C.D. (1987). Disturbance of sequential movements in patients with Parkinson's disease. *Brain, 110*, 361-379.

Berardelli, A., Dick, J.P., Rothwell, J.C., Day, B.L. & Marsden, C.D. (1986). Scaling the size of the first agonist burst during rapid wrist movements in patients with

Parkinson's Disease. *Journal of Neurology, Neurosurgery and Psychiatry, 49,* 1273-1279.

Bernstein, N. (1967). *The Co-ordination and Regulation of Movements.* Oxford, England: Pergamon.

Berridge, K.C. (1989). Substantia nigra 6-OHDA lesions mimic striatopallidal disruption of syntactic grooming chains: A neural systems analysis of sequence control. *Psychobiology, 17,* 377-385.

Berridge, K.C., & Whishaw, I.Q. (1992). Cortex, striatum and cerebellum - control of serial order in a grooming sequence. *Experimental Brain Research, 90,* 275-290.

Beuter, A., Labrie, C., & Vasilakos, K. (1991). Transient dynamics in motor control of patients with Parkinson's disease. *Chaos, 1,* 279-285.

Bizzi, E., Muzza-Ivaldi, F.A., Giszter, S. (1991). Computations underlying the execution of movement: A biological perspective. *Science, 253,* 287-291.

Bizzi, E., Hogan, N., Muzza- Ivaldi, F.A., & Giszter, S. (1992). Does the nervous system use equilibrium -point control to guide single and multiple joint movements? *Behavioral and Brain Sciences, 15,* 603-613.

Bloom, F.E., Shulmann, J.A., & Koob, G.F. (1989). Catecholamines and behaviour. In U. Trendelenburg & N. Wiener (Eds.), *Handbook of experimenta pharamcology* (Vol. 90, pp.27-88). Berlin: Springer-Verlag.

Bloxham, C.A., Mindel, T.A., & Frith, C.D. (1984)..Initiation and execution of predictable and unpredictable movements in Parkinson's disease. *Brain, 107,* 371-384.

Borrett, D.S., Yeap, T.H., & Kwan, H.C. (1992). The non-linear dynamics of connectionist networks - The basis of motor control. *Behavioral and Brain Sciences, 15,* 712-714.

Brentari, D., & Poizner, H. (1994). A phonological analysis of a deaf parkinsonian signer. *Language and Cognitive processes, 9,* 69-99.

Bronstein, A.M., Hood, J.D., Gresty, M.A., & Panagi, C. (1990). Visual control of balance in cerebellar and parkinsonian syndromes. *Brain, 113,* 767-779.

Bronstein, A.M., & Kennard, C. (1985). Predictive ocular motor control in parkinson's disease. *Brain, 108,* 925-940.

Brososki, T.J., Brown, R.M., Rosvold, H.E., & Goldman, P.S. (1979). Cognitive deficit caused by regional depletion of dopamine in prefrontal cortex of rhesus monkey. *Science, 205,* 929-931.

Brown, G.D.A., Britain, A. A., Elvevag, B., & Mitchell, I.J. (1994). A computational approach to fronto-striatal dysfunction in schizophrenia and Parkinson's disease. In M. Oaksford & G.D.A. Brown (Eds.), *Neurodynamics and Psychology.* London: Academic Press.

Brown, G.D.A., Preece, T., & Hulme, C. (1996). Oscillator based memory for serial order. Manuscript submitted.

Brown, L.L. (1992) Somatotopic organisation in rat striatum: evidence for a combinatorial map. *Proc. Natl. Acad. Sci. USA, 89,* 7403-7407.

Brown, L.L., Feldman, S.M., Divac, I., Hand, P.J. & Lidsky, T.L. (1994). A distributed network of context-dependent functional units in the rat neostriatum. In G. Percheron et al. (Eds.),*The Basal Ganglia IV.* New York: Plenum Press.

Burgess, N., & Hitch, G. J. (1992). Toward a network model of the articulatory loop. *Journal of Memory and Language, 31,* 429-460.

Caramazza, A., Miceli, G., Villa, G., & Romani, C. (1987). The role of the graphemic buffer in spelling: evidence from a case of acquired dysgraphia. *Cognition, 26,* 59-85.

Carpenter, G.A., & Grossberg, S. (1987). A massively parallel architecture for a self-organising neural pattern recognition machine. *Computer Vision, Graphics, and Image Processing, 37,* 54-115.

Changeux, J.P., & Dehaene, S. (1994). Neuronal models of cognitive functions associated with prefrontal cortex. In M. Oaksford & G.D.A. Brown, (Eds.), *Neurodynamics and Psychology,* (pp. 13-33). London: Academic Press.

Chater, N., & Conkey, P. (1994). Sequence processing with recurrent neural networks. In M. Oaksford & G.D.A. Brown, (Eds.), *Neurodynamics and Psychology.* London: Academic Press.

Chevalier, G., & Deniau, J. M. (1990). Disinhibition as a basic process in the expression of striatal functions. *Trends in Neuroscience, 13,* 277-280.

Chiodo, L.A., & Berger, T.W. (1986).Interactions between dopamine and amino acid-induced excitation and inhibition in the striatum. *Brain Research, 375,* 198-203.

Church, R.M., & Broadbent, H. (1990). Alternative representations of time, number, and rate. *Cognition, 37,* 55-81.

Cohen, J.D., Dunbar, K., & McClelland, J.L. (1990). On the control of automatic processes : A parallel distributed processing account of the stroop effect. *Psychological Review, 97,* 332-361.

Cohen, J.D., & Servan-Shreiber, D. (1989). *A PDP approach to behaviour and biology in schizophrenia.* (Tech. Rep. AIP-100). Department of psychology, Carnegie-Mellon, Pittsburgh.

Cohen, J.D., & Servan-Shreiber, D. (1992). Context, Cortex, and Dopamine: A connectionist approach to behaviour and biology in schizophrenia. *Psychological Review, 99,* 45-77.

