



This thesis entitled

Neurodynamical modeling of arbitrary visuomotor tasks

written by
Marco Loh

and directed by
Prof. Dr. Gustavo Deco

has been approved by the
Department of Information and Communication Technologies.

Barcelona, November 2007

A thesis submitted in partial fulfillment of the requirements
for the Degree of *Doctor per la Universitat Pompeu Fabra*

Dipòsit legal: B.23237-2008
ISBN: 978-84-691-4450-3

Acknowledgements

Thanks...

- to Gustavo Deco for being a very supportive supervisor and director of my thesis, even when it did not look that well at times,
- to Edmund T. Rolls for the excellent collaboration on the topic of schizophrenia,
- to Ralph G. Andrzejak and Anders Ledberg for the discussions and support on two manuscripts,
- to the rest of the computational neuroscience group,
- to Johan for suffering through most of my manuscripts and long discussions,
- to Mar and Andres for keeping me from eating all of my cookies,
- to the Boehringer Ingelheim Fonds not only for financial support, and
- to my parents for the invaluable backup during my time in Barcelona.

Abstract

Conditional visuomotor learning is a paradigm in which stimulus-response associations are learned upon reward feedback. A typical experiment is as follows: Upon a stimulus presentation, a subject has to decide which action to choose among a number of actions. After an action is selected, the subject receives reward if the action was correct. Several interacting brain regions work together to perform this task, most prominently the prefrontal cortex, the premotor cortex, the hippocampus, and the basal ganglia. Using computational modeling, we analyze and discuss the neurodynamics underlying this task. We propose the term *processing pathway* to describe how the brain performs this task and detail the roles and interactions of the brain regions. In addition, we address aberrant processing in a neurodynamical hypothesis of schizophrenia.

Keywords: Computational neuroscience, neurodynamics, modeling, cognition, neural networks, neuropsychology, visuomotor learning, association learning, reinforcement learning, schizophrenia.

Resumen

El aprendizaje visuomotor condicional es un paradigma en el que las asociaciones estímulo-respuesta se aprenden a través de una recompensa. Un experimento típico se desarrolla de la siguiente forma: cuando se presenta un estímulo a un sujeto, éste debe decidir qué acción realizar de entre un conjunto. Una vez seleccionada la acción, el sujeto recibirá una recompensa en el caso de que la acción escogida sea correcta. En este tipo de tareas interactúan distintas regiones cerebrales, entre las que destacan el córtex prefrontal, el córtex premotor, el hipocampo y los ganglios basales. El objetivo de este proyecto consiste en estudiar la dinámica neuronal subyacente a esta clase de tareas a través de modelos computacionales. Proponemos el término *processing pathway* para describir cómo ejecuta esta tarea el cerebro y explicar los roles e interacciones entre las distintas áreas cerebrales. Además, tratamos el procesamiento anómalo en una hipótesis neurodinámica de la esquizofrenia.

Palabras clave: Neurociencia computacional, neurodinámica, elaboración de modelos, cognición, redes neuronales, neuropsicología, aprendizaje visuomotor, aprendizaje asociativo, aprendizaje por refuerzo, esquizofrenia.

Contents

Abstract	v
Resumen	vii
1 Introduction and Overview	3
1.1 Conditional visuomotor learning	3
1.2 Structure of the Thesis and Publications	4
1.3 Review of visuomotor learning literature	5
1.3.1 Experimental literature	5
1.3.2 Modeling literature	13
1.4 Summary of PhD work	15
1.4.1 Neurodynamical Model	15
1.4.2 Decision making in visuomotor mappings	17
1.4.3 Prefrontal cortex activity during visuomotor learning	18
1.4.4 Dynamical systems hypothesis of schizophrenia	21
1.4.5 Neuropsychological model of novel visuomotor learning	26
1.5 Conclusion and future ideas	31
2 Cognitive flexibility and decision making	35
2.1 Introduction	36
2.2 Methods	36
2.2.1 Conditional Visuomotor Task: Experimental Paradigm	36
2.2.2 The Integrate-and-Fire Network	37
2.2.3 The Stimulus-Response Architecture	38
2.2.4 Analysis	40
2.3 Results	40
2.4 Discussion	44
3 Neurodynamics of the prefrontal cortex	51
3.1 Introduction	52
3.2 Methods	53
3.2.1 Experimental Background	53
3.2.2 Neurodynamical model	53
3.2.3 Analysis	55
3.3 Results	56
3.4 Discussion	59

4	A dynamical systems hypothesis of schizophrenia	67
4.1	Introduction	68
4.2	Results	69
4.2.1	Dynamical systems hypothesis	69
4.2.2	The depth of the basins of attraction: flow analysis	70
4.2.3	Concept of stability in network simulations	72
4.2.4	Stability	74
4.2.5	Distractibility	75
4.2.6	Signal-to-noise ratio	76
4.3	Discussion	77
4.4	Methods	82
4.4.1	Attractor Framework	82
4.4.2	Analysis	83
5	Reward-learning, Hebbian-learning, and Strategy	87
5.1	Introduction	88
5.2	Results	89
5.3	Discussion	95
5.4	Methods	100
A	Neural Network Model	103
A.1	Neural and synaptic dynamics	103
A.2	Mean-field Formulation	105
B	Connection Matrices for Chapter 2	107
C	Connection Matrices for Chapter 3	109
D	Connection Matrices for Chapter 4	111
	Bibliography	112

Chapter 1

Introduction and Overview

1.1 Conditional visuomotor learning

The central topic of this thesis is conditional visuomotor learning. It describes the ability of learning by trial and error to map a stimulus onto a particular response upon reinforcing rewards. This paradigm has also been termed arbitrary visuomotor learning, conditional visuomotor mapping, conditional discrimination, or association learning. A typical experiment is as follows (see Figure 1.1): A visual stimulus is presented. Upon the presentation of this stimulus, the subject has to decide which action to choose. After an action is executed, the subject receives a reward if the action was correct. For each stimulus only one of the possible responses is rewarded. Thus, the subject has to learn by trial-and-error which combinations of stimuli and actions yield reward.

Although this experimental setup seems to be very specific, it contains many functions which are of general importance for cognitive behavior: Identification of a visual stimulus, decision making, working memory, trial-and-error learning, reward processing, motor control. An everyday example could be the human behavior at traffic lights: Green indicates go and red stop. Also, the ability to read aloud might belong to that category, since words written on paper are mapped onto motor movements of the speech system. Thus, conditional visuomotor learning is a central building block of cognition and higher-level abilities of both humans and animals. It contains elements of both classical conditioning (since a stimulus is associated with a reward) and instrumental learning (since the action causes, in certain cases, the emission of a reward). Whereas these two basics paradigms have been addressed intensively both in experiment and theory, the more complex visuomotor learning paradigm has been addressed to a lesser extent. Nevertheless, in the last years, also during the course of the PhD, several




Figure 1.1: Exemplary time course of a visuomotor learning trial which includes a delay period. After an initial fixation period, the stimulus appears. The following delay period forces the use of working memory. If the correct response is issued, e.g., by a saccadic eye movement, a reward is received, e.g., in form of a liquid.

experimental and theoretical papers have been published on this paradigm (e.g. Pasupathy & Miller (2005), Brasted et al. (2005), Buch et al. (2006), Hadj-Bouziane et al. (2006), Fusi et al. (2007)), which highlights the current interest. The main interest of the field regards the contribution of various brain areas to learning, although there have also been recent contributions regarding the behavior. Using computational modeling we contributed both to the understanding of the neurodynamics underlying visuomotor associations and to a more critical assessment of the behavioral data by proposing a model for the neuropsychology.

Associations are also affected in mental illnesses. In schizophrenia, wrongly build up associations lead e.g. to delusions and persecution complex. Schizophrenia is a complex mental illness, which affects about 1% of the population, and only a part of the patients can be treated effectively with medications. A better understanding of the illness is crucial. While working on dopamine, a neuromodulator, which is a main target in the treatment of schizophrenia and is envisioned to be involved in learning, we set up a computational hypothesis, which might help to improve the understanding of that illness.

1.2 Structure of the Thesis and Publications

This thesis consists of an introductory chapter and publications by the doctoral student, which are printed in the chapters thereafter.

In the introductory section, we give first a review of the literature of conditional visuomotor learning. We focus on the specific experiments used to study this paradigm. This view has not been extensively considered in the literature so far, but it is central to the manuscript presented in Chapter 5: We demonstrate that the behavior of two experiments yields conflicting properties, which suggests that the specific experiments used to study visuomotor learning must be considered more carefully. After a short overview of the modeling literature, we present an extensive summary of the work done during the PhD. We conclude the introductory chapter with an outlook and future ideas for further modeling work.

After the introduction, we present in four chapters manuscripts by the doctoral student which are either accepted/published in international peer-reviewed journals (chapters 2,3,4) or in preparation (chapter 5). They provide detailed information including all results and in depth discussions of the related literature. They are ordered by their level of detail ranging from biophysical to behavioral models:

In Chapter 2, we present a neurophysiological model of visuomotor mappings which contains populations of neurons which are selective to the stimuli, the motor responses and the associations (published in the *European Journal of Neuroscience* (Loh & Deco 2005)). It is meant to cover the whole processing path from sensory to motor neurons. This model is based on earlier work by Deco & Rolls (2003), and we extended it by analyzing the influence of fluctuations and by applying a mean field technique to the model. The most important contribution of this work is to highlight the importance of error trials for modeling. In newer experimental studies such as Pasupathy & Miller (2005) or Buch et al. (2006), the activity of error trials is presented.

We further study the model of Loh & Deco (2005) by focusing just on the association selective neurons in Chapter 3. Here, we take a hypothesis driven approach in which we try to combine theories about the prefrontal cortex function in a computational model and compare it with the experimental data. We find a mismatch of the data in both the delay period and the error trials. We show that, at least in this visuomotor learning experiment,

the prefrontal cortex does not seem to have a strong working memory characteristic and that error trials happen due to processing errors, which is reflected by a separate attractor state. This work is accepted for publication in the *Journal of Cognitive Neuroscience* (Loh et al. 2008).

In Chapter 4, we present a neurodynamical hypothesis of schizophrenia. While working on dopamine, which is suggested to be involved in learning, we found that the effects of dopamine on working memory circuits are still a matter of debate. In addition, a direct connection of dopamine to the reinforcement learning model used in Chapter 5 has been proposed (Schultz 2006). Since dopamine is the major therapeutic target in schizophrenia, we investigated which alternations in a neurodynamical system could correspond to the symptoms of schizophrenia. We propose a hypothesis, which is consistent with the current knowledge of the effects of dopamine in cortical circuits, and thereby we unified the different types of symptoms of the illness. False or bizarre associations are one of the major symptoms of schizophrenia leading to delusions and mania. This work is published in *PLoS Computational Biology* (Loh et al. 2007).

Chapter 5 contains a neuropsychological modeling study in which we model and compare the behavioral data of visuomotor experiments. To our knowledge, this is the first cross-study model in visuomotor mappings, which tries to address two or more experiments at the same time. We compare evidence from experiments conducted by two groups and show that they are inconsistent. In addition, we propose that the hippocampus might contribute a Hebbian learning type component in visuomotor learning.

1.3 Review of visuomotor learning literature

Arbitrary visuomotor mappings are one of the most complex learning tasks studied with non-human primates. There are several other kinds of experimental paradigms, which are similar to arbitrary visuomotor mappings. For example, the stimulus could be the target of the action, such as in grasping a door handle (standard mapping), or the stimuli could provide target information e.g. via an arrow (transformational mapping). On a wider scope, one could not only map a visual stimulus onto a motor response, but also an auditory or olfactory stimulus onto other stimuli (stimulus-stimulus mapping, match-to-sample task) or rewards (stimulus-reward mapping). It is still a matter of discussion, if these different kinds of tasks are performed by similar brain networks and mechanisms. We emphasize that even the same brain region could play different roles in different tasks and thus drawing conclusions from the function of one brain region in one task to the function of that region in another task might not be valid. We take a conservative approach and confine our review strictly to arbitrary visuomotor mappings.

1.3.1 Experimental literature

The experimental studies on this paradigm have been reviewed several times, mainly addressing the contribution of different brain regions (Wise et al. 1996, Wise & Murray 1999, Murray et al. 2000, Wise & Murray 2000, Passingham et al. 2000, Hadj-Bouziane et al. 2003). Variations within the experimental paradigm have been largely neglected in the literature. The experiments vary in several respects such as

- the number of mappings during a session,

- the duration of presentation and delay periods,
- the use of familiar mappings (mappings which are fixed during the whole experiment),
- the use of correction trials after a wrong answer.

We suggest that the experimental paradigms used in the experiments have a substantial influence on the behavioral properties and therefore the underlying neural dynamics is most likely also different. We show which different paradigms have been used to study visuomotor learning and discuss possible effects on the behavioral and neural results. As the most notable difference, we will use the number of response options to organize the mini-review. We focus on primate studies.

Two response options

The mapping of two stimuli to two responses is the most basic and easiest form of visuomotor learning and has been applied in several variations (see Figure 1.2). We show which variations have been used and which results were gained. Thereby we discuss critically the relation to the experimental variations. There are two main issues in the literature: the effects of hippocampal lesions related to spatial vs. non-spatial responses and the use of strategies.

The effects of spatial/non-spatial learning and the effects of a fornix transection are studied by Rupniak & Gaffan (1987), Gaffan & Harrison (1989)/Gaffan & Harrison (1988), respectively. A fornix transection disconnects the output structure of the hippocampus. Since the spatial version of the task was affected by the fornix transection, but the non-spatial version was not, it was suggested that the hippocampus contributed to spatial learning. This contrasts the finding in non-spatial mapping tasks with three response options (Brasted et al. 2003), in which an effect of fornix transection was identified. Ridley & Baker (1997) also investigated lesions of the hippocampal system with two protocols of a 2:2 mapping tasks. They either presented the stimuli in a pseudo-random order or in batches, i.e. 10 times stimulus A followed by 10 times stimulus B. The hippocampal lesion affected the first version of the task, but not the second. Especially the last experimental finding suggests, that the experimental paradigm has a substantial influence on the how the brain processes that task. It might depend on the task, if the abilities of the hippocampus are used and thus an effect is seen in lesion studies. The factor of the spatial/non-spatial contribution might not describe the hippocampus function completely, as several other factors have also been suggested in the literature (see e.g. Brasted et al. (2003) for a detailed discussion). In Chapter 5, we suggest that the hippocampus might contribute to Hebbian-learning type mechanisms and show how this can be assessed in experimental data.

The use of strategies has first been pointed out by Wise & Murray (1999). There are two possible strategies, which can be applied: win-stay/lose-shift and repeat-stay/change-shift. The repeat-stay/change-shift strategy is able to solve a simple 2:2 mapping task perfectly: If the stimulus from one trial to the next stays the same, then the same response yields a reward. If the stimulus changes, then there is just one alternative response option, which is then correct (provided that there are just two stimuli). Applying this strategy after one correct trial would itself yield a performance of 100% in a 2:2 mapping.

Wang et al. (2000) studied a two-response mapping task and addressed in their analysis the use of strategies. They studied with reversible lesions, using infusion of bicuculline, the

Article	Task Characteristics	Experimental notes
Petrides (1982)	WGTA, correction trials	Periarculate area (premotor) lesion - impairment, principalis region (PFCdl) - mild impairment
Petrides (1985)	WGTA, non-spatial response	Periarculate area (premotor) - impairment
Rupniak & Gaffan (1987)	Spatial response, go/no-go response	Fornix transection - impaired, sulcus principalis - not impaired
Gaffan & Harrison (1988)	Non-spatial response, tap/hold response, correction trials	Fornix transection - not impaired
Gaffan & Harrison (1989)	WGTA, spatial response, 4:2 mapping	Fornix transection - impaired, sulcus principalis - not impaired
Canavan et al. (1989)	Pull-turn responses, correction trials	Motor thalamus - impaired
Eacott & Gaffan (1992)	Non-spatial, tapping, correction trials	Inferotemporal-frontal disconnection affects the learning of visuo-motor mappings
Ridley & Baker (1997)	WGTA, either random order or batches (10 times stimulus A followed by 10 times stimulus B).	Random order task was affected by the hippocampal lesion, batch task was not.
Asaad et al. (1998)	Mix of familiar and novel associations, delay period, no correction trials, saccadic eye movements	Recordings in PFCdl
Wang et al. (2000)	Correction trials, two protocols: mix of familiar and novel associations, two sets of novel associations.	Analysis of strategies, two sets of novel associations are learned slower than mix of one novel and one familiar, PFCv - impaired novel, but not familiar, mainly due to strategy.
Nixon et al. (2004)	Corrections trials, two sets of novel associations (to avoid strategies)	Cross-lesion of the basal-ganglia and the premotor cortex - impaired familiar, but not novel
Pasupathy & Miller (2005)	see Asaad et al. (1998)	Recordings in PFCdl and basal ganglia
Fusi et al. (2007)	see Asaad et al. (1998)	Further analysis and modeling, error reset

Figure 1.2: Overview of visuomotor experiments using two response options. If not otherwise specified, the brain regions in the experimental notes column represent the target of the lesion in the experiment. WGTA: modified Wisconsin General Test Apparatus, PFCdl: dorsolateral prefrontal cortex, PFCv: ventral prefrontal cortex.

contribution of the prefrontal cortex. They found that infusions in the ventral prefrontal cortex affected the learning of novel associations, but it did not alter the performance of familiar mappings. They analyzed the use of the repeat-stay/change-shift and lose-shift strategies and found that the effects of the ventral prefrontal cortex lesion can be explained by the inability to apply strategies. This finding might shed light on older experiments such as by Petrides (1982), in which a lesion of the sulcus principalis regions yielded a mild impairment of visuomotor learning. It could be that this impairment was also due to the elimination of strategy abilities. In addition, Wang et al. (2000) used two different protocols: A mix of familiar and novel association and the mix of two novel associations. They found that associations are learned slower when two novel associations are learned concurrently. This contrasts the finding by Fusi et al. (2007) in which it was reasoned that mappings are learned independently.

Fusi et al. (2007) analyzed experimental data by Asaad et al. (1998) and Pasupathy & Miller (2005), who, like Wang et al. (2000), used a task in which two novel stimuli were presented in parallel with two familiar stimuli. In addition, Asaad et al. (1998) used reversals: In each session, the novel stimuli were first learned until a criterion of 90% correct in 10 trials was reached. After exposing the monkey to another 20-30 trials, the mapping between the stimuli was reversed: If stimulus A mapped to action right (R) and stimulus B to action left (L), then the mappings was A to L and B to R after the reversal. This allowed the authors to study in detail the re-learning of mappings. Two familiar stimuli, whose associations never changed, were randomly intermingled. These were presented half as often as the novel stimuli.

Although the monkeys could have efficiently used strategies to solve this task, they did not. Neither did the behavioral data show a repeat-stay/change-shift pattern nor did the behavior after the reversal yield an immediate switch to the novel mapping. The task mainly differed from Wang et al. (2000) in that it had a delay period in the experiments (working memory demands) and in that no correction trials were given, i.e., no stimulus was repeated until the correct answer was given. It is an open question, which of these factors are involved in the elimination of the use of strategy. Another possible source is the reversal: Fusi et al. (2007) investigated their data further and identified a reset after each error. If the monkeys made an error, they forgot all learned associations and started with trial and learning again. The frequent reversals might have sensitized the monkeys to errors and they developed a reset strategy. These considerations highlight the importance of the exact experimental task, since relatively small differences seem to determine if strategies are employed or not.

Nixon et al. (2004), aware of the possibility of strategy use, mixed two cue-response pairs in one session, explicitly to avoid strategies. They tested the effect of a cross-lesion procedure disconnecting the basal ganglia and the premotor cortex. In a cross lesion procedure, one brain region is lesioned in one hemisphere and the other in the contralateral one to ensure that in both hemispheres there is no communication between the affected brain parts. However, all other connections remain intact in one of the two hemispheres. They found that the retention of well-established visuomotor associations seems to depend on the connection between the basal ganglia and the premotor cortex, while the ability to learn novel ones was unaffected by the lesions.

Other experimental evidence in the literature shows less controversy. Petrides (1982, 1985) studied lesions of the periarculate area and found that visuomotor learning is impaired. Lesions of the motor thalamus and Inferotemporal-frontal disconnection also affect the learning of visuomotor mappings (Canavan et al. 1989, Eacott & Gaffan 1992). In these experiments, only novel associations were learned.

Article	Task Characteristics	Experimental notes
Murray & Wise (1996)	Spatial response, correction trials, novel and familiar	Hippocampal system - novel impairment, but not familiar
Wise & Murray (1999)	see Murray & Wise (1996)	Analysis of strategy effects, Hippocampal system - strategy not impaired
Murray et al. (2000)	Review and additional data of a PFC lesion	PFC - strategy and novel impaired
Bussey et al. (2001)	Three and four response options task, spatial response, correction trials	PFC - strategy and novel impaired
Inase et al. (2001)	Spatial response, delay period, substitution of familiar stimuli with novel ones	Recordings in globus pallidus
Brasted et al. (2002)	Non-spatial response, short-hold/long-hold/tap, correction trials	Fornix transection - novel impaired
Brasted et al. (2003)	see Brasted et al. (2002)	Fornix transection - novel impaired
Brasted et al. (2005)	see Brasted et al. (2002)	Detailed analysis, Fornix transection - one-trial learning impaired
Genovesio et al. (2005)	Strategy task, repeat-stay/change-shift independent of associations	Recordings in PFC, strategy-selective neurons
Genovesio et al. (2006)	see Genovesio et al. (2005)	Detailed analysis, neural representations of previous and future goals
Inase et al. (2006)	see Inase et al. (2001)	Recordings in medial PFC

Figure 1.3: Overview of visuomotor experiments using three response options. If not otherwise specified, the brain regions in the experimental notes column represent the target of the lesion in the experiment. PFC: prefrontal cortex.

Overall, the task using just two-response options has its difficulties and which factors lead to which behavior is not clear. Nevertheless, recently the problem of strategy use in this task has been recognized and explicitly avoided by task constraints (Nixon et al. 2004). We did not address the differences within the control conditions, as a comparison of the different measures would require explicit computational modeling. This is subject to future work and might contribute to a more detailed understanding of the neuropsychological studies.

Three response options

Mappings of three stimuli onto three responses have been investigated in a series of experiments by Steve Wise and collaborators (see Figure 1.3). They target the same issues as the experiments with two response options: Lesions of the hippocampus and use of strategy. However, the three mapping task allows to distinguish between strategy and learning which is not possible with the two mapping task. When the stimulus changes from one trial to the next, the probability of a correct answer according to the strategy is 50% (and not 100% like in the two mappings task). This fact can be used to distinguish learning and strategy.

Murray & Wise (1996), Wise & Murray (1999), Murray et al. (2000) present data from monkeys with either hippocampus, amygdala or prefrontal cortex lesion. They use a 3:3

mapping with no delay periods and up to two correction trials. Correction trials are trials, which are given after an incorrect response. The monkeys had up to two more trials with the same stimulus to make a correct response. They found that the monkeys were using a strategy to improve performance: When the stimulus stayed the same in a subsequent trial, they could stick with the correct answer. When the stimulus changed, the monkeys chose one of the two responses that they had not chosen earlier. This repeat-stay/change-shift strategy could improve the performance in repeat-trials to 100% and in change-trials to 50% without any learning. Note that for this strategy to apply, the monkeys needed a correct answer. The correction trials might facilitate the build up of these strategies.

Lesion of the amygdala did not change the performance of the monkeys (Murray & Wise 1996). However, lesions of the hippocampus greatly reduced the ability to learn novel stimuli whereas the ability to perform well-established familiar mappings was retained. In addition, the monkeys were still able to apply the repeat-stay/change-shift strategy. Lesions of the ventral and orbital prefrontal cortex impaired both the ability to learn novel mappings and to apply strategies within the task (Murray et al. 2000, Bussey et al. 2001). In this task, the monkeys consistently used a strategy.

To investigate the strategy use further, Genovesio et al. (2005) designed an experiment to specifically address the repeat-stay/change-shift strategy. The monkeys were just rewarded when applying the strategy: When the stimulus was repeated, the monkeys had to stick to the previous answer. When the stimulus was changed, one of the other responses was rewarded. In this experiment, the associations changed constantly and just the rule stayed the same. Neural recordings in the prefrontal cortex revealed activity selective to the repeat-stay and change-shift strategy indicating the importance of the prefrontal cortex for strategy use.

The hippocampus is believed to play a key role in spatial learning and orientation. Results of two response mapping tasks showed that spatial versions of the task are affected by lesions of the hippocampal system (Rupniak & Gaffan 1987, Gaffan & Harrison 1989) whereas non-spatial versions are not (Gaffan & Harrison 1988). In a set of experiments, Wise and collaborators established the importance of the hippocampus in visuomotor learning beyond the spatial domain (Brasted et al. 2002, 2003, 2005). The monkeys were shown visual stimuli, but instead of giving a spatial response with a joystick, they had to respond with certain touch sequences on a touch-screen. These responses were 8 taps, a short touch (2-4 sec long), and a long touch (4-8s long). This is analogous to the 2:2 mapping task used by Gaffan & Harrison (1988). The stimulus stayed on the screen until the response was identified. A fornix transection (disruption of the output structure of the hippocampus) showed impairment in the learning of new associations suggesting a role of the hippocampus beyond the spatial domain. This is in contrast with the results of the two response mapping task. It is not clear, why one non-spatial task was affected and the other one was not. A possible answer could be the use of strategies. If the monkey of Gaffan & Harrison (1988) used strategies to solve this task, then a hippocampal lesion could show no effect. In this non-spatial version, the monkeys did not adopt a response strategy as they had done in the spatial version.

A detailed analysis of the data revealed, that the hippocampus had an especially high impact on the very fast one-trial learning. In addition, errors made prior to the correct response slowed down the learning of control monkeys. This suggests that monkeys do not learn by errors in this task, but that they are rather impeding learning.

Inase et al. (2001) used a task mixing familiar and novel mappings in one set. The monkeys first learned a 3:3 mapping to a high level. Then, one of the three problems was substituted by a novel one with which they studied the novel learning and contrasted it with the familiar

trials. In neural recordings, they identified learning related activity in the globus pallidus interna, an output nucleus of the basal ganglia. The same paradigm was used in another study (Inase et al. 2006), in which they recorded in the medial prefrontal cortex. There, they identified neurons selective to either novel or familiar cues. No use of strategy was reported in these studies, although a detailed analysis of the data in this respect would be interesting.

The results of the three response mapping tasks add to the picture gained in the experiments of the two mappings task: The task and probably also the training has a substantial influence on the behavior and the effects of lesions. We characterize the mapping task by Brasted et al. (2003, 2005) in detail in a computational model in chapter 5, which yields more insights in this version of the task.

Four response options

Visuomotor experiments with four response options were almost exclusively used together with neurophysiological recordings (see Figure 1.4). In the corresponding articles, the focus is usually on the neural data with just a brief discussion of the behavioral data.

The few neuropsychological studies indicate that strategies are also used in this type of experiments. Bussey et al. (2001) mixed three and four response option experiments in their analysis and has been mentioned above. They identified both repeat-stay/change-shift and win-stay/lose-shift patterns in their data. This is the only mentioning of strategy usage in visuomotor mappings. Nevertheless, there are some other studies with similar paradigms:

Mitz et al. (1991) recorded neurons in the premotor cortex while primates performed a 4:4 mappings task with 4 joystick movements: left, right, down, no-go. The tasks featured correction trials, a short delay period (100ms), and both novel and familiar associations. A session would either contain only familiar or only novel stimuli. Due to the great similarity of this paradigm to the one studied by Bussey et al. (2001), a use of strategy is likely also in this task, although this possibility was only recognized several years later.

Hadj-Bouziane & Boussaoud (2003) adapted their task design from Mitz et al. (1991). They use joystick movements as responses, both familiar and novel associations, and no delay period. The sessions also contained either novel or familiar stimuli. The stimuli were presented in blocks of 4 in which all stimuli were presented in a pseudorandom order. They found that the associations were learned sequentially. Learning the first associations was also quicker than learning later ones. This finding was detailed in a modeling study (Hadj-Bouziane et al. 2006). They reported neural activity in the striatum.

Williams & Eskandar (2006) also recorded in the striatum, but in addition stimulated neurons with electrical impulses. This increased the behavioral learning rate. They mixed 2 familiar and 2 novel associations, which mapped onto the 4 targets. As in Hadj-Bouziane & Boussaoud (2003), they presented the stimuli in blocks of 4 in which they were pseudo-randomly ordered and the response was given by joystick movements. Here, correction trials allowed several stimulus presentations until a correct response was given. The setup of both experiments would allow the build up of strategies, although there appear to be no reports in the literature.

The other experimental paradigm with four response options map more than four stimuli onto four responses which further complicates the learning task. However, this makes it more likely that the learning of visuomotor associations itself is studied without confounds of strategies: Chen & Wise (1995*a,b*, 1996) mixed novel and familiar stimuli and a session would contain 3-4 familiar stimuli and 1-3 novel stimuli, adapted to the monkeys' learning

Article	Task Characteristics	Experimental notes
Mitz et al. (1991)	Correction trials, novel and familiar, joystick response	Recordings in premotor cortex
Chen & Wise (1995a)	Delay period, correction trials, saccadic eye movements, mix of novel and familiar associations (3-4 familiar and 1-3 novel stimuli)	Recordings in supplementary eye field (SEF)
Chen & Wise (1995b)	see Chen & Wise (1995a)	Recordings in frontal eye field (FEF)
Chen & Wise (1996)	see Chen & Wise (1995a)	Detailed analysis, direction selectivity of neurons changes during learning
Bussey et al. (2001)	Three and four response option task, spatial response, correction trials	PFC - strategy and novel impaired
Bussey et al. (2002)	8 stimuli mapped onto 4 responses, familiar and novel, no delay period, touch-screen, correction trials	Crosslesion of IT and PFCo+v - impaired both novel and retention.
Hadj-Bouziane & Bousaoud (2003)	Task adapted from Mitz et al. (1991), stimuli were presented in blocks of 4	Recordings in striatum, associations were learned sequentially
Wirth et al. (2003)	Delay period, mix of novel and familiar stimuli, saccadic eye movements, natural scene stimuli	Recordings in hippocampus
Brasted & Wise (2004)	4 familiar and 4 novel stimuli, delay period, touch-screen, correction trials	Recordings in PMd and putamen
Buch et al. (2006)	see Brasted & Wise (2004)	see Brasted & Wise (2004)
Williams & Eskandar (2006)	2 familiar and 2 novel stimuli mapped onto 4 targets, joystick response, delay period, correction trials	Recordings and stimulation in the caudate nucleus

Figure 1.4: Overview of visuomotor experiments using four response options. If not otherwise specified, the brain regions in the experimental notes column represent the target of the lesion in the experiment. PFCo+v: orbito and ventral prefrontal cortex, IT: inferotemporal cortex, PMd: dorsal premotor cortex.

rate. The recordings were carried out in the supplementary eye field (SEF, considered to be the premotor region for eye movements) and the frontal eye field (FEF, which might have the role of the primary motor cortex). They identified less learning related changes in the FEF. Chen & Wise (1996) analyzed the same data and found that the direction selectivity of the learning-dependent neurons changes during learning.

Two recent papers by Brasted & Wise (2004) and Buch et al. (2006) compared the activity of the premotor cortex and the striatum. The experiment used an 8:4 design in which 4 familiar and 4 novel stimuli were mapped onto 4 responses. The novel stimuli appeared twice as often. The recordings in both the premotor cortex and associated striatal areas tested the hypothesis that the two areas are interconnected in cortical-basal ganglionic loops. It was confirmed that both areas changed their activity at approximately the same rate closely related to the monkey's learning curve. The activity for novel mappings reached the one of the familiar ones as learning progressed. Moreover, the activity in the putamen (part of the striatum), showed highly selective activity for the association during the pre-reward period, which suggests that it might be important during learning.

Bussey et al. (2002) also used a 8:4 task design with correction trials and, typical for lesions studies, no delay period. They studied the disconnection of the orbito-frontal and ventral PFC with IT (crosslesions procedure), and found that it is important for both learning and retention of visuomotor associations. Since also a match-to-sample task was affected, it might also contribute to higher order rules. Wirth et al. (2003) reported learning related activity in the hippocampus, which is in line with the lesion studies carried out by Steve Wise and collaborators. As stimuli, they used natural scenes instead of letter combinations.

Overall, in some experiments, which used four response options, it seems possible that strategies were involved in learning. However, many used more than four stimuli, which makes it unlikely that response strategies were developed and also the gain of using strategies would be lower. Nevertheless, an explicit analysis would be interesting. In addition, a more detailed analysis of the behavior after an error such as the one conducted by Fusi et al. (2007) could yield insights: If after an error all learned associations would reset and the learning start anew, then the performance should be quite low. This could also be addressed using computational modeling.

1.3.2 Modeling literature

The level of detail in modeling determines which questions are addressed. Here, we distinguish three types of modeling studies: Connectionist, neural, and behavioral modeling.

Connectionist models feature abstract representations of brain activity and can be applied to study the relations of different types of activity and the overall dynamics of large systems. In addition, they are easy to simulate which makes them a good starting point before going into more detailed models.

A connectionist model has been applied to the experiment by Mitz et al. (1991). Fagg & Arbib (1992) proposed a model of the premotor cortex consisting of motor selection columns. A motor selection column consisted of a feature detector, which represented the sensory input, and a voting unit, which chose a specific motor program and represented the decision-making component. The latter was modulated by noise and a threshold modulator. The model also used a reinforcement learning based learning rule and could already reproduce behavioral features plus a number of single unit properties of the premotor cortex.

Dominey et al. (1995) and Dominey & Boussaoud (1997) built up multi-area models which

contained modules for different parts of the basal ganglia and the cortex. They addressed the functional role of cortico-striatal loops and showed on a global scope how not only oculomotor associations could be processed (Dominey et al. 1995), but also sequences and contexts. The review by Houk & Wise (1995) gives an overview of multi-area interactions.

Detailed biophysical neural models have the advantage of featuring exact time constants of the underlying neural activity. This allows discussing the specific time course of the experiments. In addition, a realistic description of noise enables to address the decision-making characteristics, which is envisioned to play a crucial role in the action selection of visuomotor learning.

Deco & Rolls (2003) presented a model based on a working memory network (Brunel & Wang 2001) featuring a detailed description of single neurons including synaptic channels. The model served as a first proof-of-principle for the application of detailed neural simulations for complex cognitive tasks. Deco & Rolls (2005*b*) added a neural mechanism for the context switch based on calcium-channels. Such a switching has been observed in learning stimulus-reward associations (Thorpe et al. 1983) but not stimulus-response associations (Asaad et al. 1998). Loh & Deco (2005) extended the conditional visuomotor mapping model by studying fluctuations and multistability. We introduced multiple spiking simulations with subsequent clustering of the trials, which allowed for an assessment of the influence of fluctuations and the identification of multiple attractor states. We complemented the spiking simulations by a mean-field analysis to study the influence of the parameters on the neurodynamical system thoroughly. The major contribution of this study was to highlight the importance of error trials and suggest them as a major constraint for modeling.

Loh et al. (2008), based on the same framework, proposed a model specifically for the dynamics of the prefrontal cortex. We implemented a hypothesis-driven approach in which several hypotheses of prefrontal dynamics were integrated in one model, thereafter comparing with the neurodynamical data by Pasupathy & Miller (2005). We show that the main characteristic of the prefrontal cortex dynamics is not working memory, and that the processing follows a path from sensory to motor areas. In this work, we directly compared our model data to neural data, which is just possible if a detailed neural model is chosen.

Behavioral models address in an abstract way the learning of visuomotor mappings, since they do not explicitly take the neural substrate into account. However, since the learning might be complicated and the underlying system might span several brain regions, it may be necessary to gain insights with such models.

Fusi et al. (2007) presented a model which combines both neural and behavioral modeling. As a decision making module, they used a neural model of decision making (Wang 2002) and complemented it with plastic synapses. The learning rule changed the conductances of these synapses upon the reinforcing rewards and featured two different time scales to account for both novel and familiar mappings. In addition, Fusi et al. (2007) reasoned that the slow learning component could have a second utility: the auto balancing of the network and thereby a solution for the fine-tuning problem. They also identified a reset in the behavioral data: After each error, the monkeys seem to reset their behavior and begin at chance performance.

In chapter 5, we build on the basis of Fusi et al. (2007) and discuss in more detail a learning model focused on novel, fast learning. However, we do not model explicitly the neural decision-making network, but just use the corresponding output function in terms of action-selection probabilities. Among other things, this allows us to study experiments, which feature more than two response options. We fit our model to the visuomotor learning experiment by Brasted et al. (2003, 2005) and compare it to the reset behavior identified by

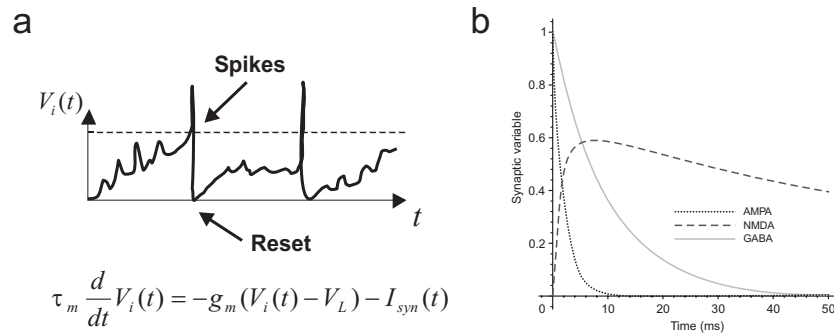


Figure 1.5: Overview of the neural level of the neurodynamical model. a. Basic behavior of the integrate-and-fire neurons. Once a threshold is reached, a spike is emitted and the membrane potential is set to a resting value. b. Synaptic variable as a function of time for the three synaptic channels AMPA, NMDA, and GABA. The detailed mathematical descriptions are found in the Appendix.

Fusi et al. (2007). We find that they are incompatible. In addition, we account for lesion data of the hippocampus and add a strategy module, which is envisioned to be implemented in the prefrontal cortex.

Another behavioral modeling study was recently presented by Hadj-Bouziane et al. (2006). They address the fact that the associations are learned sequentially in their visuomotor experiments and adjust a basic reinforcement model with ideas from viability theory to account for this data.

Recent experimental and modeling work shows that the field of visuomotor mappings is far from being understood and even on the behavioral side, there are many effects, which still need to be addressed. Behavioral modeling, based on classical reinforcement learning theory, is gaining importance again and will be complemented in the future with neural models to explain the interplay between several brain regions.

1.4 Summary of PhD work

The summary will resume the work of four manuscripts, which were written during the PhD. They are ordered by their level of detail ranging from biophysical to behavioral models, which is also the order in which they were developed. We will first present the basic neural network model, which has been used to investigate the cortical dynamics, before we go into the specific models for visuomotor associations.