Connolly, C.L., & Burns, J.B. (1995). A state-space striatal model. In J.C. Houk, J.L. Davis, & D.G. Beiser (Eds.), *Models of Information Processing in the Basal Ganglia,* (pp.163-179). Cambridge, MA: MIT Press.

Connor, N.P., Abbs, J.H., Cole, K.J., & Gracco, V.L. (1989). Parkinsonian deficits in serial multiarticulate movements for speech. *Brain, 112,* 998-1009.

Cooke, J.D., Brown. J.D., & Brooks, V.B. (1978). Increased dependence on visual information for movement control in patients with Parkinson's disease. *Canadian Journal of Neurological Science, 5,* 413-415.

Cools, A.R., Jaspers, R., Schwartz, M., Sontag, K.H., Vrijmoed, M., Van den Bercken, J. (1984). Basal ganglia and switching motor programs. In. J.S. McKenzie, R.E. Kemm, & L.N. Wilcock (Eds.),*The Basal Ganglia: Structure and Function.* New York: Plenum Press.

Cottrell, G.W., Munro, P.W., & Zipser, D. (1987). Image compression by back propagation: A demonstration of extensional programming. In N.E. Sharkey (Ed.), *Advances in cognitive science (Vol. 2).* Chichester, England: Ellis Horwood.

Crawford, T.J., Henderson, L., & Kennard, C. (1989). Abnormalities of nonvisually guided eye movements in Parkinson's disease. *Brain, 112,* 1573-1586.

Cruse, H., Dean, J., Heuer, H., & Schmidt, R.A. (1990). Utilisation of sensory information for motor control. In O. Neumann & W. Prinz, (Eds.), *Relationships between Perception and Action.* Berlin: Springer-Verlag.

Crutcher, M.D., & Delong, M.R. (1984). Single cell studies of the priumate putamen. II. Relations of directions of movements and patterns of muscular activity. *Experimental Brain Research, 53,* 244-258.

Curran, T. (1995). On the neural mechanisms of sequence learning. *PSYCHE, 2*(12).

Day, B.L., Dick, J.P.R., & Marsden, C.D. (1984). Patients with Parkinson's Disease can employ a predictive motor strategy. *Journal of Neurology, Neurosurgery and Psychiatry, 47,* 1299-1306.

De France, J. (1980). Convergence of excitatory amygdaloid and hippocampal inputs to the nucleus accumbens of the rat brain. *Brain Research, 185,* 183-196.

Dehaene, S. (1993). Temporal oscillations in human perception. *Psychological Science, 4,* 264-270.

Dell, G.S. (1986). A spreading activation theory of retrieval in sentence production. *Psychological Review, 93,* 283-321.

Dell, G.S., & Reich, P.A. (1981). Stages in sentence production: An analysis of speech error data. *Journal of verbal learning and verbal behaviour, 20,* 611-629.

Delong, M.R. (1990). Primate models of movement disorders of basal ganglia origin. *Trends in Neuroscience, 13,* 281-285.

Delong, M R., & Georgopoulos, A.P. (1979). Motor functions of the basal ganglia as revealed by studies of single cell activity in the behaving primate. In L. G. Poirier, T. L. Sourkes, & P. J. Bedard, *Advances in Neurology Vol. 24,* New York: Raven.

Denny-Brown, D. (1962). *The Basal Ganglia and their Relation to Disorders of Movement.* Oxford, England: Oxford University Press.

Dick, J.P.R., Rothwell, J., Day, B.L., Cantello, R., Buruma, O., Gioux, M., Benecke, R., Berardelli, A., Thompson, P.D., & Marsden, C.D. (1989). The bereitshaftspotential is abnormal in Parkinson's disease. *Brain, 112,* 233-244.

Donders, F.C. (1869). On the speed of mental processes. *Attention and Performance II: Acta Psychologica,* 1969, *30,* 412-431.

Dudai, Y. (1989). *The Neurobiology of Memory - Concepts, Findings, Trends.* Oxford, England: Oxford University Press.

Elman, J.L. (1990). Finding structure in time. *Cognitive Science, 14,* 179-211.

Elman, J.L., & Zipser, D. (1988). Discovering the hidden structure of speech. *Journal of the Acoustical Society of America, 83,* 1615-1626.

Estes, W.K. (1972). An associative basis for coding and organisation in memory. In A.W. Melton & E. Martin, (Eds.), *Coding processes in human memory* (pp. 161-190). Wahington, DC: Winston.

Feldman, A.G. (1966). Functional tuning of the nervous sytem during control of movement or maintenance of steady posture. II. Controllable parameters of the muscle. *Biophysics, 11,* 565-578.

Filion, M., Tremblay, L., & Bedard, P.J. (1989). Excessive and unselective responses of mdeial pallidal neurons to both passive movement and striatal stimulation in monkeys with MPTP- induced parkinsonism. In A.R. Crossman, & M.A. Sambrook (Eds.), *Neural mechanisms of movement disorders.* John Libbey & Co.

Flaherty, A.W., and Graybiel, A.M. (1993). Two input systems for body representations in the primate striatal matrix - Experimental evidence in the squirrel monkey. *Journal of Neuroscience, 13,* 1120-1137.

Flanders, M., & Soechting, J.F. (1992). Kinematics of typing: parallel activation of two hands. *Journal of Neurophysiology, 67,* 1264-1274.

Flash, T., & Henis, E. (1991). Arm trajectory modifications during reaching towards visual targets. *Journal of Cognitive Neuroscience, 3,* 220-230.

Flash, T., Henis, E., Inzelberg, R., & Korczyn, A.D. (1992). Timing and sequencing of human arm trajectories: Normal and abnormal motor behaviour. *Human Movement Science, 11,* 83-100.

Flash, T., Jordan, M.I., & Arnon, Y. (1994). A model of learning of arm trajectories from spatial targets. *Journal of Cognitive Neuroscience.* Manuscript submitted.

Flowers, K.A. (1976). Visual closed-loop and open-loop characteristics of voluntary movement in Parkinson's disease. *Brain, 112,* 233-244.

Flowers, K.A. (1978). Lack of prediction in the motor behaviour of Parkinsonism. *Brain, 101,* 35-52.

Foote, S.L., Freedman, R., & Oliver, A.P. (1975). Effects of putative transmitters on neuronal activity in monkey auditory cortex. *Brain Research, 86,* 229-242.

Frith, C.D. (1992). *The cognitive neuropsychology of schizophrenia.* Hillsdale, NJ: Erlbaum.

Funahashi, S., Bruce, C.J., & Goldman-Rakic, P.S. (1989). Mnemonic coding of visual space in monkey dorsolateral prefrontal cortex. *Journal of Neurophysiology, 61,* 334-349.

Gallistel, C. (1980). *The Organisation of Action: A new synthesis.* Hillsdale, NJ: Erlbaum.

Gallistel, C. (1990). *The Organisation of Learning.* Cambridge, MA: MIT Press.

Gentner, D.R. (1987). Timing of skilled motor performance: Tests of the proportional duration model. *Psychological Review, 94,* 255-276.

Gentner, D.R., Grudin, J., & Conway, E. (1980). *Finger movements in transcription typing,* (Tech. Rep. No. 8001), La Jolla, CA: University of California, San Diego, Center for Human Information Processing.

Georgopoulos, A.P. (1991). Higher Order Motor Control. *Annual Review of Neuroscience, 14,* 361-377.

Gerfen, C.R., Engber, T., Mahan, L.C., Susel, Z., Chase, T.N., Monsma, F.J., & Sibley, D.R. (1990). D1 & D2 dopamine receptor-regulated gene expression of striato-nigral and striato-pallidal neurons, *Science, 250,* 1429-1432.

Gibson, J.J. (1979). *The Ecological Approach to Visual Perception.* Boston: Houghton Mifflin.

Giszter, S. (1992). Spinal movement primitives and motor programs: A necessary concept for motor control. *Behavioral and Brain Sciences, 15,* 744-745

Goldberg, G. (1985). Supplementary motor area structure and function: Review and hypotheses. *Brain and Behavioral Sciences, 8,* 567-616.

Goldman-Rakic, P.S. (1988). Topography of cognition: Parallel distributed networks in primate association cortex. *Annual Review of Neuroscience, 11,* 137-156.

Goldman-Rakic, P. S. (1995). Toward a circuit model of working memory and the guidance of voluntary action. In J.C. Houk, J.L. Davis, & D.G. Beiser (Eds.), *Models of Information Processing in the Basal Ganglia,* (pp.131-149). Cambridge, MA: MIT Press.

Goldman-Rakic, P.S., & Selemon, L. (1990). New frontiers in basal ganglia research. *Trends In Neuroscience, 13,* 1-13.

Gottlieb, G.L. Latash, M.L., Corcos, D.M., Lubinskas, T.J., & Agarwal, G.C. (1992). Organising principles for single joint movements: agonist v. antagonist interactions. *Journal of Neurophysiology, 67,* 1417-1427.

Grace, A.A. (1983). *Nigral dopamine neurons: Identification, characterisation and effects on autoreceptor stimulation.* Unpublished doctoral dissertation, Yale.

Gray, J.A., Feldon, J., Rawlins, J.N.P., Hemsley, D.R., & Smith, A.D. (1991). The Neuropsychology of Schizophrenia. *Behavioral and Brain Sciences, 14,* 1-84.

Graybiel, A.M. (1990). Neurotransmitters and Neuromodulators in the basal ganglia. *Trends in Neuroscience, 13*, 244-254.

Graybiel, A.M. (1991). Basal ganglia - Input, neural activity, and relation to cortex. *Current Opinion in Neurobiology, 1*, 644-651.

Graybiel, A.M., & Kimura, (1995). Adaptive neural networks in the basal ganglia. In J.C. Houk, J.L. Davis, & D.G. Beiser (Eds.), *Models of Information Processing in the Basal Ganglia*, (pp.103-117). Cambridge, MA: MIT Press.

Grossberg, S. (1978). Behavioral contrast in short-term memory: Serial binary memory models or parallel continuous models? *Journal of Mathematical Psychology, 3*, 199-219.

Hallet, M., & Khoshbin, S. (1980). A physiological mechanism of bradykinesia. *Brain, 103*, 301-314.

Halsband, U., & Freund, H-J. (1993). Motor learning. *Current Opinion in Neurobiology, 3*, 940-949.

Hanson, S.J., & Kegl, J. (1987). Parsnip: A connectionist network that learns natural language grammar from exposure to natural language sentences. *Proceedings of the Ninth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.

Harley, T.A. (1984). A critique of top-down independent levels models of speech production: evidence from non-plan-internal speech errors. *Cognitive Science, 8*, 191-219.

Harrington, D.L., & Haaland, K.Y. (1991). Sequencing in Parkinson's disease: Abnormalities in programming and controlling movement. *Brain, 114*, 99-115.

Haykin, S. (1994). *Neural Networks: A Comprehensive Foundation*. New York: Macmillan.

Hedreen, J.C., & Delong, M.D. (1991). Organisation of striatopallidal, striatonigral and nigrostriatal projections in the macaque. *The Journal of Comparative Neurology, 304*, 569-595.

Henry, F.M., & Rogers, D.E. (1960). Increased response latency for complicated movements and a "memory drum" theory of neuromotor reaction. *Research Quarterly, 31*, 448-458.

Henson. R.N.A., Norris, D.G., Page, M.P.A., & Baddeley, A.D. (1996). Unchained memory: Error patterns rule out chaining models of immediate serial recall. *Quarterly Journal of Experimental Psychology, 49A*, 80-115.

Hick, W.E. (1952). On the rate of gain of information. *Quarterly Journal of Experimental Psychology, 4*, 11-26.

Hinton, G.E., & Sejnowski, T. (1986). Learning and relearning in Boltzmann machines. In J.L. McClelland & D.E. Rumelhart, (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition (Vol. 1)*. Cambridge MA: MIT Press.