1.4.1 Neurodynamical Model

As the basis for our neural network simulations, we use a recurrent network model which has been developed for the working memory properties of the prefrontal cortex (Brunel & Wang 2001). On the neural level (see Figure 1.5), the correlate is made up of integrate-and-fire neurons. These are described by

$$C_m \frac{dV(t)}{dt} = -g_m (V(t) - V_L) - I_{syn}(t), \quad (1.1)$$

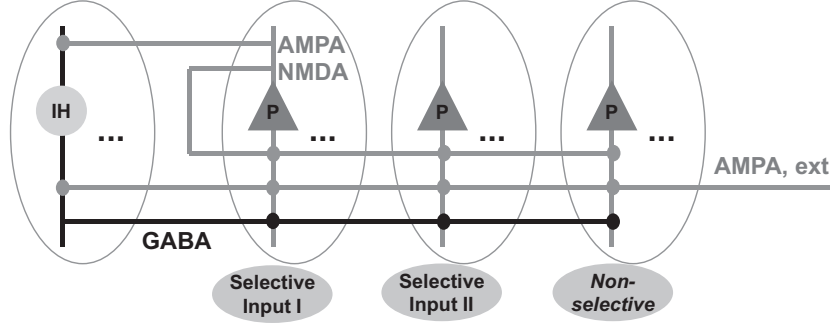


Figure 1.6: Connections within the network model. The network is all-to-all connected. Excitatory neurons emit their spikes via AMPA and NMDA channels, whereas inhibitory neurons use GABA channels. External input is given via AMPA channels. The excitatory neurons are grouped into several pools, which are either selective or non-selective. The selective pools are interconnected more strongly. The connection matrices of the studies are pictured in the Appendix.

where $V(t)$ is the membrane potential, C_m the membrane capacitance, g_m the leak conductance, and V_L the resting potential. The synaptic input of each neuron I_{syn} is made up of four parts. An external excitatory input via AMPA-type synapses to the network model and recurrent input from the other neurons of the network. The latter one consists of AMPA, NMDA and GABA currents. Thus, I_{syn} reads

$$I_{syn}(t) = I_{AMPA,ext}(t) + I_{AMPA,rec}(t) + I_{NMDA,rec}(t) + I_{GABA}(t), \quad (1.2)$$

The asynchronous external input $I_{AMPA,ext}(t)$ can be viewed as originating from 800 external neurons firing at an average rate of $\nu_{ext} = 3Hz$ per neuron, and thus has a rate of 2.4kHz in total. The synaptic currents are described by an exponential decay upon the arrival of a spike for AMPA and GABA synapses and a function including the rise time and an extra voltage dependence for the NMDA synapses. The parameters of the integrate-and-fire neurons and the synaptic channels for AMPA, NMDA and GABA are chosen in accordance with biological data.

On the network level (see Figure 1.6), the neurons are all-to-all connected and are made up of two parts: excitatory and inhibitory neurons. Whereas the excitatory neurons emit currents via AMPA and NMDA synapses, the inhibitory neurons connect via GABA synapses to the other neurons. The weights of the GABA synapses have the value one. The connections between the excitatory neurons are structured, so that pools of neurons form. Parts of the excitatory neurons group into selective pools, which respond to a certain stimulus and share stronger inner-pool connections. The parameter of this connection is typically called pool cohesion, designated w_+ . To compensate this stronger input from neurons of the same pool, the other connection strengths from the excitatory neurons are lowered to w_- . The connection strength w_- is calculated so that the average connection strength input to a neuron equals to one ($fw_+ + (1-f)w_- = 1$).

The system has stochastic properties due to both the random fluctuations of the external input and the finite size effects because of the limited number of neurons in the network. The biologically realistic time constants enable the study of the system according to experimental paradigms. To adjust the parameters and gain an overview of the dynamical states of the




Figure 1.7: Comparison of the network architectures, which are used in this thesis. a. Three layer network of stimulus-, association-, and response-selective neurons (Chapter 2). b. Network of association-selective neurons (Chapter 3). c. Two pools network to study neurodynamical properties in relation to schizophrenia (Chapter 4). The model schemes indicate the pools and important connections.

system, a mean-field formulation of the system has been developed (Brunel & Wang 2001). It is used as a first step to analyze the global dynamics of a system. The detailed mathematical descriptions and all parameter values for the studies are given in the Appendix.

In the course of the PhD, we have analyzed several architectures related to visuomotor mappings. The architecture in Figure 1.7a was very similar to the one presented by Deco & Rolls (2003) and represented the starting point of the thesis work. In the ensuing work, we simplified the model further to concentrate on specific aspects. Whereas the model depicted in Figure 1.7a contained stimulus-, response-, and association-selective neurons, we focused on the association-selective neurons in the model depicted in Figure 1.7b. For the study on schizophrenia, we focused on a system of two selective pools and analyzed detailed dynamical properties of the system (see Figure 1.7c).

In the following, we present the models ordered by level of abstraction, from detailed neurophysiological models, which address the neurodynamics of a large system, to a behavioral model based on reinforcement learning techniques. This is meant to be a short summary of the detailed studies presented in Chapters 2, 3, 4, and 5.

1.4.2 Decision making in visuomotor mappings

The starting point for the work on visuomotor learning was the article by Deco & Rolls (2003), in which they showed, using single trial simulations, how the selectivity of different neurons identified by Asaad et al. (1998) could be reproduced and related. They hypothesized the architecture depicted in Figure 1.7a, which integrates stimulus-, association-, and response-selective neurons in a coherent model. In a follow-up study, Deco & Rolls (2005b) presented a mechanism, with which the context input could be explained. The context input was hypothesized, because in the experiment the mappings switched: If stimulus A mapped onto a leftward response and stimulus B onto a rightward response, the mappings would be A to right and B to left after a reversal. Reversals were frequently applied during a session. There were two contexts: a direct context, in which A mapped to left and B to right, and a reverse context, in which A mapped to right and B to left. Deco & Rolls (2005b) hypothesized that a switch between the two contexts could be generated by calcium currents in neural assemblies.

In our work (see also Chapter 2), we use a mean-field implementation to study the dependencies of the connections between the three layers of the model: object-, association- and response-selective layer. The visuomotor association model assumes a processing pathway

from stimuli to motor circuits mediated by biasing signals which are thought to originate in the prefrontal cortex (Miller & Cohen 2001). The aim was to extract information about the connectivity, which could yield the types of neuronal behavior, which were identified in the experiment. In addition, we were interested in stability issues, which are related to decision-making. Multistability in stochastic neural systems is envisioned to be a key ingredient of decision making (Wang 2002). It was not clear if the neurodynamical model hypothesized by Deco & Rolls (2003) supports multistability. We discuss the connectivity between these abstract brain regions and thereby show the dependencies required for decision making characteristics.

Figure 1.8 shows the activity of all 8 selective pools of the system in a parameter configuration which shows multistability. Here, we categorized the simulated trials using a simple clustering approach. We identified three different types of trials: correct trials, error trials, and undecided trials. The undecided trials do not yield a clear computational answer to the given input. The distribution of the activity over the whole system in the different types of conditions represents a prediction for the experimental data. If multistable attractors support the exploration of the behavioral space, then our model predicts that the brain should respond stochastically with correct or incorrect activity to visuomotor associations until it has reached certainty. This should be visible not only in the premotor area, but also in association-selective or stimulus-selective activity in the prefrontal cortex. Most importantly, we emphasize the importance of error trials for the understanding of visuomotor associations, something that has been largely neglected in the literature so far. Usually only correctly executed trials were analyzed.

In addition, we propose an experiment in which monkeys are forced to make errors: For example, monkeys could be confronted with rapidly changing associations. They should constantly try to identify the associations and explore the behavioral options. The neural activity during the exploration could reveal the mechanisms, which underlie the response space exploration and guide the learning process.

1.4.3 Prefrontal cortex activity during visuomotor learning

During the work on the layer-based model, we saw that there were too many free parameters to make more concrete predictions about the neural activity, which go beyond general considerations. In addition, in the meantime Pasupathy & Miller (2005) published another article on the same paradigm with more detailed measurements in both the prefrontal cortex and the basal ganglia. We decided to develop a reduced model, which mainly targets the association-selective neurons, which are the main feature of the prefrontal cortex activity. We studied a system which contained four selective population for the combinations of the stimuli and responses in a 2:2 mapping (see Figure 1.7b).

As a result of our analysis, we presented a model which combined current hypotheses of the activity of the prefrontal cortex and compared it to the data by Pasupathy & Miller (2005) (see also Chapter 3). We used five hypotheses to construct the model:

- First, the model neurons should be selective to combinations of stimuli and responses, i.e. single associations in the experiment (Asaad et al. 1998). In the experimental paradigm, there are four associations between the two stimuli A,B and two responses L, R: AL, BR, AR, and BL. Thus, we included four populations which should be selective to these associations in our model (see Figures 1.7b, 2.1).

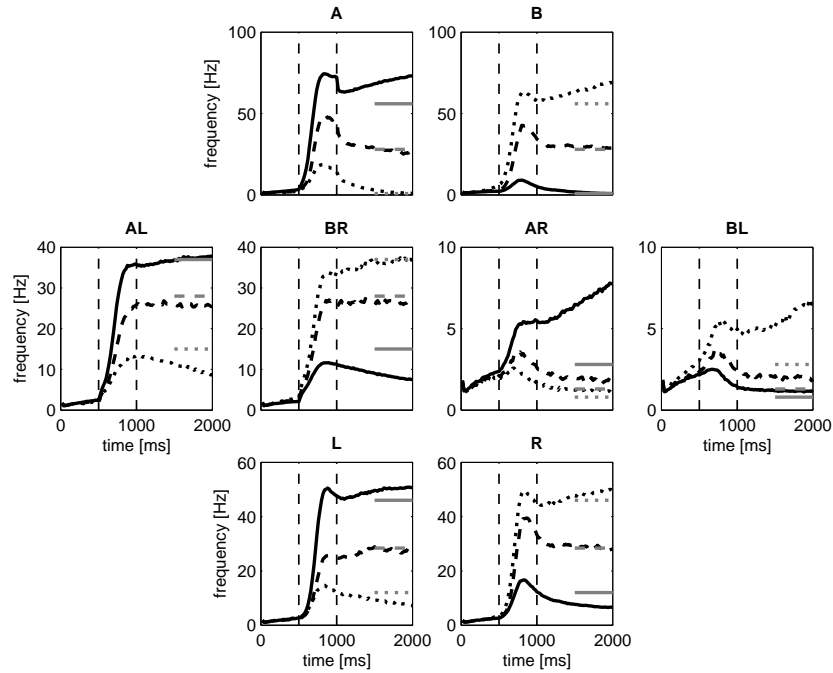


Figure 1.8: Mean spiking frequencies for each pool versus time for a parameter configuration, which yields multistability (see also Figure 2.3). The trials are clustered in three categories for correct trials (solid lines), error trials (dotted lines), and undecided trials (dashed lines). The horizontal lines between 1500 and 2000ms indicate the meanfield firing rates of the multistable states. The ordinates are not scaled uniformly.

- Second, stimulus information reaches the PFC from IT. As we are interested in the dynamics of the PFC, we model this input by an external influence on the model.
- Third, the basal ganglia map stimuli onto actions and their output influences the activity of the PFC. We model the direction input which codes for response direction left (L) or right (R) by an external input to (AL, BL) or (AR, BR) respectively. We study the effects of an increase of this external input to simulate the changes, which occur during learning. This increase is envisioned to be mediated by plasticity in the cortical-basal-ganglionic loop.
- Fourth, we assume that the response-direction input is influencing the prefrontal cortex throughout the delay period, as it has been hypothesized that the cortical-thalamic loop maintains this information (Houk & Wise 1995).
- Fifth, we hypothesize that the PFC module holds the extracted association in working memory. We use a neural network model of working memory (Brunel & Wang 2001) and focus our analysis on parameters which show delay period activity.

Figure 1.9a shows results of the analysis of the model. We show the course of the direction selectivity of the simulations. In the pre-stimulus period, the system is at rest and shows no direction selectivity. In the stimulus period (0-500ms), the direction selectivity rises almost identically for correct and error trials. It separates at the end of the delay period at which the

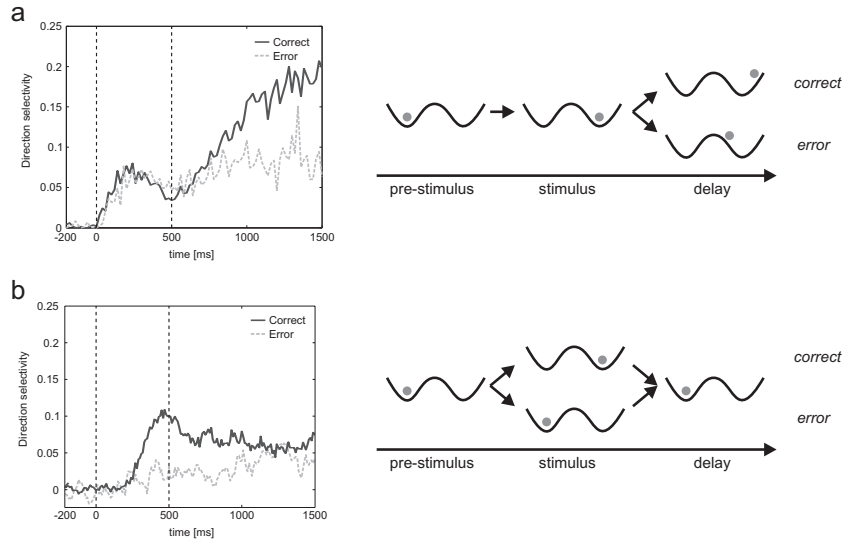


Figure 1.9: Overview of the dynamical properties of both model and experiment (see also Figures 3.2, 3.8, 3.9). a. On the left, modeling data showing direction selectivity of correct and error trials as a function of time. On the right, a sketch of the dynamical behavior of the model over time. We show in a hypothetical two attractor energy landscape for low and high activation the behavior of the model. In the prestimulus period, the system is at rest. In the stimulus period (0-500ms), the stimulus and response-direction input excite the system, so that the firings rates increase and the system jumps to a high attractor state. The distinction between correct and error trials in the delay period (500-1500ms) is caused by fluctuations in the high attractor state. b. Experimental data of PFC direction selectivity (adapted from Pasupathy & Miller (2005)) and a hypothesized sketch of the dynamical behavior.

decision must be made. In addition, we picture the dynamical mechanism in a hypothetical energy landscape. The analysis of the attractor landscape reveals that the system jumps up to a high frequency attractor in the stimulus period and then stays there in the delay period. This is not surprising, as part of our model assumptions were that the systems shows persistent activity in the delay period. The distinction between correct and error trials is done by fluctuations, which cause the activity to be either above or below a threshold.

However, the dynamics identified by the experimental data yield a different picture (Figure 1.9b). There, the direction selectivity of correct and error trials already separates in the stimulus period and reaches almost the same level at the end of the delay period. If we assume that the underlying dynamics is governed by an attractor system, we suggest that during correct trials the system jumps up to a high activity attractor state in the stimulus period, whereas during error trials this jump does not occur. This speaks in favor of computational errors implemented in the attractor landscape. Since correct and error trials reach the same direction selectivity at the end of the delay period, it is possible that the prefrontal cortex does not have an active part in the response execution. Other regions such as the premotor area might be important in this period. This is consistent with the previous concept of a processing pathway in which the prefrontal cortex comes before the premotor area.

Note that the external direction input implemented in this study is different from the context bias implemented by earlier modeling approaches (Deco & Rolls 2003, 2005b, Loh

& Deco 2005). The context input applies to the two pools, which represent the association active during an experimental trial, e.g. AL and BR. Here, we assumed that the basal ganglia generate the response-direction input, since neural recordings (Pasupathy & Miller 2005) have identified a strong direction selectivity in the basal ganglia activity and the direction selectivity also appears earlier in the basal ganglia than in the prefrontal cortex. We suggest that the main role of the prefrontal cortex is to integrate different evidence, as it is connected to almost every other brain region. It might have a leading role in learning and flexibility as its activity corresponds best with the behavioral performance.

In this modeling work, we feature a hypothesis driven approach. A priori we did not seek to fit the model to the data but wanted to test several existing hypotheses by integrating them in a computational model. Computational modeling is ideal for this approach since it allows both to integrate several hypotheses in one system and to study the contributions of each hypothesis separately. Moreover, the language of computational models is explicit and forces oneself to make concrete assumptions. We believe that the presented way of modeling, namely to test existing hypotheses instead of seeking model fits to experimental data, might contribute significantly to the theoretical understanding of brain dynamics, since it emphasizes the idea of integrating neuroscientific evidence in comprehensive models of brain functions.

1.4.4 Dynamical systems hypothesis of schizophrenia

As the next step, we wanted to go into learning of visuomotor mappings. The first three layer model did not address learning (see Section 1.4.2). The second prefrontal cortex model related a linear increase in the response-direction input to learning (see Section 1.4.3, Chapter 3). However, that study did not use any explicit learning rule either. The neuromodulator dopamine is believed to have a key role in learning, since it is among other things related to reinforcement learning theory. Dopamine neurons in the substantia nigra show firing patterns related to the reward prediction error: dopamine neurons fire when a reward is greater than previously expected. This is a key component of reinforcement learning models. However, the exact pathways and mechanisms through which dopamine acts are still unclear. We first wanted to investigate the effects of dopamine in cortical circuits, although also in this area no clear picture is available. After literature research in this area, we could not integrate the complex and at times even conflicting evidence of the effects of dopamine on cortical regions.

Most studies on dopamine take a bottom-up approach in which they describe the effects of dopamine using neurophysiological measurements. These experiments are usually *in vitro* and it seems that various experimental parameters like concentrations of buffers, temperature etc. affect the results. Inferring the main effects of dopamine in the human brain seems difficult. We implemented a different approach to the problem: We propose a top-down hypothesis (see also Chapter 4).

Schizophrenia is a major mental illness, which has a great impact on patients and their environment. The current medical treatment of schizophrenia uses dopamine antagonists, which block dopamine D2-type receptors. We aim to link the symptoms of schizophrenia and computational models of cortical function. Thereby we seek to infer which the effects of dopamine on cortical circuits might be.

In attractor networks, there are usually at least two distinct states: a spontaneous state marked by low neuronal firing rates, and a persistent state with higher firing rates in which one of the memory patterns is being maintained. Due to the influence of external inputs, or due to fluctuations caused by the stochastic spiking of the neurons, the network can switch

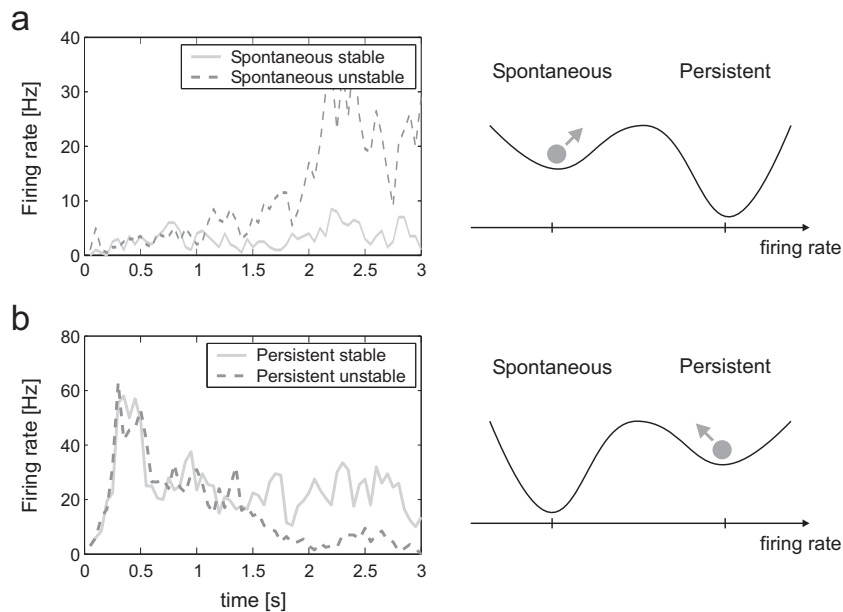


Figure 1.10: a. Behavior of an attractor network during spontaneous state simulations. Left: Sample trials, which are either stable or unstable. In the unstable spiking trial (dashed line), the firing rate escapes from the spontaneous attractor due to the influence of spiking-related statistical fluctuations. Right: The mechanism for escaping from the spontaneous attractor is pictured in a hypothetical energy landscape. We indicate for presentation purposes a ball jumping from the spontaneous state to the persistent state. The energy of each state is represented on the ordinate. b. Behavior of an attractor network during persistent state simulations. Left: The network is pushed into the persistent attractor state at the beginning, and due to the influence of statistical fluctuations the system leaves this state on some trials (dashed line). Right: This mechanism is pictured in the energy landscape as above.

between these two states. That is, if the network is firing at a low rate, an external input can move the network to a persistent state at which it maintains a high firing rate representing one of the memories stored in the network even after the external input is removed (Figure 1.10a). Fluctuations or external inputs can also switch the network back to a state of low, spontaneous firing rate (Figure 1.10b).

This behavior can be illustrated by an energy landscape (Hopfield 1982). Figure 1.10 shows such landscapes in which the attractor states (each one representing a memory) or fixed points of the network are indicated by the valleys. One can imagine a ball moving in that landscape, which is at rest at the bottom of the valleys. An extra force in terms of input or noise is needed to move the ball from one valley to another. In general, the hypothetical landscape can be multidimensional with several distinct attractor states, each one representing a different stored memory. We envision that the brain as a dynamical system has characteristics of such an attractor system including statistical fluctuations.

The stability of the persistent and spontaneous states of the network can be measured by the strength of the input needed to move the system from one state to another. One factor important in the stability is the depth of the basins of attraction. The depth depends on the strength of the synaptic coupling between neurons. If the valleys of the basins of

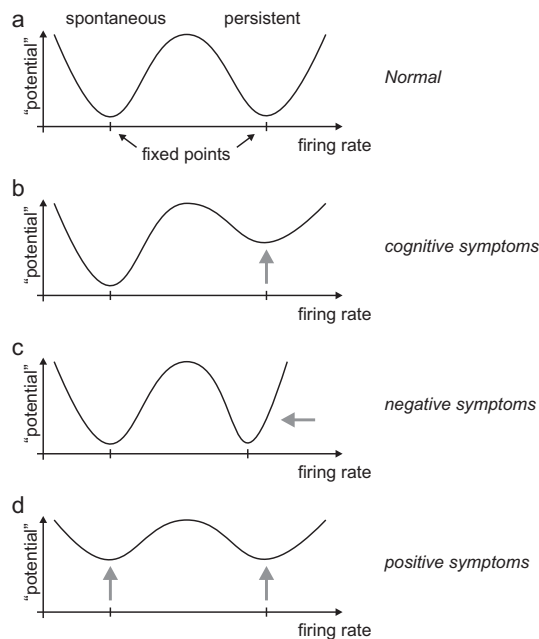


Figure 1.11: Hypothetical energy landscapes of dynamical systems related to the symptoms of schizophrenia. The abscissa shows the firing rate, and the ordinate the hypothetical potential. The potential shows the energy needed to change the firing rates of the system. a. Unchanged, normal condition. b. Cognitive symptoms might be caused by instability of the working memory attractor states, which are also called persistent states featuring high firing rates. This instability is envisioned to be related to working memory deficits, in which the system moves easily from one attractor to another, because the basins of attraction are shallow. c. Negative symptoms are related to decreases in firing rate in the orbitofrontal cortex and/or anterior cingulate cortex. d. Positive symptoms might result from shallow energy landscapes of both spontaneous and persistent states, which facilitate jumps between attractors and the spontaneous state, in semantic memory systems in the temporal lobe.

attraction are shallow as in Figure 1.11d compared to Figure 1.11a, then less force is needed to move a ball from one valley to the next. A second factor important in the stability is the noise in the system. High noise will make it more likely that the system will jump over an energy boundary from one state to another. Some of the noise in the system is due to the stochastical (probabilistic) firing of neurons whereby spikes of a number of neurons can occur relatively close together in time (within e.g. 20 ms), and this source of noise can make the system move from one state to another, depending on whether these statistical fluctuations happen to influence especially the neurons that form one of the attractor states or the spontaneous state. The probability of moving from a persistent attractor state to the spontaneous state might be different from the probability of moving from the spontaneous to a persistent attractor state (as illustrated in Figure 1.11b). We note that there will in general be many attractor states, each corresponding to a memory, in an attractor network, so that under the conditions illustrated in Figure 1.11b the system would move easily from one attractor state to another (with just one attractor basin shown in the diagram).

Our hypothesis builds upon the concept of shallow basins of attraction and contributes to an account of the different symptoms of schizophrenia (Rolls 2005). At the core we argue that

a changed attractor landscape together with the noise contributed by the statistical spiking fluctuations causes an altered signal-to-noise ratio for the transitions between the different attractors and the spontaneous firing state. The concept of an altered signal to noise ratio in schizophrenia is a concept of current interest (Winterer & Weinberger 2004), and we provide a dynamical attractor systems framework for analyzing how alterations in attractor networks might lead to an altered signal-to-noise ratio and to the symptoms of schizophrenia. We relate the three types of symptoms, cognitive dysfunction, negative symptoms, and positive symptoms (Liddle 1987, Mueser & McGurk 2004), to the dynamical systems framework as follows.

The cognitive symptoms of schizophrenia include distractibility, poor attention, and the dysexecutive syndrome (Liddle 1987, Green 1996, Mueser & McGurk 2004). At the core of the cognitive symptoms is a working memory deficit in which there is a difficulty in maintaining items in short-term memory (Goldman-Rakic 1994, 1999). We propose that these symptoms may be related to instabilities of persistent states in attractor neural networks, due to the fact that the neurons are firing at a lower frequency, leading to shallower basins of attraction, and thus a difficulty in maintaining a stable short-term memory, which normally the source of the bias in biased competition models of attention (Rolls & Deco 2002, Deco & Rolls 2005*a*). The shallower basins of attraction as illustrated in Figure 1.11b would result in distractibility, poor attention, and working memory difficulties.

The negative symptoms refer to the flattening of affect and a reduction in emotion. Behavioral indicators are blunted affect, emotional and passive withdrawal, poor rapport, lack of spontaneity, motor retardation, and disturbance of volition (Liddle 1987, Mueser & McGurk 2004). We propose that these symptoms are related to decreases in firing rates in the orbitofrontal cortex and/or anterior cingulate cortex (Rolls 2005), where neuronal firing rates and activations in fMRI investigations are correlated with reward value and pleasure (Rolls 2005, 2006). This is illustrated in Figure 1.11c by a reduced firing rate of the fixed point of the persistent attractor.

The positive symptoms of schizophrenia include bizarre (psychotic) trains of thoughts, hallucinations, and (paranoid) delusions (Liddle 1987, Mueser & McGurk 2004). We propose that these symptoms might result from a shallow energy landscape (Figure 1.11d) in the temporal lobe semantic memory networks in which the attractor states are weak and the system jumps spontaneously between attractor states. The thoughts wander loosely between weakly associated attractors, leading to bizarre associations, which may eventually over time be associated together in semantic memory to lead to false beliefs and delusions.

We analyze an attractor network, which is a component of cortical microcircuitry, focusing on the contribution of NMDA and GABA receptor-activated synaptic currents. We use an architecture of two pools (see Figure 1.7c), because it enables us not only to study the interplay between the spontaneous and the persistent state, but also the interplay of two persistent states. In particular, we used the two pool setup to study distractibility. Distractibility is a common symptom in schizophrenia.

We assessed how the stability of both the spontaneous and persistent states changes when NMDA and GABA efficacies are modulated. Figure 1.12 shows the percentage of sustained activities during spontaneous and persistent simulations. We assessed how often the system maintained the spontaneous or persistent state, assessed by the firing rate in the last second of the simulation (2–3 s) of each 3 s trial. This measures the stability of the system with no external distractor stimuli being applied. In the normal (unmodulated) condition, the spontaneous and persistent states are maintained in a high percentage of cases (around 90%).

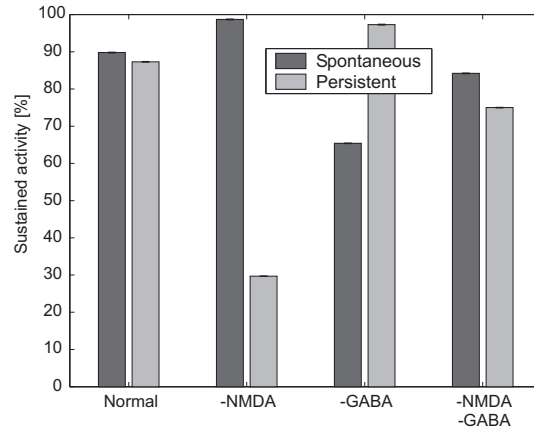


Figure 1.12: Stability of the spontaneous and persistent state as a function of the modulations of the synaptic efficacies. We assessed how often in 1000 trials the average activity during the last second (2–3 s) stayed above 10 Hz. A modulation of the synaptic currents shown as -NMDA and -GABA corresponds to a reduction of 4.5% and 9% respectively in their efficacies. The sustained activity refers to the percentage in which the 1000 trials stayed in the respective state, i.e. a high activity (above 10 Hz) for the persistent simulations, and a low activity (below 10 Hz) for the spontaneous simulations. The standard deviations of the percentage of sustained activity (shown above each bar) were approximated with the binomial distribution and are below 0.1% for all conditions.

A reduction of the NMDA conductance (-NMDA) reduces the stability of the persistent state drastically, while slightly increasing the stability of the spontaneous state (see Figure 1.12). We hypothesized that such a pattern might be related to the cognitive symptoms, since it shows a reduced stability of the working memory properties (see Figure 1.11b). (The core of the cognitive symptoms is a working memory deficit (Goldman-Rakic 1994, 1999).)

A reduction of GABA shows the opposite pattern: A slight reduction in the stability of the spontaneous state, and an increased stability of the persistent state (see Figure 1.12).

When both NMDA and GABA are reduced one might think that these two counterbalancing effects (excitatory and inhibitory) would either cancel each other out or yield a tradeoff between the stability of the spontaneous and persistent state. However, this is not the case. The stability of both the spontaneous and the persistent state is reduced (see Figure 1.12). We relate this pattern to the positive symptoms of schizophrenia, in which both the spontaneous and attractor states are shallow, and the system merely jumps by the influence of statistical fluctuations between the different attractor states (Figure 1.11d).

We relate the negative symptoms to a reduction of the mean firing rate of the persistent state of networks in for example the orbitofrontal cortex. A reduction of the firing rate was produced when the NMDA current was reduced (see also Brunel & Wang (2001)). Thus, the cognitive and negative symptoms of our hypothesis can be related to the same synaptic mechanism, namely a reduction of NMDA conductance. An additional reduction of the GABA current leads to change in activity of the network that models the positive symptoms according to our hypothesis.

Overall, our approach shows that the cognitive and negative symptoms could be caused by the same synaptic mechanism, namely a reduction in the NMDA conductance, which reduces the stability of the persistent attractors, and reduces the firing rates of the neurons compared

to the reference baseline condition. The positive symptoms could be accounted for in part by the same mechanism, namely a reduction of NMDA, but in addition reduction of the GABA conductance facilitates movement of the activity from the spontaneous state to a persistent state. This could be a mechanism that would, in the temporal lobe, lead to intrusive thoughts.

Schizophrenia is treated with dopamine receptor D2 antagonists which mainly alleviate the positive symptoms, whereas the cognitive and negative symptoms persist, especially for the typical neuroleptics (Mueser & McGurk 2004). We found that the state corresponding to the positive symptoms (-NMDA, -GABA) and the one corresponding to the cognitive/negative symptoms (-NMDA) differ in the modulation of the GABA conductance. We reason that the net effect of neuroleptics might be an increase in the GABA conductance. This is also consistent with experimental work: it has been found that D2 receptors decrease the GABA contribution (Seamans et al. 2001, Trantham-Davidson et al. 2004).

Although positive symptoms can be treated effectively in many patients, the negative and cognitive ones typically persist. To ameliorate the cognitive symptoms, the persistent state according to our hypothesis needs to be stabilized (see Figure 1.11b). In the computational model investigated, the stability of the working memory state is linked to NMDA receptor mediated currents (see also Durstewitz et al. (2000b), Wang (2001), Deco (2006)). A possible pathway to increase the NMDA currents could be via the D1 receptor (Durstewitz & Seamans 2002, Seamans & Yang 2004).

We emphasize that our hypotheses are not dependent on any particular interpretation of the biological causes of schizophrenia, which might be described at several levels of abstraction in the brain. These range from the single neuron level up to systems neuroscience approaches (Carlsson 2006). The latter address complex networks often involving cortical and subcortical regions and including neuromodulators such as dopamine and serotonin, which affect the dynamics of the interconnected levels in distinct ways. We hypothesize that the biological changes in schizophrenia eventually cause the effects described above in a dynamical attractor system. The overall goal is to investigate the pathways and mechanisms that lead to the hypothesized alterations in the dynamical attractor system.

We investigated one possible cause, namely alterations in the NMDA and GABA conductances, in more detail. The results are consistent with experimental and modeling work of dopamine, which further supports our neurodynamical hypothesis. We have shown what the net effects of dopamine on cortical microcircuits might be: D2 antagonists might act on inhibitory currents, whereas D1 might act on the excitatory NMDA contribution. As an interesting modeling result, we demonstrate that inhibition and excitation are not merely antagonistic, but have distinct effects on the attractor landscape. A reduction of the NMDA and GABA contributions destabilizes the persistent attractor and spontaneous states respectively. This highlights the importance of the statistical fluctuations caused by the probability of spiking of neurons in the modeling of brain activity.

1.4.5 Neuropsychological model of novel visuomotor learning

After the excursion to the field of schizophrenia, we go back to the original topic and study the learning of visuomotor associations. As a starting point, we used the reinforcement learning theory, which was the original motivation to start the project on dopamine and schizophrenia. This project has a different motivation than the previous projects, as we do not work with a neurophysiologically realistic model and therefore do not address neural data (see also Chapter 5).

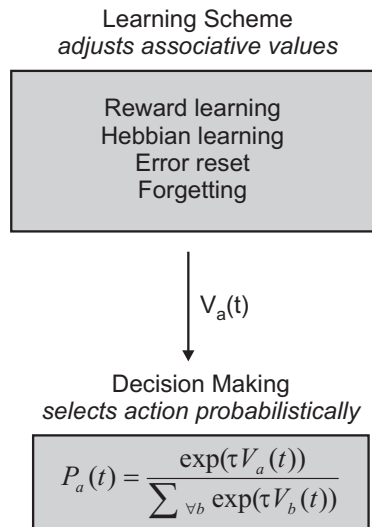


Figure 1.13: Overview of the two step algorithm, which was used to model visuomotor learning. The learning scheme adjusts the associative values $V_a(t)$ upon the received rewards (for details see Figure 1.14). Upon the associative values, the decision-making component selects an action given by a softmax algorithm.

A computational model should not only account for the data of one scientific article, but also integrate a variety of findings. In addition, if the articles in question all address the same ability, namely visuomotor learning, the underlying data should be comparable. Since all articles of visuomotor learning publish behavioral data, we sought to construct a behavioral model, which could account for several visuomotor learning experiments. Moreover, there is currently no model, which tries to address the great amount of neuropsychological data. In lesion studies, a part of the brain is damaged to investigate the contribution of this part of the brain in the execution of the task. Due to the high interconnectivity of the brain, the interpretation of lesion studies is difficult and a refined computational approach might give more insights.

As a starting point for our model, we used classical reinforcement learning theory, which is able to solve the associative learning problem (Sutton & Barto 1998). In reinforcement learning, the algorithms are made up of two parts: Action value estimation and action selection. In the action value estimation, the value of each action $V_a(t)$ is estimated upon receiving the rewards. As the update rule for the action values, we use an incremental form, which uses a reward prediction error.

$$V_a(t+1) = V_a(t) + \alpha(r(t) - V_a(t)), \quad (1.3)$$

where $V_a(t)$ is the value of action a at step t , α is the learning parameter and $r(t)$ is the reward which is received (either 1 for a reward or 0 for no reward). The term $(r(t) - V_a(t))$ represents the reward-prediction error. The values $V_a(t)$ are updated during learning for each association separately. This update mechanism has been related to the reward-prediction error signal of dopamine neurons and to the theory of classical conditioning (Rescorla & Wagner 1972). In order to select an action based upon these values, we use softmax action selection (see Figure 1.13). It chooses the action with the highest action value with the highest probability.

	Correct trial	Error trial
Executed response	$c_{rew} = \alpha_{corr}(1 - V_a(t))$ $c_{hebb} = \alpha_{hebb}(1 - V_a(t))$ $c_{forg} = -\alpha_{forg}V_a(t)$	$c_{res} = -\alpha_{rew}V_a(t)$ $c_{hebb} = \alpha_{hebb}(1 - V_a(t))$ $c_{forg} = -\alpha_{forg}V_a(t)$
Not executed response(s)	$c_{forg} = -\alpha_{forg}V_a(t)$	$c_{res} = -\alpha_{rew}V_a(t)$ $c_{forg} = -\alpha_{forg}V_a(t)$

Figure 1.14: Summary of the learning scheme. We depict which rules apply in which condition. The trials are divided in either correct or incorrect trials. The table shows which rules apply to the executed and the not executed response depending on the reward outcome. The contributions of reward-based learning c_{rew} , hebbian-based learning c_{hebb} , error reset c_{res} , and forgetting c_{forg} are later summed to adjust the associative values $V_a(t)$ of the actions.

We use this basic two-step process in our model. Our literature research showed that the learning rules are more complex in visuomotor learning. There is evidence for four learning mechanisms (see Figure 1.14):

- Reward-based learning is necessary to integrate the information about the reward. It increases the action value or associative value of the executed response when a reward is received.
- Hebbian-based learning has been hypothesized to slow down the learning after an error (Brasted et al. 2005). The Hebbian-learning component is independent of the reward and increases the associative value of the executed stimulus-response pair.
- Fusi et al. (2007) analyzed the behavioral data of Asaad et al. (1998) further, and found that after an error, the behavior showed a reset: The learned associations were forgotten and the learning started anew. After an error, we decrease the associative values of all responses, which belong to an action, regardless if it is executed or not.
- Learned associations were also forgotten when trials intervened between stimulus presentations. This mechanism takes place between trials and applies to all associations, which are stored in memory.

We implemented these four mechanisms and the softmax action selection in an agent, which can interact with different environments. These different environments are variations in the experimental paradigm, which have been used in publications. We first fit our model parameters to the experimental data by Brasted et al. (2005) in which 3 stimuli were mapped onto 3 responses with correction trials: The monkey saw a stimulus and then gave one of three possible responses. If the response was correct, a reward was given. If the response was incorrect, no reward was given and the same stimulus appeared on the next (correction)

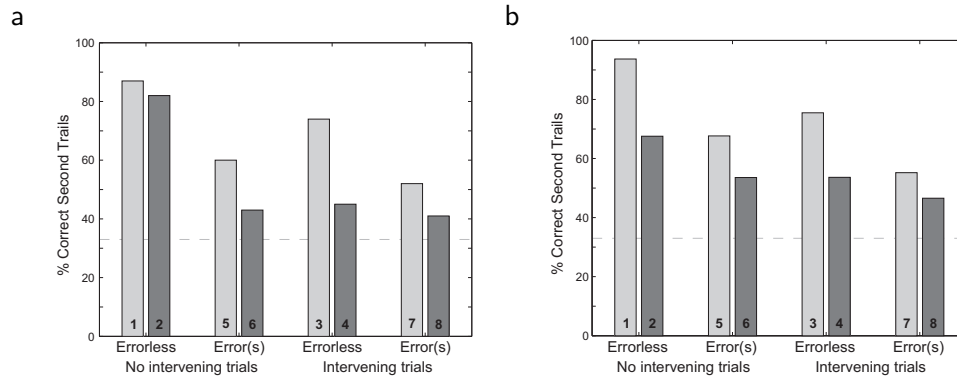


Figure 1.15: Average correct responses of the second presentation of a stimulus regardless when it occurs during a session and excluding correct trials (see also Figure 5.1). The trials can be categorized into two classes: first, if an error occurred during the first presentation of that stimulus (errorless / error trials) and, second, if other stimuli were presented between the first and second presentation (no intervening / intervening trials). (a) Experimental data by Brasted et al. (2005), see their Figure 3a (reproduced with authors' permission). The white bars denote the control condition; the gray bars the fornix transection condition. (b) Model fit with parameters $\alpha_{corr} = 0.1$, $\alpha_{hebb} = 0.4$, $\alpha_{res} = 0.7$, and $\alpha_{forg} = 0.05$. In the gray bars, the Hebbian learning parameter is reduced ($\alpha_{hebb} = 0.1$).

trial until the answer was correct. Figure 1.15a shows the average correct responses for the second presentation of a given stimulus, regardless of when it occurred during a session (Brasted et al. 2005). In the simplest case, the monkey responded by chance correctly on the first stimulus presentation, which marks an "errorless" trial. When the same stimulus was presented in the following trial again, there were "no intervening trials" between the two stimulus presentations. For this condition, the percentage of correct responses on the second presentation is given in Bar 1 in Figure 1.15a. If the monkey made an error on the first presentation, the trial would be in the "error" condition (Bar 5, Figure 1.15a). Note that the first trial always ended with a correct response due to the correction trials. Overall, errors slow down learning and worsen the performance of the second trial. Brasted et al. (2005) reasoned that Hebbian learning mechanisms might be responsible for that by strengthening the associations of executed actions even if they yielded an error (McClelland 2001).

Our model contains four parameters, which are the learning rates of the four learning mechanisms. We found a good fit to the data with the values: Reward learning $\alpha_{corr} = 0.1$, Hebbian learning $\alpha_{hebb} = 0.4$, error reset $\alpha_{res} = 0.7$, and forgetting $\alpha_{forg} = 0.05$. Figure 1.15b shows the resulting values for the model. The Hebbian learning rate is much higher than the Reward learning rate. This is surprising, as the Hebbian mechanism does not take the reward into account. In addition, there is a substantial learning when the positive reward learning is at 0. This is obvious for the errorless case (as no errors have been made). In the error case, the Hebbian learning strengthens all performed associations while the error reset eliminates the erroneously increased associations. This causes a positive performance, although no positive reward learning is present.

The next step is to work with this model fit and compare it to other experimental evidence. We discuss it in three ways. First, we check if the model indeed shows the reset properties described in Fusi et al. (2007), to see if the two qualitative findings can be integrated in our

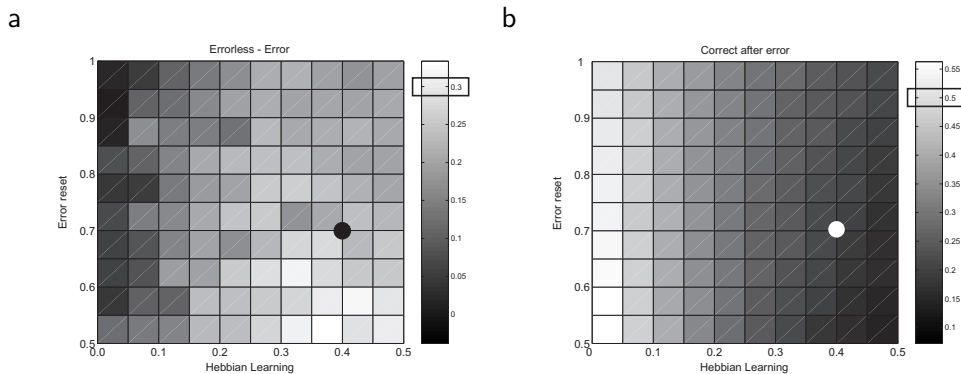


Figure 1.16: Using our model, we compare the measures of Brasted et al. (2005) and Fusi et al. (2007) in terms of dependence on the Hebbian learning parameter α_{hebb} and the error reset parameter α_{res} (see also Figure 5.3). a. Difference between the errorless and error condition (no-intervening trials) of second trial performance (Bar 1–Bar 5, Figure 1.15). It should be around 30% according to the experimental data. b. Performance after an error computed with the paradigm investigated by Fusi et al. (2007). It should be at 50%. The boxes at the color bars indicate the values according to the experimental data. The dots denote the parameter values of the reference condition.

model and maybe adjust the error reset parameter. Second, we address the main issue of the experiment by Brasted et al. (2003, 2005), namely the lesion of the hippocampal system. Third, former studies have shown that the monkey used a strategy to solve the task. If the stimulus repeated, they stayed with the previous response. If the stimulus changed, they picked another response (repeat-stay/change-shift). We added another module to the system, which complements the learning system with strategy characteristics. The strategy model is more detailed in Chapter 5.

Fusi et al. (2007) identified in their experimental data an error reset, which set the probabilities of a correct response to chance level after an error. The parameter, which is opposing the error reset, is the Hebbian learning parameter, because it reinforces the stimulus-response association when an error is made (see Figure 1.14). We computed the difference between the errorless and error case for the paper by Brasted et al. (2005) (which is achieved due to Hebbian learning) and the performance after an error for the experiment analyzed by Fusi et al. (2007). These two measures are contrasted in Figure 1.16. For achieving a difference between the errorless and the error case, a strong Hebbian learning and a weak error reset is required (Figure 1.16a). For achieving a reset after an error, we need a weak Hebbian learning and a high error reset value (Figure 1.16b): A strong Hebbian component, which slows down learning, is incompatible with the error reset. If an error reset would have indeed governed the behavior of Brasted et al. (2005), the performance of the errorless and error case should have been equal. However, they were significantly different. In conclusion, the experimental data from Brasted et al. (2005) and Fusi et al. (2007) show contradicting properties in their learning mechanisms.

The main focus of the studies by Brasted et al. (2003, 2005) were the effects of lesions of the hippocampal system. They performed a fornix transection, which impaired the learning of novel associations. The hippocampal formation has been related to the fast learning of associations. It has been hypothesized that it mainly works as a pattern associator which

binds items such as stimuli and responses (McClelland et al. 1995, McClelland 2001). This is reflected by the Hebbian learning component of our model. Thus, does a reduction in the Hebbian learning component reflect the effects of the fornix transection in the experimental data? We reduce the Hebbian learning parameter by 75% (from 0.4 to 0.1) to account for the fornix transection. Figure 1.15 shows the comparison of the experimental data and the model for the performance of the second trial. We find a good fit with the experimental data (except in the condition with no errors and no intervening trials). The absence of response strategies in our model might partially account for the differences (Brasted et al. 2005). In terms of learning, a reduction of the Hebbian learning component can account for large parts of the effects of a fornix transection, which implies that a lesion of the hippocampus could eliminate much of the associative capability of learning.

Overall, we proposed a behavioral model for learning of novel associations and fitted it to the experimental data by Brasted et al. (2005). Thereafter we compared the model to evidence by Fusi et al. (2007). We found inconsistencies in the published data and it will require a further analysis of the experimental data to identify the causes. Furthermore, we showed that the hippocampus might drive the Hebbian-based stimulus-response learning. The model itself stands as a specific prediction and can easily be implemented and compared to other experimental paradigms. We envision that it can serve as a reference to compare different experimental setups. Differences in the learning properties caused by different experimental paradigms could address learning components to various degrees. Getting an insight into this scheme might be important to discuss the underlying neural correlates in both neurophysiology and neuropsychology.

1.5 Conclusion and future ideas

The learning of visuomotor mappings is processed in a system, which covers several interacting brain regions. We have introduced the term "processing pathway" (Loh & Deco 2005) to describe the way we envision the brain performs this task. Visual information originates in the inferotemporal cortex (IT), which is known to encode stimulus identities, and the processing ends at the premotor cortex (PMC), where the final decision and motor selection is made. In between, a system comprised of the prefrontal cortex (PFC), the hippocampus, and the basal ganglia are engaged in learning the mappings (see Figure 1.17). We distinguish systems for novel and familiar learning, which can also be separated neuropsychologically.

The system for the novel learning consists of the PFC, the basal ganglia, and the hippocampus. We studied a possible influence of the basal ganglia on the PFC in Loh et al. (2008). We suggest that the basal ganglia propose a decision about the upcoming action to the prefrontal cortex and influence its dynamics with a direction-selective bias via the thalamus. This direction-selective input emerges shortly after the presentation of the stimulus and is transient. The hippocampus implements a Hebbian-based mechanism, which does not take the reward into account. Our behavioral modeling (see Chapter 5) shows that this component could be strong at least in some experiments. We envision that the prefrontal cortex integrates influences of both the basal ganglia and the hippocampus. The resulting activity influences the PMC, which makes the decision of the upcoming action. In case of an error trial, the activity of the prefrontal cortex shows a much lower direction selectivity (Pasupathy & Miller 2005), which suggests that the influence of the prefrontal cortex is lower: The PMC has less information upon which to take a decision and the stochastic dynamics "guesses"




Figure 1.17: Overview of the brain regions involved in visuomotor learning and their participation in both novel and familiar associations. PFC: prefrontal cortex, PMC: premotor cortex, HC: hippocampus, BG: basal ganglia, IT: inferotemporal cortex.

a choice. Trial-and-error learning is a key ingredient of visuomotor learning. If a reward is received, the system receives an increased input from dopamine neurons, which strengthens the connections related to the executed association. An error might in some cases cause a reset of the system (Fusi et al. 2007), although the exact mechanism of this reset is not clear yet. An additional mechanism, which supports the learning, is the use of strategies. It is envisioned that the prefrontal cortex supports this high level function, which has also been substantiated by lesion studies (Wang et al. 2000, Bussey et al. 2001). The strategy-selective activity might directly influence the decision making process of the PMC, as implemented in our model in Chapter 5. Our results indicate that the strategy mechanism might work in addition to the learning mechanism and does not substitute it.

The system, which supports the execution of familiar mappings, is made up of the IT, the basal ganglia, and the premotor cortex. Lesion studies have shown just a minor participation of the prefrontal cortex in familiar association (Wang et al. 2000). We propose that the IT provides the basal ganglia directly with the stimulus information. A cross-lesion study of the basal ganglia and the premotor cortex has shown that this system is crucial for the execution of familiar associations (Nixon et al. 2004). We suggest that a loop between the PMC and the basal ganglia builds up a memory for learned associations and takes part in the decision making process.

What happens if these systems are out of balance? We discussed aberrant signaling of both NMDA and GABA conductances and showed instabilities in a neurodynamical framework. We relate these instabilities to the symptoms of schizophrenia (Loh et al. 2007). The symptoms of schizophrenia are diverse, which can also be seen in the complex system of visuomotor learning. If instabilities affect one part of the interconnected system, other parts will also show aberrant dynamics. The strength of our approach is that we relate the types of symptoms of schizophrenia to different instabilities in attractor networks, which are not confined to specific characteristics of one brain region. We can on the one hand describe how the different symptoms are related and explain their main properties, but on the other hand also address the heterogeneity of the illness. This might contribute to a deeper understanding of schizophrenia.

Computational modeling has shown to be an important tool in the analysis and discussion of brain dynamics. First of all, it forces one to make concrete model assumptions and hypotheses, since these are implemented in a mathematical model. The effects of these assumptions and hypotheses can be then studied in detail in the model, which might also result in adjustments of hypotheses. Second, computational modeling allows the direct comparison

of experimental data to the model. This requires a sufficiently detailed biophysical model. This is increasingly important, since the measures used in experimental publications get more and more complex, and gaining an insight into the effects of these measures is crucial for the discussion of the results. Third, computational models can serve both as a specific prediction and as a reference to compare several experimental papers. The direct comparison is sometimes difficult, as experiments might use different paradigms or setups. A computational model can integrate several bodies of evidence and show inconsistencies, which can later be addressed by experimental researchers. Fourth, the concepts used in computational models generate new ideas and concepts. The concept of attractor landscape and its relation to schizophrenia is a recent example presented in this thesis.

In spite of the advances in research, the field of conditional visuomotor mappings is very diverse and even the behavioral properties seem not to be completely understood. Future modeling work could apply the behavioral model in Chapter 5 to more neuropsychological data. Earlier experimental data such as by Rupniak & Gaffan (1987), Gaffan & Harrison (1988, 1989) could serve as a starting point. In addition, the behavioral data of recent neurophysiological articles could be analyzed further. In neurophysiological articles, the behavior is typically reported only briefly and there is further potential to get insights into the underlying mechanisms. Not all parameters have been included in the behavioral model in Chapter 5 so far. For example, we completely neglected the time course of the experiment. There are also interesting projects regarding the neurophysiology: Genovesio et al. (2005) and Genovesio et al. (2006), presented detailed neural data on the repeat-stay/change-shift strategy. The relationship between the different kinds of neurons is still unknown. A neural model, which could integrate the different kinds of neurons and perform the underlying task, could give new insights into the experimental data. Here, the level of a connectionist model would be sufficient as a first approximation to describe the relationship between the types of neurons.

So far we strictly confined our research to primate studies. However, the overall goal is to understand the human brain. Are the behavioral properties of humans comparable to monkey data? What can we learn about the human brain from the research on animal behavior? The research actually speaks strongly against comparing animal and human behavior: Not even the animal behavior is consistent in this paradigm. After a more detailed characterization of the animal data, a characterization of the human behavior would also be interesting. Hadj-Bouziane & Boussaoud (2003) already suggested a common effect in animal and human behavior. For this purpose, we already implemented a web-based project with which such basic behavior can be studied. We used standard web technology to create a simple visuomotor learning experiment, which can be played in every web browser. This could facilitate the acquisition of experimental data, since everybody can play it over the web. As a technological basis, we used on the server side JAVA technology consisting of a Tomcat server and the database JavaDB. We implemented a standard model-view-controller architecture using JavaBeans, JavaServerPages, and Servlets. On the client side, we opted for standard technology of XHTML, JavaScript, and CCS. This is subject to future research.

In the field of schizophrenia, we presented a general neurodynamical model, which covers the different types of symptoms. Future research could go in two directions. On the one hand, one could look at specific symptoms and relate them in a more detailed manner to the framework presented. For example, the auditory hallucinations could be addressed by investigating in detail the effects of neuroleptics and the evidence of brain imaging and relate them to shallow basins of attraction. On the other hand, the concept of a lower signal-to-noise ratio is central to the hypothesis. A lower signal-to-noise ratio has been identified

in imaging experiments (Winterer et al. 2000, 2004, 2006). Since the modeling work so far is just a proof of principle regarding this aspect, a detailed modeling of the corresponding experiments including the exact time course would be important to further substantiate the hypothesis.

Overall, the current work in the fields of visuomotor learning and schizophrenia yielded insights and new ideas for research. However, it has also raised more questions, which need to be addressed in future work.

Chapter 2

Cognitive flexibility and decision making

*The work presented in this chapter is published in the *European Journal of Neuroscience* (Loh & Deco 2005).*

Abstract

Arbitrary visuomotor associations map a stimulus onto a particular response upon reinforcing rewards. Changes in the associations between stimuli and responses require the neural networks to discard the already-learned mappings and build new ones. This is a key aspect of cognitive flexibility. In addition, learning within this experimental paradigm requires a trial-and-error exploration strategy of the available stimulus-response associations. A system performing this task must be able to both build up mappings for stimulus-response associations and at the same time perform non-deterministically to explore the behavioral space until it reaches certainty. We hypothesize an integrate-and-fire network model that accounts for the neurophysiological data of a conditional visuomotor association task and is able to show non-deterministic behavior. We aim at identifying multistable attractor regimes in the network dynamics which intrinsically enable the system to make errors and thereby to perform trial-and-error exploration. Our model combines cognitive flexibility with multistable attractors in neurodynamical systems, believed to be the basis of decision making. If multistable attractors indeed support the exploration of the behavioral space, then our model predicts that the brain should respond stochastically with correct or incorrect activity to visuomotor associations until it has reached certainty. This should be visible not only in the corresponding activity in the premotor area, but also in activity corresponding to other associations or even other stimuli in the prefrontal cortex.

2.1 Introduction

Arbitrary stimulus-response associations build the basis for many complex behaviors of animals and humans. Due to the basic nature of this paradigm, it has been studied intensively in neuropsychology, brain imaging, and neurophysiology (for reviews see Murray et al. (2000), Passingham et al. (2000), Hadj-Bouziane et al. (2003)) and it incorporates a wide variety of sub-tasks: Identification of the stimuli, retrieval of associations, decision making, motor response, and processing of the rewards. Several imaging studies have identified a network that carries out this task and spans both cortical and subcortical areas (Passingham et al. 2000, Boettiger & D’Esposito 2005). Electrophysiological recordings specify the characteristics of these brain areas further, namely the premotor cortex (Mitz et al. 1991, Brasted & Wise 2004), the supplementary eye field (SEF) (Chen & Wise 1995*a,b*), prefrontal areas (Asaad et al. 1998, 2000, Pasupathy & Miller 2005), the hippocampus (Cahusac et al. 1993, Wirth et al. 2003) and the basal ganglia (Tremblay et al. 1998, Inase et al. 1999, Brasted & Wise 2004, Pasupathy & Miller 2005). Clinical studies in humans also show that excisions in prefrontal or premotor areas severely impair the learning of visuomotor mapping tasks (Petrides 1985, 1997). We investigate a neurodynamical model which addresses this type of task.

The neurodynamical model for the conditional visuomotor learning task (Deco & Rolls 2003) covers both the basic behavior of neurons in the prefrontal cortex as measured experimentally (Asaad et al. 1998) and a mechanism which keeps track of contextual contingencies upon the responses of the reward system (Deco & Rolls 2005*b*). Decision making is a crucial part of forming and executing conditional visuomotor associations. At the behavioral level, it is a well studied topic in psychology. The so-called diffusion models explain a wide range of experimental results (Ratcliff et al. 1999). It has also been suggested that the prefrontal cortex is involved in decision making (Passingham 1993, Passingham & Toni 2001). The synaptic, cellular, and network mechanisms of decision making can be understood in the context of neural networks exhibiting multistability (Wang 2002). It is not clear if the neurodynamical model hypothesized by Deco & Rolls (2003) supports multistability. We aim here at providing this link.

Furthermore, the organization of the system performing conditional visuomotor tasks is an open question, especially in terms of which brain regions contribute to specific functions. The visuomotor association model (Deco & Rolls 2003) assumes a processing pathway from stimuli to motor circuits mediated by biasing signals which are thought to originate from the prefrontal cortex (Miller & Cohen 2001). In this sense, one can view the levels of the conditional visuomotor association model as brain regions along the processing path. We discuss the connectivity between these abstract brain regions and thereby show the dependencies required for decision making characteristics.

2.2 Methods

2.2.1 Conditional Visuomotor Task: Experimental Paradigm

We model the experimental paradigm of the conditional visuomotor task used by Asaad et al. (1998), Asaad et al. (2000) and Pasupathy & Miller (2005). In this task, stimuli were presented foveally and monkeys had to make an oculomotor response after a subsequent delay period. The response, a left- or rightward eye movement, was learned on a trial-and-error basis by giving the monkey a liquid reward for performing the correct response. The contingencies were

reversed frequently after the monkey performed at a high level. For example, after having learned the pairings “object A, saccade right” and “object B, saccade left” (direct rule), the associations were reversed such that “object A, saccade left” and “object B, saccade right” were rewarded (reverse rule).

Three types of neurons were reported by Asaad et al. (1998) in the dorsal prefrontal cortex. The neurons were either object, direction or object-and-direction selective. First, object selective neurons showed selectively higher activity to one of the objects but no effect with regard to the saccade direction. Second, direction selective neurons showed a significant effect of the saccade direction but not of the presented object. Third, cells were selective to combinations of object and direction. Particularly interesting for the decision making process is that the direction-selective activity reflected the forthcoming saccade, irrespectively of whether the response was correct or not. Asaad et al. (1998) argue that this activity could be due to a premotor signal or a shift in visual attention preceding the eye movement. We focus on whether the proposed neurodynamical model can indeed reflect activity of direction-selective neurons, not only in terms of a correctly computed association, but also in terms of a dynamic decision making behavior. That is, the system can respond stochastically to explore the response space and thereby identify the correct associations.

2.2.2 The Integrate-and-Fire Network

We use a neural network model based on integrate-and-fire neurons and complement it with the necessary concepts and techniques to study higher cognitive tasks. In this section, we describe the single neuron level of the network model. The neural network basis is adapted from Brunel & Wang (2001) and thus we refer to their work for a more detailed description of the basic neurodynamical properties. The detailed mathematical description is provided in the Supplementary material.

We use non-linear leaky integrate-and-fire neurons to model both excitatory and inhibitory neurons. The basic behavior of a single neuron is reflected in the membrane potential which is modeled by an electrical circuit consisting of a capacitor C_m connected in parallel with a resistor R_m . This circuit describes the progression of the membrane potential $V(t)$ in time depending on external currents entering from other neurons. When the threshold membrane potential V_{thr} is reached, the neuron is set to the reset potential V_{reset} at which it is kept for a refractory period τ_{ref} and the action potential is propagated to the other neurons. Both excitatory and inhibitory neurons have a resting potential $V_L = -70mV$, a firing threshold $V_{thr} = -50mV$ and a reset potential $V_{reset} = -55mV$. The membrane parameters are different for both types of neurons: Excitatory neurons are modeled with a membrane capacitance $C_m = 0.5nF$, a leak conductance $g_m = 25nS$, a membrane time constant $\tau_m = 20ms$, and a refractory period $t_{ref} = 2ms$, whereas the inhibitory neurons share $C_m = 0.2nF$, $g_m = 20nS$, $\tau_m = 10ms$, and $t_{ref} = 1ms$. Values are extracted from McCormick et al. (1985).

The incoming synaptic influences are both excitatory and inhibitory. The excitatory neurons transmit their action potentials via the glutamatergic receptors AMPA and NMDA which are both modeled with exponential terms. We neglect the rise time of the current mediated by the AMPA channel, because it is typically very short ($\ll 1ms$), and just model the decay period with $\tau_{AMPA} = 2ms$ (Hestrin et al. 1990, Spruston et al. 1995). The NMDA channel is modeled with a rise term ($\tau_{NMDA,rise} = 2ms$), a decay term ($\tau_{NMDA,decay} = 100ms$) (Hestrin et al. 1990, Spruston et al. 1995), and an extra voltage dependence controlled by the extracellular magnesium concentration ($C_{Mg^{2+}} = 1mM$, Jahr & Stevens (1990)). The

inhibitory postsynaptic potential is mediated by the GABA_A receptor and is described by a decay term ($\tau_{GABA} = 10ms$) (Salin & Prince 1996, Xiang et al. 1998).

The network is fully connected with $N_E = 1600$ excitatory neurons and $N_I = 400$ inhibitory neurons, which is consistent with the observed proportions of the pyramidal neurons and interneurons in the cerebral cortex (Braitenberg & Schütz 1991, Abeles 1991). The connection strengths are adjusted using meanfield analysis (Brunel & Wang 2001), so that the excitatory and inhibitory neurons exhibit a spontaneous activity of 3Hz and 9Hz, respectively (Koch & Fuster 1989, Wilson et al. 1994). The recurrent excitation mediated by the AMPA and NMDA receptors is assumed to be dominated by the NMDA current to provide a more robust behavior during the delay periods.

This recurrent connectivity represents the baseline for the structuring of the network. We assume that the structuring of the network connections arises due to Hebbian learning (Hebb 1949). The coupling strength between two neurons is stronger if they have correlated activity, and weaker otherwise. As a consequence of this, neurons that share the same selective response properties are stronger mutually connected and thus are clustered into pools. All inhibitory neurons are in a common pool so that there is global inhibition throughout the network.

In addition, the network is exposed to external Poisson input spikes mediated by AMPA receptors of $N_{ext} = 800$ external neurons at average rate of 3Hz per neuron, consistently with the spontaneous activity observed in the cerebral cortex (Wilson et al. 1994, Rolls & Treves 1998).

2.2.3 The Stimulus-Response Architecture

The network architecture is based on the neurophysiological evidences of Asaad et al. (1998) and the processing concepts of cognitive control and biased competition (Miller & Cohen 2001, Rolls & Deco 2002). It has certain features in common with the ones used by Deco & Rolls (2003) and Deco & Rolls (2005*b*). However, there are fundamental differences not only in the model but also its analysis.

We assume three different types of selective neural populations (pools) as identified by Asaad et al. (1998) (see Figure 2.1). The first type shows object-selective responses to presented stimuli (A, B). The second type responds to a combination of a particular object followed by a particular response, in specific, a neuron of this type would for example respond selectively to object A followed by a rightward saccade R (AR), but not to other combinations of stimuli and responses (AL, BL, BR). The third type shows direction-selective responses of the forthcoming saccade (L, R). Each of these eight pools contains 80 neurons, and the rest of the excitatory neurons comprise a non-specific pool. Neurons in the latter show spontaneous firing rates and introduce noise in the system, which aids in generating Poisson like firing patterns throughout the system.

Figure 2.1 gives a conceptual view of the architecture. The selective pools are arranged along a processing pathway from object to direction selective pools (Miller & Cohen 2001). Furthermore, we assume selective connections along the pathway in accordance with the Hebbian learning paradigm. These connections could be the results of a learning process and reflect the combinations to which the monkey has been exposed to during learning. We assume selective associations between the object selective pool A and the object-and-direction selective pools AL and AR, analogously between B and BL/BR. The same applies between the object-and-direction selective pools and the direction selective pools, i.e., L with AL, BL

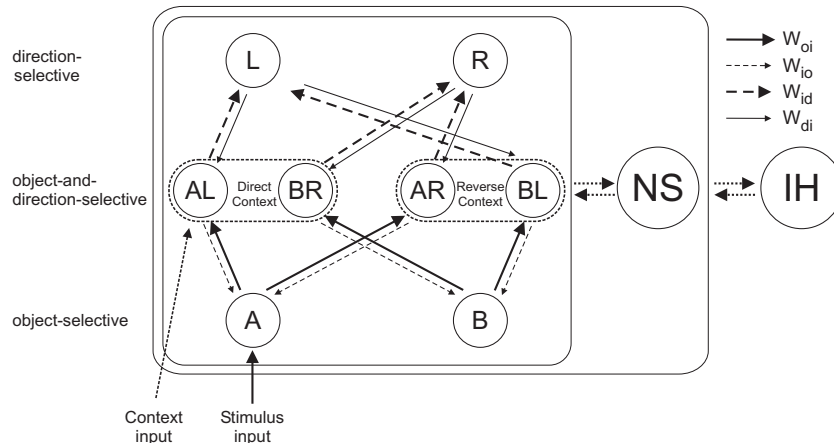


Figure 2.1: The cortical network model. The network is fully connected and the figure highlights the discussed inter-pool connection strengths (see legend). The excitatory neurons are divided into selective ones, which build the neural substrate of the experimental findings, and nonselective ones (NS). The inhibitory neurons are clustered in the pool IH. The network receives external input from 800 neurons which is specifically increased for stimulus and context (direct or reverse). The Supplementary material contains the full connection matrices.

and R with AR, BR. Accordingly, the inter-pool connections yield four different connection parameters between the layers denoted by w_{oi} , w_{io} , w_{id} and w_{di} (o: object, i: intermediate, d: direction, see Figure 2.1). In addition, networks contains a parameter for the intra-pool connection strength w_+ of the selective pools. These are the connections between neurons of the same pool. The intra-pool connections are like the inter-pool connections motivated by Hebbian learning, since neurons which show highly correlated activity are assumed to have a strong connection strength. Accordingly, the intra-pool connections should be stronger than the inter-pool connections.

The other connections between the selective pools are calculated by $w_- = (1 - f_s w_+) / (1 - f_s)$, where $f_s = 0.05$ is the fraction of neurons in a selective pool in relation to all excitatory neurons. The non-specific pool is connected to itself with $w_+ = 1$, since these neurons do not share common selective features. The connections from the non-selective pool to the selective pools w_{nj} are calculated, so that the average input to each selective pool is 1 after having added all additional connections ($w_{nj} = (1 - \sum_{i \neq n} f_i w_{ij}) / f_n$). These calculations regard both the AMPA and NMDA connections. All inhibitory GABA connections are 1 both to themselves and to all other neurons. The complete connection matrices are given in the Supplementary material.

The network receives two inputs: the stimulus and the context input. Both are added to the network by increasing the frequency of the external input of the corresponding pools. This external input originates from 800 external neurons firing at 3Hz. The stimulus input λ_{stim} is added to one of the two object-selective pools (A, B) by increasing the firing rate of the external neurons from $\nu_{ext} = 3Hz$ to $\nu_{stim} = \nu_{ext} + \lambda_{stim}$. The context input acts analogously with $\nu_{context} = \nu_{ext} + \lambda_{context}$ on the two corresponding object-and-direction selective pools. These are AL and BR for the direct context and AR and BL for the reverse context (Figure 2.1). This input influences the competition between the processing pathways and favors the stimulus-response mappings associated with reward. This is consistent with the theories of

cognitive control and biased competition (Miller & Cohen 2001, Rolls & Deco 2002). For a detailed discussion of the nature of context input see Deco & Rolls (2005b).

2.2.4 Analysis

To identify and analyze parameter configurations which show multistable behavior, we used two techniques, namely spiking trial and meanfield simulations. Spiking trial simulations integrate the neural and synaptic dynamics over time. Each simulation started with a pre-period of 500 ms to allow for network stability. According to the experimental paradigm (Asaad et al. 1998), this period was followed by the cue period in which both the cue and context input were added for 500ms. We used $\lambda_{stim} = 0.1Hz$ ($80Hz$ for total external input). The strength of the context input was varied. Afterward, the cue input was turned off and the network was run for another 1000ms which corresponds to the delay period. The context input was present until the end of the simulation. All simulations were performed with the presentation of stimulus A and the direct context. Note that it is sufficient to test one combination of stimulus and context, since the network is symmetric and the other combinations yield analogous results.

Spiking simulations are time consuming and they cannot be done for many parameter configurations of the network. Therefore we used a meanfield approach to calculate the attractor states of the network for the delay period (Brunel & Wang 2001). The attractor states represent the behavior of the network in terms of neural firing rates to which the system would converge in the absence of fluctuations and external changes. We targeted our analysis to identify regions, in which multiple attractors are present and fluctuations could cause the system to switch to one or the other attractor. The initial configuration of the Euler-algorithm of the meanfield calculation was set to 3Hz for all excitatory pools and 9Hz for the inhibitory pool. These values correspond to the approximate values of the spontaneous attractors when the network is not driven by stimulus specific inputs. In addition, an object-selective pool had a higher initial value (40Hz) to account for the excitation of these neurons during the preceding cue period. To get a first approximation for the multistability behavior, we performed several explorations for the same configuration with different starting values for the intermediate pools. Altogether we performed 5 explorations: one with all intermediate pools at 3Hz and four other, each with one intermediate pool at a starting value of 20Hz.

However, multistable attractors do not automatically cause multistable behavior, since the fluctuations must be able to drive the system to the one or the other state. Therefore it is possible that there are multiple attractors at a given parameter configuration, but one attractor is dominating the behavior completely. The percentages to which the system runs into the different attractors have to be assessed with multiple spiking trial simulations.

2.3 Results

On the basis of the meanfield analysis, we identified parameter configurations with three attractors. Figure 2.2 shows histograms of firing rates of the delay period (averaged between 1500ms and 2000ms) of 2000 trials of such a multistable parameter configuration. The averaged firing rates correspond to the outcome of the computation of the network. The three attractors, which have been identified by the meanfield analysis, are indicated by the tickmarks at the bottom of the diagrams. In addition, three peaks can be identified in the histogram, in particular in the sub-plot for pool A: One peak at 0-3Hz, a smaller one at around

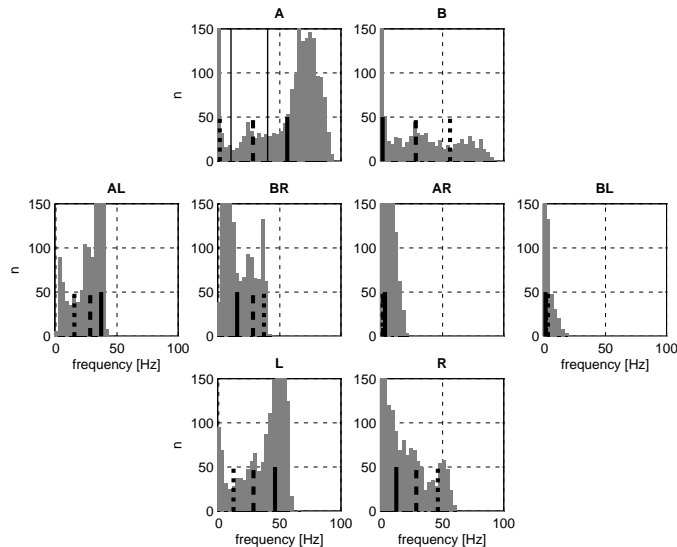


Figure 2.2: Histograms of the firing rates averaged over the second half of the delay period (between 1500ms and 2000ms) of 2000 spiking trial simulations ($w_+ = 3.0$, $\nu_{context} = 3.1Hz$, $w_{oi} = 1.1$, $w_{io} = 2.6$, $w_{ip} = 2.6$, $w_{pi} = 0.9$). The eight sub-plots show the result of the selective pools of the model: upper row - object selective pools, middle row - object-and-direction selective pools, lower row - direction selective pools. The delay period averages were counted in 3Hz bins between 0 and 99Hz. The symbol n represents the number of trials in a particular bin (peaks are higher than shown). The tickmarks correspond to the attractor states identified by meanfield analysis. The dashed lines in sub-graph A show the partitioning used to cluster the trials for the temporal analysis.

25Hz and another significant peak at around 70Hz. The same qualitative characteristic is also visible in pool B and weaker in the pools AL, BR and L. It is obvious, that an average case analysis would not yield interesting results, because we would lose important information about the dynamics. Therefore, we choose to cluster the trials according to the data of pool A in three parts as indicated in Figure 2.2A: Trials in the histogram of below 10Hz ($n=301$), between 10Hz and 40Hz ($n=259$) and above 40Hz ($n=1440$). The latter clearly contains the majority of the trials (about 70%).

Note that figure 2.2 does not reveal a correspondence between the attractors indicated by the tickmarks at the bottom of the figures, because we lost the relationship between the single trials in the histograms due to the binning of the frequencies. With the clustering of trials, we aim at resolving this correspondence. Although the borders of the clustering are merely picked on visual basis, this is just done for pool A and defines the clusters of the other ones. If the clustering chosen on the basis of pool data A indeed corresponds to the attractor states, then it should also yield a clustering of the data of the other pools which corresponds to the attractor states.

Figure 2.3 shows the averages of the clustered trials. As proposed, the three cluster averages show a strong correlation with the identified attractors of the meanfield analysis. Moreover, the frequencies of the averaged trials match approximately the ones predicted by the meanfield calculation. More sophisticated clustering approaches might yield better results. The symmetry in frequency between neighboring pools of the architectural setup is

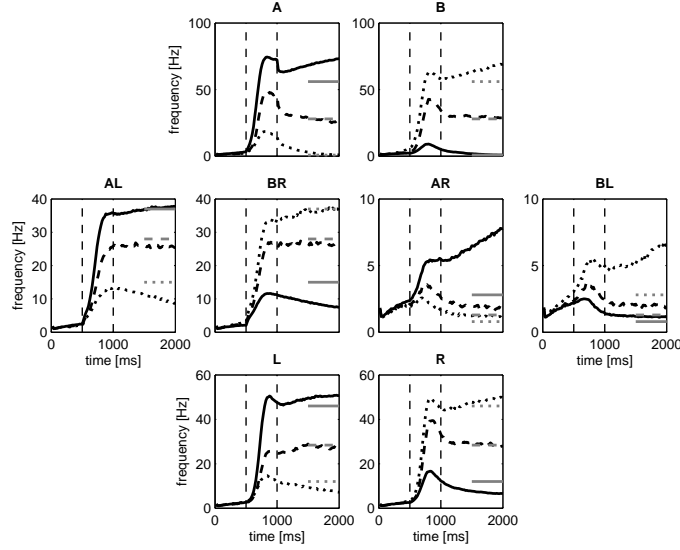


Figure 2.3: Mean spiking frequencies for each pool versus time for parameter configuration ($w_+ = 3.0$, $\nu_{context} = 3.1Hz$, $w_{oi} = 1.1$, $w_{io} = 2.6$, $w_{ip} = 2.6$, $w_{pi} = 0.9$, see Figure 2.2). The trials are clustered in three categories (see vertical dotted lines, Figure 2.2A): Below 10Hz dashed line (n=301), between 10Hz and 40Hz thin line (n=259), and above 40Hz thick line (n=1440). The simulation is conducted consistently with the actual experiment: 500ms pre-trial period (stabilization of the network), 500ms stimulus presentation with $\nu_{stim} = 3.1Hz$ for pool A, followed by a 1000ms delay period. The stimulus input is present from 0 to 500ms, the context input from 500ms until the end of the simulation (direct context: AL, BR). The context input is present from 500ms on (direct context: AL, BR). The horizontal lines between 1500 and 2000ms indicate the meanfield firing rates of the multistable states (see also tickmarks Figure 2.2). The ordinates are not scaled uniformly.

striking, which was not apparent in the trial histogram (Figure 2.2) and also not implicitly contained in our clustering technique. Two patterns are distinguishable: The thick and the dotted lines seem to reflect oppositional attractors with one resulting in a correct response (higher activation in L than in R) and the other one in an incorrect response. Thus, these two attractors represent decision making properties of the network biased by the context input toward the correct response. According to our clustering, 1440 of the 2000 trials belong to the correct mapping, while 301 trials belong to the incorrect one. The remaining 259 trials belong to the second pattern reflected by the thin line. It shows no directional preference on any stage of the processing pathway and resembles the spontaneous state of the network under the influence of the context input. In addition, we observe an almost linear ramping activity in the pools AR and BL, Figure 2.3. This could be an indicator for a linear integrator, but further investigation is needed to discuss this phenomenon.

Beyond the existence of the multistable behavior, we aim at identifying the dependencies which lead to its existence. We perform this analysis using the meanfield technique. First, we discuss the influence of different context inputs on the attractor states and use as reference the parameter configuration of Figures 2.2, 2.3. Figure 2.4 shows the development of the activity difference between the two object-selective pools of the three attractors, which have been shown in the spiking simulation analysis above. Interestingly, the trace of the attractor states




Figure 2.4: Evolution of the three attractors discussed in the spiking simulation analysis with varying context inputs $\nu_{context}$. As before, the remaining values are set to ($w_+ = 3.0$, $w_{oi} = 1.1$, $w_{io} = 2.6$, $w_{ip} = 2.6$, $w_{pi} = 0.9$). Recall that the parameter configuration of the spiking simulations is at $\nu_{context} = 3.1$. The line coding (solid, dashed, dotted) corresponds to the one in figure 3. The frequency values show the difference between the two direction-selective pools (L-R). A difference of 0Hz implies that there is no difference in activation between the two pools and thus no decision between the two directions. This is the case with no context input (0Hz) for all attractors. With increasing context input, both the correct and the incorrect response emerge, at about 3.08Hz complemented by the non responsive attractor. The attractor values are obtained with the meanfield technique using different starting values as initial conditions (for details see Methods section).

does not look as symmetric as suggested by Figure 2.3. Several non-linearities appear when the context input is increased and symmetry is just observed in configurations of low context inputs (below 3.075Hz) and within a small area between 3.1Hz and 3.11Hz. In general, the figure suggests that classical decision making behavior, i.e. two oppositional attractors, is effective at lower context inputs. Within the decision making regime, an increasing input leads to a higher difference between the two attractors. Too strong context input leads to nonlinearities and finally to a non-responsive system (here above 3.14Hz).

Note that the attractor analysis does not reveal to which extend fluctuations can cause the system to run into the corresponding attractors. This is reflected in the percentages to which the system ends up in the one or the other attractor and can just be assessed using multiple spiking trial simulations and a subsequent clustering analysis. Nevertheless, the existence of multiple attractors is essential for decision making properties.