Hinton, G.E., & Shalllice, T. (1991). Lesioning an attractor network: Investigations of __. acquired dyslexia. *Psychological Review, 98*, 74-95.

Hikosaka, O. (1991). Basal ganglia - possible role in motor cordination and learning. *Current Opinion in Neurobiology, 1*, 638-643.

Hikosaka, O., & Wurtz, R.H. (1985). Modification of saccadic eye movements by GABA-related substances. II. Effects of muscimol in the monkey substantia nigra pars reticulata. *Journal of Neurophysiology, 53*, 292-308.

Hollerbach. J.M., & Atkeson, C.G. (1987). Deducing planning variables from experimental arm trajectories: Pitfalls and possibilities. *Biological Cybernetics, 56*, 279-292.

Hoover, J.E., & Strick, P.L. (1993). Multiple output channels in the basal ganglia. *Science, 259*, 819-821.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. USA, 79*, 2554-2558.

Hore, T., Meyer-Lohmann, J., & Brooks, V. B. (1977). Basal ganglia cooling disables learned arm movements of monkeys in the absence of visual guidance. *Science, 195*, 584-586.

Horstink, M., Berger, H., van Spaendonck, K., van den Bercken, J., & Cools, A. (1990). Bimanual simultaneous motor performance and impaired ability to shift attention in Parkinson's disease. *Journal of Neurology, Neurosurgery and Psychiatry, 53*, 685-690.

Houghton, (1990). The problem of serial order: A neural network model of sequence learning and recall. In R. Dale, C. Mellish & M, Zock, (Eds.), *Current research in natural language generation*. London: Academic Press.

Houghton, G. (1994). Inhibitory control of neurodynamics: Opponent mechanisms in sequencing and selective attention. In M.Oaksford & G.D.A. Brown (Eds.), *Neurodynamics and Psychology*. London: Academic Press.

Houghton, G., Shallice, T., & Glasspool, G.W. (1994). Spelling and serial recall: Insights from a competitive queueing model. In G.D.A. Brown & N.C. Ellis (Eds.), *Handbook of Spelling* (pp. 365-404). Chichester, England: Wiley.

Houghton, G., & Hartley, T. (1995). Parallel models of serial order: Lashley revisited. *Psyche,2,25*.

Houk, J.C.(1992). *Learning in modular networks,* (Tech. Rep. No.7). Chicago, Illinois: Northwestern University Mdeical School.

Hounsgaard, J. M., & Midtgaard, J. (1989). Dendrite processing in more ways than one. *Trends in Neuroscience, 12*, 313-315.

Iberall, T. (1987). *A neural model of human prehension.* Unpublished doctoral dissertation, Department of Computer and Information Science, University of Massachusetts, Amherst,

Ito, M. (1989). *Neural Programming*. Tokyo: Japan Scientific Sciences Press.

Jacobs, R.A., Jordan, M.I., Nowlan, S., & Hinton, G.E. (1991). Adaptive mixtures of local experts. *Neural Computation, 3,* 79-87.

Jackson, S., Marocco, R., & Posner, M.I. (1994). Networks of anatomical areas controlling visuospatial attention, *Neural Networks, 7,* 925-944.

Jackson, S., & Houghton, G. (1995). Sensorimotor selection and the basal ganglia: A neural network model. In J.C. Houk, J.L. Davis, & D.G. Beiser (Eds.), *Models of Information Processing in the Basal Ganglia,* (pp.337-369). Cambridge, MA: MIT Press.

Jahanshahi, M., Brown, R.G., & Marsden, C.D. (1992). Simple and choice reaction time and the use of advance information for motor preparation in Parkinson's disease. *Brain, 115,* 539-564.

Jeannerod, M. (1988). *The Neural and Behavioural Organisation of Goal-Directed Movements*. Oxford, England: Clarendon.

Jennings, P.J. (1995). Evidence of incomplete programming in Parkinson's disease. *Journal of Motor Behaviour, 27,* 310-324.

Jordan, M.I. (1985). *The learning of representations for sequential performance*. Unpublished doctoral dissertation, University of California, San Diego.

Jordan, M.I. (1986a). *Serial order: A parallel distributed processing approach* (Tech. Rep. No. 8604). La Jolla: University of California, San Diego.

Jordan, M.I. (1986b). Attractor dynamics and parallelism in a connectionist sequential machine. *Proceedings of the Eighth Annual Conference of the Cognitive Science Society*, Hillsdale, NJ: Erlbaum.

Jordan, M.I. (1988). *Supervised learning and systems with excess degrees of freedom* (COINS Tech. Rep. 88-27). Amherst, MA: University of Massachusetts, Computer and Information Sciences.

Jordan, M.I. (1990). Motor Learning and the Degrees of Freedom Problem. In M. Jeannerod (Ed.), *Attention and Performance XIII: Motor Representation and Control*. Hillsdale, N.J: Lawrence Erlbaum.

Jordan, M.I. (1992). Constrained supervised learning. *Journal of Mathematical Psychology, 36,* 396-425.

Jordan, M.I., & Jacobs, R.A. (1990). Learning to control an unstable system with forward modeling. In D. Touretsky (Ed.), *Advances in Neural Information Processing Systems* (vol. 2). San Mateo, CA: Morgan Kaufmann.

Jordan, M.I., & Rumelhart, D.E. (1992). Forward models: supervised learning with a distal teacher. *Cognitive Science, 16,* 307-354.

Kato, M., & Kimura, M. (1992) The effects of reversible blockade of basal ganglia on a voluntary arm movement. *Journal of Neurophysiology, 62,* 1516-1534.

Kawato, M., Furukawa, K. & Susuki, R. (1987) A hierarchical neural network model for control and learning of voluntary movement. *Biological Cybernetics, 57,* 169-185.

Keele, S.W. (1968). Movement control in skilled motor performance. *Psychological Bulletin, 70,* 387-404.