Furthermore, we analyzed the dependence of the behavior on the connections strength between the three layers. We asked whether the system could perform the mapping task with all inter-pool connection strengths being equal ($w_{oi}=w_{io}=w_{id}=w_{di}$). Using this constraint, we were not able to find any parameter configuration with responsive properties. Therefore there has to be some kind of structuring in the connection strengths between the layers. For example, biasing the connection strength to stronger in direction toward the direction-

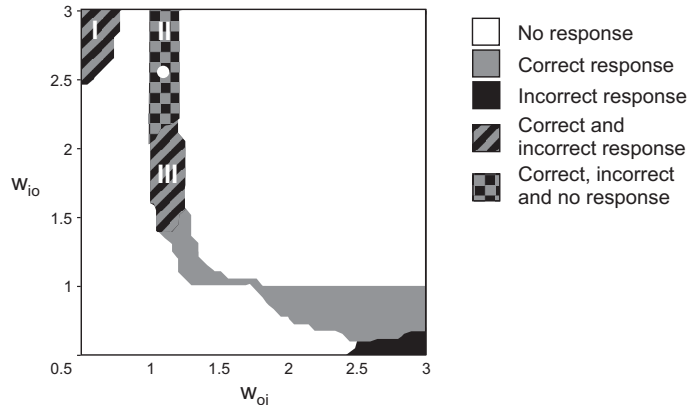


Figure 2.5: Categorization for response direction of a parameter scan. The results of the meanfield calculation are categorized using the difference in spiking frequency between the two direction-selective pools. No response corresponds to a difference below 30% between the two direction-selective pools, while correct(incorrect) response corresponds to a 30% higher(lower) frequency in the L pool than in the R pool. The connection strengths between the object-selective and intermediate pools are varied (w_{oi} , w_{io}), while the remaining parameters are fixed at ($w_+ = 3.0$, $\nu_{context} = 3.1Hz$, $w_{ip} = 2.6$, $w_{pi} = 0.9$). Three multistable region are identified: Regions I and III show two different attractors, one of the correct and one of the incorrect response. Region II shows additionally a third attractor yielding no responsive behavior. The parameter configuration discussed earlier (Figures 2.2, 2.3) is marked by the white dot. The sampling of the parameter space is $\Delta w_{oi} = 0.05$ and $\Delta w_{io} = 0.05$.

selective neurons ($w_{oi} > w_{io}$ and $w_{id} > w_{di}$) allows for parameter configurations which show the correct response upon stimulus presentation (Deco & Rolls 2003, 2005b).

However, the parameter configurations used by Deco & Rolls (2003, 2005b) do not show evidence of multistability and are not capable of contributing to the response space exploration. Figure 2.5 shows the response as functions of the connection strengths between the object-selective and the intermediate layer. The response direction is assessed by calculating the relative difference between the two direction-selective pools and using a threshold value of 30%. Other threshold values such as 20% or 40% yield similar results. Several aspects are noteworthy: First, the regions in which the network responds to the task seem to be well bounded and our data suggest that fine-tuning is needed. Second, multistable regions are present in an area with stronger connectivity from the intermediate to the object-selective layer. This balance seems to be necessary to achieve multistable behavior and thus response space exploration properties. The parameter configuration used by Deco & Rolls (2003, 2005b) is located in the single attractor region of the correct response and does not show multistable behavior.

2.4 Discussion

We have presented a detailed analysis of a generic neurodynamical model for an arbitrary visuomotor task. As shown by Deco & Rolls (2005b), this model is not only applicable to stimulus-response associations (Asaad et al. 1998), but also to stimulus-reward associations

(Thorpe et al. 1983) and probably to a variety of other experimental paradigms which require a mapping between two aspects upon previously learned rules. Working memory is important to bridge the period between cue presentation and response in delayed response tasks. We extended the working memory model of Brunel & Wang (2001) through a structured architecture which is motivated by the neurophysiological data (Asaad et al. 1998) and processing principles, the latter being cognitive control and biased competition (Miller & Cohen 2001, Rolls & Deco 2002). We show that multiple attractor states exist in this extended working memory model and that fluctuations are sufficient to drive the system to these attractors. This is the main characteristic of dynamical systems exhibiting decision making properties and is essential for trial-and-error learning of stimulus-response tasks.

The research described here adds in crucial points to the work of Deco & Rolls (2003) and Deco & Rolls (2005*b*). First, on the conceptual level we extended the conditional visuomotor mapping model by studying fluctuations and multistability. Second, we analyzed the neurodynamical behavior in more detail: Beyond the single spiking trial analysis used by Deco & Rolls (2003) and Deco & Rolls (2005*b*), we introduced extensive multiple spiking simulations with subsequent clustering of the trials. This allows for an assessment of the influence of fluctuations and the identification of multiple attractor states. However, these simulations are computationally expensive so that only a limited number of parameter configurations can be assessed in detail. Thus, we complemented the spiking simulation analysis by a meanfield analysis to study the influence of the parameters on the neurodynamical system thoroughly. In particular, we used this approach to identify parameter values for which the system shows fluctuation-dependent multistable behavior. Hence, we go significantly beyond the proof of principle presented by Deco & Rolls (2003) and Deco & Rolls (2005*b*). In particular, we showed that asymmetries of the inter-pool connection strengths along the processing pathway are essential for the computation of a response by the network. In future work, further constrains of the parameter space could be derived from additional biological data and theoretical considerations. In addition, improved theoretical techniques could promote the analysis of the networks dynamics.

Earlier models of visuomotor associations addressed learning using different degrees of biological realism. Based on experimental data by Mitz et al. (1991), Fagg & Arbib (1992) proposed a model of the premotor cortex consisting of motor selection columns. While the hierarchical structure of these columns is similar to the work presented here, the model of Fagg & Arbib (1992) is less refined at the neural level and more of an algorithmic nature. Dominey et al. (1995) presented a multiregional connectionist model using a mean approach on the neural level for learning of visuomotor associations by corticostriatal plasticity. It would be interesting to analyze such a system using spiking neurons. The work presented here is an important step in this direction, since the decision making properties discovered in our model are required for the exploration and learning of associations. Chialvo & Bak (1999) and Bak & Chialvo (2001) focus on learning mechanisms themselves and put forward the ideas of extremal dynamics and learning by mistakes. They highlight interesting mechanistic aspects of learning dynamics which should be considered in future work.

A basic feature of our conditional visuomotor association model is the hierarchical processing pathway. We emphasize that our model is of an abstract neurodynamical nature. While our model reflects the data of neurophysiological recordings, we do not make any specific claims about how the structure could be implemented in the brain. However, there are anatomical and physiological reasons, which suggest such an organization from object-selective to response-selective neurons. First, the cortical organization in several inter-connected layers

suggests such an organization. For example, in the visual system feed-forward connections originate mainly from cells in layer III and terminate mainly in layer IV, whereas feedback projection between the stages of processing originate from cells in layer V and VI and terminate in layers I/II and VI (Rockland & Pandya 1979). Hence, the connectivity between cortical layers could result in a hierarchical organization with asymmetric connections. The model of motor selection columns by Fagg & Arbib (1992) is related to this principle in the sense of computations in the cortex within structurally organized cortical columns.

Second, the neurophysiological properties of different frontal areas can be interpreted in terms of a processing pathway. The ventrolateral prefrontal cortex receives direct inputs from temporal visual areas and might select the information which is relevant for the task, i.e. the stimulus identity (see Rushworth & Owen (1998), Passingham et al. (2000)). Hence, the object-selective layer could model functions of the ventrolateral prefrontal cortex. The dorsolateral prefrontal cortex area contributes to the manipulation of information (Rushworth & Owen 1998). The stimulus information and the learned rules are combined to the corresponding association. Consistently, the dorsolateral prefrontal area contains neurons, which are not only selective to the stimulus and the response direction but also to combinations of both (Boussaoud & Kermadi 1997, Asaad et al. 1998). This object-and-direction selected activity is shown by the intermediate layer of the model. Finally, the dorsal premotor cortex contains a high proportion of direction-selective neurons (Mitz et al. 1991, Boussaoud & Kermadi 1997) and is also known to be crucial for arbitrary visuomotor tasks (Petrides 1982, Halsband & Passingham 1982). Therefore it might take part in the final stages of processing and is described best by the direction-selective layer in our model. Taken together, a pathway for the processing of visuomotor associations could be implemented in the frontal cortex from the ventrolateral area to the premotor area. Moreover, based on imaging data Koechlin et al. (2003) hypothesize a cascade organization of frontal regions reflecting different levels of processing depth, which further supports the assumption of a hierarchical structuring within the frontal cortex.

In contrast to lesions of the ventrolateral prefrontal cortex (Bussey et al. 2001), lesions of the dorsolateral part cause no or only a mild deficit in visuomotor mapping tasks (Petrides 1982, Gaffan & Harrison 1989). In addition, imaging studies have indicated no major contribution of this area (Deiber et al. 1997, Toni, Rushworth & Passingham 2001), at least in simple versions of this task (Boettiger & D'Esposito 2005). This suggests that the pathway including the dorsolateral prefrontal cortex might play a minor role (Hadj-Bouziane et al. 2003). Correspondingly, as an alternative to the cortical pathway a pathway involving the corticostriatal system has been proposed by Houk & Wise (1995), Toni & Passingham (1999), Hadj-Bouziane et al. (2003). This corticostriatal pathway involves the following brain regions: (i) the ventrolateral prefrontal cortex might hold the presented stimuli in working memory and select the relevant information, (ii) the striatum integrates sensory, motor and reward information, and (iii) the premotor cortex selects the motor response. The integrative property of the striatum, which is implemented by the intermediate layer of the model, is the key aspect of the context detection (Houk & Wise 1995, Dominey & Boussaoud 1997, Djurfeldt et al. 2001, Pasupathy & Miller 2005). The reward information might originate from the midbrain dopamine system, which sends projections to medium spiny neurons in the striatum (Schultz 2000). Furthermore, Toni et al. (2002) showed by structural equation modeling on imaging data that the effective connectivity in fronto-striatal circuits increases during learning. This also supports a processing pathway, which includes the striatum. Like the striatum, the orbitofrontal cortex contains cells, which code reward information (Watanabe 1996, Rolls 2004)

and might also be important for the processing of arbitrary visuomotor associations. The involvement of the orbitofrontal cortex is addressed by Deco & Rolls (2005*b*).

We emphasize that we do not assign properties to the specific layers in our model explicitly. Their properties emerge due to the various influences which act on the layers. For example, the selectivity of the neurons in the intermediate layer does not arise because of the network structure itself. Rather, the connectivity, the activations in neighboring layers, and the context input must act together. This yields an alternative interpretation of the neural activity in brain regions. The neural activity in a brain region does not arise on its own but by interaction with neighboring regions. Thus questions concerning the functions of brain regions might as well be formulated in terms of interactions of that region with its neighbors. This view could shed light on apparently contrary functions of a given brain region which might have a common theme in terms of interregional interactions.

The context input represents an abstract rule because it gives input to those two associations connected with the given rule. For example, the direct rule activates the pools AL and BR. This corresponds to the assumption that rules instead of single associations are learned: After a rule switch, the monkeys detect a change by a wrong answer to one stimulus. Consequently, the next answer to the other stimulus should already be correct or at least have a higher probability of being correct. Optimally the monkey could perform a rule switch after a single error. The stimulus-reward association experiment by Thorpe et al. (1983) shows evidence of such behavior. Because synaptic plasticity would take a few trials to adapt to the new contingencies, neurodynamical mechanisms are assumed to account for this behavioral adaptation (Deco & Rolls 2005*b*). However, the conditional stimulus-response association experiment (Asaad et al. 1998) does not show evidence for one-trial reversal, although the monkeys detect changes rapidly. Brasted et al. (2005) recently targeted this question of one-trial learning in stimulus-response associations. They found that prior error trials impair one-trial learning. This is important for experiments by Asaad et al. (1998) and Pasupathy & Miller (2005), because the monkeys could only detect a change in the rule by making errors. Accordingly, this made it more difficult for the monkeys to learn new associations quickly. However, this does not resolve the issue of why the monkeys needed several dozens of trials to reach a plateau in performance. Clearly more research is needed to identify the crucial factors underlying one-trial reversal learning.

The systems responsible for familiar and novel associations have different properties: For familiar associations errors are rare, and therefore the system is not multistable, i.e., it almost always ends up in the attractor state of the correct association. We identified such a single attractor dynamics in the correct response region in Figure 2.5. Experimental evidence also supports that familiar association might be executed by different neurodynamical mechanisms, because different brain regions and pathways are involved. Nixon et al. (2004) interrupted the connection between the basal ganglia and the central premotor cortex using a cross lesion method whereby a unilateral lesion of the internal pallidum in one hemisphere was combined with a premotor lesion in the opposite hemisphere. The learning of novel associations appeared relatively unaffected by the disconnection whereas the monkeys were impaired in re-learning pre-operatively learned associations and made three-times as many errors once they succeeded. These results suggest that the connections between the basal ganglia and the premotor cortex are crucial for the execution of familiar visuomotor associations but not for learning of novel associations. Thus the transition from a novel to a familiar association might go along with a shift to another pathway that show multi- and uni-stable characteristics, respectively. Moreover, the experiment by Bussey et al. (2001) has revealed pathways

with different learning rates. After bilateral ablation of the orbital and prefrontal cortex, the learning of novel associations was no longer possible within one session but rather required several daily sessions. Thus fast learning was impaired whereas a slow learning was spared. This is compatible with the cascade learning model proposed by Fusi et al. (2005) in the sense of learning on different time scales. In our model, we envision fast changing context inputs (Deco & Rolls 2005b) and slow Hebbian network structuring for novel and familiar stimuli, respectively.

We derive predictions from our model. First, the existence of a system which encodes abstract rules could be indirectly tested. Figure 2.3 shows the spiking activity during trials in which stimulus A and the direct rule was presented. The activities associated with the active rule (AL and BR) are always higher than the ones of the inactive rule (AR and BL). In particular, note that the pool BR had a higher activation than AR, although the pool AR received direct input from the object selective pool A. This is possible because not only AL but also BR received an input associated with the active context and thereby establishing an abstract rule. This abstract rule reflects the direct context. Hence, we predict that if the system uses a rule mechanism then the activity of the object-and-direction selective neurons belonging to the active rule should be higher than the activity of neurons belonging to the other rule.

Second, error trials represent a valuable source of information, because they allow a different view on the system. If error trials are implemented in a neurodynamical system such as the one proposed here, then the monkey should make exploratory errors. This is best illustrated in Figure 2.3. In most trials (70%), the activity is high in the pools A, AL and L (solid lines), correctly responding to stimulus A and the direct rule. However, in some trials (15%), the network responds to the presentation of stimulus A and the direct context with a high activity in the pools B, BR and R (dotted lines, Figure 2.3). Thus, the system responds with an incorrect answer although it has been exposed to the correct input. It is already known that the activity of direction-selective neurons in the prefrontal cortex indicates the impending response regardless of whether it is correct (Asaad et al. 1998). We predict on the basis of our model that the error activity should also be reflected in the activity of the object-and-direction selective neurons and maybe even in the object selective neurons. Specifically, when stimulus A is presented during the direct context and the monkey makes an error (response R), we predict high activity for neurons selective to the association BR and to the object B such as picture in Figure 2.3 (dotted lines). This type of error is not due to processing errors in parts of the system but due to properties of the whole stimulus-response system. In addition, our simulations suggest a second cause of errors, namely the indecisive state (dashed lines, Figure 2.3). Although the stimulus A and the direct context are influencing the network, the activities of the pools L and R remain at the same level, and analogously the activities of the pools AL, BR and A, B. Consequently, the monkey has to guess. Both patterns, the exploratory and the guessing errors, could support the response space exploration during the learning phase. However, response space exploration is only one of different possible causes of error trials.

In addition, we emphasize that the response space exploration could also be targeted experimentally. We envision experimental setups in which monkeys are forced to make exploratory errors. For example, monkeys which are well trained on learning and performing a visuomotor association task could be confronted with rapidly changing associations. They should constantly try to identify the associations and explore the behavioral options. The neural activity during the exploration could reveal the mechanisms, which underlie the re-

sponse space exploration and guide the learning process.

Acknowledgments

We thank Dr. Ralph G. Andrzejak and Dr. Anders Ledberg for discussions on earlier versions of this paper.

Abbreviations

AMPA, alpha-amino-3-hydroxy-5-methyl-4-isoxazolepropionic acid; NMDA, N-methyl-D-aspartate; GABA, gamma-aminobutyric acid.

Chapter 3

Neurodynamics of the prefrontal cortex

The work presented in this chapter is accepted for publication in the Journal of Cognitive Neuroscience (Loh et al. 2008).

Abstract

The prefrontal cortex is believed to be important for cognitive control, working memory and learning. It is known that it plays an important role in the learning and execution of conditional visuomotor associations, a cognitive task in which stimuli have to be associated with actions by trial-and-error learning. In our modeling study, we sought to integrate several hypotheses on the prefrontal cortex function in a computational model and compare the results to experimental data. We constructed a module of prefrontal cortex neurons which were exposed to two different inputs which we envision to originate from the inferotemporal cortex and the basal ganglia. We found that working memory properties do not describe the dominant dynamics in the prefrontal cortex, but the activation seems to be transient, probably progressing along a pathway from sensory to motor areas. During the presentation of the cue, the dynamics of the prefrontal cortex is bistable yielding a distinct activation for correct and error trials. We find that a linear change in network parameters relates to the changes in neural activity in consecutive correct trials during learning, which is an important evidence about the underlying learning mechanisms.

3.1 Introduction

Humans and animals can learn to associate stimuli with arbitrary responses, recall these associations and adapt them when changes in the behavioural context occur. This type of learning is called conditional visuomotor learning (Passingham 1993). Several brain structures take part in this task (for reviews see Murray et al. (2000), Passingham et al. (2000), Hadj-Bouziane et al. (2003)). These are the prefrontal cortex, the premotor cortex, the hippocampus, and the basal ganglia. In our modelling study, we focus on the prefrontal cortex.

The prefrontal cortex (PFC) is important for the processing of conditional visuo-motor mappings. Frontal lesions in humans cause deficits (Petrides 1985, 1990, 1997). The ventrolateral prefrontal cortex is crucial for the fast learning of novel associations (Murray et al. 2000, Wang et al. 2000, Bussey et al. 2001). The interaction with the inferotemporal cortex (IT) also contributes to the processing, as a disconnection using cross-lesion procedures of IT and PFC (Gaffan & Harrison 1988, Parker & Gaffan 1998, Bussey et al. 2002) and transection of the uncinate fascicle (Eacott & Gaffan 1992) causes deficits in learning. However, lesions of the dorsolateral prefrontal cortex cause no or mild impairments (Petrides 1982, Gaffan & Harrison 1989, Wang et al. 2000). This is consistent with human imaging studies which show an involvement of the ventrolateral PFC but not of the dorsolateral PFC (Toni & Passingham 1999, Toni, Ramnani, Josephs, Ashburner & Passingham 2001), at least in simple versions of this task (Boettiger & D'Esposito 2005). Electrophysiological recordings have identified neurons which were selective to the stimulus identity, the response and to combinations of both, i.e., the associations (Asaad et al. 1998). Simultaneous recordings in the caudate nucleus and the PFC have shown that the neural activity which reflected the response appeared earlier in the caudate nucleus than in the PFC (Pasupathy & Miller 2005). In addition, the activity in the caudate nucleus changed more rapidly during learning compared to the slower changes in the PFC.

Based on the literature outlined above, we sought to integrate and discuss the following findings and hypotheses about the PFC activity in a computational model: First, the model neurons should be selective to combinations of stimuli and responses, i.e. single associations in the experiment (Asaad et al. 1998). Second, stimulus information reaches the PFC from IT. The connection of the two brain regions is crucial for learning visuomotor associations (Gaffan & Harrison 1988, Parker & Gaffan 1998, Bussey et al. 2002). Third, information about movement direction influences the PFC from the basal ganglia. The basal ganglia are connected to the prefrontal cortex in a cortico-basal ganglionic loop (Houk & Wise 1995, Graybiel 1998) and have a dominant role in the learning of new associations (Houk & Wise 1995, Bar-Gad et al. 2003, Pasupathy & Miller 2005). Fourth, we assume that the directional information is influencing the prefrontal cortex throughout the delay period, as it has been hypothesized that the cortical-thalamic loop maintains this information (Houk & Wise 1995). Fifth, one of the most prominent properties of the PFC is working memory related neural activity (Fuster & Alexander 1971, Kubota & Niki 1971). We model the PFC module using a recurrent network model of working memory to maintain the extracted association during the delay period.

Our goal is to integrate these aspects in a computational model and draw conclusions about the processes underlying arbitrary visuomotor mappings by comparing the results to experimental data.

3.2 Methods

3.2.1 Experimental Background

We base our modelling approach on experiments featuring electrophysiological recordings in the PFC (Asaad et al. 1998, Pasupathy & Miller 2005). In both experiments, monkeys were presented a cue (A or B) for 500ms and after a subsequent delay period of 1000ms they had to make a leftward or rightward saccadic eye movement (L or R). Selected combinations of stimuli and responses were rewarded and thereby reinforced. To enforce learning, the associations between stimuli and responses were reversed after the monkeys had reached high performance. For example, if the stimuli A and B were associated with a leftward and rightward eye movement (L and R) respectively, then the association of A with a rightward (R) and B with a leftward (L) eye movement yielded reward after the reversal. Because the monkeys did not switch instantaneously, it was possible to investigate the neural activity during learning.

3.2.2 Neurodynamical model

We use a standard recurrent network model as the basis for working memory in the prefrontal cortex (Brunel & Wang 2001). It consists of 1000 integrate-and-fire neurons which are all-to-all connected. The integrate-and-fire neurons are described by

$$C_m \frac{dV(t)}{dt} = -g_m(V(t) - V_L) - I_{syn}(t), \quad (3.1)$$

where $V(t)$ is the membrane potential, C_m the membrane capacitance, g_m the leak conductance and V_L the resting potential. 800 of these neurons are excitatory and 200 are inhibitory. The synaptic input I_{syn} is made up of four parts. An external excitatory input via AMPA-type synapses to the network model and recurrent input from the other neurons of the network. The latter one consists of AMPA, NMDA and GABA currents. Thus, I_{syn} reads

$$I_{syn}(t) = I_{AMPA,ext}(t) + I_{AMPA,rec}(t) + I_{NMDA,rec}(t) + I_{GABA}(t), \quad (3.2)$$

The asynchronous external input $I_{AMPA,ext}(t)$ can be viewed as originating from 800 external neurons firing at an average rate of $\nu_{ext} = 3Hz$ per neuron, and thus has a rate of 2.4KHz in total. The recurrent input currents are summed over all neurons with weights depending on the connectivity. The synaptic dynamics is described by an exponential decay upon the arrival of a spike for AMPA and GABA currents and alpha-function including a rise time and an extra voltage dependence for the NMDA current. The detailed mathematical description and parameters are provided in the Supplementary Material. The parameters of the integrate-and-fire neurons and the synaptic channels for AMPA, NMDA and GABA are chosen according to biological data. We gain a dynamical system which has biologically realistic properties and can have several distinct attractor states. These are envisioned to correspond to items stored in working memory.

We incorporate the five cornerstones of our modeling as mentioned in the introduction as follows: First, the model neurons should show non-linear object-and-direction selective activity. In the experimental paradigm of Asaad et al. (1998) and Pasupathy & Miller (2005), there are four associations between the two stimuli A,B and two responses L, R: AL, BR, AR,

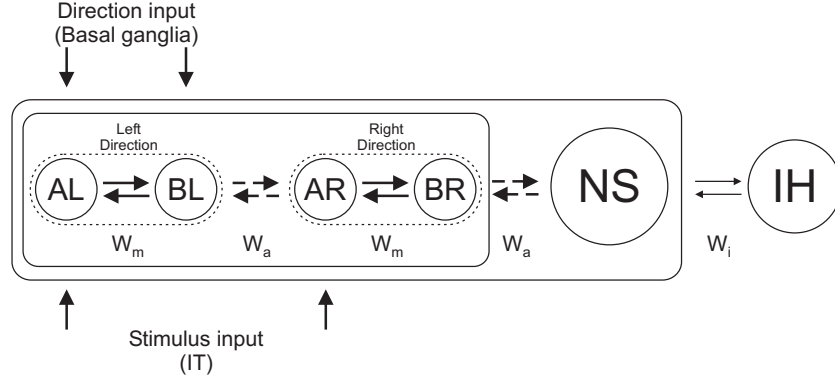


Figure 3.1: The cortical network model. The network is fully connected and the connection parameters between parts of the network are shown. The excitatory part of the network consists of selective pools (AL, BL, AR, BR) and a non-selective one (NS). The inhibitory neurons are clustered in the pool IH. The each neuron receives external input of Poisson spike trains firing at 2.4KHz, which can be views as originating from 800 neurons firing at 3Hz. This external input is selectively increased for the stimulus and response-direction input for the corresponding pools. The input for the stimulus A and the direct context is indicated as used in the simulations.

and BL. Accordingly, our model contains four populations of 100 neurons each (Figure 3.1). The rest of the excitatory neurons comprise the non-specific pool (NS). These neurons show spontaneous firing rates and introduce noise in the system, which aids in generating Poisson like firing patterns throughout the system.

Second, stimulus information reaches the PFC from IT. As we are interested in the dynamics of the PFC, we model this input by an external influence to the model. The external firing rate is increased from $\nu_{ext} = 3Hz$ to $\nu_{stim} = \nu_{ext} + \lambda_{stim}$ for the two pools associated with a stimulus, stimulus A to AL and AR, stimulus B to BL and BR. We use $\lambda_{stim} = 0.1Hz$ throughout the simulations.

Third, the basal ganglia map stimuli onto actions and their output influences the activity of the PFC. We model the direction input which codes for response direction left (L) or right (R) by increasing the external input to (AL, BL) or (AR, BR) respectively ($\nu_{dir} = \nu_{ext} + \lambda_{dir}$). In contrast to the stimulus input, we do not assume λ_{dir} to be constant during learning, since the associations change during the blocks. That is we increase the external input λ_{dir} from $0.0Hz$ to $0.1Hz$ to simulate the changes which occur during learning.

Fourth, we assume that the input coding the response-direction is influencing the pre-frontal cortex throughout the delay period, as it has been hypothesized that the cortical-thalamic loop maintains this information (Houk & Wise 1995). This is implemented by applying the response-direction input from the cue period on throughout the delay period (0-1500ms), whereas the stimulus input is applied only in the cue period (0-500ms).

Fifth, we hypothesize that the PFC module holds the extracted association in working memory. Thus we use a neural network model constructed to show working memory properties (Brunel & Wang 2001) and target our analysis to parameters which show delay period activity. In our model, the stimulus-response associations are held in working memory (Asaad et al. 1998).

The connection weights between the neurons of the same population are called intra-pool connection strength w_+ (Figure 3.1). In addition, we group the weights of the inter-pool connections into two parameters. The connection weights between neurons coding the same response-direction, i.e. between the pools AL-BL and AR-BR, are denoted w_m . The second parameter w_a applies to all other connections between excitatory pools including the non-specific pool. The connection strength w_a is calculated using the other two connection parameters w_+ and w_m , so that the average connection strength input to a neuron equals to 1 ($f w_+ + f w_m + (1 - 2f) w_a = 1$, where $f = 0.1$ is the fraction of the number of neurons in a selective pool in respect to all excitatory neurons). The excitatory connections regard both the AMPA and NMDA synaptic currents. All inhibitory GABA connections are 1 both to themselves and to all other neurons. The complete connection matrices are pictured in the Supplementary Material.

3.2.3 Analysis

To analyze the network, we use two kinds of techniques which complement each other: Spiking and meanfield simulations. Spiking simulations calculate the dynamics of every neuron in time and yield a detailed temporal evolution of the system including fluctuations. We start the spiking simulations with a pre-cue period of 500ms. In this pre-cue period neither the stimulus nor the response-direction input is applied to the PFC module. It represents the period before the stimulus onset in which the monkey does not have any information about the stimulus. Thereafter we activate the stimulus input for 500ms by increasing the external input ν_{ext} by 0.1 Hz (80Hz in total). The response-direction input ν_{dir} is also present from the beginning of the cue period but lasts throughout the delay period until 1500ms after stimulus onset with varying strength (0 - 0.1Hz).

We use the meanfield formulation (Brunel & Wang 2001) to assess the dependencies of the network parameters and identify the region of interest. The meanfield approach calculates the attractor states of the network. The attractor states represent the behavior of the network in terms of neural firing rates to which the system would converge to in the absence of fluctuations and external changes. We focus our meanfield analysis on the delay period of the system. During the delay period (500-1500ms) the response-direction input is present, so we apply this input ν_{dir} in the mean-field simulation. The stimulus input has been active during the cue period and has increased the firing rates of the pools (AL, AR) or (BL, BR). Thus we start the search for the attractor state at a higher initial frequency (40Hz) for these pools. The other pools have an initial frequency of 3Hz corresponding to the spontaneous firing.

We calculated the direction selectivity using a linear regression model. As dependent variable Y , we used the data of the pools AL and AR as these have the same stimulus selectivity. The data consisted of the firing rates (averages of the spiking data in 20ms bins for each trial) of 100 simulated trials for each parameter configuration and point in time. The independent variable X consisted of a vector denoting the response-direction (e.g. 1 for L and -1 for R). In Figure 3.2, the data was additionally categorized in correct and error trials to allow a separate analysis of these trial types.

$$Y = Xb + e \tag{3.3}$$

We calculated the least-squares estimate of the linear regression coefficients,

$$b = (X^T X)^{-1} X^T Y \tag{3.4}$$

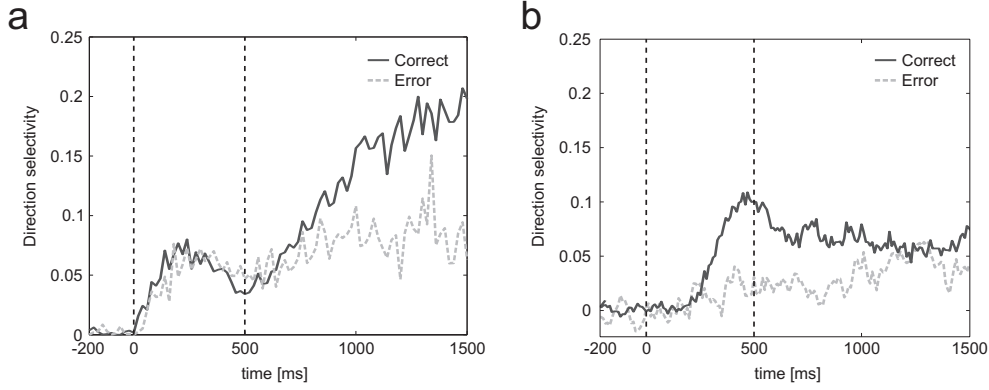


Figure 3.2: Direction selectivity of correct and error trials as a function of time. a. Modeling data using a threshold criterion to categorize the trials: A trial is identified as correct if the average activity of the last 500ms of the delay period (1000-1500ms) is above 10Hz for pool AL and below 10Hz for all other pools. The values of the network parameters are fixed at $w_+ = 1.8$, $w_m = 0.7$, $w_a = 0.9375$, and $\lambda_{dir} = 0.06Hz$ which yields a performance of 69%. b. Experimental data of PFC direction selectivity (adapted from Pasupathy & Miller (2005)).

and used this estimate to calculate the proportion of the explainable variance by the sum of squares:

$$PEV_{dir} = \frac{(Xb)^T Xb}{Y^T Y} \quad (3.5)$$

The proportion of the explainable variance PEV_{dir} represents the direction selectivity as also used in the experimental study (Pasupathy & Miller 2005).

3.3 Results

We used the meanfield technique to identify parameter regions which fit our hypothesis. We looked for parameter configurations which correspond to the properties of object-and-direction selective neurons. For example, when stimulus A is associated with a leftward eye movement, the neurons in the pool AL should respond with a higher firing rate, but not the neurons of all other combinations (AR, BL, BR). We used a threshold of 10Hz to determine high or low activation. Note that the network is symmetric and therefore it is sufficient to check just one combination of stimulus and response-direction. We used stimulus A and a leftward response-direction for all simulations.

Figure 3.3 shows six slices of the three-dimensional parameter space (w_+ , λ_{dir} , w_m) along the w_m -axis. The parameter w_+ connects neurons which have the same selectivity and are thus highly correlated. The parameter w_m represents the connection weight between populations of neurons, which share the same directional selectivity. The populations which are connected by w_a are of opposite directional selectivity (see Figure 3.1). Within the gray region, the neurons show the property of object-and-direction selective neurons, i.e., high activation (>10Hz) in one pool and low activation (<10Hz) in all other pools. The interdependency on the fourth parameter w_a , which is calculated using w_+ and w_m , shows that in all parameter configurations, in which the network neurons are "object-and-direction selective", w_+ is larger than w_a , and w_a is larger than w_m ($w_+ > w_a > w_m$). This order of the connection strengths




Figure 3.3: Meanfield analysis for object-and-direction selective patterns. The six panels show this property for slices of the parameter space with different values of w_m while varying w_+ and λ_{dir} . The meanfield analysis was performed for the parameter space w_+ ranging from 1.0 to 3.0, w_m from 0.5 to 1.0 and the response-direction input from 0.0 to 0.4 Hz. The parameter w_a was calculated as described above. We show just the part in which firing rates of the pools correspond to object-and-direction selective neurons. The dashed line corresponds to the parameters which are discussed using spiking simulations ($w_+ = 1.8$, $w_m = 0.7$, λ_{dir} from 0Hz to 0.1Hz).

seems to be necessary to achieve the object-and-direction selective activity as demanded in our model description.

To evaluate the influence of the response-direction input, we choose a line in the parameter space along the direction input axis by fixing the two other parameters $w_+ = 1.8$ and $w_m = 0.7$ (see dashed line in Figure 3.3). We analyzed these parameters using multiple spiking trial simulations (n=100 per parameter configuration).

Figure 3.4 shows an example of two simulated trials using the same parameter configuration ($w_+ = 1.8$, $w_m = 0.7$, and $\lambda_{dir} = 0.05Hz$). Due to noise in the system, trials can end up in different attractor states. One trial shows a high activation in AL and the other one a low activation. Using a threshold criteria at 10Hz to the end of the delay period, the two trials can be categorized as hit and miss, respectively.

First, we discuss the modeling data during the cue-period. To facilitate the comparison to the experimental data, we also use a window of 300ms, here between 100 and 400ms after cue onset which we call analogously to the experimental work "peri-cue period". Pasupathy & Miller (2005) used a different offset, i.e. the period between 300ms and 600ms after cue onset, due to the time delay until the visual information reaches the PFC. The Figure 3.5a shows the average activity (and standard errors) during the peri-cue period as a function of the response-direction input of the pools AL and AR. These two pools represent the two conflicting directions for the stimulus A. The difference in firing rate increases with stronger response-direction input. Moreover, the activity for AL is steadily increasing whereas the activity of AR stays at the same level. An analogous observation is made in the experimental data during the learning phase in the experiment in consecutive correct trials (see Figure

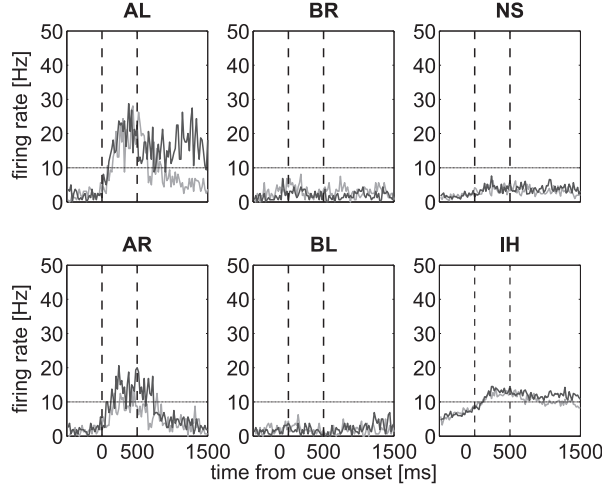


Figure 3.4: Two selected trials with same parameter configuration ($w_+ = 1.8$, $w_m = 0.7$, $\lambda_{dir} = 0.05Hz$). The six graphs show the activity of all pools in time. The 10Hz threshold is indicated and used to assess if a trial would identify an association correctly.

3.5b).

To further quantify the activity in the peri-cue period, we calculated the average direction selectivity. Figure 3.6a shows that the dependency of the average direction selectivity on the response-direction input is approximately linear. This is consistent with the linear correlation between direction selectivity and consecutive correct trials as observed in the experimental data (Figure 3.6b). Note that both in Figure 3.5 and Figure 3.6, we already implicitly related the strength of the response-direction input to the correct trials in the experiment. The changes occurring during the cue period in consecutive correct trials in the experimental data correspond well to a linearly increasing response-direction input in the model.

The experimental data show that the monkeys behavior correlates much better with the slower gradual changes in the neural activity in the PFC than with the almost bistable changes in the basal ganglia. Pasupathy & Miller (2005) concluded that the behavior is related more closely to the PFC activity. Motivated by this finding, we introduced a behavioral measure in our model: An association is correctly identified, if the mean of the last 500ms of the simulation (1000-1500ms after cue onset) of pool AL is above 10Hz and of the other pools below 10Hz, analogously to the meanfield analysis. Figure 3.7 shows the evolution of the percent correctly identified associations with increasing response-direction input. With no direction input (0Hz), the association cannot be identified, and all pools stay at low activation levels. An increasing direction input leads to a higher percentage of trials in which the AL pool is activated until the percentage reaches almost 100% at around 0.1Hz. Note that when the system has not identified an association the monkey could still guess. Thus the performance could be at 50% (choice between L,R) when the identification level is at 0%.

We use this measure to compare the direction selectivity between correctly and incorrectly identified associations. Figure 3.2a shows the direction selectivity for a single direction bias (λ_{dir}), separated in correct and error trials. The progression of both cases almost identical until 1000ms after cue onset. After a rise due to the stimulus and response related input, the direction selectivity drops shortly around the end of the cue period (500ms) before it increases

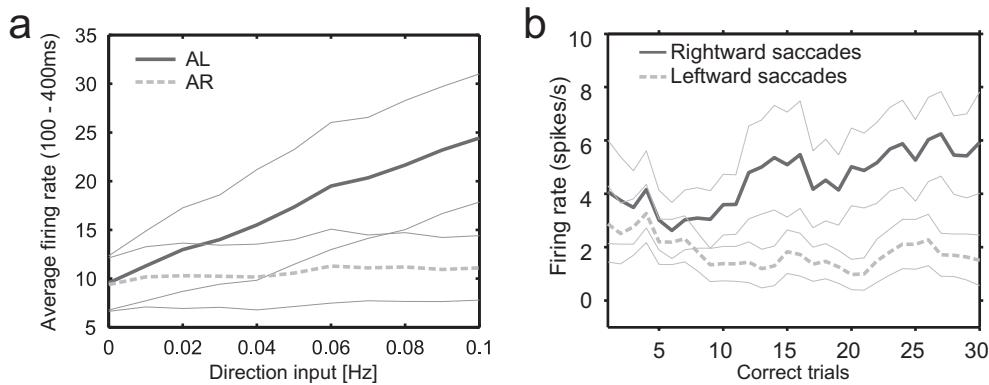


Figure 3.5: Average activity (and SEM) during peri-cue period for both model and experiment. a. The average activity of the model calculated for the pools AL and AR between 100-400ms as a function of direction input strength. The two pools show the same stimulus selectivity and code for the direction when stimulus A is presented. The values of the network parameters are fixed at $w_+ = 1.8$, $w_m = 0.7$, and $w_a = 0.9375$. b. Average activity as a function of the correct trial number for the experimental data (adapted from Pasupathy & Miller (2005)).

again. The underlying mechanism of this phenomenon is based on the network dynamics. Stimulus and response-direction input are applied to excitatory populations. These drive the inhibitory population which increases its activity with some time delay. Once the inhibitory effect increases, it reduces the direction selectivity of the network. After a transient low, the direction selectivity rises again, because the system settles in its stationary attractor which has a high direction selectivity. After 1000ms from cue onset, the progression of the correct and error trials separates. This is because we categorize the trials in correct and error using the average firing rate between 1000 and 1500ms. We used this period, since the monkeys had to issue their response at the end of the delay period and at this point in time the information about the response-direction must be available.