Keele, S.W., Cohen, A., & Ivry, R. (1990). Motor programs: Concepts and Issues. In M. Jeannerod, (Ed.), *Attention and Performance XIII: Motor Representation and Control,* Hillsdale, N.J: Lawrence Erlbaum.

Kelso, J.A.S. (1982). Epilogue: Two strategies for investigating action. In J.A.S. Kelso, (Ed.), *Human Motor Behaviour: An Introduction.* Hillsdale, NJ: Erlbaum.

Kelso, J.A.S., & Ding, M. (1993). Fluctuations, intermittency and controllable chaos in biological coordination. In K. M. Newell & D. M. Corcos, (Eds.), *Variability in Motor Control.* Champaign: Human Kinetics Publishers.

Kermadi, I., Jurquet, Y., Arzi, M., & Joseph, J.P. (1993). Neural activity in the caudate nucleus of monkeys during spatial sequencing. *Experimental Brain Research, 94,* 352-356.

King, R., Raese, J.D., & Barchas, J.D. (1981). Catastrophe theory of dopamine transmission: A revised dopamine hypothesis of schizophrenia. *Journal of Theoretical Biology, 92,* 373-400.

Kita, H. (1992). Responses of Globus Pallidus neurons to cortical stimulation: Intracellular study in the rat. *Brain Research, 589,* 84-90.

Kolb, B., & Whishaw, I.Q. (1990). *Fundamentals of human neuropsychology.* New York: Freeman.

Kopell, N. & Ermentrout, G.B. (1990). Structure and function in an oscillating neural network. In E.L. Schwartz, (Ed.), *Computational Neuroscience* (pp. 201-209). Cambridge, MA: MIT Press.

Kornblum, (1965). Response competition and / or inhibition in two choice reaction time. *Psychonomic Science, 2,* 55-56.

Kornhuber, H.H. (1971). Motor functions of the cerebellum and basal ganglia. The cerebello saccadic (ballistic) clock, the cerebellonuclear hold regulator, and the basal ganglia ramp (voluntary speed smooth movement) generator. *Kybernetik, 8,* 157-173.

Kugler, P.N., Kelso, J.A.S., and Turvey, M.T. (1982). On the control and coordination of naturally developing systems. In J.A.S. Kelso & J.E. Clark (Eds.), *The development of movement control and coordination* (pp. 5-78), New York: Wiley.

Kunzle, H. (1975). Bilateral projections from precentral motor cortex to the putamen and other parts of the basal ganglia. An autoradiographic study in *macaca fascicularis. Brain Research, 88,* 195-209.

Lashley, K.S. (1917). The acuracy of movement in the absence of excitation from the moving organ. *American Journal of Physiology, 43,* 169-194.

Lashley, K.S. (1951). The problem of serial order in behaviour. In L.A. Jefress (Ed.), *Central Mechanisms in Behaviour*. New York, Wiley.

Levine, D.S. (1986). A neural network theory of frontal lobe function. In *Proceedings of the 8th Annual Conference of the Cognitive Science Society*, (pp.716-727), Amherst, MA.

Ljungberg, T., Apicella, P., & Schultz, W. (1991). Responses of monkey midbrain dopamine neurons during delayed alternation performance. *Brain Research, 567*, 337-341.

Lorenz, K. (1977). *Behind the Mirror: a search for a natural history of human knowledge*. London: Methuen.

Marsden, C.D. (1982). The mysterious motor function of the basal ganglia: The Robert Wartenberg lecture. *Neurology, 32*, 514-539.

Marteniuk, R.G. MacKenzie, C.L., Jeannerod, M., Athenes, S., & Dugas, C. (1987). Constraints on human arm movement trajectories. *Canadian Journal of Psychology, 41*, 365-378.

Marteniuk, R.G. (1992). Issues in goal directed motor learning: Feedforward control, motor equivalence, specificity and artificial neural networks. in G. E. Stelmach, & J. Requin (eds.), *Tutorials in motor behaviour(II)*. B.V.:Elsevier Science.

Martin, J.P. (1967). *The Basal Ganglia and Posture*. London: Pitman Medical.

Masson, G., & Pailhous, J. (1992). Locomotion, oscillating dynamic systems and stiffness regulation by the basal ganglia. *Behavioural and Brain Sciences, 15*, 778-779.

Matsumura, M., Sawaguchi, T., & Kubota, K. (1992). GABAergic inhbition of neuronal activity in primate motor and premotor cortex during voluntary movement. *Journal of Neurophysiology, 67*, 692-702.

Matsumura, M., Sawaguchi, T., Oishi, T., Ueki, K., & Kubota, K. (1991). Behavioural deficits induced by local injection of bicuculline and muscimol into primate motor and premotor cortex. *Journal of Neurophysiology, 65*,1542-1553.

McClelland, J.L. (1979). On the time relations of mental processes: An examination of systems of processes in cascade. *Psychological Review, 86*, 287-330.

McClelland, J.L. (1992). Toward a theory of information processing in graded, random and interactive networks. In D.E. Meyer & S. Kornblum (Eds.), *Attention and Performance XIV: A synthesis of experimental psychology, artificial intelligence and cognitive neuroscience*, (pp. 655-688). Cambridge, MA: MIT Press.

McClelland, J.L., & Rumelhart, D. E. (1986). Amnesia and distributed memory. In J.L. McClelland, & D.E. Rumelhart (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition (Vol. 2)*, (pp. 503-527). Cambridge, MA: MIT Press.

McClelland, J.L., & Rumelhart, D.E. (1988). *Explorations in Parallel Distributed Processing: A Handbook of Models, Programs, and Exercises*. Cambridge, MA: MIT Press.

McCloskey, M., & Cohen, N.J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. In G.H. Bower (Ed.), *The Psychology of Learning and Motivation,* (Vol. 24) New York: Academic Press.

Mink, J.W., & Thach, W.T. (1991a). Basal ganglia motor control. I. Nonexclusive relation of pallidal discharge to five movement modes. *Journal of Neurophysiology, 65,* 273-300.

Mink, J.W., & Thach, W.T. (1991b). Basal ganglia motor control. II. Late pallidal timing relative to movement onset and inconsistentpallidal coding of movement parameters. *Journal of Neurophysiology, 65,* 301-329.

Mink, J.W., & Thach, W.T. (1993). Basal ganglia intrinsic circuits and their role in behaviour. *Current Opinion in Neurobiology, 3,* 950-957.

Mitchell, I. J. (1990). Striatal outputs and dyskinesia. *Journal of Psychopharmacology, 4,* 188-197.

Morasso, P., & Sanguineti, V. (1992). Neural models of distributed motor control. In G. E. Stelmach, & J. Requin (Eds.), *Tutorials in Motor Behaviour* (II),Amsterdam: Elsevier Science.

Mozer, M.C. (1993). Neural net architectures for temporal sequence processing. In A. Weigend & N. Gershenfeld (Eds.), *Predicting the future and understanding the past.* Redwood City, CA: Addison-Wesley Publishing.

Mozer, M.C., & Behrmann, M. (1990). On the interaction of selective attention and lexical knowledge: A connectionist account of neglect dyslexia. *Journal of Cognitive Neuroscience,2*(2), 96-123.

Mpsitos, G.J., & Soinila, S. (1993). In search of a unified theory of biological organisation: What does the motor system of a sea slug tell us about human motor integration: In K.M. Newell & D.M. Corcos (Eds.),*Variability in Motor Control.* Champaign: Human Kinetics Publishers.

Mulder, G., Smid, H.G., & Mulder, L.J.M. (1992). On the transfer of partial information between perception and action. In D.E. Meyer & S. Kornblum (Eds.), *Attention and Performance XIV: A synthesis of experimental psychology, artificial intelligence and cognitive neuroscience,* (pp. 567-588). Cambridge, MA: MIT Press.

Neafesy, E.J., Hull, C.D., & Buchwald, N.A. (1978). Preparation for movement in the cat. II. unit activty in the basal ganglia and thalamus. *Electroencephalography and Clinical Neurophysiology, 44,* 714-723.

Newell, K. M. (1991). Motor skill acquisition. *Annual Review of Psychology, 42,* 213-237.

Norman, D.A., & Shallice, T. (1986). Attention to action: Willed and automatic control of behaviour. In R.J. Davidson, G.E. Schwartz, & D. Shapiro (Eds.), *Consciousness and self regulation: Advances in research (Vol. 4).* New York: Plenum Press.

Parent, A. (1990). Extrinsic connections of the basal ganglia. *Trends in Neuroscience, 13,* 254-258.

Parent, A., & Hazrati, L. (1993). Anatomical aspects of information processing in primate basal ganglia. *Trends in Neuroscience, 16,* 111-116.

Pastor, M.A., Artieda, J., Jahanshahi, M., Obeso, J.A. (1992). Time estimation and reproduction is abnormal in Parkinson's disease. *Brain, 115,* 875-891.

Pastor, M.A., Jahanshahi, M., Artieda, J., Obeso, J.A. (1992). Performance of repetitive wrist movements in Parkinson's disease, *Brain, 115,* 875-891.

Phillips, J.G., Stelmach, G.E., & Teasdale, N. (1991). What can indices of handwriting quality tell us about parkinsonian handwriting? *Human Movment Science, 10,* 301-314.

Plaut, D.C. (1991). *Connectionist neuropsychology: The breakdown and recovery of behavior in lesioned attractor networks.* Unpublished doctoral thesis. Carnegie Mellon University, Pittsburgh, PA.

Requin, J. (1992). From action representation to movement control. In G. E. Stelmach, & J. Requin (Eds.), *Tutorials in Motor Behaviour (II),* Amsterdam: Elsevier Science.

Requin, J., Riehle, A., & Seal, J. (1992). Neuronal networks for movement preparation. In D.E. Meyer & S. Kornblum (Eds.), *Attention and Performance XIV: A synthesis of experimental psychology, artificial intelligence and cognitive neuroscience,* (pp.746-769). Cambridge, MA: MIT Press.

Robertson, C., & Flowers, K.A. (1990). Motor set in Parkinson's Disease. *Journal of Neurology, Neurosurgery and Psychiatry, 53,* 583-592.

Robertson, H.A. (1992). Dopamine receptor interactions: some implications for the treatment of Parkinson's disease. *Trends in Neuroscience, 15,* 201-206.

Roland, P.E., Larsen, B., Lassen, N.A., & Skinho, E. (1980). Supplementary motor area and other cortical areas in organisation of voluntary movements in man. *Journal of Neurophysiology, 43,* 118-136.

Romo, R., & Shultz, W. (1992). Role of primate basal ganglia and frontal cortex in the internal generation of movements. 3. Neuronal activity in the supplementary motor area. *Experimental Brain Research, 91,* 296-407.

Rosenbaum, D.A. (1991). *Human Motor Control.* San Diego: Academic Press.

Rumelhart, D.E., Hinton, G.E., & Williams, R.J. (1986a). Learning internal representations by error backpropagation. In J.L. McClelland, & D.E. Rumelhart (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition (Vol. 1).* Cambridge, MA: MIT Press.

Rumelhart, D.E., Hinton, G.E., & Williams, R.J. (1986b). Learning error representations by backpropagating errors. *Nature, 323,* 533-536.

Rumelhart, D.E., & McClelland, J.L. (1986). On learning the past tenses of English verbs. In J.L. McClelland, & D.E. Rumelhart (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition (Vol. 2).* Cambridge, MA: MIT Press.

Rumelhart, D.E., & Norman, D.A. (1982). Simulating a skilled typist: A study of skilled cognitive-motor performance. *Cognitive Science, 6,* 1-36.

Saint-Cyr, J.A., & Taylor, J.E. (1992). The mobilisation of procedural learning: The 'key signature' of the basal ganglia. In L. Squires & N. Butters (Eds.) *Neuropsycholoy of Memory, 2.* (pp 188-202), New York: Guilford Press.