When comparing the time course of the direction selectivity of our model, which is based on hypotheses of the PFC, to the experimental data (Figure 3.2), we see that correct and error trials have a different time course. Whereas the direction selectivity in the model separates in the middle of the delay period, this separation appears in the experimental data already in the cue period and disappears at the end of the delay period. After the good fit of the model data in the cue period, this result is surprising but also interesting. By comparing experimental data and hypotheses, we can discuss the hypotheses on which the computational model was constructed.

3.4 Discussion

In this work we implemented a hypothesis driven approach to modeling. Based on experimental evidences and hypotheses we selected five cornerstones about the properties of a PFC module in arbitrary visuomotor mappings and integrated them in a computational model. Thereafter we compared the results to experimental data. Note that a priori we did not seek to fit our model to experimental data. Rather we wanted to test the validity of the hypotheses

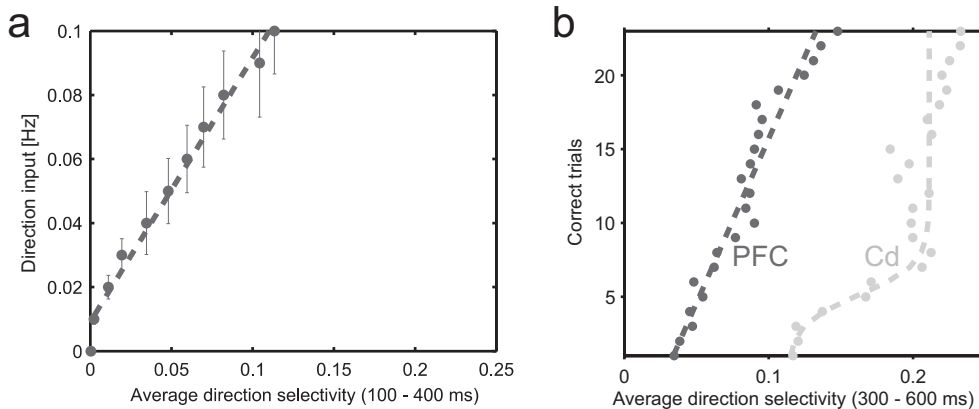


Figure 3.6: Average direction selectivity. a. Average peri-cue period direction selectivity (100 - 400ms) calculated from model data as a function of direction input. The dashed line shows a linear fit. The values of the network parameters are fixed at $w_+ = 1.8$, $w_m = 0.7$, and $w_a = 0.9375$. b. Average direction selectivity of experimental data during peri-cue epoch (adapted from Pasupathy & Miller (2005)). As in the model data, the progression of the PFC selectivity is almost linear.

about the PFC by comparing the modeling results to experimental data. In the following, we will first picture the dynamics of our computational model and compare it to the possible neurodynamics underlying the experimental data.

We sketch the dynamical behavior of the model system in simplified cartoons of a hypothetical energy landscape (Figure 3.8). The energy landscape shows two valleys which corresponds to two attractors, one at low and one at high firing rates. By influence of fluctuations or external input, the system behavior can switch from one attractor state to another, analogously to a ball moving in the landscape which is at rest in the bottoms of the valleys and can be moved by external forces. Figure 3.8 shows the behavior over the three phases of the simulation: prestimulus, stimulus and delay period. In the prestimulus period the system is at rest in the low attractor state. Due to the stimulus and response-direction input, the system jumps up to a high attractor state both for correct and error trials. This can be seen by the increase in direction selectivity in the cue period (see Figure 3.2). When the direction selectivity increases, the difference in firing rate between two pools (in our simulations AL and AR) increases. Since we set up the network as a single attractor system, the system jumps to a high attractor state (in our simulations a high activity state of AL). After the cue period, the direction selectivity stays at high levels for both correct and error trials indicating that in both cases the system stays in the high attractor state. The distinction between correct and error trials come about due to fluctuations among the fixed point of this attractor state. The system is stochastic and the firing rate can be lower or higher from trial to trial. The fluctuations in the system cause a higher direction selectivity in some trials and thus the association to be extracted or not.

The analogous analysis of the experimental data yields a different view of the neurodynamics. The direction selectivity of correct and error trials already parts in the stimulus period (Figure 3.2b). This suggests that the system either jumps up to a high attractor state or stays in the low attractor state in the stimulus period depending on a correct or an error trial (Figure 3.9). Thereafter the direction selectivity of the correct trials drops and eventually




Figure 3.7: Analysis of multiple spiking simulations. The values of the parameters are fixed at $w_+ = 1.8$, $w_m = 0.7$, and accordingly $w_a = 0.9375$. The direction input strength is varied. The graph shows the percentage of correctly identified association. The mean of the last 500ms of the delay period is calculated for every pool. An association is identified correctly if this mean of the pool connected to it (e.g., AL) has a value of above 10Hz and the one of other three selective pools of below 10Hz. For every data point, 100 spiking simulations were performed.

reaches the same level as the error trials at the end of the delay period. Thus, the dynamical system drops back to the low activity state in the delay period for the correct trials.

Overall, there are differences in the neurodynamical properties. First, we will review the five assumptions which we stated in the introduction and discuss their validity. Note that we achieved a particularly good fit to the correct trials in the cue period (see Figures 3.5, 3.6) which suggests a relation between the strength of the response-related input and the learning in correct trials.

First, the model neurons should show non-linear object-and-direction selective activity (Asaad et al. 1998). Using the meanfield technique, we specifically selected parameters which yielded an object-and-direction selective activity. We find that a specific ordering of the con-




Figure 3.8: Sketch of the dynamical behavior of the model over time. We show in a hypothetical two attractor energy landscape for low and high activation the behavior of the model. In the prestimulus period, the system is at rest. The stimulus and response-direction input excite the system so that the firings rates increase by jumping to a high attractor state. The distinction between correct and error trials is caused by fluctuations in the high attractor state.




Figure 3.9: Sketch of the dynamical behavior of the experimental data over time based on the Figure 3.2b. We show in a hypothetical two attractor energy landscape for low and high activation the behavior of the model. In the prestimulus period, the system is at rest. The stimulus and the response-direction input excite the system in a way that only correct trials jump to a high activity state whereas error trials remain in the low activity state. Since the difference in direction selectivity disappears at the end of the delay period, we hypothesize that the correct trials jump out of the high attractor.

nection strength, namely $w_+ > w_a > w_m$ is necessary to achieve this property. We analyzed a widespread parameter space which contained all possible orderings of connection weights. The non-linear object-and-direction selective activity is responsible for the directional selectivity observed in the model.

Second, stimulus information reaches the PFC from IT. IT is implicated in the processing of visual information about objects and anatomical studies suggests that projections to the PFC (Pandya & Kuypers 1969). Moreover a disconnection of PFC and IT using cross-lesion procedures (Gaffan & Harrison 1988, Parker & Gaffan 1998, Bussey et al. 2002) or transaction of the uncinate fascicle (Eacott & Gaffan 1992) cause deficits in learning arbitrary visuomotor mappings. Stimulus information itself is also present in the PFC during arbitrary visuomotor mappings (Asaad et al. 1998). Thus the stimulus input to our module might originate either directly from IT or from other PFC neurons or both. Since we do not model the originating brain regions explicitly, the exact origin of this input is not crucial for the simulation results. We assume this input to be constant in respect to learning. The perceived objects remained constant within the blocks and thus we also assumed their perception to be constant. Furthermore, imaging studies do not indicate a change in correlation between IT and PFC. The good match between our model and the experimental data during the cue period for the correct trials also supports this hypothesis.

Third, the basal ganglia work as a context detector and their output influences the activity of the PFC. The basal ganglia could generate a signal reflecting the response-direction (Hadj-Bouziane et al. 2003). This would also be consistent with the idea of context extraction by the basal ganglia in terms of action planning (Houk & Wise 1995). The basal ganglia are connected to the PFC in a cortico-basal ganglionic loop (Middleton & Strick 1994, Houk & Wise 1995, Graybiel 1998). There are two hypothesis about the role of the basal ganglia in learning: One is that the basal ganglia have a dominant role in the learning of new associations (Houk & Wise 1995, Bar-Gad et al. 2003, Pasupathy & Miller 2005), whereas the other one suggests that they are responsible for consolidating associations and thus habits (Graybiel 1998, Packard & Knowlton 2002). We build upon the first hypothesis. We do not model the basal ganglia explicitly, but include the characteristics of the activity as identified by experimental results in our modeling process. Thus we focus on the input characteristics to the prefrontal cortex and less on how the basal ganglia and the thalamus generate this input.

The response-direction input is applied at the same time as the stimulus input (onset at 0ms). The simultaneous onset of direction and stimulus input is a first approximation, since the exact timing of the two potential inputs is not known. We suggest that stimulus information might reach both the PFC and the basal ganglia directly from IT. Thereafter the basal ganglia extract the direction and transmit it via the thalamus back to the PFC. In addition, we assumed the onset of the response-direction input to be constant with learning. Pasupathy & Miller (2005) showed that the onset of the direction selectivity in the basal ganglia moved to an early position rapidly as learning progressed and reached a steady point in time after a few trials. The direction selectivity identified in the basal ganglia could be the origin of a response-direction input signal mediated by the thalamus. We varied the strength of that input. The underlying hypothesis is that slow synaptic plasticity in the synapses which target PFC neurons from the thalamus could enhance the connection strength causing an increase in the response-direction input (Houk & Wise 1995). This is consistent with observations of imaging studies: Both prefrontal areas and the basal-ganglia were identified to take part in the arbitrary visuomotor associations (Toni & Passingham 1999, Toni, Ramnani, Josephs, Ashburner & Passingham 2001, Toni, Rushworth & Passingham 2001) and, more importantly, the correlation between the two areas increased during learning (Toni et al. 2002, Boettiger & D'Esposito 2005). The linear increase in the direction bias causes changes which correspond well to the experimental observations during learning. Especially in the cue period, both the direction selectivity and average firing rate show good fits to experimental data.

Conceptually, the direction bias is different from the context bias implemented by earlier modeling approaches (Deco & Rolls 2003, 2005*b*, Loh & Deco 2005). The context input would go to two pools which represent the association active during an experimental trial, e.g. AL and BR. This would mean that the two associations would not be learned independently. Learning the association AL would also favor BR. In the extreme case, the monkey could perform a one-trial reversal, i.e., the monkey would switch context after the first error. The experiment by Thorpe et al. (1983) shows this evidence which is modeled by Deco & Rolls (2005*b*). Although such rule would greatly facilitate the execution of the task, there is no clear evidence in this conditional visuomotor experiment that the monkeys actually did make use of this relation.

Fourth, we assume that sustained cortical-thalamic loop feedback serves as a working memory of the directional information extracted by the basal ganglia. Houk & Wise (1995) explain the mechanism as follows: Striatal neurons in the basal ganglia receive input from the several cortical regions. By virtue of reinforcement signals originating from midbrain dopamine neurons, the striatal neurons learn to recognize complex activity patterns such as the association in a conditional visuomotor task. Via disinhibitory mechanism implemented by in the striatum-pallidum-thalamus connection, a detected direction might cause a positive feedback in the reciprocal cortical-thalamic loop. This activity could be self-sustained and thereby guide future actions such as an upcoming motor response to a stimulus. Our modeling data shows that the stimulus and the response-direction input compute the correct association, so that the correct response can be issued. Furthermore, the association is maintained in the PFC module (see Figure 3.2). However, the experimental data show a significant decrease in direction selectivity after 1000ms. These could be due to a variety of reasons which are also related to the fifth hypothesis.

Fifth, we hypothesize that the PFC module holds the extracted association in working memory. There is much evidence that the PFC is involved in working memory related

processes implemented by delayed neural activity (Kubota & Niki 1971, Fuster & Alexander 1971). Based on the fact that any neural system that contributes to visuomotor learning must retain information about the IS-response mapping (Buch et al. 2006), we reasoned that the PFC would be in the ideal position. The underlying neural model is setup up for working memory properties (Brunel & Wang 2001) and we chose the network parameter accordingly (see Methods section). However, the direction selectivity in the experimental data drops out in the middle of the delay period. This suggests that the information about the movement direction is not kept in the prefrontal area or that relatively few neurons may be needed to maintain this information during the delay period. Thus, either the working memory hypothesis of the prefrontal cortex or the sustained activity in the cortico-thalamic loop could be partly incorrect assumptions.

The dynamical behavior could also be explained by other neurodynamical mechanisms. For example, the activity could propagate to other regions of the brain involved in the response. This idea is also supported by experimental evidences which show that the number of object selective neurons decreases in the prefrontal cortex from the beginning to the end of a trial whereas the number of direction selective neurons increases (Asaad et al. 1998). A propagation between layers has been observed in multilayer dynamical systems as studied by (Deco & Rolls 2003, Loh & Deco 2005). In addition, in these models the activity of correct and error trials differ from the beginning of the cue period, which is consistent with the progression of the direction selectivity identified by Pasupathy & Miller (2005). Loh & Deco (2005) related the multistability in a multilayer neural network to response-space exploration in arbitrary visuomotor tasks. Thus, the errors might not be caused by working memory failures in the prefrontal cortex but by computational properties of the complete system supporting arbitrary visuomotor associations. In this sense, one could also view our model as representation of a larger dynamical system compressed in one layer.

The model in this paper was specifically constructed to target the activity in the prefrontal cortex upon influences from connecting brain areas. Fusi et al. (2007) presented a model based on the same experimental data (Asaad et al. 1998, Pasupathy & Miller 2005) which features a learning rule and a reversal mechanism. We did not address this in our model. The two models make different claims about how the activity in the prefrontal cortex arises. The model by Fusi et al. (2007) features one input to the prefrontal cortex which is modulated by learning. The learning rule is responsible for acquiring the associations and the reset after the reversals. We propose that the activity of the PFC arises due to the influence of several brain areas, the IT and the basal ganglia. Thereby we propose how the object-and-direction selective neurons might come about. If the input from the basal ganglia to the PFC indeed exists, could be tested directly experimentally by disconnecting the basal ganglia from the PFC. This experiment was proposed by Nixon et al. (2004).

Our analysis shows that the dynamics of the prefrontal cortex can arise due to several external influences, namely the stimulus bias and the direction bias which we envision to originate from the IT area and the basal-ganglionic-thalamic-cortical loop, respectively. Neuropsychological and physiological evidences point in this direction. Furthermore, the dominant dynamics in the PFC seem to be transient, since the comparison of model and experimental data do not suggest that a great population of neurons showing stable working memory properties exist. This does not imply that the PFC does not have working memory capabilities, but merely that in this task it is not used to a great extent. The activity propagates along one or multiple processing pathways (Hadj-Bouziane et al. 2003, Loh & Deco 2005) and thereby passing through the PFC. What is the role of the PFC in the processing pathway? Lesion

studies suggest that the prefrontal cortex is important for the fast learning (Murray et al. 2000, Wang et al. 2000, Bussey et al. 2001) of arbitrary visuomotor associations. Maybe in conjunction with hippocampal regions, the PFC adds a fast learning component and thereby speeds up learning. The proposed plasticity in the cortico-thalamic loop (Houk & Wise 1995) is a possible candidate. A second processing pathway from IT and the basal ganglia to the premotor cortex might provide slower learning mechanisms as disconnection of the basal ganglia and the premotor cortex affects the execution of well-established associations but spares fast learning mechanisms (Nixon et al. 2004). An experiment disconnecting the basal ganglia and the prefrontal cortex could clarify this issue.

In our modeling we feature a hypothesis driven approach. This means that a priori we did not seek to fit model data but wanted to test several existing hypotheses by integrating them in a computational model. Computational modeling is ideal for this approach since it allows to both integrate several hypotheses in one system and to study the contributions of each hypothesis separately. Moreover, the language of computational models is explicit and thus forces one to make concrete assumptions. We believe that the presented way of modeling, namely to test existing hypotheses instead of seeking model fitting to experimental data, might contribute significantly to the theoretical understanding of brain dynamics, since it emphasizes the idea of integrating neuroscientific evidences towards comprehensive models of brain functions.

Acknowledgments

We thank Ralph G. Andrzejak and Anders Ledberg for discussions and insights which improved our work considerably.

Chapter 4

A dynamical systems hypothesis of schizophrenia

The work presented in this chapter is published in PLoS Computational Biology (Loh et al. 2007).

Abstract

We propose a top-down approach to the symptoms of schizophrenia based on a statistical dynamical framework. We show that a reduced depth in the basins of attraction of cortical attractor states destabilizes the activity at the network level due to the constant statistical fluctuations caused by the stochastic spiking of neurons. In integrate-and-fire network simulations, a decrease in the NMDA receptor conductances, which reduces the depth of the attractor basins, decreases the stability of short term memory states and increases distractibility. The cognitive symptoms of schizophrenia such as distractibility, working memory deficits or poor attention could be caused by this instability of attractor states in prefrontal cortical networks. Lower firing rates are also produced, and in the orbitofrontal and anterior cingulate cortex could account for the negative symptoms including a reduction of emotions. Decreasing the GABA as well as the NMDA conductances produces not only switches between the attractor states, but also jumps from spontaneous activity into one of the attractors. We relate this to the positive symptoms of schizophrenia including delusions, paranoia, and hallucinations, which may arise because the basins of attraction are shallow and there is instability in temporal lobe semantic memory networks, leading thoughts to move too freely round the attractor energy landscape.

4.1 Introduction

Schizophrenia is a major mental illness, which has a great impact on patients and their environment. One of the difficulties in proposing models for schizophrenia is the complexity and heterogeneity of the illness. We propose that part of the reason for the inconsistent symptoms may be a reduced signal-to-noise ratio and increased statistical fluctuations in different cortical brain networks. The novelty of the approach described here is that instead of basing our hypothesis purely on biological mechanisms, we develop a top-down approach based on the different types of symptoms and relate them to instabilities in attractor neural networks (Rolls 2005).

The main assumption of our hypothesis is that attractor dynamics are important in cognitive processes (O'Reilly 2006). Our hypothesis is based on the concept of attractor dynamics in a network of interconnected neurons which in their associatively modified synaptic connections store a set of patterns, which could be memories, perceptual representations, or thoughts (Hopfield 1982, Amit 1989, Rolls & Deco 2002). The attractor states are important in cognitive processes such as short term memory, attention, and action selection (Deco & Rolls 2005a). The network may be in a state of spontaneous activity, or one set of neurons may have a high firing rate, each set representing a different memory state, normally recalled in response to a retrieval stimulus. Each of the states is an attractor in the sense that retrieval stimuli cause the network to fall into the closest attractor state, and thus to recall a complete memory in response to a partial or incorrect cue. Each attractor state can produce stable and continuing or persistent firing of the relevant neurons. The concept of an energy landscape (Hopfield 1982) is that each pattern has a basin of attraction, and each is stable if the basins are far apart, and also if each basin is deep, caused for example by high firing rates and strong synaptic connections between the neurons representing each pattern, which together make the attractor state resistant to distraction by a different stimulus. The spontaneous firing state, before a retrieval cue is applied, should also be stable. Noise in the network caused by statistical fluctuations in the stochastic spiking of different neurons can contribute to making the network transition from one state to another, and we take this into account by performing integrate-and-fire simulations with spiking activity, and relate this to the concept of an altered signal-to-noise ratio in schizophrenia (Winterer et al. 2000, 2004, 2006).

Schizophrenia is characterized by three main types of symptom: cognitive dysfunction, negative symptoms, and positive symptoms (Liddle 1987, Baxter & Liddle 1998, Mueser & McGurk 2004). We consider how the basic characteristics of these three categories might be produced in a neurodynamical system, as follows. Dysfunction of working memory, the core of the cognitive symptoms, may be related to instabilities of persistent attractor states (Durstewitz et al. 2000b, Wang 2001) which we show can be produced by reduced firing rates in attractor networks, in brain regions such as the prefrontal cortex. The negative symptoms such as flattening of affect or reduction of emotions may be caused by a consistent reduction in firing rates of neurons in regions associated with emotion such as the orbitofrontal cortex (Rolls 2005). These hypotheses are supported by the frequently observed hypofrontality, a reduced activity in frontal brain regions in schizophrenic patients during cognitive tasks (Ingvar & Franzen 1974, Kircher & Thienel 2005, Scheuerecker et al. 2007). The positive symptoms are characterized by phenomenologically overactive perceptions or thoughts such as hallucinations or delusions which are reflected for example by higher activity in the temporal lobes (Shergill et al. 2000, Scheuerecker et al. 2007). We relate this category of symptoms to a spontaneous appearance of activity in attractor networks in the brain and more generally

to instability of both the spontaneous and persistent attractor states. We do not try to account for every detail of schizophrenic symptoms which are diverse and vary among patients, but instead show how these three main categories of symptoms could be related to changes in the stability of dynamical attractor systems in the brain, and how the changes in the stability might be produced by changes at the level of the currents being passed through receptor-activated ion channels at synapses. There are specific symptoms such as aberrant eye movements which cannot be accounted for by this general scheme.

In particular we were interested in how these symptoms are related. Negative and cognitive symptoms typically precede the first psychotic episode (Lieberman et al. 2001, Hafner et al. 2003). Positive symptoms can be treated in most cases with neuroleptics whereas negative and cognitive symptoms persist, at least for typical neuroleptics. Can a mapping onto a dynamical system help to understand these relations? After proposing a dynamical systems hypothesis for the different symptoms of schizophrenia, we study a standard neural network model (Brunel & Wang 2001) of cortical dynamics specifically in relation to our hypothesis. We were especially interested in how excitation and inhibition implemented by NMDA and GABA synapses affect the network dynamics. Alterations in the efficacies of the NMDA and GABA channels have been identified in the pathology of schizophrenia (Coyle et al. 2003, Lewis et al. 2005), and transmitters such as dopamine influence the currents in these receptor-activated channels (Seamans & Yang 2004). Do NMDA and GABA currents have antagonistic effects or do they have a special role in the network dynamics? How could this be related to our hypothesis of schizophrenia? Building upon the current body of neural network research, we specifically conduct neural network simulations to substantiate our dynamical systems hypothesis of schizophrenia. While focussing on NMDA and GABA synapses in the Results, in the Discussion we consider how altered transmission at D1 and D2 receptors by modulating NMDA and GABA conductances could not only influence working memory, which has been investigated previously (Durstewitz et al. 1999, 2000*a*, Brunel & Wang 2001, Durstewitz & Seamans 2002), but could in particular influence the different symptoms of schizophrenia.

4.2 Results

4.2.1 Dynamical systems hypothesis

We relate the three types of symptoms of schizophrenia to the dynamical systems attractor framework described at the end of the Introduction as follows (Rolls 2005).

The cognitive symptoms of schizophrenia include distractibility, poor attention, and the dysexecutive syndrome (Liddle 1987, Green 1996, Mueser & McGurk 2004). The core of the cognitive symptoms is a working memory deficit in which there is a difficulty in maintaining items in short term memory (Goldman-Rakic 1994, 1999), which could directly or indirectly account for a wide range of the cognitive symptoms. We propose that these symptoms may be related to instabilities of persistent states in attractor neural networks, consistent with the body of theoretical research on network models of working memory (Durstewitz et al. 2000*b*). The neurons are firing at a lower rate, leading to shallower basins of attraction of the persistent states, and thus a difficulty in maintaining a stable short term memory, normally the source of the bias in biased competition models of attention (Rolls & Deco 2002, Deco & Rolls 2005*b*). The shallower basins of attraction would thus result in working memory deficits, poor attention, distractibility, and problems with executive function and action selection (Deco & Rolls 2003, 2005*b*).

The negative symptoms refer to the flattening of affect and a reduction in emotion. Behavioural indicators are blunted affect, emotional and passive withdrawal, poor rapport, lack of spontaneity, motor retardation, and disturbance of volition (Liddle 1987, Mueser & McGurk 2004). We propose that these symptoms are related to decreases in firing rates in the orbitofrontal cortex and/or anterior cingulate cortex (Rolls 2005), where neuronal firing rates and activations in fMRI investigations are correlated with reward value and pleasure (Rolls 2005, 2006). Consistent with this, imaging studies have identified a relationship between negative symptoms and prefrontal hypometabolism, i.e. a reduced activation of frontal areas (Wolkin et al. 1992, Aleman & Kahn 2005).

The positive symptoms of schizophrenia include bizarre (psychotic) trains of thoughts, hallucinations, and (paranoid) delusions (Liddle 1987, Mueser & McGurk 2004). We propose that these symptoms are related to shallow basins of attraction of both the spontaneous and persistent states in the temporal lobe semantic memory networks and to the the statistical fluctuations caused by the probabilistic spiking of the neurons. This could result in activations arising spontaneously, and thoughts moving too freely round the energy landscape, loosely from thought to weakly associated thought, leading to bizarre thoughts and associations, which may eventually over time be associated together in semantic memory to lead to false beliefs and delusions. Consistent with this, neuroimaging studies suggest higher activation especially in areas of the temporal lobe (Weiss & Heckers 1999, Shergill et al. 2000, Scheuerecker et al. 2007).

To further investigate our hypothesis, we use an attractor network, as this is likely to be implemented in many parts of the cerebral cortex by the recurrent collateral connections between pyramidal cells, and has short term memory properties with basins of attraction which allow systematic investigation of stability and distractibility. The particular neural network implementation we adopt includes channels activated by AMPA, NMDA and GABA_A receptors and allows not only the spiking activity to be simulated, but also a consistent mean-field approach to be used (Brunel & Wang 2001).

4.2.2 The depth of the basins of attraction: flow analysis

First we introduce an analytical approach to the concepts of how changes in transmitters could affect the depth of the basins of attraction in networks in ways that may be related to the symptoms of schizophrenia.

The depth of the basins of attraction can be assessed by calculating the flow towards the attractors using the meanfield reduction (see Methods and Mascaró & Amit (1999)). The flow is the force that draws the dynamical system to one of the attractors. Figure 4.1 shows the flow between the spontaneous and persistent state in a network featuring one selective pool. The curve for (Normal) shows the flow for the reference baseline condition. The crossing of the curve and the 0-flow axis represent the fixed points, which are either stable (negative derivative) or unstable (positive derivative). We use this line as a reference to assess the relative depth of the energy landscape. The system flows either into a spontaneous firing state (of approximately 2 Hz), or a high firing attractor state.

A reduction of NMDA (-NMDA) shows a stronger flow than the unchanged condition at low firing rates towards the spontaneous attractor (at about 2 Hz). The absolute values of the function are higher compared to the normal condition until the first unstable fixed point (at around 6–7 Hz). The basin of attraction towards the persistent attractor at high firing rates yields the reverse picture. Here the (-NMDA) curve is clearly below the unchanged condition




Figure 4.1: Flow of a single attractor network. We assesses the flow of an attractor network with one selective pool using the meanfield analysis for different modifications of synaptic conductances (see Methods section). The flow represents the force which drives the system towards one of the stable attractors. The stable/unstable attractor states are at crossings with the flow=0 value on the flow axis with a negative/positive derivative respectively. A modulations of the synapses labelled as (-NMDA) and (-GABA) corresponds to a reduction of 4.5% and 9% respectively of the efficacies. A pool cohesion of $w_+ = 1.6$ and a selective pool size of 80 neurons were used for the simulations.

and the flow towards the attractor is smaller. Overall, the basin of attraction is deeper for the spontaneous state and shallower for the persistent state compared to the unchanged condition. This pattern fits to the cognitive symptoms of schizophrenia as proposed in our hypothesis. Also note that the firing rate of the persistent fixed point is reduced in the (-NMDA) condition (crossing with the flow=0 axis), which is consistent with the hypothesis for the negative symptoms.

A reduction of the GABA conductance (-GABA) yields the opposite pattern to that in the reduced NMDA condition. Here the basin of attraction of the persistent state is deeper. This is not a condition that we suggest is related to the symptoms of schizophrenia.

However, in the condition in which both the NMDA and the GABA conductances are reduced (-NMDA, -GABA), the persistent state basin of attraction is shallower, and the spontaneous state basin is a little shallower. This condition corresponds to the proposed landscape for the positive symptoms as considered above. In particular, in the (-NMDA, -GABA) condition, the system would be less stable in the persistent state, tending to move to another attractor easily, and less stable in the spontaneous state, so tending to move too readily into an attractor from spontaneous activity.

Overall, the flow analysis suggests that both the cognitive and negative symptoms could be related to a decrease in the NMDA conductances. This is consistent with the fact that these two symptoms usually appear together. The flow analysis suggests that the positive symptoms are related to a reduction in both NMDA and GABA. Thus, the transition from the cognitive and negative symptoms to the positive, psychotic, symptoms might be caused by an additional decrease in the GABA conductance. It is notable that excitation and inhibition




Figure 4.2: Example trials of the integrate-and-fire attractor network simulations. a. Trials in which no recall stimulus was applied to S1, and spontaneous firing continued until the end of the trial in the ‘spontaneous’ simulation condition. However, on one of the trials illustrated, the firing increased as the network climbed into one of the attractors because of the statistical fluctuations caused by the spiking dynamics. b. Trials in which a recall stimulus was applied to S1 at 0–500 ms, and continuing firing normally occurred until the end of the trial in the ‘persistent’ simulation condition. However, on one of the trials illustrated, the firing decreased as the network fell out of the attractor because of the statistical fluctuations caused by the spiking dynamics. In these simulations the network parameter is $w_+ = 2.1$ and the modulation of the synapses is (-NMDA) and (-GABA). This corresponds to a reduction of 4.5% and 9% respectively in the currents produced through these synapses. The average firing rate of all the neurons in the S1 pool is shown.

do not cancel each other out as assumed by many models, but have distinct influences on the network dynamics.

4.2.3 Concept of stability in network simulations

The flow analysis provides insight into how the depth of the basins of attraction and the firing rates are influenced by changes in the conductivities of the channels activated via NMDA and GABA receptors. However the overall stability of the different attractors is affected not only by the depth of the basins of attraction, but also by the breadth and distance apart of the basins, and by the statistical noise generated by the randomness of the spiking of the different neurons. These statistical fluctuations play a role in the way in which the system moves from one state to another, for these statistical fluctuations can cause hills in the energy landscape to be crossed stochastically. Since the meanfield analyses do not take these properties into account, we investigate the system further using large-scale integrate-and-fire network simulations and measuring the statistics of the network behavior.

To clarify the concept of stability, we show examples of trials of spontaneous and persistent simulations in which the statistical fluctuations have different impacts on the temporal dynamics. Figure 4.2 shows the possibilities, as follows.

In the spontaneous state simulations, no cue is applied, and we are interested in whether




Figure 4.3: The simulation protocols. Stimuli to either S1 or S2 are applied at different times depending on the type of simulations. The spontaneous simulations include no input. The persistent simulations assess how stable a stimulus is retained by the network. The distractor simulations add a distractor stimulus to further address the stability of the network activity.

the network remains stably in the spontaneous firing state, or whether it is unstable and on some trials due to statistical fluctuations enters one of the attractors, thus falsely retrieving a memory. Figure 4.2a shows an example of a trial on which the network correctly stays in the low spontaneous firing rate regime, and another trial (labelled spontaneous unstable) in which statistical spiking-related fluctuations in the network cause it to enter a high activity state, moving into one of the attractors even without a stimulus.

In the persistent state simulations, a strong excitatory input is given to the S1 neuronal population between 0 and 500 ms (Figure 4.3). Two such trials are shown in Figure 4.2b. In one, the S1 neurons (correctly) keep firing at approximately 30 Hz after the retrieval cue is removed at 500 ms. However, due to statistical fluctuations in the network related to the spiking activity, on the trial labelled persistent unstable the high firing rate in the attractor for S1 was not stable, and the firing decreased back towards the spontaneous level, in the example shown starting after 1.5 s. This trial illustrates the failure to maintain a stable short term memory state, even when no distractor is applied.

In Figure 4.2 the transitions to the incorrect activity states are caused by statistical fluctuations in the spiking activity of the integrate-and-fire neurons, and the depth of the basins of attraction which has been reduced in the simulations shown by reducing both the NMDA and the GABA currents, as indicated in the Figure legend. We hypothesize that such instabilities are related to the symptoms of schizophrenia. We note that there are two sources of noise in the spiking networks that cause the statistical fluctuations: the randomly arriving external Poisson spike trains, and the statistical fluctuations caused by the spiking of the neurons in the finite sized network. The magnitude of these fluctuations increases as the number of neurons in the network becomes smaller (Mattia & Del Giudice 2004).

For our investigations, we selected $w_+ = 2.1$, which with the default values of the NMDA and GABA conductances yielded stable dynamics, that is, a stable spontaneous state if no retrieval cue was applied, and a stable state of persistent firing after a retrieval cue had been applied and removed. To investigate the effects of changes (modulations) in the NMDA and GABA conductances, we chose for demonstration purposes a reduction of 4.5% and 9%, respectively, as these could cause instabilities, as illustrated in Fig. 4.2. However, the exact values are not crucial to observe the effects described. The magnitudes of these reductions are smaller than those that can be produced experimentally (Durstewitz & Seamans 2002, Seamans & Yang 2004). A strength of our approach is that we show that even quite small




Figure 4.4: Stability of the spontaneous and persistent state. We assessed how often in 100 trials the average activity during the last second (2–3 s) stayed above 10 Hz. The value shows how often it stayed more in the respective state than in the unmodulated reference state. A negative percentage means that the system was less stable than in the reference state. A modulation of the synapses shown as (-NMDA) and (-GABA) corresponds to a reduction of 4.5% and 9% respectively in their efficacies. We assessed with the Binomial distribution the statistical significance of the effects observed, with $P < 0.01$ relative to the reference state marked by **, and $P < 0.02$ by *.

reductions in the synaptic currents can alter the global behaviour of the network, e.g. the stability of its attractors.

4.2.4 Stability

We assessed how the stability of both the spontaneous and persistent states changes when NMDA and GABA efficacies are modulated. Specifically we ran multiple trial integrate-and-fire network simulations and counted how often the system maintained the spontaneous or persistent state, assessed by the firing rate in the last second of the simulation (2–3 s) of each 3 s trial. Figure 4.4 shows the stability of the spontaneous and persistent attractor relative to the unmodulated reference state (Normal). A negative percentage means that the system was less stable than in the unmodulated state.

A reduction of the NMDA conductance (-NMDA) reduces the stability of the persistent state drastically, while slightly increasing the stability of the spontaneous state (see Figure 4.4). We hypothesized that this type of change might be related to the cognitive symptoms, since it shows a reduced stability of the working memory properties. A reduction of GABA shows the opposite pattern: a slight reduction in the stability of the spontaneous state, and an increased stability of the persistent (i.e. attractor) state (see Figure 4.4).

When both NMDA and GABA are reduced one might think that these two counterbalancing effects (excitatory and inhibitory) would either cancel each other out or yield a tradeoff between the stability of the spontaneous and persistent state. However, this is not the case. The stability of both the spontaneous state and the persistent state is reduced (see Figure 4.4). We relate this pattern to the positive symptoms of schizophrenia, in which both the spontaneous and attractor states are shallow, and the system merely jumps by the influence




Figure 4.5: Wandering between attractor states. We simulated a single long trial (60 s) in the spontaneous test condition for the synaptic modification (-NMDA, -GABA). The two curves show the activity of the two selective pools over time smoothed with a 1 s sliding averaging window. The activity moves noisily between the attractor for the spontaneous state and the two persistent states S1 and S2 by virtue of statistical fluctuations caused by the randomness of the spiking activity.

of statistical fluctuations between the different (spontaneous and attractor) states.

To investigate more directly the wandering between spontaneous and several different persistent attractor states, we simulated the condition with decreased NMDA and GABA conductances over a long time period in which no cue stimulus input was given. Figure 4.5 shows the firing rates of the two selective pools S1 and S2. The high activity switches between the two attractors due to the influence of fluctuations, which corresponds to spontaneous wandering in a shallow energy landscape, corresponding for example to sudden jumps between unrelated cognitive processes. These results are consistent with the flow analysis and demonstrate that the changes in the attractor landscape influence the behavior at the stochastic level.

4.2.5 Distractibility

As distractibility is directly related to the symptoms of schizophrenia, we ran simulations specifically to assess this property using persistent and distractor simulations (see Figure 4.3). A distractor strength of 0 Hz corresponds to the persistent condition described in the preceding section (stability). Figure 4.6 shows the stability and distractibility for reductions of NMDA and GABA currents. The reference state is labelled ‘Normal’. In this state, pool S1 continued to maintain its attractor firing without any distractor (distractor strength = 0 Hz) throughout the delay period on almost 90 % of the trials. In both conditions which reduce the NMDA current (labelled -NMDA) the network was less and less able to maintain the S1 attractor firing as the distractor stimulus strength was increased through the range 0–80 Hz. The stability of the persistent state was reduced, and also the distractibility was increased, as shown by the fact that increasing distractor currents applied to S2 could move the attractor away from S1. The implication therefore is that a reduction of the NMDA currents could




Figure 4.6: Stability and distractibility as a function of the distractor strength and modulations of the synaptic efficacies. We assessed how often in 100 trials the average activity during the last second (2–3 s) stayed above 10 Hz in the S1 attractor. The modulation of the synapses (-NMDA) and (-GABA) corresponds to a reduction of 4.5% and 9%, respectively. The strength of the distractor stimulus applied to S2 is an increase in firing rate above the 2.4 kHz background activity which is distributed among 800 synapses per neuron. The lower the sustained activity in S1 the higher is the distractibility. The standard deviations were approximated with the binomial distribution.

cause the cognitive symptoms of schizophrenia, by making short term memory networks less stable and more distractible, thereby reducing the ability to maintain attention. Reducing only the GABA currents (-GABA) reduces the distractibility for low distractor strengths and coincides with the reference (Normal) condition at high values of the distractor strengths.

4.2.6 Signal-to-noise ratio

We further investigated the signal-to-noise ratio in relation to the changes in synaptic conductances. The signal-to-noise ratio denotes the level of a signal relative to the level of background noise. In an attractor network, a high signal-to-noise ratio indicates that the network will maintain the attractor stably, as it will be unlikely to be disrupted by spiking-related statistical fluctuations that are the source of the noise in the network. Figure 4.7 shows the signal-to-noise ratio of a measure related to the fMRI BOLD signal. (This measure described in the legend to Figure 4.7 and below was used because the experimental data with which we wish to compare the simulation results use fMRI measures (Winterer et al. 2000, 2004, 2006). The index we used of the activity of the network was the total synaptic current of selective pool 1 averaged over the whole simulation time of 3 s, to take the temporal filtering properties of the BOLD signal into account, given the typical time course which lasts for several seconds of the haemodynamic response function (Deco et al. 2004). Further, we subtracted the averages of the spontaneous trial simulations which represent the baseline activity from the persistent trial simulation values. The signal to noise ratio was calculated from the mean of this index across trials divided by the standard deviation of the index, both measured using 1000 simulation trials. If the network sometimes had high activity, and




Figure 4.7: Signal-to-noise ratio of a measure related to the BOLD signal. We computed the mean and standard deviation of averages of the synaptic currents of the selective pool over the whole simulation period of a persistent condition simulation. The mean of the spontaneous baseline condition was subtracted. We conducted 1000 simulated trials. The signal-to-noise ratio is calculated by division of the mean synaptic current by the standard deviation measured using 1000 trials. The error bars indicate an estimate of the standard deviation measured over 20 epochs containing 50 trials each. A modulation of the synapses shown as (-NMDA) and (-GABA) corresponds to a reduction of 4.5% and 9% respectively in their conductances.

sometimes low, then the signal-to-noise measure gave a low value. If the network reliably stayed in the high persistent firing states, then the signal-to-noise ratio measure was high.) As shown in Figure 4.7, we found that in all the cases in which the NMDA, or the GABA, conductance, or both, are reduced, the signal-to-noise ratio, computed by the mean divided by the standard deviation, is also reduced. This relates to recent experimental observations which show a decreased signal-to-noise ratio in schizophrenic patients (Winterer et al. 2000, 2004, 2006). Here we directly relate a decrease in the signal-to-noise ratio to changes (in this case decreases) in receptor activated synaptic channel conductances. Given these results, it would be of interest in future studies to model the exact paradigm used by Winterer et al. (2006).