Saint-Cyr, J.A., Taylor, J.E., & Lang, A.E. (1988). Procedural learning and neostriatal dysfunction in man. *Brain, 111,* 941-959.

Saltzman, E. (1995). Dynamics and coordinate systems in skilled sensorimotor activity. In R.F. Port, & T. van Gelder (Eds.), *Mind as motion: Explorations in the dynamics of cognition* (pp.149-174). Cambridge, MA: MIT Press.

Saltzman, E., & Kelso, J.A.S. (1987). Skilled actions: a task dynamic approach. *Psychological Review, 94,* 84-106.

Sandefur, J.T. (1990). *Discrete Dynamical Systems: Theory and Applications.* Oxford: Clarendon Press.

Schell, G.R. & Strick, P.L. (1984). The origin of thalamic inputs to the arcuate premotor and supplementary motor areas. *Journal of Neuroscience, 4,* 539-560.

Schmidt, R.A. (1975). A schema theory of discrete motor skill learning. *Psychological review, 82,* 225-260.

Schmidt, R.A., Hawkins, B., Franck, J.S., & Quinn, J.T. (1979). Motor-output variability: a theory for the accuracy of rapid motor acts. *Psychological Review, 86,* 415-451.

Schmidt, R.A. (1988). *Motor Control and Learning,* (2nd ed.), Champaign, IL: Human Kinetics Publishers.

Schmidt, R.C., Shaw, B.K., & Turvey, M.T. (1993). Coupling dynamics in interlimb coordination. *Journal of Experimental Psychology, 19*(2), 397-415.

Seidenberg, M., & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review, 96,* 523-568.

Seitz, R.J., Roland, P.E., Bohm, C., Greitz, T., & Stone-Elander, S. (1990). Motor learning in man: a positron emission tomographic study. *NeuroReport, 1*(1 (Sampler)), 17-20.

Selemon, L.D. & Goldman-Rakic, P.S. (1985). Longitudinal topography and interdigitation of corticostriatal projections in the rhesus monkey. *Journal of Neuroscience, 5,* 776-794.

Semjen, A., & Gottsdanker, R. (1992). Plans and programs for short movement sequences. In G. E. Stelmach, & J. Requin (Eds.), *Tutorials in Motor Behaviour (II),* Amsterdam: Elsevier Science.

Servan-Schreiber, D., Printz, H., & Cohen, J. D. (1990). A Network Model of Catecholamine Effects: Gain, Signal-to-Noise Ratio, and Behaviour. *Science, 249,* 892-895.

Shaffer, L.H. (1975). Control processes in typing. *Quarterly Journal of Experimental Psychology, 27,* 419-432.

Shaffer, L.H. (1992). Motor programming and control. In G. E. Stelmach, & J. Requin (Eds.), *Tutorials in Motor Behaviour (II),* Amsterdam: Elsevier Science.

Shallice, T. (1988). *From Neuropsychology to Mental Sructure.* Cambridge, England: Cambridge University Press.

Shastri, L., & Ajjangadde, V. (1993). From simple associations to systematic reasoning: A connectionist representation of rules, variables, and dynamic bindings using temporal synchrony. *Behavioral and Brain Sciences, 16,* 427-451.

Shepherd, G. M. (1988). *Neurobiology.* New York: Oxford University Press.

Sheridan, M.R., & Flowers, K.A. (1990). Movement variability and bradykinesia in Parkinson's disease. *Brain, 113,* 1149-1161.

Sheridan, M.R., Flowers, K.A., & Hurrell, J. (1987). Programming and execution of movement in Parkinson's disease. *Brain, 110,* 1247-1271.

Schöner, G., & Kelso, J.A.S. (1988). Dynamic pattern generation in behavioural and neural systems. *Science, 239,* 1513-1520.

Schultz, W. (1992). Activity of dopamine neurons in the behaving primate. *Seminars in the Neurosciences, 4,* 129-138.

Shultz, W., Apicella, P., & Ljungberg, T. (1993). Response of monkey dopamine neurons to reward and conditioned stimuli during successive steps of learning a delayed response task. *Journal of Neuroscience, 13,* 900-913.

Shultz, W., Apicella, P., Scarnati, E., & Ljungberg, T. (1992). Neuronal activity in monkey ventral striatum related to expectation of reward. *Journal of Neuroscience, 12,* 4595-4610.

Shultz, W., Apicella, P., Romo, R., & Scarnati, E. (1995). Context-dependent activity in primate striatum reflecting past and future behavioural events. In In J.C. Houk, J.L. Davis, & D.G. Beiser (Eds.), *Models of Information Processing in the Basal Ganglia,* Cambridge, MA: MIT Press.

Shultz, W., & Romo, R. (1992). Role of primate basal ganglia and frontal cortex in the internal generation of movements. 1. Preparatory activity in the anterior striatum. *Experimental Brain Research, 91,* 296-407.

Smith, A. D., & Bolam, J. P. (1990). The neural network of the basal ganglia as revealed by the study of synaptic connections of identified neurones. *Trends in Neuroscience, 13,* 259-265.

Stelmach(Ed). (1976). *Motor Control: Issues and Trends.* London: Academic Press.

Stelmach G.E., & Castiello, U. (1992). Functional force control in Parkinson's disease patients. In G. E. Stelmach & J. Requin (Eds.), *Tutorials in Motor Behaviour (II),* Amsterdam: Elsevier Science.

Stelmach, G.E., Worringham, C.J., & Strand, E.A. (1986). Movement preparation in Parkinson's disease. *Brain, 109,* 1179-1194.

Sternberg, S. (1969). The discovery of processing stages: Extensions of Donder's method. *Acta Psychologica, 30,* 276-315.

Sternberg, S., Knoll, R.L., & Turock, D.L. (1990). Hierarchical control in the execution of action sequences: Tests of two invariance properties. In M. Jeannerod (Ed.), *Attention and Performance XIII: Motor Representation and Control,* (pp. 1-56). Hillsdale, NJ: Erlbaum.

Swerdlow, N.R., & Koob, G.F. (1987). Dopamine, schizophrenia, mania, and depression: Toward a unified hypothesis of cortico-striato-pallido-thalamic function. *Behavioral and Brain Sciences, 10,* 197-245.

Tanji, J., & Kurata, K. (1985). Contrasting neuronal activity in supplementary and precentral motor cortex of monkeys. I. Responses to instructions determining motor responses to forthcoming modalities. *Journal of Neurophysiology, 53,* 129-141.

Teasdale, N., Phillips, J., & Stelmach, G.E. (1990). Temporal movement control in patients with Parkinson's disease. *Journal of Neurology, Neurosurgery and Psychiatry, 53,* 862-868.

Teuber, H. L. (1976). Complex functions of the basal ganglia. In M.D. Yahr (Ed.), *The Basal Ganglia,* (pp. 151-168). New York: Raven.

Teulings, H-L., & Stelmach, G.E. (1992), Simulation of impairment of force amplitude and force timing in Parkinsonian handwriting. In G. E. Stelmach & J. Requin (eds.), *Tutorials in Motor Behaviour (II),* Amsterdam: Elsevier Science.

Thompson, J.M.T., & Stewart, H.B. (1986). *Nonlinear dynamics and chaos.* New York: Wiley.

Tipper, S.P., Lortie, C., & Baylis, G.C. (in press). Selective reaching: Evidence for action centred attention. *Journal of Experimental Psychology: Human Perception and Performance.*

Treisman, M., Cooke, N., Naish, P.L.N., & McCrone, J.K. (1994). The internal clock: Electroencephalographic evidence for oscillatory processes underlying time perception. *Quarterly Journal of Experimental Psycholgy, 47A,* 241-189.

Tuller, B & Kelso, J.A.S. (1990). Phase transitions in speech and their perceptual consequences. In M. Jeannerod (Ed.), *Attention and Performance XIII: Motor Representation and Control,* (pp. 429-452). Hillsdale, NJ: Erlbaum.

Turvey, M.T. (1977). Preliminaries to a theory of action with reference to vision. In R. Shaw & J. Bransford (Eds.), *Perceiving, acting and knowing: Toward an ecological psychology.* Hillsdale, NJ: Erlbaum.

Turvey, M.T., & Carello, C. (1995). Some dynamical themes in perception and action. In R.F. Port, & T. van Gelder (Eds.), *Mind as motion: Explorations in the dynamics of cognition* (pp.149-174). Cambridge, MA: MIT Press.

Viallet, F., Trouche, E., Beaubaton, D., & Legallet, E. (1987). The role of visual reafferents during a pointing movement: comparative study between open-loop and closed-loop performances in monkeys before and after unillateral electrolytic lesion of the substantia nigra. *Experimental Brain Research, 65,* 399-410.

Viallet, F., Trouche, E., Beaubaton, D., Legallet, E., & Khalil, R. (1986). Visual feedback and motor performance in human and animal basal ganglia dysfunction. In J. S. Schneider, & T. Lidsky (eds), *Basal Ganglia and Behaviour: Sensory Aspects of Motor Functioning,*Toronto: Hans Huber.

Von Hofsten, C. (1992). The gearing of early reaching to the environment. In G. E. Stelmach, & J. Requin (Eds.), *Tutorials in Motor Behaviour* (II), Amsterdam: Elsevier Science.

Von Holst, E. (1939). Die Relative Koordination als Phänomenonun als methode zentral-nervöser funktionsanalyse. *Engebrisse der Physiologie, Biologischen Chemie und Experimentellen Pharmakologie, 42,* 228-306. (Translation: Martin, 1973).

Von Holst, E. Coordination as a phenomenon and as a method of analysis of central nervous system functions. In R.Martin (Trans.). *The behavioural physiology of animal and man: The collected papers of Erich Von Holst (vol 1).* London: Methuen.

Vousden, J. (1996). *Serial control of phonology in speech production.* Unpublished doctoral dissertation, University of Warwick, Coventry, England.

Wahl, (1932). Neue Untersuchungen über das zeitgedachtnis der bienen. *Zeitschrift für vergelichende Physiologie, 18,* 709-717.

Weardon, J. (1994). Prescriptions for models of biophysical time. In M. Oaksford & G.D.A. Brown (Eds.) *Neurodynamics and Psychology.* London: Academic Press.

Wickelgren, W.A. (1965). Acoustic similarity and intrusion errors in short-term memory. *Journal of Experimental Psychology, 70,* 102-108.

Wickens, J.R. (1993a). *A Theory of the Striatum.* Oxford, England: Pergamon.

Wickens, J.R. (1993b). Corticostriatal interactions in neuromotor programming. *Human Movement Science, 12,* 17-35.

Wickens, J.R., Alexander, M.E. & Miller, R. (1991). Two dynamic modes of striatal function under dopaminergic-cholinergic control: a model and computer simulation. *Synapse, 7,* 181-192.

Wiles, J., & Elman, J. (1995). learning to count without a counter: A case study of dynamics and activation landscapes in recurrent networks. *Neural Computation, 7,* 482-487.

Wing, A.M. (1992). The uncertain motor system: Perspectives on the variability of movement. In D.E. Meyer & S. Kornblum (Eds.), *Attention and Performance XIV: A synthesis of experimental psychology, artificial intelligence and cognitive neuroscience,* (pp.710-744). Cambridge, MA: MIT Press.

Wing, A.M., & Kristofferson, A.B. (1973). Response delays and the timing of discree motor responses. *Perception and Psychophysics, 14,* 5-12.

Wright, et al., (1990). Generalised motor programs: Reexamining claims of effector independence in writing. In M. Jeannerod (Ed.), *Attention and Performance XIII: Motor Representation and Control.* Hillsdale, NJ: Erlbaum.