4.3 Discussion

We have proposed a hypothesis that relates the cognitive, negative, and positive symptoms of schizophrenia (Liddle 1987, Mueser & McGurk 2004) to the depth of basins of attraction and to the stability properties of attractor networks caused by statistical fluctuations of spiking neurons. This assumes that some cognitive processes can be understood as dynamical attractor systems, which is an established hypothesis in areas such as working memory, but has also been used many other areas (Rolls & Deco 2002, O'Reilly 2006). Our approach applies this concept to mental illnesses (Bender et al. 2006). Due to the diversity of schizophrenia, our general hypothesis is meant to serve as a heuristic of how the different kinds of symptoms might arise and are related. We investigated the hypothesis empirically in a computational

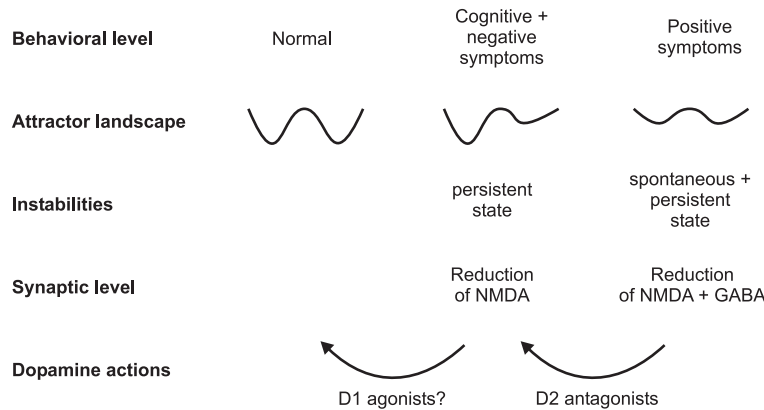


Figure 4.8: Summary of attractor hypothesis of schizophrenic symptoms and simulation results (see text). The first basin (from the left) in each energy landscape is the spontaneous state, and the second basin is the persistent attractor state. The vertical axis of each landscape is the energy potential.

attractor framework to capture an important aspect of cortical functionality. Figure 4.8 summarizes our hypothesis and its relation to the investigations of a neural network.

The middle column in Figure 4.8 shows the overview for the cognitive and negative symptoms. The core of the cognitive symptoms is a failure of working memory and attentional mechanisms. Working memory activity is related to the ongoing (i.e. persistent) firing of neurons during the delay period of cognitive tasks (Goldman-Rakic 1994, 1999). This could be implemented by associatively modifiable synapses between the recurrent collateral synapses of cortical pyramidal cells (Rolls & Treves 1998, Durstewitz et al. 2000*b*, Renart et al. 2001, Wang 2001). We propose that the cognitive symptoms of schizophrenia could arise because the basins of attraction of the persistent states in the prefrontal cortex become too shallow. This leads in combination with the statistical fluctuations due to randomness of the spiking activity to either a fall out of an active working memory state or to a shift to a different attractor state, leading to a failure to maintain attention and thereby impairing executive function. The hypofrontality in schizophrenia, that is less activation in frontal brain regions during working memory tasks (Ingvar & Franzen 1974, Carter et al. 1998), is in line with our hypothesis, since the firing rates of the persistent state are lower in the reduced NMDA condition (Figure 4.1), and the system spends on average less time in the persistent state, since it is less stable than in the normal condition (Figure 4.4). In addition, a reduced signal-to-noise ratio as shown in our simulations (Figure 4.7) has also been identified in imaging studies (Winterer et al. 2000, 2004, 2006). Our simulations suggest that a reduction in NMDA conductance at the synaptic level (see Figure 4.8) can account for this phenomenon. This is in line with previous work on the stability of working memory networks (Wang 1999, Durstewitz et al. 2000*a*, Wang 2001).

A reduction of the NMDA conductance also results in a reduction of the firing rates of the neurons in the persistent state (see Figure 4.1 and Brunel & Wang (2001)). We relate this, following (Rolls 2005), to the negative symptoms which include flattening of affect, a reduction in emotion, emotional and social withdrawal, poor rapport, passive withdrawal, lack of spontaneity, motor retardation, apathy, and disturbance of motivation. These symptoms are related to decreases in activity in the orbitofrontal cortex and/or anterior cingulate cortex

(Wolkin et al. 1992, Aleman & Kahn 2005), both of which are implicated in emotion (Rolls 2005, 2006, 2005). The emotional states represented in the orbitofrontal cortex and anterior cingulate cortex include states elicited both by rewards and punishers. Our hypothesis is that both would be reduced by the mechanism described. Correspondingly, motivation would be reduced in the same way, in that motivation is a state in which we work to obtain goals (rewards) or avoid punishers (Rolls 2005).

Both the negative and cognitive symptoms thus could be caused by a reduction of the NMDA conductance in attractor networks. The proposed mechanism links the cognitive and negative symptoms of schizophrenia in an attractor framework and is consistent with a close relation between the cognitive and negative symptoms: blockade of NMDA receptors by dissociative anesthetics such as ketamine produces in normal subjects schizophrenic symptoms including both negative and cognitive impairments (Malhotra et al. 1996, Newcomer et al. 1999); agents that enhance NMDA receptor function reduce the negative symptoms and improve the cognitive abilities of schizophrenic patients (Goff & Coyle 2001); and the cognitive and negative symptoms occur early in the illness and precede the first episode of positive symptoms (Lieberman et al. 2001, Hafner et al. 2003, Mueser & McGurk 2004). Consistent with this hypothesized role of a reduction in NMDA conductances being involved in schizophrenia, postmortem studies of schizophrenia have identified abnormalities in glutamate receptor density in regions such as the prefrontal cortex, thalamus and the temporal lobe (Goff & Coyle 2001, Coyle et al. 2003), brain areas that are active during the performance of cognitive tasks.

The dopamine D1 receptor has been shown to modulate the performance of working memory tasks (Sawaguchi & Goldman-Rakic 1991, 1994, Goldman-Rakic 1999, Castner et al. 2000). An increase in D1 receptor activation has been shown to increase the NMDA current (Durstewitz & Seamans 2002, Seamans & Yang 2004), and modeling studies have shown that this increase is related to the stability of working memory states (Durstewitz et al. 1999, 2000*a*, Brunel & Wang 2001). Imaging data also support the importance of the D1 receptor in schizophrenia (Okubo, Suhara, Sudo & Toru 1997, Okubo, Suhara, Suzuki, Kobayashi, Inoue, Terasaki, Someya, Sassa, Sudo, Matsushima, Iyo, Tateno & Toru 1997). We therefore suggest that an increased activation of D1 receptors might alleviate both the cognitive and the negative symptoms of schizophrenia (Goldman-Rakic et al. 2004, Miyamoto et al. 2005), by increasing NMDA receptor mediated synaptic currents (Figure 4.8). Atypical neuroleptics might use this mechanism by not only blocking D2 receptors, but also by increasing the presynaptic release of dopamine which in turn would increase the activation of the extrasynaptic D1 receptors (Castner et al. 2000, Moller 2005).

Taken together, we suggest that the cognitive and negative symptoms could be caused by the same synaptic mechanism, namely a reduction in the NMDA conductance, which reduces the stability and increases the distractibility of the persistent attractors, and reduces the activity (firing rates) of neurons (Figure 4.8, middle column). The reduced depth of the basins of attraction can be understood in the following way. (Hopfield 1982) showed that the recall state in an attractor network can be thought of as the local minimum in an energy landscape, where the energy would be defined as

$$E = -\frac{1}{2} \sum_{i,j} w_{ij} (y_i - \langle y \rangle)(y_j - \langle y \rangle) \quad (4.1)$$

where y_i and y_j are the firing rates of the i 'th and j 'th neurons in the network, which are connected by synaptic weight w_{ij} . In general neuronal systems do not admit an energy function.

Nevertheless, we can assume an effective energy function: in fact, the flow picture shown in Figure 4.1 resulting from the mean-field reduction associated with the spiking network here analyzed, can be viewed as an indirect description of an underlying effective energy function. From this equation, it follows that the depth of a basin of attraction is deeper if the firing rates are higher and if the synaptic strengths that couple the neurons that are part of the same attractor are strong. (The negative sign results in a low energy, and thus a stable state, if the firing rates of the neurons in the same attractor and their synaptic coupling weights are high.) If we reduce the NMDA receptor activated channel conductances, then the depth of the basins of attraction will be reduced both because the firing rates are reduced by reducing excitatory inputs to the neurons, and because the synaptic coupling weights are effectively reduced because the synapses can pass only reduced currents.

The positive symptoms (Figure 4.8, right column) of schizophrenia include delusions, hallucinations, thought disorder, and bizarre behavior. Examples of delusions are beliefs that others are trying to harm the person, impressions that others control the person's thoughts, and delusions of grandeur. Hallucinations are perceptual experiences, which are not shared by others, and are frequently auditory but can affect any sensory modality. These symptoms may be related to activity in the temporal lobes (Liddle 1987, Epstein et al. 1999, Mueser & McGurk 2004). The attractor framework approach taken in this paper hypothesizes that the basins of attraction of both spontaneous and persistent state are shallow (Figure 4.8). Due to the shallowness of the spontaneous state, the system can jump spontaneously up to a high activity state causing hallucinations to arise and leading to bizarre thoughts and associations. This might be the cause for the higher activations in temporal lobe areas which are identified in imaging experiments (Shergill et al. 2000, Scheuerecker et al. 2007).

We relate the positive symptoms to not only a reduction in NMDA conductance, but also to a reduction in GABA conductance. This is consistent with the fact that the positive symptoms usually follow the cognitive and negative ones and represent a qualitative worsening of the illness (Mueser & McGurk 2004). Alterations in GABA receptors have been identified in schizophrenia (Wang et al. 2004, Lewis et al. 2005).

D2 receptor antagonism remains a main target for antipsychotics (Seeman & Kapur 2000, Leuner & Muller 2006). Dopamine receptor D2 antagonists mainly alleviate the positive symptoms of schizophrenia, whereas the cognitive and negative symptoms persist, especially for the typical neuroleptics (Mueser & McGurk 2004). Together with the simulations, our hypothesis suggests that an increase in the GABA current in the state corresponding to the positive symptoms (-NMDA, -GABA) might have the same effect as D2 antagonists. The therapeutic effect of D2 antagonists might thus be caused by an increase in GABA currents. Indeed, it has been found that D2 receptors decrease the efficacy of the GABA system (Seamans et al. 2001, Trantham-Davidson et al. 2004). (For example, Seamans et al. (2001) found that the application of D2 antagonists prevented a decrease in eIPSC amplitude produced by dopamine.) Thus D2 antagonists would, in a hypersensitive D2 receptor state (Seeman et al. 2005, 2006), increase GABA inhibition in the network, and this we suggest could increase the stability of attractor networks involved in the positive symptoms of schizophrenia, and thus ameliorate the positive symptoms. Since the concentration of dopamine in the cortex depends on cortical-subcortical interactions (Carlsson 2006), the causes of the described changes could also result from subcortical deficits. A detailed analysis of these feedback loops would require specific modeling.

Earlier accounts of the relation of dopamine and schizophrenia in the cortex (Seamans et al. 2001, Seamans & Yang 2004) have suggested two distinct states of dopamine modulation.

One is a D2-receptor-dominated state in which there is weak gating and information can easily affect network activity. The other is a D1-receptor-dominated state in which network activity is stable and maintained. We have proposed a more detailed account for stability and discussed this separately for the spontaneous and persistent attractor states. This allows us to account for the dichotomy between the cognitive/negative and positive symptoms. We emphasize that, in biophysically realistic network simulations, excitation and inhibition are not merely antagonistic but implement different functions in the network dynamics. Thereby our modeling approach provides a missing link between the symptoms of schizophrenia and network models of working memory and dopamine (Durstewitz et al. 1999, 2000a, Brunel & Wang 2001). This approach is not meant to be a detailed discussion of specific symptoms of schizophrenia. Further research would be needed to relate specific symptoms to this general scheme presented here in this paper.

We concentrated in our study on the basic effects in local cortical dynamics based on biophysically realistic spiking networks. One can extend the model to study the interaction between large scale networks involving more than one location in the cortex and/or subcortical regions. This increases the dimensionality of the system and adds to the issues of stability (Kaneko 1997, 1998). Processes such as bubbling transitions and chaotic attractors also add to the notion of stability (Ashwin et al. 1994, 1996). For high-dimensional systems, Kaneko uses the concept of a Milnor attractor (Milnor 1985) which offers a way to formally describe the stability of an attractor. This approach measures the stability of an attractor against perturbations by introducing the concept of ‘return probability’. By sampling over random perturbation and orbit positions, the return probability is defined as the fraction of trials in which the system returns back to the original point. This measures the strength of the attractor. In our work, we wanted to stay close to paradigms of working memory and address mechanisms involved in cognitive processes, attention, and distractibility. Consequently, our quantitative measures of stability are intrinsically related to these paradigms and therefore utilize concepts such as the escape probability instead of the return probability. Indeed, we introduced these measures in order to relate the quantitative concept of stability more directly with the cognitive symptoms. Nevertheless, in a detailed investigations of specific symptoms of schizophrenia, the concepts mentioned above should be kept in mind as they might play a role in neural dynamics.

Cohen and collaborators focused in a series of computational and experimental studies on the cognitive symptoms of schizophrenia. Their connectionist models try to account for specific experimental data on context-processing deficits using working memory, gating and attentional selection as mechanisms (Servan-Schreiber et al. 1990, Cohen & Servan-Schreiber 1992, Braver et al. 1999). The context processing hypothesis is compatible with the hypothesis of working memory deficits in schizophrenia as context processing relies on the ability of stable working memory: a deficit in working memory would consequently lead to a deficit in context processing. Our model is conceptually at another level featuring biophysically plausible single neuron properties and specific synaptic currents. It is not intended to account for specific experimental data, but provide insights at a more generic and also biological level. In this sense, the two models complement each other.

We discussed a possible cause for the proposed alterations of the attractor landscape related to schizophrenia, namely changes in NMDA and GABA conductance as these are directly related to schizophrenia (Coyle et al. 2003, Lewis et al. 2005). We did not investigate changes in AMPA conductance. In this particular model the contribution of the AMPA current is relatively small (Brunel & Wang 2001). A more detailed investigation could also in-

clude AMPA conductance, especially because it is known to be influenced by NMDA synaptic plasticity (Bagal et al. 2005). Indeed, if reduced NMDA currents led in turn by synaptic plasticity to reduced AMPA currents, this would amplify the effects we describe. The proposed alterations in the attractor landscape could have a variety of causes at the neurobiological level: abnormalities in glutamate and GABA receptors and signaling, modulations in synaptic plasticity, aberrant dopamine signaling, reduced neuropil, genetic mechanisms, and brain volume reduction (Goldman-Rakic 1999, Winterer & Weinberger 2004, Mueser & McGurk 2004, Stephan et al. 2006). Besides cortical mechanisms, cortical-subcortical dynamics could also cause the proposed alterations in the cortical attractor landscape for example via neuromodulatory influences such as dopamine or serotonin or cortical-subcortical feedback loops (Capuano et al. 2002, Carlsson 2006). Our general hypothesis regarding the attractor landscape is meant to describe the aberrant dynamics in cortical regions which could be caused by several pathways. Future work could analyze further how changes of different factors such as regional differences, subcortical-cortical networks or even more detailed neural and synaptic models might influence the stability of the type of neurodynamical system described here. We envision that our hypothesis could serve as a useful guideline for further theoretical work.

4.4 Methods

4.4.1 Attractor Framework

Our aim is to investigate stability and distractibility in a biophysically realistic attractor framework, so that the properties of receptors, synaptic currents and the statistical effects related to the probabilistic spiking of the neurons can be part of the model. We use a minimal architecture, a single attractor or autoassociation network (Hopfield 1982, Amit 1989, Hertz et al. 1991, Rolls & Treves 1998, Rolls & Deco 2002). We chose a recurrent (attractor) integrate-and-fire network model which includes synaptic channels for AMPA, NMDA and GABA_A receptors (Brunel & Wang 2001). These synaptic receptors are important since there is evidence that alterations in synaptic currents are related to the symptoms of schizophrenia, and many of the drugs used to treat schizophrenia act on these receptor-mediated currents, either directly or indirectly (Coyle et al. 2003, Lewis et al. 2005).

Both excitatory and inhibitory neurons are represented by a leaky integrate-and-fire model (Tuckwell 1988). The basic state variable of a single model neuron is the membrane potential. It decays in time when the neurons receive no synaptic input down to a resting potential. When synaptic input causes the membrane potential to reach a threshold, a spike is emitted and the neuron is set to the reset potential at which it is kept for the refractory period. The emitted action potential is propagated to the other neurons in the network. The excitatory neurons transmit their action potentials via the glutamatergic receptors AMPA and NMDA which are both modeled by their effect in producing exponentially decaying currents in the postsynaptic neuron. The rise time of the AMPA current is neglected, because it is typically very short. The NMDA channel is modeled with an alpha function including both a rise and a decay term. In addition, the synaptic function of the NMDA current includes a voltage dependence controlled by the extracellular magnesium concentration (Jahr & Stevens 1990). The inhibitory postsynaptic potential is mediated by a GABA_A receptor model and is described by a decay term.

The single attractor network contains 400 excitatory and 100 inhibitory neurons, which is consistent with the observed proportions of pyramidal cells and interneurons in the cerebral




Figure 4.9: The attractor network model. The excitatory neurons are divided into two selective pools S1 and S2 (with 40 neurons each) with strong intra-pool connection strengths w_+ and one non-selective pool (NS) (with 320 neurons). The other connection strengths are 1 or weak w_- . The network contains 500 neurons, of which 400 are in the excitatory pools and 100 are in the inhibitory pool IH. The network also receives inputs from 800 external neurons, and these neurons increase their firing rates to apply a stimulus or distractor to one of the pools S1 or S2. The Appendix contains the synaptic connection matrices.

cortex (Abeles 1991, Braitenberg & Schütz 1991). The connection strengths are adjusted using meanfield analysis (Brunel & Wang 2001), so that the excitatory and inhibitory neurons exhibit a spontaneous activity of 3 Hz and 9 Hz, respectively (Wilson et al. 1994, Koch & Fuster 1989). The recurrent excitation mediated by the AMPA and NMDA receptors is dominated by the NMDA current to avoid instabilities during the delay periods (Wang 2002).

Our cortical network model features a minimal architecture to investigate stability and distractibility, and consists of two selective pools S1 and S2 (Figure 4.9). We use just two selective pools to eliminate possible disturbing factors. The non-selective pool NS models the spiking of cortical neurons and serves to generate an approximately Poisson spiking dynamics in the model (Brunel & Wang 2001), which is what is observed in the cortex. The inhibitory pool IH contains the 100 inhibitory neurons. The connection weights between the neurons of each pool or population are called the intra-pool connection strengths w_+ . The increased strength of the intra-pool connections is counterbalanced by the other excitatory connections (w_-) to keep the average input constant.

The network receives Poisson input spikes via AMPA receptors which are envisioned to originate from 800 external neurons at an average spontaneous firing rate of 3 Hz from each external neuron, consistent with the spontaneous activity observed in the cerebral cortex (Wilson et al. 1994, Rolls & Treves 1998). A detailed mathematical description is provided in the Appendix.

4.4.2 Analysis

Our analysis is targeted to investigate the stability and distractibility with respect to NMDA and GABA receptor modulations. We use two different techniques: Multiple trial spiking simulations and meanfield simulations. Spiking trial simulations integrate the complete neural

and synaptic dynamics over time including statistical components of the network model. Therefore the spiking simulations are needed to assess the stability and distractibility of the dynamical system, for this depends in part on the statistical fluctuations that occur in a network of spiking neurons (Deco & Rolls 2006). This is done by simulating a network configuration for several trials each run with different random seeds and running a statistical analysis on the data.

We simulate three different conditions: the spontaneous, persistent, and distractor conditions (see Figure 4.3).

In spontaneous simulations, we run spiking simulations for 3 s without any extra external input. The aim of this condition is to test whether the network is stable in maintaining a low average firing rate in the absence of any inputs, or whether it falls into one of its attractor states without any external input.

In persistent simulations, an external cue of 120 Hz above the background firing rate of 2400 Hz is applied to each neuron in pool S1 during the first 500 ms to induce a high activity state and then the system is run for another 2.5 s. The 2400 Hz is distributed across the 800 synapses of each S1 neuron for the external inputs, with the spontaneous Poisson spike trains received by each synapse thus having a mean rate of 3 Hz. The aim of this condition is to investigate whether once in an attractor short term memory state, the network can maintain its activity stably, or whether it falls out of its attractor, which might correspond to an inability to maintain attention.

The distractor simulations start off like the persistent simulations with a 500 ms input to pool S1 to start the S1 short term memory attractor states, but between 1 s and 1.5 s we apply a distracting input to pool S2 with varying strengths. The aim of this condition is to measure how distractible the network is. The degree of distractibility is measured parametrically by the strength of the input to S2 required to remove the high activity state of the S1 population. These simulation protocols serve to assess the generic properties of the dynamical attractor system rather than to model specific experimental data obtained in particular paradigms.

We used a mean-field approach (described in the Appendix) to calculate the stationary attractor states of the network for the delay period (Brunel & Wang 2001). These attractor states are independent of any simulation protocol of the spiking simulations and represent the behavior of the network by mean firing rates to which the system would converge in the absence of statistical fluctuations caused by the spiking of the neurons and by external changes. Therefore the mean-field technique is suitable for tasks in which temporal dynamics and fluctuations are negligible. It also allows a first assessment of the attractor landscape and the depths of the basin of attraction which then need to be investigated in detail with stochastic spiking simulations. Part of the utility of the mean-field approach is that it allows the parameter region for the synaptic strengths to be investigated to determine which synaptic strengths will on average produce stable activity in the network, for example of persistent activity in a delay period after the removal of a stimulus. For the spontaneous state, the initial conditions for numerical simulations of the mean-field method were set to 3 Hz for all excitatory pools and 9 Hz for the inhibitory pool. These values correspond to the approximate values of the spontaneous attractors when the network is not driven by stimulus-specific inputs. For the persistent state, a selective pool was set to a higher initial value (30 Hz) to account for the excitation of these neurons during the preceding cue period.

In addition, we used the meanfield technique to assess the flow between the attractor states. The flow is the force, which drives the system towards the attractor given a parameter value in phase space, i.e. the firing rates of the pools. Since we were interested in the depth

of a single attractor we used a setup with just one selective pool. This was done by fixing the value of the firing rate of the selective pool and letting the other values converge to their fixed point. Afterwards we computed the flow with this configuration (Mascaro & Amit 1999).

Acknowledgments

Support was provided by the Oxford McDonnell Centre for Cognitive Neuroscience.

Chapter 5

Reward-learning, Hebbian-learning, and Strategy

Manuscript in preparation.

Abstract

Conditional visuomotor learning is a paradigm in which stimulus-response associations are learned upon reward feedback. This complex learning task incorporates features of both instrumental learning and classical conditioning. Several learning properties have been identified in experiments. Based on reinforcement learning theory, we implement a computational model, which includes reward-based learning, Hebbian-based learning, error reset and forgetting mechanisms. In addition, we add a strategy component to our model, which carries out a repeat-stay/change-shift response strategy. Do different experiments in the field of visuomotor learning have similar contributions of these learning components? How do these components interact? We analyze in detail one conditional visuomotor experiment and show that it contains a strong Hebbian-based learning component. Reducing this component resembles the effects of hippocampal lesions. We show that adding the strategy component on top of the learning component can account for the degree of strategy usage. This yields insights in the interaction of these two mechanisms. In addition, we compare the resulting model to a different experiment and find inconsistencies in the behavioral data. We conclude that the exact experimental paradigm in complex learning tasks is crucial in assessing the results. Modeling provides a way to compare different experiments and helps to discuss at first sight inconsistent data.

5.1 Introduction

Conditional visuomotor learning is the ability of animals and humans to associate a visual stimulus with a motor response and learn these mappings by trial-and-error (Passingham 1993). In the experiments, a set of visual stimuli (typically 2 to 4 in primate studies) is associated with the same number of motor responses. A reward after the execution of the motor response signals if the given answer was correct. This paradigm has been studied using various experimental setups in neuropsychology, brain imaging and neurophysiology (for reviews, see Wise et al. (1996), Murray et al. (2000), Passingham et al. (2000), Hadj-Bouziane et al. (2003)). Not only the neural data measured in several brain areas such as the prefrontal cortex, premotor cortex, basal ganglia and hippocampus yield a very complex and diverse picture, also the behavioral data show an immense richness: Novel associations are processed differently than familiar ones. Brain lesion can affect either the learning of novel mappings (Murray & Wise 1996, Wise & Murray 1999, Wang et al. 2000, Brasted et al. 2003) or only the execution of familiar ones (Nixon et al. 2004). Response strategies such as the repeat-stay/change-shift or the win-stay/lose-shift pattern support the learning of associations in some but not all experiments (Wise & Murray 1999, Wang et al. 2000, Bussey et al. 2001, Genovesio et al. 2005). Errors have been reported to have specific effects from slowing down learning (Brasted et al. 2005) to a reset of learned associations (Fusi et al. 2007).

All these studies supposedly address the same learning ability, namely conditional visuomotor learning. How do the behavioral properties of the experiments compare? And if they do not, which different learning components are used to which extent in a particular experiment? Moreover, how do these different learning components interact and produce the behavioral data? These questions are not only interesting from a neuropsychological point of view to assess the behavioral data and the effect of lesions, but also for neurophysiological experiments. The setup of the experimental paradigm could determine which learning components are addressed and thereby influence the underlying neural dynamics.

By experimental means, these questions are difficult to address since the training of monkeys on different paradigms is demanding and even this training might influence the resulting behavior. Computational models might be ideal to give at least a first idea about the different learning components and influencing factors. They can be easily implemented for several paradigms and measures to compare experimental data. Manipulation of model parameters could account for the effects of lesions. We focus on the learning of novel associations. Based on classical reinforcement learning theory (Sutton & Barto 1998), we implement such a model and complement it with Hebbian learning (Hebb 1949, McClelland 2001) and strategy mechanisms. We start by fitting our model to experimental data by Brasted et al. (2003, 2005), which is a 3:3 mapping task and features non-spatial responses. Then we use this as a reference to compare it to two other sets of experiments: The experiments by Asaad et al. (1998), Pasupathy & Miller (2005), Fusi et al. (2007) implemented a 2:2 mapping task in which the associations were reversed after some time. The experiments by Murray & Wise (1996), Wise & Murray (1999) used a very similar paradigm than Brasted et al. (2003), but they clearly identified the usage of strategy in the behavioral data. We were interested if the characteristics of the experimental data can be fitted with the same computational model. Thereby we discuss contributions of the different learning components: Reward-based learning, Hebbian-based learning, error reset mechanism and forgetting. Our goal is to show the importance of the exact experimental paradigm and provide a way of a how to discuss different setups.

5.2 Results

Our model is based on a classical reinforcement learning theory and uses associative values $V_a(t)$ to denote the associative strength between a stimulus and a response. We assume that the associations of each stimulus are learned separately (Fusi et al. 2007). We found indicators for four learning mechanisms in the literature: reward-based learning, reset upon an error, Hebbian-based learning, and forgetting. These four mechanisms change the associative values $V_a(t)$ upon which the model chooses an action using a sigmoidal function such as also used in models of decision making. In the following, we fit the learning parameters to the data of Brasted et al. (2005) and study the dependencies among these mechanisms.

We discuss the experimental results of a mapping task by Brasted et al. (2005) in which 3 stimuli were mapped onto 3 responses with correction trials: The monkey saw a stimulus and then gave one of three possible responses. If the response was correct, a reward was given. If the response was incorrect, no reward was given and the same stimulus appeared on the next (correction) trial as long as the answer was correct. Figure 5.1a shows the average correct responses for the second presentation of a given stimulus, regardless of when it occurred during a session (Brasted et al. 2005). In the simplest case, the monkey responded by chance correctly on the first stimulus presentation, which marks an "errorless" trial. When the same stimulus was presented in the following trial again, there were "no intervening trials" between the two stimulus presentations. For this condition, the percentage of correct responses on the second presentation is given in Bar 1 in Figure 5.1a. If the monkey made an error on the first presentation, the trial would be in the "error" condition (Bar 5, Figure 5.1a). Note that the first trial always ended with a correct response due to the correction trials. Overall, errors slow down learning and worsen the performance of the second trial. Brasted et al. (2005) reasoned that Hebbian learning mechanisms might be responsible for that by strengthening the associations of executed actions even if they yielded an error (McClelland 2001). We implement this mechanism in a computational model and first discuss the two mechanisms, which are important for the learning of new associations: Reward-based learning and Hebbian learning. The reward based-learning increases the associative strengths between a stimulus and a response when a reward is received. The Hebbian learning mechanism increases the associative strengths when the response is executed regardless of the outcome. We set the error reset to a high value (0.7) as proposed by Fusi et al. (2007) and the forgetting to a low value (0.05). Since there are no established associations at the beginning, the forgetting value plays a minor role. Figure 5.2ab shows the dependency of the percentage of the second trial correct responses on the learning rates of these two mechanisms. We are looking for a high value of about 90% for the errorless case and a value of about 60% for the error case. This can be found in a region in with strong Hebbian learning ($\alpha_{hebb} = 0.4$) and low Reward based learning ($\alpha_{corr} = 0.1$). At first sight, this is surprising, as the Hebbian mechanism does not take the reward into account, which is crucial to learn the task. There is also a substantial learning when the positive reward learning is at 0. This is obvious for the errorless case (as no errors have been made). In the error case, the Hebbian learning strengthens all executed associations whereas the error reset eliminates the erroneously increased associations. This causes a positive performance, although no positive reward learning is present. We emphasize that we were not interested in exactly fitting the data, but in showing the dependencies of the parameters to present a complete picture of the dynamics.

As next step, we look at the dependency of the remaining two parameters: The error reset and the forgetting. The error reset sets all associations related to the stimulus to a

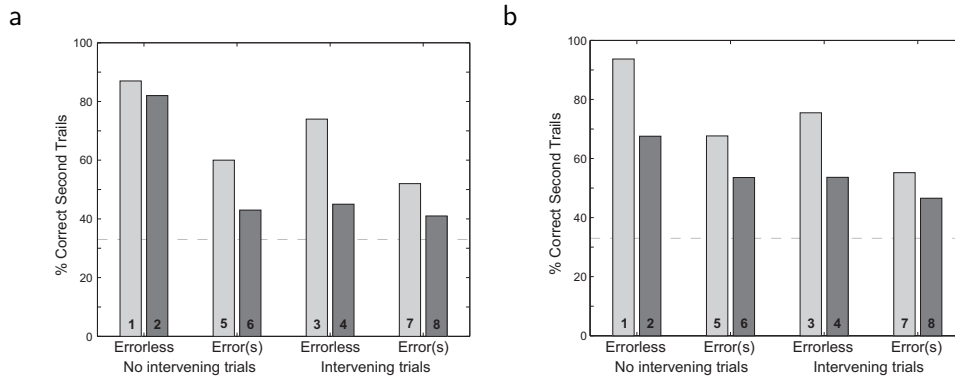


Figure 5.1: Average correct responses of the second presentation of a stimulus regardless when it occurs during a session and excluding correct trials. The trials can be categorized into two classes: first, if an error occurred during the first presentation of that stimulus (errorless / error trials) and second other stimuli were presented between the first and second, if presentation of a stimulus. (a) Experimental data by Brasted et al. (2005), see their Figure 3a (reproduced with authors' permission). The white bars denote the control condition the gray bars the fornix transection condition. (b) Model fit with parameters $\alpha_{corr} = 0.1$, $\alpha_{hebb} = 0.4$, $\alpha_{res} = 0.7$, and $\alpha_{forg} = 0.05$. In the gray bar, the Hebbian learning parameter is reduced ($\alpha_{hebb} = 0.1$).

lower value upon an error, whereas the forgetting gradually decays all associative strengths slowly in between every trial. Figure 5.2cd shows the performance difference between the no intervening trials and intervening trials case. This is about 10% for both cases in the experimental data (see Figure 5.1a, Bar 1–Bar 3, Bar 5–Bar 7). This difference depends mainly on the forgetting learning value which we set to $\alpha_{forg} = 0.05$. We set the error reset parameter α_{res} to a high value (0.7) as discussed by Fusi et al. (2007) in order to achieve the reset of the performance.

The model parameters are set with the performance of the second trial. Do these parameter also yield a comparable performance on a session basis? Brasted et al. (2003) assessed the performance on the session scope using the measure "errors to criterion". The criterion was 10 errors in a session of 100 trials (not counting correction trials). We computed the same measure for the model and found a value of 12.2 errors to criterion. This is in the range of the one observed in the experiment of 15 ± 4 errors (see Figures 3a, 8b, Brasted et al. (2003)). Thus, our model also reflects the experimental data on a session basis.

How does our model reflect the findings in other experiments? We check if the model indeed shows the reset properties described (Fusi et al. 2007) to maybe adjust the error reset parameter. The parameter, which is opposing the error reset, is the Hebbian learning parameter, because it also reinforces the stimulus-response association when an error is made. We computed the difference between the errorless and error case for the paper by Brasted et al. (2005) (which is achieved due to Hebbian learning) and the performance after an error for the experiment analyzed by Fusi et al. (2007). These two measures are contrasted in Figure 5.3. For achieving a difference between the errorless and the error case, a strong Hebbian learning and a weak error reset is required (Figure 5.3a). For achieving a reset after an error, we need a weak Hebbian learning and a very strong error reset value (Figure 5.3b): Strong Hebbian learning which slows down learning is incompatible with the error reset. If

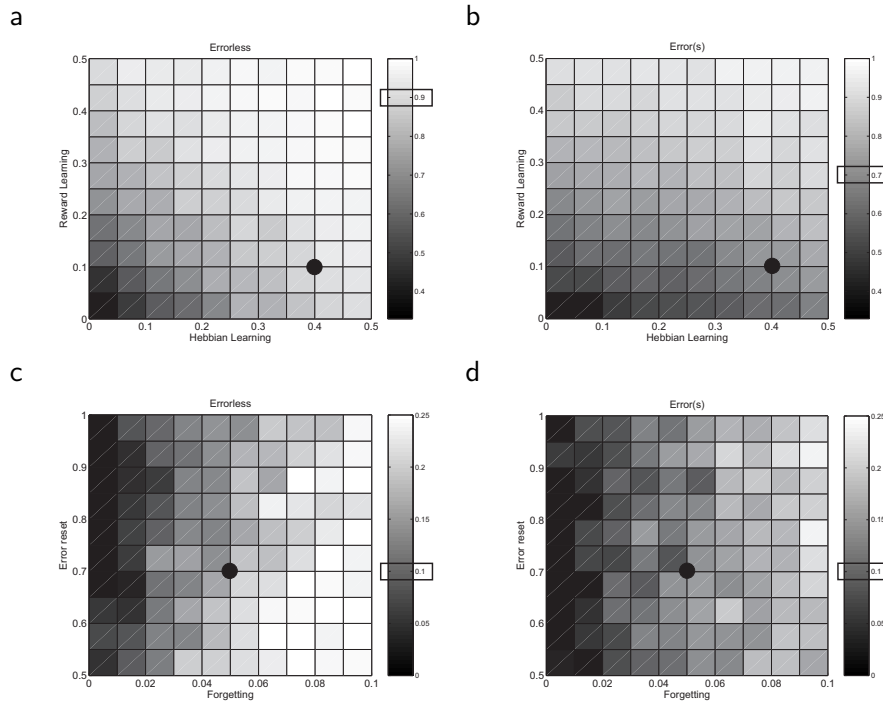


Figure 5.2: Overview of the dynamics of the model based on the performance of the second stimulus presentation (see Figure 5.1) as a function of the learning parameters. (a) Percentage correct of second stimulus presentation, errorless/no-intervening trials condition (Bar 1, Figure 5.1). (b) Percentage correct of second stimulus presentation, error(s)/no-intervening trials condition (Bar 5, Figure 5.1). (c,d) Difference of the no-intervening and intervening trials (c) for errorless condition (Bar 1–Bar 3, Figure 5.1) and (d) for the error condition (Bar 5–Bar 7, Figure 5.1). The dots denote the parameter values of the reference condition. The boxes at the color bars denote the values of the experimental data (Brasted et al. 2005).

an error reset would indeed govern the behavior of Brasted et al. (2005), the performance of the errorless and error case should be equal. However, they are significantly different. In conclusion, the experimental data from Brasted et al. (2005) and Fusi et al. (2007) show contradicting properties in their learning mechanisms. We will detail possible causes in the Discussion section.

While studying the model in relation to the experiment investigated by Fusi et al. (2007), we found that as in the experimental data, our model has a maximal performance of about 90% before the reset (Figure 5.4). We investigated this phenomenon further to see which parameters govern this behavior. We found that either strong Hebbian learning or a strong error reset yield a low maximal performance before the reset (data not shown). This shows that the underlying learning mechanisms themselves could be responsible for this low performance and not only distracting factors unrelated to the experiment such as proposed and implemented by Fusi et al. (2007).

The main focus of the studies by Brasted et al. (2003, 2005) were the effects of lesions of the hippocampal system. They performed a fornix transection, which impaired the learning of novel associations. The hippocampal formation has been related to the fast learning of associations. It has been hypothesized that it mainly works as a pattern associator which

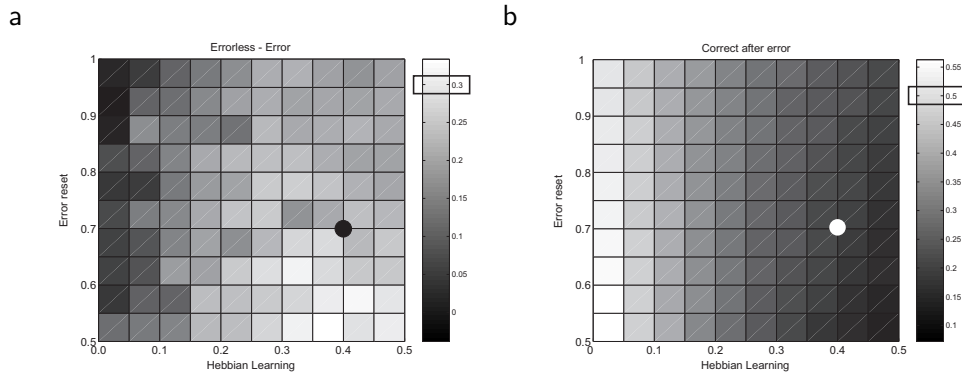


Figure 5.3: Comparison of measures of Brasted et al. (2005) and Fusi et al. (2007) using our model in dependence on the Hebbian learning parameter α_{hebb} and the error reset parameter α_{res} . (a) Difference between the errorless and error condition (no-intervening trials) of second trial performance (Bar 1–Bar 5, Figure 5.1). It should be around 30% according to the experimental data. (b) Performance after an error computed with the paradigm investigated by Fusi et al. (2007). It should be at 50%. The boxes at the color bars indicate the values according to the experimental data. The dots denote the parameter values of the reference condition.

binds items such as stimuli and responses (McClelland et al. 1995, McClelland 2001). This is reflected by the Hebbian learning component of the model. Does a reduction in the Hebbian learning component in the model reflect the effects of the fornix transection in the experimental data? We reduce the Hebbian learning parameter to 25% from 0.4 to 0.1 to account for the fornix transection. Figure 5.1 shows the comparison of the experimental data and the model for the performance of the second trial. We find a good fit with the experimental data (except in the condition with no errors and no intervening trials). In addition, we computed the errors to criterion as above. We found a value of around 140 errors which is higher than for the trials 18 months after surgery (70 ± 8 , Brasted et al. (2003)), but still in the range of the values which have been observed during the whole experiment (see Figure 3a, Brasted et al. (2003)). The absence of response strategies in our model might partially account for the differences (Brasted et al. 2005). In terms of learning, a reduction of Hebbian learning component can account for large parts of the effects of a fornix transaction, which implies that a lesion of the hippocampus could eliminate much of the associative capability of learning.

Earlier experiments (Murray & Wise 1996, Wise & Murray 1999) have shown that monkeys spontaneously adopt response strategies to solve this task. Specifically, they used a repeat-stay/change-shift strategy: If the stimulus repeated in the next trial and the previous answer was correct, the monkeys could stay with the previously given response (repeat-stay). If the stimulus changed, then the same response given on the previous trial would not yield a reward as each response was always mapped to only one stimulus. Thus one of the other responses would be rewarded (change-shift). If a trial did not end with a correct response, the monkeys employed a loose-shift response in which they avoided a previously erroneous response. Due to the usage of correction trials (in this experiment only up to 2 correction trials were given), the monkeys almost always (between 92-97%) ended a trial including correction trials with a correct answer (Wise & Murray 1999). The low error rate of on average 5% is even lower than expected by applying the loose-shift strategy perfectly (7.4%), which suggests that the




Figure 5.4: Behavioral performance at the reversal using the experimental paradigm by Asaad et al. (1998) and the model with reference parameters. As in the experimental data, there is a plateau in the maximal performance which is caused by the learning dynamics (here primarily due to the Hebbian learning component (data not shown)). Note the lack of the reset after the error caused by the reversal. (see Figure 2b, Asaad et al. (1998)).

monkeys also kept track of the response two stimuli before. In our model, we implemented a loose-shift strategy, which yielded by interference with the learning component and the probabilistic decision making a final error rate of 9.0%. This is close to the experimental data. The two mechanisms, the learning and the strategy, conflict each other in different parts of the experiment: At the beginning, the strategy usage would yield a higher reward and speed up learning, since it can be used without any knowledge of the specific mappings. In case of a 3:3 mapping, the strategy alone can yield a maximal performance of 66% (100% in repeat trials and 50% in change trials). Learning is needed to improve the performance further. How do these two mechanisms interact? In principle, there are two possibilities: First, the additional information given by the strategy component is added on top of the learning process, e.g., due to higher input to the decision making module. Or second, there is a certain trade off between strategy and learning component and once there is strategy information the influence of the learning component is reduced. We used that the experiments by Murray & Wise (1996), Wise & Murray (1999) and Brasted et al. (2003, 2005) are very similar – except that in the first set of experiments a repeat-stay/change-shift strategy was adopted by the monkeys and in the latter at least not "unambiguously" (Brasted et al. 2003). Since the overall performance of the two experiments was comparable (see Figure 5, Brasted et al. (2005)), we took the fit of the learning component to Brasted et al. (2003, 2005) above and investigated what happens if learning and strategy are combined. We implemented the strategy component by associative values $S_a(t)$: In a repeat-stay trial, $S_a(t)$ of the previously executed action a is 1, the other actions 0. In a change-shift-trial, $S_a(t)$ of the previously executed action is 0, the other actions 1. If an error trial occurred and the stimulus stayed the same, $S_a(t)$ of the previously executed action is 0, the other actions 1 (loose-shift). The associative strength of the learning $V_a(t)$ and the strategy $S_a(t)$ are added by $lc \cdot V_a(t) + sc \cdot S_a(t)$. This value is entered in the decision making module (see Figure 5.7, Methods section).

Figure 5.5 shows both the strategy score and the performance at the end of a session for different contributions of the strategy component on top of the learning component ($lc = 1$).




Figure 5.5: Strategy score as a function of the contribution of the strategy component sc on top of the learning component ($lc = 1$). The strategy score is defined by $ss = c - r/0.5$, where c is the percentage of errors in change trials and r in repeat trials. 0.5 is the difference when the repeat-stay/change-shift strategy is applied perfectly (Bussey et al. 2001). In addition, we picture the maximal performance (average of trials 40-50) and the strategy score by Murray & Wise (1996).

The strategy score has been introduced by Bussey et al. (2001). It reflects the degree of strategy usage during the first 8 trials of a session. It is 1 when the strategy is applied perfectly and 0 when there is no difference in performance between repeat and change trials. Could the strategy component just be added on top of the learning component? In Figure 5.5, the strategy score starts at around 0.15 without the strategy component ($sc = 0$) and then reaches a plateau of around 0.55 at $sc > 0.3$. Interestingly, the experiment by Murray & Wise (1996) has also a strategy score of around 0.55 (see Figure 4, Bussey et al. (2001)). This could indicate that the strategy component is mostly added on top of the learning component. If the contribution of the learning component is decreased ($lc < 1$), then the strategy score increases while the performance decreases (data not shown). Computationally, different and competing evidences might add up at the decision making module which selects the action (see Figure 5.7). The contributions between the learning and the strategy component might also change during learning. Here we just considered a fixed ratio, which is sufficient to account for the experimental data. Genovesio et al. (2005) did neurophysiological measurements while monkeys performed a repeat-stay/change-shift task, which we will discuss for the neural correlate in the Discussion section.

Finally, we consider the case in which a repeat-stay/change-shift strategy applies and the learning is reduced by a hippocampal lesion (Murray & Wise 1996, Wise & Murray 1999). We simulate this condition by applying the strategy component ($sc = 0.3$) and reduce at the same time the Hebbian learning component as above ($\alpha_{hebb} = 0.1$). Figure 5.6 shows the comparison of the model and the experiment. There is a significant difference in the repeat trials in the experimental data before and after the lesion, but not in the model. Thus the lesions of the hippocampus also reduces the performance of the repeat trials, which could either be explained by a general deficit after the lesion or a minor contribution of the hippocampus to the strategy processing. On the other hand, the change trials reach almost the level of the repeat trials in the experiment for the pre-operation condition, whereas the model performance of the change trials of the normal condition has not reached the level of

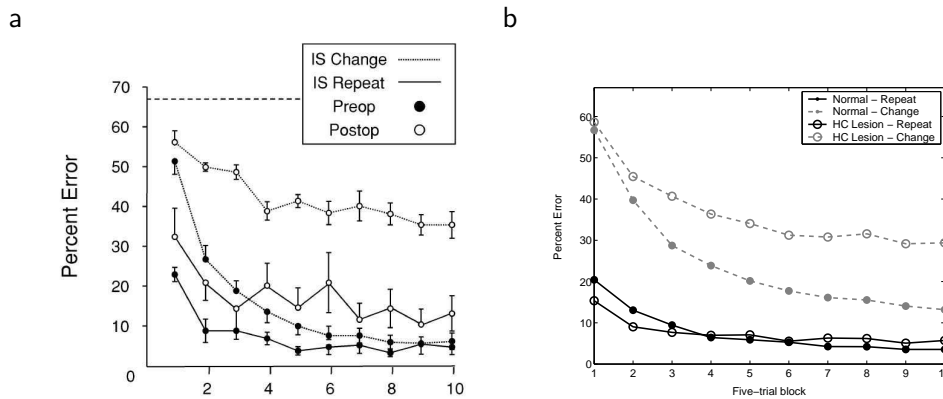


Figure 5.6: Percentage of errors for the paradigm by Murray & Wise (1996) separated by repeat/change trials and control/hippocampal lesion conditions. (a) Experimental data (Wise & Murray (1999), reproduced with authors' permission) (b) Model data using strategy component ($sc = 0.3$), normal: reference condition, hc lesion: $\alpha_{hebb} = 0.1$.

the repeat trials after 50 trials. Either stronger learning parameters or a weaker strategy component could resolve this issue. Overall, our model captures both the higher performance of the repeat trials and the different learning curves of the change trials in the control and hippocampal lesion condition.

5.3 Discussion

We have constructed a behavioral model for the learning of arbitrary visuomotor mappings. We fitted the parameters using data by Brasted et al. (2003, 2005) and compared the resulting model to the work by Murray & Wise (1996), Wise & Murray (1999) and Asaad et al. (1998), Pasupathy & Miller (2005), Fusi et al. (2007). We found that the Hebbian learning and error reset mechanisms are incompatible based on our current model. In our model, we assume that the error reset and the Hebbian learning, the two opposing mechanisms, are independent. They change the same associative strength $V_a(t)$ based on its previous value. Other possibilities of implementation are that either the learning rules interact in one trial, e.g. by acting in a particular order, or the two mechanisms are completely independent and act on different associative values. If the two mechanisms act consecutively, the Hebbian learning should be first as the execution of a stimulus-response association is before the reward is delivered. After an error, a previous strengthening due to the Hebbian mechanism would be eliminated: Errors would not slow down learning as in Brasted et al. (2005). If the two mechanisms would act on different connections, there would be no interaction between the mechanisms. The error reset would just reset the connections, which have been built up due to the reward learning. The Hebbian learning would remain untouched and errors would slow down learning: There would be no error reset. Thus, although the type of the interaction might change, there remains in any case a conflict between the Hebbian and error reset mechanism. How could the different experimental results be explained?

Our simulations suggests that the primates use different learning mechanisms in the two experiments by Brasted et al. (2003) and Asaad et al. (1998). This could have a variety of reasons. First, the paradigms in the two experiments are different: A 3:3 mapping with un-

limited correction trials versus a 2:2 mapping interleaved with familiar mappings and reversals after reaching a high performance. Since there were only two stimulus-response associations in the experiment by Asaad et al. (1998), there could have been less interference due to Hebbian learning. Second, there was a delay period between the stimulus and the response to measure the activity of neurons related to working memory in the paradigm by Asaad et al. (1998), whereas the stimulus was displayed continuously until the response was given in Brasted et al. (2003). This could also increase the distance between stimulus and response in Asaad et al. (1998) and thereby decrease the effect of Hebbian learning. Third, the type of the response was different in both experiments (non-spatial touching response vs. saccadic eye movements). Forth, the interleaving familiar stimuli could have obstructed the learning in Asaad et al. (1998). The 2:2 mapping experiment could have been solved perfectly by the described repeat-stay/change-shift strategy (without the interleaving trials). Nevertheless, the reset after an error could also be viewed as a strategy caused by the reversals in the experiments. Since errors often marked the complete change in the associations (reversals), the monkeys could have been sensitized to errors, which then let them to reset their current knowledge (since the more elaborate strategy to switch the associations was not used (Deco & Rolls 2005*b*), this is the next best strategy). The repeat-stay/change-shift strategy was identified in the experiment by Murray & Wise (1996), but not in Brasted et al. (2003) although the experiments are very similar. The main differences are the type of the response (joystick movements vs. non-spatial touching response) and the number of correction trials (2 vs. unlimited). Thus, various kinds of differences in the experiments on top of individual difference of the primates might impose the usage of different learning components. Further analysis would be needed to assess which of the factors described above lead to the usage of a particular learning component.

We reasoned that the Hebbian learning component might be associated with the Hippocampal formation. The role of the hippocampal formation in visuomotor learning has been addressed by many articles (Murray & Wise 1996, Wise & Murray 1999, Brasted et al. 2003, 2005). Several possibilities have been discussed: General mapping function over all domains, working memory contribution, relational scene, episodic memory, spatial and idiothetic memory. Especially, the spatial/non-spatial contribution of the hippocampus has been addressed by Brasted et al. (2003, 2005). We do not detail these hypotheses further as they have been discussed in the articles mentioned above. However, we would like to add another possibility, which has not been looked at so far: In our computational model, we could reproduce the experimental data quite well by reducing the Hebbian learning component to 25%. This still leaves a residual ability of learning arbitrary visuomotor mappings functional, which is in line with experimental data (Wise & Murray 1999). This ability is envisioned to be carried out by the basal ganglia, which is also related to the earliest phases of learning (Pasupathy & Miller 2005) and the prefrontal cortex. We reason that the hippocampus might be a major part of the brain pattern-associator network. In addition, we show that the Hebbian component might be used to various degrees in different experiments, since it is even incompatible with the experimental results by Asaad et al. (1998), Fusi et al. (2007). In conclusion, the controversial results of lesions of the hippocampal system (for the sake of discussion we do not make a difference between aspiration lesions of the hippocampus and fornix-transections) might be caused by differences in the experimental paradigms, which address the Hebbian learning component to different degrees. If the Hebbian learning component was used to a small extent in an experiment, then lesions of the hippocampal system would also have a small effect. As discussed above, small differences in the experimental paradigm can lead

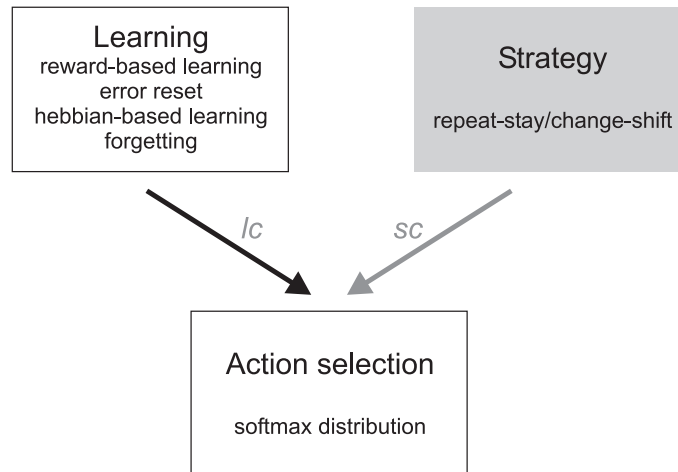


Figure 5.7: Overview of the model components. The learning component is made up of four rules which adjust the associative strength $V_a(t)$ during learning. The strategy component uses a memory of the last executed trial to generate temporary associative strengths $S_a(t)$ for the repeat-stay/change-shift behavior. The strategy component is just used to model the experiment by Murray & Wise (1996) showed such a pattern. The decision-making components select an action upon incoming evidence of the learning (and strategy) component.

to a different behavior. Most prominently, mapping tasks using either spatial (Rupniak & Gaffan 1987) or non-spatial consequents (Gaffan & Harrison 1988) yielded different effects on the fornix-transection. We hypothesize that the non-spatial version of the experiments might have addressed the Hebbian component to a smaller extent and thereby the fornix-transection had no effect on the behavior. The performance after error trials such as discussed by Brasted et al. (2005) could give an estimate about the involvement of the Hebbian learning component. Thus, we suggest an alternative view on the role of the hippocampus, which can be further tested in the experimental data.

Along the with role of the Hippocampus, the neurophysiological basis of visuomotor learning in general is one of the main objectives of the field (Murray et al. 2000, Hadj-Bouziane et al. 2003). Our schematic model yields some constraints and ideas about the underlying neural mechanisms. The model components are pictured in Figure 5.7. As already mentioned, the Hippocampus might be responsible for the Hebbian learning and especially for the fast novel learning, which is reflected in the strong contribution in relation to the reward-based component. This is also backed up by neurophysiological recordings of the hippocampus (Cahusac et al. 1993, Wirth et al. 2003), which identified selective neural activity even before learning occurred. The reward-based learning component might be associated with the basal ganglia and the prefrontal cortex. Reward signals of dopaminergic neurons have been identified in the basal ganglia (Tremblay et al. 1998, Schultz 2000) which have also been related to reinforcement learning theory (Schultz 2006) such as implemented in our model. Several experiments have identified a pivotal role of the basal ganglia (Hadj-Bouziane et al. 2003) in learning both due to the involvement early during learning (Pasupathy & Miller 2005) and the neural activity shortly before the reward period (Buch et al. 2006). However, following Houk & Wise (1995), Pasupathy & Miller (2005), the actual learning might take place in the prefrontal cortex which is connected by loops to the basal ganglia. The basal ganglia

might have a leading role in learning and thereby "train" the cortical networks. We envision a similar mechanism for the interaction between the hippocampus and the prefrontal cortex (McClelland et al. 1995): The hippocampus is connected via the fornix to the cortex and might thereby train the prefrontal networks. Cortical regions might be the point of integration for the different evidences which are generated by the basal ganglia, the hippocampus and the cortical regions themselves (Miller & Cohen 2001). This integration is crucial for our model, since we investigate the interaction between the Hebbian-based learning and the reward-based learning. This interaction is especially important for the error reset mechanism (Fusi et al. 2007), since it supposedly resets the learned connections of the whole system, which contains the basal ganglia, hippocampus, premotor cortex and prefrontal cortex. In addition, an interaction and integration also occurs between the learning components and the strategy. These also have conflicting characteristics, since at the beginning of a trial the strategy itself will yield a higher performance than chance. However, after learning, the learning component performs better. Neural recordings (Genovesio et al. 2005) and lesion studies (Wang et al. 2000, Bussey et al. 2001) suggest the prefrontal cortex as a possible source of the signals which lead to the rule-based behavior (see later in this section). Finally, all the different evidences must be integrated in a decision-making component which then selects stochastically an action upon the present evidences. We envision that two areas could play this role, maybe in conjunction. The basal ganglia is in the optimal position to integrate various evidences as it receives input from almost every cortical region (Houk & Wise 1995). In addition, the binary characteristic of the striatal spiny neurons in the striatum could facilitate the decision-making. A second important region for the decision-making component could be the premotor cortex. Learning related activity has been identified in the premotor cortex (Mitz et al. 1991, Brasted & Wise 2004, Buch et al. 2006) and the supplementary eye field (Chen & Wise 1995*a,b*). In addition, it contains a large portion of neurons, which directly reflect the executed action. As the premotor cortex and the basal-ganglia are strongly interconnected and have a similar learning related activity (Brasted & Wise 2004), there is probably not only a single region responsible for the decision making. Neuropsychological work has also identified the basal ganglia and the premotor cortex as crucial for the execution of familiar associations (Nixon et al. 2004) which fits well in this picture: The associations might be automatized within the decision making component which is trained by the flexible learning system made up of basal ganglia, prefrontal cortex and hippocampus. Modeling work has shown that the decision making process can be distributed among several processing layers (Deco & Rolls 2003, Loh & Deco 2005). Given the interplay of several cognitive processes, an interaction at several stages seems likely. In addition, such a processing pathway might be reflected in time by an activity wave from sensory to motor areas (Buch et al. 2006, Loh et al. 2008). Overall, the fast learning components might be guided by the basal ganglia and the hippocampus and converge in prefrontal regions, which might also generate the strategy signals. An integration of all signals could occur in a system made up of the basal ganglia and the premotor cortex.

The strategy signal was implemented in the model by adding an associative strength to the related associations. In specific, we implemented a repeat-stay/change-shift strategy, if a reward was received on the previous trials, and a loose-shift strategy, if no reward was received. What is the neurophysiological correlate of this process? Genovesio et al. (2005, 2006) identified neurons which are selective to either the repeat trials or change trials. In order to apply a repeat-stay/change-shift strategy correctly, one has to integrate the previous and present stimulus information to identify the trial type (repeat/change) and then select

upon the previous action either a stay or shift action. In terms of neural circuitry, this computation is very complex and will be pursued in future modeling work. In our model, we directly implemented the outcome of the computation and increased the associative strength of the actions corresponding to the strategy. We discussed how different contributions of the learning and strategy component relate to the strategy score. We found that in our model, adding the strategy signal on top of the learning signals yields a comparable strategy score as in the experiment. On this basis, we conclude that the contributions of the strategy component can be static during the learning of novel associations at least within a session. However, Genovesio et al. (2005) found little strategy related activity during the execution of a familiar mapping task, which suggests that at least for highly familiar associations the strategy component contributes little and is down-regulated.

Fusi et al. (2007) presented a model which combines both a neural decision making model and a reinforcement learning based learning role. They fitted a simple softmax function such as used in our model to the probabilistic output pattern of a two-pool decision-making model made up a biophysically plausible neural network. In specific, they showed that the conductance difference between the inputs to the two pools is the crucial parameter for the decision making outcome. They related the associative strengths (our $V_a(t)$) to the proportion of potentiated synapses to a neural synapses. Thus, they provide the connection of the type of model used in our work to the underlying neurophysiological correlate. In addition, they combined both novel and familiar associations in a model based on the idea of learning on different time scales (Fusi et al. 2005). We focused in more detail on the learning of novel associations and split this part into several sub-components: Reward-based learning, Hebbian-based learning, error reset and forgetting. In addition, we expanded our model to more than two actions and added a strategy component. Overall, the two modeling approaches complement each other. A similar model for conditional visuomotor learning, also based on reinforcement learning theory, was presented by Hadj-Bouziane et al. (2006). They focused on a specific effect found in a 4:4 mapping experiment in which the stimuli were presented in a particular order: The associations were learned sequentially and there seemed to an exponential increase in the learning speed as a function of the number of stimuli learned. They connected this effect to the viability theory of economics and a dimension reduction in learning. A detailed analysis of our model in respect to their data could be interesting.

Overall, we proposed a behavioral model for learning of novel associations and fitted it to the experimental data by Brasted et al. (2005). Thereafter we compared the model to evidences by Fusi et al. (2007) and Murray & Wise (1996). We found inconsistencies in the published data, and it will require a further analysis of the experimental data to identify the causes. Furthermore, we showed that the hippocampus might drive Hebbian-based stimulus-response learning and that the strategy component probably implemented in pre-frontal circuits might add evidences on top of the learning signal to the decision-making areas. The model itself stands as a specific prediction and can easily be implemented and compared to other experimental paradigms. We envision that it can serve as a reference to compare different experimental setups. Differences in the learning properties caused by different experimental paradigms could address learning components to various degrees. Getting an insight in this scheme might be important to discuss the underlying neural correlate in both neurophysiology and neuropsychology.




Figure 5.8: Probability of an action a to be selected $P_a(t)$ in dependence on its associative strength $V_a(t)$ for different numbers of actions following Equation 5.1. We assume that the $V_b(t)$ of all other actions b are 0.

5.4 Methods

The basic model for learning stimulus-response-reward associations contains two components: A decision-making component, and a learning component.

The decision-making or action selection component selects an action upon the learned associative strength for each stimulus. We assume as a first approximation that all stimuli are learned independently (Fusi et al. 2007). The model description stands for the associations of each stimulus separately and thus we do not write the index of the stimulus explicitly. We use a softmax distribution to convert the associative strength into probabilities of actions, a standard method in reinforcement learning (Sutton & Barto 1998):

$$P_a(t) = \frac{\exp(\tau V_a(t))}{\sum_{\forall b} \exp(\tau V_b(t))}, \quad (5.1)$$

where $P_a(t)$ is the probability of action a , τ is the temperature of the softmax or Boltzman distribution, and $V_a(t)$ the associative strength. Fusi et al. (2007) fitted such a function to a biophysically realistic neural model of decision making for two decisions. They use a network model of integrate-and-fire neurons and two bistable populations to discuss the action selection behavior of the network upon changing input. We use $\tau = 0.15$ and generalize for more than two possible actions. Figure 5.8 shows the probability of the selection of an action as a function of the associative strength $V_a(t)$ for different numbers of actions (all other action values are assumed to be 0).

The learning component modifies the associative strength $V_a(t)$ upon which the decision making component generates the actions. We implement four learning rules: Reward-based learning, Hebbian-based learning, error reset and forgetting. After the reward is received, the value of the associative strength associated with the presented stimulus is changed by $V_a(t + \frac{1}{2}) = V_a(t) + c_{rew} + c_{hebb}$, where c_{rew} represents the reward-based learning or the error reset in case of a correct or an error trial, respectively. c_{hebb} reflects the Hebbian learning component. The summation of these values assumes different processes which act on the same set of synapses and thereby allows for an interaction of the learning types. This is necessary especially for the error reset, which is assumed to set the system back to an equilibrium state.

First we define the reward depending learning rules: In the case of a correct trial (i.e. when a reward is received), the associative strength of the executed action a is changed by

$c_{rew} = \alpha_{corr}(1 - V_a(t))$, where 1 is the maximal associative strengths (Rescorla & Wagner 1972, Sutton & Barto 1998). In case of an error trial (i.e. when no reward is received), the associative strength of all actions of the presented stimulus are changed by $c_{rew} = -\alpha_{res}V_a(t)$, for all a . This reset after an error has been identified by Fusi et al. (2007) in data of an associative visuomotor mapping experiment (Asaad et al. 1998, Pasupathy & Miller 2005). Second, we assume that the executed stimulus-response mappings are strengthened independently of the reward due to a Hebbian mechanism: The associative strength of each action is changed by $c_{hebb} = \alpha_{hebb}(1 - V_a(t))$. This strengthening is independently of the reward and also applies to error trials.

We chose to include the forgetting term separately as the forgetting supposedly takes place in the time between the trials. It decays the associative strengths exponentially between the trials: $c_{forg} = -\alpha_{forg}V_a(t)$. Thus, the associative strength of the next trial is computed by $V_a(t + 1) = V_a(t + \frac{1}{2}) + c_{forg}$.

In neural models, the associative strengths $V_a(t)$ can be viewed as the fraction of potentiated synapses which associate a stimulus with a response. An increase in $V_a(t)$ could reflect an increase in the strong synaptic connections between stimulus and response representations. This increase in the fraction of potentiated synapses might result in a higher current input to the decision-making module upon stimulus presentation.

In addition to the learning model, we include a strategy module for experiments in which a repeat-stay/change-shift behavior has been observed (Wise & Murray 1999, Murray et al. 2000). We implement this by associative values $S_a(t)$, which are increased or decreased upon the last stimulus-response association if the last trial ended with a correct answer: In a repeat trial (i.e. the stimulus is the same as in the previous trial), $S_a(t)$ of the last action a is 1 whereas the value of the other actions is 0. In a change trial (i.e. the stimulus is different from the previous trial), the $S_a(t)$ of the last action a is 0, whereas the values of the other actions are 1. If a trial ended with an incorrect answer and the stimulus stayed the same, we implemented a loose-shift strategy: The $S_a(t)$ of the last action a is 0, whereas the value of the other actions is 1. To combine the strategy with the learning module, we add the values of $S_a(t)$ and $V_a(t)$ by $A_a(t) = lc \cdot V_a(t) + sc \cdot S_a(t)$ for all a , whereas lc is the pre-factor of the learning component and sc of the strategy component. The $A_a(t)$ values are entered in the decision making module as above. We discuss the pre-factors in our analysis. An overview of the model components is pictured in Figure 5.7.

We discuss in this paper the experimental paradigms by Murray & Wise (1996), Brasted et al. (2002) and Asaad et al. (1998) (only first publications of the paradigm cited). The paradigm by Brasted et al. (2002) uses a 3:3 mapping with unlimited correction trials. We chose the order of the stimuli pseudorandomly and compute 100 blocks of 5000 trials (the 5000 trials are used for the "errors to criterion" measure. The paradigm by Murray & Wise (1996) used also a 3:3 mapping but only two correction trials. Here we compute 1000 blocks of 50 trials. The paradigm by Asaad et al. (1998) used 2:2 mappings with no correction trails but both familiar and novel mappings. First, we trained our model on a familiar mappings followed by 50 sessions. In these 50 sessions, the two familiar stimuli were presented pseudorandomly in 1/3 of the cases and the novel stimuli in 2/3 of the cases. During a session the novel mappings were reversed after 100 presentations of the novel stimuli. We computed 1000 trials/9 reversals per session.

Appendix A

Neural Network Model

A.1 Neural and synaptic dynamics

We use the mathematical formulation of the integrate-and-fire neurons and synaptic currents described in Brunel & Wang (2001). Here we provide a brief summary of this framework.

The dynamics of the sub-threshold membrane potential V of a neuron are given by the equation:

$$C_m \frac{dV(t)}{dt} = -g_m(V(t) - V_L) - I_{syn}(t), \quad (\text{A.1})$$

Both excitatory and inhibitory neurons have a resting potential $V_L = -70mV$, a firing threshold $V_{thr} = -50mV$ and a reset potential $V_{reset} = -55mV$. The membrane parameters are different for both types of neurons: Excitatory (Inhibitory) neurons are modeled with a membrane capacitance $C_m = 0.5nF(0.2nF)$, a leak conductance $g_m = 25nS(20nS)$, a membrane time constant $\tau_m = 20ms(10ms)$, and a refractory period $t_{ref} = 2ms(1ms)$. Values are extracted from McCormick et al. (1985).

When the threshold membrane potential V_{thr} is reached, the neuron is set to the reset potential V_{reset} at which it is kept for a refractory period τ_{ref} and the action potential is propagated to the other neurons.

The network is fully connected with $N_E = 400$ excitatory neurons and $N_I = 100$ inhibitory neurons, which is consistent with the observed proportions of the pyramidal neurons and interneurons in the cerebral cortex (Braitenberg & Schütz 1991, Abeles 1991). The synaptic current impinging on each neuron is given by the sum of recurrent excitatory currents ($I_{AMPA,rec}$ and $I_{NMDA,rec}$), the external excitatory current ($I_{AMPA,ext}$) the inhibitory current (I_{GABA}):

$$I_{syn}(t) = I_{AMPA,ext}(t) + I_{AMPA,rec}(t) + I_{NMDA,rec}(t) + I_{GABA}(t). \quad (\text{A.2})$$

The recurrent excitation mediated by the AMPA and NMDA receptors, inhibition by GABA receptors. In addition, the neurons are exposed to external Poisson input spike trains mediated by AMPA receptors at a rate of 2.4kHz. These can be viewed as originating from $N_{ext} = 800$ external neurons at average rate of 3Hz per neuron, consistently with the spontaneous activity observed in the cerebral cortex (Wilson et al. 1994, Rolls & Treves 1998). The currents are defined by:

$$I_{AMPA,ext}(t) = g_{AMPA,ext}(V(t) - V_E) \sum_{j=1}^{N_{ext}} s_j^{AMPA,ext}(t) \quad (A.3)$$

$$I_{AMPA,rec}(t) = g_{AMPA,rec}(V(t) - V_E) \sum_{j=1}^{N_E} w_{ji}^{AMPA} s_j^{AMPA,rec}(t) \quad (A.4)$$

$$I_{NMDA,rec}(t) = \frac{g_{NMDA}(V(t) - V_E)}{1 + [Mg^{++}] \exp(-0.062V(t))/3.57} \times \sum_{j=1}^{N_E} w_{ji}^{NMDA} s_j^{NMDA}(t) \quad (A.5)$$

$$I_{GABA}(t) = g_{GABA}(V(t) - V_I) \sum_{j=1}^{N_I} w_{ji}^{GABA} s_j^{GABA}(t) \quad (A.6)$$

where $V_E = 0$ mV, $V_I = -70$ mV, w_j are the synaptic weights, s_j 's the fractions of open channels for the different receptors and g 's the synaptic conductances for the different channels. The NMDA synaptic current depends on the membrane potential and the extracellular concentration of Magnesium ($[Mg^{++}] = 1$ mM, Jahr & Stevens (1990)). The values for the synaptic conductances for excitatory neurons are $g_{AMPA,ext} = 2.08$ nS, $g_{AMPA,rec} = 0.208$ nS, $g_{NMDA} = 0.654$ nS and $g_{GABA} = 2.50$ nS and for inhibitory neurons $g_{AMPA,ext} = 1.62$ nS, $g_{AMPA,rec} = 0.162$ nS, $g_{NMDA} = 0.516$ nS and $g_{GABA} = 1.946$ nS. These values are obtained from the ones used by Brunel & Wang (2001) by correcting for the different numbers of neurons. The conductances were calculated so that in an unstructured network the excitatory neurons have a spontaneous spiking rate of 3 Hz and the inhibitory neurons a spontaneous rate of 9 Hz. The fractions of open channels are described by:

$$\frac{ds_j^{AMPA,ext}(t)}{dt} = -\frac{s_j^{AMPA,ext}(t)}{\tau_{AMPA}} + \sum_k \delta(t - t_j^k) \quad (A.7)$$

$$\frac{ds_j^{AMPA,rec}(t)}{dt} = -\frac{s_j^{AMPA,rec}(t)}{\tau_{AMPA}} + \sum_k \delta(t - t_j^k) \quad (A.8)$$

$$\frac{ds_j^{NMDA}(t)}{dt} = -\frac{s_j^{NMDA}(t)}{\tau_{NMDA,decay}} + \alpha x_j(t)(1 - s_j^{NMDA}(t)) \quad (A.9)$$

$$\frac{dx_j(t)}{dt} = -\frac{x_j(t)}{\tau_{NMDA,rise}} + \sum_k \delta(t - t_j^k) \quad (A.10)$$

$$\frac{ds_j^{GABA}(t)}{dt} = -\frac{s_j^{GABA}(t)}{\tau_{GABA}} + \sum_k \delta(t - t_j^k), \quad (A.11)$$

where $\tau_{NMDA,decay} = 100$ ms is the decay time for NMDA synapses, $\tau_{AMPA} = 2$ ms for AMPA synapses (Hestrin et al. 1990, Spruston et al. 1995) and $\tau_{GABA} = 10$ ms for GABA synapses (Salin & Prince 1996, Xiang et al. 1998); $\tau_{NMDA,rise} = 2$ ms is the rise time for NMDA synapses (the rise times for AMPA and GABA are neglected because they are typically very short) and $\alpha = 0.5$ ms⁻¹. The sums over k represent a sum over spikes formulated as δ -Peaks $\delta(t)$ emitted by presynaptic neuron j at time t_j^k .

The equations were integrated numerically using a second order Runge-Kutta method with step size 0.02 ms. The Mersenne Twister algorithm was used as random number generator for the external Poisson spike trains and different trials for equal parameter configurations were run with different random seeds (as the only difference).

A.2 Mean-field Formulation

The meanfield approximation used in the present work was derived by Brunel & Wang (2001), assuming that the network of IF neurons is in a stationary state. In this formulation the potential of a neuron is calculated as:

$$\tau_x \frac{dV(t)}{dt} = -V(t) + \mu_x + \sigma_x \sqrt{\tau_x} \eta(t) \quad (\text{A.12})$$

where $V(t)$ is the membrane potential, x labels the populations. The symbol τ_x is the effective membrane time constant, μ_x the mean value the membrane potential would have in the absence of spiking and fluctuations, σ_x measures the magnitude of the fluctuations and η is a Gaussian process with absolute exponentially decaying correlation function with time constant τ_{AMPA} . The quantities μ_x and σ_x^2 are given by:

$$\mu_x = \frac{(T_{ext}\nu_{ext} + T_{AMPA}n_x^{AMPA} + \rho_1 n_x^{NMDA})V_E + \rho_2 n_x^{NMDA}\langle V \rangle + T_I n_x^{GABA}V_I + V_f}{S_x} \quad (\text{A.13})$$

$$\sigma_x^2 = \frac{g_{AMPA,ext}^2 (\langle V \rangle - V_E)^2 N_{ext}\nu_{ext}\tau_{AMPA}^2 \tau_x}{g_m^2 \tau_m^2}. \quad (\text{A.14})$$

where $\nu_{ext} = 3$ Hz ($+\lambda_{cue}$), ν_I is the spiking rate of the inhibitory pool, $\tau_m = C_m/g_m$ with the values for the excitatory or inhibitory neurons depending of the pool considered and the other quantities are given by:

$$S_x = 1 + T_{ext}\nu_{ext} + T_{AMPA}n_x^{AMPA} + (\rho_1 + \rho_2)n_x^{NMDA} + T_I n_x^{GABA} \quad (\text{A.15})$$

$$\tau_x = \frac{C_m}{g_m S_x} \quad (\text{A.16})$$

$$n_x^{AMPA} = \sum_{j=1}^p f_j w_{jx}^{AMPA} \nu_j \quad (\text{A.17})$$

$$n_x^{NMDA} = \sum_{j=1}^p f_j w_{jx}^{NMDA} \psi(\nu_j) \quad (\text{A.18})$$

$$n_x^{GABA} = \sum_{j=1}^p f_j w_{jx}^{GABA} \nu_j \quad (\text{A.19})$$

$$\psi(\nu) = \frac{\nu \tau_{NMDA}}{1 + \nu \tau_{NMDA}} \left(1 + \frac{1}{1 + \nu \tau_{NMDA}} \sum_{n=1}^{\infty} \frac{(-\alpha \tau_{NMDA, rise})^n T_n(\nu)}{(n+1)!} \right) \quad (\text{A.20})$$

$$T_n(\nu) = \sum_{k=0}^n (-1)^k \binom{n}{k} \frac{\tau_{NMDA, rise} (1 + \nu \tau_{NMDA})}{\tau_{NMDA, rise} (1 + \nu \tau_{NMDA}) + k \tau_{NMDA, decay}} \quad (\text{A.21})$$

$$\tau_{NMDA} = \alpha \tau_{NMDA, rise} \tau_{NMDA, decay} \quad (\text{A.22})$$

$$T_{ext} = \frac{g_{AMPA, ext} \tau_{AMPA}}{g_m} \quad (\text{A.23})$$

$$T_{AMPA} = \frac{g_{AMPA, rec} N_E \tau_{AMPA}}{g_m} \quad (\text{A.24})$$

$$\rho_1 = \frac{g_{NMDA} N_E}{g_m J} \quad (\text{A.25})$$

$$\rho_2 = \beta \frac{g_{NMDA} N_E (\langle V_x \rangle - V_E) (J - 1)}{g_m J^2} \quad (\text{A.26})$$

$$J = 1 + \gamma \exp(-\beta \langle V_x \rangle) \quad (\text{A.27})$$

$$T_I = \frac{g_{GABA} N_I \tau_{GABA}}{g_m} \quad (\text{A.28})$$

$$\langle V_x \rangle = \mu_x - (V_{thr} - V_{reset}) \nu_x \tau_x, \quad (\text{A.29})$$

where p is the number of excitatory pools, f_x the fraction of neurons in the excitatory x pool, $w_{j,x}$ the weight of the connections from pool x to pool j , ν_x the spiking rate of the x excitatory pool, $\gamma = [Mg^{++}]/3.57$ and $\beta = 0.062$.

The spiking rate of a pool as a function of the defined quantities is then given by:

$$\nu_x = \phi(\mu_x, \sigma_x), \quad (\text{A.30})$$

where

$$\phi(\mu_x, \sigma_x) = \left(\tau_{rp} + \tau_x \int_{\beta(\mu_x, \sigma_x)}^{\alpha(\mu_x, \sigma_x)} du \sqrt{\pi} \exp(u^2) [1 + \text{erf}(u)] \right)^{-1} \quad (\text{A.31})$$

$$\alpha(\mu_x, \sigma_x) = \frac{(V_{thr} - \mu_x)}{\sigma_x} \left(1 + 0.5 \frac{\tau_{AMPA}}{\tau_x} \right) + 1.03 \sqrt{\frac{\tau_{AMPA}}{\tau_x}} - 0.5 \frac{\tau_{AMPA}}{\tau_x} \quad (\text{A.32})$$

$$\beta(\mu_x, \sigma_x) = \frac{(V_{reset} - \mu_x)}{\sigma_x} \quad (\text{A.33})$$

where $\text{erf}(u)$ the error function and τ_{rp} the refractory period which is considered to be 2 ms for excitatory neurons and 1 ms for inhibitory neurons. To solve the equations defined by (A.30) for all x we integrate numerically (A.29) and the differential equation below, which has fixed point solutions corresponding to equations A.30:

$$\tau_x \frac{d\nu_x}{dt} = -\nu_x + \phi(\mu_x, \sigma_x). \quad (\text{A.34})$$

The equations were integrated using the Euler method with step size 0.2 and 8000 iterations, which allowed for convergence.

Appendix B

Connection Matrices for Chapter 2

Fraction of pool sizes f_i

Values are relative to all neurons, not only excitatory the portion.

A	B	AL	BR	AR	BL	L	R	NS	IH
0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.04	0.48	0.2

Connection matrix for AMPA and NMDA – [from, to]

	A	B	AL	BR	AR	BL	L	R	NS	IH
A	w_+	w_-	w_{oi}	w_-	w_{oi}	w_-	w_-	w_-	1	1
B	w_-	w_+	w_-	w_{oi}	w_-	w_{oi}	w_-	w_-	1	1
AL	w_{io}	w_-	w_+	w_-	w_-	w_-	w_{ip}	w_-	1	1
BR	w_-	w_{io}	w_-	w_+	w_-	w_-	w_-	w_{ip}	1	1
AR	w_{io}	w_-	w_-	w_-	w_+	w_-	w_-	w_{ip}	1	1
BL	w_-	w_{io}	w_-	w_-	w_-	w_+	w_{ip}	w_-	1	1
L	w_-	w_-	w_{pi}	w_-	w_-	w_{pi}	w_+	w_-	1	1
R	w_-	w_-	w_-	w_{pi}	w_{pi}	w_-	w_-	w_+	1	1
NS	w_n	w_n	w_n	w_n	w_n	w_n	w_n	w_n	1	1
IH	0	0	0	0	0	0	0	0	0	0

Appendix C

Connection Matrices for Chapter 3

Fraction of pool sizes f_i

Values are relative to all neurons, not only excitatory the portion.

AL	BR	AR	BL	NS	IH
0.08	0.08	0.08	0.08	0.48	0.2

Connection matrix for AMPA and NMDA – [from, to]

	AL	BR	AR	BL	NS	IH
AL	w_+	w_m	w_a	w_a	1	1
BR	w_m	w_+	w_a	w_a	1	1
AR	w_a	w_a	w_+	w_m	1	1
BL	w_a	w_a	w_m	w_+	1	1
NS	w_a	w_a	w_a	w_a	1	1
IH	0	0	0	0	0	0

Connection matrix for GABA – [from, to]

	AL	BR	AR	BL	NS	IH
AL	0	0	0	0	0	0
BR	0	0	0	0	0	0
AR	0	0	0	0	0	0
BL	0	0	0	0	0	0
NS	0	0	0	0	0	0
IH	1	1	1	1	1	1

Appendix D

Connection Matrices for Chapter 4

Fraction of pool sizes f_i

Values are relative to all neurons, not only the excitatory portion.

S1	S2	NS	IH
0.08	0.08	0.64	0.2

Connection matrix for AMPA and NMDA – [from, to]

	S1	S2	NS	IH
S1	w_+	w_-	1	1
S2	w_-	w_+	1	1
NS	w_-	w_-	1	1
IH	0	0	0	0

where $w_- = \frac{0.8 - f_{S1}w_+}{0.8 - f_{S1}}$.

Connection matrix for GABA – [from, to]

	S1	S2	NS	IH
S1	0	0	0	0
S2	0	0	0	0
NS	0	0	0	0
IH	1	1	1	1

Bibliography

- Abeles, M. (1991), *Corticonics*, Cambridge University Press, New York.
- Aleman, A. & Kahn, R. S. (2005), ‘Strange feelings: do amygdala abnormalities dysregulate the emotional brain in schizophrenia?’, *Prog Neurobiol* **77**(5), 283–298.
- Amit, D. J. (1989), *Modeling Brain Function. The World of Attractor Neural Networks.*, Cambridge University Press, Cambridge.
- Asaad, W., Rainer, G. & Miller, E. K. (1998), ‘Neural activity in the primate prefrontal cortex during associative learning’, *Neuron* **21**, 1399–1407.
- Asaad, W., Rainer, G. & Miller, E. K. (2000), ‘Task-specific neural activity in the primate prefrontal cortex’, *Journal of Neurophysiology* **84**, 451–459.
- Ashwin, P., Buesch, J. & Stewart, I. (1994), ‘Bubbling of attractors and synchronization of chaotic attractors.’, *Physics Letters A* **193**, 126–139.
- Ashwin, P., Buesch, J. & Stewart, I. (1996), ‘From attractor to chaotic saddle: A tale of transverse stability.’, *Nonlinearity* **9**, 703–737.
- Bagal, A. A., Kao, J. P. Y., Tang, C.-M. & Thompson, S. M. (2005), ‘Long-term potentiation of exogenous glutamate responses at single dendritic spines.’, *Proc Natl Acad Sci U S A* **102**(40), 14434–9.
- Bak, P. & Chialvo, D. R. (2001), ‘Adaptive learning by extremal dynamics and negative feedback’, *Phys Rev E Stat Nonlin Soft Matter Phys* **63**, 031912.
- Bar-Gad, I., Morris, G. & Bergman, H. (2003), ‘Information processing, dimensionality reduction and reinforcement learning in the basal ganglia’, *Prog Neurobiol* **71**(6), 439–473.
- Baxter, R. D. & Liddle, P. F. (1998), ‘Neuropsychological deficits associated with schizophrenic syndromes’, *Schizophrenia Research* **30**(3), 239–49.
- Bender, W., Albus, M., Moller, H. J. & Tretter, F. (2006), ‘Towards systemic theories in biological psychiatry’, *Pharmacopsychiatry* **39 Suppl 1**, S4–9.
- Boettiger, C. A. & D’Esposito, M. (2005), ‘Frontal networks for learning and executing arbitrary stimulus-response associations’, *J Neurosci* **25**(10), 2723–2732.
- Boussaoud, D. & Kermadi, I. (1997), ‘The primate striatum: neuronal activity in relation to spatial attention versus motor preparation’, *Eur J Neurosci* **9**(10), 2152–2168.

- Braitenberg, V. & Schütz, A. (1991), *Anatomy of the Cortex*, Springer Verlag, Berlin.
- Brasted, P. J., Bussey, T. J., Murray, E. A. & Wise, S. P. (2002), 'Fornix transection impairs conditional visuomotor learning in tasks involving nonspatially differentiated responses', *J Neurophysiol* **87**(1), 631–3.
- Brasted, P. J., Bussey, T. J., Murray, E. A. & Wise, S. P. (2003), 'Role of the hippocampal system in associative learning beyond the spatial domain', *Brain* **126**(Pt 5), 1202–23.
- Brasted, P. J., Bussey, T. J., Murray, E. A. & Wise, S. P. (2005), 'Conditional motor learning in the nonspatial domain: effects of errorless learning and the contribution of the fornix to one-trial learning', *Behav Neurosci* **119**(3), 662–676.
- Brasted, P. J. & Wise, S. P. (2004), 'Comparison of learning-related neuronal activity in the dorsal premotor cortex and striatum', *Eur J Neurosci* **19**(3), 721–740.
- Braver, T. S., Barch, D. M. & Cohen, J. D. (1999), 'Cognition and control in schizophrenia: a computational model of dopamine and prefrontal function.', *Biol Psychiatry* **46**(3), 312–28.
- Brunel, N. & Wang, X. J. (2001), 'Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition', *Journal of Computational Neuroscience* **11**, 63–85.
- Buch, E. R., Brasted, P. J. & Wise, S. P. (2006), 'Comparison of population activity in the dorsal premotor cortex and putamen during the learning of arbitrary visuomotor mappings', *Exp Brain Res* **169**(1), 69–84.
- Bussey, T. J., Wise, S. P. & Murray, E. A. (2001), 'The role of ventral and orbital prefrontal cortex in conditional visuomotor learning and strategy use in rhesus monkeys (*Macaca mulatta*)', *Behav Neurosci* **115**(5), 971–82.
- Bussey, T. J., Wise, S. P. & Murray, E. A. (2002), 'Interaction of ventral and orbital prefrontal cortex with inferotemporal cortex in conditional visuomotor learning', *Behav Neurosci* **116**(4), 703–715.
- Cahusac, P. M., Rolls, E. T., Miyashita, Y. & Niki, H. (1993), 'Modification of the responses of hippocampal neurons in the monkey during the learning of a conditional spatial response task', *Hippocampus* **3**(1), 29–42.
- Canavan, A. G., Nixon, P. D. & Passingham, R. E. (1989), 'Motor learning in monkeys (*Macaca fascicularis*) with lesions in motor thalamus.', *Exp Brain Res* **77**(1), 113–26.
- Capuano, B., Crosby, I. T. & Lloyd, E. J. (2002), 'Schizophrenia: genesis, receptorology and current therapeutics', *Curr Med Chem* **9**(5), 521–548.
- Carlsson, A. (2006), 'The neurochemical circuitry of schizophrenia', *Pharmacopsychiatry* **39** Suppl 1, S10–4.
- Carter, C. S., Perlstein, W., Ganguli, R., Brar, J., Mintun, M. & Cohen, J. D. (1998), 'Functional hypofrontality and working memory dysfunction in schizophrenia', *Am J Psychiatry* **155**(9), 1285–1287.

- Castner, S. A., Williams, G. V. & Goldman-Rakic, P. S. (2000), 'Reversal of antipsychotic-induced working memory deficits by short-term dopamine D1 receptor stimulation', *Science* **287**(5460), 2020–2022.
- Chen, L. L. & Wise, S. P. (1995a), 'Neuronal activity in the supplementary eye field during acquisition of conditional oculomotor associations', *J Neurophysiol* **73**(3), 1101–1121.
- Chen, L. L. & Wise, S. P. (1995b), 'Supplementary eye field contrasted with the frontal eye field during acquisition of conditional oculomotor associations', *J Neurophysiol* **73**(3), 1122–1134.
- Chen, L. L. & Wise, S. P. (1996), 'Evolution of directional preferences in the supplementary eye field during acquisition of conditional oculomotor associations.', *J Neurosci* **16**(9), 3067–81.
- Chialvo, D. R. & Bak, P. (1999), 'Learning from mistakes', *Neuroscience* **90**(4), 1137–1148.
- Cohen, J. D. & Servan-Schreiber, D. (1992), 'Context, cortex, and dopamine: a connectionist approach to behavior and biology in schizophrenia', *Psychol Rev* **99**(1), 45–77.
- Coyle, J. T., Tsai, G. & Goff, D. (2003), 'Converging evidence of NMDA receptor hypofunction in the pathophysiology of schizophrenia', *Ann N Y Acad Sci* **1003**, 318–327.
- Deco, G. (2006), 'A dynamical model of event-related fmri signals in prefrontal cortex: predictions for schizophrenia', *Pharmacopsychiatry* **39 Suppl 1**, S65–S67.
- Deco, G. & Rolls, E. T. (2003), 'Attention and working memory: A dynamical model of neuronal activity in the prefrontal cortex', *European Journal of Neuroscience* **18**, 2374–2390.
- Deco, G. & Rolls, E. T. (2005a), 'Attention, short term memory, and action selection: a unifying theory', *Progress in Neurobiology* **76**, 236–256.
- Deco, G. & Rolls, E. T. (2005b), 'Synaptic and spiking dynamics underlying reward reversal in the orbitofrontal cortex', *Cereb Cortex* **15**(1), 15–30.
- Deco, G. & Rolls, E. T. (2006), 'A neurophysiological model of decision-making and Weber's law', *European Journal of Neuroscience* **24**, 901–916.
- Deco, G., Rolls, E. T. & Horwitz, B. (2004), 'What' and 'where' in visual working memory: a computational neurodynamical perspective for integrating fMRI and single-neuron data', *Journal of Cognitive Neuroscience* **16**, 683–701.
- Deiber, M. P., Wise, S. P., Honda, M., Catalan, M. J., Grafman, J. & Hallett, M. (1997), 'Frontal and parietal networks for conditional motor learning: a positron emission tomography study', *J Neurophysiol* **78**(2), 977–991.
- Djurfeldt, M., Ekeberg, Ö. & Graybiel, A. M. (2001), 'Cortex-basal ganglia interaction and attractor states', *Neurocomputing* **38–40**, 573–579.
- Dominey, P. F., Arbib, M. A. & Joseph, J. P. (1995), 'A model of corticostriatal plasticity for learning associations and sequences', *J. Cog. Neurosci* **7**, 311–336.

- Dominey, P. F. & Boussaoud, D. (1997), 'Encoding behavioral context in recurrent networks of the fronto-striatal system: a simulation study', *Brain Res Cogn Brain Res* **6**(1), 53–65.
- Durstewitz, D., Kelc, M. & Gunturkun, O. (1999), 'A neurocomputational theory of the dopaminergic modulation of working memory functions', *J Neurosci* **19**(7), 2807–2822.
- Durstewitz, D. & Seamans, J. K. (2002), 'The computational role of dopamine D1 receptors in working memory', *Neural Netw* **15**(4-6), 561–572.
- Durstewitz, D., Seamans, J. K. & Sejnowski, T. J. (2000a), 'Dopamine-mediated stabilization of delay-period activity in a network model of prefrontal cortex', *J Neurophysiol* **83**(3), 1733–1750.
- Durstewitz, D., Seamans, J. K. & Sejnowski, T. J. (2000b), 'Neurocomputational models of working memory', *Nat Neurosci* **3 Suppl**, 1184–91.
- Eacott, M. J. & Gaffan, D. (1992), 'Inferotemporal-frontal disconnection: The uncinate fascicle and visual associative learning in monkeys', *Eur J Neurosci* **4**(12), 1320–1332.
- Epstein, J., Stern, E. & Silbersweig, D. (1999), 'Mesolimbic activity associated with psychosis in schizophrenia. Symptom-specific PET studies', *Ann N Y Acad Sci* **877**, 562–574.
- Fagg, A. H. & Arbib, M. A. (1992), 'A model of primate visual-motor conditional learning', *Adaptive Behavior* **1**, 1–37.
- Fusi, S., Asaad, W. F., Miller, E. K. & Wang, X. J. (2007), 'A neural circuit model of flexible sensorimotor mapping: learning and forgetting on multiple timescales', *Neuron* **54**(2), 319–33.
- Fusi, S., Drew, P. J. & Abbott, L. F. (2005), 'Cascade models of synaptically stored memories', *Neuron* **45**(4), 599–611.
- Fuster, J. & Alexander, G. (1971), 'Neuron activity related to short-term memory', *Science* **173**, 652–654.
- Gaffan, D. & Harrison, S. (1988), 'Inferotemporal-frontal disconnection and fornix transection in visuomotor conditional learning by monkeys', *Behav Brain Res* **31**(2), 149–163.
- Gaffan, D. & Harrison, S. (1989), 'A comparison of the effects of fornix transection and sulcus principalis ablation upon spatial learning by monkeys', *Behav Brain Res* **31**(3), 207–220.
- Genovesio, A., Brasted, P. J., Mitz, A. R. & Wise, S. P. (2005), 'Prefrontal cortex activity related to abstract response strategies', *Neuron* **47**(2), 307–20.
- Genovesio, A., Brasted, P. J. & Wise, S. P. (2006), 'Representation of future and previous spatial goals by separate neural populations in prefrontal cortex.', *J Neurosci* **26**(27), 7305–16.
- Goff, D. C. & Coyle, J. T. (2001), 'The emerging role of glutamate in the pathophysiology and treatment of schizophrenia', *Am J Psychiatry* **158**(9), 1367–1377.
- Goldman-Rakic, P. (1994), 'Working memory dysfunction in schizophrenia', *Journal of Neuropsychology and Clinical Neuroscience* **6**, 348–357.

- Goldman-Rakic, P. S. (1999), 'The physiological approach: functional architecture of working memory and disordered cognition in schizophrenia', *Biol Psychiatry* **46**(5), 650–661.
- Goldman-Rakic, P. S., Castner, S. A., Svensson, T. H., Siever, L. J. & Williams, G. V. (2004), 'Targeting the dopamine D1 receptor in schizophrenia: insights for cognitive dysfunction', *Psychopharmacology (Berl)* **174**(1), 3–16.
- Graybiel, A. M. (1998), 'The basal ganglia and chunking of action repertoires', *Neurobiol Learn Mem* **70**(1-2), 119–136.
- Green, M. F. (1996), 'What are the functional consequences of neurocognitive deficits in schizophrenia?', *Am J Psychiatry* **153**(3), 321–30.
- Hadj-Bouziane, F. & Boussaoud, D. (2003), 'Neuronal activity in the monkey striatum during conditional visuomotor learning', *Exp Brain Res* **153**(2), 190–196.
- Hadj-Bouziane, F., Frankowska, H., Meunier, M., Coquelin, P.-A. & Boussaoud, D. (2006), 'Conditional visuo-motor learning and dimension reduction', *Cogn Process* **7**(2), 95–104.
- Hadj-Bouziane, F., Meunier, M. & Boussaoud, D. (2003), 'Conditional visuo-motor learning in primates: a key role for the basal ganglia', *J Physiol Paris* **97**(4-6), 567–579.
- Hafner, H., Maurer, K., Loffler, W., an der Heiden, W., Hambrecht, M. & Schultze-Lutter, F. (2003), 'Modeling the early course of schizophrenia', *Schizophr Bull* **29**(2), 325–340.
- Halsband, U. & Passingham, R. (1982), 'The role of premotor and parietal cortex in the direction of action', *Brain Res* **240**(2), 368–372.
- Hebb, D. (1949), *The Organization of Behavior - A Neurophysiological Theory*, John Wiley, New York.
- Hertz, J., Krogh, A. & Palmer, R. G. (1991), *Introduction to the Theory of Neural Computation*, Addison Wesley, Wokingham, U.K.
- Hestrin, S., Sah, P. & Nicoll, R. (1990), 'Mechanisms generating the time course of dual component excitatory synaptic currents recorded in hippocampal slices', *Neuron* **5**, 247–253.
- Hopfield, J. J. (1982), 'Neural networks and physical systems with emergent collective computational abilities', *Proc. Nat. Acad. Sci. USA* **79**, 2554–2558.
- Houk, J. C. & Wise, S. P. (1995), 'Distributed modular architectures linking basal ganglia, cerebellum, and cerebral cortex: their role in planning and controlling action', *Cereb Cortex* **5**(2), 95–110.
- Inase, M., Li, B. M., Takashima, I. & Iijima, T. (2001), 'Pallidal activity is involved in visuomotor association learning in monkeys', *Eur J Neurosci* **14**(5), 897–901.
- Inase, M., Li, B. M., Takashima, I. & Iijima, T. (2006), 'Cue familiarity is represented in monkey medial prefrontal cortex during visuomotor association learning.', *Exp Brain Res* **168**(1-2), 281–6.

- Inase, M., Tokuno, H., Nambu, A., Akazawa, T. & Takada, M. (1999), 'Corticostriatal and corticosubthalamic input zones from the presupplementary motor area in the macaque monkey: comparison with the input zones from the supplementary motor area', *Brain Res* **833**(2), 191–201.
- Ingvar, D. H. & Franzen, G. (1974), 'Abnormalities of cerebral blood flow distribution in patients with chronic schizophrenia', *Acta Psychiatrica Scandinavica* **50**(4), 425–462.
- Jahr, C. & Stevens, C. (1990), 'Voltage dependence of NMDA-activated macroscopic conductances predicted by single-channel kinetics', *Journal of Neuroscience* **10**, 3178–3182.
- Kaneko, K. (1997), 'Dominance of minor attractors and noise-induced selection in a multi-attractor system.', *Physical Review Letters* **78**, 2736–2739.
- Kaneko, K. (1998), 'On the strength of attractors in a high-dimensional system.', *Physica D* **124**, 322–344.
- Kircher, T. T. & Thienel, R. (2005), 'Functional brain imaging of symptoms and cognition in schizophrenia', *Prog Brain Res* **150**, 299–308.
- Koch, K. W. & Fuster, J. M. (1989), 'Unit activity in monkey parietal cortex related to haptic perception and temporary memory', *Exp. Brain Res.* **76**, 292–306.
- Koechlin, E., Ody, C. & Kouneiher, F. (2003), 'The architecture of cognitive control in the human prefrontal cortex', *Science* **302**(5648), 1181–1185.
- Kubota, K. & Niki, H. (1971), 'Prefrontal cortical unit activity and delayed alternation performance in monkeys', *J Neurophysiol* **34**(3), 337–347.
- Leuner, K. & Muller, W. E. (2006), 'The complexity of the dopaminergic synapses and their modulation by antipsychotics', *Pharmacopsychiatry* **39 Suppl 1**, S15–20.
- Lewis, D. A., Hashimoto, T. & Volk, D. W. (2005), 'Cortical inhibitory neurons and schizophrenia', *Nat Rev Neurosci* **6**(4), 312–324.
- Liddle, P. F. (1987), 'The symptoms of chronic schizophrenia: a re-examination of the positive-negative dichotomy', *British Journal of Psychiatry* **151**, 145–151.
- Lieberman, J. A., Perkins, D., Belger, A., Chakos, M., Jarskog, F., Boteva, K. & Gilmore, J. (2001), 'The early stages of schizophrenia: speculations on pathogenesis, pathophysiology, and therapeutic approaches.', *Biol Psychiatry* **50**(11), 884–97.
- Loh, M. & Deco, G. (2005), 'Cognitive flexibility and decision making in a model of conditional visuomotor associations', *European Journal of Neuroscience* **22**(11), 2927–2936.
- Loh, M., Pasupathy, A., Miller, E. K. & Deco, G. (2008), 'Neurodynamics of the prefrontal cortex during conditional visuomotor associations', *Journal of Cognitive Neuroscience* (in press).
- Loh, M., Rolls, E. T. & Deco, G. (2007), 'A dynamical systems hypothesis of schizophrenia', *PLoS Computational Biology* **3**(11), e228.

- Malhotra, A. K., Pinals, D. A., Weingartner, H., Sirocco, K., Missar, C. D., Pickar, D. & Breier, A. (1996), 'NMDA receptor function and human cognition: the effects of ketamine in healthy volunteers', *Neuropsychopharmacology* **14**(5), 301–307.
- Mascaro, M. & Amit, D. J. (1999), 'Effective neural response function for collective population states', *Network* **10**(4), 351–373.
- Mattia, M. & Del Giudice, P. (2004), 'Finite-size dynamics of inhibitory and excitatory interacting spiking neurons', *Physical Review E* **70**, 052903.
- McClelland, J. L. (2001), *Mechanisms of Cognitive Development: Behavioral and Neural Approaches.*, Lawrence Erlbaum Associates, Mahwah, NJ, chapter Failures to learn and their remediation: A Hebbian account., pp. 197–211.
- McClelland, J. L., McNaughton, B. L. & O'Reilly, R. C. (1995), 'Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory.', *Psychol Rev* **102**(3), 419–57.
- McCormick, D., Connors, B., Lighthall, J. & Prince, D. (1985), 'Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons in the neocortex', *Journal of Neurophysiology* **54**, 782–806.
- Middleton, F. A. & Strick, P. L. (1994), 'Anatomical evidence for cerebellar and basal ganglia involvement in higher cognitive function', *Science* **266**(5184), 458–461.
- Miller, E. K. & Cohen, J. D. (2001), 'An integrative theory of prefrontal cortex function', *Annu Rev Neurosci* **24**, 167–202.
- Milnor, J. (1985), 'On the concept of attractor.', *Communications in Mathematical Physics* **99**, 177–195.
- Mitz, A. R., Godschalk, M. & Wise, S. P. (1991), 'Learning-dependent neuronal activity in the premotor cortex: activity during the acquisition of conditional motor associations', *J Neurosci* **11**(6), 1855–1872.
- Miyamoto, S., Duncan, G. E., Marx, C. E. & Lieberman, J. A. (2005), 'Treatments for schizophrenia: a critical review of pharmacology and mechanisms of action of antipsychotic drugs', *Molecular Psychiatry* **10**(1), 79–104.
- Moller, H. J. (2005), 'Antipsychotic and antidepressive effects of second generation antipsychotics: two different pharmacological mechanisms?', *Eur Arch Psychiatry Clin Neurosci* **255**(3), 190–201.
- Mueser, K. T. & McGurk, S. R. (2004), 'Schizophrenia', *Lancet* **363**(9426), 2063–2072.
- Murray, E. A., Bussey, T. J. & Wise, S. P. (2000), 'Role of prefrontal cortex in a network for arbitrary visuomotor mapping', *Exp Brain Res* **133**(1), 114–129.
- Murray, E. A. & Wise, S. P. (1996), 'Role of the hippocampus plus subjacent cortex but not amygdala in visuomotor conditional learning in rhesus monkeys', *Behavioral Neuroscience* **110**, 1261–1270.

- Newcomer, J. W., Farber, N. B., Jevtovic-Todorovic, V., Selke, G., Melson, A. K., Hershey, T., Craft, S. & Olney, J. W. (1999), 'Ketamine-induced NMDA receptor hypofunction as a model of memory impairment and psychosis', *Neuropsychopharmacology* **20**(2), 106–118.
- Nixon, P. D., McDonald, K. R., Gough, P. M., Alexander, I. H. & Passingham, R. E. (2004), 'Cortico-basal ganglia pathways are essential for the recall of well-established visuomotor associations', *Eur J Neurosci* **20**(11), 3165–3178.
- Okubo, Y., Suhara, T., Sudo, Y. & Toru, M. (1997), 'Possible role of dopamine D1 receptors in schizophrenia', *Molecular Psychiatry* **2**(4), 291–292.
- Okubo, Y., Suhara, T., Suzuki, K., Kobayashi, K., Inoue, O., Terasaki, O., Someya, Y., Sassa, T., Sudo, Y., Matsushima, E., Iyo, M., Tateno, Y. & Toru, M. (1997), 'Decreased prefrontal dopamine D1 receptors in schizophrenia revealed by PET', *Nature* **385**(6617), 634–636.
- O'Reilly, R. C. (2006), 'Biologically based computational models of high-level cognition', *Science* **314**(5796), 91–94.
- Packard, M. G. & Knowlton, B. J. (2002), 'Learning and memory functions of the basal ganglia', *Annu Rev Neurosci* **25**, 563–593.
- Pandya, D. N. & Kuypers, H. G. (1969), 'Cortico-cortical connections in the rhesus monkey', *Brain Res* **13**(1), 13–36.
- Parker, A. & Gaffan, D. (1998), 'Memory after frontal/temporal disconnection in monkeys: conditional and non-conditional tasks, unilateral and bilateral frontal lesions', *Neuropsychologia* **36**(3), 259–271.
- Passingham, R. E. (1993), *The frontal lobes and voluntary action.*, Oxford University Press, Oxford.
- Passingham, R. E. & Toni, I. (2001), 'Contrasting the dorsal and ventral visual systems: guidance of movement versus decision making', *Neuroimage* **14**, 125–131.
- Passingham, R. E., Toni, I. & Rushworth, M. F. (2000), 'Specialisation within the prefrontal cortex: the ventral prefrontal cortex and associative learning', *Exp Brain Res* **133**(1), 103–113.
- Pasupathy, A. & Miller, E. K. (2005), 'Different time courses of learning-related activity in the prefrontal cortex and striatum', *Nature* **433**(7028), 873–876.
- Petrides, M. (1982), 'Motor conditional associative-learning after selective prefrontal lesions in the monkey', *Behav Brain Res* **5**(4), 407–413.
- Petrides, M. (1985), 'Deficits on conditional associative-learning tasks after frontal- and temporal-lobe lesions in man', *Neuropsychologia* **23**(5), 601–614.
- Petrides, M. (1990), 'Nonspatial conditional learning impaired in patients with unilateral frontal but not unilateral temporal lobe excisions', *Neuropsychologia* **28**(2), 137–149.
- Petrides, M. (1997), 'Visuo-motor conditional associative learning after frontal and temporal lesions in the human brain', *Neuropsychologia* **35**(7), 989–997.

- Ratcliff, R., Zandt, T. V. & McKoon, G. (1999), 'Connectionist and diffusion models of reaction time', *Psychol Rev* **106**(2), 261–300.
- Renart, A., Moreno, R., Rocha, J., Parga, N. & Rolls, E. T. (2001), 'A model of the IT–PF network in object working memory which includes balanced persistent activity and tuned inhibition', *Neurocomputing* **38–40**, 1525–1531.
- Rescorla, R. A. & Wagner, A. R. (1972), *A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement*, Appleton-Century Company, New York.
- Ridley, R. M. & Baker, H. F. (1997), 'Evidence for a specific information processing deficit in monkeys with lesions of the septo-hippocampal system.', *Cortex* **33**(1), 167–76.
- Rockland, K. S. & Pandya, D. N. (1979), 'Laminar origins and terminations of cortical connections of the occipital lobe in the rhesus monkey.', *Brain Res* **179**(1), 3–20.
- Rolls, E. T. (2004), 'The functions of the orbitofrontal cortex', *Brain Cogn* **55**(1), 11–29.
- Rolls, E. T. (2005), *Emotion Explained*, Oxford University Press, Oxford.
- Rolls, E. T. (2006), The neurophysiology and functions of the orbitofrontal cortex, in D. H. Zald & S. L. Rauch, eds, 'The Orbitofrontal Cortex', Oxford University Press, Oxford, chapter 5, pp. 95–124.
- Rolls, E. T. & Deco, G. (2002), *Computational Neuroscience of Vision*, Oxford University Press, Oxford.
- Rolls, E. T. & Treves, A. (1998), *Neural Networks and Brain Function*, Oxford University Press.
- Rupniak, N. M. & Gaffan, D. (1987), 'Monkey hippocampus and learning about spatially directed movements', *Journal of Neuroscience* **7**, 2331–2337.
- Rushworth, M. F. S. & Owen, A. M. (1998), 'The functional organization of the lateral frontal cortex: conjecture or conjuncture in the electrophysiology literature?', *Trends in Cognitive Sciences* **2**(2), 46–53.
- Salin, P. & Prince, D. (1996), 'Spontaneous GABA-A receptor mediated inhibitory currents in adult rat somatosensory cortex', *Journal of Neurophysiology* **75**, 1573–1588.
- Sawaguchi, T. & Goldman-Rakic, P. S. (1991), 'D1 dopamine receptors in prefrontal cortex: Involvement in working memory', *Science* **251**, 947–950.
- Sawaguchi, T. & Goldman-Rakic, P. S. (1994), 'The role of d1-dopamine receptor in working memory: local injections of dopamine antagonists into the prefrontal cortex of rhesus monkeys performing an oculomotor delayed-response task', *J Neurophysiol* **71**(2), 515–528.
- Scheuerecker, J., Ufer, S., Zipse, M., Frodl, T., Koutsouleris, N., Zetzsche, T., Wiesmann, M., Albrecht, J., Bruckmann, H., Schmitt, G., Moller, H. J. & Meisenzahl, E. M. (2007), 'Cerebral changes and cognitive dysfunctions in medication-free schizophrenia - An fMRI study', *J Psychiatr Res* .

- Schultz, W. (2000), 'Multiple reward signals in the brain', *Nat Rev Neurosci* **1**(3), 199–207.
- Schultz, W. (2006), 'Behavioral theories and the neurophysiology of reward.', *Annu Rev Psychol* **57**, 87–115.
- Seamans, J. K., Gorelova, N., Durstewitz, D. & Yang, C. R. (2001), 'Bidirectional dopamine modulation of GABAergic inhibition in prefrontal cortical pyramidal neurons', *J Neurosci* **21**(10), 3628–3638.
- Seamans, J. K. & Yang, C. R. (2004), 'The principal features and mechanisms of dopamine modulation in the prefrontal cortex', *Prog Neurobiol* **74**(1), 1–58.
- Seeman, P. & Kapur, S. (2000), 'Schizophrenia: more dopamine, more D2 receptors', *Proc Natl Acad Sci U S A* **97**(14), 7673–5.
- Seeman, P., Schwarz, J., Chen, J. F., Szechtman, H., Perreault, M., McKnight, G. S., Roder, J. C., Quirion, R., Boksa, P., Srivastava, L. K., Yanai, K., Weinschenker, D. & Sumiyoshi, T. (2006), 'Psychosis pathways converge via D2 high dopamine receptors', *Synapse* **60**(4), 319–346.
- Seeman, P., Weinschenker, D., Quirion, R., Srivastava, L. K., Bhardwaj, S. K., Grandy, D. K., Premont, R. T., Sotnikova, T. D., Boksa, P., El-Ghundi, M., O'Dowd, B. F., George, S.-R., Perreault, M. L., Mannisto, P. T., Robinson, S., Palmiter, R. D. & Talerico, T. (2005), 'Dopamine supersensitivity correlates with D2 High states, implying many paths to psychosis', *Proc Natl Acad Sci U S A* **102**(9), 3513–3518.
- Servan-Schreiber, D., Printz, H. & Cohen, J. D. (1990), 'A network model of catecholamine effects: gain, signal-to-noise ratio, and behavior', *Science* **249**(4971), 892–895.
- Shergill, S. S., Brammer, M. J., Williams, S. C., Murray, R. M. & McGuire, P. K. (2000), 'Mapping auditory hallucinations in schizophrenia using functional magnetic resonance imaging', *Arch Gen Psychiatry* **57**(11), 1033–1038.
- Spruston, N., Jonas, P. & Sakmann, B. (1995), 'Dendritic glutamate receptor channel in rat hippocampal CA3 and CA1 pyramidal neurons', *Journal of Physiology* **482**, 325–352.
- Stephan, K. E., Baldeweg, T. & Friston, K. J. (2006), 'Synaptic plasticity and disconnection in schizophrenia', *Biol Psychiatry* **59**(10), 929–939.
- Sutton, R. S. & Barto, A. G. (1998), *Reinforcement learning*, MIT Press, Cambridge, MA.
- Thorpe, S. J., Rolls, E. T. & Maddison, S. (1983), 'The orbitofrontal cortex: neuronal activity in the behaving monkey', *Exp Brain Res* **49**(1), 93–115.
- Toni, I. & Passingham, R. E. (1999), 'Prefrontal-basal ganglia pathways are involved in the learning of arbitrary visuomotor associations: a PET study', *Exp Brain Res* **127**(1), 19–32.
- Toni, I., Ramnani, N., Josephs, O., Ashburner, J. & Passingham, R. E. (2001), 'Learning arbitrary visuomotor associations: temporal dynamic of brain activity', *Neuroimage* **14**(5), 1048–1057.
- Toni, I., Rowe, J., Stephan, K. E. & Passingham, R. E. (2002), 'Changes of cortico-striatal effective connectivity during visuomotor learning', *Cereb Cortex* **12**(10), 1040–1047.

- Toni, I., Rushworth, M. F. & Passingham, R. E. (2001), 'Neural correlates of visuomotor associations. spatial rules compared with arbitrary rules', *Exp Brain Res* **141**(3), 359–369.
- Trantham-Davidson, H., Neely, L. C., Lavin, A. & Seamans, J. K. (2004), 'Mechanisms underlying differential D1 versus D2 dopamine receptor regulation of inhibition in prefrontal cortex', *J Neurosci* **24**(47), 10652–9.
- Tremblay, L., Hollerman, J. R. & Schultz, W. (1998), 'Modifications of reward expectation-related neuronal activity during learning in primate striatum', *J Neurophysiol* **80**(2), 964–977.
- Tuckwell, H. (1988), *Introduction to Theoretical Neurobiology*, Cambridge University Press, Cambridge.
- Wang, M., Zhang, H. & Li, B. M. (2000), 'Deficit in conditional visuomotor learning by local infusion of bicuculline into the ventral prefrontal cortex in monkeys', *Eur J Neurosci* **12**(10), 3787–3796.
- Wang, X. J. (1999), 'Synaptic basis of cortical persistent activity: the importance of NMDA receptors to working memory', *J Neurosci* **19**(21), 9587–9603.
- Wang, X. J. (2001), 'Synaptic reverberation underlying mnemonic persistent activity', *Trends Neurosci* **24**(8), 455–463.
- Wang, X. J. (2002), 'Probabilistic decision making by slow reverberation in cortical circuits', *Neuron* **36**(5), 955–968.
- Wang, X. J., Tegner, J., Constantinidis, C. & Goldman-Rakic, P. S. (2004), 'Division of labor among distinct subtypes of inhibitory neurons in a cortical microcircuit of working memory', *Proc Natl Acad Sci U S A* **101**(5), 1368–1373.
- Watanabe, M. (1996), 'Reward expectancy in primate prefrontal neurons', *Nature* **382**(6592), 629–632.
- Weiss, A. P. & Heckers, S. (1999), 'Neuroimaging of hallucinations: a review of the literature', *Psychiatry Res* **92**(2-3), 61–74.
- Williams, Z. M. & Eskandar, E. N. (2006), 'Selective enhancement of associative learning by microstimulation of the anterior caudate.', *Nat Neurosci* **9**(4), 562–8.
- Wilson, F., Scialoja, S. & Goldman-Rakic, P. (1994), 'Functional synergism between putative gamma-aminobutyrate-containing neurons and pyramidal neurons in prefrontal cortex', *Proceedings of the National Academy of Science* **91**, 4009–4013.
- Winterer, G., Coppola, R., Goldberg, T. E., Egan, M. F., Jones, D. W., Sanchez, C. E. & Weinberger, D. R. (2004), 'Prefrontal broadband noise, working memory, and genetic risk for schizophrenia', *Am J Psychiatry* **161**(3), 490–500.
- Winterer, G., Musso, F., Beckmann, C., Mattay, V., Egan, M. F., Jones, D. W., Callicott, J. H., Coppola, R. & Weinberger, D. R. (2006), 'Instability of prefrontal signal processing in schizophrenia', *Am J Psychiatry* **163**(11), 1960–1968.

- Winterer, G. & Weinberger, D. R. (2004), 'Genes, dopamine and cortical signal-to-noise ratio in schizophrenia', *Trends Neurosci* **27**(11), 683–690.
- Winterer, G., Ziller, M., Dorn, H., Frick, K., Mulert, C., Wuebben, Y., Herrmann, W. M. & Coppola, R. (2000), 'Schizophrenia: reduced signal-to-noise ratio and impaired phase-locking during information processing', *Clin Neurophysiol* **111**(5), 837–849.
- Wirth, S., Yanike, M., Frank, L. M., Smith, A. C., Brown, E. N. & Suzuki, W. A. (2003), 'Single neurons in the monkey hippocampus and learning of new associations', *Science* **300**(5625), 1578–1581.
- Wise, S. P. & Murray, E. A. (1999), 'Role of the hippocampal system in conditional motor learning: mapping antecedents to action', *Hippocampus* **9**(2), 101–17.
- Wise, S. P. & Murray, E. A. (2000), 'Arbitrary associations between antecedents and actions', *Trends Neurosci* **23**(6), 271–6.
- Wise, S. P., Murray, E. A. & Gerfen, C. R. (1996), 'The frontal cortex-basal ganglia system in primates', *Crit Rev Neurobiol* **10**(3-4), 317–56.
- Wolkin, A., Sanfilipo, M., Wolf, A. P., Angrist, B., Brodie, J. D. & Rotrosen, J. (1992), 'Negative symptoms and hypofrontality in chronic schizophrenia', *Arch Gen Psychiatry* **49**(12), 959–965.
- Xiang, Z., Huguenard, J. & Prince, D. (1998), 'GABA-A receptor mediated currents in interneurons and pyramidal cells of rat visual cortex', *Journal of Physiology* **506**, 715–730.